Ca' Foscari University of Venice

Department of Economics

Working Paper

Paolo Pellizzari

Learning Whether to Be Informed in an Agent-Based Evolutionary Market Model

ISSN: 1827-3580
No. 03/WP/2024

Working Papers
Department of Economics
Ca' Foscari University of Venice
No. 03/WP/2024
ISSN 1827-3580

Learning Whether to Be Informed in an Agent-Based Evolutionary Market Model

Paolo Pellizzari
Ca' Foscari University of Venice

Abstract
Can traders in a financial market learn whether to be informed and which information to use in their demand for risky assets? We describe in this paper an agent-based model where heterogeneous traders seek short-term profits and differ in their choices to use or discard some signals. In the model, a vector of fresh news/signals is available at every period and some (but not all) the signals affect the stochastic payoff of the stock.
Under an evolutionary dynamics favouring higher myopic returns we find that, in equilibrium, traders mostly end up in either discarding all signals or being (perfectly) informed using all the relevant signals (paying the related costs). Moreover, the rate of use of information strongly depends on the "complexity" of the market: an excessively large abundance of signals to be screened or a high volatility of the market, result in large shares of passive agents who overestimate the market's risk; conversely, low market complexity is associated with a more intense use of information and aggressiveness of informed traders. Evolutionary models and Agent-based models and Information in financial markets

Keywords
Financial markets w information, agent-based models, evolutionary game theory, equity premium puzzle

JEL Codes
G17, D83, D91

Address for correspondence:
Paolo Pellizzari
Department of Economics
Ca' Foscari University of Venice
Cannaregio 873, Fondamenta S.Giobbe
30121 Venezia - Italy
e-mail: paolop@unive.it

The Working Paper Series
is available only on line
(http://www.unive.it/pag/16882/)
For editorial correspondence, please contact:
wp.dse@unive.it

Department of Economics
Ca' Foscari University of Venice
Cannaregio 873, Fondamenta San Giobbe
30121 Venice Italy
Fax: ++39 041 2349210

LEARNING WHETHER TO BE INFORMED IN AN AGENT-BASED EVOLUTIONARY MARKET MODEL

PAOLO PELLIZZARI

Abstract. Can traders in a financial market learn whether to be informed and which information to use in their demand for risky assets? We describe in this paper an agent-based model where heterogeneous traders seek short-term profits and differ in their choices to use or discard some signals. In the model, a vector of fresh news/signals is available at every period and some (but not all) the signals affect the stochastic payoff of the stock.

Under an evolutionary dynamics favouring higher myopic returns we find that, in equilibrium, traders mostly end up in either discarding *all* signals or being (perfectly) informed using *all* the relevant signals (paying the related costs). Moreover, the rate of use of information strongly depends on the "complexity" of the market: an excessively large abundance of signals to be screened or a high volatility of the market, result in large shares of passive agents who overestimate the market's risk; conversely, low market complexity is associated with a more intense use of information and aggressiveness of informed traders. Evolutionary models and Agent-based models and Information in financial markets

1. Introduction

Many investors acquire information on their investments and try to make some sense of the markets' situations and prospects. We refer, in what follows, especially to "fundamental" information regarding what is typically believed in economic textbooks to be relevant to explain to some extent the movements of equity prices such as, say, interest rates, inflation, GDP growth in developed end emerging markets, geopolitical events, international imbalances and breaking news on firms or events of (potential) broad impact[1].

[1]We discard "technical" information, mostly derived from time-series and historical data. Many traders may use such "information" but the model has little to say in this respect as no past observation is used, see [Lo et al., 2000] for an evergreen examination of technical trading.

As an example, on April 29th 2023 the most common Italian financial newspaper, "Il Sole 24 ore" printed that:

- Wall Street "bets" on (forthcoming) interest rates cuts of 300 basis points (first page);
- Taipei has denounced the intrusion of 38 Chinese military airplanes in its space. Chinese authorities pleaded that they were "monitoring" one US military fighter flying in the area (first page);
- Jerome Powell, chairman of the FED, declared "We will not reduce the interest rates in 2023" (page 7);

The previous news (or signals) have the capability to provide valuable investment insights and, yet, it is hard to pick the most significant or decide whether to use them all. It is even more difficult to unambiguously interpret the news. There are basically opposite statements on the trend of US interest rates: an hypothetical personification of Wall Street expects a drop of 3%, but Powell stated this is not going to happen. Well, at least in 2023! In principle, exploiting the lack of a clear timeline, both news could be correct as rates' cuts may come in 2024. Indeed, there is a good joke stating that wise forecasters should never provide a number *and* a date... Geopolitical strained relations involving China, Taipei and USA are hinted at, with no clear implications on asset prices. Whether and how to use such information is an interesting, as well as far from trivial, issue.

