Journal of Economic Interaction and Coordination
https://doi.org/10.1007/s11403-024-00424-0

REGULAR ARTICLE

From zero-intelligence to Bayesian learning: the effect of rationality on market efficiency

Daniele Giachini[1] · **Shabnam Mousavi**[2] · **Matteo Ottaviani**[3,4]

Received: 15 February 2024 / Accepted: 5 August 2024

Abstract
In this paper, we investigate the relationship between individual rationality and price informative efficiency studying a prediction market model where agents repeatedly bet on the occurrence of a binary event following their subjective beliefs. We define individual rationality in terms of the amount of past observations used to update beliefs. In this way, a wide spectrum of rationality levels emerges, ranging from zero-intelligence to Bayesian learning. We show that the relationship between individual rationality and price informative efficiency is nonlinear and U-shaped. We argue that the results emerge from the particular interaction of two evolutionary forces operating at different levels: the market selection mechanism that moves wealth toward more accurate agents and the individual learning process that moves posterior probabilities over models depending on observed realizations.

Keywords Information efficiency · Updating beliefs · Structure of the market · Evolution of markets · Market selection

JEL Classification C60 · D53 · D81 · D83 · G11 · G12

Shabnam passed away on October 12th, 2023. Before that, we had carried on the present work altogether. We dedicate this manuscript to her.

✉ Matteo Ottaviani
matteo.ottaviani@sns.it

1 Institute of Economics & Department L'EMbeDS, Scuola Superiore Sant'Anna, Piazza Martiri della Libertà 33, 56127 Pisa, PI, Italy

2 Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany

3 Scuola Normale Superiore, Piazza dei Cavalieri 7, 56126 Pisa, PI, Italy

4 German Centre for Higher Education Research and Science Studies (DZHW), Schützenstraße 6a, 10117 Berlin, Germany

Published online: 29 August 2024

1 Introduction

Market efficiency is one of the most debated topics in the field of finance. Testing whether the efficient market hypothesis (Fama 1970; LeRoy 1989) holds, identifying possible violations, and proposing theories to explain them have been occupying the research agendas of many financial economists and drove the development of the field of behavioral finance (see, e.g., Barberis and Thaler 2003; Hirshleifer 2015). A widespread perspective is that market efficiency is a consequence of full information and rational expectations, with the equilibrium price identifying the efficient one. Hence, within this framework, the use of heuristics to form expectations moves away the equilibrium price from the efficient one (see, e.g., Barberis et al. 1998; Daniel et al. 1998; Bottazzi and Giachini 2022; Antico et al. 2024). It follows that, as we move further away from full rationality, the larger the loss in terms of market efficiency.

Such a simple connection between individual rationality and aggregate outcomes, such as pricing and market efficiency, is not accurate. Indeed, simple fast and frugal rules may actually provide superior outcomes in complex environments (Gigerenzer and Brighton 2009; Gigerenzer and Gaissmaier 2011; Kirman 2010; Mousavi et al. 2015). Gode and Sunder (1993) prove that a double-auction market composed by zero-intelligence agents respecting budget constraints produces efficiency results that match competitive equilibrium allocations. Giachini (2021) shows that, under heterogeneous and incorrect beliefs, the evolutionary market selection forces operating in an economy populated by agents following a simple trading rule (as in Bottazzi et al. 2018, 2019) can generate a pricing performance more in line with efficient (fully rational) levels than the one obtained in an economy where agents inter-temporally maximize their utility (as in Sandroni 2000; Blume and Easley 2006, 2009). In heterogeneous agent models of financial markets, rational traders (fundamentalists) usually act as a price stabilizing force, while chartists tend to destabilize (Hommes 2006). However, Gardini et al. (2022) show that price destabilization always emerges when fundamentalists have a different perception of the fundamental price. In addition, Naimzada and Ricchiuti (2009) show that price destabilization also emerges in a model where only fundamentalists with different beliefs are active. At the same time, Naimzada and Ricchiuti (2014) show that belief heterogeneity can have a positive effect on price stability if the market maker adopts a multiplicative price setting mechanism. Such an ambiguous relationship between individual characteristics, market structure, and aggregate outcomes also emerges from the analysis of Bargigli (2021), where an optimizing monopolistic market maker is assumed in an otherwise standard model à la Brock and Hommes (1998).

A general picture on the relationship between individual rationality and aggregate outcomes is offered by Fehr and Tyran (2005). They report five ways in which markets can somehow compensate for deviations from full rationality at the individual level and eventually produce efficient aggregate outcomes. The first is aggregation: random individual deviations from rationality cancel out in the aggregate. The second is learning: agents learn from experience and eventually become rational. The third concerns interaction rules and constraints: market structure and budget constraints (such as in Gode and Sunder 1993) let rational aggregate outcomes emerge from zero-intelligence agents. The fourth is market selection: in competitive markets, rational individuals gain wealth at the expenses of irrational individuals; hence, the latter ones are driven out of the market and do not influence prices. The fifth is the relative position in the curves of demand and supply: while irrational agents tend to hold extreme positions, rational individuals hold marginal ones; thus, prices reflect rational choices. In this paper, we merge the second, third, and fourth mechanism (learning, market interaction, and market selection) to show that the relationship between individual rationality and price informative efficiency (that is, market efficiency) is nonlinear and U-shaped.

In particular, we study a repeated prediction market model in discrete time where agents repeatedly bet on the occurrence of a binary stochastic event following the Kelly rule, i.e., they "bet their beliefs" (Beygelzimer et al. 2012; Bottazzi and Dindo 2013; Kets et al. 2014; Bottazzi and Giachini 2017, 2019a, b). While agents follow the same investment behavior, they have heterogeneous beliefs that are generated by means of a learning process over two misspecified models: one optimist and one pessimist, that are agent-specific. In our framework model, misspecification represents the complexity of the environment in which agents operate: It captures the fact that in a complex world, the models agents have in mind are simplistic representations of the true data generating process (Bottazzi et al. 2023). We assume that the learning protocol of each agent depends on the number n of previous data points (i.e., states of nature), it can observe to update conditional probabilities. Such a number, equal for each agent, represents the degree of individual rationality, since a low value of n means that agents are actually discarding observations to form their beliefs, while a large value of n means that agents use a lot of past observations to inform their choices. More precisely, in the extreme case $n = 0$, agents do not learn and maintain the initial beliefs that the true probability can be obtained as the arithmetic average of the two models. In the opposite case, $n = \infty$, agents are Bayesian learners (the cornerstone of rational learning) and are able to asymptotically recognize which one between the two models is the most accurate (Berk 1966). We study the emerging prices and we put into relationship the degree of rationality n with a measure of price informative efficiency. In particular, given the probabilistic nature of prediction market prices, the natural measure of price informative efficiency is the opposite of the average relative entropy of prices with respect to the true probability (see, e.g., Bottazzi and Giachini 2019b). Mixing analytical results with extensive numerical simulations in a setting where models are randomly assigned to agents, we observe a robust U-shaped relationship between individual rationality and average price informative efficiency. In line with the literature on fast and frugal heuristics, the zero-intelligence scenario presents performances that are close to the rational (Bayesian) one.

