

Selection in asset markets: the good, the bad, and the unknown

Giulio Bottazzi · Pietro Dindo

Published online: 19 May 2013
© Springer-Verlag Berlin Heidelberg 2013

Abstract In this paper, we use a series of simple examples to illustrate how wealth-driven selection works in a market for Arrow securities. Our analysis delivers both a good and a bad message. The good message is that, when traders invest constant fractions of their wealth in each asset and have equal consumption rates, markets are informationally efficient: the best informed agent is rewarded and asset prices eventually reflect this information. However, and this is the bad message, when asset demands are not constant fractions of wealth but dependent upon prices, markets might behave sub-optimally. In this case, asymptotic prices depend on preferences and beliefs of the whole ecology of traders and do not, in general, reflect the best available information. We show that the key difference between the two cases lies in the local, i.e. price dependent, versus global nature of wealth-driven selection.

Keywords Market selection · Evolutionary Finance · Informational efficiency · Asset pricing · CRRA preferences

JEL Classifications D50 · D80 · G11 · G12

1 Introduction

The aim of this paper is to illustrate, using simple examples, the good, the bad, and the unknown about wealth-driven selection in asset markets. For this purpose, we consider the simplest setting in which the problem can be investigated, a repeated

G. Bottazzi
Istituto di Economia, Scuola Superiore Sant'Anna, Piazza Martiri della Libertà 33, 56127 Pisa, Italy

P. Dindo (✉)
Dipartimento di Economia e Management, Università di Pisa, Via C.Ridolfi 10, 56124 Pisa, Italy
e-mail: p.dindo@ec.unipi.it

market for two Arrow securities. Each asset represents a different way to transfer current wealth to future consumption or re-investment. Whereas aggregate wealth is fixed ex-ante and depends only on the realized state of Nature, prevailing asset prices and the distribution of wealth among agents are endogenously determined and depend also on agents' investment decisions. Even in our stylized framework, whether the market is capable of rewarding the best informed agents or, more in general, the nature of the long-run prices and wealth distributions is still an open issue.

Despite this lack of knowledge, most models in economics and finance postulate the ability of markets to select for agents whose equilibrium prices reflect the correct, or the best available, information about asset fundamentals. They assume so in much more complex settings than the one analyzed here. The underlying idea is that agents who have inferior information would lose wealth over time in favor of the better informed. If this were true, the market would converge to a long-run equilibrium in which the best informed agent has all wealth and prices would reflect this information. The market could be said to be informationally efficient, as it efficiently conveys, through prices, the best available information about asset fundamentals.

The examples of this paper are meant to illustrate when the above idea is correct and when, and why, it is not. We show that the class of investment behaviors considered crucially influences the efficiency of the wealth-driven selection process. To illustrate the point, we consider investment behaviors or "portfolio rules" derived from the one period¹ maximization of an expected log or power utility, that is, Constant Relative Risk Aversion (CRRA) preferences, given relative risk aversion γ and subjective probabilities about the occurrence of states of Nature. When $\gamma = 1$, that is with log utility, the demand is expressed as a constant fraction of wealth to be invested in each asset. When $\gamma \neq 1$, that is with power utility, the fraction of wealth becomes price dependent. We analyze the asymptotic states of our market dynamics and compare results for the two cases. When the fraction of wealth invested in each asset is constant, the agent whose subjective probabilities are the "closest" to the true probability distribution of the states of Nature dominates the economy and takes all the wealth in the long run.² Since the agent possessing the best information about the fundamental process is rewarded by the market, prices are set to reveal this information. This is the "good" side of selection in asset markets.

On the other hand, when portfolio rules are price dependent, long-run prices do not need to represent the best available information. In general, it is not possible to judge ex-ante who will be rewarded by the market and to what extent the latter will be able to reveal the correct information about asset fundamentals. In fact, it is easy to construct examples in which prices exhibit endogenous fluctuations and

¹Since we are primarily interested in analyzing the selective capability of markets, we assume that agents are myopic and drop the assumption of perfect foresight or rational-expectations. These assumptions would indeed sterilize the effect of the trading-induced wealth reallocation, as all the problems would be reduced to an ex-ante identification of possible equilibria.

²As we shall see, the appropriate quantification of the "distance" between probability distributions is provided by the relative entropy.

never converge or, depending on initial conditions and realizations of states of Nature, reveal the worst available information. This is the “bad” side of selection in asset markets. Interestingly, the failure to select the best informed agents can be related to analogous issues pointed out in Nelson and Winter (2002) within the firm selection literature, namely: variety, behavioral continuity, limited path dependence, and profit induced growth. We refer to the same classification at the end of the paper when we summarize our findings.

Technically, the dynamics of agents wealth and asset prices can be formalized as a discrete-time random dynamical system. Given agents wealth fractions, portfolio rules, and asset prices at time t , the market dynamics is a random map which depends on the next realized state of Nature. Different states of Nature give rise to different reallocation of resources at time $t + 1$. The newly determined wealth distribution, together with the updated portfolio rules, possibly changed by the information revealed by previous trading, will, in turn, determine the next prices, and the process is repeated over and over again. The analysis of long run wealth and price distributions becomes the ultimate object of interest. In Bottazzi and Dindo (2010) we focus on the analytical investigation of the random dynamical system representing the market dynamics and are able to derive local stability conditions of those states in which only one trader has all the wealth. In the present contribution, relying on those results, we discuss selection in asset markets using simple examples and with the help of a graphical tool, thus avoiding technical proofs. First, by plotting agents portfolio rules against a normalization of the supply, the so-called Equilibrium Market Curve (EMC, see Section 4), we find the possible long-run selection equilibria as the short-run Walrasian equilibria where only one agent has positive wealth. Second, comparing the distance of each individual rule to a benchmark rule, the Kelly rule, so-called after Kelly (1956), we discuss the local stability of single survivor equilibria. This characterization of the long-run dynamics enables us to assess the selective capability of the market and its informational efficiency.

The outline of this paper is as follows. In the next section, we provide a brief introduction to the recent works that have investigated wealth-driven selection in asset markets. In Section 3, we introduce the asset market model and derive agents' demand from the (myopic) maximization of a CRRA utility function. In Section 4, we consider the case of constant portfolio rules, characterize the outcome of the long-run market dynamics and show how, in this case, the market rewards the best informed agent. The analysis begs the question of what happens with more general demand functions. We answer this question in Section 5, where we introduce price dependence in agents investment decisions. As the examples in Sections 5.1 and 5.2 show, the market selective capability in favor of the best informed agent turns from working globally to working only locally so that the convergence of prices toward informationally efficient values is not granted anymore. In Section 6, we show that the local nature of market selection leads to the impossibility of ordering rules according to their mutual survivability. In Section 7, we conclude by linking our findings with those summarized by Nelson and Winter (2002) for the firm selection literature, by presenting the open issues, and by setting a possible agenda for future research.

