

How Market Ecology Explains Market Malfunction

Maarten P. Scholl^{a,b,1}, Anisoara Calinescu^{b,1}, and J. Doyné Farmer^{a,c,d,2}

^aInstitute for New Economic Thinking, Oxford Martin School, University of Oxford; ^bComputer Science Department, University of Oxford; ^cMathematical Institute, University of Oxford; ^dSanta Fe Institute

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, i.e. 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 explains 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 (3–8) 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 (9). 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. (10–13). 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 efficiency are in some sense small, but they persist


Significance Statement

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 three 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 three strategies make the same average returns. However, this happens so 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.

The authors declare no conflict of interest.

¹ E-mail: maarten.scholl@cs.ox.ac.uk

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.



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.

A. Model Description. The structure of the model is schematically summarized in Figure 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, 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 (14), taking $g = 2\%$ per year for the growth rate of the dividend with a volatility of $\sigma = 10\%$ (see reference (15), 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 $\tau = 1$ day.

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* $W(t)$, i.e. 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 ensure that the excess demand is bounded and differentiable. A strategy's excess demand for the stock is

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

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 $\frac{p(t)S(t)}{W(t)} = \bar{\lambda} |\tanh(c \cdot \phi(t)) + \frac{1}{2}|$.

B. 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, i.e. 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 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 the investors in our model expect for the additional risks associated with the stock. We follow (16) and use a fixed discount rate $k = 2\%$, based on the average rate of return implied by historical data.

The valuation $V^{\text{VI}}(t)$ used by the value investors is made by estimating the mean growth rate g from past data using the classical dividend discount model (17) $\mathbb{E}[D(t)] = D(0)(1+g)^t$. This ignores the autocorrelation coefficient ω and so in general $V^{\text{VI}}(t) \neq V(t)$.

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

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

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 (18).

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, 11, 19–21). There is evidence to suggest that trend-based investment strategies are profitable over long time horizons, and reference (22) 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 $p(t-2)$ and $p(t-1)$ time steps in the past. It always buys when prices trend upwards and sells when they trend downwards.

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

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 represent non-professional 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

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

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 years, in accordance with the values estimated by Bouchaud (23). 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 U.S. stock market.

The parameters of the model are summarized in Table 5. We have chosen them for an appropriate tradeoff between realism and conceptual interest, e.g. 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

C. 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.

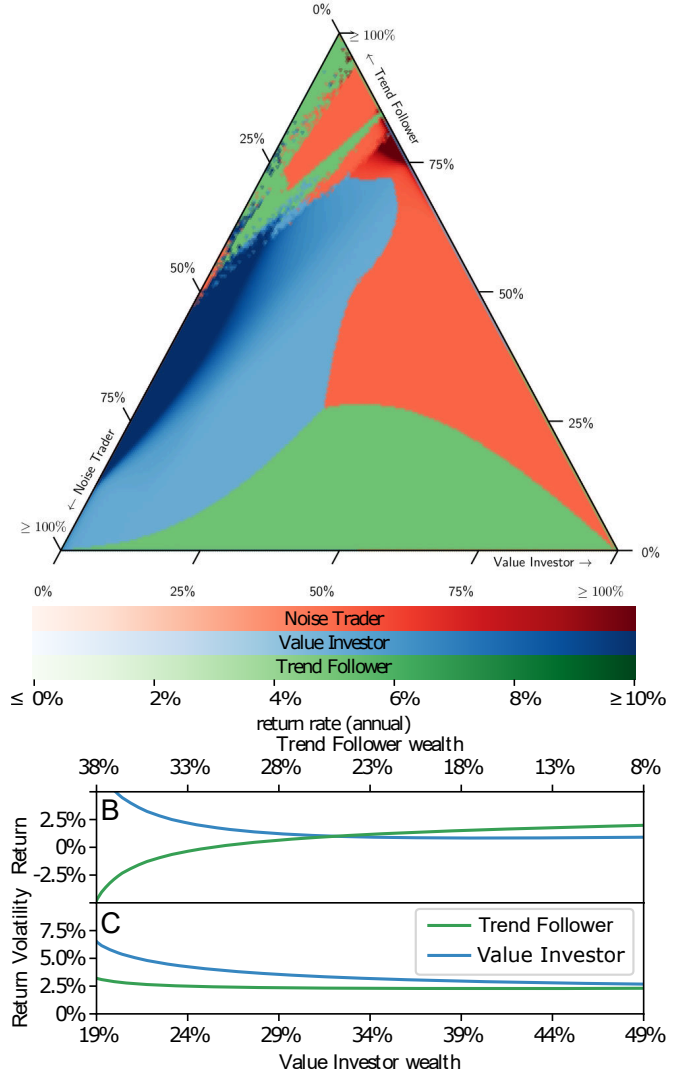


Fig. 2. The profitability of the dominant strategy in the wealth landscape. Panel A is a ternary plot which displays the returns achieved by the strategy with the largest return. The axes correspond to the relative wealth invested in each strategy. The top corner is pure trend followers, the left corner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

When the core ideas in this paper were originally introduced in reference (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 non-linearities that are important for understanding market ecology.

With market clearing there is strong density dependence. This is evident in Figure 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 $\mathbf{W} = (W_{\text{NT}}, W_{\text{VI}}, W_{\text{TF}})$. We somewhat arbitrarily let the total wealth $W_T = W_{\text{NT}} + W_{\text{VI}} + W_{\text{TF}} = 3 \times 10^8$. 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.

A quantitative snapshot of the average returns and volatility is given in Figure 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.