The key intuition driving our results is the interaction between two evolutionary forces, individual learning, and market selection, under the market interaction protocol. Indeed, following Blume and Easley (1993), Beygelzimer et al. (2012), and Giachini (2021), the structure of subsequent interactions generates a form of Bayesian learning via market selection, where markets try to "learn" which agent is the most accurate. That is, when the event occurs, the agents whose beliefs make them heavily bet on the occurrence of the event, will see their wealth increase at the expenses of those betting against the event. If the event does not occur, the opposite happens. Such a redistribution process let wealth concentrate in those agents that are the most correct on average. Moreover, given the particular structure of portfolios, prices emerge as wealth-weighted convex combinations of individual beliefs. Thus, under the market learning metaphor, the agent's wealth share resembles the probability that such an agent is the one with the most accurate model, while prices resemble the market estimates of the probabilities attached to states of nature. When agents keep their beliefs fixed, the evolutionary learning process performed by market selection operates among models belonging to the same family with respect to the truth, since it follows an independent and identically distributed (i.i.d.) process. Thus, given this form of misspecification, the price informative efficiency results rather high. When agents start using the last (few) observations to update beliefs, they induce some spurious probabilistic structure in prices that are weakly related to the true data generating process. This mismatch results that detrimental for informative efficiency and selection cannot compensate for it. As agents start using a sufficiently high number of last observations, they end up giving (almost) full weight to the model closest to the truth (in relative entropy) between the optimist and the pessimist one. Indeed, in the Bayesian limit, they eventually converge on the most accurate between the two. In that case, the evolutionary market selection process ends up, again, operating among models belonging to the same family with respect to the truth and high efficiency is restored.

Our paper is organized in the following way. In Sect. 2, we present a repeated prediction market model, which is the general framework of our study, and the limited memory Bayesian learning, that is the key learning process of our investigation. In Sect. 3, we derive the price informative efficiency of a prediction market that we will employ as the performance measure for each degree of rationality investigated. In Sect. 4, we analytically explore our measure of price informative efficiency for three benchmark levels of rationality. In Sect. 5, we explore numerically the generic cases of different degrees of rationality. In Sect. 6, we present and discuss the results. Section 7 concludes.

2 The model

Following Beygelzimer et al. (2012), Kets et al. (2014), Bottazzi and Giachini (2017, 2019a, 2019b), we study a repeated prediction market model in discrete time where N agents repeatedly bet on the occurrence of a binary event. The realization of the event at time t is indicated with the random variable $s_t \in \{1, 0\}$: $s_t = 1$ means that the event has occurred while $s_t = 0$ means that it has not occurred. The probability of observing the event is constant over time, that is $Pr\{s_t = 1\} = \pi^*$ with $\pi^* \in (0, 1)$. The market is complete in the sense that two securities are available to wager in every period. The first security traded at time t pays 1 dollar if $s_t = 1$ and 0 otherwise. The second security traded at time t, instead, pays 1 dollar if $s_t = 0$ and 0 otherwise. Each agent i has an initial wealth equal to $W_0^i > 0$ that becomes W_t^i at the end of time t due to trading. The total initial wealth in the market is normalized to one, and all the available wealth is invested. Defining α_t^i as the share of wealth invested by agent i at time t in the security paying 1 if $s_t = 1$, agent i buys an amount $W_{t-1}^i \alpha_t^i / P_{1,t}$ of the first security and an amount $W_{t-1}^i (1 - \alpha_t^i)/P_{2,t}$ of the second security, where $P_{1,t}$ and $P_{2,t}$ are the prices at time t of the two securities. Hence, the timeline of events in each period t is as follows: at the beginning of the period each agent i has a wealth W_{t-1}^i from previous rounds; every agent decides how to allocate its wealth between the two securities; markets open and prices are set; the event is realized and securities pay off. It follows that the wealth of agent i evolves according to

$$W_t^i = \begin{cases} W_{t-1}^i \frac{\alpha_t^i}{P_{1,t}} & \text{if } s_t = 1, \\ W_{t-1}^i \frac{1-\alpha_t^i}{P_{2,t}} & \text{if } s_t = 0. \end{cases} \tag{1}$$

Prices are fixed in temporary equilibrium according to market clearing conditions and assuming unitary supply. Thus, we have

$$P_{1,t} = \sum_{i=1}^{N} \alpha_t^i W_{t-1}^i \text{ and } P_{2,t} = \sum_{i=1}^{N} (1 - \alpha_t^i) W_{t-1}^i. \tag{2}$$

The total wealth in the market is one in every period, and we can simplify the notation defining $P_t = P_{1,t}$ and setting $P_{2,t} = 1 - P_t$. In so doing and given the payoff structure of the securities, P_t can be interpreted as the probability the "market" assigns to the realization of the event at time t. Finally, we assume that agents follow the *Kelly rule* to bet. That is, agent i invest in each security proportionally to the probability it assigns to the state in which the security pays out. In formal terms, we define π_t^i the probability agent i assigns at the beginning of time t to the realization of the event $s_t = 1$, such that Kelly betting simply means $\alpha_t^i = \pi_t^i$. Alternatively, we may get the same investing behavior assuming that, in each period, agents myopically maximize their expected logarithmic utility under their subjective probabilities as in Bottazzi and Dindo (2013). We further assume that each agent i does not know the data generating process (i.e., π^*) and tries to learn it by means of two misspecified models. A model is a probability distribution $(\pi, 1 - \pi)$ on $\{1, 0\}$ and we assume that agent i relies on an *optimist* model, $(\pi_o^i, 1 - \pi_o^i)$, and a *pessimist* model, $(\pi_p^i, 1 - \pi_p^i)$. Those are obtained assuming $0 < \pi_p^i < \pi^* < \pi_o^i < 1$. Moreover, individual models are heterogeneous across the population: each agent has its own couple of optimistic and pessimist models.[1] Agents attempt to learn the true probability π^* by adopting a limited memory Bayesian approach. In the next paragraph, we provide a formalization of such a learning process.

