

УДК 336.012.23

An Evolutionary Model of Financial Market Efficiency with Costly Information

Aleksei Pastushkov

National Research University Higher School of Economics,
11, Pokrovsky Bulvar, Moscow, 109028, Russian Federation.
E-mail: apastushkov@hse.ru

A classical paper by [Grossman, Stiglitz, 1980] showed that asset prices in equilibrium necessarily contain some degree of inefficiency when information is costly. Moreover, the inefficiency should decrease with decreasing information costs. However, a number of recent empirical studies cast some doubt on the postulate that real financial markets are becoming more efficient, despite the ostensible advances in technology and radical increases in the availability of information. The pricing of financial assets strongly depends on the relative wealth shares of heterogeneous groups of investors interacting in the market and the recent field of evolutionary finance has produced a number of studies showing how asset price dynamics develop under the influence of endogenously changing populations of heterogeneous traders. However, very few studies exist examining how the cost of information affects prices of assets in an evolutionary context. We therefore construct an evolutionary agent-based model of a financial market with boundedly rational traders who learn from experience. We conduct a number of computational experiments with varying information costs and show that even under zero information costs uninformed traders can survive and even dominate the market in finite time. Thus, for an initially low level of information costs, a marginal decrease in them does not necessarily lead to increased price efficiency. For higher initial levels of information costs our results however agree with those of [Grossman, Stiglitz, 1980] in that an increasing information cost generally leads to informed traders being driven out of the market and asset prices becoming less efficient. Implications of our findings for financial market regulation are discussed.

Key words: market efficiency; bounded rationality; multiagent model; evolutionary finance; costly information; agent-based modelling.

JEL Classification: C63, D82, G14.

Aleksei Pastushkov – research assistant, International Laboratory for Experimental and Behavioural Economics.

The article was received: 19.12.2023/The article is accepted for publication: 10.05.2024.

DOI: 10.17323/1813-8691-2024-28-2-276-301

For citation: Pastushkov A. An Evolutionary Model of Financial Market Efficiency with Costly Information. *HSE Economic Journal*. 2024; 28(2): 276–301.

Introduction

Financial market efficiency is one of the cornerstones of the theory of asset pricing as well as a practical concern related to the ability of financial markets to perform their societal function of allocating capital to its most productive uses. With regards to its theoretical importance, it can be pointed out that market efficiency, in the form of the no-arbitrage condition, is one of the key assumptions still underlying most standard models of asset pricing (see e.g. [Black, Scholes, 1973; Cox et al., 1979; Merton, 1973; Cochrane, 2009; Duffie, 2010]). With regards to the multiplication of societal welfare, as [Fama, 1970] points out, “the primary role of the capital market is allocation of ownership of the economy's capital stock. In general terms, the ideal is a market in which prices provide accurate signals for resource allocation [...]”. Therefore, from a practical standpoint, investigating whether the assumption of efficiency generally holds, and if not, what the contributing factors leading to inefficiency are, is of utmost importance for understanding when conventional asset pricing models may break down and what actions, if any, regulators can take to increase the efficiency of capital allocation.

Notably, in a more recent review of the literature on market efficiency [Fama, 1991] concedes that an extreme form of market efficiency is only possible in a frictionless world. By contrast, in the real world, with positive information and transaction costs, perfect efficiency almost surely cannot be attained. This intuitive idea is formally proved in a classical paper by [Grossman, Stiglitz, 1980], who explicitly take into account costly information and show that in this setting a financial market necessarily has some degree of inefficiency in equilibrium, and a state in which prices fully reflect costly information (i.e. are fully efficient) cannot be an equilibrium state. Another prominent result of [Grossman, Stiglitz, 1980] is that the cost of information is positively related to the size of asset mispricing.

The publication of the [Grossman, Stiglitz, 1980] paper inspired, on the one hand, a stream of research seeking to document empirical manifestations of capital market inefficiency, and on the other, a large theoretical literature investigating how information gets incorporated into market prices under various assumptions. To name a few prominent contributions, the former strain of papers is represented by [Hasbrouck, 1993; Heston et al., 2010] who seek to spot price inefficiencies in the historical market data using econometric techniques, while some early representatives of the latter are [Ho, Roni, 1988; Figlewski, 1982] who introduce variable information costs and information diversity, respectively, and show how these assumptions affect a possible asymptotic convergence of the market towards an informationally-revealing equilibrium.

More recently, however, it has been pointed out by multiple researchers that despite ostensible technological improvements and an increase in availability (or equivalently, a decrease in cost) of information that financial markets have witnessed over the recent decades, there is no clear indication that markets are steadily becoming more efficient. More specifically, [Lo, 2004] points out that the first-order autocorrelation of monthly S&P 500 returns exhibits a cyclical pattern in a large sample comprising more than 100 years of data, with markets being less auto-correlated in certain periods in the 1950s than in the early 1990s. [Alvarez-Ramirez et al., 2012]

propose an entropy-based measure of stock market efficiency and study an 80-year sample of daily returns of the Dow Jones Industrial Average. They find a time-varying pattern of efficiency, with the market notably becoming less efficient from 2006 onwards. [Lux, 2009] points out that there is “no indication that the shape of the return distribution has undergone any remarkable changes over the past decades reflecting an increase in information transmission”, despite the advances in electronic trading. [Lim, Brooks, 2011] survey a number of empirical studies examining the evolution of market efficiency over time and summarize the findings as inconclusive, with some showing a gradual improvement in efficiency and others failing to find it. Finally, the recent experience of notable market practitioners speaks against an improvement in market efficiency as well¹.

The evidence discussed above casts doubt on the simple relationship between the cost of information and the degree of market efficiency postulated by [Grossman, Stiglitz, 1980]. To be fair, [Grossman, Stiglitz, 1980] point out that the cost of information is not the only factor determining the degree of market efficiency. Clearly, the ability of the price to reflect information also depends on the quality of the information signal and the informativeness of the price system [Grossman, Stiglitz, 1980]. The informativeness of the price system, in turn, presumably depends on the employed market mechanism and the relative wealth shares of the informed and uninformed traders. In their investigation [Grossman, Stiglitz, 1980] choose to focus solely on the effect of the information cost, leaving aside considerations pertaining to the signal quality, market mechanism and the relative shares of the informed and uninformed traders.

Clearly, however, the relative shares of the informed and uninformed traders depend on the excess profit opportunities available to the informed traders in the market. This introduces a feedback loop between the degree of market efficiency and the relative wealth shares of the two groups of traders. In the original paper by [Grossman, Stiglitz, 1980] the authors circumvent this difficulty by assuming perfect rationality of agents who are able to evaluate *ex ante* the expected utilities of the informed and uninformed strategies. The relative shares of the informed and uninformed agents are then instantly optimized, such that the marginal utilities of being informed and uninformed are equal. This obviates the need to explicitly model the feedback loop between the degree of market inefficiency and the relative growth rates of wealth controlled by the informed and uninformed agents.

However, it is important to recognize that the utility of individual traders following either the informed or the uninformed strategy depends on the behavior of the other traders, who for their part face the same issue. It is far from obvious that in real markets each trader has accurate and mutually consistent beliefs about the preferences, wealth levels and strategies of all the other traders. We argue that in the absence of this common knowledge assumption the only way for a trader to find out the relative value of being informed or uninformed is by *learning* it. More specifically, by learning it from own experience of interacting with the market.

The above considerations suggest that to capture possible interactions between the cost of information, the degree of inefficiency and the informativeness of the price system a model should incorporate two additional features: endogenous wealth accumulation dynamics and learning performed by traders. The former is a feature typically encountered in evolutionary models of financial markets while the latter is a characteristic trait of bounded rationality models.

¹ See e.g. [Wigglesworth, 2023].

