



How market ecology explains market malfunction

Maarten P. Scholl^{a,b,1} , Anisoara Calinescu^b , and J. Doyne Farmer^{a,c,d} 

^aInstitute for New Economic Thinking, Oxford Martin School, University of Oxford, Oxford OX1 3QD, United Kingdom; ^bComputer Science Department, University of Oxford, Oxford OX1 3QD, United Kingdom; ^cMathematical Institute, University of Oxford, Oxford OX1 3QD, United Kingdom; and ^dSanta Fe Institute, Santa Fe, NM 87501

Edited by Simon Asher Levin, Princeton University, Princeton, NJ, and approved March 24, 2021 (received for review September 16, 2020)

Standard approaches to the theory of financial markets are based on equilibrium and efficiency. Here we develop an alternative based on concepts and methods developed by biologists, in which the wealth invested in a financial strategy is like the abundance of a species. We study a toy model of a market consisting of value investors, trend followers, and noise traders. We show that the average returns of strategies are strongly density dependent; that is, they depend on the wealth invested in each strategy at any given time. In the absence of noise, the market would slowly evolve toward an efficient equilibrium, but the statistical uncertainty in profitability (which is calibrated to match real markets) makes this noisy and uncertain. Even in the long term, the market spends extended periods of time away from perfect efficiency. We show how core concepts from ecology, such as the community matrix and food webs, give insight into market behavior. For example, at the efficient equilibrium, all three strategies have a mutualistic relationship, meaning that an increase in the wealth of one increases the returns of the others. The wealth dynamics of the market ecosystem explain how market inefficiencies spontaneously occur and gives insight into the origins of excess price volatility and deviations of prices from fundamental values.

market ecology | market efficiency | agent-based modeling

Why do markets malfunction? According to the theory of market efficiency, markets always function perfectly. Prices always reflect fundamental values and change only when new information affects fundamental values. Thus, by definition, any problems with price setting are caused by factors outside the market. Empirical evidence suggests otherwise. Large price movements occur even when there is very little new information (1), and prices often deviate substantially from fundamental values (2). This means that we need to go beyond the theory of market efficiency to understand how and why markets malfunction.

Here we build on earlier work (3–7)* and develop the theory of market ecology, which provides the necessary alternative. This approach borrows concepts and methods from biology and applies them to financial markets. Financial trading strategies are analogous to biological species. Plants and animals are specialists that evolve to fill niches that provide food; similarly, financial trading strategies are specialists that evolve to exploit market inefficiencies. Trading strategies can be classified into distinct categories, such as technical trading, value investing, market making, statistical arbitrage, and many others. The capital invested in a strategy is like the population of a species. Trading strategies interact with one another via price setting, and the market evolves as the wealth invested in each strategy changes through time, as regulations change, and as old strategies fail and new strategies appear.

The theory of market ecology emerges from the inherent contradictions in the theory of market efficiency. A standard argument used to justify market efficiency is that competition for profits by arbitrageurs should cause markets to rapidly evolve to an equilibrium where it is not possible to make excess profits based on publicly available information. But, if there are no profits to be made, there are no incentives for arbitrageurs, so there

is no mechanism to make markets efficient. This paradox suggests that, while markets may be efficient in some approximate sense, they cannot be perfectly efficient (8). In contrast, under the theory of market ecology, trading strategies exploit market inefficiencies, but, as new strategies appear and as the wealth invested in each strategy changes, the inefficiencies change as well. To understand how the market functions, it is necessary to understand how each strategy affects the market and how the interactions between strategies cause market inefficiencies to change with time. The theory of market ecology naturally addresses a different set of problems than the theory of market efficiency and can be viewed as a complement rather than a substitute.

Our study here builds on a large body of work on agent-based models of financial markets (e.g. refs. 9–12). The theory of market ecology provides a conceptual framework for understanding such models. Our goal is not to construct a better model of financial markets, but rather to show how ideas from ecology can be used to interpret market phenomena and predict market behavior.

Here we study a stylized toy market model with three trading strategies. We approach the problem in the same way that an ecologist would study three interacting species. We study how the average returns of the strategies depend on the wealth invested in each strategy, how their wealth evolves through time under reinvestment, and how their endogenous time evolution causes the market to malfunction.

We show that, with realistic parameters, evolution toward market efficiency is very slow. The expected deviations from

Significance

We develop the mathematical analogy between financial trading strategies and biological species and show how to apply standard concepts from ecology to financial markets. We analyze the interactions of stereotypical trading strategies in ecological terms, showing that they can be competitive, predator-prey, or mutualistic, depending on the wealth invested in each strategy. The deterministic dynamics suggest that the system should evolve toward an efficient state where all strategies make the same average returns. However, this happens slowly and the evolution is so noisy that there are large fluctuations away from the efficient state, causing bursts of volatility and extended periods where prices deviate from fundamental values. This provides a conceptual framework that gives insight into the reasons why markets malfunction.

Author contributions: M.P.S., A.C., and J.D.F. designed research, performed research, contributed new analytic tools, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

¹To whom correspondence may be addressed. Email: maarten.scholl@cs.ox.ac.uk.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2015574118/-DCSupplemental>.

Published June 25, 2021.

*B. LeBaron, Building the Santa Fe artificial stock market. www2.econ.iastate.edu/tesfatsi/blake.sfishum.pdf (2002).

efficiency are, in some sense, small, but they persist even in the long term, and cause extended deviations from fundamental values and excess volatility (which, in extreme cases, becomes market instability). Our study provides a simple example of how analyzing markets in these terms and tracking market ecosystems through time could give regulators and practitioners better insight into market behavior.

Model Description

The structure of the model is schematically summarized in Fig. 1. There are two assets, a stock and a bond. The bond trades at a fixed price and yields $r = 1\%$ annually in the form of coupon payments that are paid out continuously. The stock pays a dividend $D(t)$ at each time step that is modeled as a discrete time-autocorrelated geometric Brownian motion, of the form

$$\begin{aligned} D(t) &= D(t-1) + gD(t-1) + \sigma D(t-1)U(t), \\ U(t) &= \omega U(t-\tau) + (1-\omega^2)Z(t), \end{aligned} \quad [1]$$

where g is the average rate of dividend payments, σ is the variance, ω is the autocorrelation parameter of the process, and Z is Gaussian noise. The auxiliary process $U(t)$ introduces persistence. We choose parameters so that one time step is roughly equal to a day. We use estimates from market data by LeBaron (13), taking $g = 2\%$ per year for the growth rate of the dividend with a volatility of $\sigma = 10\%$ (see ref. 14, for example, for a review of the empirical evidence on dividends). The dividend process is positively autocorrelated through the auxiliary process $U(t)$ with lag $\tau = 1$ d.

We use market clearing to set prices. The stock has a fixed supply Q , but the excess demand $E(t)$ for the stock by each trading strategy varies in time. We allow the trading strategies to take short positions and to use leverage (i.e., to borrow in order to take a position in the stock that is larger than their wealth). We impose a strategy-specific leverage limit $\bar{\lambda}$. Because we use leverage and because the strategies can have demand functions with unusual properties, market clearing is not always straightforward—see *Materials and Methods*.