2 Selection in asset markets

The investigation of the ability of markets to redistribute resources in favor of the rational and better informed agents started with Blume and Easley (1992) (see also Blume and Easley 2010). They study the relative wealth dynamics among different investors repeatedly exchanging Arrow securities in a temporary equilibrium framework. Trading occurs because agents have different beliefs on the likelihood that a security will pay its dividend and/or different risk preferences. They find that market selection works in that there exists an investment rule that dominates against any other rule and drives asset prices to their correct values. We draw upon this result in Section 4 where we present the “good” side of selection in asset markets. More in general, we show that, within portfolio rules that prescribe investing a fixed fraction of wealth in each asset, informational efficiency holds and asset prices converge as close as possible to their fundamentals. However, Blume and Easley (1992) also construct examples where, in absence of the global dominating rule, market selection does not work and leads to asset miss-pricing. We clarify the origin of this form of inefficiency in Section 5 where we present the “bad” side of selection in asset markets.

The following literature has tried to generalize the hypothesis under which market selection does work. Evolutionary Finance (see Evstigneev et al. 2009 for a recent survey) has focused on more general market structures, while keeping agents behaviors in the class of constant rules we analyse in Section 4. Incomplete markets and generic Markovian processes driving the states of Nature have been studied: Amir et al. (2005) and Evstigneev et al. (2008) extended the analysis for short and long-living asset, respectively. Although partly different results have been found with respect to the global dominating strategy, differences which led to the introduction of the so-called generalized Kelly rule (see Sections 4 and 5), in all these works it is confirmed that, as long as the global dominating rule is trading, market informational efficiency holds.

A different group of works has focused on generalizing the class of rules to encompass any investment decision explicitly coming from utility maximization, so that asset demand is not necessarily expressed as a fixed, belief dependent, fraction of wealth. Assuming perfect foresight on realized prices and market completeness Sandroni (2000) and Blume and Easley (2006) find that, no matter the functional form of the utility function they maximize, the agent whose beliefs are “nearest” the correct ones is selected in the long run. Sandroni (2005) generalizes this result to specific types of incomplete markets. Although these results have strong implications, they are based on the very demanding assumption that agents are able to coordinate on having perfect foresight on realized prices. This assumption does not only imply that agents are able to coordinate their expectations, but also that they can guess *ex-ante* the correct price. In all these works, market informational efficiency is a by-product of allocative efficiency which, in turns, is granted by the perfect foresight assumption.

In the context of financial market models with heterogeneous agents, the general finding is that the interaction of boundedly rational agents may lead, through the linking of past market performances to present returns, to endogenous price fluctuations

and asset mis-pricing (see Hommes 2006; LeBaron 2006, for a review). In particular, in our previous work (see e.g. Anufriev et al. 2006; Anufriev and Bottazzi 2010; Anufriev and Dindo 2010) we have studied wealth-driven market selection on the class of price dependent investment rules in a standard model with a risky and a risk-free asset. Those works are, however, based on an essentially deterministic framework and do not discuss the informational efficiency issue we are interested in here.

The effort to investigate wealth-driven selection in an stochastic framework with non constant portfolio rules started with our contribution, (Bottazzi and Dindo 2010). There we take a standard Evolutionary Finance market model with short-lived assets and add price dependent portfolio rules to the analysis. We find that the comparison of the relative performance of portfolio rules is only possible at given prices, explaining why there exist cases in which no rule gains all wealth in the long run or in which different rules gain all wealth for different initial conditions or different realizations of the fundamental process. We draw upon this contribution in Sections 5 and 6 where we present the “bad” consequences of selection in asset markets.

3 The model

Assume that there are 2 possible states of Nature, state 1 and state 2. Their unfolding in time can be represented by a Bernoulli process $\omega = (\omega_1, \dots, \omega_t, \dots)$, with $\omega_t \in \{1, 2\}$ and $\omega_t = 1$ with probability π for every $t \in \mathbb{N}$. We denote with \mathbb{N}^+ the set $\mathbb{N} \cup \{0\}$. In order to transfer wealth inter-temporally, and thus being able to consume in the future, I agents can trade in two short-lived Arrow securities. Asset $k = 1, 2$, exchanged in period t at price $p_{k,t}$, pays one unit of the consumption good in period $t + 1$ if $\omega_{t+1} = k$, and zero otherwise. The consumption good is the numéraire of the economy. Each agents' demand is expressed as a fraction of wealth, that is, the demand of agent i for asset k at time t is given by $\alpha_{k,t}^i w_t^i / p_{k,t}$, where $\alpha_{k,t}^i$ is agent i invested fraction in asset k in period t , and w_t^i is agent i wealth in period t . We refer to the vector $(\alpha_{1,t}^i, \alpha_{2,t}^i)$ as the portfolio, or investment³, rule of agent i in period t .

Asset markets open at time 0 and close at an arbitrary large final period T . We assume that consumption occurs only in this final period. All intermediate wealth is thus saved and re-invested in assets⁴. Agent i budget constraint implies that her portfolio rule obeys to $\alpha_{1,t}^i + \alpha_{2,t}^i = 1$ for every $t \in \mathbb{N}^+$. Moreover we impose that $\alpha_{k,t}^i \geq 0$, $k = 1, 2$, for every $t \in \mathbb{N}^+$, as the agent who violates this constraint will

³Portfolio rules are distinct from investment rules when intermediate consumption is considered, in which case the portfolio rule specifies the fraction of wealth to be allocated to each asset whereas the investment rule also specifies the fraction of wealth to be saved and the fraction of wealth to be consumed. See also footnote 4 below.

⁴This is equivalent to assuming a constant and homogeneous consumption rate for all agents, bar a renormalization of price levels. As discussed in Blume and Easley (1992), the introduction of heterogeneous consumption rates weakens markets informational efficiency when better informed traders have a higher propensity to consume. Since the effect is well understood, we have chosen not to consider it in the present analysis.

have negative wealth in finite time. Both constraints imply that each agent i period t investment rule belongs to Δ_2 , the unitary simplex of \mathbb{R}^2 .

The price of asset $k = 1, 2$ in period t , $p_{k,t}$, is fixed by Walrasian market clearing. Assuming unitary assets supply, period t market clearing can be written both in terms of units

$$1 = \sum_{i=1}^I \frac{\alpha_{k,t}^i w_t^i}{p_{k,t}}, \quad k = 1, 2$$

or, more conveniently for our purposes, in terms of market values

$$p_{k,t} = \sum_{i=1}^I \alpha_{k,t}^i w_t^i, \quad k = 1, 2. \tag{1}$$

At time $t + 1$, if the event $\omega_{t+1} \in \{1, 2\}$ is realized, the wealth of agent i , w_{t+1}^i , is equal to the amount of shares of asset $k = \omega_{t+1}$ bought at time t , or

$$w_{t+1}^i = \frac{\alpha_{\omega_{t+1},t}^i w_t^i}{p_{\omega_{t+1},t}}. \tag{2}$$

By summing up the previous equation over all the I agents and using Eq. 1, it is clear that, no matter the initial wealth, the total wealth in the economy is then equal to one, the total assets dividend payment. Since portfolio rules are constrained to add up to one, this implies that

$$p_{1,t} + p_{2,t} = 1 \quad \text{for every } t \in \mathbb{N}^+. \tag{3}$$

The latter equation can be used to normalize prices so that we shall use $p_{1,t} = p_t$ and $p_{2,t} = 1 - p_t$ in what follows.