In this paper we construct a model combining bounded rationality and features of evolutionary financial models to examine the impact of information costs on the efficiency of a financial market. In our model, agents learn the optimal information acquisition strategy by a simple form of reinforcement learning while simultaneously being subjected to market selection forces via endogenous wealth dynamics. As the complex non-linear interactions between the model's components make mathematical analysis of the model intractable, we turn to the computational agent-based modelling paradigm.

The advantages of computational agent-based models are on the one hand their ability to model complex systems such as financial markets “from the bottom up”, allowing macro-level properties to emerge from the micro-level behavior of the constituent parts, and thus avoiding the mathematical difficulties associated with aggregation. [Axtell, Farmer, 2022] On the other hand, computational models allow us to examine whether the dynamics of agent interactions lead to any steady state, in cases where the existence of such a steady state cannot be *a priori* assumed (see [Arthur, 1995; Tesfatsion, 2002]). In cases where asymptotic convergence to a steady state is assured, these models allow us to examine pre-asymptotic behaviour occurring in finite time.

To our knowledge there currently exists only a scarce amount of papers examining the relationship between information costs and stock market efficiency in an evolutionary setting with boundedly rational learning traders. The few papers that do tackle this issue, e.g. [Sciubba, 2005; Gong, Diao, 2022], either make the assumption that traders learn from prospective equilibrium prices, which has been criticized by [Hellwig, 1982], or define traders' profits as capital gains from changing market prices of assets, without tying them to any exogenous fundamental value. In our model we avoid both of these assumptions and show that in this setting decreasing information costs do not necessarily lead to greater market efficiency, particularly when these costs are initially low, and uninformed traders can survive even when the information cost is zero.

Before proceeding to the description of our model, we familiarize the reader with the evolutionary finance paradigm, highlighting its most prominent results as well as modelling assumptions and approaches. Similarly, a short section is dedicated to a brief introduction of key concepts in reinforcement learning which are relevant to our modelling of agents' learning.

The remainder of the paper is organized as follows: in Section 1.1 we briefly review the extant results on the incorporation of costly information into asset prices obtained in the static equilibrium paradigm. Section 1.2 provides a short introduction to the key ideas of evolutionary finance, while Section 1.3 introduces reinforcement learning and, in particular, its subfield concerned with the so-called multi-armed bandit problems. In Section 2 we describe the structure of the computational multiagent model used here as well as its key parameters. In Section 3 we provide the results of our computational experiments highlighting the discovered relationships between the degree of market efficiency and the information costs. Section 4 suggests prospective avenues for further research and concludes.

1. Literature Review

1.1. Financial Market Efficiency

Following the publication of the [Grossman, Stiglitz, 1980] paper, [Diamond, Verrechia, 1981] pointed out that the model proposed by [Grossman, Stiglitz, 1980], although providing a

justification for partially revealing prices, does not actually illuminate the role of market prices as aggregators of diverse pieces of information. Indeed, in the model of [Grossman, Stiglitz, 1980] there is just one piece of information that investors may choose to purchase, and if several investors choose to become informed, they are essentially duplicating their effort in acquiring information to receive a prediction identical to that received by other informed agents. [Diamond, Verrechia, 1981] thus expand upon this framework and investigate how market prices reflect information when its diverse pieces are distributed among individual traders. In doing so, they also circumvent an additional conceptual difficulty of the [Grossman, Stiglitz, 1980] model: while in the market of [Grossman, Stiglitz, 1980], in order to incentivize investors to acquire private information there has to be some exogenous supply noise preventing the price from revealing all information, in the market considered by [Diamond, Verrechia, 1981] there is no need to introduce an exogenous noise variable. The price is prevented from being perfectly correlated with a given investor's private information by the existence of divergent pieces of information reflected in the price via demands of other traders. In other words, an investor can use the price as an information signal and still have an incentive to acquire additional private information. A major result of the paper is that in the presence of various divergent pieces of information the price cannot be fully revealing, even if there is no exogenous noise.

[Hellwig, 1982] notes that the [Grossman, Stiglitz, 1980] model faces another conceptual difficulty: namely, investors in the model learn from a *prospective* equilibrium price. As [Hellwig, 1982] writes, "Grossman and Stiglitz do not assume that agents observe which price clears the market and then determine their demands. Instead, the price clears the market after agents have already used the information contained in the fact that it is an equilibrium price. In announcing their demands to the «auctioneer», agents must already consider what it would mean if the price called by the «auctioneer» happened to clear the market." This, according to [Hellwig, 1982], is the likely reason behind the major result obtained by [Grossman, Stiglitz, 1980] that market prices are bounded away from full efficiency. In the model proposed by [Hellwig, 1982] agents learn from past realizations of equilibrium prices. The result obtained is that when the time between successive transactions is sufficiently small, the market prices can approximate full efficiency arbitrarily closely.

[Verrechia, 1982], using the framework introduced by [Diamond, Verrechia, 1981], demonstrates the existence of a rational expectations competitive equilibrium in which the amount of costly diverse information acquisition is endogenously determined. As in the original [Grossman, Stiglitz, 1980] paper, they show that the level of price informativeness increases with decreases in the noise levels, the cost of information and the aggregate risk aversion of traders.

In contrast to the models reviewed above which model the capital market as perfectly competitive, [Kyle, 1989] takes an altogether different approach by considering a market in which traders have a non-negligible impact on the price and are aware of it. The traders in this model thus act strategically. [Kyle, 1989] justifies this assumption on the empirical grounds that informed traders in real markets tend to be large and are presumably aware of the impact of their own trading. Additionally, as he notes, treating the market as perfectly competitive leads to the very paradox highlighted in the [Grossman, Stiglitz, 1980] paper, in which there is an interactive effect between the informativeness of the price and the decision of traders to acquire information, implying the no-fully-revealing-equilibrium result.

[Kyle, 1989] shows that when each informed trader trades against an upward-sloping residual supply curve, he restricts the quantity of risky asset he trades compared to the perfectly

competitive level, acting as a monopsonist. The main result of the paper is that when noise trading vanishes, there exists an equilibrium in which the price is not fully revealing and yet the profits provided by private information are driven to zero, as the supply curves against which informed traders trade are too steep, or in other words, the market is too illiquid. Thus, the price does not even have to reach full efficiency for information gathering to cease being profitable when traders are strategic.

More recently, [Berk, 1997] considers an intertemporal rational expectations model of a market with a finite number of trading periods. He shows that in this model, a fully revealing equilibrium exists in which all traders pay a strictly positive amount to acquire information. Additionally, the possibility of an equilibrium is demonstrated, in which agents purchase information despite this rendering the whole agent population, including the informed ones, strictly worse off. This result is thus in contradiction to the one obtained by [Kyle, 1989]. Arguably, however, the assumption of finiteness of transaction periods on which the model of [Berk, 1997] is built is not a realistic one.

Finally, we can mention some experimental literature that considers the question of informational market efficiency from a slightly different angle: namely, whether a larger amount of private information that an investor has always leads to better trading outcomes as expressed in higher profits. [Huber, 2007] conducts an experiment in which different groups of traders are communicated private information with different time lags, such that some of them have a timing advantage in access to information. He finds that in this setting, while the best informed traders demonstrate superior performance, the performance of the average informed traders is worse than that of the least informed. Thus, this experimental market demonstrates J-shaped returns to information.

In a similar vein, in an experiment conducted by [Huber et al., 2008] it is demonstrated that when information is cumulative, it is only the best informed traders that significantly outperform all others, while the difference in performance between traders with various levels of information below the full information is not significant. The result suggests that the marginal returns to information are not necessarily strictly positive.