The size of a trading strategy is given by its wealth $W(t)$, that is, the capital invested in it at any given time. In ecology, this corresponds to the population of a species, which is also called its abundance. Unless otherwise stated, the wealth of each strategy varies proportional to its cumulative performance. Letting $\pi_i(t)$ be the return of strategy i at time t , the wealth changes according to

$$W_i(t+1) = (1 + f\pi_i(t))W_i(t). \quad [2]$$

The reinvestment rate f models investor flows of capital. The default value $f = 1$ corresponds to passive reinvestment, and $f > 1$ means that profitable strategies attract additional capital and unprofitable strategies lose additional capital.

A trading strategy is defined by its trading signal $\phi(t)$, which can depend on the price $p(t)$ and other variables, such as dividends and past prices. We modify ϕ by a tanh function to




Fig. 1. The three trading strategies correspond to noise traders, value investors and trend followers. They invest their capital in a stock and a bond. The mixture for each strategy changes with time as strategies accumulate wealth based on their historical performance.

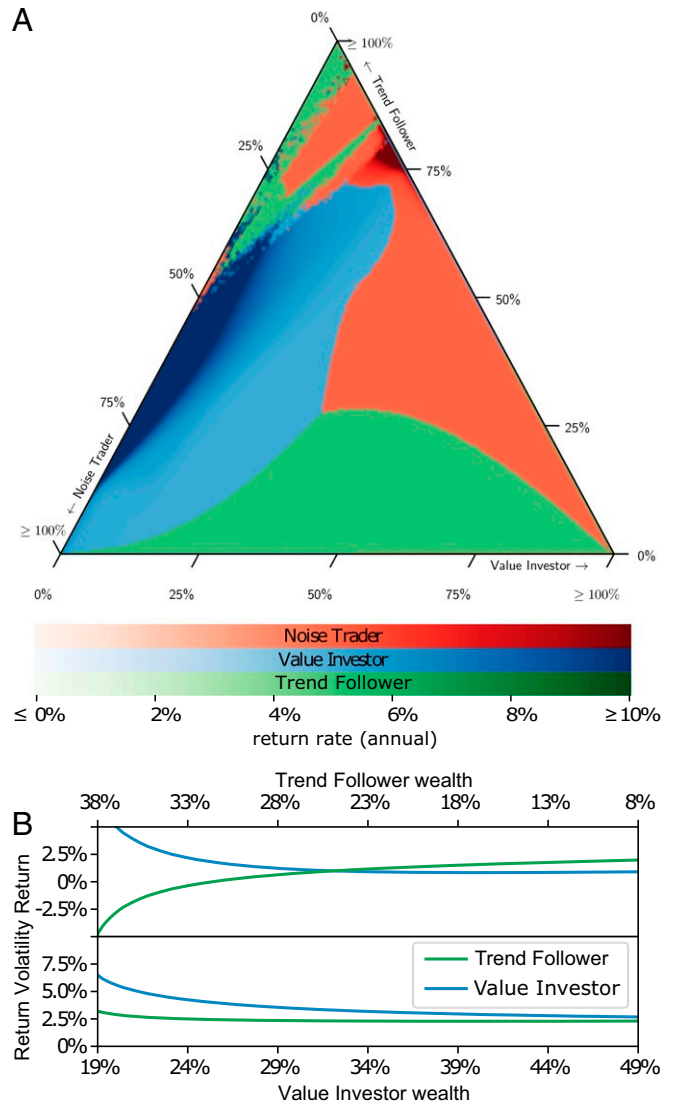


Fig. 2. The profitability of the dominant strategy in the wealth landscape. **A** is a ternary plot which displays the returns achieved by the strategy with the largest return. The axes correspond to the relative wealth invested in each strategy. The top corner is pure trend followers, the left corner is pure noise traders, and the right corner is pure value investors. The color indicates the strategy with the highest returns at a given relative wealth vector \mathbf{w} . The regions colored in red correspond to the noise traders, blue regions correspond to value investors, and green regions correspond to the trend followers. The intensity of the color indicates the size of the average return. **(B) Upper** shows the average returns to value investors (blue) and trend followers (green) while holding the noise trader wealth at its equilibrium value of 42%. **Lower** shows the volatility in the returns of each strategy. The horizontal axis is the relative wealth of the trend follower (top axis) and value investor (bottom axis).

ensure that the excess demand is bounded and differentiable. A strategy's excess demand for the stock is

$$E(t) = \frac{W(t-1)\bar{\lambda}}{p(t)} \left(\tanh(c \cdot \phi(t)) + \frac{1}{2} \right) - S(t-1), \quad [3]$$

where $S(t-1)$ is the number of shares of the stock held at the previous time step, and $\bar{\lambda}$ is the strategy-specific leverage limit. The parameter $c > 0$ determines the aggressiveness of the response to the signal ϕ , and is strategy specific. When the signal of the strategy is zero, the agent is indifferent between the

stock and the bond and splits its portfolio equally between the two (hence the term of $1/2$). The leverage $\lambda(t)$ of a strategy at any given time is $p(t)S(t)/W(t) = \bar{\lambda} |\tanh(c \cdot \phi(t) + 1/2)|$. As in the usual Walrasian process, $p(t)$ is set such that aggregate excess demand $\sum_i E_i(t)$ over all agents i is zero.

Investment Strategies. We study three typical trading strategies, which we call value investors, trend followers, and noise traders. We intentionally make all strategies boundedly rational, that is, they work from limited information, and their strategies are not optimal. We use a representative agent hypothesis, treating each strategy as though it were only used by a single fund; however, these should be thought of as representing all investors using these strategies. We now describe each strategy in turn.

Value investors. Value investors observe the dividend process and use a model to derive the value of the stock. They seek to

hold more of the stock when it is undervalued and hold more of the bond when the stock is overvalued. The parameters of their model are estimated based on historical dividends.

The fundamental value $V(t)$ of the stock at $t=0$ is given as the discounted expected future dividends,

$$V(0) = \mathbb{E} \left[\sum_{t=1}^{\infty} \frac{D(t)}{(1+k)^t} \right]. \quad [4]$$

The parameter k is a discount rate called the *required rate of return*, with $k \geq r$; k is the sum of the risk-free rate r and a risk premium all of the investors in our model expect for the additional risks associated with the stock. We follow ref. 15 and use a fixed discount rate $k = 2\%$, based on the average rate of return implied by historical data.