Limited memory Bayesian learning: Broadly speaking, we define the limited memory Bayesian learning as a form of Bayesian learning in which the agents

[1] At the technical level, we will produce such a form of heterogeneity randomly and uniformly drawing one optimistic and one pessimistic model for each agent.

cannot store in their memory all the available information (i.e., previous realizations of the binary event) but only a limited amount.[2] The informative set affects the way in which beliefs are computed and updated and this, in turn, affects the dynamics of wealth and prices.

In line with the literature on Bayesian learning (see e.g., (Epstein et al. 2010; Beygelzimer et al. 2012; Massari 2020; Bottazzi et al. 2023), we assume that agents build their probabilistic predictions for the occurrence of the event at time t as a weighted average of the probabilities suggested by the two models. Moreover, since the individual probability of agent i at time t, π_t^i, is now a function of its rationality level n, we make explicit such a dependence adding the superscript n. Hence, one has

$$\pi_t^{i,n} = w_{t-1}^{i,n}\pi_o^i + (1 - w_{t-1}^{i,n})\pi_p^i, \tag{3}$$

where $w_t^{i,n}$ is the probability agent i assigns at the end of time t to the optimist model to be the correct one given the observation of the last n realizations of the process $\{s_t\}$. Since there are only two models, it follows that the probability agent i assigns at the end of time t to the pessimist model to be true given the observation of the last n realizations of $\{s_t\}$ is $1 - w_t^{i,n}$. To formally define the dynamics of such probabilities, we introduce the Bayesian posterior probability of the optimist model to be the true one given a prior $w \in (0,1)$ and the realization $s \in \{1,0\}$. It reads

$$\mathcal{B}(w,s) = s\frac{w\pi_o^i}{w\pi_o^i + (1-w)\pi_p^i} + (1-s)\frac{w(1-\pi_o^i)}{1 - w\pi_o^i - (1-w)\pi_p}, \tag{4}$$

such that, calling $\sigma_{t,n} = (s_{t-n+1}, \dots, s_t)$, we have the recursive definition of the limited memory Bayesian posterior given the observation of $\sigma_{t,n}$, that is

$$\mathcal{B}^n(\sigma_{t,n}) = \mathcal{B}\big(\mathcal{B}^{n-1}\big(\sigma_{t-1,n-1}\big), s_t\big), \tag{5}$$

with $\mathcal{B}^0(\sigma_{t-n,0}) = w_0^{i,n}$ acting as initial prior. In what follows, we shall assume a non-informative initial prior: $w_0^{i,n} = 0.5\ \forall i, n$. Thus, the dynamics of $w_t^{i,n}$ can be described as

$$w_t^{i,n} = \begin{cases} \mathcal{B}(w_{t-1}^{i,n}, s_t) = w_{t-1}^{i,n}\left(s_t\frac{\pi_o^i}{\pi_t^{i,n}} + (1-s_t)\frac{1-\pi_o^i}{1-\pi_t^{i,n}}\right) & \text{if } t \le n, \\ \mathcal{B}^n(\sigma_{t,n}) & \text{if } t > n, \\ 0.5 & \text{if } n = 0. \end{cases} \tag{6}$$

The number n of previous observations used to update beliefs may be considered as a measure of rationality in the learning process. The case $n = 0$ represents zero-intelligence learners. Indeed, in such a case, no learning occurs at all: the information set is empty and the non-informative initial priors are never updated. As a

[2] Bottazzi et al. (2023) introduce the idea of limited memory Bayesian learning focusing on the special case in which agents can use only the last realization of the data generating process.

consequence, the predicting probability is constant over time: $\pi_t^i = (\pi_o^i + \pi_p^i)/2\ \forall t$. As n grows, agents use an increasing number of past observations to update their beliefs and approach the Bayesian case for $n = \infty$.

3 Learning performance, market selection, and informative efficiency

As previously highlighted and according to the literature (see, e.g., Arrow et al. 2008), in a prediction market prices can be thought as probability distributions. Hence, in line with the literature (Blume and Easley 2009; Bottazzi and Giachini 2019b; Dindo and Massari 2020), we choose the opposite of the average relative entropy of prices with respect to the truth as our measure of informative efficiency. In formal terms, given a distribution $(\pi, 1-\pi)$, its relative entropy with respect to the true distribution $(\pi^*, 1-\pi^*)$ is defined as

$$D(\pi^*||\pi) = \pi^* \log\frac{\pi^*}{\pi} + (1-\pi^*)\log\frac{1-\pi^*}{1-\pi}.$$

This quantity measures the amount of information lost when approximating $(\pi^*, 1-\pi^*)$ with $(\pi, 1-\pi)$. Hence, it is a measure of "how different" a distribution is with respect to another one. The relative entropy (also known as the Kullback Leibler divergence) is not a distance – it is not symmetric and it does not respect the triangular inequality – but it is always non-negative and reaches zero if and only if $\pi = \pi^*$. Hence, $-D(\pi^*||\pi)$ can be considered a measure of similarity that increases as π gets “closer” to π^*. Following the literature (Dindo and Massari 2020), in a dynamic setting where the distribution changes over time, that is $\{(\pi_t, 1-\pi_t), t = 1, 2, \dots\}$, one should consider the infinite time average of the relative entropy. We define it as

$$\overline{D}(\pi^*||\pi) = \lim_{t\to\infty}\frac{1}{t}\sum_{\tau=1}^{t}\left(\pi^*\log\frac{\pi^*}{\pi_\tau} + (1-\pi^*)\log\frac{1-\pi^*}{1-\pi_\tau}\right). \tag{7}$$