Similarly to the [Grossman, Stiglitz, 1980] model, the research summarized above employs the perfect rationality paradigm and thus circumvents the need to model the influence of evolutionary dynamics and learning on market efficiency under costly information. As pointed out above, in real markets to evaluate the utility of becoming informed a trader has to have access to the information sets of the other traders, who jointly with him determine market prices, net asset payoffs and hence his returns. In the absence of an explicit mechanism that explains how a population of traders can come to this common knowledge equilibrium, a more realistic approach to modelling would be to assume that traders learn the utility of being informed from experience. As experience takes time to acquire, traders' wealth evolves in parallel to the learning process by way of accumulating profits and losses. [LeBaron, 2011] terms these parallel processes "active" and "passive" learning, respectively. Arguably, however, while the former process represents learning in the conventional sense, the latter may be more narrowly defined as evolutionary selection.

Models of financial markets incorporating the evolutionary element of endogenous wealth growth in a heterogeneous agent setting are studied in the relatively novel field of evolutionary finance. Over the last two decades this field has delivered a number of non-trivial results on market stability, portfolio selection and asset price dynamics, however, applications of evolution-

ary modelling to the study of financial market efficiency under costly information are less common. The next subsection introduces the key elements of evolutionary finance and highlights the few studies that study financial market efficiency in an evolutionary setting.

1.2. Evolutionary Models of Financial Markets

The evolutionary approach to finance is built upon the notions of heterogeneity and selection, whereby the selection in the context of financial markets proceeds by way of some traders accumulating more wealth and dominating the market and others going bankrupt and leaving the market.

One strain of evolutionary finance research, espoused by its pioneers, such as Thorsten Hens and Igor Evstigneev, does away with the notion of individual rationality and optimization (see e.g. [Evstigneev et al., 2009]) and studies which market strategies survive in an evolutionary competition. Hereby, the unit of selection is a strategy itself, that is, a trading behavior, and not an agent who has well-specified preferences and beliefs. By making this shift, evolutionary models of this type circumvent the need for elaborate specification of the model of an individual trader and allow for a wide range of strategies, including heuristic ones, to be considered as candidates for survival. The key notion in these models is the concept of an evolutionarily stable strategy (ESS), that is, a strategy that asymptotically drives any other strategies out of the market (for more on the notion of ESS, see [Smith, Price, 1973]).

The results obtained in this type of models are often asymptotic, describing an equilibrium ecology of strategies (which may be represented by one or many strategies), which forms in the long run. These models, however, do not make a prediction by which trajectory the market ecology settles into an equilibrium. Nor does it say anything about the reasonableness of assuming the existence of a given candidate strategy in the initial pool of strategies in the first place.

Another strain of evolutionary finance research proceeds in a more traditional way by specifying a model of an individual trader endowed with a certain kind of preferences and/or beliefs (which may be of either the traditional neoclassical kind or of the more unorthodox behavioral kind). It often utilizes a computational approach, whereby once the models of individual heterogeneous traders are specified, an artificial market is constructed and populated by these traders in various proportions and a number of stochastic simulations are carried out to examine the arising market dynamics and possible equilibrium states.

Recent representatives of this second approach are the papers by [Scholl et al., 2021; Lussange et al., 2021]. The former examines how the relative returns of three groups of traders (value investors, trend followers and noise traders) depend on wealth shares of these groups in the market. The authors identify different types of relationships that may exist between any two groups of traders, e.g. a mutualistic relationship, where both groups benefit from coexisting, or a predator-prey relationship, where one group of traders benefits at the expense of the other.

The paper by [Lussange et al., 2021], although mainly focussing on replicating a number of stylized facts of financial markets, also incorporates an evolutionary element, as the wealth of agents in this agent-based simulation is endogenously driven by their market performance.

As the question of informational efficiency arises in the context of asymmetrically informed traders, which is one type of agent heterogeneity, and, intuitively, the more wealth a given group of agents controls, the larger is the impact that it should have on market prices, the study of market efficiency can naturally be considered in an evolutionary context. There exists a limited number of studies examining various aspects of market efficiency in an evolutionary paradigm.

For instance, [Sciubba, 2005] develops an evolutionary model in the classical setting of [Grossman, Stiglitz, 1980] and obtains an asymptotic result: although informed traders accumulate wealth at faster rates when they are a minority in the market, this relative advantage decreases when informed traders control large proportions of wealth and is eventually offset by the cost of information, which means that in the limit uninformed traders survive, i.e. are not driven out of the market, as long as the information costs are strictly positive. However, the model of [Sciubba, 2005] is subject to the same critique leveled by [Hellwig, 1982]: in this model, analogously with [Grossman, Stiglitz, 1980], uninformed traders imperfectly learn the payoff signal from observing the equilibrium price. But the equilibrium price is partially the result of their own trading actions, hence they cannot learn from it prior to determining their own risky asset demands. This issue is avoided in our model as the information set of the uninformed traders includes only the past realizations of risky asset payoffs, which are publicly available.

[Gong, Diao, 2022] construct an evolutionary model of the market incorporating information asymmetry, two types of trading strategies (fundamentalist and chartist), various degrees of rationality (conceptualized as the propensity to switch between the fundamentalist and chartist strategies) and a liquidity factor. They find that in general, the price almost never reflects the true fundamental value of an asset, but rather a weighted average of fundamental value estimates of heterogeneous groups of traders. Additionally, the authors characterize the influence exerted on market prices by other factors in the model, such as the liquidity factor and the rationality factor.

The approach we take in this paper is methodologically similar to the one taken in [Gong, Diao, 2022] as well as in [Scholl et al., 2021] in that we examine the dynamics of market efficiency by way of constructing a computational model with boundedly rational investors. However, our model is different from these predecessors in that [Scholl et al., 2021] only examine evolutionary dynamics of 3 groups of traders with fixed behavior (fundamentalist, chartist and noise traders) and do not model the influence of information costs either on their relative wealth levels or on market efficiency. Our model is also different from that of [Gong, Diao, 2022] in several ways. First, the agents in our model are risk-neutral. Second, we use the market clearing mechanism of the so-called temporary equilibrium (more on this in Section 2.4) to avoid additional complexity which comes from introducing an ad hoc parameter of market liquidity, as done in [Gong, Diao, 2022]. Third, even the informed traders in the model of [Gong, Diao, 2022] only observe an imperfect signal of the fundamental value, which may contribute to their finding that uninformed trading survives. As was shown by [Grossman, Stiglitz, 1980] the quality of the informed trader's signal is positively related to market efficiency. In our model, informed traders receive a perfect signal of the payoff value and hence the quality of the signal is not a factor determining market efficiency or relative performance of informed and uninformed traders. Finally, profits that traders can earn in our model explicitly result from the difference between prices and fundamentals, in contrast to [Gong, Diao, 2022], whose traders in fact profit from differences in market prices in successive periods. This way of determining the profits can arguably represent a separate destabilizing factor which causes the price to deviate from the fundamental value in their model.

It is hoped that our model, given the assumptions elucidated, more clearly illuminates the relationship between the cost of information, the evolutionary dynamics of the informed and the uninformed traders and the asset mispricing in a setting which explicitly rewards traders for correctly estimating the fundamental value, in line with the original framework of [Grossman, Stiglitz, 1980]. The next section introduces the basic features of reinforcement learning relevant for our modelling of agents' bounded rationality.

1.3. Reinforcement Learning

Reinforcement learning describes the process of goal-directed learning by interacting with the environment. In general, the goal is to learn an optimal *policy*, that is the optimal action in a given situation. In contrast to supervised learning where the learner may be presented with a sample of actions and corresponding outcomes, in reinforcement learning the learner must obtain such a sample by iteratively performing actions in the environment and receiving rewards (see e.g. [Sutton, Barto, 2018]). A measure of the efficiency of reinforcement learning is how quickly the learner converges to the optimal policy, since clearly the sooner the agent learns to select the optimal action in each situation the greater his cumulative reward will be.

The above considerations introduce a trade-off between exploration and exploitation. By tending to exploit the actions that have been found beneficial in the past the agent runs the risk of getting stuck in a local optimum and failing to find the globally optimal policy. On the other hand, by tending to explore too much the agent might forego receiving the rewards he could have received if the locally optimal policy turned out to be the global optimum.