D. 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 Figure 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 \mathcal{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 Figure 3A. Tracking a few individual trajectories, as we do in Figure 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 Figure 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 years, as shown in Figure 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 Figure 3B spans 200 years 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 reference (3). In the ideal case of a stationary market and I.I.D. normally distributed returns, the time required to detect excess performance ΔS with a statistical significance of s standard 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 $\mathcal{S} = 0.5$. For a strategy whose annualized Sharpe ratio is superior by $\Delta S = 0.1$, a 20% improvement over a buy and hold, roughly 400 years are required to confirm its superior performance with two standard deviations. Furthermore, as shown in reference (24), 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 Supplementary Materials (section D). 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 Figure 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

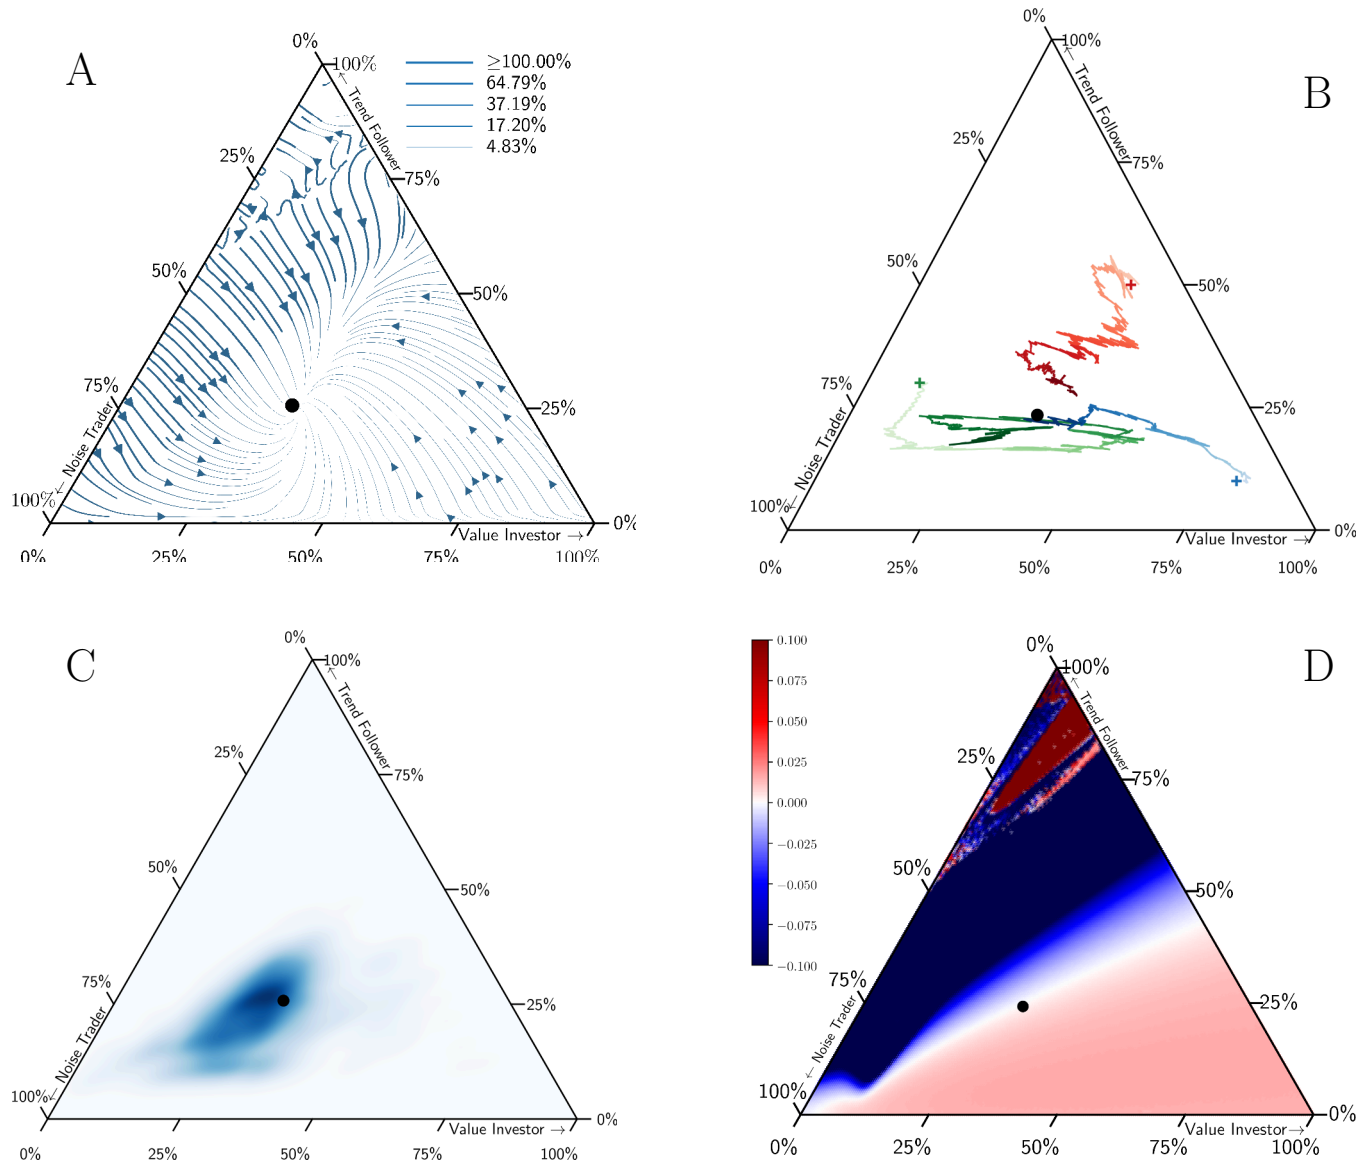


Fig. 3. Profit dynamics as a function of wealth. **A** shows how wealth evolves *on average* through time under reinvestment. The intensity of the color denotes the rate of change. **B** shows sample trajectories for a few different initial values of the wealth vector, making it clear that the trajectories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

process. The wealth of trend followers fluctuates, even at very long times, and consequently the degree of autocorrelation in price returns fluctuates as well.

