Endogenous fluctuations in an evolutionary learning model with adverse selection

Clemens Buchen[†] and Alberto Palermo[‡]

[†] *EBS Business School, Wiesbaden*

[‡] *EBS Business School, Wiesbaden and Panthéon-Assas University, Paris*

Preliminary version, please don't cite or disseminate.

Abstract

We set-up an evolutionary learning model including information asymmetry. Principals who face an adverse selection problem are heterogeneous in that they can have different beliefs about the distribution of agents' types in the population. A fraction of principals starts out with unbiased beliefs and a fraction of principals starts out with biased beliefs. In addition, principals form a belief about the total quantity in the market. Periodically, they are allowed to observe profits of others with the same agent match. Then the probability of changing beliefs depends on the prevalence of the belief and the size of the difference in payoffs. This type of imitation gives rise to a version to the replicator dynamic. The resulting nonlinear dynamical system is studied, showing the existence of multiple equilibria involving both monomorphic and bimorphic populations. The local and global bifurcation analysis is performed to show the existence of endogenous fluctuations driven by no changes in the fundamentals. The global behaviour is non-trivial and it involves cycling patterns and chaotic behaviour.

JEL Classification: C73, D82, D83, E32,

Keywords: Evolutionary game theory, heterogeneous beliefs, adverse selection, bifurcation theory

1 Introduction

Recently, a large and growing interest has been devoted to the role of evolutionary mechanisms in economics. Since the pioneering work by Alchian (1950) several important theoretical contributions have been developed. Models in this tradition can be roughly grouped along two lines. On the one hand, some of them come directly from evolutionary game theory and its attempt to study social interactions when the common assumption of rationality is relaxed. An introduction and overview can be found in Sandholm (2011), Weibull (1995), or Gintis (2009). On the other hand, a large class of learning models have been developed with the aim of understanding what drives economic decisions depending on different degrees of sophistication of players (Fudenberg and Levine, 1998; Young, 2004). These strands are interrelated in numerous ways, as evolutionary models are used to model learning processes. Our paper follows the latter tradition.

In a simple set-up, the aim of our paper is to introduce the role of evolutionary learning through imitation in a market, in which production is the result of a principal-agent relationship with adverse selection and some principals' beliefs about the distribution of agent types is biased.

Usually in mechanism design models players have a subjective probability distribution over a set of possible elements or outcomes, which represents information privately known to other players. More specifically, in a principal-agent relationship the principal does not know the type of agent that she is matched with, but the distribution of types is common knowledge. In the basic model with adverse selection the principal knows that she faces an agent with a given ability drawn from a distribution of types. Agents can be efficient meaning they have a comparatively small marginal cost of production, or they can be inefficient with a high marginal cost. On the basis of the distribution of efficient and inefficient agents, principals write court-enforcable contracts and agents self-select.

Against this backdrop of the standard model we introduce a limitation on the part of the principals, which is based on the observation that economic decisions are the result of interactions where the context may not be perfectly forecast and understood by the players. We assume that a population of principals is heterogeneous in that they have different beliefs about the distribution of types. Hence, there exists a population of agents with different abilities, and a population of principals having different beliefs about the distribution of these abilities.

To fix ideas, we assume that the population of principals is characterized such that a fraction of them is optimistic about the distribution. This means that they believe that the proportion, and not the abilities[1], of efficient types is larger than it really is. Both types of principals offer contracts to different ability types of agents, but the contracts of the optimistic principals will differ from the regular contract. Then, the profit of principals are different depending on the belief a principal holds. Hence, we study a polymorphic population characterized by unbiased and optimistic principals. While the role of possibly biased agents and optimal contracts has been studied (see e.g. Santos-Pinto (2010) and the bibliography therein), to the best of our knowledge, our paper is the first to introduce the issue assuming that the uninformed side has a bias.

Aside from a belief about the distribution of types any principal in the model must also form

[1]A model where principals are biased concerning the abilities and not the distribution as in our model would lead to different conclusions.

an expectation about the total quantity in the market. Since we want principals to be less than perfectly informed about the market structure in terms of the beliefs of other principals, it would be detrimental to our efforts if principals had rational expectations about the total quantity in the market. Instead, we use a simple naive expectation in that when writing contracts principals expect the total quantity to be the same as in the previous period. This implies that there are two variables that change over time, the fraction of principals with the correct belief about the distribution and the total quantity in the market, which is a function of the distribution of principals.

First we show that principals with optimistic beliefs earn a higher profit if matched with an efficient agent. On the other hand, it is not clear which type earns a higher profit if matched with the inefficient agent, as it depends on the quantity in the market. If the increase of the quantity in the market is too sudden from one period to the next, it can be that the optimistic principal is again better off. For small changes or any decrease of the quantity the principal with the unbiased belief earns a higher profit when matched with an inefficient agent. The latter result is connected to the production quantities that the contracts stipulate. In general, optimistic principals decrease the quantity for both types of agents and as a result efficient agents matched with them earn a smaller rent. As the marginal cost of production of the agent increases, i.e. the agent becomes inefficient, the marginal benefit from a smaller rent is outweighed by the marginal benefit of a higher quantity, and then unbiased principals earn a higher profit with the inefficient agent.