Applying it to the sequence of prices generated by the prediction market and changing its sign, we obtain our measure of price informative efficiency. In particular, calling P^n the sequence of prices obtained from a prediction market where agents observe the last n realizations to build their beliefs, the informative efficiency of such an economy is

$$\mathcal{E}^n = -\overline{D}(\pi^*||P^n). \tag{8}$$

To characterize $\mathcal{E}^n$, we exploit the fact that prediction market prices have a natural probabilistic interpretation and the market selection process resembles Bayesian learning (Beygelzimer et al. 2012; Giachini 2021). Hence, define

$$\mathcal{P}_t^n = \prod_{\tau=1}^{t} \left(s_\tau P_\tau^n + (1 - s_\tau)(1 - P_\tau^n) \right), \tag{9}$$

such a quantity can be interpreted as the likelihood the market assigns to the particular sequence of states that have been realized up to time t. Along the same lines, we define the likelihood assigned to the realizations up to t by agent i as

$$\Pi_t^{i,n} = \prod_{\tau=1}^{t} \left(s_\tau \pi_\tau^{i,n} + (1 - s_\tau)(1 - \pi_\tau^{i,n}) \right).$$

Using $W_t^{i,n}$ to indicate the wealth at time t of agent i in an economy with rationality level n, its evolution can be rewritten as

$$W_t^{i,n} = \frac{s_t \pi_t^{i,n} + (1 - s_t)(1 - \pi_t^{i,n})}{s_t P_t^n + (1 - s_t)(1 - P_t^n)} W_{t-1}^{i,n}. \tag{10}$$

Hence, from (9) and the market clearing condition (2), one has

$$\mathcal{P}_t^n = \mathcal{P}_{t-1}^n \sum_{i=1}^{N} W_{t-1}^{i,n} \left(s_t \pi_t^{i,n} + (1 - s_t)(1 - \pi_t^{i,n}) \right)$$

and, iteratively substituting with (10), one obtains

$$\mathcal{P}_t^n = \sum_{i=1}^{N} W_0^{i,n} \Pi_t^{i,n}. \tag{11}$$

Equation (11), together with equation (10), perfectly resembles one of the standard ways of representing the Bayesian learning process (see Massari 2017, 2020; Marinacci and Massari 2019; Bottazzi et al. 2023). Thus, following the procedure used by Massari (2020) and Bottazzi et al. (2023), call $\overline{\Pi}_t^n = \max_{i \in \{1,\dots,N\}} \Pi_t^{i,n}$ and $j_{n,t}$ one of the agents such that $\Pi_t^{j_{n,t},n} = \overline{\Pi}_t^n$, in this way one has

$$W_0^{j_{n,t},n} \overline{\Pi}_t^n \le \mathcal{P}_t^n \le \overline{\Pi}_t^n.$$

Dividing on any side by the true likelihood $\Pi_t = (\pi^*)^{t_1} (1 - \pi^*)^{t - t_1}$ (with t_1 the number of times the event has occurred), taking logarithms, and multiplying on any side by $-t^{-1}$, it is

$$-\frac{1}{t} \log W_0^{j_{n,t},n} \ge \frac{1}{t} \log \frac{\Pi_t}{\mathcal{P}_t^n} - \frac{1}{t} \log \frac{\Pi_t}{\overline{\Pi}_t^n} \ge 0.$$

Taking the limit for $t \to \infty$ and invoking Lemma A.1 of Bottazzi et al. (2023), under the assumption that the limits involved exist, one almost surely obtains[3]

[3] Cf. Proposition 3.1 of Bottazzi et al. (2023).

$$0 = \lim_{t \to \infty} \left(\frac{1}{t} \log \frac{\Pi_t}{\mathcal{P}^n_t} - \frac{1}{t} \log \frac{\Pi_t}{\overline{\Pi}^n_t} \right) = \overline{D}(\pi^* || P^n) - \overline{D}(\pi^* || \pi^{i^\star_n, n})$$

where $i^\star_n$ indicates an agent whose (infinite time) average relative entropy equals the minimum level in the population. Finally, we can express the price informative efficiency of a prediction market characterized by a rationality level n as

$$\mathcal{E}^n = -\overline{D}(\pi^* || \pi^{i^\star_n, n}).$$

Notice that such a conclusion is a direct consequence of the evolutionary market selection dynamics taking place among heterogeneous agents. Indeed, the previous computations mirror the process of wealth reallocation operating in the market. Wealth moves toward those who are able to make the best predictions, hence, asymptotically, only the agents who have the most accurate beliefs survive (see, e.g., Blume and Easley 2009; Beygelzimer et al. 2012; Kets et al. 2014; Bottazzi and Giachini 2019b). Such an accuracy is transmitted to prices that eventually are as accurate as the best agents in the market.

4 Mathematical exploration of specific cases

Although characterizing the average relative entropy of the agents is difficult for a generic n and we will rely on numerical exercises to do it, there are some specific (but significant) cases in which one can derive analytical results: $n = 0, 1, \infty$.

In the case of zero-intelligence agents ($n = 0$), for any agent i it is

$$\overline{D}(\pi^* || \pi^{i,0}) = D\Big(\pi^* || (\pi^i_o + \pi^i_p)/2\Big),$$

almost surely, thus, one has

$$\mathcal{E}^0 = -D\Big(\pi^* || (\pi^{i^\star_0}_o + \pi^{i^\star_0}_p)/2\Big).$$

In the case $n = 1$, agents use only the last realized state to update their weights from the initial prior. Hence, one has

$$w^{i,1}_t = \begin{cases} \dfrac{\pi^i_o}{\pi^i_o + \pi^i_p} & \text{if } s_t = 1, \\ \dfrac{1 - \pi^i_o}{2 - \pi^i_o - \pi^i_p} & \text{if } s_t = 0, \end{cases}$$

that implies

$$\pi_t^{i,1} = \begin{cases} \frac{(\pi_o^i)^2+(\pi_p^i)^2}{\pi_o^i+\pi_p^i} & \text{if } s_{t-1}=1, \\ \frac{(1-\pi_o^i)\pi_o^i+(1-\pi_p^i)\pi_p^i}{2-\pi_o^i-\pi_p^i} & \text{if } s_{t-1}=0. \end{cases}$$