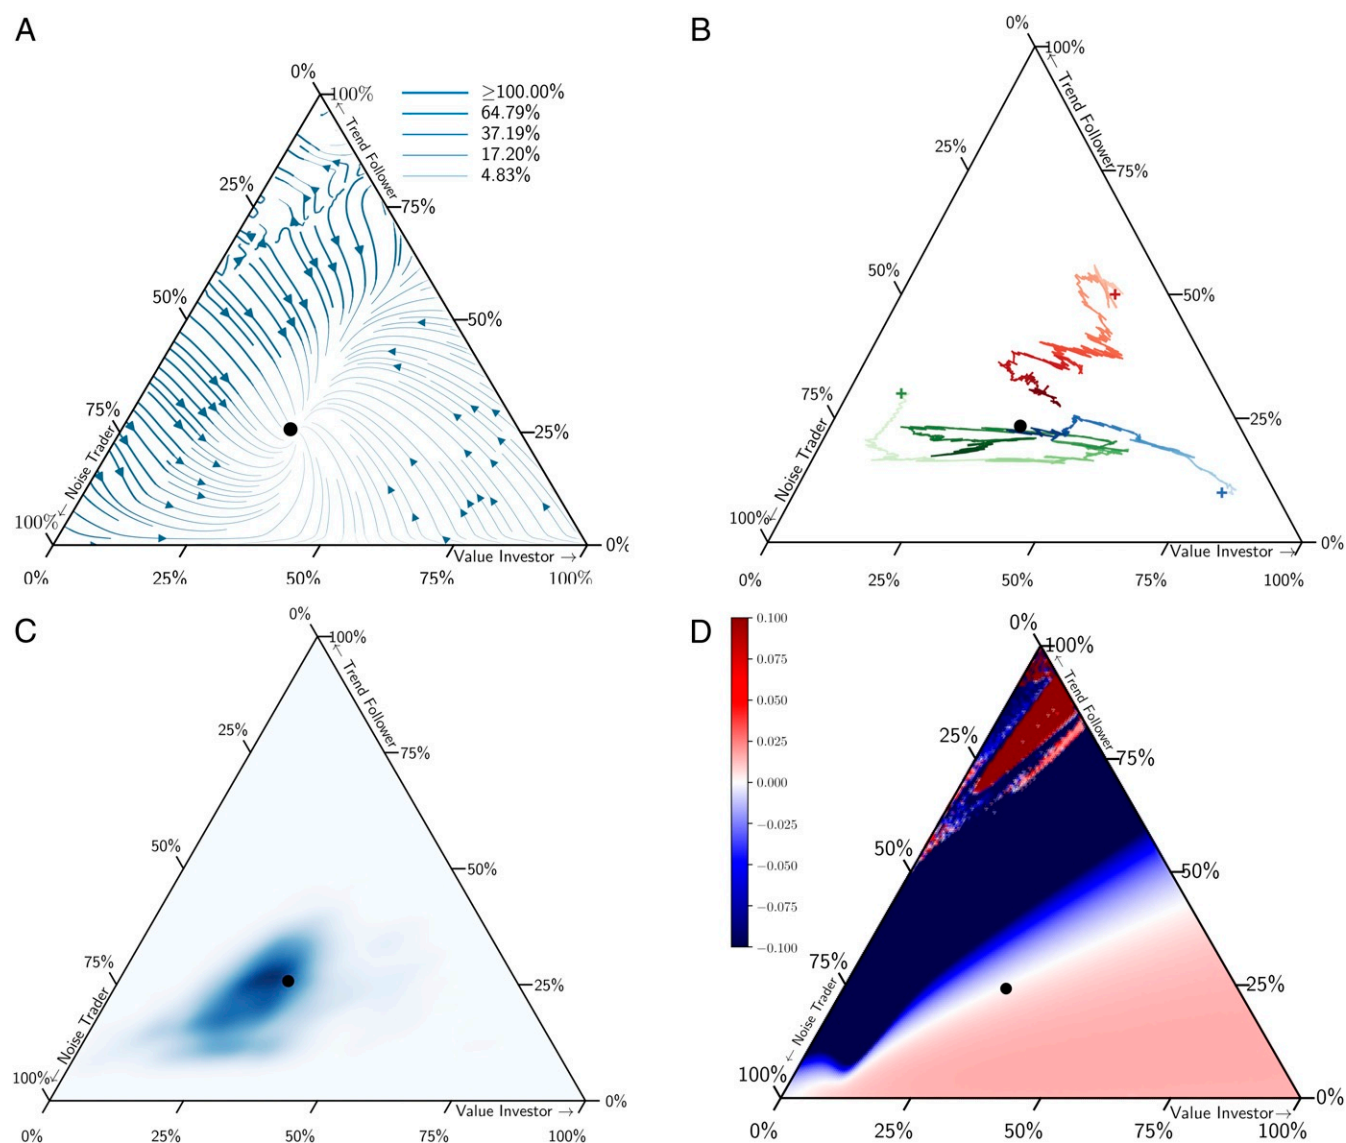


Fig. 3. Profit dynamics as a function of wealth. **A** shows how wealth evolves, on average, through time under reinvestment. The intensity of the color denotes the rate of change. **B** shows sample trajectories for a few different initial values of the wealth vector, making it clear that the trajectories are extremely noisy due to statistical uncertainty, so that the deterministic dynamics of **A** are a poor approximation. The visualization displays three different initial wealth vectors, each color coded. The marker “+” indicates the initial wealth. The trajectories with the same color follow the system for $T = 200$ y, and color saturation increases with time. Starting from uniformly distributed initial conditions, **C** displays a density map of the asymptotic wealth distribution after 200 y. The system is initialized at random with a uniformly distributed wealth vector and then allowed to freely evolve for 200 y. The color’s darkness is proportional to density. **D** displays the one-time-step autocorrelation of price returns. The black dot is the equilibrium point from **A**.

Table 1. Estimated community matrix near the equilibrium at $w = (NT = 0.43, VI = 0.34, TF = 0.23)$

G_{ij}	NT, %	VI, %	TF, %
NT	-0.89	0.89	0.82
VI	26.6	-10.6	22.4
TF	11.1	15.2	-19.3

NT, noise traders; VI, value investors; TF, trend followers.

The valuation $V^{VI}(t)$ used by the value investors is made by estimating the mean growth rate g from past data using the classical dividend discount model (16) $\mathbb{E}[D(t)] = D(0)(1 + g)^t$. This ignores the autocorrelation coefficient ω , and so, in general, $V^{VI}(t) \neq V(t)$.

We define the trading signal for the value investor as the difference in log prices between the estimated fundamental value $V^{VI}(t)$ and the market price.

$$\phi_{VI}(t) = \log_2 V^{VI}(t) - \log_2 p(t). \quad [5]$$

This strategy enters into a long position when the proposed price is lower than the estimated fundamental value and enters into a short position when the proposed price is higher than the estimated fundamental value. The use of the base two logarithm means that the value investor employs all of its assets when the stock is trading at half the perceived value (17).

Trend followers. Trend followers expect that historical trends in returns continue into the short-term future. Several variants exist in the literature, including the archetypal trend follower that we use here (3, 10, 18–20). There is evidence to suggest that trend-based investment strategies are profitable over long time horizons, and ref. 21 argues that investors earn a premium for the liquidity risk associated with stocks with high momentum (momentum trading is a synonym for trend following).

The trend-following strategy extrapolates the trend in price between recently realized prices $p(t-1)$ and $p(t-2)$ time steps in the past. It always buys when prices trend upward and sells when they trend downward.

$$\phi_{TF}(t) = \log_2 p(t-1) - \log_2 p(t-2). \quad [6]$$

The trend followers' demand is a decreasing function of price. Unlike the value investor, who cannot observe autocorrelations in dividends, the trend follower can exploit the autocorrelation that dividends impart to prices. Because the positive autocorrelation in market prices is, at most, $\omega = 0.1$, we multiply the signal by $1/\omega$.

Noise traders. Noise traders represent nonprofessional investors who do not track the market closely. Their transactions are mostly for liquidity, but they are also somewhat aware of value, so they are slightly more likely to buy when the market is undervalued and slightly more likely to sell when the market is overvalued (this is necessary so that their positions do not diverge over long periods of time).