A full reinforcement learning problem arises when the actions performed by the agent not only influence his immediate rewards but also the transition probabilities between states in which he consecutively finds himself. A simpler subset of reinforcement learning are the so-called multiarmed bandit problems in which either the transition probabilities between the states are not affected by the agent's actions but the agent can still find himself in different states (the so-called contextual bandits), or there is only one state with several actions to choose from such that the issue of transition probabilities does not arise. Whether an agent can recognize the state he is in, and therefore whether a given problem must be modelled as a contextual or a simple multiarmed bandit problem, depends on the agent's information set.

In the context of financial markets, it can be argued that the only reliable information a trader has access to are the publicly available market prices of assets, historical fundamentals and his own trading history. Given the competitive nature of financial markets, in which successful strategies are held in secrecy, it is reasonable to assume that a trader cannot observe other traders' strategies, capitals or trading histories. As there is no general principled way for an agent to partition this information set into discrete states and the theory presented in [Grossman, Stiglitz, 1980] and subsequent literature reviewed in Section 1.1 is not contextual, i.e. it does not assume that a trader decides whether to become informed or uninformed based on some temporary state of environment, we follow the convention by assuming that an agent believes that there is a *globally* optimal strategy, either the informed or the uninformed one. Therefore, the agent faces a non-contextual multiarmed bandit problem, in which he must learn the optimal strategy in only one state via interacting with the market. The next section describes our model of a financial market in full detail.

2. The Model

2.1. Assets

Consider a financial market with two assets: a risky short-lived asset in fixed finite supply of N infinitely divisible shares, and a risk-free bond in infinite supply. The risk-free bond pays a fixed interest rate in each period. The annual interest rate is set to 1%, the daily interest rate is thus $1/252\%$.

The risky short-lived asset can be interpreted as short-term stock options, futures or, more generally, any asset that provides a payoff determined by random exogenous factors and becomes worthless afterwards. The use of short-lived assets for the purpose of modelling evolutionary dynamics in financial markets is common in evolutionary finance (see e.g. [Evstigneev et al., 2009; Evstigneev et al., 2002; Hens, Schenk-Hoppé, 2005]).

The risky asset in our model lives for one period, provides a payoff and is then reissued in the next period. The intrinsic value of the risky asset is equal to its discounted payoff, which is determined by:

$$(1) \quad F_{t+1} = F_t \cdot (1 + r_r),$$

where r_r is the stochastic period-by-period change in the payoff, governed by a white noise process with a normal distribution $\sim N(\mu, \sigma)$.

The market price of the risky asset is endogenously determined by the agents' trading. The exact trading mechanism is described in Section 2.4.

In each simulation run, the initial value of the payoff is set somewhat arbitrarily to $F_0 = 30$ and a stochastic path is simulated for T steps following Equation 1.

2.2. Traders

As in the model of [Grossman, Stiglitz, 1980] traders in our model can be of two types: informed and uninformed. In our model traders endogenously decide on their type in each period of the simulation based on a simple multiarmed-bandit form of reinforcement learning which will be described shortly.

Informed and uninformed traders are described by their information sets. Uninformed traders can costlessly observe the full history of the payoff process, up to and including the time t , however, they do not know the realization of the stochastic payoff process in the next period $t + 1$. Informed traders, on the other hand, receive a perfect signal of the payoff value at time $t + 1$, i.e. they know with certainty what the realization of the payoff will be in the next period. For that, however, they incur a fixed cost of acquiring information.

Uninformed traders use the full history of the payoff process up to the time t which is by assumption available to them to predict the value of the payoff in time $t + 1$. This uninformed estimate is obtained by calculating the average payoff change during the whole observed history of the payoff process and extrapolating this change to the period $t + 1$, using the value of the payoff in time t as the base value. Forecasts $Z_{k,t+1}$ of payoffs obtained by the informed and uninformed traders are thus:

$$(2) \quad \begin{aligned} Z_{k,t+1} &= F_{t+1} \text{ if } k = 'Inf', \\ Z_{k,t+1} &= F_t + (1 + E_{k,t}(r_r)) \text{ if } k = 'Uninf', \end{aligned}$$

where

$$(3) \quad E_{k,t}(r_r) = \frac{\sum_{t=1}^T \left(\frac{F_t}{F_{t-1}} - 1 \right)}{t}.$$

Thus, although the uninformed traders' estimate of the next period's payoff converges to the true value as the observed sample of the payoff process grows with t , it still remains imprecise in finite time and the uninformed traders can over- or underestimate the value of the payoff.

The traders' profits when trading the risky asset result from the difference between the payoffs they receive and the market price. When the equilibrium price at which the asset is purchased is below the discounted value of the payoff, the traders that purchased the asset generate a positive return on investment. When the equilibrium price is above the discounted payoff value, the traders that purchased it experience a loss. As is common in evolutionary finance models (see e.g. [Evstigneev et al., 2009]), short sales and borrowing are not allowed and traders can only profit by exploiting underpricing of the risky asset, not overpricing.

Traders are risk-neutral. Both the informed and the uninformed traders produce point estimates of the next period's payoff, with the only difference that in the case of the informed traders this estimate corresponds to the actual future realization of the payoff, while in the case of the uninformed ones it only approximates it. Once the traders have produced their estimate of the payoff, they determine their demands for the shares of the risky asset. As all traders have an alternative opportunity to receive the risk-free return by keeping all of their capital invested in the risk-free asset, they discount their estimate of the payoff at the risk-free rate to determine the highest price they are willing to pay for the risky asset. In fact, both the informed and the uninformed traders invest all of their capital into the risky asset as long as the equilibrium market price is below or equal to their estimate of the payoff discounted at the risk-free rate. If the price is higher than the discounted estimate of the payoff, the traders keep their capital invested in the risk-free bond.

All traders are initialized at period $t = 0$ holding random amounts of initial wealth drawn from a uniform distribution.

The evolution of the wealth of each trader i is given by the following equation:

$$(4) \quad W_{it} = \alpha_{it} W_{it-1} \left(\frac{F_t}{P_t^*} \right) + (1 - \alpha_{it}) W_{it-1} (1 + r_f) - C,$$

where α is the percentage of the trader i 's wealth invested in the risky asset, $1 - \alpha$ is the percentage invested in the risk-free asset, F_t is the actual value of the risky asset's payoff at period t , P_t^* is the equilibrium market price, r_f is the risk-free rate and C is the constant denoting the cost of acquiring information. Uninformed traders do not incur an information cost and thus for them, $C = 0$. For informed traders, the cost C is a monetary constant that does not change throughout a simulation run. We have conducted computational experiments with different values of C (see more on this in Section 2.5).

2.3. Learning

As mentioned earlier, traders in our model are boundedly rational and adaptive. Each trader learns the value of being informed and uninformed by interacting with the market and observing own performance. We model the traders' learning as a multiarmed bandit problem in which a trader at the start of each period t chooses whether to acquire information about the

risky asset's payoff. If he chooses to become informed he pays the information cost C , observes the perfect signal of the payoff value in period $t + 1$, decides whether and how much to invest in the risky asset and receives the payoff. If the trader chooses to be uninformed, he estimates the value of the next period's risky asset payoff using the observed history of payoffs as described above and similarly decides on his allocation to the risky and risk-free assets. In both cases, the returns that the trader receives are the returns of following either the informed or the uninformed strategy. In the language of reinforcement learning these are the agent's rewards that he seeks to maximize.

