

Bottazzi, Giulio; Dindo, Pietro

Working Paper

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

LEM Working Paper Series, No. 2011/11

Provided in Cooperation with:

Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies

Suggested Citation: Bottazzi, Giulio; Dindo, Pietro (2011) : Selection in asset markets: The good, the bad, and the unknown, LEM Working Paper Series, No. 2011/11, Scuola Superiore Sant'Anna, Laboratory of Economics and Management (LEM), Pisa

This Version is available at:

<https://hdl.handle.net/10419/89384>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Laboratory of Economics and Management

Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)

Tel. +39-050-883-343 Fax +39-050-883-344

Email: lem@sssup.it Web Page: <http://www.lem.sssup.it/>

LEM

Working Paper Series

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

Giulio Bottazzi
Pietro Dindo

LEM and CAFED, Scuola Superiore Sant'Anna, Pisa, Italy

2011/11

May 2011

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

Giulio Bottazzi Pietro Dindo

May 17, 2011

LEM and CAFED, Scuola Superiore Sant'Anna,
Piazza Martiri della Libertà 33, 56127 Pisa, Italy

Abstract

We provide simple examples to illustrate how wealth-driven selection works in asset markets. Our examples deliver both good and bad news. The good news is that if individual assets demands are expressed as a fractions of wealth to be invested in each asset, e.g. because traders maximize an expected Constant Relative Risk Aversion utility with unitary coefficient, then market rewards the best informed agent. As a result asset prices eventually reflect this information and the market can be said informationally efficient. However, and this is the bad news, when asset demands are expressed as price dependent fractions, e.g. they are derived from the maximization of expected Constant Relative Risk Aversion utility with non unitary coefficients, anything can happen and the informational content of long-run prices strongly depend on the ecology of traders' preferences and beliefs. Our examples 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
JEL Classification: D50, D80, G11, G12

*This work benefited from feedbacks by participants to workshops and seminars held in (order of appearance) Paris, Amsterdam, Pisa, Udine, London, Budapest, Rotterdam, Leuven, Cornell University, Greqam Marseille, Florence. 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.

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 draw on the analysis of the simplest economic framework where the problem is well understood and some clear answers are available: a repeated market for two Arrow securities. Each asset represents a different way to transfer current wealth to future consumption or investment. Whereas aggregate wealth is fixed ex-ante and depends only on the specific state of Nature which is realized, prevailing asset prices and the distribution of wealth among agents, being endogenously determined, depend also on the ecology of behaviors. In fact realized returns and the implied distribution of wealth are determined both by chance (Nature) and, through prices, by agents investment decisions. Each agent behavior is modeled through an explicit asset demand function. This function can represent a behavioral rule or an investment rule derived from the maximization of an expected utility, given individual preference and beliefs about the unfolding of future states of Nature. Since we are primarily interested in analyzing the selective capability of markets we drop the assumption of perfect foresight by the part of agents or rational-expectation. These assumptions would indeed sterilize the effect of the trading-induced wealth reallocation, as all the problem would be reduced to an ex-ante identification of possible equilibria. In a dynamic framework, the key issue is to understand which kind of asymptotic states are possible for the economy. Do prices and wealth distributions converge toward a unique equilibrium and, if so, which properties it displays? Which type of beliefs, preferences, behaviors, are rewarded by the trading? Is the market capable of rewarding the best informed agents? Relatedly, are long-run prevailing prices revealing the best available information? Above all, what is the role of market selection in shaping aggregate economic outcomes?

Despite general results are still missing, also for the simple framework with Arrow securities discussed in our examples, most economic and finance models rely on the market ability to select for agents using rational expectations or, more generally, for agents whose equilibrium prices reflect the correct, or the best available, information about the probability distribution of the states of Nature. 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 indeed converge to a long-run equilibrium where the best informed agent has all wealth and prices reflect this information. In this case the market could be said informationally efficient, as it efficiently conveys, through prices, all the information available about fundamentals.

The examples of this paper are meant to illustrate when the above idea is correct and when, and why, it is not. The general message is as follows. On the one hand, when the fraction of wealth invested in each asset, also named portfolio rule, do not depend on prices (if agents are intertemporal maximizers this corresponds to having a Constant Relative Risk Aversion utility with unitary risk aversion coefficient) the dominating agent is the one whose equilibrium prices are indeed closest to the probability distribution of the states of Nature. Thus the agent possessing the best information about the fundamental process is rewarded by the market, and prices are set so as to reveal this information. This is the good side of the story.

On the other hand, when portfolio rules are price dependent, for example when the utility to be maximized is still of the Constant Relative Risk Aversion (CRRA) type but with a risk aversion coefficient different from one, long-run prices do not need to represent the best available information. Wealth-driven selection may lead market dynamics to get stuck in suboptimal, informationally inefficient, equilibria or even display endogenous fluctuations. It

is easy to construct examples where, depending on initial conditions and the realization of states of Nature, prices do or do not reveal the best available information, or where prices and wealth keep fluctuating indefinitely. More in general, not knowing the types of demand expressed by agents, 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 the fundamentals. This is the bad side of the story.

In our analysis we shall present cases where market selection works and cases where it doesn't. Technically the dynamics of agents wealth and asset prices can be formalized as a random dynamical system. Given agents' wealth fractions, agents portfolio rules, and asset prices at time t , the market dynamics is a random map whose actual form is fixed once the uncertainty about the next realized state of nature is resolved. Different realizations of the market map 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 agents' trading about the realized state of nature and condition of market, will in turn determine the next prices, until the next state of nature is revealed. This process is repeated in time and can be described as a discrete-time stochastic process. The analysis of the long run wealth and price distributions, our ultimate object of interest, becomes thus rather technical. In our previous work Bottazzi and Dindo (2010) we have focused on the analytical investigation of the random dynamical system representing the market dynamics and have been able to derive local stability conditions. A contribution of this paper is that, relying on those results, we are able to discuss this seemingly complicated issue using simple example and with the help of a graphical tool. First, by plotting asset's agents portfolio rules against a normalization of the supply, we are able to characterize the possible long-run selection equilibria as the short-run Walrasian equilibria where only one agent has positive wealth to invest in the assets. Second, comparing the distance of each individual rule to a reference, relative entropy minimizing, rule we are able to characterize long-run single agents equilibria local stability.

The outline of this paper is as follows. In the next section we provide a brief introduction to the series of recent works whose results provided the basis of the present paper. 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 rules, or maximization of a CRRA utility with unitary risk aversion coefficient, and characterize the outcome of the long-run market dynamics. This analysis illustrate the good side of the story, as we show that the market rewards the best informed agent. However, the analysis begs the question of what happens when risk aversion is higher or lower than one. In Section 5 we study the market dynamics when agents demand are derived from the maximization of a CRRA utility for any given risk aversion γ . This will have the effect that investment rules will not be fixed, but dependent on prices. As a result the market selective capability in favor of the best informed agent turns from global into local, and the convergence of the market toward informationally efficient prices is not granted anymore. Examples in Sections 5.1 and 5.2 will point to the possible sources of problems. In Section 6 we show that the local nature of market selection leads to the impossibility to order rules according to their mutual survivability. Section 7 concludes and set the agenda for future research. An extension of our model to consider non-myopic agents is presented in Appendix A.