The signal function of our noise traders contains the product of the value estimate $V^{VI}(t)$ (which is the same as for the value investors) and a stochastic component $X(t)$,

$$\phi_{NT}(t) = \log_2 X(t) V^{VI}(t) - \log_2 p(t). \quad [7]$$

The noise process $X(t)$ is a discretized Ornstein–Uhlenbeck process, which has the form

$$X(t) = X(t-1) + \rho(\mu - X(t-1)) + \gamma\epsilon. \quad [8]$$

This process reverts to the long-term mean $\mu = 1$ with reversion rate $\rho = 1 - \sqrt[6 \times 252]{0.5}$, meaning the noise has a half life of 6 y, in accordance with the values estimated by Bouchaud et al. (22). The variable ϵ is a standard normal random variable, and $\gamma = 20\%$ is a volatility parameter, chosen so that noise traders generate volatility in excess of the volatility of dividends, matching the level observed in the US stock market.

The parameters of the model are summarized later (see Table 5). We have chosen them for an appropriate trade-off between realism and conceptual interest, for example, so that each strategy has a region in the wealth landscape where it is profitable. We will see that there are a few ways in which the properties of these strategies do not match those observed in real markets (the trend strategy is very short term, the value investor uses somewhat high leverage, and, at equilibrium, the noise trader has a surprisingly low return-to-risk ratio). Nonetheless, they are realistic enough to make our key points.

Results

Density Dependence. An ecosystem is said to have density dependence if its characteristics depend on the population sizes of the species, as is typically the case. Similarly, a market ecosystem is density dependent if its characteristics depend on the wealths of the strategies (both its own wealth and that of other strategies). The toy market ecosystem that we study here is strongly density dependent.

When the core ideas in this paper were originally introduced in ref. 3, prices were formed using a market impact function, which translates the aggregate trade imbalance at any time into a shift in prices. This can be viewed as a local linearization of market clearing. The use of a market impact function suppresses density dependence and neglects nonlinearities that are important for understanding market ecology.

With market clearing, there is strong density dependence. This is evident in Fig. 2, which shows which strategy makes the highest profits as a function of the relative size of each of the three strategies. To control the size of each strategy, we turn off reinvestment, and instead replenish the wealth of each strategy at each step as needed to hold wealth constant. We then systematically vary the wealth vector $W = (W_{NT}, W_{VI}, W_{TF})$. We somewhat arbitrarily let the total wealth $W_T = W_{NT} + W_{VI} + W_{TF} = 3 \times 10^8$. For convenience of interpretation, we plot the relative wealth $w_i(t) = W_i/W_T$. The results shown are averages over many long runs; to avoid transients, we exclude the first 252 time steps, corresponding to one trading year.

Roughly speaking, the profitability of the dominant strategy divides the wealth landscape into four distinct regions. Trend followers dominate at the bottom of the diagram, where their wealth is small. Value investors dominate on the left side of the diagram, where their wealth is small, and noise traders dominate on the right side of the diagram, where their wealth is small. There is an intersection point near the center where the returns of all three strategies are the same. In addition, there is a complicated region at the top of the diagram, where no single strategy dominates. The turbulent behavior in this region comes about because the wealth invested by trend followers is large, and the price dynamics are unstable. We do not regard this region as realistic, except perhaps in rare extreme market conditions.

Table 2. Estimated community matrix near $w = (NT = 0.26, VI = 0.55, TF = 0.19)$

G_{ij}	NT, %	VI, %	TF, %
NT	-0.46	0.40	0.36
VI	8.94	-0.77	-1.89
TF	6.81	6.87	-9.65

A quantitative snapshot of the average returns and volatility is given in Fig. 2B, where we hold the size of the noise traders constant at 42%, corresponding to their wealth at the intersection point, and vary the wealth of the value investors and trend followers. The average return to both trend followers and value investors decreases monotonically as their wealth increases. The volatility of the returns of both strategies, in contrast, is a monotonic function of the wealth of the trend followers alone—higher trend follower wealth implies higher volatility. Although this is not shown here, the average return of the value investors increases strongly with the wealth of the noise traders; in contrast, the average return of the trend followers is insensitive to it.

Adaptation and the Slow Approach to Market Efficiency. We now investigate the dynamics of the market ecosystem. To understand how the wealth of the strategies evolves through time, we allow reinvestment and plot trajectories corresponding to the average return at each wealth vector \mathbf{w} . This is done by averaging over many different runs. The result is shown in Fig. 3A. Most of the wealth trajectories in the diagram evolve toward a fixed point where the wealths of the strategies no longer change. There is also a region at the top of the diagram where the dynamics are more complicated (due to instabilities) and a region in the lower left corner where the ecosystem evolves toward the boundary of the simplex.

At the fixed point, the annual returns to the three strategies are all equal to 2.05%, which, to an investor, is statistically indistinguishable from the 2% return from simply buying and holding the stock. We will loosely refer to this fixed point as the *efficient equilibrium*. We say “loosely” because the volatilities of the strategies are 4.07% for noise traders, 6.76% for value investors, 4.62% for trend followers, and 9.09% for a buy and hold, so that, with a more sophisticated model of fund flows, the equilibrium might shift slightly in favor of strategies with less risk. A common way to measure the performance of a trading strategy is in terms of the ratio of the mean to the standard deviation of its returns, which is called the Sharpe ratio, S . The corresponding Sharpe ratios (without subtracting the risk-free rate) are 0.50, 0.31, and 0.44 for the three strategies, and 0.22 for a buy and hold. These Sharpe ratios and their variation are reasonable numbers for investment funds, indicating that our model ecosystem is roughly as efficient (or inefficient) as a real market (if anything, the Sharpe ratios are a bit low).

The large central region of initial conditions that are attracted to the efficient equilibrium gives a misleading impression of a smooth evolution toward a state of market efficiency. In fact, the dynamics are noisy and stray far from the deterministic dynamics shown in Fig. 3A. Tracking a few individual trajectories, as we do in Fig. 3B, demonstrates that the dynamics are dominated by noise, due to the statistical uncertainty in the performance of the strategies. The typical trajectories bear little correspondence to the deterministic trajectories of Fig. 3A, and the convergence to the efficient equilibrium is weak.

To get a feeling for the asymptotic wealth distribution, we sample the space of initial wealth uniformly, simulate the ecosystem dynamics under reinvestment with $f = 1$, and record the final wealth after 200 y, as shown in Fig. 3C. The deviations from the efficient equilibrium point are substantial, often more than 20%. Furthermore, the evolution toward the asymptotic distribution is exceedingly slow: Each trajectory in Fig. 3B spans 200 y of simulated time. There are substantial changes in the relative wealth taking place over time scales that are longer than a century.

The long time scale to reach efficiency observed here matches with the estimate made by Farmer in ref. 3. In the ideal case of a stationary market and independent and identically distributed (I.I.D.) normally distributed returns, the time required to detect excess performance ΔS with a statistical significance of s stan-

dard deviations is approximately $(s/\Delta S)^2$. To take an example, a buy and hold of the S&P index has a Sharpe ratio of roughly $S = 0.5$. For a strategy whose annualized Sharpe ratio is superior by $\Delta S = 0.1$, a 20% improvement over a buy and hold, roughly 400 y are required to confirm its superior performance with 2 SD. Furthermore, as shown in ref. 23, the approach to market efficiency follows a power law of the form $t^{-\alpha}$, where $0 \leq \alpha \leq 1$. For large times, this is much slower than an exponential. This happens because the approach to efficiency slows down as the market becomes more efficient. Near the efficient equilibrium, the dynamics are dominated by the noise.

