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

Anonymous Authors[1]

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 strategies using machine learning algorithms and defining research tasks.

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. 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 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 investment dynamics and short interest. 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.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

2. Model

We consider a population of n agents, acting as investment funds, who trade shares of a single representative asset. Every period-day t, the funds can buy, sell and short-sell shares of the asset in constant supply. Asset shares pay daily dividends $D(t)$ following an autocorrelated Geometric Brownian Motion. Cash 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, as well as adding a new strategy evolution component described in Section 5.

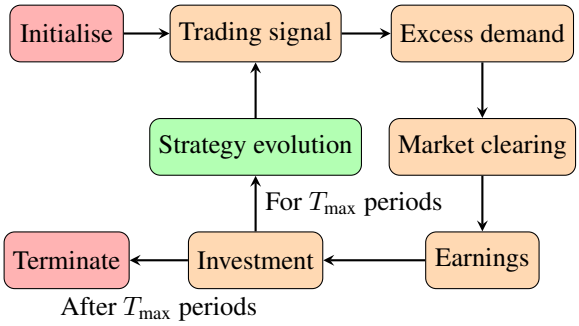

Figure 1. Visual summary of the model components.

2.1. Trading strategies and signals

Like real markets, our financial market model features a diverse sample of stylized versions of the most common funds' trading strategies (Scholl et al., 2021), adding heterogeneous sub-strategies. *Noise traders* (NT) trade on individual-specific, mean-reverting Ornstein Uhlenbeck processes $X(t)$, mimicking crowd dynamics not driven by financial fundamentals. We calibrate the parameters of the Ornstein Uhlenbeck process to match empirical excess volatility (Scholl et al., 2021). *Value investors* (VI) form heterogeneous subjective valuations V_i of the asset based on discounted sums of dividends. *Trend followers* (TF) trade

on the existence of trends in the asset price over various time horizons. Agents' trading strategies are represented by their *trading signals* $\phi(t)$.

$$\phi_i^{\mathrm{NT}}(t) = X_i(t) \tag{1}$$

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

$$\phi_i^{\mathrm{TF}} = \log_2\left(p(t)/p(t-\theta_i)\right) \tag{3}$$

2.2. Excess demand

The funds' daily trading signals are inputs of the *excess demand function* for the asset. The excess demand function expresses the demand of the fund for the asset as a function of the unknown price $p(t)$. The $\tanh$ function smooths and bounds the trading signal in the feasible trading range $[-1, 1]$. In the case of indifference ($\phi(t) = 0$), the 0.5 term targets an equal split of wealth between asset shares and cash. Fund wealth W is the sum of agent cash, present value of asset shares and liabilities.

$$E(t, p(t)) = \frac{W(t, p(t-1))}{p(t)} \tanh\left(\phi(t) + \frac{1}{2}\right) - S(t) \tag{4}$$

2.3. Market-clearing

The *market-clearing* process finds the root of the aggregate excess demand, i.e. the price for which the sum of the funds' excess demands equals 0, demand matching supply. When this is not possible, we minimize the difference between demand and supply.

2.4. Earnings

After computing the clearing price, funds execute the resulting excess demand orders. Agents receive earnings: the dividends $D(t)$ and interest r corresponding to their new positions. The excess demand can be negative: agents then (short) sell shares.

2.5. Investment

The actions of external investors play an essential role in the wealth dynamics of mutual funds. Depending on the performance of the funds, 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.

2.6. 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.

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 set model input parameters (short position size and investment) using real data to reduce the dimensionality of the parameter space. We *validate* the model showing that it reproduces the stylised financial properties of asset returns (Cont, 2001).

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 all main stylised facts (Cont, 2001). Our log prices display intermittency and volatility clustering: irregular bursts in volatility that tend to cluster together, separated by periods of low volatility. The log price returns do not show significant linear autocorrelations past trivial frequencies. Returns themselves show heavy tail distributions, displaying 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.

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 the profitability of a given strategy, including a realistic mapping from returns to investment flows is thus necessary.

To understand the empirical relationship between fund performance measured by returns and redemption flows, we analysed reporting data from the Securities and Exchange Commission. Investment companies are subject to reporting requirements of their assets, redemptions, sales, and other indicators through the form N-PORT. Net redemption flows appear negative for good funds and positive for funds with negative returns at the one-quarter horizon. Although some previous literature identified a positive association (Chevalier & Ellison, 1997), other analyses have highlighted strong contrarian or profit-taking behaviours in mutual fund investors (Goetzmann & Massa, 2002; Rakowski & Yamani, 2021). The sign of the flows changes non-linearly with respect to the lagged return horizon and fund's performance, following an S-shape approximated by a sigmoid function. We model the share of redemption flows over the net value of

the assets as being distributed as a double Gumbel distribution. Our investment module is calibrated on those empirical distributions to provide a realistic mapping between funds' returns and external investments.

3.3. Calibrating Short Interest

The size of short positions for the stock impacts ecology wealth dynamics. Some strategies might accumulate substantial short positions, be vulnerable to price movements and become extinct. The objective of this calibration of short interest is thus to bound the financial system in a realistic operating window. We ensure that our single asset is representative of the most popular stocks traded in empirical markets, considering short interest: how many shares of a stock are sold short compared to the total amount of shares. Using stock exchange data sourced by Yahoo Finance, we found that the average short ratio days to cover of the ten most popular NASDAQ stocks (i.e. with the highest market capitalisation) is 1.124. Compared to the number of shares available for trading, the average short-selling volume represents 1.17% of the float. In our model, we have restricted the maximum size of funds' short positions accordingly.

4. Results

We present the dynamics of the Evology ABM: 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 share of noise traders goes down to small 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 totally 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 very detrimental to the other species. This stability and diversity are desirable: actual markets feature a diversity of profitable strategies.