2 Selection in asset markets: the known

The investigation of the ability of markets to redistribute resources in favor of the rational or better informed agents has started with Blume and Easley (1992), see also Blume and Easley (2010) They study wealth-driven selection on a class of portfolio rules that do not depend upon prices and are thus expressed as fixed fraction of wealth to be invested in each asset. They find that average wealth growth differences are equal to differences of relative entropies of subjective state of Nature probabilities with respect to the correct probabilities. The 'closer', in relative entropy terms, an agent is to the correct probabilities the higher her average wealth growth, the higher the impact she will have on prices in the long run. As a consequence the agent with lowest relative entropy will dominate all others and set long-run prices and the market does seem to reward some type of rational behaviors, or preferences, over others. Informational efficiency holds and price converge to the best available information. We draw upon this contribution in Section 4 where we present the good side of assets market selection.

The literature that has followed has been devoted to generalize the hypothesis under which informationally efficiency holds. Evolutionary Finance (see Evstigneev et al., 2009, for a recent survey) has taken the direction of generalizing the market structure, while considering agents behavior in the same class of constant rules. Generalizations involve the asset payoff matrix, also incomplete markets being analyzed; the process ruling states of Nature, extended to be Markovian; and the nature of assets, long-lived assets have been studied alongside short-lived assets. Two representative papers are Amir et al. (2005), where the extended analysis is performed for short-lived assets, and Evstigneev et al. (2008), which instead concentrates on long-lived assets. No significance differences with results from the seminal Blume and Easley (1992) analysis are found: relative entropy still rules and the "best" informed agent dominates. It is important to mention that both Evolutionary Finance and the original Blume and Easley contribution also consider a generalization of agents' rules in terms of learning, that is, allowing agents to learn about assets payoffs, or state of Nature, processes. Both find that, as long as the learning algorithm is at least as fast as the Bayesian, and has the truth in its support, informational efficiency still holds.

A different group of works has instead focused on generalizing the class of rules considered to encompass any investment decision explicitly coming from utility maximization, so that assets demand is not necessarily expressed as a fixed, beliefs 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 for 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 strong assumption that agents are able to coordinate on having perfect foresight on realized prices. This not only implies that agents should coordinate their prices expectations, but also that they should guess the correct price. Technically perfect foresight is essential because it leads to allocative efficiency which in turns it is shown to imply informational efficiency. Moreover, whereas the previous contributions follow an explicit approach, in that market equilibrium prices and the resulting wealth fractions are explicitly derived and their distribution is studied as time passes, both Sandroni (2000, 2005) and Blume and Easley (2006) follow an implicit approach, in that they derive the long run distribution of wealth as a consequence of first order conditions, market equilibrium, and perfect foresight, but explicit prices and wealth fraction are never computed.

The effort to investigate wealth-driven selection for a general class of rules, but dropping perfect foresight, has started with our recent contribution Bottazzi and Dindo (2010), where

we take a standard Evolutionary Finance market model with short-lived assets and add price dependent portfolio rules to the analysis. We find that market selection operates locally in the prices space. Comparison of rules is thus possible only at given prices, making possible to build examples where no rule gains all wealth in the long run or where two rules gain all wealth for different initial conditions or different states of Nature process realizations. The reason is that the relative entropy of investment rules with respect to the underlying dividend payment process depend on realized prices. Hence, the same is true for their average relative wealth growth rate. Wealth-driven selection is a local, that is, price dependent, mechanism. Whereas with constant rules the local and global analysis coincide so that market selection operates on an effective global scale and is able to achieve a global optimum, when rules become price dependent, local and global results do not necessarily coincide, so that market selection may not be able to push the market to the global optimum. Given two rules, it may well happen that the first rule has a higher wealth growth rates at the prices determined by the second, while the second has a higher wealth growth rate at the prices determined by the first, so that none can prevail. Alternatively it may well happen that both rules have higher wealth growth at their prices so that both might, in principle, prevail. Market prices do not need to converge to the level reflecting the “best” available information, as they may instead keep fluctuating around it or, even, converge to a suboptimal, worse-information-revealing, level. We draw upon this contribution in Sections 5-6 where we present the good side of assets market selection.

3 The model

Assume that there are 2 possible states of Nature, state 1 and state 2. Their unfolding in discrete time can be represented by a Bernoulli process $\omega = (\omega_1, \dots, \omega_t, \dots)$, with $\omega_t \in \Omega = \{1, 2\}$ and $\omega_t = 1$ with probability π for every $t \in \mathbb{N}$. In order to transfer wealth intertemporally, and thus being able to consume in the future, I agents can trade in two short-lived Arrow securities. Asset $k = 1, 2$, exchanged at period t at price $p_{k,t}$, pays one unit of the consumption good at period $t + 1$ if $\omega_t = k$, zero otherwise. The consumption good is used as the numéraire of the economy. Each agent demand is denoted as fractions 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 her chosen fraction and w_t^i her wealth. We refer to the vector α_t^i of fractions $\alpha_{k,t}^i$ as the portfolio 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 at this final period. All intermediate wealth is thus saved and re-invested in assets. Adding the assumption that short positions are not considered, this implies that portfolio rules belong to the set $\{\alpha \in \mathbb{R}^2 \mid \alpha_1 + \alpha_2 = 1 \text{ and } \alpha_k > 0, k = 1, 2\}$.

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

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

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

$$p_{k,t} = \sum_{i=1}^I \alpha_{k,t}^i w_t^i. \tag{3.1}$$

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 $k = \omega_{t+1}$ she has bought at time t , or

$$w_{t+1}^i = \frac{\alpha_{\omega_{t+1},t}^i w_t^i}{p_{\omega_{t+1},t}}. \quad (3.2)$$

By summing up the previous equation over all the I agents and remembering the definition of price in (3.1), it is clear that, no matter the initial wealth, in all subsequent periods the total wealth in the economy is equal to one. Since portfolio rules are constrained to be positive and not greater than one, this implies that

$$p_{1,t} + p_{2,t} = 1 \quad \text{for every } t > 0. \quad (3.3)$$

This relation can be used to normalize prices so that will use the price of the first asset only, sometimes simply denoted as p_t , to quantify the exchange.

For definiteness and illustrative purposes in this contribution we choose to restrict our analysis to the class of rules derived from the maximization of the expected utility of next period wealth, with the utility function of CRRA type. However, as it will become clear in the following discussion, the results presented 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 would derive the vector rule (α_1, α_2) by solving the problem

$$\text{Max} \left\{ \pi^e U \left(\frac{\alpha_1}{p} \right) + (1 - \pi^e) U \left(\frac{\alpha_2}{1 - p} \right) \right\}, \quad \text{s.t.} \quad \alpha_1 + \alpha_2 = 1 \quad \text{and} \quad \alpha_1, \alpha_2 \geq 0,$$

where

$$U(w) = \begin{cases} \frac{w^{1-\gamma}}{1-\gamma} & \gamma \in (0, 1) \cup (1, +\infty) \\ \log(w) & \gamma = 1 \end{cases}. \quad (3.4)$$

Notice that choices have been constrained to meet the intertemporal budget constraint (3.2) and the price normalization in (3.3).¹ Standard first order conditions leads to the optimal rule. The fraction of wealth to be invested in the two assets reads

$$\alpha_1(p; \pi^e, \gamma) = \frac{(\pi^e(1-p)^{1-\gamma})^{\frac{1}{\gamma}}}{(\pi^e(1-p)^{1-\gamma})^{\frac{1}{\gamma}} + ((1-\pi^e)p)^{\frac{1}{\gamma}}}, \quad (3.5)$$

$$\alpha_2(p; \pi^e, \gamma) = 1 - \alpha_1(p; \pi^e, \gamma).$$

When $\gamma = 1$, the previous solution collapses to the price independent rule

$$\alpha_1(p; \pi^e, 1) = \pi^e, \quad \alpha_2(p; \pi^e, 1) = 1 - \pi^e. \quad (3.6)$$

As we will see in the following sections, the analysis of market dynamics, and especially its asymptotic characterization, critically depends on the nature of portfolio rules.