To demonstrate that statistical uncertainty is the dominant factor determining the approach to efficiency, we did a series of experiments that are described in detail in *SI Appendix, section 4*. As the reinvestment rate f in Eq. 2 varies in the range $0.1 \leq f \leq 3$, the rate of approach to the asymptotic wealth distribution remains roughly the same. In contrast, varying the noise trader volatility, which affects the statistical uncertainty in the performance of all three strategies, has a substantial effect.

The absence of autocorrelation in price returns is an indicator of market efficiency. Efficient price returns should have an autocorrelation that is reasonably close to zero (close enough that it is not possible to make statistically significant excess profits). In Fig. 3D, we plot the one-time-step autocorrelation of returns across the wealth landscape. There is a striking white band across the center of the simplex, corresponding to zero autocorrelation. This happens when trend followers invest about 40% of the total wealth, thereby eliminating the autocorrelation coming from the dividend process. The wealth of trend followers fluctuates, even at very long times, and, consequently, the degree of autocorrelation in price returns fluctuates as well.

Community Matrix. The community matrix is a tool used in ecology to describe the pairwise effects of the population of species

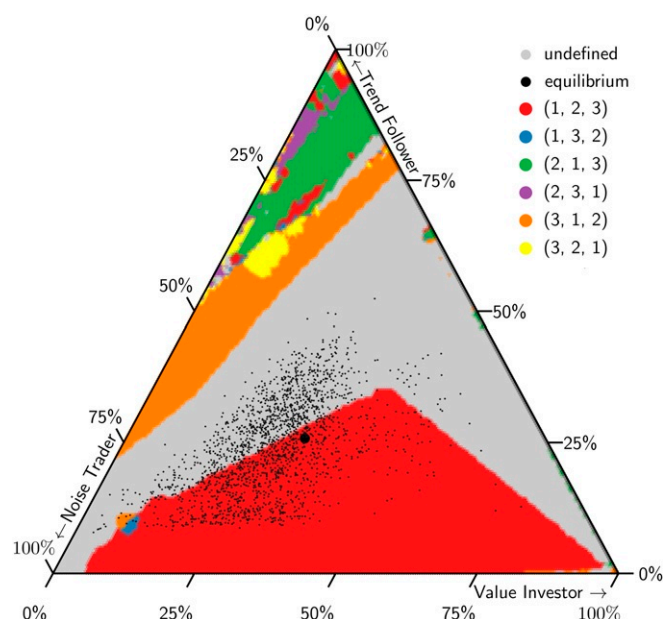


Fig. 4. A survey of the trophic levels across the wealth landscape. We color the diagram according to the increasing ordering of the trophic levels of the three strategies, in the order (noise trader, value investor, trend follower), as indicated in the legend. (The numbers now indicate the ordering rather than the precise trophic level). The dominant zone where the ordering is (1, 2, 3) is colored red. In the gray region, there are cycles where the trophic levels become undefined. The black dots correspond to samples of the wealth vector after 200 y, as shown in Fig. 3C. The system spends most of its time in the gray and red zones.