E. Community matrix. The community matrix is a tool used in ecology to describe the pairwise effects of the population of species j on the population growth rate of species i (25, 26). 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 , i.e. $\pi_i = \lim_{T \rightarrow \infty} (\prod_{t=1}^T \pi_i(t) - 1)^{1/T}$. 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 the predator and j the prey; and if both G_{ij} and G_{ji} are positive, then there is a mutualistic interaction (27).

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 Figure 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 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 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 (though 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.

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%

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%

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 to 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 assumption 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.

F. 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 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 \mathbf{W} to those when strategy j is removed, i.e. when $W_j = 0$ but all 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.

Equations (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 the Supplementary Material. 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 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 Figure 4.




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 (NT, VI, TF), 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 grey region there are cycles where the trophic levels become undefined. The black dots correspond to samples of the wealth vector after 200 years, as shown in Figure 3(c). The system spends most of its time in the grey and red zones.

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 the noise traders exploit trend followers (who generate excess volatility), to complete a cycle. When this happens equation (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 Figure 4 to Figure 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.

G. 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*, i.e. deviations of prices from fundamental values. Both in real markets (2) and in the agent-based models mentioned earlier (including our model

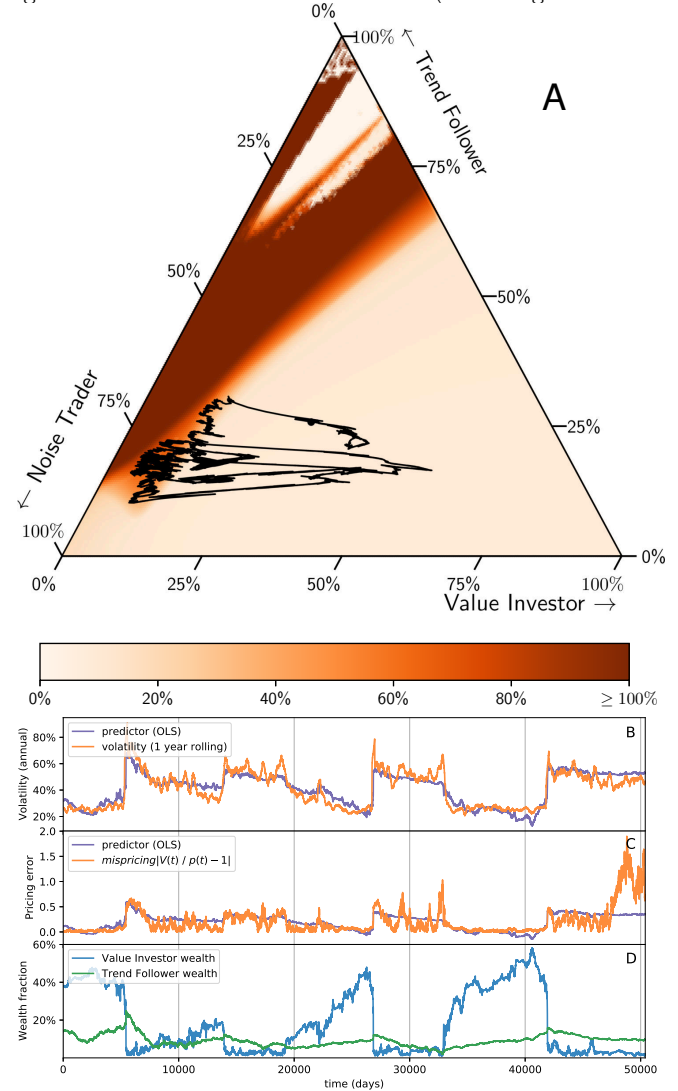


Fig. 5. How ecosystem dynamics cause market malfunctions. Panel A gives a color map of the price volatility over the wealth landscape; the volatility is low and constant throughout the lower right part of the diagram, where the system spends most of its time, but there is a high volatility region running across the upper left. A sample trajectory spanning 200 years 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. Panel B shows the volatility of this trajectory as a function of time, plotted against the predicted volatility (see Eq. (12)). Panel C shows the actual mispricing plotted against the predicted mispricing. Panel D shows the wealth of the value investors and trend followers.

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 (10, 11, 28) and leverage (29). 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.

Figure 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.

Figure 5A shows a sample trajectory that begins at the efficient equilibrium and spans 200 years. This sample is representative of trajectories that fluctuate around the equilibrium point indefinitely. We choose a 200-year 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.

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

volatility		$R^2 = 0.79$	observations: 50,397
independent variable	coefficient	t-statistic	
noise trader	2.4	10	
value investor	-68	-249	
trend follower	107	169	
mispricing		$R^2 = 0.33$	observations: 50,397
independent variable	coefficient	t-statistic	
noise trader	-0.15	-18	
value investor	-1.02	-107	
trend follower	1.5	69	

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 Figure 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 Figures 5B and 5C we compare a time series of the predicted volatility and predicted mispricing against the actual values. The predictions are very good.

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

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 Figure 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, there can nonetheless 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 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 (30, 31). 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 is difficult for most researchers to obtain.

Regulators potentially have access to the balance sheets of all market participants, which can allow them to track the ecology of the markets they regulate in detail. 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,

Table 4. This table details the balance sheet items used by all funds. All securities use the most recent market value in valuation.