Drawing the usual analogy to the standard game theoretic terminology, the different beliefs that the principals hold about the distribution of types can be interpreted as the*strategies*, while the *payoffs* are simply the monetary gains from the principal-agent relationship. In general one would expect that over time, in a repeated interaction, principals can learn from the strategies played by other principals. We model this in the following way. At the beginning of a period, principals write a contract based on their beliefs. Then they are matched with an agent, so that at the end of a period principals earn a profit from the production of the agent. Some principals are then randomly drawn to be able to observe the strategy (belief) and payoff of a different randomly drawn principal. Then we define a probability of switching to the other principal's belief. The probability of switching is always zero if the other principal either has the same belief (obviously) or if she was matched with a different type of agent. This means that we restrict positive probabilities of switching beliefs to those cases where the other principal comes from the defined reference groups (Selten and Ostmann, 2001). In our game the reference group comprises all principles, who were matched with the same type of agent. Only in this case the profit of a fellow principal is comparable enough such that a principal might consider switching beliefs.

The natural question in our evolutionary set-up is: where does the economy converge in terms of different beliefs? If we assumed rational players able to revise beliefs as Bayesian updaters with an appeal to the law of large number the answer would be straightforward. As more information accumulates over time with perfect memory Bayesian updating would ultimately lead to all principles playing the same, unbiased strategy. However, in many applications Bayesian updating makes too great demands on the the quality of the information set and cognitive abilities of agents (Young, 2004).

Given this discussion, the contributions of the paper are aimed to be twofold. On the one hand, we wish to make a contribution to a literature of learning by studying the evolution of contracts in a situation where information can be outright wrong. The existence and stability of equilibria are crucial in this respect. A related research program is developed by Young (1998). He develops the notion of adaptive play where individuals act on a restricted sample of plays. Most importantly, Young makes the distinction between asymptotic stability and stochastic stability of an equilibrium, where in the latter case an equilibrium survives a series of subsequent disruptions. Disruptions of an equilibrium come from mistakes individuals make. Our approach is different, because we focus on finite time horizons. It is novel in the sense that we explicitly allow for ignorance of a subset of individuals (in our case principals), which can vanish from the population or even proliferate or survive. Once the model is set up we can examine conditions for fast or slow dissipation of information through the economy solely on the basis of imitation of better practice.

Our approach is different and somewhat complementary to other attempts to model adjustment processes in principal-agent settings with adverse selection. Arifovic and Karaivanov (2010) start from an adverse selection model evolving over time as we do assuming that principals are unable to solve the correct maximization problem. Principals start with a randomly assigned choice of contracts and are allowed to update their strategy by observing randomly played strategies by other players. In a simulation the authors then test under which conditions principals can work towards the "correct" contracts by replication of other strategies. In contrast to that we model principals that are able to solve the maximization problem, but some of the principals have the wrong fundamental expectation about the true distribution of types of agents. So rather than model a cognitive limitation with respect to the maximization problem we model a limitation that comes from a lack of information or a misguided type of optimism. Our model is in the tradition of a literature on imitation in different market settings. For example, Vega-Redondo (1997) studies equilibria in a Cournot game, where players can imitate other players. The model is extended and applied to other market types in a number of other papers (see citations in Ania et al. (2002) andAlós-Ferrer (2004)). Moreover, our approach is similar in spirit to the information sharing literature (Vives, 2001). There, principals can play a strategy to disclose information learned about agent's cost functions, because they were matched in an employment relationship. Of course, the difference in our model is that any information sharing occurs unwittingly and not as a strategic play.

In addition, our model can be understood as a stable cobweb model, where the source of the heterogeneity of firms is different than what is usually assumed in the literature. As will become clear as the model progresses our model can be easily written as a model where firms have different expectations about the price in the next period. The seminal paper on this is Brock and Hommes (1997). There, the heterogeneity comes from the fact that some firms invest more to better predict expected quantity than others. Since the incentive to invest in prediction changes as the fluctuations of the price change, chaotic and unpredictable behaviour of time series and bifurcations of variables can be shown. Our model gives an alternative explanation of heterogeneity, which comes from different beliefs about the ability level in the population of agents, holding the prediction ability of firms equal for all. We are therefore able to explain an unpredictable time series of the price (or the quantity) with a micro foundation that is the source of different production levels.

The paper proceeds as follows: the next section introduces the model in two steps. First the stage game is described, which derives the principals' payoffs based on the agents' decisions. Then, the population game is introduced. This entails defining an mechanism that allows principals to modify their beliefs and therefore change their contract offers in the stage game. This gives a dynamical system, which we extensively study in section 3. Finally we discuss our results in chapter 4 and give some concluding remarks.