For definiteness and illustrative purposes, in this contribution we restrict our analysis to the class of investment rules derived from the maximization of the expected CRRA utility of next period wealth. However, as it will become clear later, the phenomena we discuss have a general character and do not depend in any respect on this assumption. An agent who assigns the subjective probability π^e to the realization of state 1 and who possesses a relative risk aversion parameter γ derives period t portfolio rule $(\alpha_{1,t}, \alpha_{2,t})$ by maximizing the expected utility of wealth in period $t + 1$, that is, using Eq. 2, by solving

$$\operatorname{argmax}_{(\alpha_{1,t}, \alpha_{2,t}) \in \Delta_2} \left\{ \pi^e u \left(\frac{\alpha_{1,t} w_t}{p_t} \right) + (1 - \pi^e) u \left(\frac{\alpha_{2,t} w_t}{1 - p_t} \right) \right\},$$

where u is a log or power utility function

$$u(x) = \begin{cases} \frac{x^{1-\gamma}}{1-\gamma} & \text{if } \gamma \in (0, 1) \cup (1, +\infty), \\ \log(x) & \text{if } \gamma = 1. \end{cases} \tag{4}$$

It is straightforward to get the explicit solution of the maximization problem: when $\gamma \neq 1$, the portfolio rule depends on p_t but not on w_t and reads

$$\alpha_{1,t} = \alpha(p_t; \pi^e, \gamma) = \frac{(\pi^e(1-p_t)^{1-\gamma})^{\frac{1}{\gamma}}}{(\pi^e(1-p_t)^{1-\gamma})^{\frac{1}{\gamma}} + ((1-\pi^e)(p_t)^{1-\gamma})^{\frac{1}{\gamma}}}, \tag{5}$$

$$\alpha_{2,t} = 1 - \alpha(p_t; \pi^e, \gamma).$$

When $\gamma = 1$, the previous solution collapses to constant portfolio rule

$$\alpha_{1,t} = \alpha(p_t; \pi^e, 1) = \pi^e, \text{ and } \alpha_{2,t} = 1 - \alpha(p_t; \pi^e, 1) = 1 - \pi^e. \tag{6}$$

We will denote with $\alpha^i(p; \pi^e, \gamma)$ the portfolio rule of agent i . As we shall see in the following sections, the analysis of the market dynamics, and especially its asymptotic characterization, critically depends on the nature of these rules.

4 The good: wealth-driven selection with constant portfolio rules

In this section, we consider the benchmark case of constant portfolio rules, that is, agents invest constant fractions of their wealth in each asset, irrespectively of their wealth or assets prices. As shown in the previous section, this rule corresponds to the maximization of the expected logarithmic utility of next period wealth.⁵ The fraction of wealth invested in each asset is equal to the subjective probability of realization of the associated state of Nature, see Eq. 6.

Consider first the case of a market with a single investor. Let w_t be the investor wealth and α the share of wealth invested in asset 1. We assume that investors assign a positive probability to both states, so that $\alpha \in (0, 1)$. The market dynamics in Eqs. 1 and 2 becomes

$$w_{t+1} = \begin{cases} \frac{\alpha w_t}{p_t} & \text{if } \omega_{t+1} = 1, \\ \frac{(1-\alpha)w_t}{1-p_t} & \text{if } \omega_{t+1} = 2, \end{cases} \tag{7}$$

where the price of the first asset is

$$p_t = \alpha w_t. \tag{8}$$

Substituting Eq. 8 in Eq. 7, it is immediately clear that, by investing a positive amount in both assets, the trader will secure all future wealth no matter the realization of the state of Nature. Wealth evolution does not pay any role and $p_t = \alpha$ for every t .

We can use a standard supply and demand plot to visualize the market equilibrium price. In the left panel of Fig. 1 we plot both the supply market value p , and the demand market value α . Their crossing, E_1 , fixes the equilibrium price, which in this case is trivially equal to α . Consistently with our previous works, we name the diagonal supply curve in this plot the Equilibrium Market Curve (EMC). Indeed,

⁵Were the reader, for any reason, averse to the expected utility framework, he or she is free to consider the constant rules as behavioral rules, that is, as mere descriptions of agents behavior. Our results and the overall analysis remain the same.




Fig. 1 *Left Panel:* EMC plot representing a market with a single investor using $\alpha = 0.4$. The only possible outcome of the dynamics is E_1 . *Right Panel:* EMC plot representing a two-agent market. Investment shares are $\alpha^1 = 0.4$ and $\alpha^2 = 0.7$, respectively. E_1 and E_2 are the two single survivor long-run equilibria

we shall show that the EMC is the locus of all possible long-run equilibria of our economy.

In order to study the role of the market in redistributing wealth among agents, and the ensuing asset price behavior, we need to add a second investor. Let the investment rules be α^1 and α^2 , respectively. The wealth of agent $i = 1, 2$ evolves according to

$$w_{t+1}^i = \begin{cases} \frac{\alpha^i w_t^i}{p_t} & \text{if } \omega_{t+1} = 1, \\ \frac{(1-\alpha^i)w_t^i}{1-p_t} & \text{if } \omega_{t+1} = 2, \end{cases} \tag{9}$$

where the price of the first asset is given by the combination of both rules

$$p_t = \alpha^1 w_t^1 + \alpha^2 (1 - w_t^1), \tag{10}$$

and we have used $w_t^2 = 1 - w_t^1$. If, in a given period, agent i possesses all the wealth, then $p = \alpha^i$ and agent i will possess all the wealth in all subsequent periods. If otherwise, both agents have positive wealth, prices are between α^1 and α^2 and returns depend on the realization of the state of Nature. Notice that, in each period, the market rewards the agent with the highest stake in the dividend paying asset.