Hence, similarly to the case investigated in Bottazzi et al. (2023), agents' beliefs show a Markov structure. Invoking the Strong Law of Large Numbers, for any agent i one obtains, almost surely,

$$\overline{D}(\pi^*||\pi^{i,1}) = \pi^* D\left(\pi^* \,\middle\|\, \frac{(\pi_o^i)^2+(\pi_p^i)^2}{\pi_o^i+\pi_p^i}\right) + (1-\pi^*) D\left(\pi^* \,\middle\|\, \frac{(1-\pi_o^i)\pi_o^i+(1-\pi_p^i)\pi_p^i}{2-\pi_o^i-\pi_p^i}\right)$$

and this implies

$$\mathcal{E}^1 = -\pi^* D\left(\pi^* || \frac{(\pi_o^{i_1^\star})^2+(\pi_p^{i_1^\star})^2}{\pi_o^{i_1^\star}+\pi_p^{i_1^\star}}\right) - (1-\pi^*) D\left(\pi^* || \frac{(1-\pi_o^{i_1^\star})\pi_o^{i_1^\star}+(1-\pi_p^{i_1^\star})\pi_p^{i_1^\star}}{2-\pi_o^{i_1^\star}-\pi_p^{i_1^\star}}\right).$$

If $n=\infty$, then agents are Bayesians. Thus, according to Proposition 3.1 of Bottazzi et al. (2023), agents are asymptotically as accurate as the most accurate model between the optimist and the pessimist one. Hence, one generically and almost surely has

$$\overline{D}(\pi^*||\pi^{i,\infty}) = \min_{k\in\{o,p\}} D(\pi^*||\pi_k^i)$$

and this directly implies that

$$\mathcal{E}^\infty = -\min_{k\in\{o,p\}} D(\pi^*||\pi_k^{i_\infty^\star}).$$

As one can notice, our measure of price informative efficiency changes significantly in the three cases. Moreover, ranking the different economies with respect to such a quantity can be difficult, since it may depend on the assumptions on the number of agents, the value of the true probability, and the distribution of models in the population. Hence, to provide some reference points, we investigate some special and opposite case.

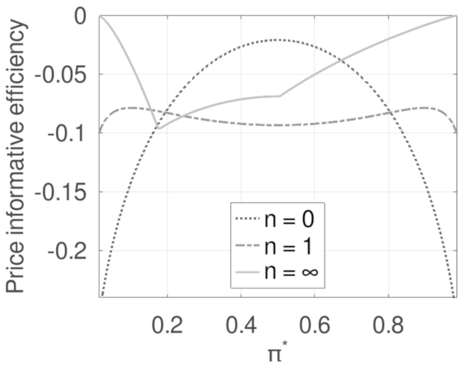

Fig. 1 Expected price informative efficiency for $\pi^* \in [0.015, 0.985]$ and $n \in \{0, 1, \infty\}$ in the case $N = 1$

4.1 Very large versus single-agent economies

Here, we consider two extreme cases. The first case is the one in which N is so large that any combination of the optimist and the pessimist model can be found in the population. Under such an assumption, there exist an agent i such that $(\pi^i_o + \pi^i_p)/2 = \pi^*$. That implies $\mathcal{E}^0 = 0$: with zero-intelligence agents, such an economy reaches the maximum possible level of price informative efficiency. Along the same lines, in the case $n \to \infty$ we have that there exists an agent i such that either π^i_o or π^i_p is very close to π^*. Thus, $\mathcal{E}^\infty \simeq 0$: the economy approaches the maximum level of price informative efficiency also with Bayesian agents. In the Markovian case $n = 1$, in order to reach the maximal efficiency performance, there exists an agent i such that

$$\begin{cases} \frac{\pi^i_o}{\pi^i_o+\pi^i_p}\pi^i_o + \frac{\pi^i_p}{\pi^i_o+\pi^i_p}\pi^i_p = \pi^*, \\ \frac{1-\pi^i_o}{2-\pi^i_o-\pi^i_p}\pi^i_o + \frac{1-\pi^i_p}{2-\pi^i_o-\pi^i_p}\pi^i_p = \pi^*. \end{cases}$$

Notice that the system of equations has solution if and only if $\pi^i_o = \pi^i_p$. Thus, assuming that there exists an agent such that $\pi^i_o \simeq \pi^* \simeq \pi^i_p$, one can obtain $\mathcal{E}^1 \simeq 0$. These simple computations show that, in an extremely large population where any combination of optimist and pessimist model can be found, we expect the price informative performances to converge to full efficiency. At the same time, it is worth to notice that, while in the zero-intelligence economy, it is enough to have one agent whose models are correct on average, in the Bayesian economy, there exists an agent with one between the optimist and the pessimist model (almost) correct and in the Markovian economy, there exists an agent with both the optimist and the pessimist model (almost) correct.

The opposite extreme case is the one in which there is a single agent in the economy. Looking at the efficiency measures, one immediately notices that, depending on the relative positions of the pessimist and optimist models with respect to the truth, several possible scenarios are possible. Hence, we assume that π_p^1 and π_o^1 are randomly drawn from uniform distributions over, respectively, $(0.01, \pi^* - 0.001)$ and $(\pi^* + 0.001, 0.99)$. In this way, we can obtain a measure of average performance looking at the expected price informative efficiency for different values of π^* and n. To do that, we numerically evaluate

$$\int_{0.01}^{\pi^*-0.001} \int_{\pi^*+0.001}^{0.99} \frac{\mathcal{E}^n}{(\pi^* - 0.011)(0.989 - \pi^*)} d\pi_o^1 \, d\pi_p^1$$

for $\pi^* \in [0.015, 0.985]$ and $n \in \{0, 1, \infty\}$. We report the results in Fig. 1. As one can notice, a stable ranking among the three economies does not appear. For extreme values of π^* exploiting more information–that is, letting the agents be more rational–produces an advantage in price informative efficiency on average. This is expected: as π^* approaches the boundaries, drawing a model very close to the truth becomes more likely. This benefits Bayesian learning and disadvantages zero-intelligence: while a Bayesian learner selects an (almost) correct model, a zero-intelligence agent continuously mixes an (almost) correct model with a (possibly heavily) incorrect one. The situation is rather different for values of π^* around 0.5. In such a case, a non-linear pattern appears: The zero-intelligence economy is the most efficient (on average), the Bayesian economy occupies the second position, and the Markovian one shows the worst performance. The use of little information, as in $n = 1$, produces the best result in the average price informative efficiency only in a narrow region at the left of $\pi^* = 0.2$. Overall, it appears to show rather stable average efficiency levels across the support of π^*.