¹Since the market lasts an arbitrary large number of periods T , it would be more consistent to derive demands from the maximization of period T expected wealth. The task cannot be achieved without modeling periods beliefs about the probability distribution of the states of Nature as well as price expectations. Since both, but especially the latter, are hard tasks, we restrict here to a myopic optimizer. Despite this limitation, we derive in Appendix an answer to the T periods maximization problem that, under some simplifying assumptions, can be used in our framework. In the Appendix, we also show that in case of log utility, or unitary risk aversion $\gamma = 1$ in (3.4), the found answer is general and does not depend on our simplifying assumptions.

4 The good: wealth-driven selection and market dynamics with constant portfolio rules.

In this section we consider the benchmark case of agents who use constant rules, that is, agents who invest fixed fractions of their wealth in each asset, irrespectively of their wealth or assets prices. As shown in the previous section, this investment rule corresponds to the maximization of the expected logarithmic utility of next period wealth with CRRA coefficient $\gamma = 1$.²

Consider first the case of a market with a single investor. Denoting her wealth simply by w_t and her rule as α , her belief about state 1 occurrence, the market dynamics in (3.1-3.2) can be written as

$$w_{t+1} = \begin{cases} \frac{\alpha w_t}{p_t} & \text{if } \omega_{t+1} = 1 \\ \frac{(1-\alpha)w_t}{1-p_t} & \text{if } \omega_{t+1} = 2 \end{cases}, \quad (4.1)$$

where the price of the first asset is

$$p_t = \alpha w_t,$$

and we have used the fact that both prices and portfolio fractions add up to one. No matter the realization of the state of Nature, the agent, by investing a positive amount in both assets, will secure all future wealth (we are assuming that the agent puts a positive probability on both states, or $\alpha \in (0, 1)$). Wealth evolution does not pay any role and market equilibrium fixes the price of the first asset at her belief α , i.e.

$$p_t = \alpha \quad \text{for every } t. \quad (4.2)$$

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 left hand side (supply market value) and the right hand side (demand market value) of (4.2). Their crossing, E_1 , fixes the equilibrium price, which in this case is trivially equal to α . Consistently with our previous work (see e.g. Anufriev et al., 2006; Anufriev and Bottazzi, 2010; Anufriev and Dindo, 2010) we shall rename the supply curve, the diagonal line of this plot, the Equilibrium Market Curve (EMC), in that it is the locus of all possible long-run equilibria of our economy. The EMC plot will provide a useful graphical tool to illustrate our results.

In order to study the role of the market in redistributing wealth among agents, and the ensuing asset prices behavior, we need to add a second agent. Rules, or beliefs, are now given by α^1 and α^2 respectively. As before, each agent wealth evolves according to

$$w_{t+1}^i = \begin{cases} \frac{\alpha^i w_t^i}{p_t} & \omega_{t+1} = 1 \\ \frac{(1-\alpha^i)w_t^i}{1-p_t} & \omega_{t+1} = 2 \end{cases} \quad i = 1, 2, \quad (4.3)$$

but now the price of the first asset is given by the combination of both agents rules

$$p_t = \alpha^1 w_t^1 + \alpha^2 w_t^2. \quad (4.4)$$

²Were the reader, for any reason, adverse 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 will of course remain the same.




Figure 1: Left Panel: EMC supply and demand plot representing a single agent market investing $\alpha^1 = 0.4$. The only possible outcome of the dynamics is E_1 . Right Panel: EMC supply and demand plot representing a two agents market. Demands are represented by $\alpha^1 = 0.4$ and $\alpha^2 = 0.7$. E_1 and E_2 are the two single survivor long-run equilibria.

Since the total wealth is one, the price lies between the beliefs of the first agent and the beliefs of the second agent, with weights given by their wealth shares. If one agent has all the wealth at a given period then prices will be fixed by her investment rule and she will have all the wealth at all subsequent periods. So if $w_t^1 = 1$ or $w_t^2 = 1$ the market behaves as in the single agent case, where uncertainty and wealth evolution do not play any role, as in (4.1). 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. At each period the market rewards the agent with an higher stake in the 'lucky' asset, the asset that turns out to pay the dividend. It is this process of wealth dynamics that interests us.

We can use the EMC plot to illustrate a two-agent market, as in the right panel of Fig. 1. Points E_1 and E_2 are the crossing of each agent rule with the EMC. They are also the long-run market equilibrium, that is the outcome of the market dynamics, if only agent 1 or agent 2 have positive wealth when $T \rightarrow \infty$. For this reason they are named single survivor long-run equilibria. Technically, points E_1 and E_2 are deterministic fixed points of the random dynamical system specified by (4.3): no matter the realization of state of nature, the market dynamics starting in E_1 (or E_2) will stay there for all subsequent periods. A first characterization of the long-run behavior of the market depends upon the stability of these two fixed points. More in general, the issue we want to address is whether the market will converge toward a single agent equilibrium, and if so which one among the two, or keep fluctuating between E_1 and E_2 .

In order to investigate the market dynamics when the initial condition is different from E_1 or E_2 , we can use the evolution of the wealth ratio. From (4.3), knowing that $\omega_{t+1} = 1$ with probability π , one can derive

$$\frac{w_{t+1}^1}{w_{t+1}^2} = \begin{cases} \frac{\alpha^1 w_t^1}{\alpha^2 w_t^2} & \text{with probability } \pi, \\ \frac{1-\alpha^1 w_t^1}{1-\alpha^2 w_t^2} & \text{with probability } 1-\pi. \end{cases} \quad (4.5)$$

The wealth of the agent with the higher stake in asset 1, that is the one with the larger 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

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

so that, 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

$$\lim_{T \rightarrow \infty} \frac{1}{T} \log \frac{w_T^1}{w_T^2} = \left(\pi \log \frac{\alpha^1}{\alpha^2} + (1-\pi) \log \frac{1-\alpha^1}{1-\alpha^2} \right) \quad (4.6)$$

Define now the relative entropy of the investment strategy $(\alpha, 1-\alpha)$ with respect the true probability measure as

$$I_\pi(\alpha) = \pi \log \frac{\pi}{\alpha} + (1-\pi) \log \frac{1-\pi}{1-\alpha}. \quad (4.7)$$

$I_\pi(\alpha)$ is always non-negative and is zero if and only if $\alpha = \pi$. The relative entropy is a measure of information loss: the lower it is, the greater the agreement of beliefs $(\alpha, 1-\alpha)$ with the true Bernoulli distribution $(\pi, 1-\pi)$. Substituting the expression for the relative entropy in (4.6) 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 $I_\pi(\alpha^2) > I_\pi(\alpha^1)$ then when T increases $w_T^2 \rightarrow 0$ at an exponential rate. Agent 1 dominates, that is, gains all the wealth in the long-run. If instead $I_\pi(\alpha^2) < I_\pi(\alpha^1)$ then when T increases $w_T^1 \rightarrow 0$ at an exponential rate. In this case agent 2 dominates. We have established that the agent with the lowest relative entropy dominates. Since here rules can be interpreted as beliefs, this implies that the agent whose beliefs have the lowest relative entropy with respect to the true probabilities will eventually dominates the economy and, due to Walrasian market clearing, long-run prices will reflect these beliefs, that is, in terms of relative entropy, will move as close as possible to the probability distribution $(\pi, 1-\pi)$.