Many analogous examples can be drawn, virtually any day, from other financial newspapers or websites, official and informal reports by public and private institutions, central banks' statements, and various intelligence from advisory firms and respected professionals or gurus. We present in this model a stylised depiction of investors who are similarly flooded with information and have little guidance on how to use such body of insights. They try to use and interpret a stream of signals in order to decipher how to behave, most of them attempt to select relevant information, weight it properly and discard irrelevant news, are aware there is no easy recipes and are willing to imitate strategies or practices put in place by others.

We assume that traders are boundedly rational and learn to change their investment strategy by imitating other agents who had better (i.e., higher) returns. In other words, they copy the pattern of use of information made by more successful peers. A strategy is a vector of bits (bit-string) where 1 means that the news is used and 0 means that it is discarded. We include mutation allowing a small fraction of agents to occasionally flip one of their bits at random. In our simple setup, given their strategies, all agents have to decide in every period is how much to buy of a risky asset (while the rest of their endowments will be put in a safe bond paying a constant interest rate). In this sense, there is a one-to-one correspondence between a strategy (a bit-string listing signals to be used or discarded) and a demand function for the risky asset (as the demand depends on the used information/signals).[2] Broadly speaking, demands are based on perceived mean and variance of returns (the demand of equity is directly/inversely proportional to the mean/variance). Once all agents submit their demands, a noisy clearing price can be computed in this market and all transactions will occur at this local-in-time equilibrium price, realising gains and losses. Actual returns are also used to assess the quality of the strategies and fuel learning through a very simple mechanism: couples of agents will be matched, they compare the realised returns and the worst performer copies the strategy (or demand function, if you wish) of the best performer, beginning to use it in the following period when everyone's endowment is replenished[3].

Having defined a population of agents, a game that is repeatedly played by traders and a process to revise old strategy (or learn better ones), our agent-based model can be interpreted as a canonical evolutionary model. Such models were first introduced in biology, where

[2]Admittedly, the agents in our model learn in a very basic way, as they have no memory or expectations and update their behaviour based on a single random match. A discussion of more sophisticated reinforcement learning approaches (with an extended bibliography) is in [Huang et al., 2024], where a form of collective intelligence is built to maximise returns. In contrast, we assume agents are selfish and myopic.

[3]It is useful to add to the description of what our agents do, a list of things they do not do: they do not explicitly maximise any utility function, they do not have memory, they do not search in a set of possible alternative strategies or, if they do so, they may need several periods in which they compare the outcomes with a single strategy, they do not try to anticipate the equilibrium price based on the shares of strategies in the population, they do not save or accumulate wealth strategically. In a nutshell, they keep a strategy till they stumble on concrete evidence that someone else makes higher returns and occasionally flips some bits.

genotypes are inherited and not chosen by individuals, [Smith, 1982], but were increasingly applied to social and economic environments where strategies are selected consciously in such a way that the ones with greater payoff tend to prevail, see the classic [Sandholm, 2010], or [Newton, 2018] where recent applications are surveyed. We aim at identifying the set of strategies that will thrive in the long-run, simulating the market for many periods and examining the final shares resulting from this evolutionary dynamics.

In brief, our evolutionary model robustly show two main results: first, most traders end up either in being passive (i.e., discard all signals) or being (wholly) informed (i.e., acquire all the relevant signals); second, information usage depends on the "complexity" of the market, as measured by its volatility or by the quantity of the information that traders have to screen and process. Overall, the combination of the above outcomes explains why only some of the relevant information is used by the agents in the market, with the informed traders holding notably riskier positions than passive ones. Several of these findings, driven by short-term evolutionary pressure and inability to deal with the overwhelming complexity of the market, appear to have a realistic flavour that is somewhat difficult to get in standard rational expectation equilibrium models where, for instance, it is difficult to justify why relevant information should be discarded.

The paper is organised as follows. The next section describes the model of the market. An example is used to illustrate the flow of decisions/actions, from strategies/demands to prices/profits and evolution through learning, that are executed in every period. In Section 3, the main results are presented and some conclusive reflections are given in Section 4.