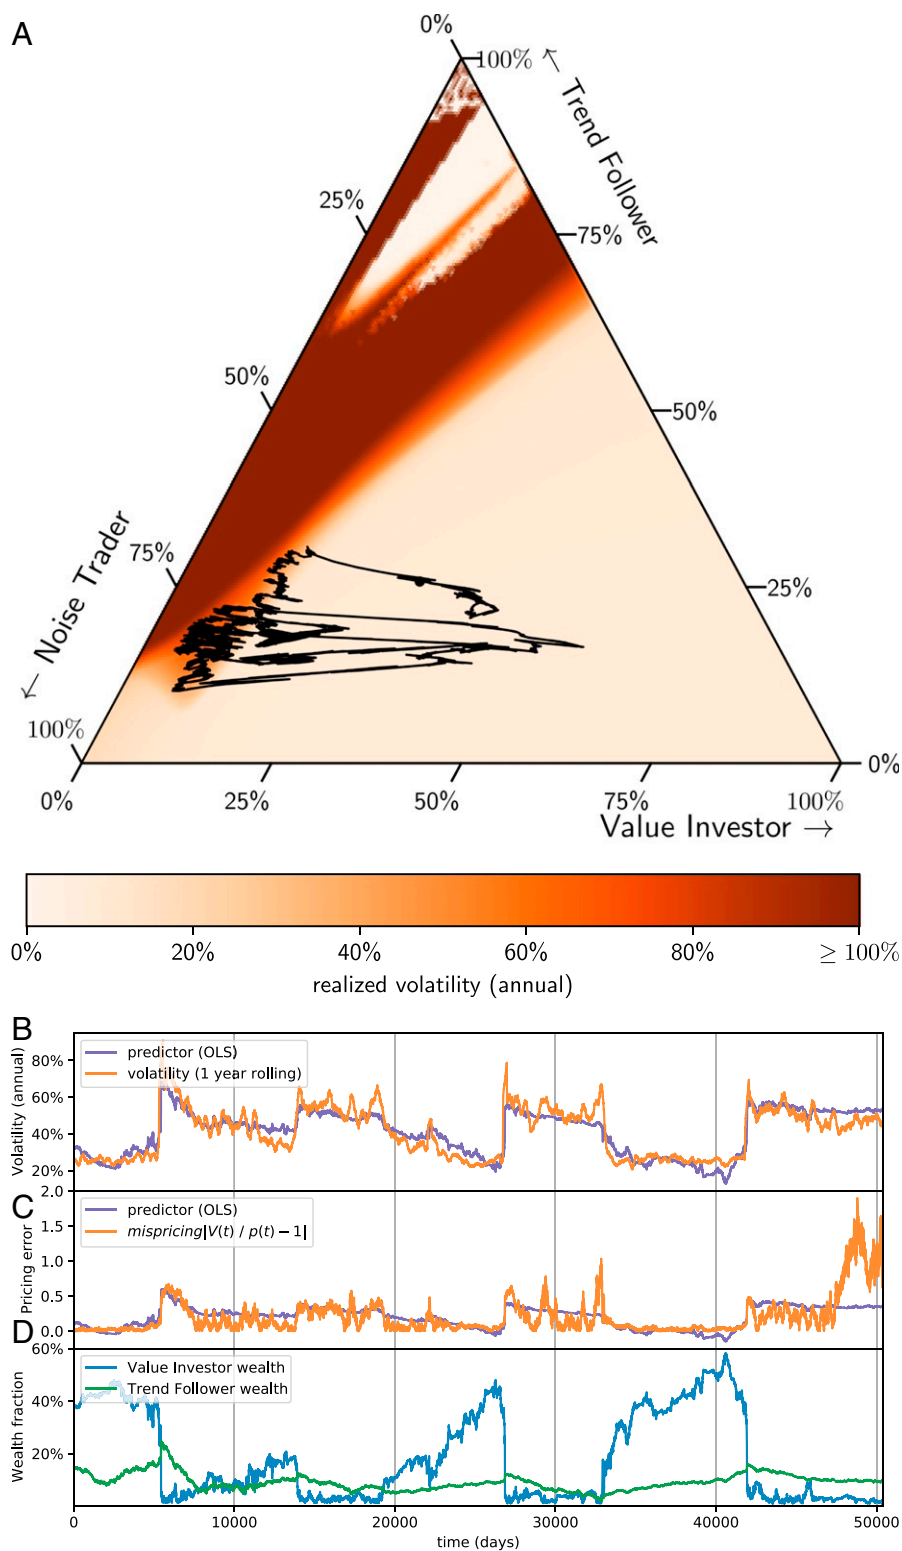


Fig. 5. How ecosystem dynamics cause market malfunctions. *A* gives a color map of the price volatility over the wealth landscape; the volatility is low and constant throughout the lower right part of the diagram, where the system spends most of its time, but there is a high-volatility region running across the upper left. A sample trajectory spanning 200 y, beginning at the efficient equilibrium, is shown in black. The noise caused by statistical fluctuations in performance causes large deviations from equilibrium and excursions into the high volatility region. *B* shows the volatility of this trajectory as a function of time, plotted against the predicted volatility (see Eq. 12). *C* shows the actual mispricing plotted against the predicted mispricing. *D* shows the wealth of the value investors and trend followers.

j on the population growth rate of species i (24, 25). As originally pointed out by Farmer (3), who called it the gain matrix, an analogous quantity is also useful for interpreting the behavior of market ecosystems.

Let π_i be the average return of strategy i , that is, $\pi_i = \lim_{T \rightarrow \infty} \left(\prod_{t=1}^T \pi_i(t) - 1 \right)^{1/T}$. The analogue of the community matrix for market ecology is

$$G_{ij} = \frac{\partial \pi_i}{\partial w_j}. \quad [9]$$

This has units of one over time. The wealth $w_i(t)$ invested in strategy i replaces the population size of a species. The possible pairwise interactions between strategies can be classified according to the sign of G_{ij} . If both G_{ij} and G_{ji} are negative, then strategies i and j are competitive; if G_{ij} is positive and G_{ji} is negative, then there is a predator–prey interaction, with i as the predator and j as the prey; and, if both G_{ij} and G_{ji} are positive, then there is a mutualistic interaction (26).

Because we do not have a differentiable model for our toy market ecosystem, we compute the community matrix numerically using finite differences (see *Materials and Methods*). The community matrix is strongly density dependent. If we compute the community matrix near the equilibrium point in the center of the simplex, we get the result shown in Table 1.

The diagonal entries of the community matrix are all negative, indicating that the strategies are competitive with themselves. This means that their average returns diminish as the strategy gets larger, causing what is called *crowding* in financial markets. We already observed this in Fig. 2. Interestingly, however, the size of the diagonal terms varies considerably, from -0.89 for noise traders to -19.3 for trend followers. This means that we should expect trend followers to experience crowding much more strongly than noise traders.

All of the other entries of the community matrix are positive, indicating mutualism. This implies that every strategy benefits from an increase in the wealth of any of the other strategies. While we initially found it surprising that all of the strategies could have mutualistic interactions with each other, on reflection, this makes sense: The ecosystem is, by definition, efficient at the equilibrium, and driving any of the strategies away from equilibrium creates an inefficiency that provides a profit opportunity for the other two strategies.

The community matrix is density dependent. If we compute the community matrix at the wealth vector given in Table 2, where the value investors are dominant, there is a shift in the pairwise community relations. As before, all of the terms in the row corresponding to the noise traders are small, indicating that the noise traders are not strongly affected by other strategies, and that they compete only weakly with themselves. This should not be surprising—the noise traders' strategy is mostly random, and is less influenced by prices than the other two strategies. Value investors, who have the majority of the wealth in this case, still strongly benefit from an increase in the wealth of noise traders (although less so than at the equilibrium). However, there is now a negative (2, 3) off-diagonal term, while the opposite (3, 2) term remains positive. In other words, returns to value investors drop if the wealth of trend followers increases, but not vice versa, implying that trend followers now have a predator–prey relationship with value investors. Other variations in community relationships can be found at different points in the wealth landscape, illustrating density dependence.

The Lotka–Volterra equations, which describe how the populations in an idealized predator–prey system evolve through time, are perhaps the most famous equations in population biology. Their surprising result is that, at some parameter values, they have solutions that oscillate indefinitely. Using the assump-

tion of no density dependence, Farmer derived Lotka–Volterra equations for market ecology (3). Our results here indicate that the density dependence in this system is so strong that simple Lotka–Volterra equations are a poor approximation, at least for this system. The existence of oscillating solutions in financial ecosystems remains an open question.

Food Webs and Trophic Level. The food web provides an important conceptual framework for understanding the interactions between species. If lions eat zebras, and zebras eat grass, then the population of lions is strongly affected by the density of grass, and, similarly, the density of grass depends on the population of lions, even though lions have no direct interactions with grass. The trophic level of a species is, by definition, one level higher than what it eats, so, in this idealized system, grass has trophic level one, zebras have trophic level two, and lions have trophic level three.

The existence of animals with more-complicated diets, such as omnivores and detritivores, means that real food webs are never this simple. If we let A_{ij} be the share of species j in the diet of species i , then the trophic level T_i of species i can be computed by the relation

$$T_i = 1 + \sum_j A_{ij} T_j. \quad [10]$$

The resulting trophic levels are typically not integers, but they still provide a useful way to think about the role that a given species plays in the ecosystem.

We can also compute trophic levels for the strategies in an ecosystem. We define the analogous quantity A_{ij} as the fraction of the returns of strategy i that can be attributed to the presence of strategy j . We do this by simply comparing the returns of strategy i at wealth W to those when strategy j is removed, that is, when $W_j = 0$ but all of the other wealths remain the same. In mathematical terms,

$$A_{ij} = \frac{\max[0, \pi_i(W_1, \dots, W_j, \dots, W_N) - \pi_i(W_1, \dots, 0, \dots, W_N)]}{\sum_k \max[0, \pi_i(W_1, \dots, W_k, \dots, W_N) - \pi_i(W_1, \dots, 0, \dots, W_N)]}, \quad [11]$$

and $A_{ij} = 0$ when the denominator is zero. Note also that, in order for the diagonal entries A_{ii} to be defined, we can not set wealth to zero but set it to the smallest possible value, $1/256$, which is determined by the simulation grid size.

Eqs. 10 and 11 allow us to compute trophic levels. At the efficient equilibrium, the trophic levels are (noise trader = 1.00, value investor = 2.00, trend follower = 2.99). The proximity to integer values is because there are only three strategies with similar wealth levels: The noise trader cannot profit from either the value investor or the trend follower, the value investor profits from the noise trader, and the trend follower profits just a little from the noise trader but quite a lot from the value investor—see the discussion in *SI Appendix*. However, away from the equilibrium, where the wealths of the three strategies are substantially unequal, this changes. In order to better understand the density dependence, we compute trophic levels at

Table 3. Multivariate regressions with volatility and mispricing as dependent variables, and the funds' wealth as independent variables

Independent variable	Coefficient	t statistic
Volatility, $R^2 = 0.79$; 50,397 observations		
Noise trader	2.4	10
Value investor	−68	−249
Trend follower	107	169
Mispricing, $R^2 = 0.33$; 50,397 observations		
Noise trader	−0.15	−18
Value investor	−1.02	−107
Trend follower	1.5	69

Table 4. Details of the balance sheet items used by all funds

Assets	Liabilities
Cash C	Equity Capital K
Margin M	Debt Loans L
trading securities S ⁺	borrowed securities S ⁻

All securities use the most recent market value in valuation.

each point in the wealth landscape. For three strategies, there are 3! = 6 possible orderings of the trophic levels. We display the ordering of the trophic levels across the wealth landscape in Fig. 4.