Assets		Liabilities	
cash	C	Equity	
		capital	K
margin	M	Debt	
		loans	L
trading securities	S^+	borrowed securities	S^-

e.g. 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.

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 U.S. 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* †. The model and code to run the experiments in this paper is available at <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 TAILOR, a project funded by EU Horizon 2020 research and innovation programme under GA No 952215.

Table 5. This table provides a listing 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 day	dividend autocorrelation lag
f	1	reinvestment rate
ρ	$1 - 6^{-252} \sqrt{\frac{1}{2}}$	noise trader mean reversion rate
σ^{NT}	20% annual	noise trader volatility
$\bar{\lambda}_{\text{NT}}, \bar{\lambda}_{\text{VI}}, \bar{\lambda}_{\text{TF}}$	1, 8, 1	leverage limit
$C_{\text{NT}}, C_{\text{VI}}, C_{\text{TF}}$	5, 10, 4	signal scale

- DM Cutler, JM Poterba, LH Summers, What moves stock prices? *The J. Portfolio Manag.* **15**, 4–12 (1989).
- RJ Shiller, *Market volatility*. (MIT Press), p. 463 (1989).
- JD Farmer, Market force, ecology and evolution. *Ind. Corp. Chang.* **11**, 895–953 (2002).
- B LeBaron, Building the Santa Fe artificial stock market. (2002).
- A Lo, The adaptive market hypothesis: market efficiency from an evolutionary perspective. *The J. Portfolio Manag.* **30**, 15–29 (2004).
- T Hens, KR Schenk-Hoppé, *Handbook of financial markets : dynamics and evolution*. (North Holland), p. 584 (2009).
- JD Farmer, S Skouras, An ecological perspective on the future of computer trading. *Quant. Finance* **13**, 325–346 (2013).
- N Beale, R Gunton, K L Bashe, H S Battey, R S MacKay, Dynamics of Value-Tracking in Financial Markets. *Royal Soc. Open Sci.* **11**, 1–22 (2019).
- SJ Grossman, JE Stiglitz, On the Impossibility of Informationally Efficient Markets, Technical Report 3 (1980).
- WB Arthur, JH Holland, BD LeBaron, RG Palmer, P Tayler, Asset Pricing Under Endogenous Expectations in an Artificial Stock Market in *The Economy as an Evolving Complex System II*. (CRC Press), pp. 15–44 (1997).
- WA Brock, CH Hommes, A Rational Route to Randomness. *Econometrica* **65**, 1059 (1997).
- T Lux, M Marchesi, Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature* **397**, 498–500 (1999).
- C Chiarella, The dynamics of speculative behaviour. *Annals Oper. Res.* **37**, 101–123 (1992).
- B LeBaron, Empirical Regularities From Interacting Long-and Short-Memory Investors in an Agent-Based Stock Market, Technical Report 5 (2001).
- EF Fama, KR French, Dividend yields and expected stock returns. *J. Financial Econ.* **22**, 3–25 (1988).
- B LeBaron, Agent-based computational finance in *Handbook of computational economics*, eds. L Tesfatsion, KL Judd. (Elsevier) Vol. 2, 1 edition, pp. 1187–1233 (2006).
- MJ Gordon, E Shapiro, Capital Equipment Analysis: The Required Rate of Profit (1956).
- F Black, M Scholes, The effects of dividend yield and dividend policy on common stock prices and returns. *J. Financial Econ.* **1**, 1–22 (1974).
- WA Brock, CH Hommes, Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *J. Econ. Dyn. Control.* **22**, 1235–1274 (1998).
- JD Farmer, S Joshi, The price dynamics of common trading strategies. *J. Econ. Behav. Organ.* **49**, 149–171 (2002).
- B LeBaron, Wealth dynamics and a bias toward momentum trading. *Finance Res. Lett.* **9**, 21–28 (2012).
- CS Asness, et al., Value and Momentum Everywhere. *The J. Finance* **LXVIII** (2013).
- JPP Bouchaud, et al., Black Was Right: Price Is Within a Factor 2 of Value, Technical report (2017).
- D Cherkashin, JD Farmer, S Lloyd, The reality game. *J. Econ. Dyn. Control.* **33**, 1091–1105 (2009).
- RM May, Will a large complex system be stable? *Nature* **238**, 413–414 (1972).
- 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).
- JL Bronstein, Conditional outcomes in mutualistic interactions. *Trends Ecol. Evol.* **9**, 214–217 (1994).
- R Franke, F Westerhoff, Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *J. Econ. Dyn. Control.* **36**, 1193–1211 (2012).
- S Thurner, JD Farmer, J Geanakoplos, Leverage Causes Fat Tails and Clustered Volatility (2011).
- A Kirilenko, AS Kyle, M Samadi, T Tuzun, The Flash Crash: High-Frequency Trading in an Electronic Market. *The J. Finance* **72**, 967–998 (2017).
- F Musciotto, L Marotta, J Pilo, RN Mantegna, Long-term ecology of investors in a financial market. *Palgrave Commun.* **4**, 92 (2018).
- LW Johnson, DR Scholz, On Steffensen's Method, Technical Report 2 (1968).
- D Cooke, D Kuhn, Imperial College London MEng Final Year Project Investigating General Equilibria in Artificial Stock Markets, Technical report (2012).
- RJ Hogan, Fast reverse-mode automatic differentiation using expression templates in C+++. *ACM Transactions on Math. Softw.* **40**, 1–16 (2014).
- J Chevallier, G Ellison, Risk taking by mutual funds as a response to incentives. *J. Polit. Econ.* **105**, 1167–1200 (1997).