2. The model

Consider M agents in a market with two assets, a risk-less bond with unit cost and payoff $R > 1$ and a risky stock that for a price p_t^*, to be endogenously determined based on the demands of the traders, will pay a random payoff $\tilde{D}_t$. Agents are endowed with a constant periodic endowment $w_{0t} \equiv w_0$, care about return and variance of their portfolio, and have to decide how many units of the stock to buy or sell in any period (what is not spent in stocks will be invested in the bond).

A stream of N news $\theta_t = (\theta_1, \theta_2, ..., \theta_N)_t \in \mathbf{R}^N$ is available to traders in any period $t = 1, ..., T$. We assume that each of the N signals is identically and independently distributed as $N(0, v_\theta)$. The careful reader should notice the lapse between *texts*, such as the ones listed in the introduction, and a vector θ_t of *numeric values*. For simplicity, we just suppose that some judgemental or mechanic procedure translates news (sentences, comments...) into a vector of estimates.

As it will be made clear below, only $S \leq N$ individual signals θ_j will truly affect $\tilde{D}_t$. With no loss of generality, we will assume in this treatment that the first signals $\theta_1, ..., \theta_S$ are relevant. This simplifies the exposition but is unknown to agents who must decide by trial-and-error whether to use signals at all and which ones to use. To keep track of this *learning process*, each traders has a strategy $b_{it} = (b_1, b_2, ..., b_N)_{it}$, where each bit $b_j, j = 1, ..., N \in \{0, 1\}$ denotes if the j-th signal is used: a value of 1/0 means the signal is used/discarded. Equipped with the strategy b_{it} at time t, the demand schedule of the i-th agent is

$$x_{it}(p) = \frac{d + b_{it}\theta'_t - pR}{av_{it}}, \tag{1}$$

where $d > R$, v_{it} is an individual assessment of the variance of D_t, and a can be thought as a risk-aversion coefficient. For tractability, we assume that the deterministic component d of the payoff is exogenous and known to agents and a is constant across them. Given the price p, the demand (1) is, essentially, a ratio of expected excess return of one unit of stock (in excess of pR that could be gained with the bond) over perceived risk. Hence, the expected $\tilde{D}_t^{(i)}$ for agent i is given by

$$d + b_{it}\theta'_t = d + \sum_{j=1}^{N} b_{ijt}\theta_{jt},$$

and depends on which bits are switched-on in the strategy b_{it}. The individual demand function $x_{it}(p)$ is then readily obtained.

A unique transaction price p_t^* is determined matching the demand functions of all the agents and solving

$$\sum_{i=1}^{M} x_{it}(p) = 0. \tag{2}$$

Eq. (2) is linear in p and the solution p^*_t can be numerically computed to allow the agents to determine the realised purchases/sales of the stock, $x^*_{it} = x_{it}(p^*_t | b_{it}, v_{it})$, where we stress that the quantity (plainly) depends also on the individual strategy b_{it} and on v_{it}.

Let the realised stock's payoff be determined "by nature" as

$$\tilde{D}_t = d + b^* \theta'_t + \epsilon_t, \tag{3}$$

where $\epsilon_t \sim N(0, v_\epsilon)$ and, as said before, b^* has the first S bits set to one, $(\underbrace{1, ..., 1}_{S \text{ bits}}, \underbrace{0, ..., 0}_{N-S \text{ bits}})$, so that only S bits out of N affects the payoff.

The profits of agents can now be computed and used to evolve a new population of strategies, or demand functions, moving from $\mathcal{P}_t = \{b_{1t}, ..., b_{Mt}\}$ to $\mathcal{P}_{t+1}$ that differs from the old one because some agents are offered the chance to imitate and mutate their strategy. In detail, denote the profit of i-th agent at t as:

$$w_{it} = x^*_{it} D_t + (w_0 - x^*_{it} p^*_t) R - cost \cdot b_{it} \mathbf{1}',$$

where *cost* is the cost of acquiring or processing one signal, $\mathbf{1}$ is the vector with N ones, $(1, ..., 1)$, and $b_{it}\mathbf{1}'$ is the number of used bits. The amount w_{it} is consumed or spent elsewhere (and, therefore, agents start afresh, in terms of wealth, in the next period). Evolution and competition in the market occurs matching h couples of agents, comparing the profits and changing the strategic profile (b_i, v_i). When, say, traders r and s are matched:

$$\text{If } w_{rt} \geq w_{st} \text{ then } \begin{cases} b_{s,t+1} & = & b_{rt}, \\ v_{s,t+1} & = & v_{rt} \tilde{U}(1-q, 1+q), \end{cases} \tag{4}$$

where $q > 0$ is a small number and $U(a, b)$ is a uniform random variable in $]a, b[$. Formula (4) describes how, if agent r outperformed agent s, the latter copies the strategy of the former and replaces his v_s with a random multiplicatively shocked v_r. Observe that each pair (r, s) of agents is randomly formed and, hence, there is no deliberate attempt to imitate or cherry-pick successful traders. Moreover, the straight comparison of revenues in (4) is entirely justified by the constant risk-aversion parameters. In each period, we also allow a single mutation, flipping a random bit of a random trader's strategy.

Among the 2^N strategies that can be evolved, two turn out to be prominent in the following: we will refer to agents with $b_i = b^*$ as *informed*, in that they come to know and use in their demand all the relevant bits in Eqs. (1,3); we call *passive* the agents with $b_i = (0, 0, ..., 0)$ as they do not use any signal and resort to a very simple constant expected value for D, namely d.

The following example will clarify the mechanics: at (the beginning of) a given time t (omitted in the sequel), with $N = 5, cost = 0.01, d = 1.1, R = 1.01, a = 2$ and $b^* = (1, 1, 1, 0, 0)$, the first and second agents have $v_1 = 0.025, b_1 = (1, 1, 1, 0, 0), v_2 = 0.055$ and $b_2 = (0, 0, 0, 0, 0)$. The first agent is informed, employing all relevant bits and paying a total cost of 0.03 per period, whereas the second agent does not use any information (and has null cost). If the vector of signals is $\theta_t = (0.03, -0.05, 0.07, 0.01, -0.05)$ and, based on the demand schedule of all agents, $p_t^* = 1.11$, accordingly the realised demands are

$$
\begin{aligned}
x_1^* &= \frac{d + b_1\theta_t' - p_t^* R}{a v_1} = \frac{1.1 + 0.05 - 1.11 \cdot 1.01}{2 \cdot 0.02} = 0.578, \\
x_2^* &= \frac{d + b_2\theta_t' - p_t^* R}{a v_2} = \frac{1.1 + 0 - 1.11 \cdot 1.01}{2 \cdot 0.055} = -0.192,
\end{aligned}
$$

meaning that, at the equilibrium price 1.11 prevailing at time t, the informed trader buys 0.578 units of the stock and the passive one sells 0.192 units. Now, ϵ_t is drawn and payoffs can be computed: let the random value be, for instance, $\epsilon_t = -0.1$ so that

$$D_t = d + b^*\theta_t' + \epsilon_t = 1.1 + 0.05 - 0.1 = 1.05.$$

Observe that the payoff is smaller than the price and, as a consequence, net buyers/sellers will experience a loss/gain. Indeed,

$$
\begin{aligned}
w_{1t} &= x_{1t}D_t + (w_0 - x_{1t}p_t^*)R - 3 \cdot cost \\
&= 0.578 \cdot 1.05 + (1 - 0.578 \cdot 1.11) \cdot 1.01 - 0.03 = 0.939, \\
w_{2t} &= x_{2t}D_t + (w_0 - x_{2t}p_t^*)R - 0 \cdot cost \\
&= -0.192 \cdot 1.05 + (1 + 0.192 \cdot 1.11) \cdot 1.01 = 1.024.
\end{aligned}
$$

Hence, due to the (relatively large) negative ϵ and to other "unlucky" events, the first agent happens to record a loss and the second agent a gain. We stress that this outcome holds at time t, due to the values taken by the random variables involved in this period (i.e., θ_t, ϵ_t) and to the shares of different strategies in the population that ultimately contribute to determine *the current* $p_t^* = 1.11$. Other realisations would obviously have produced different w_{1t} and w_{2t} for the two agents in our example. Assume now that in the learning phase the previous two agents are randomly matched: the first trader (painfully) realises that the second trader outperformed him by about 8% in period t and, therefore, he imitates the other strategy and variance, so that $b_{1,t+1} = (0,0,0,0,0)$ and his variance will move to $v_{1,t+1} = v_{2t}\tilde{U}$, for a random $\tilde{U}$. The second trader does not change in any way his strategy/parameters and is ready to begin period $t+1$.

In a standard application of evolutionary game theory, we are interested in looking at the features of the stationary population $\mathcal{P}_t$, for $t \to \infty$.