The right panel of Fig. 1 illustrates a two-agent market using the ‘‘EMC plot’’. E_1 and E_2 are the crossing of agent rules with the EMC. They are deterministic fixed points of the random dynamical system specified by Eq. 9: no matter the realized states of Nature, the market dynamics starting in E_1 (or E_2) will stay there for ever. E_1 and E_2 are the single survivor equilibria of the market dynamics since they occur when one agent possesses the entire wealth. The question we want to answer is whether for generic initial conditions the market will converge toward a single survivor equilibrium or keep fluctuating between them. For this purpose, we use the evolution of the wealth ratio. From Eq. 9, knowing that $\omega_{t+1} = 1$ with probability π , one has

$$\frac{w_{t+1}^1}{w_{t+1}^2} = \begin{cases} \frac{\alpha^1 w_t^1}{\alpha^2 w_t^2} & \text{with probability } \pi, \\ \frac{1-\alpha^1}{1-\alpha^2} \frac{w_t^1}{w_t^2} & \text{with probability } 1 - \pi. \end{cases} \tag{11}$$

The wealth of the agent with the highest stake in asset 1, that is, the agent with the larger belief α (agent 2 in the example in the right panel of Fig. 1), increases with probability π and decreases with probability $1 - \pi$. In T periods, denoting with T_1 the number of times state 1 is realized, one has

$$\frac{w_T^1}{w_T^2} = \left(\frac{\alpha^1}{\alpha^2}\right)^{T_1} \left(\frac{1 - \alpha^1}{1 - \alpha^2}\right)^{T - T_1} \frac{w_0^1}{w_0^2},$$

and, taking the log,

$$\log \frac{w_T^1}{w_T^2} = T_1 \log \frac{\alpha^1}{\alpha^2} + (T - T_1) \log \frac{1 - \alpha^1}{1 - \alpha^2} + \log \frac{w_0^1}{w_0^2}.$$

Since, for the Law of Large Numbers, $T_1/T \rightarrow \pi$ as $T \rightarrow \infty$, one obtains

$$\lim_{T \rightarrow \infty} \frac{1}{T} \log \frac{w_T^1}{w_T^2} = \left(\pi \log \frac{\alpha^1}{\alpha^2} + (1 - \pi) \log \frac{1 - \alpha^1}{1 - \alpha^2}\right). \tag{12}$$

To interpret Eq. 12, notice that the investment rule of each agent can be seen as a probability measure $(\alpha, 1 - \alpha)$ defined over the states of Nature. Consider now the relative entropy of this measure with respect the true probabilities $(\pi, 1 - \pi)$ defined as

$$I_\pi(\alpha) = \pi \log \frac{\pi}{\alpha} + (1 - \pi) \log \frac{1 - \pi}{1 - \alpha}. \tag{13}$$

$I_\pi(\alpha)$ is always non-negative and is zero if and only if $\alpha = \pi$. The relative entropy is a measure of information loss: the lower its value, the greater the agreement of the beliefs $(\alpha, 1 - \alpha)$ with the true Bernoulli distribution $(\pi, 1 - \pi)$. Substituting the expression for the relative entropy in Eq. 12, one obtains

$$\lim_{T \rightarrow \infty} \frac{1}{T} \log \frac{w_T^1}{w_T^2} = I_\pi(\alpha^2) - I_\pi(\alpha^1).$$

If $I_\pi(\alpha^2) > I_\pi(\alpha^1)$ the limit is positive, which implies that $w_T^2 \rightarrow 0$ as $T \rightarrow \infty$. Agent 1 dominates, that is, gains all the wealth in the long-run. If instead $I_\pi(\alpha^2) < I_\pi(\alpha^1)$, then $w_T^1 \rightarrow 0$ as $T \rightarrow \infty$. In this case, agent 2 dominates. We have established that the agent whose investment rule or beliefs have the lowest relative entropy with respect to the process generating the sequence of states of Nature dominates. At the same time, long-run prices will reflect the best beliefs and will move as close as possible to the probability distribution $(\pi, 1 - \pi)$. In this sense, the market is informationally efficient.

It is useful to restate this result using the EMC plot. For this purpose, we add the line π to the two-agent EMC plot in the right panel of Fig. 1 to obtain Fig. 2.




Fig. 2 EMC plot representing a two agents market with $\pi = 1/2$, $\alpha^1 = 0.4$, and $\alpha^2 = 0.7$. The two possible long-run outcomes of the dynamics are S_1 , where agent 1 dominates, and U_2 , where agent 2 dominates. Since the line α^1 is closer to π than α^2 , S_1 is globally stable and U_2 unstable

Notice that the relative entropy of a strategy α is monotonically related to its distance from the line π . Indeed, the expression in Eq. 13 is monotonically decreasing in α if $\alpha < \pi$ and monotonically increasing if $\alpha > \pi$. Moreover, when $\pi = 1/2$, as in all the examples of this paper, the function $I_\pi(\alpha)$ is symmetric around π , so that the relative entropy can be directly derived from the Euclidean distance in the EMC plot.⁶ Thus, since in Fig. 2 the π line is closer to α^1 than to α^2 , if both agents start with positive initial wealth, no matter how much in favor of agent 2, agent 1 will gain all wealth in the long run and the price of the first asset converges to E_1 . Technically, we have just established that the only (globally) stable equilibrium is E_1 .

From the previous discussion it is clear that the best possible constant portfolio rule, the rule that gains all wealth when trading with (different) constant rules, is the rule with minimal entropic distance from the process governing the succession of the states of Nature. The rule prescribes investing a fraction of wealth π in the first asset and a fraction $1 - \pi$ in the second asset. This rule, to invest proportionally to the probabilities, is named the Kelly rule after Kelly (1956). Although it has its origin in models with exogenous prices (exogenous odds in the original betting framework, see e.g. Breiman 1961), it provides the average highest growth rate of wealth also in our model. According to our analysis, when present in the market, the Kelly rule beats any other different constant rule and bring prices to the probability revealing values π and $1 - \pi$. In terms of the maximization problem, the Kelly rule corresponds to an investor possessing correct beliefs about the process governing the state of

⁶When $\pi \neq 1/2$, one can use the function $I_\pi(\alpha)$ to rescale the vertical axis in the EMC plot so that the entropic distance can be still inferred by visual inspection.

Nature and having CRRA preferences with unitary relative risk aversion coefficient, i.e. logarithmic utility, a principle known as log-optimality.⁷

In the EMC plot, the Kelly rule coincides with the probability line π , so that comparing distances with respect to the probability line amounts to comparing distances with respect to the Kelly rule. Thus, we can restate the global stability results by saying that the rule closest to the Kelly rule is the one that takes all the wealth in the long-run.

Similar conclusions can be reached in the general case of many assets and many agents, as originally shown in Blume and Easley (1992) (see also Evstigneev et al. 2009, for a survey of other possible extensions). In all cases, the agent gaining all wealth in the long-run is the one using the rule with the lowest relative entropy with respect to the invariant distribution of the dividend generating process. By gaining all wealth, this agent drives prices toward the single survivor fixed point “closest” to the probability measure ruling the states of Nature process. In this case, the market is informationally efficient and prices assets as close as possible to their probability revealing values. This is not, however, the end of the story.

5 The bad: wealth-driven selection with price dependent portfolio rules

In this section, we investigate market behavior under wealth-driven selection when agents have values of the relative risk aversion coefficient γ different from one. As is clear from Eq. 5, this implies that portfolio rules depend on prices. In Fig. 3, different rules are displayed, both for values of γ close to one (left panel) and far from it (right panel). The shapes of CRRA portfolio rules are consistent with the notion of risk aversion. When γ is close to zero, agents are close to being risk-neutral, and choose to invest all their wealth in the asset with the highest (subjective) expected return, so that extremal values of α become more likely. On the contrary, more risk adverse agents tend to split their wealth proportionally to asset prices in order to achieve a less risky, but smaller, return. In this case the value of α remains closer to the EMC for a wide price range.