* The Electronic Data Gathering, Analysis, and Retrieval system (EDGAR) <https://www.sec.gov/edgar.shtml>. Securities and Exchange Commission

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

Supplementary Materials

A. Statistical Properties of the Model

Dividends. The dividend process has the following properties. Consider the auxiliary autocorrelated process $U(t)$ from equation 1, which is a modified AR (1) model.

$$\begin{aligned}\text{Var}(U(t)) &= \text{Var}(\omega U(t-2) + (1-\omega^2)Z(t)) \\ &= \omega^2 \text{Var}(U(t-2)) + \text{Var}((1-\omega^2)Z(t)) \\ &= 1\end{aligned}\quad [13]$$

Therefore the dividend process underlying the autocorrelated process $U(t)$ has variance 1. This is useful in the construction of the dividend process, as from this it follows that the dividend process' log returns have a variance of σ .

We now investigate the autocorrelation of the process $U(t)$. We proceed by induction, and it is useful to know that in the base case, $\text{Cov}(U(t), U(t)) = 1$. The lag τ autocorrelation function of the process $U(t)$ is given by

$$\begin{aligned}\text{Cov}(U(t), U(t-\tau)) &= \text{Cov}(\omega U(t-2) + (1-\omega^2)Z(t), U(t-\tau)) \\ &= \omega \text{Cov}(U(t-2), U(t-\tau)) + (1-\omega^2)\text{Cov}(Z(t), U(t-\tau)) \\ &= \omega \text{Cov}(U(t-2), U(t-\tau)) \\ &= \begin{cases} \omega^{\tau-1} & \tau > 1 \\ 0 & \tau = 1 \end{cases}\end{aligned}\quad [14]$$

Market Prices. We illustrate aspects of market prices generated by the model in figure 6.




Fig. 6. A brief summary of aspects of market prices. Panel A: market prices track fundamental over long periods, are often within a factor of two, yet are noisy. Panel B: the returns distribution of market prices has heavy tails, as is evidenced by the deviation of the empirical distribution (markers) from the return quantiles expected by a process with normally distributed returns (line). Panel C: market prices display cluster volatility, with large price moves likely followed by large moves resulting in clustering in time.

B. Market Clearing

Prices are set by a price setter who chooses prices such that demand and supply match as close as possible. The excess demand of agent a for the stock is defined in equation 3. The market excess demand curve for the stock is the aggregate of the excess demand of all agents a . As in the classical Walrasian setting, the price setter seeks to match demand and supply, so that aggregate excess demand is zero for each investment, by finding a root of the market excess demand curve.

$$\sum_{a \in \mathcal{A}} E_a(p) = 0. \quad [15]$$

However, if no solution is found through the root-finding process, we must fall back to a heuristic that seeks for the best solution that

only partially clears the market. We interpret the goodness of a solution as the extent to which the solution minimizes demand and supply *mismatch*. We here use the square of excess demand, and this way the root-finding problem is transformed into the corresponding minimization problem:

$$\begin{aligned}\underset{p}{\text{minimize}} \quad & \left(\sum_{a \in \mathcal{A}} E_a(p) \right)^2 \\ \text{subject to} \quad & \text{finite } p.\end{aligned}\quad [16]$$

In practice, we rarely encounter situations where no clearing price is found, as they only occur in the chaotic regions where trend followers have amassed most wealth, and the system does not often visit that region. For univariate root-finding problems, methods such as Steffensen's method (32) can quickly find a root of a continuous function using the derivative of the function. For the types of demand and supply functions we permit, this process is not guaranteed to find the nearest market-clearing price, nor is it guaranteed to find a clearing price at all, since there is no guarantee of existence of such a solution. If the root-finding procedure fails, we fall back on minimizing the clearing error (equation 16). In the spirit of Walras' auction, we use quasi-Newton methods to gradually move into the direction that decreases the clearing error fastest. Limited-memory BFGS(L-BFGS) is chosen as the minimization algorithm. As with many non-convex optimization problems, convergence is not guaranteed, though the chosen method compares favorably to others (33). The method requires at each iteration the gradient in numerical form. The starting point for the minimization algorithm is the market-clearing prices for the previous time point. Quasi-Newton methods that are widely used in numerical optimization rely on reconstructing a gradient of the function, and an analytic expression for the gradient is often not available. If the method uses a numerical scheme to reconstruct the gradient, such as the central difference method, this requires multiple evaluations of the function, and these schemes are prone to estimation errors. Automatic differentiation is used to obtain both the value, gradient and Jacobian in as few evaluations of the function as possible. We use the software library presented in (34), which uses C++ meta-programming rules to rewrite equation 3 to compute not only the excess demand but also the gradient and Jacobian. With an eye towards future work, we develop a general solution that can deal with multiple assets.

C. Analysis of the trophic levels at different coordinates in the wealth space

The trophic levels at the equilibrium point in Figure 4 are 1.00, 2.00, and 2.99 for the noise trader, value investor and trend follower respectively. These values are obtained from Eqs. (10-11) as follows: First, we compute $\Delta_{ij} = \max[0, \pi_i(W_1, \dots, W_j, \dots, W_N) - \pi_i(W_1, \dots, 0, \dots, W_N)]$ and obtain the values in Table 6. The diagonal elements Δ_{ii} are computed by letting strategy i have infinitesimal wealth while holding the wealth of the other strategies at their original finite values W_j . The returns for all three strategies are higher when their wealth is zero, as indicated by the negative entries on the diagonal.