3. Results

We run 100 simulations for each $N = 5, 10, 15$, using $T = 10000$ periods and setting the other parameters as in Table 1.

Table 1. Values and description of the parameters.

Param	Value	Description	Param	Value	Description
v_ϵ	$\{\frac{2}{100}, \frac{3}{100}, \frac{4}{100}\}$	Variance of ϵ	v_θ	0.01	Variance of θ
R	1.01	Bond yield	d	1.10	Stock yield
a	2	Risk aversion	q	0.1	Variance adjustment
S	$\{1, ..., 5\}$	# of bits	M	1000	# of agents
$w_{0i} \equiv w_0$	1	Endowment	cost	0.01	Cost of information

b_i	Initialised with random bits, then subject to learning and mutation
v_i	Initialised at v_ϵ, then subject to random shocks

Both N and v_ϵ can, to some extent, quantify the complexity of informational extraction in a market: a large N corresponds to situations where agents are exposed to many signals and, especially for small values of S, this means that relatively few relevant signals must be carefully picked (out of the possible N). Besides being a direct measure of the volatility of the payoff, v_ϵ affects the signal-to-noise ratio that is proxied by $\frac{v_\theta}{v_\epsilon}$ or, in other words, *ceteris paribus*, signals are expected to be more valuable when v_ϵ is smaller and, in such sense, v_ϵ can be thought as an "adjusted price" of information.

TABLE 2. Number of agents using the most common strategies (All figures are averages over 100 simulations).

	$N = 5$	$N = 10$	$N = 15$
Top 1 strat.	387.20	375.53	373.19
Top 2 strat.	574.74	563.17	577.09
Top 5 strat.	765.86	694.56	691.06
Inf + Pas	549.64	554.37	561.34

Generally speaking, the strategies surviving at equilibrium (i.e., in $\mathcal{P}_T$ that proxies $\mathcal{P}_\infty$)) are only a tiny fraction of the possible ones and this is a robust finding holding in all parametrizations. Table 2 shows the number of agents using the most common, the two and the five most common strategies. For instance, when $N = 5$, the two most frequent strategies are used on average by 575 agents (out of 1000). In other words, a fraction of $2/32 \approx 6\%$ of the strategies account for 57.5% of agents in equilibrium. Usually, the most used 5 strategies are taken by about 70% of agents or more. The concentration into very few strategies is striking if one thinks that there are $2^{10} = 1024$ and $2^{15} = 32768$ strategies when $N = 10$ and 15.

The last row of Table 2 shows the average number of agents who selected either the informed ($b_i = b^*$) or the passive strategy ($b_i = \mathbf{0}$). A comparison of the second and the 'Inf + Pas' row reveals that, essentially, the two most used strategies are precisely the informed and the passive one. Hence, evolution drives most of the agents to pick exactly one between these two strategies, despite the availability of tens (or hundreds or thousands) of alternatives.

Not only agents concentrate on 2 (or very few) strategies but use less information than may be naively expected. From the seminal work in [Grossman and Stiglitz, 1980] we know that enough informed traders should in principle allow the others to deduce or "smell" what is needed even with no direct access to the information itself. However, the extent to which this happens is probably surprising.

TABLE 3. Use of information: shares of relevant bits set and overall share (All figures are averages over 100 simulations.

		$N = 5$		$N = 10$		$N = 15$	
		Overall	Relevant	Overall	Relevant	Overall	Relevant
	0.02	0.241	0.225	0.251	0.216	0.256	0.214
v_ϵ	0.03	0.165	0.149	0.187	0.156	0.189	0.152
	0.04	0.154	0.141	0.166	0.141	0.165	0.133

Table 3 shows the fraction of bits set to 1 by the whole population of traders and the fraction of relevant bits[4], as a function of N and v_ϵ. The number of used bits does not depend much on N but it is quite sensitive to the variance (price of information): regardless of N, the fraction of 1-bits drops, say, from one-quarter to about 16% as v_ϵ varies from 0.02 to 0.04.

Put differently, the model points to a low use of information that is discarded by many traders in our noisy setup where strategies compete on short-term profits. Table 4, showing the fraction of informed and passive agents, provides additional details.

TABLE 4. Percentage of informed and passive traders. All figures are averages over 100 simulations.