The computation of trophic levels is complicated by the fact that, for some wealth vectors, there are cycles in the food web. For example, for $w = (0.05, 0.15, 0.80)$, value investors exploit noise traders (who cause reversion to fundamental value), trend followers exploit value investors (who induce autocorrelations), and noise traders exploit trend followers (who generate excess volatility), to complete a cycle. When this happens, Eq. 9 does not converge, and the trophic levels become undefined. Cycles in the trophic web are not unique to markets—they can also occur in biology, for example, due to cannibalism or detritivores.

A comparison of Figs. 4 and 3C makes it clear that, at long times (after transients have died out), the system divides its time between the region in which the trophic levels are ordered as (noise trader, value investor, trend follower), as they are at the equilibrium point, and the region where there are cycles, where the trophic levels are undefined.

How Ecosystem Dynamics Cause Market Malfunction. The wealth dynamics of the market ecosystem help explain why the market malfunctions and illuminate the origins of excess volatility and mispricing, that is, deviations of prices from fundamental values. Both in real markets (2) and in the agent-based models mentioned earlier (including our model here), volatility and mispricings change endogenously with time—there are eras where they are large and eras where they are small. Volatility varies intermittently, with periods of low volatility punctuated by bursts of high volatility, called *clustered volatility*. The standard explanations for clustered volatility are fluctuating agent populations (9, 10, 27) and leverage (28). We focus here on the first explanation; while we also observe that clustered volatility increases with increasing leverage, we have not investigated this, in detail, here.

Fig. 5A presents the variation of the volatility across the wealth landscape. The landscape can roughly be divided into two regions. On the lower right, there is a flat “low-volatility plain” occupying most of the landscape. On the upper left, there is a high-volatility region, with a sharp boundary between the two. As we will now show, excursions into the high-volatility region cause clustered volatility. A similar story holds for mispricing.

Fig. 5A shows a sample trajectory that begins at the efficient equilibrium and spans 200 y. This sample is representative of trajectories that fluctuate around the equilibrium point indefinitely. We choose a 200-y time span to show the scale at which spikes in volatility and mispricing occur. The statistical fluctuations in the performance of the three strategies act as noise, causing large excursions away from equilibrium. The trajectory mostly remains on the volatility plain, but there are several epochs where it ventures into the high-volatility region, causing bursts of high volatility.

The wealth dynamics have strong explanatory power for both mispricing and volatility. This is illustrated in Table 3, where we perform regressions of the strategies’ wealth against volatility using daily values for the time series shown in Fig. 5A. For volatility, $R^2 = 0.79$, and, for mispricing, $R^2 = 0.33$. In both cases, the

value investor’s wealth and the trend follower’s wealth have large coefficients (in absolute value), and the fit is overwhelmingly statistically significant. The noise trader is also highly statistically significant, but the coefficients and the t statistics are more than an order of magnitude smaller. In Fig. 5B and C, we compare a time series of the predicted volatility and predicted mispricing against the actual values. The predictions are very good.

$$\begin{aligned} \hat{v} &= -68w_{vi} + 107w_{tf} + 2.4w_{nt} \\ \hat{m} &= -1.02w_{vi} + 1.5w_{if} - 0.15w_{nt} \end{aligned} \quad [12]$$

Note that, in both cases, the coefficient for trend followers is positive, indicating that they drive instabilities, and the coefficient for value investors is negative, indicating their stabilizing influence.

Nonetheless, due to their effect on the population of value investors, the net effect of the trend followers on market malfunctions is not obvious. In Fig. 5D, we plot the wealth of value investors and trend followers. The strong mutualism predicted by the community matrix is clearly evident from the fact that the wealth of trend followers and value investors rise and fall together. However, their dynamics are quite different—there are several precipitous drops in the value investors’ wealth, whereas the trend followers tend to take more-gradual losses. As predicted, the highest-volatility episodes happen when the value investors’ wealth drops sharply while the trend followers’ wealth is high.

Discussion

Our analysis here demonstrates how understanding fluctuations of the wealth of the strategies in the ecosystem can help us predict market malfunctions such as mispricings and endogenously generated clustered volatility. The toy model that we study here is simple and highly stylized, but it illustrates how one can import ideas from ecology to better understand financial markets. Our analysis of this model illustrates several properties of market ecosystems that we hypothesize are likely to be true in more general settings.

Concepts from ecology give important insights into how deviations from market efficiency occur and how they affect prices. While the market may be close to efficiency in the sense that the excess returns to any given strategy are small, nonetheless, there can be substantial deviations in the wealth of different strategies, that can cause excess volatility and market instability.

Market ecology is a complement rather than a substitute for the theory of market efficiency. There are circumstances, such as pricing options, where market efficiency is a useful hypothesis. Market ecology, in contrast, provides insight into how and why markets deviate from efficiency, and what the consequences of this are. It can be used to explain the time dependence in the returns of trading strategies, and, in some cases, it can be used to explain market malfunctions. One of our main

Table 5. List of the model parameters and their values

Parameter	Value	Description
r	1% annual	Risk-free rate
g	1% annual	Dividend growth rate
k	2% annual	Cost of equity
σ	10% annual	dividend growth Volatility
ω	0.1	Dividend autocorrelation parameter
τ	1 d	Dividend autocorrelation lag
f	1	Reinvestment rate
ρ	$1 - \sqrt[6-252]{\frac{1}{2}}$	Noise trader mean reversion rate
σ^{NT}	20% annual	Noise trader volatility
$\bar{\lambda}_{NT}, \bar{\lambda}_{VI}, \bar{\lambda}_{TF}$	1, 8, 1	Leverage limit
C_{NT}, C_{VI}, C_{TF}	5, 10, 4	Signal scale

innovations here is to demonstrate how to compute the community matrix and the trophic web, which provide insight into the interactions of strategies. Surprisingly, at the efficient equilibrium, we find that all three strategies have mutualistic relationships to one another.

There are, so far, only a few examples of empirical studies of market ecology (29, 30). This is because such a study requires counterparty identifiers on transactions in order to know who traded with whom. Trying to study a market ecosystem without such data is like trying to study a biological ecosystem in which one can observe that an animal ate another animal, without any information about the types of animals involved. Unfortunately, for markets, such data are difficult for most researchers to obtain.

Regulators potentially have access to the balance sheets of all market participants, which can allow them to track, in detail, the ecology of the markets they regulate. Ideas such as those presented here could provide valuable insight into when markets are in danger of failure, and make it possible to construct models for the ecological effect of innovations, for example, the introduction of new types of assets such as mortgage-backed securities.

One of our most striking results is that the approach to efficiency is highly uncertain and exceedingly slow. As already pointed out, this should be obvious from a straightforward statistical analysis, but it is not widely appreciated. Our results demonstrate this dramatically, and they indicate that, even in the long term, we should expect large deviations from efficiency.