The off-diagonal entries in the matrix Δ_{ij} can be understood by referring to figure 3A and projecting to the face of the simplex where $W_j = 0$, i.e. the face opposite the vertex where the strategy has 100% of the wealth, and then determining whether the flow lines increase the wealth of strategy i at this point. The noise trader does not increase in wealth when either the value investor or the trend follower are removed. The value investor benefits from the presence of noise traders, but only mildly, as indicated by the positive element $\Delta_{FV,NT} = \Delta_{21}$. When we remove the noise traders the projection of the equilibrium point on the right edge has about 60% value investors and 40% trend followers. Inspecting sample price realizations shows that in this case prices stay close to the fundamental. This means that without noise traders mispricing is small, even if there are abundant trend followers. There is little mispricing for the value investor to exploit, hence the low value.

To understand the third row it is instructive to examine Figure 3D, which shows the daily autocorrelation in returns.

Removing the value investor corresponds to projecting onto the left edge of the simplex, where there is a strong negative autocorrelation for the trend follower to exploit (hence the trend follower's returns go up). There is also slight increase in the size of the autocorrelation when the noise trader is removed, but only slight, hence this term is small. This is not surprising given that the noise trader is much less efficient than the value investor in translating autocorrelation of the dividend payments into the price. The returns of the trend follower are thus mainly dependent on the value investor. This is visible in Tables 6, where approximately 99% of the return of the trend follower at the equilibrium point is attributable to the value investor.

The fact that the trophic levels at the equilibrium point take values 1.00, 2.00 and 2.99 is mainly by design of the simulation. During the development of the model we fixed the wealth distribution to $\mathbf{W} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and chose archetypal strategies that might interact in the way as described above. We then chose parameters (signal scale, leverage limit) such that the three strategies were all somewhat profitable, and found the nearby equilibrium point where profitability was equal for all three strategies.

This changes as we move away from the equilibrium point. Consider, for example, the point $\mathbf{W} = (20\%, 66\%, 14\%)$. Here, the relationship between the trend follower and noise trader has changed sign. The trend follower is much smaller, and therefore is not able to exploit the noise trader (and value investor) as much. This is visible in Figure 3D, which makes it clear that this point has significant positive autocorrelation. The result is a cycle of predation: value investors prey on noise traders, trend followers prey on value investors, and the noise traders in turn obtain a benefit of 0.78% due to the presence of trend followers. The result is that the trophic levels are no longer defined. The values are summarized in Table 7.

Suppose we move into the region of the simplex where prices begin to diverge from fundamental values, such as the point $\mathbf{W} = (50\%, 5\%, 45\%)$. Here, trend followers start to reinforce their own profitability, as an increase in trend followers pushes the market prices away from fundamental value even faster. The result is large oscillations around the fundamental value, during which the value investor quickly reaches its leverage limit and must close its position at a loss. The noise trader in turn benefits from the trend follower, because it is less aggressive and therefore does not have to sell at a loss, but is able to profit a small amount from the mispricing over very long time spans. The resulting trophic levels are (5.34, 1.00, 4.34) respectively, and these values are increased as the trend follower is self-reinforcing.

Δ	NT	FV	TF	A	NT	FV	TF
NT	-2.21%	-5.61%	-2.48%	NT	0.00	0.00	0.00
FV	0.24%	-2.64%	-7.12%	FV	1.00	0.00	0.00
TF	0.06%	7.20%	-9.78%	TF	0.01	0.99	0.00

Table 6. The left hand table is $\Delta_{ij} = \max[0, \pi_i(W_1, \dots, W_j, \dots, W_N) - \pi_i(W_1, \dots, 0, \dots, W_N)]$ around the equilibrium point. The right hand side is normalized as indicated in Eq. (11). See section C for an interpretation of these values.

Δ	NT	FV	TF	A	NT	FV	TF
NT	-0.64%	-6.61%	0.78%	NT	0.00	0.00	1.00
FV	1.77%	-0.26%	-0.20%	FV	1.00	0.00	0.00
TF	-0.36%	3.65%	5.02%	TF	0.00	0.42	0.58

Table 7. The entries of the matrix Δ_{ij} and its normalization A_{ij} at the coordinate $\mathbf{W} = (20\%, 66\%, 14\%)$

Δ	NT	FV	TF	A	NT	FV	TF
NT	-4.64%	-2.99%	2.30%	NT	0.00	0.00	1.00
FV	-0.66%	-1.29%	-9.14%	FV	0.00	0.00	0.00
TF	-0.48%	4.40%	10.29%	TF	0.00	0.30	0.70

Table 8. The entries of the matrix Δ_{ij} and its normalization A_{ij} at the coordinate $\mathbf{W} = (50\%, 5\%, 45\%)$

D. Long-term dynamics

We investigate the distribution of system states $\mathbf{W}(t)$ at time t , as $t \rightarrow \infty$. Let $\mathcal{P}(t)$ and $\mathcal{W}(t)$ be discrete probability distributions defined on the space \mathbf{W} . \mathcal{W} will be the distribution of the system states of the model as it is presented in the main text. $\mathcal{P}(t)$ will be the distribution under a change of model parameters, defined by a series of experiments that are defined later.

Probability distributions such as \mathcal{W} will be estimated using N Monte-Carlo samples

$$\mathcal{W}_t \approx \frac{1}{N} \sum_{w \in \mathbf{W}} 1_{W(t)=w}.$$

A measure to describe the difference between the asymptotic distribution \mathcal{W}_∞ and the distribution over states under an altered parameter choice \mathcal{P} is the Kullback-Leibler divergence, in this case given by

$$D_{\text{KL}}(\mathcal{W}_\infty || \mathcal{P}) = - \sum_{w \in \mathbf{W}} \mathcal{W}_\infty(w) \log \left(\frac{\mathcal{P}(w)}{\mathcal{W}_\infty(w)} \right).$$

A property of this measure is that identical distributions have a distance of $D_{\text{KL}} = 0$, and that the measure is unbounded; at our initial condition we find that $D_{\text{KL}} \approx 0.98$.