		$N = 5$		$N = 10$		$N = 15$	
		Informed	Passive	Informed	Passive	Informed	Passive
	0.02	0.440	0.067	0.402	0.101	0.369	0.111
v_ϵ	0.03	0.346	0.232	0.299	0.272	0.291	0.291
	0.04	0.253	0.325	0.221	0.391	0.245	0.380

First, scanning the table horizontally, it can be seen that the fraction of informed traders decreases with increasing complexity, as measured by N (for any level of v_ϵ). The effect is more pronounced when v_ϵ is low or medium.

Second, the fraction of informed agents sharply decreases with v_ϵ. For instance, when $N = 10$, doubling v_ϵ roughly halves the fraction of informed traders (from 40.2% to 22.1%). This hints at the fact that (the same) information is less useful when embedded in noisier markets or, if you wish, when it is more expensive in relative terms.

[4]The "overall" number of set bits is (# of set bit) / NM and the number of "relevant" bits is (# of set bits / SM).

Third, a similar portrait surfaces looking at the number of passive traders, the ones who decide not to use *any* signal at equilibrium. Their number is inversely proportional to that of informed traders and, hence, in general there are more such agents for large N and v_ϵ. The intuition is that passive agents are better equipped to survive in more volatile markets, flooded with plenty of information. Simply put, in such "difficult" environments, discarding all signals and avoiding any cost is often the most commonly evolved strategy (picked in equilibrium by nearly 40% of agents in some cases). Interestingly, there are several values of the parameters (depicting somewhat realistic markets) where the passive traders outnumber the informed one, a finding that goes against the conventional wisdom that using (good) information should be better than discarding it.

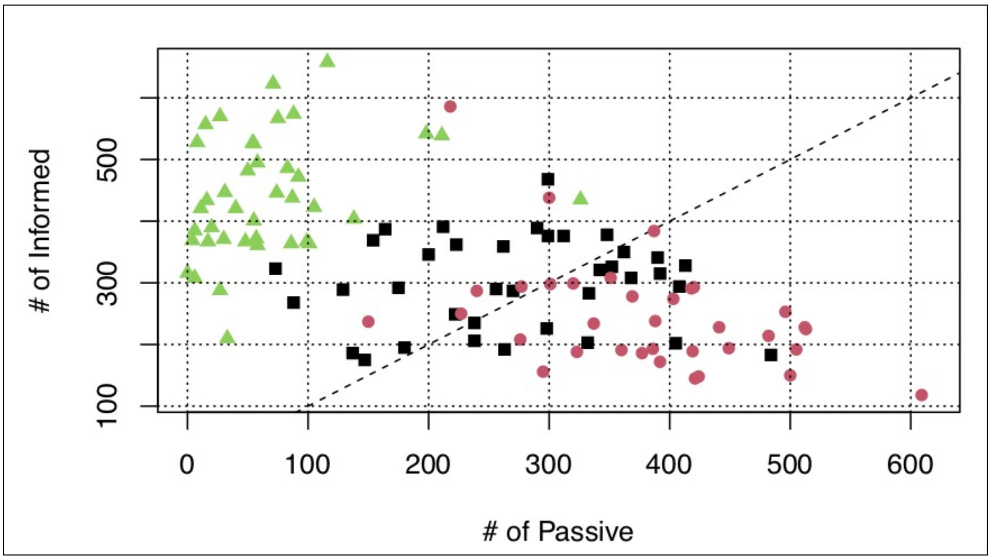

FIGURE 1. Number of passive (on the x-axis) and informed traders (on the y-axis) in three different markets: low informational complexity on the top-left corner, for $N = 5, v_\epsilon = 0.02$, and high informational complexity moving down to the bottom-right corner, where $N = 15, v_\epsilon = 0.04$. The dashed diagonal line is where the two numbers are the same.

Figure 1 depicts the number of passive and informed agents (out of 1000) and visually reinforces the previous claims: on the top left corner there is one green triangle for each simulation, with $N = 5, v_\epsilon = 0.02$ (low complexity), and few passive agents are outnumbered by many informed ones. At the other extreme, red circles show that the situation reverses when $N = 15, v_\epsilon = 0.04$ (high complexity). Black squares depict simulations with $N = 10, v_\epsilon = 0.03$ in which members of the two subpopulations are roughly equal in number. This is in good accordance with the fractions exhibited in the central cell of Table 4.