2 The Model

We consider two large but finite populations of principals and agents with an equal size. Making the set-up general, we refer to the usual assumptions of the general principal-agent framework, which for completeness we briefly go through in their standard notation (see e.g. Laffont and Martimort, 2002). Each principal wants to delegate a task to an agent in order to produce a quantity q. Agents are heterogeneous with regard to their ability to produce the quantity in that some are more efficient than others. Agents have a linear cost function defined as $C(q, \theta) = \theta q$. As it is standard, we assume that $\theta \in \left\{\underline{\theta}, \overline{\theta}\right\}$, where $\overline{\theta} > \underline{\theta}$. The proportion of efficient agents with marginal cost $\underline{\theta}$ is $v \in (0, 1)$, which implies that the proportion of inefficient agents with marginal cost $\underline{\theta}$ is $1 - v$. In the next step the stage game will be defined.

2.1 Stage game

Principals are heterogeneous in that they hold different beliefs about the distribution of agents' abilities. We assume that each principal can have a belief $\phi \in \{\rho, v\}$ with $\rho > \nu$. In other words, some principals believe that the proportion of efficient agents in the population is larger than it really is. We will sometimes refer to those biased principals as overly optimistic.

The agent's production provides a benefit to a principal i, which is measured by a function $S(q_t^i, \tilde{q}_t)$, where q_t^i is the quantity produced by the agent working for principal i in period t and $\tilde{q}_t$ a sufficient statistic of the total quantity in the market. The precise definition of $\tilde{q}_t$ will be given below. The presence of the term $\tilde{q}_t$ implies that there is a interdependency between the quantities produced by principals.

Timing

Since most of economic activities take place in discrete time intervals, we will use this assumption to set up the model. At the beginning of a generic period t, the fraction of principals who write contracts on the basis of belief v is denoted by α_t and then the fraction of principals using ρ is given by $1 - \alpha_t$. Principals write contracts according to their beliefs $\{v, \rho\}$ and on the basis of a belief about the benefit they will obtain at the end of the period. In what follows, with a slight abuse of notation, we indicate realized variables at the end of a generic period with $t + 1$, which will determine the variables for the beginning of the next period. Next we give a functional form to the benefit function. We use $S_{t+1}\left[q_{t+1}, \mathbb{E}_t\left(\tilde{q}_{t+1}\right)\right] = \beta q_{t+1} - \frac{(q_{t+1})^2}{2} + \delta q_{t+1} \mathbb{E}_t\left(\tilde{q}_{t+1}\right)$ as the benefit a principal expects to gain at the end of the period. The functional form is often used in information sharing models and is sure to give a closed-form solution to the adverse selection problem[2]. We assume β to be a positive constant and $\delta \in (-1, 0)$, which is a measure of the degree of substitutability between principals' outputs. The variable $\mathbb{E}_t(\tilde{q}_{t+1})$ denotes the expectation, which a principal forms in the beginning of a period about the value of $\tilde{q}$ at the end of the period, when all production is carried out. We make the following assumption about the expectation.

Assumption 1. *In each period t, each principal has a naive expectation about the sufficient statistic of the quantity in the market: $\mathbb{E}_t(\tilde{q}_{t+1}) = \tilde{q}_t$.*

The rationale for this assumption is the following. Principals in our model do not know the salient characteristics of the market, simply because they are not aware of the fact that there are different beliefs present, which have an impact on the quantity the principals produce. Rational expectations about the quantity would run contrary to this view of the role of principals and therefore are not useful in this respect. The simplest version of adaptive expectations are naive expectations as in our assumption, which take only one preceding period into account (Evans and Honkapohja, 2001). The timing for each period can be summarized as follows:

1. Each agent realizes his type.
2. Principals write contracts according to their beliefs about the distribution of types and according to naive expectations about the quantity $\tilde{q}_{t+1}$.
3. Each agent is matched with a principal and accepts the contract. There is no unemployment.
4. Contracts are executed: quantities are produced, profits and payments to agents are realized.

Contracts

At the beginning of each period, principals write contracts which entail a rent for each quantity observed at the end of the period. Given the possibly different beliefs, the quantities produced are indicated by $\underline{q}^{\phi}$ for the efficient type and $\overline{q}^{\phi}$ for the inefficient, where $\phi \in \{v, \rho\}$. The individual quantities are also a function of the expected quantity $\tilde{q}$ in the market. Accordingly, we denote the rent for the efficient agent with $\underline{U}^{\phi}$ and for the inefficient type with $\overline{U}^{\phi}$, respectively. When the contracts are executed, the realized quantity implies a profit for each principal, which then depends only on her quantity and the aggregate $\tilde{q}$.

We restrict our analysis to direct revelation mechanisms that are truthful. This can be done since the agent's rent is only a function of his principal's contract and it does not depend on the total quantity in the market. Moreover, we assume that the matching is type-independent avoiding therefore any sort of competition among principals[3].

[2]See Vives (2001) and for a recent application Piccolo and Pagnozzi (2013).