The analysis presented here highlights that, on the one hand, we should expect an increase in price informative efficiency as the number of agents populating the economy grows. On the other hand, when N is small, a U-shaped relationship emerges for values of π^* around 0.5 but vanishes as we move toward the boundaries. Indeed, a non-stable ranking among different levels of n can be observed. This consideration, however, is based on the special case in which a single agent populates the economy. Hence, in such a case, the market selection mechanism has been shut down. This may have an effect on the relative performances since, as mentioned above, for generic cases, it is more "likely" to have an agent with the average model (almost) correct or with one (almost) correct model than finding an agent with both models (almost) correct. Such a difference should emerge when N is sufficiently large (but not too large): The evolutionary wealth dynamics taking place among heterogeneous agents should let a U-shaped relationship between the level of rationality n and the price informative efficiency emerge for a wider set of true probabilities. We test for such a conjecture in the next section by means of numerical exercises.

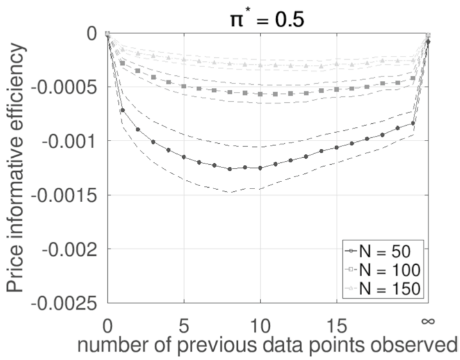

Fig. 2 Montecarlo estimates of the average price informative efficiency as a function of the number of previous data points observed n for $N = 50, 100, 150$. Dashed lines represent 95% confidence bounds. The symbol ∞ on the horizontal axis indicates the Bayesian case

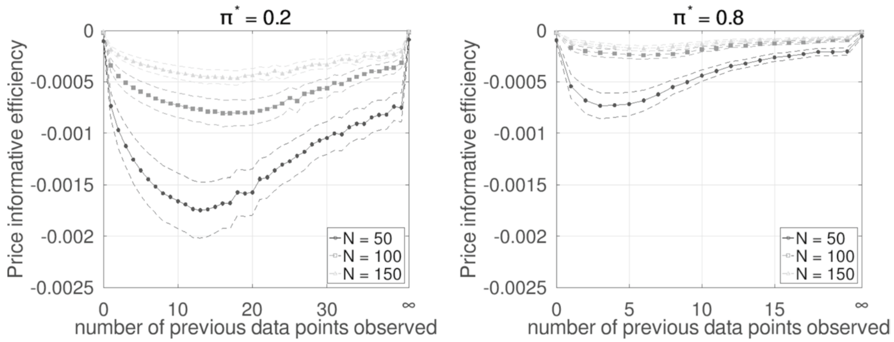

Fig. 3 Montecarlo estimates of the average price informative efficiency as a function of the number of previous data points observed n for $N = 50, 100, 150$. Left: $\pi^* = 0.2$. Right: $\pi^* = 0.8$. Dashed lines represent 95% confidence bounds. The symbol ∞ on the horizontal axis indicates the Bayesian case

5 Numerical exploration of generic cases

To test the relationship between rationality level and price informative efficiency by means of numerical simulations, we proceed in the following way. We fix $\pi^* = 0.5$ and, for each level of $N \in \{50, 100, 150\}$, we randomly draw N pairs of pessimist and optimist model probabilities from their respective supports, $(0, \pi^*)$ and $(\pi^*, 1)$ (one pair per each agent). Then, the dynamics in eq. (6) are evolved for $T = 10000$ time steps for every agent and for every $n = 1, 2 \ldots, 20$. We record the average relative entropy of each agent and select the smallest one. The negative of such a quantity is an estimate of the price informative efficiency along that trajectory.[4] We add to those quantities the price informative efficiency levels computed analytically for the cases $n = 0$ and $n = \infty$. The procedure is repeated for $MC = 250$ independent replicas such that, averaging over the runs, we can get rid of the variability induced by the random choice of models and obtain an estimate of the average price informative efficiency. In Fig. 2, we report the results of such an exercise, surrounding the

[4] Extensive numerical simulations show that the chosen value of T is sufficient to obtain reliable estimates.

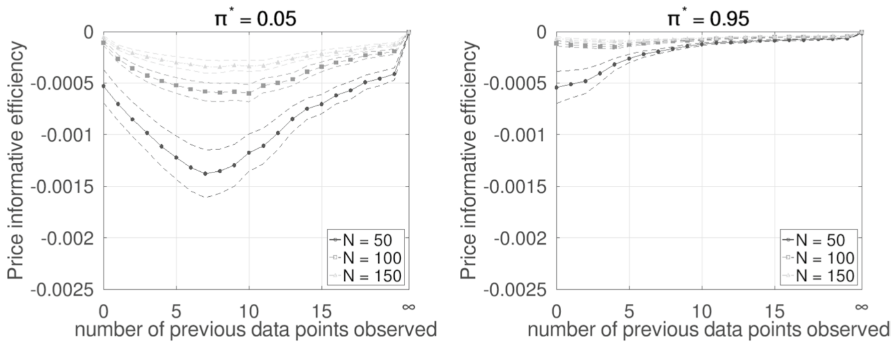

Fig. 4 Montecarlo estimates of the average price informative efficiency as a function of the number of previous data points observed n for $N = 50, 100, 150$. Left: $\pi^* = 0.05$. Right: $\pi^* = 0.95$ Dashed lines represent 95% confidence bounds. The symbol ∞ on the horizontal axis indicates the Bayesian case

estimates with 95% confidence bounds. The estimated average price informative efficiency in the Bayesian case is added as the last point and indicated by ∞ on the horizontal axis. An asymmetrical U-shaped pattern clearly emerges. The highest average price informative efficiency levels are showed in the case of zero-intelligence and full rationality, with the estimates in the two cases resulting not significantly different. Considering $n = 1$, one notices a drop in average efficiency and the decreasing pattern continues until it reaches its lowest point around $n = 10$. Afterward average efficiency starts to climb up. As expected, we notice an improvement in performance as the number of agents N increases. Comparing Fig. 2 with Fig. 1, one notices that the average efficiency levels reached here are much higher than those recorded for $\pi *= 0.5$ in the single-agent economy for any rationality level considered. That is a direct consequence of the selection process operated by the market.