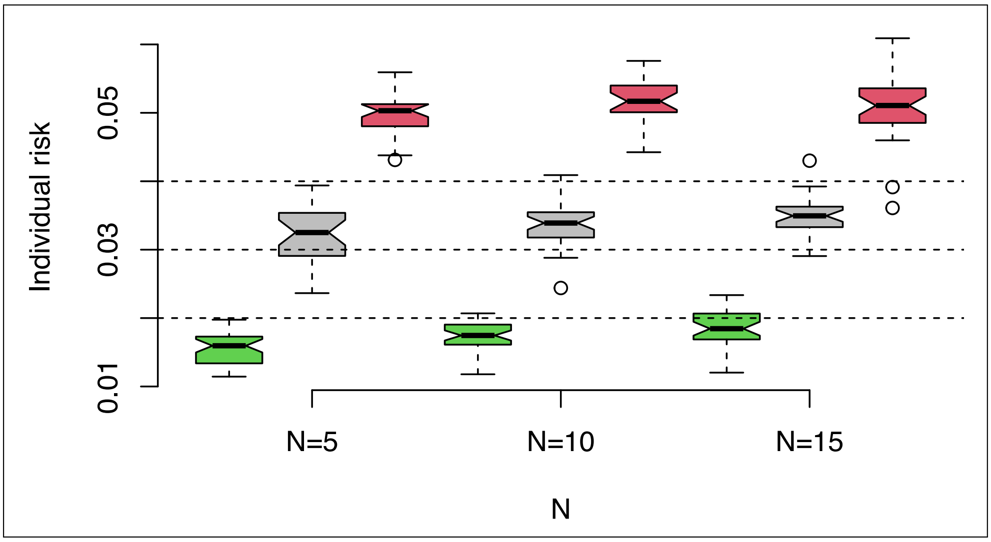

FIGURE 2. Distributions of individual risk for different N and levels of $v_\epsilon = 0.02, 0.03, 0.04$ in red, green and black, respectively.

The equilibrium shares of strategies are characterised by bit-strings $b_i, i = 1, ..., 1000$, but also by the individual risk assessments v_i that appear in the denominator of Eq. (1). Figure 2 represents through box-plots the distributions of the set of $v_i, i = 1, ..., 1000$, for markets with $N = 5, 10, 15$ (from left to right) and $v_\epsilon = 0.02$ (green), $v_\epsilon = 0.03$ (grey) and $v_\epsilon = 0.04$ (red). For instance, the median v_i when $N = 10, v_\epsilon = 0.03$ is about 0.035, as shown by the black line in the central grey box-plot, and most values are in an interval whose lower/upper extremes are slightly smaller/bigger than 0.03/0.04, respectively. The figure shows that the risk perceived by agents is mildly increasing with N, for whatever v_ϵ.

More importantly, it is worth noticing that when $v_\epsilon = 0.02$ the perceived risk is, on average, smaller than 0.02, whereas for higher v_ϵ agents on average evolve a much higher assessment. This is particularly true for (large) v_ϵ such as 0.04, as the red box-plots show substantially large medians about 0.05. In other words, when the volatility of the markets is low, agents learn to slightly underestimate the noise level in the market; conversely, when the volatility is higher, and especially for $v_\epsilon = 0.04$, they learn to overestimate the riskiness of the stock. As a consequence, because the v_i are in the denominator of the demand function, they take larger equity positions (than would perhaps be expected) when the volatility of the market is low, and reduce the risky component of their investments when the volatility is high. This behaviour, aggressive as well as cautious in different cases, appears to curb the probability of financial extinction of agents' current strategy in our setup where sustained evolutionary competition is present.

As seen before, the number of passive traders is relatively large in markets with high v_ϵ and, consistently with this fact, passive traders evolve higher v_i. Hence, they not only discard information in building their portfolio but demand less units of the risky stock (other things being fixed) in an attempt to take into account the residual risk of not being informed. Symmetrically, informed agents use all the relevant information *and* boost their demand through lower v_is[5].

4. Conclusion

The model described in this paper depicts a market where only S pieces of news out of N affects the stochastic payoff of a risky stock. News/information are available at a cost in any period t and boundedly rational agents must figure out (or learn through evolution) which signal to use in forming their demand for the risky asset. Agents' demands (phenotypes) are driven by strategies (genotypes) prescribing the signal to use/discard and by an adjustable assessment of risk. Updates occur through pairwise comparisons (or tournaments) aiming at favouring the strategies yielding the highest payoffs. A small rate of mutation ensures that adequate diversity is preserved in learning.