Each trader estimates the return $Q_{k,n}$ of following each strategy k by averaging the returns he receives when choosing this particular strategy which is a widely used approach in reinforcement learning (see e.g. Chapter 2 of [Sutton, Barto, 2018]). On the n -th iteration the estimate $Q_{k,n}$ after the strategy k has been tried $n - 1$ times is thus:

$$(5) \quad Q_{k,n} = \frac{R_1 + R_2 + \dots + R_{n-1}}{n-1},$$

where R_n is the return the trader has received on the n -th iteration of following the strategy k . At the start of the simulation all traders' estimates of $Q_{k,n}$ are set to zero and the initial strategy is chosen randomly by each trader. As the traders' estimates of $Q_{k,n}$ are necessarily uncertain, an element of exploration must be incorporated into the adaptation process, such that each trader tries each strategy. This is achieved by introducing a dependence of the optimal strategy k on the uncertainty associated with the value $Q_{k,n}$ of this strategy:

$$(6) \quad k_t = \arg \max_k \left[Q_{k,t} + c \sqrt{\frac{\ln t}{N_{k,t}}} \right],$$

where $N_{k,t}$ is the number of times the strategy k has been tried up to the time t and c is a strictly positive constant regulating the degree of exploration. The above formulation of the optimal choice at each time t is termed the upper confidence bound method which has been found to outperform other methods of balancing exploration and exploitation (see [Sutton, Barto, 2018] for the results). Clearly, as the simulation time t progresses, traders tend to choose the underexplored strategy as measured by $N_{k,t}$, but the degree to which the underexplored strategy is preferred declines over time thanks to the logarithmic function of time in the numerator. This is helpful, as the uncertainty about the true values of alternative strategies is highest at the start of the learning process and gradually declines over time, as the strategies are tried over and over and larger samples of returns to each strategy are obtained.

2.4. Market Mechanism and Equilibrium Price Determination

Multiagent models of financial markets usually use one of the four types of market clearing mechanisms described in [LeBaron, 2006]: a price impact function (used e.g. in [Gong, Diao,

2022]), random matching, a replicated order book or temporary market equilibrium. In this paper we opt for the latter mechanism, i.e. the temporary market equilibrium, as the other three give rise to additional complexity that we seek to avoid in this stylized model². The temporary equilibrium mechanism is described below.

Both informed and uninformed traders produce their forecasts $Z_{k,t+1}$ of the risky asset's payoff value at time $t+1$. These predicted payoff values discounted at the risk-free rate (due to the risk-neutrality assumption) represent the maximum price each of the trader types k is willing to pay for the risky asset at time t :

$$(7) \quad P_{k,t}^{\max} = \frac{Z_{k,t+1}}{(1+r_f)}.$$

Below the maximum price $P_{k,t}^{\max}$ each trader of a given type k is willing to invest all of his wealth into the risky asset, such that if there was only one trader i controlling the total market wealth the equilibrium price P_t^* would be determined according to:

$$(8) \quad P_t^* = \min\left(P_{i,t}^{\max}, \frac{W_{it}}{N}\right),$$

where N is the constant exogenous supply of the risky asset, as described in Section 2.1.

Clearly, when $P_{i,t}^{\max} < \frac{W_{it}}{N}$ the trader would not be able to invest all of his wealth into the risky asset and his share of wealth invested in it would be given by:

$$(9) \quad \alpha = \frac{P_t^* \cdot N}{W_{it}}.$$

In the general case we have the total market wealth controlled by two groups of traders, informed and uninformed. All traders within the same group make identical payoff predictions which greatly simplifies the equilibrium price setting process. Consider Fig. 1.

We have two groups of traders whose payoff predictions, and therefore $P_{k,t}^{\max}$, are different. Because neither group of traders would buy the risky asset above its estimate of the payoff, the demand functions D_1 and D_2 both have a sharp cutoff at the top. Below the highest estimate of the payoff (in this case corresponding to P_1) three cases are possible.

² For instance, a replicated order book necessarily introduces additional complexity related to the microstructural aspects of the market (e.g. different types of orders, times when orders are active in the order book etc.), which are not the focus of our model but have been found to influence the price formation process. The price impact function mechanism, on the other hand, requires an introduction of a liquidity parameter, which is usually chosen in a somewhat ad hoc fashion. Finally, the random matching mechanism introduces search costs which would make the attribution of our results less straightforward.




Fig 1. An illustrative example of the market mechanism

Two trivial cases arise when the fixed supply of the risky asset equals either the demand of the group of traders with the higher payoff estimate (in this case the equilibrium price is P_2) or the sum of the demands of both groups (in this case the equilibrium price is set at P_4).

A more interesting situation arises when $D_1 < S < D_{1+2}$. This is the case with supply equal to S_3 in Fig. 1. In this case both groups of traders are willing to buy the asset at price P_3 but the supply of the asset is less than the total demand. In this case the equilibrium price is set at P_3 and the number of shares of the risky asset allocated to each trader is determined according to:

$$(10) \quad n_{it} = \alpha \cdot W_{it},$$

where

$$(11) \quad \alpha = \frac{P_t^* \cdot N}{\sum_{i=1}^M W_{it}}.$$

Thus each willing buyer receives the amount of the risky asset proportionate to his share in the total market wealth.

Finally, in the case where the supply of the asset is less than the demand at the highest payoff estimate (P_1 in Fig. 1), the price is set equal to this estimate and the amount of wealth of the willing buyers invested in the risky asset at this price is determined according to Equation 11.

Appendix describes the algorithm governing the behavior of the simulation at each time step t .

2.5. Simulation Parameters and Computational Experiments

We conduct 10 computational experiments with a varying cost of information, C . In each experiment, M agents are initialized with a random amount of capital drawn from a continuous uniform distribution $U(0,1000]$. In each simulation run, the initial payoff value of the risky asset is set to 30\$. All agents are endowed with the same learning mechanism described in Section 2.3. The initial choice of strategy k at $t = 0$ for each agent is determined randomly:

$$(12) \quad \begin{aligned} k_{it} &= 'Inf' \text{ if } U(0,1] > 0.5, \\ k_{it} &= 'Uninf' \text{ if } U(0,1] \leq 0.5. \end{aligned}$$

In the subsequent steps the agents learn the optimal strategy using the RL learning mechanism.

As the uninformed agents need observations of a history of payoff values to make their predictions, the first 50 periods of the fundamental value process are taken as the initialization period and the actual trading starts at the 51st period.

As our model does not include borrowing, and since the cost of information is a monetary constant, agents cannot become informed if their capital level is below the cost of information. Thus, to become informed an agent needs to simultaneously fulfill two conditions:

$$(13) \quad Type_{it} = 'Inf', \text{ if } k_{Uninf,t} < k_{Inf,t}, W_{it} > C.$$

Indeed, we find that for certain values of the cost of information some agents reach quite low values of capital and therefore the constraint in Equation (13) becomes relevant. In cases where $k_{Uninf,t} = k_{Inf,t}$, ties are broken randomly.

The simulation parameters for the 10 computational experiments are presented in Table 1. In total, 500 simulation runs were carried out, 50 for each value of the information cost.

Table 1.

Simulation parameters

Parameters	Values
Cost of information, C	{0; 0.02; 0.04; 0.06; 0.08; 2; 4; 6; 8; 10}
Number of traders, M	100
Number of shares, N	1000
Initial capital of trader i^3	$U(0,1000]$
Initial value of the payoff, F_0	30
Risk-free rate, r_f	0.01/252
Average growth rate of the payoff, μ	0.1/252
St. deviation of the payoff value growth, σ	0.01
Number of simulation steps, T	2200
Exploration constant, c	0.001

³ $U()$ denotes a continuous uniform distribution on a given interval.

3. Results and Discussion

The results of the 10 computational experiments are reported in Fig. 2–11. Fig. 2a–11a show the average mispricing of the risky asset relative to the true payoff on the left vertical axis (the solid black line; the surrounding grey margins show the 10th and 90th percentiles, respectively). As the mispricing is very volatile, resulting in the plot being too jagged and difficult to read, we calculated 300-period moving averages of mispricings in each simulation run and then averaged the resulting data across simulation runs to get the solid line in the charts. The percentiles are calculated on the same moving average time series. Importantly, prior to the moving average transformation, we first calculated the absolute levels of per-period mispricings to avoid negative mispricings offsetting positive ones when the averages are taken.