[3]In absence of this assumption, the revelation principle might be invalid, see e.g. Pavan and Calzolari (2010). It might be argued that this makes the agents's role somewhat overly passive, because in our model agents are not able to realize that there are different rents to be earned dependent on the principal. Then agents would have a clear preference of a match, which would make the matching type-dependent. One could introduce congestion effects of matching to ensure that all types of principals still receive matches as in Guerrieri et al. (2010). Since we focus entirely on the dynamic of the belief of agents and the quantities, we refrain from this in order to keep the model as simple as possible.

Summarizing in a more formal way, in a generic time t, each principal writes contracts according to $\{v, \rho\}$ and the expectation of $\tilde{q}$. While the reference period is t, we specify the realized payoffs and the realized quantity with an index $t+1$. The reason for this comes directly from our assumption about the timing. While the contracts are stipulated in t, only when the principal has observed the realized quantity by the agent the contract is honored. Hence, in t each principal defines a mechanism $\langle q_{t+1}(\theta), w_{t+1}(\theta)\rangle$ which entails a transfer w_{t+1} for each observed quantity in $t+1$. Formally, each principal maximizes expected profits given the usual incentive (IC) and participation constraints (PC):

$$\underset{q_{t+1}, w_{t+1}}{Max}\ \phi\left\{S\left[\underline{q}_{t+1}, \mathbb{E}_t\left(\tilde{q}_{t+1}\right)\right] - \underline{w}_{t+1}\right\} + (1-\phi)\left\{S\left[\overline{q}_{t+1}, \mathbb{E}_t\left(\tilde{q}_{t+1}\right)\right] - \overline{w}_{t+1}\right\} \tag{1}$$

s.t.

$$U\left(q_{t+1}(\theta), w_{t+1}(\theta), \theta\right) = w_{t+1}(\theta) - \theta q_{t+1}(\theta) \geqslant 0 \qquad \forall\, \theta \in \left\{\underline{\theta}, \overline{\theta}\right\} \qquad (ICs)$$

$$U\left(q_{t+1}(\underline{\theta}), w_{t+1}(\underline{\theta}), \underline{\theta}\right) \geqslant U\left(q_{t+1}\left(\overline{\theta}\right), w_{t+1}\left(\overline{\theta}\right), \underline{\theta}\right) \qquad \left(PC\left(\underline{\theta}\right)\right)$$

$$U\left(q_{t+1}\left(\overline{\theta}\right), w_{t+1}\left(\overline{\theta}\right), \overline{\theta}\right) \geqslant U\left(q_{t+1}(\underline{\theta}), w_{t+1}(\underline{\theta}), \underline{\theta}\right) \qquad \left(PC\left(\overline{\theta}\right)\right)$$

Given this set-up the following proposition is straightforward.

Proposition 1. *Given different beliefs and the same naive expectations about $\tilde{q}$, the quantities for the efficient types are equal, or $\underline{q}^{v} = \underline{q}^{\rho}$, while for the inefficient types we have $\overline{q}^{v} > \overline{q}^{\rho}$. The rent for the inefficient type is the same, or $\overline{U}^{v} = \overline{U}^{\rho}$, while for the efficient type we have $\underline{U}^{v} > \underline{U}^{\rho}$.*

A formal derivation of the values of the variables and a proof can be found in the appendix. The proof of Proposition 1 should be understood in the following way. For the efficient agent both types of principals stipulate the same, first-best quantity. However, the ρ-principal offers a smaller rent, because she mistakenly believes that there are more efficient agents than there really are. For the inefficient type both contracts offer the same rent, but the v-principal stipulates a bigger quantity. This is so, because the odds of being matched with an efficient agent appear too large for the ρ-principal. Since the quantity for the inefficient type is smaller the smaller the odds of being matched with one, the quantity of the optimistic principal is set too low.

Quantities

At the end of each period contracts are executed and the total quantity is produced in the market. Profits of principals are realized on the basis of their produced quantities and the total quantity. We denote the expected quantity at the end of a period produced for a principal with belief ϕ by $\mathbb{E}_\theta\left[q^{\phi}_{t+1}(\theta)\right] = v\underline{q}^{\phi} + (1-v)\overline{q}^{\phi}$ and with $\tilde{q} = \int_i q_i di$ the average quantity in the market. Given the different proportions of principals with different beliefs, an informal appeal to the Law of Large

Numbers allows us to write the average total quantity at the end of a period as:

$$\tilde{q}_{t+1} = \int_i q^i_{t+1} di = \alpha_t \mathbb{E}_\theta \left[q^v_{t+1}(\theta)\right] + (1-\alpha_t)\mathbb{E}_\theta \left[q^\rho_{t+1}(\theta)\right] \tag{2}$$

Profits of principals

For contracts stipulated in t, the realized profits for each $\theta \in \{\underline{\theta}, \overline{\theta}\}$ and for a given belief $\phi \in \{v, \rho\}$ are functions $\underline{\pi}^\phi_{t+1}(\tilde{q}_{t+1}, \tilde{q}_t, \phi, \underline{\theta})$ and $\overline{\pi}^\phi_{t+1}(\tilde{q}_{t+1}, \tilde{q}_t, \phi, \overline{\theta})$, while the expected profits are $\mathbb{E}_\theta \left[\pi^\phi_{t+1}(\tilde{q}_{t+1}, \tilde{q}_t, \phi, v)\right]$. In the last term, the presence of $\tilde{q}_t$ comes from the fact that each quantity $q_{t+1}(\theta)$ is defined as a function of the $\tilde{q}$ in the market one period earlier, which comes from Assumption 1. The following proposition concerns the differential profits for each type of principal given the type of agent.