Going back to our example of Fig. 1, we can use the relative entropy to analyze the global stability of single survivor fixed points E_1 and E_2 . First, we add the line π to the two-agent EMC plot presented in the right panel to obtain Fig. 2. Second, we notice that the relative entropy of a strategy α is monotonically related to its distance from the line π . Indeed the expression in (4.7) is monotonically decreasing in α if $\alpha < \pi$ and monotonically increasing if $\alpha > \pi$. If two rules α^1 and α^2 are such that $\alpha_1 > \alpha_2 > \pi$ or $\alpha_1 < \alpha_2 < \pi$, then rule α^2 , having smaller distance from π , has also smaller relative entropy so that agent 2 dominates in the long run. If two rules are such that $\alpha_1 > \pi > \alpha_2$ then the monotonicity argument does not directly apply. However, when $\pi = 1/2$, as in all the examples of this paper, the function $I_\pi(\alpha)$ is symmetric around π , so that the distance in terms of relative entropy can be directly derived from the Euclidean distance.³ So, since in Fig. 2 the π line is closer to α^1 than to α^2 ,

³In general one can use the function $I_\pi(\alpha)$ to rescale the vertical axis of the EMC so that entropic distance can be simply inferred by visual inspection




Figure 2: EMC supply and demand 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.

we can conclude that as long as both agents start with positive initial wealth, no matter how bigger in favor of agent 2, agent 1 will gain all wealth in the long run and, at the same time, the price of the first asset converge to the investment rule, or belief, α^1 . Technically we have just established that, in the case depicted in Fig. 2 the only (globally) stable equilibrium is E_1 . In general, the market results informationally efficient as the wealth dynamics led prices to converge as close as possible, in terms of relative entropy, to the probability distribution $(\pi, 1 - \pi)$.

From the previous discussion it is clear that the best possible constant rule, the rule which can never be beaten, is the rule with minimal entropic distance from the process governing the succession of the states of the word. Namely, the rules that prescribe to invest a fraction of wealth π in the first asset and a fraction $1 - \pi$ in the second asset. This simple rule, invest proportionally to the probabilities, is named the Kelly rule after Kelly (1956). In terms of the maximization of logarithmic utility, the Kelly rule corresponds, not surprisingly, to an investor possessing correct beliefs about the process governing the state of Nature. When it is present in the market, the Kelly rule beats any other different rule and bring prices to assets expected values, π and $1 - \pi$ respectively.

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

Analogous 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 also set long-run prices at her investment rule, or at that deterministic fixed point “closest” to the probability line π , making the market informationally efficient. The market will price assets exactly at their




Figure 3: Portfolio rules derived by the maximization of a CRRA utility of final wealth, as given in (3.5). Left panel: CRRA coefficients are 1 and values close to it. Right panel. CRRA coefficients are 1 and extreme values.

probability revealing prices. This is not, however, the end of the story.

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

The reassuring results presented in the previous section are very much dependent on the fact that portfolio rules are constant and do not depend on prices, that is, in the context of CRRA maximizing agents, on the fact that risk aversion coefficient γ is exactly one for all traders.

In this section we investigate market behavior under wealth-driven selection when agents have values of γ that differ from one. As it is clear from (3.5) this causes rules to 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 plots are consistent with the intuitive notion of risk aversion. When γ is close to zero, agents are close to be 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 sure, but unitary, return. In these cases the value of α remains closer to the EMC for a wider price range.

When rules depend on prices agents' wealth evolve according to

$$w_{t+1}^i = \begin{cases} \frac{\alpha^i(p_t)w_t^i}{p_t} & \omega_{t+1} = 1 \\ \frac{(1-\alpha^i(p_t))w_t^i}{1-p_t} & \omega_{t+1} = 2 \end{cases}, \quad i = 1, 2 \quad (5.1)$$

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

$$p_t = \alpha^1(p_t)w_t^1 + \alpha^2(p_t)w_t^2. \quad (5.2)$$

In the rest of the paper we assume that there exists a unique price that solves the market equilibrium equation, and that the implicit function theorem can be applied in its neighborhood so that an explicit local map from the wealth distribution to prices can be derived.⁴ The

⁴Conditions ensuring that this property holds are those usually ensuring existence and regularity of a Walrasian equilibrium.

evolution of wealth ratios still reads

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

but now, due the price dependence of the α s, the value of the ratio $\frac{w_T^1}{w_T^2}$ after T time steps depends on the price history. At each time step, the expected log-growth rate of the wealth ratio can still be written as the difference of the relative entropy of the two strategies

$$\mathbb{E} \left[\log \frac{w_{t+1}^1}{w_{t+1}^2} - \log \frac{w_t^1}{w_t^2} \right] = \pi \log \frac{\alpha^1(p_t)}{\alpha^2(p_t)} + (1 - \pi) \log \frac{1 - \alpha^1(p_t)}{1 - \alpha^2(p_t)} = I_\pi(\alpha^2(p_t)) - I_\pi(\alpha^1(p_t)),$$

but the relative entropy of the latter depend on the prices prevailing in the market as their distance from the Kelly rule in the EMC plot varies with the value of p at which it is computed. This implies that market selection trough entropy minimization operates only locally, that is, for given prices, rather than globally. As a result there may exist prices where a given rule is the nearer to the Kelly rule flat line, thus having the lowest relative entropy and gaining wealth shares, and prices where the opposite happens.

The fact that rules relative performance becomes a price dependent, local property implies that a global stability analysis similar to the one done in the previous section is not feasible. We can however still analyze the occurrence of the asymptotic states in which a single agent dominates the market. Indeed in this case prices converge to constant levels, determined by the survivor's rule, possibly fluctuating around it in the process. If it turns out that these fluctuations do not ruin the actual tendency of the system toward the constant price level, then we can conclude that the dominance of the said agent has positive probability, at least for a set of initial conditions. Thus our strategy will be, first, to identify single survivor equilibria, where all the wealth is in the hands of one agent and prices are set by her rule α , and then check whether market dynamics is stable for prices close to α .⁵ For example, in the two agents market dynamics illustrated in the left panel of Fig. 4, local stability analysis amounts to study 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 at the solution of $p = \alpha^2(p)$. If two rules have a common intersection \bar{p} with the EMC, that is $\alpha^1(\bar{p}) = \alpha^2(\bar{p}) = \bar{p}$, then it is possible to have fixed points where more agents survive, i.e. both w^1 and w^2 are different from zero. At these fixed points, named multiple survivor equilibria, both agents behave as a representative single survivor investing $\alpha^1(\bar{p})$.