To check for the robustness of the findings, we repeat the exercise considering $\pi = 0.2$ and $\pi^* = 0.8$. The results are shown in Fig. 3. We observe the same U-shaped pattern and the tendency of average efficiency to grow with N, but some differences need to be highlighted. For $\pi^* = 0.8$, the rate of convergence appears higher than the case $\pi^* = 0.5$. Moreover, the degree of efficiency seems to be higher than before for the vast majority of n considered and the minimum level appears to be reached around $n = 5$. In the case $\pi^* = 0.2$, we increase the number of rationality levels to 40. Indeed, the U-shaped pattern needs more data points to emerge, with the minimum level of average efficiency reached between $n = 10$ and $n = 20$, Moreover, the convergence toward full rationality levels appears slower and the number of agents matters more than the previous cases. Comparing Fig. 3 with Fig. 1, we observe the effect of market selection as mentioned above. Indeed, with $N = 1$ and $\pi^* = 0.8$, we have that full rationality outperforms zero-intelligence and Markov beliefs. With $N = 1$ and $\pi^* = 0.2$, we observe that zero-intelligence and Markov beliefs present the best performances, while full rationality falls behind. Here, instead, full rationality and zero-intelligence deliver high and almost identical efficiency performances. The partial use of past information, instead, produces a marked drop in efficiency, especially when we pass from $n = 0$ to $n = 1$. This is in line with the intuition suggested in advance about the contribution of interaction and selection in widening the region where the U-shaped relationship is observed.

We further explore the relationship between n and (average) $\mathcal{E}^n$ investigating its behavior for values of π^* close to the boundaries. As previously mentioned, for $\pi^* \to 0$ or $\pi^* \to 1$, one of the two models becomes (almost) correct, thus, misspecification basically disappears, Bayesian learning is favored and zero-intelligence is at a disadvantage. As a consequence, in those cases, the U-shaped relationship may disappear. In Fig. 4, we repeat the numerical exercise considering $\pi^* = 0.05$ and $\pi^* = 0.95$. As one can notice, while the U-shaped relation still holds for $\pi^* = 0.05$, it breaks down for $\pi^* = 0.95$. This is particularly evident in the case $N = 50$. Indeed, as N grows, the estimates flatten out, showing a slightly increasing behavior. Interestingly, a hint of U-shaped behavior reappears in the case $N = 100$; however, it does not seem very significant. This final exercise, while proving that the relation can fall short when we are close to a certain or impossible event, provides further support to our intuition on the effect of interaction and selection: the direct comparison between Figs. 1 and 4 in the case $\pi^* = 0.05$ confirms that market selection contributes to generate the U-shaped relationship.

6 Discussion of results

The economic mechanism that underlies our results can be understood considering the interplay of two evolutionary forces under the market interaction protocol. On the one hand, there is market selection: the wealth reallocation process emerging from the market interaction among agents. On the other hand, there is learning: the posterior probability reallocation process operated by each agent over the two models. When $N = 1$, the first force is absent, while, when $n = 0$ the second force is absent (agents' beliefs are constant).

In the single-agent economy, the U-shaped relationship is observed when π^* shows values around 0.5: averaging a pessimist and an optimist model or selecting the best model is, on average, more advantageous than changing mixture with respect to the last observation. Indeed, in the case, $n = 1$ individual probabilistic predictions follow a Markov process and this imposes a Markovian probabilistic structure in prices which is not present in the true data generating process. Such a spurious probabilistic structure is what generates the difference in the "likelihood to have an agent with (almost) correct beliefs" mentioned at the end of Sect. 4.

When the number of agents in the economy is larger than one, the two forces interact with each other. A first notable outcome is the increase of average efficiency for any rationality level considered with respect to the single-agent economy. Indeed, the expansion in heterogeneity together with the selection process operated by markets results beneficial. This is particularly evident in the zero-intelligence case, where the action of the market interaction (that is, the *structure*) is quite effective in selecting the best combination of models and, given the high number of different combinations, efficiency results high. As agents start to use information to update beliefs, efficiency plummets. As noted above, incorporating the last (few) data points introduces some spurious structure in agents' beliefs which is only weakly related to the true characteristics of the data generating process. Selection interferes with that, since the inferior performance of using few observations to learn with respect to zero-intelligence appears for values of π^* that do not generate the same relative outcome in the single-agent case. That is, the positive effect generated by market selection in the zero-intelligence case may not be equally large when agents use little information. A reason for that is the fact that, while in the zero-intelligence case, market selection operates over models belonging to the same family with respect to the truth, when n is low selection operates over models belonging to a different family with respect to the true one. For instance, in the case, $n = 1$ it operates over Markovian models, while the true process is i.i.d.. As agents start using a sufficiently high amount of past observations to learn, they approach Bayesian learning: They give (almost) full posterior probability to the most accurate model between the optimist one and the pessimist one. Hence, in such a case, the two forces are aligned: Agents select the best model they have, while the market selects the best agent. Hence, market selection ends up operating again over models belonging to the same family of the true one. Prices converge to the most accurate model in the set of all the models used by agents, and this is reflected by the high level of informative efficiency achieved in such a case.

As the number of agents becomes very large, individual learning does not matter anymore. Indeed, there exist an agent with (almost) correct beliefs for any value of n that accrues all the wealth and lets the market be efficient.

7 Conclusions

In this paper, we study how the price informative efficiency relates to individual rationality in learning in a simple prediction market model populated by heterogeneous Kelly traders and characterized by model misspecification. The prediction market framework allows for a straightforward interpretation of prices as probability. Thus, one can measure price informative efficiency in terms of the opposite of the average relative entropy of prices.

At the individual level, agents build their conditional probabilities combining heterogeneous optimist and pessimist models by means of weights that are updated using the last observations. If agents discard all the available information and never update their initial guesses, we obtain zero-intelligence agents. At the opposite extreme, agents use all the available information and follow Bayesian learning, the cornerstone of rational learning. Gradually increasing the number of past realizations the agents use, we can vary agents' rationality level and observe the consequences on price informative efficiency.