When rules depend on prices, agents’ wealth evolve according to

$$w_{t+1}^i = \begin{cases} \frac{\alpha^i(p_t)w_t^i}{p_t} & \omega_{t+1} = 1 \\ \frac{(1-\alpha^i(p_t))w_t^i}{1-p_t} & \omega_{t+1} = 2 \end{cases}, \quad i = 1, 2, \tag{14}$$

⁷In some contributions of the Evolutionary Finance literature, the Kelly rule has been replaced by the generalized Kelly rule, that is, investing proportionally to expected assets payoffs. Despite the fact that the two rules coincide in a market with only Arrow securities, they differ under more general asset structures. In particular, the generalized Kelly rule is in general not log-optimal, although it has been shown to dominate among constant rules, see Evstigneev et al. (2009). The relation between the log-optimal Kelly rule and the generalized Kelly rule is the object of ongoing research. See also Section 5.3 of this paper and Section 5.1 in Bottazzi and Dindo (2010).




Fig. 3 Portfolio rules derived by the maximization of a CRRA utility defined in Eq. 5 and 6). *Left panel:* values of γ close to 1. *Right panel:* extremal values of γ

where p_t is a solution of the (now implicit) equation

$$p_t = \alpha^1(p_t)w_t^1 + \alpha^2(p_t)(1 - w_t^1). \tag{15}$$

Since $\alpha^i(p) \geq p$ for $p \leq \pi^{e,i}$, as can be shown directly from Eq. 5, the right hand side of Eq. 15 is greater than p_t when $p_t < \min\{\pi_1, \pi_2\}$, while lower than p_t when $p_t > \max\{\pi_1, \pi_2\}$. Being a continuous function, it has at least an intersection with the supply curve p_t in the interval $[\min\{\pi_1, \pi_2\}, \max\{\pi_1, \pi_2\}]$. Regarding uniqueness, given the analysis of market dynamics we shall employ in what follows, we need only to ensure that the equilibrium price where agent 1 or agent 2 has all wealth, π^1 and π^2 , respectively, are regular, that is, locally unique. The result follows easily by imposing either $w_t^1 = 1$ or $w_t^1 = 0$ in Eq. 15 and using again $\alpha^i(p) \geq p$ for $p \leq \pi^{e,i}$. Once an equilibrium price is fixed, the evolution of wealth ratios still reads

$$\frac{w_{t+1}^1}{w_{t+1}^2} = \begin{cases} \frac{\alpha^1(p_t) w_t^1}{\alpha^2(p_t) w_t^2} & \text{with probability } \pi, \\ \frac{1-\alpha^1(p_t) w_t^1}{1-\alpha^2(p_t) w_t^2} & \text{with probability } 1 - \pi, \end{cases}, \tag{16}$$

but now, due the price dependence of the α s, the value of the ratio w_T^1/w_T^2 after T time steps depends on the price history. At each time step we can still compute the expected log-growth rate of the wealth ratio which can be written as the difference of the relative entropies

$$\begin{aligned} E \left[\log \frac{w_{t+1}^1}{w_{t+1}^2} - \log \frac{w_t^1}{w_t^2} \right] &= \pi \log \frac{\alpha^1(p_t)}{\alpha^2(p_t)} + (1 - \pi) \log \frac{1 - \alpha^1(p_t)}{1 - \alpha^2(p_t)} \\ &= I_\pi(\alpha^2(p_t)) - I_\pi(\alpha^1(p_t)). \end{aligned}$$

The difference depends on prevailing prices, and so does the relative distance from the Kelly rule, as can be seen in the EMC plot. As a result, there may exist prices at which rule 1 is nearer to the Kelly rule, thus having the lowest relative entropy and gaining, on average, wealth shares, and prices at which rule 2 is nearer. Market selection through entropy minimization operates only locally, that is, for given prices, rather than globally.

Since the relative performance of rules is price dependent, a global stability analysis similar to the one done in the previous section is not straightforward. Thus, we

follow a different approach: we analyze the occurrence of the asymptotic states in which a single agent dominates. First, we identify the single survivor equilibria at which all the wealth is in the hands of one agent and prices are accordingly set by the agent rule. Second, we check whether market dynamics is stable for prices close to these equilibria.⁸ For example, in the two agents market dynamics illustrated in the left panel of Fig. 4, local stability analysis amounts to studying the dynamics in the neighborhoods of S_1 , where $w^1 = 1$ and the price of the first asset is $p = \alpha^1$, and of S_2 , where $w^2 = 1$ and prices are set by the solution of $p = \alpha^2(p)$. The local analysis provides sufficient insight to discuss the asymptotic behavior of prices and wealth.

In the following sections, we use portfolio rules of the types given in Eqs. 5 and 6 to illustrate the implication of the price dependency for market informational efficiency. Notice that, no matter the value of γ , all the rules with the same π^e cross the EMC at the same point, that is, $p = \pi^e$ (c.f. rules in Fig 3). As a result, irrespectively of their preferences, if agents share the same beliefs about the occurrence of states of Nature, then assets are priced at that level and all agents have the same unitary return, and constant wealth. Prices represent the homogeneous beliefs, and preferences do not matter. As we shall see, the dynamics is more interesting when agents disagree on the value of π .

5.1 Coexistence of stable long-run market equilibria

For the first example depicted in Fig. 4, left panel, we take $\pi = 1/2, \pi^{e,1} = 0.25, \gamma^1 = 1, \pi^{e,2} = 0.65$, and $\gamma^2 = 0.5$. There exist two single survivor equilibria: S_1 , where only agent 1 survives and the price of the first asset is set at $p = \pi^{e,1}$, and S_2 , where only agent 2 survives and the price of the first asset is set to $p = \pi^{e,2}$.

Computing the distances of portfolio rule α^1 and α^2 from π at the price corresponding to S_1 one finds that $|\alpha^2(\pi^{e,1}) - 1/2| > |\pi^{e,1} - 1/2| = 1/4$. Then for initial prices near to $\pi^{e,1}$ the rule used by agent 1 has a lower relative entropy compared to the rule used by agent 2, i.e. $I_\pi(\alpha^2(p_t)) > I_\pi(\alpha^1(p_t))$. The wealth share of agent 1 will, on average, increase and will eventually converge to one. As a result, if the market share of the first agent is almost 1, and we imagine forcefully moving a small fraction of wealth from agent 1 to agent 2, the dynamics of the market will bring the system back to the situation in which $w^1 = 1$. We have established that S_1 is locally stable.

A symmetric argument holds in S_2 , where $|\pi^{e,2} - 1/2| = 3/20 < |\pi^{e,1} - 1/2|$, so that agent 2 is closer to the probability line π than agent 1. Then for initial prices near to $\pi^{e,2}$, it is the wealth share of agent 2 that increases, on average, so that rule

⁸In other words, we shall first identify all the deterministic fixed points of the random dynamical system, and then perform their local asymptotic stability analysis. The local analysis is made possible by the fact that in single agent markets the equilibrium price is regular, i.e. locally unique, so that the market dynamics is well defined. Thus market dynamics can be linearized around the fixed points and agents can be described as if they were using suitably defined constant rules. As a result, the analysis of asymptotic states can proceed along the lines of Section 4. For sufficiently smooth investment rules, results from the linearized markets carry over to the original market, albeit only locally (see Bottazzi and Dindo 2010, for details).