In the following sections we use simple examples to illustrate the implication of the price dependency of rules for the market informational efficiency. For this purpose we will consider rules derived by the maximization of CRRA utility function of the next period wealth. Every agent i will be characterized by an individual belief about π and a coefficient of relative risk aversion, denoted respectively $\pi^{e,i}$ and γ^i . Notice that, no matter the value of γ , all 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

⁵In 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, around the fixed points, the system can be linearized and agents can be described as if they were using suitably defined constant rules, so that the analysis can proceed along the lines of Sec. 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) for details.




Figure 4: EMC supply and demand plots representing two agents market with price dependent rules. Left panel: $\pi = 1/2$, α^1 is a CRRA demand with $\pi^{e,1} = 0.25$ and $\gamma^1 = 1$, and $\alpha^2(p)$ is a CRRA demand with $\pi^{e,2} = 0.65$ and $\gamma^2 = 0.5$. 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 stable. Right panel: $\pi = 1/2$, $\alpha^1(p)$ is a CRRA demand with $\pi^{e,1} = 0.25$ and $\gamma^1 = 2$, $\alpha^2(p)$ is a CRRA demand with $\pi^{e,2} = 0.65$ and $\gamma^2 = 1$. Both long-run equilibria U_1 and U_2 are unstable.

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. This is an example of stable multiple survivor equilibrium. The market is efficient in the sense that prices represent the homogeneous beliefs and preferences do not count.

The dynamics is more interesting when agents disagree on the value of π . Take two agents with $\pi^{e,1} \neq \pi^{e,2}$. If they have the same risk-aversion $\gamma = 1$, then we are back to the case of constant rules, so that the best informed agent dominates in the long run (a global result) and the market is informationally efficient. But what if an agent, say the first, has a different risk-aversion? Now the distance of the two rules from the probability line, and thus their relative entropy, changes with the price. In particular it is not anymore granted that the equilibrium where the best informed agent dominates and prices converge to reflect the best available information is globally stable. The type of behavior that can be observed depends very much on the exact choices of probability estimates and risk preferences. We present two typical occurrences in what follows.

5.1 Coexistence of stable long-run market equilibria

For the first example we take $\pi = 1/2$ and agent one has $\pi^{e,1} = 0.25$ and $\gamma^1 = 1$, whereas agent two has $\pi^{e,2} = 0.65$ and $\gamma^2 = 0.5$, the same as depicted in the left panel of Fig. 4. There exists two single survivor equilibria: S_1 , where only agent one 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 the rules α^1 and α^2 from π at the price corresponding to S_1 one finds that $\alpha^2(\pi^{e,1}) - 1/2 > 1/2 - \pi^{e,1} = 1/4$. 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 at each time step and will eventually dominate. If the market share of the first agent is almost 1, and we imagine to forcefully move a small fraction of wealth from agent 1 to agent 2, the dynamics of




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

the market will bring the system back to the situation in which $w^1 = 1$. We have established that S_1 is locally stable.

Conversely, computing the distance of the two rules from π at the prices corresponding to S_2 , one has that $\pi^{e,2} - 1/2 = 3/20 < 1/2 - \pi^{e,1}$, that is 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, at each time step and rule α^2 will eventually dominates. It follows that also 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 along which the first agent dominates and asset prices converge to her beliefs, and trajectories where the second agent dominates and asset prices converge to her beliefs.

In Fig. 5 we plot the wealth dynamics 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 approaches $\pi^{e,2}$. Since this individual expectation is the nearest to the true probabilities among the strategies active in the market, in this case the latter seems informationally efficient. In the right panel, however, 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 coefficient of risk aversion, agent 2 is more distant than agent 1 from the probability line, so that 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 $\pi^{e,2}$ is “closer” to the fundamental price $\pi = 1/2$ than the long-run price $\pi^{e,1}$, there are cases where prices converge to $\pi^{e,1}$. The market is not informationally efficient anymore: even if a “better” long-run equilibrium is possible wealth-driven selection, acting locally and not globally, is not always able to drive the market dynamics there. In any case, since informational efficiency is achieved only for certain initial conditions we can rule out the hypothesis that the market is generically informationally efficient. In particular, it is not enough to have the “best” beliefs to dominate as preferences also seem to matter. Notice that if in this market the agent with the best beliefs would be at least as risk adverse as a logarithmic preferences imply, or $\gamma^1 \geq 1$, then informational efficiency would be established again. The fact that a particular type of preferences is enough to impose informational efficiency is not a robust property though, as the next example shows.




Figure 6: Price and wealth dynamics in a market with $\pi = 0.5$, $\alpha^1(p)$ is a CRRA demand with $\pi^{e,1} = 0.25$ and $\gamma^1 = 2$, $\alpha^2(p)$ is a CRRA demand with $\pi^{e,2} = 0.65$ and $\gamma^2 = 1$. Both long-run equilibria U_1 and U_2 are unstable. Left panel: typical wealth dynamics. Right panel: typical price dynamics.

5.2 Coexistence of unstable long-run market equilibria

Consider the same market and 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 different risk preferences, $\gamma^1 = 2$ and $\gamma^2 = 1$. 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 her beliefs, $p = \pi^{e,1}$; and U_2 , where agent 2 dominates and prices are fixed at her (“better”) beliefs, or $p = \pi^{e,2}$. Consider the relative distance of $\alpha^1(p)$ and $\alpha^2(p)$ from the probability line π at these two prices. It can be easily seen, just by graphical inspection, 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, in this case agent 1 is rewarded when the prices are near to those set by rule α^2 , that is when the latter is the rule with the largest wealth share. Conversely, agent 2 is rewarded when agent 1 has the largest wealth share.

Simulations of a generic wealth and price dynamics for this market 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 her beliefs $\pi^{e,2}$, 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 $(0, 1)$, with agent one dominating in some periods and agent two dominating in other periods. Prices behave accordingly and keep fluctuating in the interval $(\pi^{e,1}, \pi^{e,2})$. Also in this example the market, by rewarding the agents closest to the best rule locally, rather than globally, fails to achieve informational efficiency. The fluctuations in price level have an endogenous cause and do not represent any change in market fundamentals, which are kept constant, nor any change in agents preferences.

In fact, upon believing that assets markets as able to select the equilibrium entirely reflecting the available information, from the observation of Fig. 6 one would wrongly conclude that the asset dividend process is non-stationary. It is true, however, that the long-run average price seem 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, the market does seem to deliver a not too incorrect average

information in this case.

5.3 Some special rules

There exists a very specific price dependent rule: 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 an unique market equilibrium, in the sense that it is consistent with all prices in $(0, 1)$. For this reason, no matter the realized state of nature, it has constant and unitary return, and thus constant wealth, as can be easily derived from (3.2). Thus the market portfolio rule is special because, no matter the rule used by competing agents, 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, as can be seen, for example, in the right panel of Fig. 3. Thus agents with a high coefficient of relative risk aversion have high chances to survive in the long run. On the contrary, agents with a low coefficient of relative risk aversion tend to invest all the wealth in the under-priced asset, see for example the left 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.⁶

Also among the price dependent rules, given the simple market structure considered in this Section, the rule with the highest chance to dominate the market continues to be the Kelly rule presented in Section 4. Indeed it is still the case that the Kelly rule destabilizes all other fixed points and the fixed point it determines is never unstable. As a result, even if the Kelly rule may not be the unique surviving rule, when it is present in the market prices may not converge to level different than the informationally efficient $(\pi, 1 - \pi)$. Then one part of the good message that holds with constant rules is still true: once the Kelly rule is present in the market, it makes it informational efficient.