On the right vertical axes of Fig. 2a–11a we plot the averages of the 300-period moving averages (the same calculation method applies as the one used for the mispricing data) of wealth levels for the informed, the uninformed and both informed and uninformed traders in total.

Fig. 2b–11b show the median number of informed traders across simulation runs (solid line) as well as the 10th and 90th percentiles of this number shown in light grey margins around the line. We report these data for the last 1300 periods of simulation runs, as the numbers of informed traders converge to stable values only after this cutoff.

One can observe several notable features. First, there is a general tendency for the mispricing to increase with rising information costs. This is particularly notable when one compares Fig. 6a to Fig. 7a and 8a, where the information costs make a large jump from 0.08\$ per period per trader to 2\$ and 4\$ per period per trader, respectively. This is associated with the fact that for the information cost 2\$ it takes much longer for the informed traders to overtake the uninformed ones in terms of wealth controlled by them, compared to the case of the information cost of 0.08\$. Predictably, when the market is dominated by the uninformed traders (periods 0 to approximately 1500 of Fig. 7a) the mispricing is high. For the information cost 4\$ (Fig. 8a) the informed traders are driven out of the market completely in the first periods and the mispricing is even higher, as the market is completely controlled by the uninformed traders. This pattern is reproduced in Fig. 9a–11a, as further increasing information costs makes it even harder for the informed traders to survive (indeed they are driven out of the market even earlier). The mispricing converges to a stable value of about 0.8% when the informed traders are driven out of the market. This is not surprising, given that information costs do not affect the wealth growth of the uninformed traders and hence do not affect the price dynamics determined by their trading.

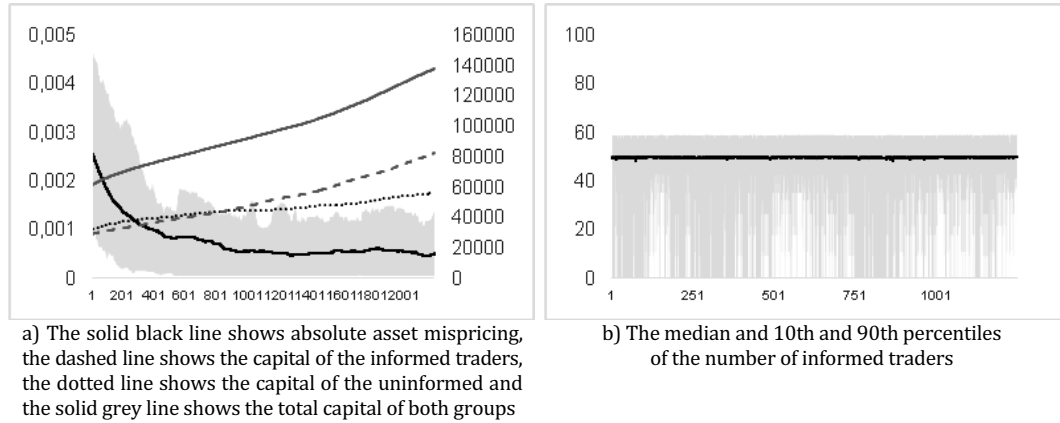


Fig. 2. Information cost $C=0$

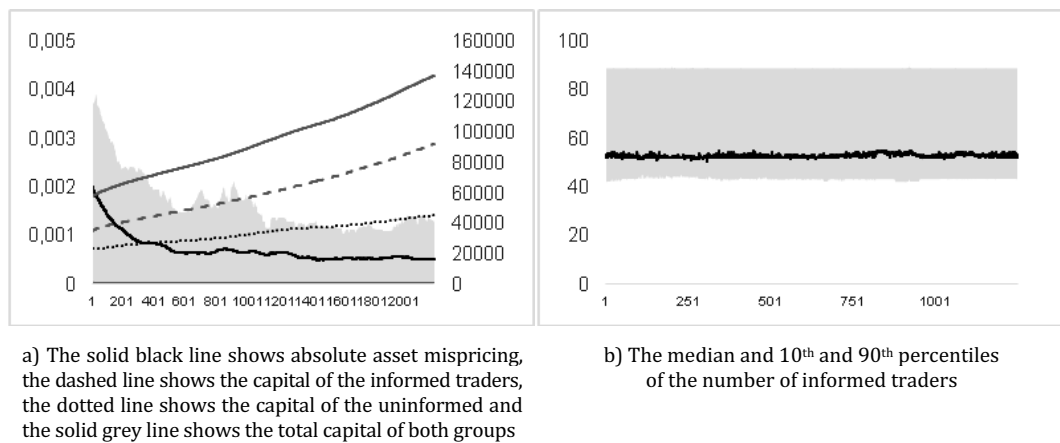




Fig. 3. Information cost $C=0.02$



a) The solid black line shows absolute asset mispricing, the dashed line shows the capital of the informed traders, the dotted line shows the capital of the uninformed and the solid grey line shows the total capital of both groups

b) The median and 10th and 90th percentiles of the number of informed traders


Fig. 4. Information cost $C=0.04$



a) The solid black line shows absolute asset mispricing, the dashed line shows the capital of the informed traders, the dotted line shows the capital of the uninformed and the solid grey line shows the total capital of both groups

b) The median and 10th and 90th percentiles of the number of informed traders


Fig. 5. Information cost $C=0.06$



a) The solid black line shows absolute asset mispricing, the dashed line shows the capital of the informed traders, the dotted line shows the capital of the uninformed and the solid grey line shows the total capital of both groups

b) The median and 10th and 90th percentiles of the number of informed traders

Fig. 6. Information cost $C=0.08$



a) The solid black line shows absolute asset mispricing, the dashed line shows the capital of the informed traders, the dotted line shows the capital of the uninformed and the solid grey line shows the total capital of both groups

b) The median and 10th and 90th percentiles of the number of informed traders

Fig. 7. Information cost $C=2$

The above observations on the survival of informed traders are confirmed by Fig. 2b–11b, which show the numbers of informed traders. With rising costs of information, not only the wealth controlled by the informed traders decreases, but also their numbers, which means that the relative wealth dynamics of the informed and uninformed groups are driven by the different survival rates of these two types and not any within-type competition.


These findings suggest that there is an “equilibrium” level of information cost, somewhere between 2\$ and 4\$ in our case, for which both the informed and the uninformed traders have the same level of fitness, such that neither of the two strategies drives the other out of the market.

However, our main and most surprising finding emerges when one compares Fig. 2a and 3a, showing the results for the information costs 0\$ and 0.02\$, respectively. We can observe that when the cost of information transitions from a positive value to zero, the capital accumulation

by the informed traders in the first half of the simulation is actually slower than that of the uninformed ones. Eventually the informed traders manage to overtake the uninformed ones in terms of accumulated wealth, but paradoxically, this happens much later compared to the case where the cost of information is slightly positive (Fig. 3a–6a). This finding is further strengthened by Fig. 2b and 3b which show the numbers of the informed traders. The 10th percentile of informed traders is actually lower for the case with the information cost of 0\$ than for the cases with $C = [0.02, 0.08]$. In fact, for the slightly positive information costs, the number of the informed traders is never zero at the 10th percentile level. The relative wealth dynamics of the informed and uninformed traders find their reflection in the mispricing dynamics, with the mispricing being higher (at least in the initial periods) for the information cost of 0\$ than 0.02\$.

These findings are in contrast to the predictions of [Grossman, Stiglitz, 1980] who argued that mispricing necessarily decreases with falling information costs. Simultaneously, our results strengthen the findings of [Sciubba, 2005; Gong, Diao, 2022] who find that uninformed traders survive as long as the information costs are bounded away from zero. We show, in fact, that uninformed traders can survive even under zero information costs. Moreover, uninformed traders can ultimately achieve higher levels of wealth under zero information costs compared to the case with slightly positive information costs. Finally, uninformed traders can even initially dominate the market under zero information costs, which is not observed for the slightly positive information costs.