Fig. 4 EMC plots representing two-agent markets with price dependent rules. *Left panel:* $\pi = 1/2$, $\alpha^1(p) = \alpha(p; 0.25, 1)$, $\alpha^2(p) = \alpha(p; 0.65, 0.5)$. The two possible long-run outcomes of the dynamics are S_1 , where agent 1 dominates, and S_2 , where agent 2 dominates. Both are locally stable. *Right panel:* $\pi = 1/2$, $\alpha^1(p) = \alpha(p; 0.25, 2)$, $\alpha^2(p) = \alpha(p; 0.65, 1)$. Both long-run equilibria U_1 and U_2 are unstable

α^2 will eventually dominate. It follows also that S_2 is locally stable. We have found that the market dynamics has two locally stable fixed points, or two possible long-run equilibria: there exist market trajectories at which the first agent dominates and asset prices converge to S_1 , and trajectories at which the second agent dominates and asset prices converge to S_2 .

In Fig. 5, we plot the wealth dynamics⁹ for two different simulations, that is, two different sequences ω , and the same initial conditions $w_0^1 = w_0^2 = 1/2$. In the left panel, despite some initial fluctuations, the second agent dominates in the long-run and the price of the first asset converges to $\pi^{e,2}$. In this case, the market behaves efficiently, as it rewards the strategy based on the most accurate beliefs. In the right panel, conversely, it is the first agent who gains all the wealth in the long run. When the value of p is relatively small, due to the lower risk aversion, agent 2 is more distant than agent 1 from the probability line, so that the market selects against the former and the price of the first asset approaches $\pi^{e,1}$.

We have shown that, despite the long-run price $\pi^{e,2}$ being “closer” to the asset fundamental value $\pi = 1/2$ than the long-run price $\pi^{e,1}$, there are cases where prices converge to the latter. In this example, the market is not informationally efficient. It is not enough to have the “best” beliefs to dominate as preferences matter, too. Finally, notice that if the agent with the best beliefs would be at least as risk adverse as logarithmic preferences imply, or $\gamma^1 \geq 1$, then informational efficiency would be established again. The fact that a particular level of risk aversion is enough to recover informational efficiency is not a general property though, as the next example shows.

5.2 Coexistence of unstable long-run market equilibria

Consider two agents with the same beliefs as in the previous example, i.e. $\pi^{e,1} = 0.25$ and $\pi^{e,2} = 0.65$, but with different risk preferences, $\gamma^1 = 2$ and $\gamma^2 = 1$.

⁹When performing simulations, we need to ensure that the market equilibrium is unique for all possible wealth distribution, and not only when a single agent has all the wealth. For the rules the wealth dynamics of which is plotted in Fig. 5 the result easily holds because their convex combination is always non-increasing and thus has a unique interception with the EMC.




Fig. 5 Simulations of the wealth dynamics for the market represented in the left panel of Fig. 4 with initial condition $w_0^1 = 1/2$ and two different realizations of ω . Left panel: agent 2 dominates. Right panel: agent 1 dominates

Notice that agent 1 is still more risk averse than agent 2. The corresponding EMC plot is depicted in the right panel of Fig. 4. Two single agent equilibria exist: U_1 , where agent 1 dominates and prices are fixed at $p = \pi^{e,1}$; and U_2 , where agent 2 dominates and prices are fixed at $p = \pi^{e,2}$. Consider the relative distance of $\alpha^1(p)$ and α^2 from the probability line π at these two prices. By graphical inspection, it can be easily seen that rule α^2 is closer to π at $\pi^{e,1}$, the price set by rule α^1 , and rule α^1 is closer to π at $\pi^{e,2}$, the price set by rule α^2 . Since the market rewards the rule with the lowest relative entropy, that is, closest to the π line, agent 1 wealth share increases, on average, when price are near to those set by rule α^2 , that is, when the latter is the rule with the larger wealth share. Conversely, agent 2 gains most wealth when agent 1 has the larger wealth share.

Simulations of a generic wealth and price dynamics for this market¹⁰ are presented in the plots of Fig. 6. At the beginning, prices are close to $\pi^{e,1}$, agent 2 is (on average) better-off and takes most wealth, thus driving prices close to $\pi^{e,2}$, where agent 1 is better-off, thus driving prices close to $\pi^{e,1}$ and so on and so forth. Both agents' wealth keeps fluctuating indefinitely in the range (0, 1). Prices behave accordingly and keep fluctuating in the interval $(\pi^{e,1}, \pi^{e,2})$, never converging to the “best” beliefs. Prices and wealth fluctuations have an endogenous cause and do not represent any change in market fundamentals, nor in agents' preferences. From the observation of Fig. 6 alone, believing that assets markets are able to set prices according to the best available information, one would wrongly conclude that the asset dividend process is non-stationary. It is true, however, that the long-run average price seems close to the payoff payment probability $\pi = 1/2$. A precise characterization of how close would require the analysis of the global dynamics, which is not performed here. Upon admitting that prices may have short-run persistent fluctuations, due to agents heterogeneity in preferences and expectations, despite the lack of

¹⁰Also in this example, when performing simulations, we need to ensure that the market equilibrium is unique for all possible wealth distribution and not only for the single survivor case. The result can be established by looking at the first derivative of the convex combination of the two rules, which has limit $+\infty$ for $p \rightarrow 0$, first decreasing and then increasing, and limit $+\infty$ for $p \rightarrow 1$. It follows that the convex combination of the two rules cannot have more than one intercept with the EMC.




Fig. 6 Price and wealth dynamics for the market represented in the right panel of Fig. 4. Left panel: typical wealth dynamics. Right panel: typical price dynamics

market informational efficiency average prices do seem to deliver a not too incorrect information in this example.

5.3 Some special rules

There exists an important price dependent rule that we have not considered because it does not fall in the class of CRRA rules: holding the market portfolio. The market portfolio rule corresponds to $\alpha^M(p) = p$, which coincides with the EMC line in a EMC plot. This rule does not define a unique market equilibrium, in the sense that it is consistent with all prices in $(0, 1)$. No matter the realized states of Nature, it has constant and unitary return, and thus constant wealth, as can be easily derived from Eq. 2. Regarding market selection the market portfolio rule is special because, no matter the rule used by competing agents, by adapting to prevailing prices it never disappears.