We have examined which strategy prevail in the long run and their shares in the population $\mathcal{P}_t, t \to \infty$ of heterogeneous agents. This can

[5]Another interpretation leads to overconfidence on the part of informed traders: in principle, once all relevant component of θ_t are used, there is no intrinsic noise other than ϵ and super-rational agents should set $v_i = v_\epsilon$. The model shows that there is evolutionary pressure to adjust downwards the individual risk assessment, or, that it pays off for the informed to be overconfident at equilibrium.

be thought as a canonical evolutionary model where fitter strategies are determined by sharp pairwise comparison of profits, and tend to grow along time.

We found that most traders evolve (or learn) either a passive strategy discarding all signals or a fully informed strategy, where all relevant signals are incorporated in the estimation of the future payoffs (sustaining the costs). Whether the passive or informed strategy takes the lion's share mostly depends on the "complexity" of the market: if the exogenous volatility or the number of news N are large, most traders will be passive and use no information whatsoever; if, instead, the volatility is low and the informational landscape has manageable size and costs, then a majority of agents will develop informed strategies. This polarity between full usage and full disregard of information fits well with models in the spirit of [Schredelseker, 2014], where it is shown that returns are U-shaped in terms of information and, hence, being entirely uninformed (passive in our setup) or fully informed is more profitable than being half-way on either side.

Overall, the model is demonstrating that it may impossible for boundedly rational agents to exploit all the information or disentangle relevant from irrelevant news in volatile market setups or when too much information is provided and must be screened. The model shows that fact-based learning by trial-and-error and imitation (plus mutation), does not allow full exploitation of relevant information. It is left to future research to investigate whether the same results hold for other learning schemes (or using other ways to spread fitter strategies). For example, the probability to adopt a better strategy in a pair may be proportional to the returns, instead of being 100%. While our switching rule may appear dummy, agents have a nice way to rationalise their behaviour as, in equilibrium, all strategies have the same median returns and, say, passive agents have a 50% chance to over-perform informed trader, i.e. they would fare first one time out of two in a race.

A byproduct of the model possibly hints at a novel way to explain why many investors appear to hold an excessively small share of equity[6], see [Robson and Orr, 2021] for an explanation based on biological evolution. In our model, for many values of the parameters, the passive traders, who demand little equity due to their large v_i's, are the majority share and this leads in aggregate to a limited (relative) share of risky holdings with respect to risk-less investments, in line with the historical simulation in [Benartzi and Thaler, 1995].

4.0.1. *Acknowledgements.* We thank the audiences at WEHIA 2022 in Catania, City University London and AAU Klagenfurt for their comments and remarks. Luca Gerotto and Marco Tolotti provided many insightful comments and ideas on related previous work.

References

[Benartzi and Thaler, 1995] Benartzi, S. and Thaler, R. (1995). Myopic Loss Aversion and the Equity Premium Puzzle. *The Quarterly Journal of Economics*, 110(1):73–92.

[Grossman and Stiglitz, 1980] Grossman, S. J. and Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408.

[Huang et al., 2024] Huang, Y., Zhou, C., Cui, K., and Lu, X. (2024). A multi-agent reinforcement learning framework for optimizing financial trading strategies based on timesnet. *Expert Systems with Applications*, 237:121502.

[Lo et al., 2000] Lo, A. W., Mamaysky, H., and Wang, J. (2000). Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, 55(4):1705–1765.

[Mehra and Prescott, 1985] Mehra, R. and Prescott, E. C. (1985). The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161.

[Newton, 2018] Newton, J. (2018). Evolutionary game theory: A renaissance. *Games*, 9(2).

[Robson and Orr, 2021] Robson, A. J. and Orr, H. A. (2021). Evolved attitudes to risk and the demand for equity. *Proceedings of the National Academy of Sciences*, 118(26):e2015569118.

[Sandholm, 2010] Sandholm, W. H. (2010). *Population Games And Evolutionary Dynamics.* Economic learning and social evolution. MIT Press.

[Schredelseker, 2014] Schredelseker, K. (2014). Pascal's wager and information. *Journal of Forecasting*, 33(6):455–470.

[6]The reason why this happens is not entirely clear and this phenomenon is known as the "equity premium puzzle", see [Mehra and Prescott, 1985].

[Smith, 1982] Smith, J. M. (1982). *Evolution and the Theory of Games*. Cambridge University Press.

DEPARTMENT OF ECONOMICS, CA' FOSCARI UNIVERSITY, VENICE, ITALY
Email address: paolop@unive.it