Using a mix of mathematical and computational techniques, our analysis reveals that a robust U-shaped relationship emerges between the two quantities. Moreover, the zero-intelligence economy reaches average efficiency levels that match those reached in the fully rational Bayesian case. Contrasting our results with two extreme cases, we argue that our results are driven by the interplay of two evolutionary forces: the market selection mechanism that moves wealth toward more accurate agents and the individual learning process that moves posterior probabilities over models depending on observed realization.

Our framework and our analysis can be extended along several dimensions. For instance, an interesting possible research avenue for expanding our analysis might be considering weighted averages in the learning process. Memory process approaches, long-term or short-term oriented, might produce unexpected positive effects in terms of market efficiency performances.[5]

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest None.

References

Antico A, Bottazzi G, Giachini D (2024) Pricing anomalies in a general equilibrium model with biased learning. Available at SSRN 4809542

Arrow KJ, Forsythe R, Gorham M, Hahn R, Hanson R, Ledyard JO, Levmore S, Litan R, Milgrom P, Nelson FD et al (2008) The promise of prediction markets. Science 320(5878):877

Barberis N, Shleifer A, Vishny R (1998) A model of investor sentiment. J Financ Econ 49(3):307–343

Barberis N, Thaler R (2003) A survey of behavioral finance. In: Constantinides G, Harris M, Stultz R (eds) Handbook of the economics of finance. Elsevier, Amsterdam

Bargigli L (2021) A model of market making with heterogeneous speculators. J Econ Interact Coord 16:1–28

Berk RH (1966) Limiting behavior of posterior distributions when the model is incorrect. Annal Math Stat 37(1):51–58

Beygelzimer A, Langford J, Pennock D M (2012) Learning performance of prediction markets with kelly bettors. In: Proceedings of the 11th international conference on autonomous agents and multiagent systems. volume 3, pp 1317–1318

Blume L, Easley D (1993) Economic natural selection. Econ Lett 42:281–289

Blume L, Easley D (2006) If you are so smart why aren't you rich? Belief selection in complete and incomplete markets. Econometrica 74:929–966

Blume L, Easley D (2009) The market organism: long-run survival in markets with heterogenous traders. J Econ Dyn Control 33:1023–1035

Bottazzi G, Dindo P (2013) Selection in asset markets: the good, the bad, and the unknown. J Evol Econ 23:641–661

Bottazzi G, Dindo P, Giachini D (2018) Long-run heterogeneity in an exchange economy with fixed-mix traders. Econ Theo 66(2):407–447

Bottazzi G, Dindo P, Giachini D (2019) Momentum and reversal in financial markets with persistent heterogeneity. Annal Finan 15(4):455–487

[5] We thank an anonymous reviewer for this interesting suggestion.

Bottazzi G, Giachini D (2017) Wealth and price distribution by diffusive pproximation in a repeated prediction market. Physica A Stat Mech Appl 471:473–479

Bottazzi G, Giachini D (2019) Betting, selection, and luck: a long-run analysis of repeated betting markets. Entropy 21(6):585

Bottazzi G, Giachini D (2019) Far from the madding crowd: collective wisdom in prediction markets. Quant Financ 19(9):1461–1471

Bottazzi G, Giachini D (2022) A general equilibrium model of investor sentiment. Econ Lett 218:110749

Bottazzi G, Giachini D, Ottaviani M (2023) Market selection and learning under model misspecification. J Econ Dyn Control 156:104739

Brock W, Hommes C (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. J Econ Dyn Control 22:1235–1274

Daniel K, Hirshleifer D, Subrahmanyam A (1998) Investor psychology and security market under-and overreactions. J Financ 53(6):1839–1885

Dindo P, Massari F (2020) The wisdom of the crowd in dynamic economies. Theor Econ 15(4):1627–1668

Epstein LG, Noor J, Sandroni A (2010) Non-bayesian learning. BE J Theor Econ 10(1):0000102202193517041623

Fama E (1970) Efficient capital markets: a review of theory and empirical work. J Financ 25:383–422

Fehr E, Tyran J-R (2005) Individual irrationality and aggregate outcomes. J Econ Perspect 19(4):43–66

Gardini L, Radi D, Schmitt N, Sushko I, Westerhoff F (2022) Causes of fragile stock market stability. J Econ Behav Org 200:483–498

Giachini D (2021) Rationality and asset prices under belief heterogeneity. J Evol Econ 31(1):207–233

Gigerenzer G, Brighton H (2009) Homo heuristicus: why biased minds make better inferences. Top Cogn Sci 1(1):107–143

Gigerenzer G, Gaissmaier W (2011) Heuristic decision making. Ann Rev Psychol 62:451–482

Gode DK, Sunder S (1993) Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. J Polit Econ 101(1):119–137

Hirshleifer D (2015) Behavioral finance. Ann Rev Financ Econ 7:133–159

Hommes C (2006) Heterogeneous agent models in economics and finance. In: Judd K, Tesfatsion L (eds) Handbook of Computational Economics Agent-Based Computational Economics. Elsevier, Amsterdam

Kets W, Pennock D M, Sethi R, Shah N (2014) Betting strategies, market selection, and the wisdom of crowds. In: Twenty-eighth AAAI conference on artificial intelligence

Kirman A (2010) Complex economics: individual and collective rationality. Routledge, London

LeRoy SF (1989) Efficient capital markets and martingales. J Econ Lit 27(4):1583–1621

Marinacci M, Massari F (2019) Learning from ambiguous and misspecified models. J Math Econ 84:144–149

Massari F (2017) Markets with heterogeneous beliefs: a necessary and sufficient condition for a trader to vanish. J Econ Dyn Control 78:190–205

Massari F (2020) Under-reaction: Irrational behavior or robust response to model misspecification? Available at SSRN 3636136

Mousavi S, Gigerenzer G, Kheirandish R (2015) Rethinking behavioral economics through fast-and-frugal heuristics. Routledge, New York

Naimzada AK, Ricchiuti G (2009) Dynamic effects of increasing heterogeneity in financial markets. Chaos, Solitons Fractals 41(4):1764–1772

Naimzada AK, Ricchiuti G (2014) Complexity with heterogeneous fundamentalists and a multiplicative price mechanism. Econ Notes Rev Bank Financ Monet Econ 43(3):233–247

Sandroni A (2000) Do markets favor agents able to make accurate predictions. Econometrica 68(6):1303–1341