Our findings are in qualitative agreement with experimental results of [Huber, 2007] who find J-shaped returns to information availability. It appears that due to the interaction of informed and uninformed traders growth levels of the informed traders' wealth are non-monotone in the level of information costs, i.e. for already very low value of information costs a marginal decrease in these costs does not necessarily lead to better market performance and a higher wealth share of informed traders and hence to lower mispricing of the assets.



a) The solid black line shows absolute asset mispricing, the dashed line shows the capital of the informed traders, the dotted line shows the capital of the uninformed and the solid grey line shows the total capital of both groups

b) The median and 10th and 90th percentiles of the number of informed traders


Fig. 8. Information cost $C=4$



a) The solid black line shows absolute asset mispricing, the dashed line shows the capital of the informed traders, the dotted line shows the capital of the uninformed and the solid grey line shows the total capital of both groups

b) The median and 10th and 90th percentiles of the number of informed traders


Fig. 9. Information cost $C=6$



a) The solid black line shows absolute asset mispricing, the dashed line shows the capital of the informed traders, the dotted line shows the capital of the uninformed and the solid grey line shows the total capital of both groups

b) The median and 10th and 90th percentiles of the number of informed traders

Fig. 10. Information cost $C=8$



a) The solid black line shows absolute asset mispricing, the dashed line shows the capital of the informed traders, the dotted line shows the capital of the uninformed and the solid grey line shows the total capital of both groups

b) The median and 10th and 90th percentiles of the number of informed traders

Fig. 11. Information cost $C=10$

One explanation for this finding may be provided by the fact that a certain level of uninformed trading in the market is actually beneficial for the informed traders' returns. To give an illustrative example, consider a situation, where the wealth of the uninformed traders is relatively large and they underprice the asset. This is clearly beneficial for the informed traders who would be able to buy the risky asset at a depressed price. On the other hand, when the wealth share of the informed traders in the market is large and that of the uninformed ones is relatively small, this does not give any advantage to the uninformed ones, because the exploitable mispricings will be eliminated by the informed trading. Hence, a higher wealth growth of the uninformed can be associated with a higher wealth growth of the informed, but not vice versa. In the terms of evolutionary finance (see e.g. [Scholl et al., 2021]), there exists a predator-prey relationship between the informed and the uninformed traders in our model, whereby one group's returns benefit from the wealth share of the other group, but not the other way around.

Since in our model traders can switch endogenously from the informed to the uninformed strategy and back based on individual learning, the exact dynamics of the two groups' wealth shares depend on their learning processes, and also on the interaction between individual learning and evolutionary selection. In the terms used by [LeBaron, 2011] complex interactions between active and passive learning determine the agents' survival dynamics, and hence the level of asset mispricing.

One interesting analysis that is not undertaken here and is left for future work is to construct a community matrix of heterogeneous traders in our model and examine quantitatively the relationships between each group's returns and the other group's density. A community matrix is a tool borrowed from ecology and introduced in finance by [Scholl et al., 2021]. The columns and rows of a community matrix represent diverse groups of traders in the market and its entries are the ratios of the growth rates of each group's returns to its own and the other group's growth rates of wealth. Based on the sign of the entries in the matrix, [Scholl et al., 2021] identify a mutualistic and a predator-prey relationships in their model, the first being a situation of mutually beneficial coexistence and the second – of one group unilaterally benefiting from the wealth share of the other. Apart from the type of relationship a community matrix is also able to measure its strength. Indeed, it would be interesting to compare the relative strengths

of relationships between the informed and uninformed traders in our model and those between fundamentalists, chartists and noise traders in the model of [Scholl et al., 2021].

4. Conclusion and Further Research Directions

In this paper we examined the efficiency dynamics of a financial market populated with boundedly rational informed and uninformed traders. The novelty of our approach consists in examining the traditional question of market efficiency in an evolutionary context, in which the wealth shares of informed and uninformed agents are driven by the evolutionary feedback loop from asset mispricing, while asset mispricing is simultaneously driven by the ecology of the market, or more precisely, by the relative and absolute levels of wealth of the informed and the uninformed traders.

Additionally, the traders in our model are boundedly rational and learn the optimal strategy (either the informed or the uninformed one) by interacting with the market, using a simple reinforcement learning algorithm.

We find that under the selection and learning dynamics described above the market price of the risky asset exhibits mispricings across all simulation runs with differing levels of the information cost. Additionally, for higher levels of information costs, the survival probability of the informed traders decreases with rising information costs. As a result, the efficiency of the market is worsened. In this respect our results are in line with the theoretical predictions of [Grossman, Stiglitz, 1980] and the literature on the impossibility of full informational efficiency inspired by them.

However, our main and most surprising finding is that for low levels of information costs a marginal decrease in them does not necessarily lead to better performance of the informed traders, and hence does not necessarily improve market efficiency. In particular, for the case of zero information costs, informed traders initially control even lower levels of capital compared to the case of low but non-zero information costs. As a result, the asset mispricing is worse under zero information costs compared to low but positive information costs. This result sharply contradicts the predictions of [Grossman, Stiglitz, 1980] of a monotone relationship between information costs and market efficiency. This result, however, agrees with some experimental literature (see e.g. [Huber, 2007]), finding J-shaped returns to the availability of information.

Our results have implications both for academic research and regulatory practice. First, the non-trivial interaction between selection dynamics and learning in our model suggest potential avenues for further theoretical research aimed at finding regularities in the joint impact of evolutionary selection and adaptive learning on the survival of heterogeneous agent types and the efficiency of financial markets. Second, our finding of a non-monotone relationship between the level of information costs and market efficiency suggests a need for further empirical research aimed at determining whether the financial markets of today have already hit the point of diminishing returns to the availability of information, as has been suggested by some industry practitioners.

For regulators, our findings provide yet another piece of evidence that it is not easy to decrease the amount of uninformed trading and thus improve market efficiency just by regulating the availability of information, for example by imposing additional disclosure requirements on listed companies. Such measures may not lead to intended outcomes, and taking into account the cost of imposing additional disclosure requirements on issuers, may even lead to deteriora-

tion of total societal welfare. Our results suggest that regulatory decisions of this kind must be based on careful empirical investigation of the market ecology in order to determine at which point increasing availability of information hits diminishing returns and actually hurts market efficiency.

Data Availability. The simulation code and data used for the analysis are available upon request.

Appendix.

The algorithm describing each simulation run

Step 1. At time $t = 0$, each trader is initialized with a random level of wealth; a stochastic path of the risky asset is generated; the information costs are set.

Step 2. The type of each trader is determined either randomly (at time $t = 0$) or according to Equation 6 (for $t > 0$).

Step 3. Both the informed and the uninformed traders make forecasts of the payoff value at $t + 1$ according to Equation 2 and determine the maximum price they are willing to pay for the risky asset according to Equation 7.

Step 4. The capitals of all informed and uninformed traders are aggregated into two groups.

Step 5. The equilibrium price of the risky asset and its allocations to individual traders are determined as described in Section 2.4.

Step 6. The wealth levels of all traders are updated according to Equation 4.

Step 7. The return estimates of the informed and uninformed strategy for each trader are updated according to Equation 5.

Step 8. Steps 2–7 are repeated until the simulation time reaches $t = T$.