Notice however that the Kelly rule is not only characterized by the correctness of its beliefs but also from the precise structure of its preferences, corresponding to a coefficient of relative risk aversion of one. As the previous examples show, the same beliefs with different preferences may not grant informational efficiency.

6 Rule ordering

In the previous section we have 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 would allow to have absolute judgments about their performance. 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 not only imply the possibility to look for the best, and the least, performing rules. It also allows for an ex-ante judgment about the asymptotic state of the market: given an ecology composed of a collection of different trading rules, irrespectively of

⁶Even more in danger are risk neutral, or risk lover, agents who, by investing all their wealth in the first (second) asset if its price is lower (higher) than π , would disappear from the market in finite time.




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

their number, it is only the best rules in the collection that will survive and set the prices in the long run.

In markets with constant rules, as those analyzed in Sec. 4 this ordering is possible. Using the relative entropy $I_\pi(\alpha)$ one can define a simple relation by postulating that rule α^+ is better than rule α^- when its relative entropy is lower or

$$\alpha^+ \succ \alpha^- \quad \text{if and only if} \quad I_\pi(\alpha^+) < I_\pi(\alpha^-) \quad (6.1)$$

It is immediate to see that this relation is anti-reflexive and anti-symmetric. It is also immediate to see that the relative entropy is a concave function of its argument.⁷ Then if $I_\pi(\alpha^+) < I_\pi(\alpha^0)$ and $I_\pi(\alpha^0) < I_\pi(\alpha^-)$ one has, for any $\lambda \in [0, 1]$,

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

and in particular $I_\pi(\alpha^+) < I_\pi(\alpha^-)$, so that the relation defined in (6.1) is transitive and, hence, an 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. This implies that when one observes a market dynamics, assuming that what he is observing is the long run, one can be sure that worst rules have been wiped out by the selection process. The long-run price is determined only by the best rules, as all the resources have been allocated to them.

Unfortunately, the possibility to build an order relation is lost when price dependent rules are taken into account. We show it by building a counterexample. Fix $\pi = 1/2$ and consider the following three CRRA portfolio rules (also plotted in Fig. 6): rule one, $\alpha^1(p)$, has $\pi^{e,1} = 0.25$ and $\gamma^1 = 2$; rule two, $\alpha^2(p)$, has $\pi^{e,1} = 0.3$ and $\gamma^1 = 1$; and rule three, $\alpha^3(p)$, has $\pi^{e,1} = 0.65$ and $\gamma^1 = 1$.

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

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

⁷It directly follows for the convexity of the log function.




Figure 8: Market with $\pi = 1/2$ and the three rules given in the text. Left panel: wealth ratio dynamics. Right panel: prices dynamics.

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 $I_\pi(\alpha^2(p)) < I_\pi(\alpha^1(p))$ for all $p \in [\pi^{e,1}, \pi^{e,2}]$, agent 2 is going to dominate 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, in that it rewards the agent with the “best” information.

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 in that it rewards the agent with the “best” information.

The transitive property, a necessary property for \succ to be a strict order relation, would now imply $\alpha^3 \succ \alpha^1$. Is it true? 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 $I_\pi(\alpha^1)$ have not 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. 6. It holds $I_\pi(\alpha^3) < I_\pi(\alpha^1)$ when rule 1 is close to have all wealth and the price of the first asset is close to $\pi^{e,1}$. 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 rather holds $\alpha^3 \sim \alpha^1$. Despite rule 3 conveys the best information it fails to dominate and prices keep fluctuating between $\pi^{e,1}$ and $\pi^{e,3}$.

It is also interesting to analyze market dynamics when all the 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 all fixed points are locally unstable, as can immediately be judged by comparing all the relative distance to the Kelly rule π in the EMC plot of Fig. 6, wealth fractions, and thus prices, keep fluctuating (see also the right panel of Fig. 6). Again the market cannot be said informational efficient: despite rule α^3 has the best information the market fails to set the prices according to its information.

7 Conclusion and Outlook

We have discussed wealth-driven selection in a simple market for Arrow securities using a series of simple examples. When rules do not depend on prices, which occurs, for instance, when agents maximize a CRRA expected utility with unitary risk aversion coefficient, the rule relying on the most accurate knowledge about the unfolding of states of Nature, as measured by the relative entropy, attracts all wealth in the long run, thus driving prices as close as possible to fundamentals. However, the picture changes when price dependent rules, employed e.g. by agents whose coefficient of relative risk aversion is different from one, are considered. In fact, 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 global best rule but might get stuck in a local optimum or, even, keep fluctuating in search of an optimum that doesn't exist. We have provided examples of both occurrences. Concerning the former, we have shown that when two locally stable fixed points with a single dominating agent exist, wealth selection may drive the market towards one or another depending on the initial condition and/or the specific realization of states of Nature. For the latter, we have shown that in a market where two rules are trading, if one rule is closest to the best rule at the prices determined by the second and vice-versa, then wealth shares and prices keep fluctuating indefinitely. In the final part we have shown that the local nature of wealth-driven selection implies the impossibility to order strategies according to their market performance. The general message is that one cannot be sure of the informational content of prices as it is not granted that prices reflect the best available information. Instead, prices may display endogenous fluctuations or, even, converge to the worst information revealing level.

How should the reader take these results? Despite in our examples the market may not reward the best informed agent, there is a rule that is rewarded by the market for every price, the so called Kelly rule, which amounts to investing according to a CRRA utility with unitary risk aversion and using the correct probabilities. The presence of the Kelly rule in a market makes it informational efficiency. At the same time, however, it implies the existence of at least one agent with specific preferences and perfect knowledge about the process governing the state of the world. Assume this information is not readily available, so that the Kelly rule is not present. The relevant question then becomes if, in this case, the Kelly rule is learnable. In this market context the answer very much depends on the type of learning considered, that is, whether agents learn probabilities, future prices, or both. So far learning have being considered as updating the subjective probabilities π^e , but only in a context of constant rules (see e.g. Evstigneev et al., 2009; Blume and Easley, 2010), or using past prices to forecast future prices (see e.g. Bottazzi and Dindo, 2010). In both cases conditions are given to the learning process in order to converge to the Kelly rule (fast enough on dividends not to be wrong for too much time, slow enough for prices not to create a form of deterministic overshooting). An investigation of the learning process in both dimensions and for general rules is however still missing. The issue of learning is crucial if one wants to drop the assumption of rational expectations but keep the inter-temporal optimization in place. In fact, in the Appendix we show that, even restricting the analysis to CRRA expected utility, rules depends on conditional probabilities and price expectations. Which class of probability updating and expectations formation grants survival, market efficiency, or convergence to the Kelly rule are all question still waiting for an answer.

There are other open issues. Firstly, one would like to consider also rule that depends on wealth. This has both a behavioral relevance, as it is natural to assume that an agent observing a constant reduction in the personal wealth level looks for and implement different

trading strategies, and a positive appeal, as it allows to study market selection for any given risk preference. Further work is also needed to characterize the global market dynamics. Here we have relied exclusively on local results, without characterizing the overall dynamics. This characterization is important in that it would provide the complete picture of possible market dynamics.