5. Trading Strategy Optimisation in Evology

5.1. The trading strategy optimisation problem

Optimisation problem 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.

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

Dynamic strategy training can lead to dynamic strategies: the function ϕ could change at every step. *Static* strategy training optimizes a single function ϕ that will generate the list of actions $\phi(t)$.

The objective function for the trading strategy optimisation problem is typically a measure of profit, 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). On a meta-level, this simulator also allows studying the performance of the optimisation objectives, checking which metrics lead to the most profitable strategies and possibly searching for an optimal combination of objectives.

Lucas' critique and changes in market conditions It has been known since the Lucas critique 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 be modelled as a fictitious agent with a constant positive excess demand for the asset.

5.2. Simulation training and time series training

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 been successful in creating 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 dis-

tribution 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.

5.3. Benchmark tasks

Here, we present some tasks to advance trading strategy optimisation in Evology.

1. **Evolve static/dynamic strategies that are better than the base strategies**. First, the optimisation algorithms should be capable of finding statistically significantly more profitable strategies than those already present in the model and some trivial strategies, such as the typical buy and hold or the interest rate on cash. The objective would be to evolve realistic, high-performing strategies without human input, e.g. the Kelly criterion. Setting a small percentage of the fund population to be adaptive, trading strategy search under *coevolution* could be an interesting task. In real markets, actively managed funds and hedge funds are adaptive. Coevolution makes trading strategy optimisation more challenging because the evolved strategies may be copied or face stronger opponents in the simulation, impacting their profitability. Setting a benchmark for simulation performance is 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.

2. **Evolve robust, interpretable strategies generalisable to real markets** For the trained strategies to be interesting, they need to be *interpretable* and *robust*. 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, obstacles that are easy to implement in the ABM. Finally, the Evology environment would be more useful if it could train strategies that could be used in real markets, an additional challenge to robustness and ABM realism.

3. **Continuous development and calibration of the simulation platform** The end 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 the development of 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 the SEC data publicly available, model calibration constitutes a benchmark task of interest.

6. Conclusion and next steps

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 and propose a set of tasks to advance ABM-based trading strategy optimisation: evolving better strategies than the base ones and under coevolution, improving strategies' interpretability and robustness.

Software and Data

The Evology ABM is available open-source[1]. 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[2].

Acknowledgements

This publication is based on work supported by the EPSRC Centre for Doctoral Training in Mathematics of Random Systems: Analysis, Modelling and Simulation

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

(EP/S023925/1). The authors thank Christian Schroeder de Witt for advice.

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.

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

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.

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

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

Goetzmann, W. N. and Massa, M. Daily momentum and contrarian behavior of index fund investors. *Journal of Financial and Quantitative Analysis*, 37(3):375–389, 2002.

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.

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.

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.

Rakowski, D. and Yamani, E. Endogeneity in the mutual fund flow–performance relationship: An instrumental variables solution. *Journal of Empirical Finance*, 64: 247–271, 2021.

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.

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

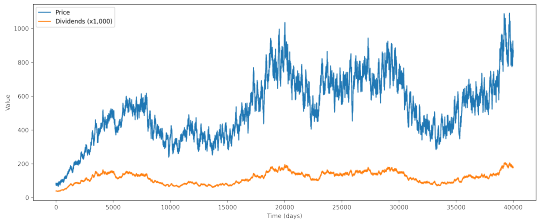

Figure 2. Price (blue) and dividends ($\times 1000$, orange) during a 40,000-day (around 150 years) simulation run, starting from the initial coordinates $[1/3, 1/3, 1/3]$.

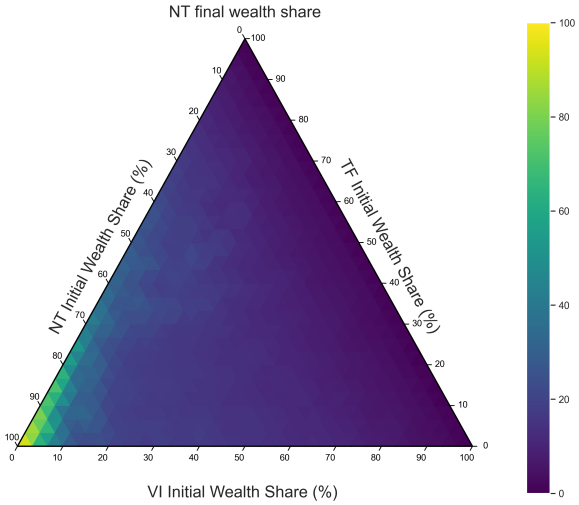

Figure 3. Noise traders' final wealth share

Figure 4. Value investors' final wealth share

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

Table 2.

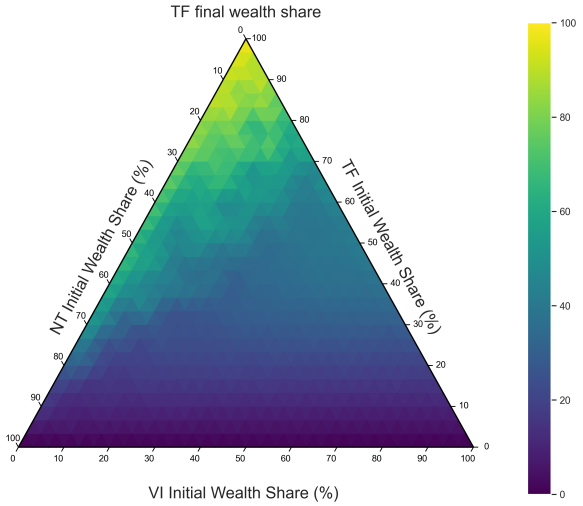

Figure 5. Trend followers' final wealth share