* *
*

References

- Alvarez-Ramirez J., Rodriguez E., Espinosa-Paredes G. (2012) Is the US Stock Market Becoming Weakly Efficient over Time? Evidence from 80-year-long Data. *Physica A: Statistical Mechanics and its Applications*, 391, 22, pp. 5643–5647.
- Arthur W.B. (1995) Complexity in Economic and Financial Markets. *Complexity*, 1, 1, pp. 20–25.
- Axtell R.L., Farmer J.D. (2022) *Agent-Based Modeling in Economics and Finance: Past, Present, and Future*. INET Oxford Working Paper, 2022-10.
- Berk J.B. (1997) The Acquisition of Information in a Dynamic Market. *Economic Theory*, 9, pp. 441–451.
- Black F., Scholes M. (1973) The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81, 3, pp. 637–654.
- Cochrane J. (2009) *Asset Pricing: Revised Edition*. Princeton University Press.
- Cox J.C., Ross S.A., Rubinstein M. (1979) Option Pricing: A Simplified Approach. *Journal of Financial Economics*, 7, 3, pp. 229–263.
- Diamond D.W., Verrecchia R.E. (1981) Information Aggregation in a Noisy Rational Expectations Economy. *Journal of Financial Economics*, 9, 3, pp. 221–235.
- Duffie D. (2010). *Dynamic Asset Pricing Theory*. Princeton University Press.
- Evstigneev I.V., Hens T., Schenk-Hoppé K.R. (2009) Evolutionary Finance. *Handbook of Financial Markets: Dynamics and Evolution*, pp. 507–566.
- Evstigneev I.V., Hens T., Schenk-Hoppé K.R. (2002) Market Selection of Financial Trading Strategies: Global Stability. *Mathematical Finance*, 12, 4, pp. 329–339.
- Fama E.F. (1991) Efficient Capital Markets: II. *The Journal of Finance*, 46, 5, pp. 1575–1617.
- Fama E.F. (1970) Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, 25, 2, pp. 383–417.
- Figlewski S. (1982) Information Diversity and Market Behavior. *The Journal of Finance*, 37, 1, pp. 87–102.
- Gong Q., Diao X. (2022) Bounded Rationality, Asymmetric Information and Mispricing in Financial Markets. *Economic Theory*, 74, 1, pp. 235–264.
- Grossman S.J., Stiglitz J.E. (1980) On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 70, 3, pp. 393–408.
- Hasbrouck J. (1993) Assessing the Quality of a Security Market: A New Approach to Transaction-Cost Measurement. *The Review of Financial Studies*, 6, 1, pp. 191–212.
- Hellwig M.F. (1982) Rational Expectations Equilibrium with Conditioning on Past Prices: A Mean-Variance Example. *Journal of Economic Theory*, 26, 2, pp. 279–312.
- Hens T., Schenk-Hoppé K.R. (2005) Evolutionary Stability of Portfolio Rules in Incomplete Markets. *Journal of Mathematical Economics*, 41, 1–2, pp. 43–66.
- Heston S.L., Korajczyk R.A., Sadka R. (2010) Intraday Patterns in the Cross-Section of Stock Returns. *The Journal of Finance*, 65, 4, pp. 1369–1407.
- Ho T.S., Roni M. (1988) Information Quality and Market Efficiency. *Journal of Financial and Quantitative Analysis*, 23, 1, pp. 53–70.
- Huber J. (2007) 'J'-Shaped Returns to Timing Advantage in Access to Information—Experimental Evidence and a Tentative Explanation. *Journal of Economic Dynamics and Control*, 31, 8, pp. 2536–2572.
- Huber J., Kirchler M., Sutter M. (2008) Is More Information Always Better?: Experimental Financial Markets with Cumulative Information. *Journal of Economic Behavior & Organization*, 65, 1, pp. 86–104.
- Kyle A.S. (1989) Informed Speculation with Imperfect Competition. *The Review of Economic Studies*, 56, 3, pp. 317–355.

- LeBaron B. (2011) Active and Passive Learning in Agent-Based Financial Markets. *Eastern Economic Journal*, 37, pp. 35–43.
- LeBaron B. (2006) Agent-Based Computational Finance. *Handbook of Computational Economics*, 2, pp. 1187–1233.
- Lim K., Brooks R. (2011) The Evolution of Stock Market Efficiency over Time: A Survey of the Empirical Literature. *Journal of Economic Surveys*, 25, 1, pp. 69–108.
- Lo A.W. (2004) The Adaptive Market Hypothesis. *The Journal of Portfolio Management*, 30, 5, pp. 15–29.
- Lussange J., Lazarevich I., Bourgeois-Gironde S., Palminteri S., Gutkin B. (2021) Modelling Stock Markets by Multi-Agent Reinforcement Learning. *Computational Economics*, 57, pp. 113–147.
- Lux T. (2009) Stochastic Behavioral Asset-Pricing Models and the Stylized Facts. *Handbook of Financial Markets: Dynamics and Evolution*, pp. 161–215.
- Scholl M.P., Calinescu A., Farmer J.D. (2021) How Market Ecology Explains Market Malfunction. *Proceedings of the National Academy of Sciences*, 118, 26, e2015574118.
- Sciubba E. (2005) Asymmetric Information and Survival In Financial Markets. *Economic Theory*, 25, pp. 353–379.
- Smith J.M., Price G.R. (1973) The Logic of Animal Conflict. *Nature*, 246, 5427, pp. 15–18.
- Tesfatsion L. (2002) Agent-Based Computational Economics: Growing Economies from the Bottom Up. *Artificial Life*, 8, 1, pp. 55–82.
- Verrecchia R.E. (1982) Information Acquisition in a Noisy Rational Expectations Economy. *Econometrica: Journal of the Econometric Society*, 50, 6, pp. 1415–1430.
- Wigglesworth R. (2023) Markets are Becoming Less Efficient, Not More, Says AQR's Clifford Asness // ft.com. Available at: <https://www.ft.com/content/813b3d76-6ef1-427d-a2e0-76540f58a510> (Retrieved: 23.04.2024)

Эволюционная модель эффективности финансового рынка с информационными издержками

Пастушков Алексей Владимирович

Стажер-исследователь, Международная лаборатория экспериментальной и поведенческой экономики, Национальный исследовательский университет «Высшая школа экономики»

11, Покровский б-р, Москва, 109028, Россия

E-mail: apastushkov@hse.ru

В своей классической работе [Grossman, Stiglitz, 1980] показывают, что цены на финансовые активы неизбежно содержат определенную степень неэффективности при ненулевых информационных издержках. Более того, с понижением информационных издержек уровень неэффективности также должен понижаться. Однако несколько современных эмпирических исследований показывают, что реальные финансовые рынки не становятся более эффективными, несмотря на очевидный технологический прогресс и радикально повышающуюся доступность информации. Рыночные цены финансовых активов в высокой мере зависят от относительных уровней капитала гетерогенных инвесторов, взаимодействующих на финансовых рынках. Относительно новое направление исследований по эволюционным финансам предоставило большое количество исследований, показывающих, как развивается динамика цен на финансовые активы под действием эндогенно изменяющихся пропорций гетерогенных трейдеров. Однако лишь очень небольшое количество исследований в эволюционных финансах моделирует влияние информационных издержек на цены финансовых активов в эволюционном контексте. В данной статье мы проводим несколько симуляционных экспериментов с различными информационными издержками и показываем, что даже при нулевых информационных издержках неинформированные инвесторы выживают на рынке, и могут быть доминирующей группой, как минимум в конечных промежутках времени. Таким образом, при изначально низких информационных издержках, маргинальное их понижение необязательно ведет к повышению эффективности рынка. При высоких информационных издержках наши результаты в целом совпадают с результатами [Grossman, Stiglitz, 1980] в том смысле, что повышающиеся издержки на информацию ведут к уходу информированных трейдеров с рынка и ухудшению эффективности. В работе также обсуждаются выводы из полученных результатов для практики регулирования финансовых рынков.

Ключевые слова: эффективность рынка; ограниченная рациональность; мультиагентная модель; эволюционные финансы; информационные издержки; агент-ориентированное моделирование.

JEL Classification: C63, D82, G14.