Another direction of investigation is the extension of the present analysis to markets for long-lived assets. Whereas with constant rules the story doesn't change (see e.g. Evstigneev et al., 2009) it is not known what happens when rules are price dependent. Would the shape of the Kelly rule, who sets the reference rule, be affected? As a consequence, what can be said about wealth-driven market selection in this case? Last, one would like to introduce production into the picture. After all, even to stick with finance application, the role of capital markets is to finance productive activities so that asset performances should influence firms performance in much the same way as the other way round. Intuitively adding this layer would amount to possible shifts in the Kelly rule, due e.g. to possible shifts in the likelihood to get a certain payoff or in their size. How the overall picture would change not having such an anchor point as the Kelly rule, is yet another question to be answered.

A CRRA portfolio rules

Here we derive the asset demand which maximizes the expected utility of final wealth, at period T , when risk preferences are described by the generic CRRA utility in (3.4). As the general answer depends on agents future prices expectations, and we do not assume perfect foresight, we will have to make a simplifying assumption in order to be able to get an answer. It will turn out that such assumption is not that strict given the local nature of the stability analysis exploited in the paper.

Given a sequence ω ruling the occurrence of states of Nature, denote as ω^t the partial history till time t , that is, $\omega^t = (\omega_1, \dots, \omega_t) \in \Omega^t$. The probability that such an event occurs is denoted by $\pi(\omega^t)$, which, having a Bernoulli process, is obviously equal to the product $\prod_{s=0}^t \pi_{\omega_s}$. Denote as $\pi^{e,i}(\omega^t)$ the belief agent i holds on the occurrence of the partial history ω^t .

In each period t , an agent chooses, given partial history ω^t , her asset demand as to maximize the utility of her final wealth, given all the intertemporal budget constraints. Following our previous notation it is convenient to consider as the choice variable the fraction of current wealth invested in the two assets at time t given history ω^t , $\alpha_{k,t}(\omega^t)$ $k = 1, 2$. As a result one can express budget constraints as

$$\alpha_{1,t}(\omega^t) + \alpha_{2,t}(\omega^t) \leq 1 \quad \text{for all } \omega^t \text{ and for all } t \leq T - 1. \quad (\text{A.1})$$

Each agent chooses the fractions $\alpha_{k,t}(\omega^t)$ in order to maximize the utility of her final wealth w_T , or

$$\text{Max} \sum_{\omega^T \in \Omega^T} \pi^e(\omega^T) U(w_T(\omega^T)).$$

Given our assets structure, the final wealth is just the number of assets bought at time $T - 1$ of type ω_T , or

$$w_T(\omega^T) = \frac{\alpha_{\omega_T, T-1}(\omega^{T-1})}{p_{\omega_T, T-1}(\omega^{T-1})} w_{T-1},$$

which, by recursion, gives

$$w_T(\omega^T) = \left(\prod_{t=1}^T \frac{\alpha_{\omega_t, t-1}(\omega^{t-1})}{p_{\omega_t, t-1}(\omega^{t-1})} \right) w_0.$$

For each history ω^T the final wealth is a function of a stream of vector choices $\alpha_t(\omega^t)$ satisfying (A.1). We assume an utility of CRRA type as in (3.4) and consider the case $\gamma = 1$ first. Using the property of the logarithmic function our problem becomes

$$\text{Max} \sum_{\omega^T \in \Omega^T} \pi^e(\omega^T) \sum_{t=1}^T \log(\alpha_{\omega_t, t-1}(\omega^{t-1}))$$

which, rearranging the terms and denoting as $\pi^e(\omega_t; \omega^{t-1})$ the probability that event ω_t occurs given that ω^{t-1} has been observed, is equivalent to

$$\text{Max} \sum_{t=1}^T \sum_{\omega^{t-1} \in \Omega^{t-1}} \pi^e(\omega_t; \omega^{t-1}) \log(\alpha_{\omega_t, t-1}(\omega^{t-1})).$$

Imposing the first order conditions with respect to the generic variable $\alpha_{k,s}(\omega^s)$, while taking into account the constraint (A.1), gives

$$\frac{\pi^e(k; \omega^s)}{\alpha_{k,s}(\omega^s)} = \lambda(\omega^s),$$

where $\lambda(\omega^s)$ is the multiplier associated with the binding constraint for the event ω^s . Since the equation holds for both $k = 1, 2$ with the same multiplier, we can find

$$\alpha_{k,s}(\omega^s) = \pi^e(k; \omega^s) \quad k = 1, 2.$$

The fraction to be invested in asset k at time s given history ω^s is equal to the belief that the event k will occur at time $s + 1$ given that ω^s has been observed.

We turn now to a general CRRA function, that is, (3.4) with $\gamma \neq 1$. The maximization problem becomes

$$\text{Max} \sum_{\omega^T \in \Omega^T} \frac{\pi^e(\omega^T)}{1 - \gamma} \left(\prod_{t=1}^T \frac{\alpha_{\omega_t, t-1}(\omega^{t-1})}{p_{\omega_t, t-1}(\omega^{t-1})} w_0 \right)^{1-\gamma}$$

where choices are constrained as in (A.1). If we consider the variable $\alpha_{k,s}(\omega^s)$, it appears only in sequences ω^T whose partial history till time s is ω^s , name such a sequence $\omega_{|\omega^s}^T$. The first order condition with respect $\alpha_{k,s}(\omega^s)$ then gives

$$\left[\sum_{\omega_{|\omega^s, k}^T} \pi^e(\omega^T) \left(\prod_{t=s+2}^T \frac{\alpha_{\omega_t, t-1}(\omega^{t-1})}{p_{\omega_t, t-1}(\omega^{t-1})} \right)^{1-\gamma} \right] \frac{(\alpha_{k,s}(\omega^s))^{-\gamma}}{(p_{k,s}(\omega^s))^{1-\gamma}} \left(\prod_{t=1}^s \frac{\alpha_{\omega_t, t-1}(\omega^{t-1})}{p_{\omega_t, t-1}(\omega^{t-1})} \right)^{1-\gamma} w_0^{1-\gamma} = \lambda(\omega^s)$$

which, when $s = T - 1$, reduces to

$$\pi^e((\omega^{T-1}, k)) \frac{1}{\alpha_{k, T-1}(\omega^{T-1})} \left(\frac{\alpha_{k, T-1}(\omega^{T-1})}{p_{k, T-1}(\omega^{T-1})} \right)^{1-\gamma} \left(\prod_{t=1}^{T-1} \frac{\alpha_{\omega_t, t-1}(\omega^{t-1})}{p_{\omega_t, t-1}(\omega^{t-1})} \right)^{1-\gamma} w_0^{1-\gamma} = \lambda(\omega^{T-1}).$$