Proposition 2. *Given different beliefs and the same naive expectations about the total quantity, for any realization $\theta \in \{\underline{\theta}, \overline{\theta}\}$, the realized profits $\pi^\phi(\cdot, \theta)$ are such that:*

$\underline{\pi}^\rho_{t+1} > \underline{\pi}^v_{t+1}$ *with*

$$\underline{\pi}^\rho_{t+1} - \underline{\pi}^v_{t+1} = (\Delta\theta)^2 \frac{\rho - v}{(1-\rho)(1-v)} \tag{3}$$

$\overline{\pi}^v_{t+1} \lesseqqgtr \overline{\pi}^\rho_{t+1}$ *with*

$$\overline{\pi}^v_{t+1} - \overline{\pi}^\rho_{t+1} = \Delta\theta \frac{\rho - v}{(1-\rho)(1-v)} \left[\delta(\tilde{q}_{t+1} - \tilde{q}_t) + \frac{1}{2}\Delta\theta \left(\frac{v}{1-v} + \frac{\rho}{1-\rho}\right)\right] \tag{4}$$

Again, the derivations of functions can be found in the appendix. Parts of the results in Proposition 2 are a direct result from Proposition 1. Since the unbiased v-principal increases the rent for the efficient agent, but produces the same quantity as the biased ρ-principal, profits must be smaller. This can be seen from equation (3). On the other hand, as can be seen in equation (4) the difference in profits for the inefficient agent depends on the change of the quantity $\tilde{q}$ from one period to the next. The v-principal makes a larger profit than the ρ-principal if the quantity decreases or is constant from one period to the next. The reverse is true if the change is positive and large enough.

To summarize, the basic stage game defines an aggregative game, in which the profit of principals in a particular period depends on the belief about the distribution of types, the specific match and the behavior of all other principals, which affects the total quantity in the market. The aggregative game is characterized by submodularity in the benefit function.

Remark about the timing and benefit function

Given our timing and the specification of our benefit function, we remark that our model is mathematically equivalent to a model where principals form expectations about the price instead of the quantity, where the relationship is given by a linear demand function. In other words, our model can be related to standard cobweb models such as discussed in Hommes (2013). This equivalent model can be summarized as follows:

1. Each principal maximizes:

$$\underset{q_{t+1},t_{t+1}}{Max}\ \phi\left\{\mathbb{E}_t\left[P_{t+1}\right]\underline{q}_{t+1}-\frac{\underline{q}_{t+1}^2}{2}-\underline{w}_{t+1}\right\}+(1-\phi)\left\{\mathbb{E}_t\left[P_{t+1}\right]\overline{q}_{t+1}-\frac{\overline{q}_{t+1}^2}{2}-\overline{w}_{t+1}\right\}$$

   under (ICs) and (PCs), and therefore a linear supply function is obtained.

2. Principals have naive expectation about the price: $\mathbb{E}_t\left[P_{t+1}\right]=P_t$.

3. The demand is linear: $Q_{t+1}=A-BP_{t+1}$.

4. Market clears: the prices are computed on the demand function.

The connection to our model is established for $\beta=\frac{A}{B}$ and $\delta=-\frac{1}{B}$. Our choice of the interval for $\delta\in(-1,0)$ defines a standard stable cobweb model (in the absence of any kind of heterogeneity of expectations about any variable). This simple observation will be useful later on to make a comparison to those models, which introduce heterogeneity into unstable cobweb models, such as Brock and Hommes (1997), Branch and McGough (2008) and Hommes (2013). These approaches model heterogeneous expectations about the price on the basis of an unstable cobweb model. Hence, it is worth pointing out that our modeling of the quantity is equivalent to modeling price time series with a linear demand. However, since we are interested in the micro foundation of the adverse selection game it is more natural to use the quantity as the time series variable.

2.2 Evolutionary learning by imitation

Our evolutionary learning model is very closely related to the deterministic evolutionary models discussed in Fudenberg and Levine (1998). We take the basic idea of the social learning model and combine this with the imitation within reference groups by Selten and Ostmann (2001). The way we derive the dynamic largely follows the exposition and notation in Sandholm (2011) and Gintis (2009).

Recall that we interpret the belief held by a principal as a sort of pure strategy in the game. Each principal "plays" the pure strategy, which can be either v or ρ and earns a profit at the end of a period. In each period some principals exit the economy and are replaced by new principals, who take over the strategies of the exiting principals. What is important is that principals have no memory about past plays. The next step is to define the conditional switch rate, which is the probability γ that at the end of a period a principal changes beliefs, either from v to ρ, or the other way around. To do that, we periodically allow some principals at the end of a period to observe the profit of a second principal from the same reference group (Selten and Ostmann, 2001). The reference group in our case are all those principals, who got matched with the same type of agent in the period that just ended. Stated differently, we make the assumption that a principal will only take a change of strategy into consideration if for her the situation of the other principal is comparable. In addition, we condition the probability to switch beliefs on the absolute difference in profits, where a larger profit differential increases the probability.