Numerical Computation. Assume observations w_1 and w_2 are distributed as multivariate Gaussians $\mathcal{N}_1 = (\mu_1, \Sigma_1)$ and $\mathcal{N}_2 = (\mu_2, \Sigma_2)$

$$D_{\text{KL}}(\mathcal{N}_1, \mathcal{N}_2) = \frac{1}{2} \left[\text{tr}(\Sigma_2^{-1} \Sigma_1) + (\mu_2 - \mu_1)^T \Sigma_2^{-1} (\mu_2 - \mu_1) - k + \log \left(\frac{|\Sigma_2|}{|\Sigma_1|} \right) \right], \quad [17]$$

where k is the number of dimensions ($k = 3$ in the case of three investment strategies).

Reinvestment Rate. In the literature, a positive relationship between fund performance and subsequent investment by external investors is found (35). This relationship however is non-linear but convex, meaning that the higher the funds performance, the larger the inflow of new capital over the next period. There are other factors that play a role, such as that new funds have different investor dynamics than more established funds.

We simplify this complex relationship by defining f to be the reinvestment rate, meaning the rate at which external investors to the fund invest additional capital when the fund performs well, and withdraw when the fund under-performs the benchmark rate of return r : $W(t+1) = W(t)(1 + f\pi(t))$.

This is achieved by depositing additional cash into the fund when $\pi(t) > 0$, and decreasing cash and stock when $\pi(t) < 0$.

The study (35) shows that for established funds, when a fund earns an annual return of 25%, the fund sees an additional inflow of 50%, corresponding with $f = 3$. However, this is an estimate of the highest reinvestment ratio, as for most cases the inflows and outflows are much smaller, as the authors showed the relationship to be convex in fund performance.

Determining the time to convergence. We approximate the asymptotic distribution by taking the average of multiple samples of long runs of the simulation. For this we choose a horizon of $T = 200$ years. We choose initial conditions different from the observed equilibrium state, so that we can tell apart transient states and stochastic steady states more clearly. The initial conditions are endowments distributed uniform in the strategy simplex.

The chosen horizon of 200 years is sufficient to bring down the distance by four fifths, after which random fluctuations dominate.

Results: Reinvestment Rate. We investigate how the reinvestment rate influences the time to converge to the steady state. For this, we vary the reinvestment rate f ranging from the values $f = 0.1, 0.3, 1.0, 3.0$ to 5.0 . The first thing that stands out is that when the reinvestment rate is 5, the system does not converge. Inspection reveals that this happens because for this parameter choice, the system almost always ends up with one of the funds in default, meaning the system moves to a point on the boundaries of the simplex far away from the equilibrium point. For smaller choices of f , we observe eventual convergence. Because the reinvestment rate affects the equilibrium state, we only measure the Kullback-Leibler divergence to the long-term $t = 100Y \dots 200Y$ distribution of trajectories with the same reinvestment rate. However, for each test we initialize the system to the same uniform distribution. The uniform distribution will have a different KL-divergence value to the different equilibrium states for each choice of f .




Fig. 7. KL-Divergence between the Monte-Carlo distribution estimate at time t versus the long-term distribution sampled over the period $t = 100Y \dots 200Y$, tested for different choices of the reinvestment rate f (see: D). For each parameter choice, we plot the divergence between the distribution at time t to their respective long-term distribution. The choice of the parameter causes the steady-state distribution to change, and therefore the KL-divergence of the distribution of endowments at $t = 0$ to this steady state has different values.

We learn that a higher rate of reinvestment leads to a mildly faster rate of convergence, however, it also leads to dynamics with a higher variance. Thus, large reinvestment rates also lead to destabilize the system. When the reinvestment rate is set to large values with $f \gg 3$, it becomes increasingly likely that the market price changes on a single day are enough to bankrupt one of the funds. For example, 20% returns on a single day for a single stock are not uncommon on the time scale of a century and a reinvestment rate of five would be enough to bankrupt a fund that has the wrong position in that single price move.

Results: Role of noise trader noise level. We repeat the same experiment as before, but now we fix $f = 1.0$ and instead vary the noise trader volatility $\gamma = 0.05, 0.20, 0.40$. The result is visualized in figure 8, where we see that higher noise trader volatility leads to slower convergence to the steady state, and slightly larger variance during the steady state.

Results: Role of fundamental value volatility. In this experiment, we alter the volatility of the dividend process by changing the parameter σ from $\sigma = 0.025, 0.10, 0.40$ in equation 1. There appears to be no significant effect on convergence rates, and the steady state is also not altered, as is evident from figure 9 where the trajectories of all Monte Carlo experiments are close.




Fig. 8. KL-Divergence between the state of the system at time t versus the long-term distribution $t = 100Y \dots 200Y$, tested for different choices of the noise trader volatility γ (see: D).




Fig. 9. KL-Divergence between the state of the system at time t versus the long-term distribution $t = 100Y \dots 200Y$, tested for different choices of the dividend volatility σ (see: D).