There are many possible extensions to this work. An obvious follow-up is to explore a larger space of strategies, or to let new strategies evolve in an open-ended way through time. Does the process of strategy innovation tend to stabilize or destabilize markets? Another follow-up is to construct a model that is empirically validated against data with counterparty identifiers. Our analysis here provides concepts and methods that could be used to interpret the behavior of real world examples.

1. D. M. Cutler, J. M. Poterba, L. H. Summers, What moves stock prices? *J. Portfolio Manag.* **15**, 4–12 (1989).
2. R. J. Shiller, *Market Volatility* (MIT Press, 1989).
3. J. D. Farmer, Market force, ecology and evolution. *Ind. Corp. Change* **11**, 895–953 (2002).
4. A. Lo, The adaptive market hypothesis: Market efficiency from an evolutionary perspective. *J. Portfolio Manag.* **30**, 15–29 (2004).
5. T. Hens, K. R. Schenk-Hoppé, *Handbook of Financial Markets: Dynamics and Evolution* (North Holland, 2009).
6. J. D. Farmer, S. Skouras, An ecological perspective on the future of computer trading. *Quant. Finance* **13**, 325–346 (2013).
7. N. Beale, R. Gunton, K. L. Bashe, H. S. Battey, R. S. MacKay, Dynamics of value-tracking in financial markets. *R. Soc. Open Sci.* **11**, 1–22 (2019).
8. S. J. Grossman, J. E. Stiglitz, On the impossibility of informationally efficient markets. *Am. Econ. Rev.* **70**, 393–408 (1980).
9. W. B. Arthur, J. H. Holland, B. D. LeBaron, R. G. Palmer, P. Tayler, “Asset pricing under endogenous expectations in an artificial stock market” in *The Economy as an Evolving Complex System II*, W. B. Arthur, S. N. Durlauf, D. Lane, Eds. (CRC, 1997), pp. 15–44.
10. W. A. Brock, C. H. Hommes, A rational route to randomness. *Econometrica* **65**, 1059 (1997).
11. T. Lux, M. Marchesi, Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature* **397**, 498–500 (1999).
12. C. Chiarella, The dynamics of speculative behaviour. *Ann. Oper. Res.* **37**, 101–123 (1992).
13. B. LeBaron, “Empirical regularities from interacting long- and short-memory investors in an agent-based stock market” in *IEEE Transactions on Evolutionary Computation* (IEEE, 2001), vol. 5, pp. 442–455.
14. E. F. Fama, K. R. French, Dividend yields and expected stock returns. *J. Financ. Econ.* **22**, 3–25 (1988).
15. B. LeBaron, “Agent-based computational finance” in *Handbook of Computational Economics*, L. Tesfatsion, K. L. Judd, Eds. (Elsevier, ed. 1, 2006), vol. 2, pp. 1187–1233.

Materials and Methods

Accounting and Balance Sheets. The funds in our model use a stylized balance sheet that is presented in Table 4. Funds are endowed with equity capital $K = W(0)$, in the form of cash C in dollars and shares of trading securities S . When $S > 0$, the fund holds this amount of securities, and when $S < 0$, it has borrowed this amount from other market participants, to create a short position. In order to guarantee that the short-selling fund can return the borrowed securities to the lender at a later time, the fund sets aside a margin amount M equal to the current market value of the borrowings, in the form of cash. The funds can use *leverage*, meaning using borrowed funds to purchase additional risky assets, by borrowing cash L . Public regulatory filings of US institutional fund managers indicate common leverage ratios between 1 and 10.[†] The interest rate that applies to cash holdings, loans, and margin is the same as the risk-free rate from holding the bond.

Wealth is calculated as $W(t) = C(t) + S(t)p(t) - L(t)$. A fund can only violate its leverage constraint when the proportion of risky assets changes faster than the amount of equity capital. This can happen due to losses. We require that all funds meet the solvency condition $W(t) > 0$. The simulation ends when one or more funds are insolvent. Table 5 lists the default parameters.

Software. The simulation in this paper builds on the Economic Simulation Library.[‡]

Data Availability. The model and code to run the experiments have been deposited in GitHub (<https://github.com/INET-Complexity/market-ecology>).

ACKNOWLEDGMENTS. We thank Klaus Schenk-Hoppé, Robert MacKay, Michael Wooldridge, Rama Cont, Alissa Kleinnijenhuis, and Cephas Svovse for enlightening discussion during the development of this work. We acknowledge funding by J.P. Morgan AI Faculty Awards, Baillie-Gifford, and the Rebuilding Macroeconomics program, funded by the Economic and Social Research Council. This research was partially supported by Trustworthy AI - Integrating Learning, Optimisation and Reasoning (TAILOR) (<https://tailor-network.eu/>), a project funded by European Union Horizon 2020 research and innovation program under Grant Agreement 952215.

[†]The Electronic Data Gathering, Analysis, and Retrieval system (EDGAR), publicly available at <https://www.sec.gov/edgar.shtml>, Securities and Exchange Commission.

[‡]An open-source library for agent-based modeling which is publicly accessible at <https://github.com/INET-Complexity/ESL>.

16. M. J. Gordon, E. Shapiro, Capital equipment analysis: The required rate of profit. *J. Manag. Sci.* **3**, 102–110 (1956).
17. F. Black, M. Scholes, The effects of dividend yield and dividend policy on common stock prices and returns. *J. Financ. Econ.* **1**, 1–22 (1974).
18. W. A. Brock, C. H. Hommes, Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J. Econ. Dynam. Contr.* **22**, 1235–1274 (1998).
19. J. D. Farmer, S. Joshi, The price dynamics of common trading strategies. *J. Econ. Behav. Organ.* **49**, 149–171 (2002).
20. B. LeBaron, Wealth dynamics and a bias toward momentum trading. *Finance Res. Lett.* **9**, 21–28 (2012).
21. C. S. Asness *et al.*, Value and momentum everywhere. *J. Finance* **68**, 929–985 (2013).
22. J.-P. P. Bouchaud *et al.*, Black was right: Price is within a factor 2 of value. arXiv [Preprint] (2017). <https://arxiv.org/abs/1711.04717> (Accessed 1 June 2018).
23. D. Cherkashin, J. D. Farmer, S. Lloyd, The reality game. *J. Econ. Dynam. Contr.* **33**, 1091–1105 (2009).
24. R. M. May, Will a large complex system be stable? *Nature* **238**, 413–414 (1972).
25. M. Novak *et al.*, Characterizing species interactions to understand press perturbations: What is the community matrix? *Annu. Rev. Ecol. Evol. Syst.* **47**, 409–432 (2016).
26. J. L. Bronstein, Conditional outcomes in mutualistic interactions. *Trends Ecol. Evol.* **9**, 214–217 (6 1994).
27. R. Franke, F. Westerhoff, Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *J. Econ. Dynam. Contr.* **36**, 1193–1211 (2012).
28. S. Thurner, J. D. Farmer, J. Geanakoplos, Leverage causes fat tails and clustered volatility. *Quant. Financ.* **12**, 695–707 (2012).
29. A. Kirilenko, A. S. Kyle, M. Samadi, T. Tuzun, The flash crash: High-frequency trading in an electronic market. *J. Finance* **72**, 967–998 (2017).
30. F. Musciotto, L. Marotta, J. Piilo, R. N. Mantegna, Long-term ecology of investors in a financial market. *Palgrave Commun.* **4**, 92 (2018).