Hence, in what follows we will define a *proportional imitation rule* (see Schlag (1998) and references therein for a discussion on different imitation rules) as the rule for principals to switch within reference groups. Given the proportion of efficient types v and the proportion α_t of principals using v, we need to define a conditional switch rate from one belief to the other. In order to do that we define the following probabilities: $P(\phi \rightsquigarrow \neg\phi) = \alpha_t (1-\alpha_t)$ is the probability that a principal with a belief ϕ meets a principal with the opposite belief. Since we assume that matching between principals and different types of agents is type-independent, the probability that two principals were matched with an efficient agent is simply $(v)^2$ and the probability that both were matched with an inefficient agents is given by $(1-v)^2$. Given these definition we define the probabilities that two principals with opposite beliefs and in the same reference group meet:

$$\gamma_t^{v\rho} = P(v \rightsquigarrow \rho) v^2 = \alpha_t (1-\alpha_t) v^2$$

$$\gamma_t^{\rho v} = P(\rho \rightsquigarrow v) (1-v)^2 = \alpha_t (1-\alpha_t) (1-v)^2$$

In words, $\gamma_t^{v\rho}$ is the probability that a v-principal would consider switching to belief ρ with a similar interpretation of $\gamma_t^{\rho v}$. Next, we need the probability of switching to the other strategy. As is standard in the literature, the simplest way to capture the idea that the higher the profit differential the more likely a principal switches to the other belief is to make the switching probability linearly dependent on the payoff difference. Formally, the switching probability can then be expressed as $\Omega\left[\pi_{t+1}^{\phi} - \pi_{t+1}^{\neg\phi}\right]$, where Ω is chosen to scale the payoff difference in such a way that it can be used as a probability. The economic interpretation is straightforward. Given a difference in payoffs, the higher Ω the higher the *propensity to switch*. In other words, it can be interpreted as the inverse of a degree of stubbornness of principals. It will be specified in more detail below.

Putting the pieces together, the dynamics over time are described by the following equation:

$$\alpha_{t+1} = \alpha_t + \gamma_t^{\rho v} \left\{\Omega\left[\overline{\pi}_{t+1}^{v} - \overline{\pi}_{t+1}^{\rho}\right]\right\} - \gamma_t^{v\rho} \left\{\Omega\left[\underline{\pi}_{t+1}^{\rho} - \underline{\pi}_{t+1}^{v}\right]\right\}$$

The equation should be read as follows. The fraction of v-principals in the next period is equal to the fraction in the previous period plus all ρ-principals, who switch to v minus all v-principals, who switch to ρ. From Proposition 2 we know that the term $\overline{\pi}_{t+1}^{v} - \overline{\pi}_{t+1}^{\rho}$ can be positive or negative depending on the magnitude and direction of fluctuations of the quantity in the market. If the term is negative the direction of proportional imitation is reversed, which means that the v-principal switches to ρ with the given probability. The resulting equation is equivalent[4]. Substituting the specific switch rates defined above we arrive at the discrete change of α from one period to the next:

$$\alpha_{t+1} = \alpha_t + \alpha_t (1-\alpha_t) \Omega \left\{(1-v)^2 \left[\overline{\pi}_{t+1}^{v} - \overline{\pi}_{t+1}^{\rho}\right] - v^2 \left[\underline{\pi}_{t+1}^{\rho} - \underline{\pi}_{t+1}^{v}\right]\right\} \tag{5}$$

It should be noted that the dynamic in equation (5) is close to but not the same as the well-known replicator dynamic. While the general derivation of the dynamic is the same as for the replicator dynamic the difference comes from the choice of the reference group. Since we allow only principals

[4]To see this, write the dynamic for $\overline{\pi}_{t+1}^{v} < \overline{\pi}_{t+1}^{\rho}$ as $\alpha_{t+1} = \alpha_t - \gamma_t^{\rho v} \left\{\Omega\left[\overline{\pi}_{t+1}^{\rho} - \overline{\pi}_{t+1}^{v}\right]\right\} - \gamma_t^{v\rho} \left\{\Omega\left[\underline{\pi}_{t+1}^{\rho} - \underline{\pi}_{t+1}^{v}\right]\right\}$, which is equivalent.

matched with the same type of agent to consider switching the curly brackets in (5) do not represent the difference in expected payoffs from the two strategies as in the standard replicator dynamic, but the difference in expected payoffs conditional on the outcome of the matching.