E. Kelly Criterion

The Kelly criterion is used to determine the optimal fraction x^* of wealth to invest on a risky bet. Suppose we have chosen a fraction x , then the wealth of our portfolio is expected to grow over one time step ahead from time t from the fraction x invested in the risky stock and the remainder $(x - 1)$ invested in the risk-free money market instrument earning the interest rate r . It follows that the wealth under investment of fraction x of total wealth $W(t)$ into the stock is given by

$$W(t+1) = xW(t) \left(\frac{p(t+1) + D(t+1)}{p(t)} - 1 \right) + (1-x)W(t)r.$$

The Investment. Recall the fundamental value equation.

$$V(0) = \sum_t \frac{D(t)}{(1+k)^t}.$$

Suppose that the autocorrelation in dividend values is small enough to be ignored, and that the market is near efficiency. Then, by the Gordon growth model

$$p(0) \approx V(0) = \frac{D(1)}{k-g},$$

from which we learn that the relation between dividend payment and market price under those assumptions is linear. This implies if dividend's returns are normally distributed, we have normal distributed log returns $\log p(t+1)/p(t)$ on the stock. Thus, near the efficient point, $\mathbb{E} \left[\frac{D(t+1)}{p(t)} \right] \propto k - g$. From this approximation we learn that the returns on the stock and dividend should be highly correlated, under the assumption of little autocorrelation and fundamental-value efficient markets.

Dividend Process. For convenience, we repeat here the expectation of the dividend process.

$$\begin{aligned} \mathbb{E} \left[\frac{D(t+1)}{D(t)} \right] &= \exp \left(g - \frac{1}{2} \sigma^2 + \sigma \mathbb{E}[U(t+1)] \right) \\ &= \exp \left(g - \frac{1}{2} \sigma^2 + \sigma \omega U(t-1) \right) \end{aligned} \quad [18]$$

Recall that the expectation of the future value of the stochastic process $\sigma\mathbb{E}[U(t+1)]$ has expectation $\sigma\omega U(t-1)$, where $U(t-1)$ can be inferred by the investor from $D(t-1)$ having correctly estimated g , σ and ω .

Solution Method. A fully rational Kelly bettor would model possible other participants in the market and their likely influence on the returns of the investment. This quickly leads to a complicated and intractable model. We will instead introduce an approximation, based on the insight that price returns correlate highly with the dividend returns as the preceding sections showed.

Recall we use π to denote returns to the fund, and x is the fraction invested in the stock. We assume that a solution can be found of the form

$$\begin{aligned} \pi(x) = x\mathbb{E}[\exp(g - \frac{\sigma^2}{2} + \sigma\omega U(t-1)) \\ \cdot \exp(\mu_Y - \frac{\sigma_Y^2}{2} + \sigma_Y O(t))] \\ + (1-x)r. \end{aligned} \quad [19]$$

In the first term in the product we recognize the dividend process from equation 18. The second component is due to an approximation Y of the stock returns. This process Y has two parameters μ_Y , the drift, and σ_Y , the volatility, that are dynamically estimated by the Kelly investor. The process is defined as:

$$\frac{Y(t+1)}{Y(t)} = \exp\left(\mu_Y - \frac{1}{2}\sigma_Y^2 + \sigma_Y O(t)\right) \quad [20]$$

with $O(t)$ a standard normal variable. To summarize, the approximation that we make is that

$$\log\left(\frac{p(t+1)}{p(t)} - 1 + \frac{D(t+1)}{p(t)}\right) \approx \log\left(\frac{D(t+1)}{D(t)} \cdot \frac{Y(t+1)}{Y(t)}\right).$$

We test whether this approximation yields the optimal investment fraction in a Monte-Carlo experiment with dividends and prices from the main model. The fraction given by the approximation will be x^* , and we will compare this with other fractions x . In Figure 10 we plot on the x-axis the different fractions invested x/x^* , and on the y-axis the expected returns. The figure shows that the highest mean returns are obtained near $x^* = x$, meaning the approximation is close to optimal.

Kelly Strategy. The Kelly bettor maximizes log returns in equation 19. This gives the following optimization problem:

$$\begin{aligned} \max_x \pi(x) = x(g + \mu_Y + \sigma\omega U(t-1) - r - 2\rho\sigma\sigma_Y) - \frac{1}{2}(x\sigma)^2 - \frac{1}{2}(x\sigma_Y)^2 \\ \text{subject to } -\bar{\lambda} \leq x \leq \bar{\lambda}, \end{aligned}$$

wherein ρ is the estimated correlation coefficient between the market prices and dividends, and $\bar{\lambda}$ is the leverage limit as in the main text. This is maximized when:

$$x^* = \max\left(-\bar{\lambda}, \min\left(\bar{\lambda}, \frac{g + \mu_Y + \sigma\omega U(t-1) - r - 2\rho\sigma\sigma_Y}{\sigma^2 + \sigma_Y^2}\right)\right). \quad [21]$$

We will assume the Kelly bettor has no information other than the dividend announcements and market prices. Therefore, the Kelly bettor will have at time t an estimate of the parameters of equation 21 based on the public information up to time t .

Survival. In previous studies on markets with different constraints than ours (6), it is shown that the Kelly criterion strategy is a survival strategy. In our discrete time model, we observe that the Kelly criterion is indeed highly likely to survive for long time spans, but still it is possible that the Kelly bettor fails. One scenario in which the Kelly bettor, using equation 21 to invest, can go extinct, is when for a long run market prices increase, followed by a sudden drop. During the build-up phase, the Kelly bettor estimates large values for g and μ_Y , and aggressively uses leverage. Depending on




Fig. 10. A Monte-Carlo experiment illustrates that the approximate Kelly strategy derived in section E is numerically close to the optimal strategy.




Fig. 11. This plot shows the estimated probability that a fund with a given strategy still has positive wealth at the time t , out of a Monte-Carlo experiment with 5000 samples.

how the investor updates the parameter estimates, a sudden drop in prices and or dividends may happen fast enough that the Kelly bettor still maintains positive leverage, leading to excessive losses and bankruptcy. The different possibilities for failure of one or more strategies also indicate that the system is more unstable with the strategy added (see figure 11).