Since $\alpha(p; \pi^e, \gamma) \rightarrow p$ as $\gamma \rightarrow \infty$, CRRA myopic maximizing agents approach the market portfolio rule when the coefficient of relative risk aversion increases to infinity (see the right panel of Fig. 3). Thus agents with high risk aversion have higher chances of surviving in the long run. By contrast, agents with low risk aversion tend to invest all the wealth in the under-priced asset (see the right panel of Fig. 3). This behavior puts survivability at great risk in the sense that, resulting in a large relative entropy for a large price range, exposes the agent at the risk of disappearing from the market even when the beliefs of the competing agents are relatively far from the truth.¹¹

A particularly interesting question is to find the best performer, that is, the portfolio rule that gains all wealth when trading against any other rule. In our context an answer to this question is not available, as it would require performing the global market dynamics analysis. However, it can be shown that the (log-optimal) Kelly rule, i.e. to invest proportionally to the correct probabilities, has only stable fixed points and destabilizes all other single survivor fixed points. In this sense, the Kelly rule is a local champion, that is, the best performing rule in the neighborhoods of fixed points

¹¹Even more in danger are risk neutral agents, or risk lovers, who, by investing all their wealth in the asset with the highest expected payoff, would disappear from the market in finite time.

where a single agent survives. When it is present in the market, prices do converge to the informationally efficient level $(\pi, 1 - \pi)$. Notice, however, that the Kelly rule is not only characterized by the correctness of its beliefs but also by the precise structure of its preferences, corresponding to a coefficient of relative risk aversion equal to one. As the previous examples show, the same beliefs with different preferences may not grant informational efficiency. In any case, the Kelly rule may not be the unique survivor.¹²

6 Rule ordering

In the previous section, we learned that the market may select different agents for different price ranges. In this section, we show the consequences of a price-dependent selection on the possibility to order rules according to their relative market performance. Ordering rules would be a desirable property in that it would allow us to express an absolute judgment about their relative virtue. If the order relation “doing better than”, meaning “gaining all wealth when trading with”, could be established, transitivity would imply that, if rule α^+ does better than rule α^- , the same rule α^+ would also do better than any other rule inferior to α^- . Having an order relation would thus allow for an ex-ante characterization of the asymptotic state of the market: given an ecology composed of a collection of different trading rules, irrespectively of their nature or number, it is only the best rules in the collection that will survive and set prices in the long run.

In markets with constant rules, as those analyzed in Section 4, a natural relation can be established using the relative entropy $I_\pi(\alpha)$. We define rule α^+ to be better than rule α^- when its relative entropy is lower or

$$\alpha^+ \succ \alpha^- \quad \text{if and only if} \quad I_\pi(\alpha^+) < I_\pi(\alpha^-). \tag{17}$$

It is immediate to see that this relation is anti-reflexive and asymmetric. Since the relative entropy is a concave function of its argument¹³, then if $I_\pi(\alpha^+) < I_\pi(\alpha^0)$ and $I_\pi(\alpha^0) < I_\pi(\alpha^-)$, for any $\lambda \in [0, 1]$, it holds

$$I_\pi(\lambda\alpha^+ + (1 - \lambda)\alpha^0) \leq \lambda I_\pi(\alpha^+) + (1 - \lambda)I_\pi(\alpha^0) < I_\pi(\alpha^0) < I_\pi(\alpha^-).$$

In particular, $I_\pi(\alpha^+) < I_\pi(\alpha^-)$, so that the relation defined in Eq. 17 is also transitive and, hence, a (strict) order relation. Among constant rules, it thus makes sense to ask which rule is the best, irrespectively of the specific rules trading in a given market. Unfortunately, the possibility to build an order relation is lost when price dependent rules are taken into account as, shown by the following example. Fix $\pi = 1/2$ and consider the following three CRRA portfolio rules (also plotted in Fig. 7): rule one,

¹²Referring to the discussion in footnote 7, in the case of price dependent rules, it is log-optimality that grants local dominance. In the context of price dependent rule and non-Arrow asset structures, the generalized Kelly rule of Amir, Hens, Evstigneev, and Shenck-Hoppeè, i.e. investing proportionally to expected payoff, is not log-optimal. Regarding both points see Section 5.1 of Bottazzi and Dindo (2010).

¹³It directly follows for the convexity of the log function.




Fig. 7 EMC plot representing a market with $\pi = 1/2$ and the three rules considered in Section 6

$\alpha^1(p)$, has $\pi^{e,1} = 0.25$ and $\gamma^1 = 2$; rule two, $\alpha^2(p)$, has $\pi^{e,2} = 0.3$ and $\gamma^2 = 1$; and rule three, $\alpha^3(p)$, has $\pi^{e,3} = 0.65$ and $\gamma^3 = 1$.

When only $\alpha^1(p)$ and $\alpha^2(p)$ are trading, the price p is the unique¹⁴ solution of

$$p_t = \alpha^1(p_t) w_t^1 + \alpha^2(p_t) (1 - w_t^1),$$

which is always between $p = \pi^{e,1} = 0.25$, when agent 1 has all the wealth, and $p = \pi^{e,2} = 0.3$, when agent 2 has all the wealth. Since, for all realized prices, i.e. prices between $\pi^{e,1}$ and $\pi^{e,2}$, rule α^2 has a lower relative entropy than rule α^1 , or $I_\pi(\alpha^2(p)) < I_\pi(\alpha^3(p))$ for all $p \in [\pi^{e,1}, \pi^{e,2}]$, agent 2 dominates for every initial condition: the wealth of the first agent converges to zero and prices converge to $\pi^{e,2} = 0.3$. We can state that $\alpha^2 \succ \alpha^1$ and the market is informationally efficient.

Compare now rule α^2 and rule α^3 . When they are trading, the price p is in between $\pi^{e,2}$, when agent 2 has all the wealth, and $\pi^{e,3}$, when agent 3 has all the wealth. Since for all these prices agent 3 has a lower relative entropy, rule α^3 dominates in the long run and prices converge to $\pi^{e,3} = 0.65$. We can then state that $\alpha^3 \succ \alpha^1$ and, again, the market is informationally efficient.

Transitivity, necessary for \succ being a (strict) order relation, would now imply $\alpha^3 \succ \alpha^1$. When only α^1 and α^3 are trading the price is fixed between $\pi^{e,1}$ and $\pi^{e,3}$, depending on the relative wealth size. Now, importantly, relative entropies $I_\pi(\alpha^3)$ and $I_\pi(\alpha^1)$ do not have the same ranking for all realized prices, as can be appreciated by comparing the relative distance from the probability line π in the plot of Fig. 7. When rule 1 is close to having all wealth and the price of the first asset is close to $\pi^{e,1}$ it is $I_\pi(\alpha^3) < I_\pi(\alpha^1)$. Otherwise, the opposite ranking occurs. As a result, local stability analysis says that both fixed points where a rule dominates are unstable, so that neither rule 1 nor rule 3 can ever be said to dominate. The proposition $\alpha^3 \succ \alpha^1$ is thus false, and upon introducing the relation \sim by saying that two rule are equivalent when none dominates, it holds that $\alpha^3 \sim \alpha^1$.

¹⁴In all the examples of this section, uniqueness of the market equilibrium price can be established by following the same reasoning presented in footnote 10.