3 Equilibria, stability and dynamics

The economy in our model is governed by equations (2) and (5). Having set up the basic model describing our economy the natural next step is to analyze the evolution of the economy. It is worth pointing out that we are not interested in explaining why a population starts with a specific initial value of the distribution of beliefs. Rather, we want to analyze how the population evolves, which are the steady states and how those depend on the initial conditions. In what follows, we will first study possible fixed points of the system. We will derive conditions for the existence of two fixed points of a monomorphic population in which each principal has the same belief (either v or ρ). The presence of two steady states comes from our assumption about the irrationality of principals. It is worth to be clear on the meaning that we give to "irrationality" here. It is true that principals act in a rational way in writing contracts according to their beliefs; the irrational behaviour comes from their inability to update their priors in a Bayesian fashion. This is so, because the information set available to principals is limited. Taken together, a steady state, in which each principal has belief ρ cannot be ruled out a priori. In addition to the two fixed points, under certain conditions, infinitely many fixed points are possible, in which the population is polymorphic. In other words, both beliefs coexist. The next step then is to analyse the stability of the fixed points. Again, the stability will crucially depend on the relationship between the beliefs v and ρ. Then we will show that in addition to the fixed points involving either bimorphic or polymorphic populations under certain conditions periodic cycling patterns and chaotic behaviour can be obtained. In relation to that the analysis will provide conditions for local bifurcations that arise for different parameter values. In general, a bifurcation is a change in the qualitative structure of a dynamic system for changes in some parameters. Due to the nonlinearity of the system the analysis of the global behavior of the system is more difficult[5]. Helped by computer-assisted proofs we are able to derive conditions for global bifurcations of limit cycles and chaotic behaviour of the system, which lead to endogenous fluctuations of the variables.

3.1 Fixed points

In order to analyze what determines the change of the variables from one period to the next, we rewrite equations (2) and (5) and rearrange using the functional forms of the quantities and profits. The algebraic derivation can be found in the appendix.

$$\alpha_{t+1} = \alpha_t + \alpha_t \left(1 - \alpha_t\right) a\Omega \left\{ (1-v)^2 \left\{ \delta \left[\beta - \overline{\theta} + (\delta - 1)\widetilde{q}_t - \left(1 - \alpha_t\right) c \right] + b \right\} - v^2 \Delta\theta \right\} \tag{6}$$

$$\widetilde{q}_{t+1} = \beta - \overline{\theta} + \delta \widetilde{q}_t - \left(1 - \alpha_t\right) c \tag{7}$$

[5]For an introduction to dynamical systems that puts an emphasis on the discussion of local and global bifurcations see Guckenheimer and Holmes (1983).

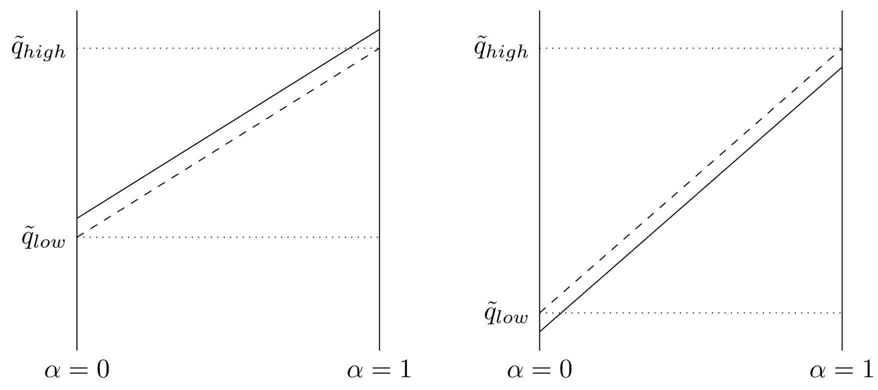

(a) $\rho < \rho_c$ (b) $\rho > \rho_c$

Figure 1: Two phase diagrams for different values of ρ with α on the horizontal and $\tilde{q}$ on the vertical axis. The dashed line shows the locus of points where $\tilde{q}$ is fixed, the unbroken lines show the locus of points where α is fixed over time. Panel (a) ((b)) shows the case where $\rho < \rho_c$ ($\rho > \rho_c$) as defined in the text. $\tilde{q}_{high}$ ($\tilde{q}_{low}$) is the fixed point-quantity associated with $\alpha = 1$ ($\alpha = 0$).

where $a = \Delta\theta \frac{\rho - v}{(1-\rho)(1-v)}$, $b = \frac{1}{2}\Delta\theta \left(\frac{v}{1-v} + \frac{\rho}{1-\rho}\right)$ and $c = \Delta\theta \frac{\rho - v}{1-\rho}$.

These last two equations define a nonlinear map Γ from a state $X_t = [\alpha_t, \widetilde{q}_t]^T$ to $X_{t+1} = \Gamma(X_t)$. Given the nonlinearity, before attempting to analyse the global behavior we start identifying the steady states and their local stability. The nullclines of the system are plotted in Figure 1. The dashed line gives the combination of points where $\tilde{q}_{t+1} = \tilde{q}_t$ from equation (6) and the unbroken line gives the locus of points where $\alpha_{t+1} = \alpha_t$ from equation (7).