Equating the expressions for $k = 1$ and $k = 2$ we find

$$\frac{\pi^e((\omega^{T-1}, 1))}{\alpha_{1, T-1}(\omega^{T-1})} \left(\frac{\alpha_{1, T-1}(\omega^{T-1})}{p_{1, T-1}(\omega^{T-1})} \right)^{1-\gamma} = \frac{\pi^e((\omega^{T-1}, 2))}{\alpha_{2, T-1}(\omega^{T-1})} \left(\frac{\alpha_{2, T-1}(\omega^{T-1})}{p_{2, T-1}(\omega^{T-1})} \right)^{1-\gamma},$$

which, using the constraint (A.1) and the fact that $\pi^e((\omega^{T-1}, k)) = \pi^e(k; \omega^{T-1})\pi^e(\omega^{T-1})$, can be solved to give

$$\alpha_{1, T-1}(\omega^{T-1}) = \frac{(\pi^e(1; \omega^{T-1})(p_{2, T-1}(\omega^{T-1}))^{1-\gamma})^{\frac{1}{\gamma}}}{(\pi^e(1; \omega^{T-1})(p_{2, T-1}(\omega^{T-1}))^{1-\gamma})^{\frac{1}{\gamma}} + (\pi^e(2; \omega^{T-1})(p_{1, T-1}(\omega^{T-1}))^{1-\gamma})^{\frac{1}{\gamma}}}. \quad (\text{A.2})$$

Having found the maximizing value of portfolio choices at time $T - 1$ we can concentrate on the problem at time $T - 2$. Taking the derivative with respect to $\alpha_{k, T-2}(\omega^{T-2})$ and using the first order conditions for $\alpha_{k, T-1}(\omega^{T-1})$ leads to the following equation

$$\lambda((\omega^{T-2}, k))\alpha_{1, T-1}((\omega^{T-2}, k)) + \lambda((\omega^{T-2}, k))\alpha_{2, T-1}((\omega^{T-2}, k)) = \lambda(\omega^{T-2})\alpha_{k, T-2}(\omega^{T-2})$$

The fact that choices are normalized allows to rewrite the previous equation first as

$$\alpha_{k,T-2}(\omega^{T-2}) = \frac{\lambda((\omega^{T-2}, k))}{\lambda(\omega^{T-2})}$$

and finally as

$$\alpha_{k,T-2}(\omega^{T-2}) = \frac{\lambda((\omega^{T-2}, k))}{\lambda((\omega^{T-2}, 1)) + \lambda((\omega^{T-2}, 2))}. \quad (\text{A.3})$$

Since both multipliers $\lambda((\omega^{T-2}, k))$ $k = 1, 2$ depend upon $\alpha_{k,T-2}(\omega^{T-2})$, (A.3) is still an implicit equation in $\alpha_{k,T-2}(\omega^{T-2})$. The solution is a function of current prices and probabilities, as dependent on the observed history ω^{T-2} , and prices and probabilities of period $T - 1$, as dependent on histories $\omega^{T-1} = (\omega^{T-2}, k)$ $k = 1, 2$.

Although it is not particularly difficult to obtain the explicit solution for period $T - 2$ demands, iterating back to period 0 is instead rather complicated. A part from that, it gives asset demand that depend on future price expectations, which are outside the scope of this contribution. The overall procedure is heavily simplified if period $T - 1$ portfolio choices do not depend on history ω^{T-1} . A necessary condition is that both time $T - 1$ estimates of the probabilities and prices do not depend on the history. This holds in our market if agents do not update their beliefs and the market is close to one of its long-run equilibria, where prices are indeed no history dependent.⁸

If prices and choices at time $T - 1$ do not depend on history, (A.3) can be written as

$$\alpha_{1,T-2}(\omega^{T-2}) = \frac{\pi^e(1; \omega^{T-2}) \left(\frac{\alpha_{1,T-2}(\omega^{T-2})}{p_{1,T-2}(\omega^{T-2})} \right)^{1-\gamma}}{\pi^e(1; \omega^{T-2}) \left(\frac{\alpha_{1,T-2}(\omega^{T-2})}{p_{1,T-2}(\omega^{T-2})} \right)^{1-\gamma} + \pi^e(1; \omega^{T-2}) \left(\frac{\alpha_{2,T-2}(\omega^{T-2})}{p_{2,T-2}(\omega^{T-2})} \right)^{1-\gamma}},$$

which can be solved to give

$$\alpha_{1,T-2}(\omega^{T-2}) = \frac{(\pi^e(1; \omega^{T-2})(p_{2,T-2}(\omega^{T-2}))^{1-\gamma})^{\frac{1}{\gamma}}}{(\pi^e(1; \omega^{T-2})(p_{2,T-2}(\omega^{T-2}))^{1-\gamma})^{\frac{1}{\gamma}} + (\pi^e(2; \omega^{T-2})(p_{1,T-2}(\omega^{T-2}))^{1-\gamma})^{\frac{1}{\gamma}}}.$$

Not surprisingly we have found the same expression as in (A.2). Since the same reasoning can be applied at all periods, denoting $\pi^e(1) = \pi^e$, and using that prices are normalized, we find the optimal fraction to be invested in the first asset as the generic period t , $\alpha_{1,t}(p)$, as a function of the parameters π^e and γ :

$$\alpha_{1,t}(p; \pi^e; \gamma) = \frac{(\pi^e(1-p)^{1-\gamma})^{\frac{1}{\gamma}}}{(\pi^e(1-p)^{1-\gamma})^{\frac{1}{\gamma}} + ((1-\pi^e)(p)^{1-\gamma})^{\frac{1}{\gamma}}}. \quad (\text{A.4})$$

The expression above is exactly the same demand used throughout the paper. Thus the results discussed can be generalized, under the assumption discussed above, to long-horizon maximizing CRRA agents.

⁸Notice however that, even if an agent believes that the underlying process is independent, imposing that probability estimates do not depend on the history implies that the agent is not able to learn about the exogenous process.

References

- Amir, R., I. Evstigneev, T. Hens, and K. Schenk-Hoppé (2005). Market selection and survival of investment strategies. *Journal of Mathematical Economics* 41, 105–122.
- Anufriev, M. and G. Bottazzi (2010). Market equilibria under procedural rationality. *Journal of Mathematical Economics* 46, 1140–1172.
- Anufriev, M., G. Bottazzi, and F. Pancotto (2006). Equilibria, stability and asymptotic dominance in a speculative market with heterogeneous agents. *Journal of Economic Dynamics and Control* 30, 1787 – 1835.
- Anufriev, M. and P. Dindo (2010). Wealth-driven selection in a financial market with heterogeneous agents. *Journal of Economic Behavior and Organization* 73, 327–358.
- Blume, L. and D. Easley (1992). Evolution and market behavior. *The Journal of Economic Theory* 58, 9–40.
- Blume, L. and D. Easley (2006). If you are so smart why aren't you rich? Belief selection in complete and incomplete markets. *Econometrica* 74, 929–966.
- Blume, L. and D. Easley (2010). Heterogeneity, selection, and wealth dynamics. *Annual Review of Economics* 2, 425–450.
- Bottazzi, G. and P. Dindo (2010). Evolution and market behavior with endogenous investment rules. LEM Working Paper 2010-20, Scuola Superiore Sant'Anna, Pisa.
- Evstigneev, I., T. Hens, and K. Schenk-Hoppé (2008). Globally evolutionary stable portfolio rules. *Journal of Economic Theory* 140, 197–228.
- Evstigneev, I., T. Hens, and K. Schenk-Hoppé (2009). Evolutionary finance. In T. Hens and K. Schenk-Hoppé (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*. North-Holland (Handbooks in Economics Series).
- Kelly, J. (1956). A new interpretation of information rates. *Bell System Technical Journal* 35, 917–926.
- 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. *Journal of Mathematical Economics* 41, 91–104.