Fig. 8 Market with $\pi = 1/2$ and the three rules considered in Section 6. *Left panel:* wealth dynamics. *Right panel:* prices dynamics

By using a counterexample, we have established that, among price-dependent rules, relative dominance is not transitive and thus it does not define an order relation. It is also interesting to analyze the market dynamics when all three rules are trading in the same market. Market clearing prices are now functions of two wealth shares, w^1 for agent 1 and w^2 for agent 2, since $w_t^3 = 1 - w_t^1 - w_t^2$. Since all fixed points are locally unstable, as can be judged by comparing the relative distances to the Kelly rule π in the EMC plot of Fig. 7, wealth fractions, and thus prices, keep fluctuating (see also the right panel of Fig. 8). Again, the market cannot be said informational efficient: despite rule α^3 having the best information, the market fails to set prices according to it.

7 Conclusion and open issues

We have discussed wealth-driven selection in a market for two Arrow securities using a series of simple examples. When portfolio rules are constant, the rule with the lowest relative entropy attracts all wealth in the long run, thus driving prices as close as possible to asset fundamentals. This happens, for instance, when all agents maximize a CRRA expected utility with unitary relative risk aversion coefficient. The picture changes when price dependent portfolio rules are considered, for instance, when the coefficient of relative risk aversion of some traders is different from one. In this broader case, wealth-driven selection works only locally, that is, for given prices, and the market may not be able to select the rule that uses the best information.

In the context of market selection of profit maximizing firms, Nelson and Winter (2002) argue that postulating the capability of markets to select for the global optimum (investor in our context, firm in their) might open four possible issues. Our examples show that the same issues play a role also in asset markets. The first issue is *variety*: selection operates exclusively on the existing competing rules, so that if the optimal rule is not trading in the market, it cannot possibly be selected. This is obviously the case in the asset markets we consider, no matter whether rules are price dependent or constant. The second issue is *behavioral continuity*: a rule that is successful in period t is not necessarily successful in period $t + 1$, when market conditions may have changed. This is precisely what happens in our example with

two locally unstable fixed points, where each rule is the most successful at the prices determined by the other. If, in a given period market, prices are close to the beliefs of rule 1, then rule 2 is (on average) the most successful. The wealth share of the latter does, on average, increase and prices are brought close to its beliefs so that rule 1 becomes now the most successful. Hence behavioral continuity does not hold, no rule is selected, and wealth shares and prices keep fluctuating indefinitely. Notice that this contrasts with the outcome of a market in which only constant portfolio rules are trading. In that case, the ranking induced by relative entropy holds for all market conditions, hence implying behavioral continuity. The third issue is *limited path dependence*, or the fact that there can be transient phases in which markets dynamics eliminates “good” rules. This is precisely what happens when we have two locally stable fixed points. Irrespectively of the ordering of long-run outcomes, wealth selection may drive the market towards either outcomes, depending on initial conditions and/or specific realizations of states of Nature. As such, there exist cases in which the best informed agent is driven out of the market in the transient. By the fourth issue, *profit-induced growth*, it is meant that, even if a rule is earning wealth, it is not granted that this wealth is in fact used to grow. This is equivalent to what happens when modeling traders with heterogeneous propensity to consume. We did not discuss explicitly the issue in this paper, but since the more is consumed the less is invested, a trade off between the quality of beliefs and the quantity of consumption determines the local stability of long-run equilibria; see Bottazzi and Dindo (2010) for more details. As these four main issues suggest, in general it is not possible to order strategies according to their market performance and one cannot be sure of the informational content of prices, as it is not granted that prices reflect the best available information.

To conclude, the unknown part of the story is probably the largest and is composed of many interrelated open questions. First, in this paper we have not dealt with any form of learning, regarding neither the exogenous dividend process, nor endogenous market prices. Both types of learning have been separately investigated. In order to have stability, the former should be not too slow and the latter not too fast (see e.g. Blume and Easley 2010; Bottazzi and Dindo 2010, respectively). Their joint investigation is however still missing. Second, one would like to consider also rules that depend on wealth. This has both a behavioral relevance, as the individual wealth level and its variation over time are natural indicators for the appropriateness of the implemented rule, and a positive appeal, as it allows the study of market selection for a broader class of risk preferences, e.g. constant absolute risk averse agents. Third, further work is needed to characterize the global market dynamics. In more complicated settings than the ones analyzed here, one cannot rely exclusively on local stability results. Finally, an important new direction of investigation is the extension of the present analysis to markets for long-lived assets. Whereas with constant rules the message doesn't change (see e.g. Evstigneev et al. 2009), it is not known what happens when more general portfolio rules are considered.

Acknowledgments We acknowledge financial support from the Institute for New Economic Thinking, INET inaugural grant #220 and the European Commission 6th FP Project DIME (Contract CIT3-CT-2005-513396). All usual disclaimers apply.

References

- Amir R, Evstigneev I, Hens T, Schenk-Hoppé K (2005) Market selection and survival of investment strategies. *J Math Econ* 41:105–122
- Anufriev M, Bottazzi G (2010) Market equilibria under procedural rationality. *J Math Econ* 46:1140–1172
- Anufriev M, Dindo P (2010) Wealth-driven selection in a financial market with heterogeneous agents. *J Econ Behav Organ* 73:327–358
- Anufriev M, Bottazzi G, Pancotto F (2006) Equilibria, stability and asymptotic dominance in a speculative market with heterogeneous agents. *J Econ Dyn Control* 30:1787–1835
- Blume L, Easley D (1992) Evolution and market behavior. *J Econ Theory* 58:9–40
- 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 (2010) Heterogeneity, selection, and wealth dynamics. *Ann Rev Econ* 2:425–450
- Bottazzi G, Dindo P (2010) Evolution and market behavior with endogenous investment rules. LEM Working Paper 2010-20. Scuola Superiore Sant'Anna, Pisa
- Breiman L (1961) Optimal gambling systems for favorable games. *Proceedings of the 4th Berkeley symposium on mathematical statistics and probability* 1:63–68
- Evstigneev I, Hens T, Schenk-Hoppé K (2008) Globally evolutionary stable portfolio rules. *J Econ Theory* 140:197–228
- Evstigneev I, Hens T, Schenk-Hoppé K (2009) Evolutionary finance. In: Hens T, Schenk-Hoppé K (eds) *Handbook of financial markets: dynamics and evolution*. North-Holland (Handbooks in Economics Series)
- Hommes C (2006) Heterogeneous agent models in economics and finance. In: Judd K, Tesfatsion L (eds) *Handbook of computational economics, vol 2. Agent-based computational economics*. North-Holland Handbooks in Economics Series, Amsterdam
- Kelly J (1956) A new interpretation of information rates. *Bell Syst Tech J* 35:917–926
- LeBaron B (2006) Agent-based computational finance. In: Judd K, Tesfatsion L (eds) *Handbook of computational economics, vol 2. Agent-based computational economics*. North-Holland (Handbooks in Economics Series)
- Nelson R, Winter S (2002) Evolutionary theorizing in economics. *J Econ Perspect* 16:23–46
- Sandroni A (2000) Do markets favor agents able to make accurate predictions. *Econometrica* 68(6):1303–1341
- Sandroni A (2005) Market selection when markets are incomplete. *J Math Econ* 41:91–104