To better describe the fixed points and their dependence on the beliefs we introduce a critical ρ_c such that $k \equiv (1-v)^2 b - v^2\Delta\theta = 0$. This ρ_c determines the relative location of the nullclines. For $\rho < \rho_c$ ($> \rho_c$) the nullcline giving the steady states for the quantity $\tilde{q}$ is below (above) the one for α (see Figure 1). The third case (not shown) is $\rho = \rho_c$, when the two diagonal nullclines overlap. From Figure 1, it is clear that the system admits two or infinitely many fixed points, where the latter occurs only when $\rho = \rho_c$. Interested only in the two non-degenerate cases (where the nullclines do not overlap) we can claim:

Lemma 1. *Whenever $k = (1-v)^2 b - v^2\Delta\theta \neq 0$, the nonlinear system always admit two hyperbolic steady states $X^0 = (0, \widetilde{q}_{low})$, $X^1 = (1, \widetilde{q}_{high})$ with*

$$\widetilde{q}_{low} = \frac{\beta - \overline{\theta}}{1-\delta} - \frac{c}{1-\delta}$$

$$\widetilde{q}_{high} = \frac{\beta - \overline{\theta}}{1-\delta}$$

The quantities $\tilde{q}_{low}$ and $\tilde{q}_{high}$ are simply the intersections of the nullcline associated with the quantity (the dashed line in Figure 1) with the vertical parts of the nullclines for α. The proof is omitted. Next, in order to study the stability the Jacobians are evaluated at the steady states:

$$J\left(X^0\right) = \begin{bmatrix} 1 + a\Omega k & 0 \\ c & \delta \end{bmatrix} \qquad J\left(X^1\right) = \begin{bmatrix} 1 - a\Omega k & 0 \\ c & \delta \end{bmatrix}$$

Using the usual definitions related to local bifurcations, from the Jacobians the following proposition immediately follows:

Proposition 3. *Given beliefs v and ρ the following holds for the system defined in (6) and (7).*

1. *The system has always either*
   (a) *a stable and an unstable fixed point, which are both hyperbolic (for $k \neq 0$), or*
   (b) *two non-hyperbolic fixed points and infinitely many fixed points (for $k = 0$).*
2. *The system has a local fold bifurcation for both fixed points if $k = 0$.*
3. *The system undergoes a transcritical bifurcation. For $k > 0$ X^1 is the stable and X^2 is the unstable fixed point. It is the other way around for $k < 0$.*

The proof for Proposition 3 involves a simple inspection of the eigenvalues of the Jacobians. The two eigenvalues are δ and $1 \pm a\Omega k$. Then, for $k = 0$ one eigenvalue crosses the unit circle for both points (fold bifurcation). Moreover, for k changing sign one point has both eigenvalues smaller then one while the other becomes a saddle point exchanging stability (transcritical bifurcation). This implies that for $\rho > \rho_c$ the point X^1 is a stable hyperbolic steady state, which corresponds to the situation depicted in Figure 1b. Accordingly, the reverse case of $\rho < \rho_c$ is in Figure 1a, in which X^2 is stable.

3.2 Normalization

At this point, two interrelated remarks are needed concerning first the value for Ω and second a possible *shutdown policy* of principals, which in the literature of mechanism design refers to a situation in which principals choose to write contracts only for the efficient type (see Laffont and Martimort, 2002, chapter 2). Concerning the latter point, we make it clear that a linear demand and supply function in the standard cobweb model with naive expectations leads necessarily to fluctuations in the quantity (and therefore price) such that negative values are unavoidable unless more restrictive assumptions are imposed. Therefore, it appears clear that high values of the quantity $\tilde{q}$ in a preceding period could lead principals (who solve maximization problem with naive expectations about $\tilde{q}$) to adopt a *shutdown policy* for the inefficient type whenever the gain from a negotiation with this type is, in expectation, negative. Turning to the first point, we remind the reader that when we introduced the propensity to switch Ω in the previous section we simply defined it as a parameter needed to bound the expression to unity such that it could be considered a probability. This standard normalization is the basis to define the well-known replicator dynamic and the assumption there is innocuous, because it does not affect the global dynamic (Weibull, 1995). However, here we deal with a discrete set-up and the same does not necessarily apply unless we are close enough to the fixed points such that the system is topologically equivalent to its linearization. To show this, we start by reporting some important tools that are useful for our discussion. The following definitions and some results can be found in the literature on nonlinear maps and they are to large extent taken from Kuznetsov (1998) and Hommes (2013). The latter represents a cornerstone for the study of chaotic behaviour in economic modeling and in which several applications of bifurcation analysis to economics can be found. We start by defining the stable and unstable manifolds for our dynamical system.

Definition 1. *Given a map $\Gamma(x)$ and a fixed point x_0, the stable and unstable manifolds $W^s(x_0)$, $W^u(x_0)$ are*

$$W^s(x_0) = \{x : \Gamma^n(x) \to x_0, n \to \infty\}$$

$$W^u(x_0) = \{x : \Gamma^n(x) \to x_0, n \to -\infty\}$$

where n is integer "time" and $\Gamma^n(x)$ denotes the nth iterate of x under Γ.

On the basis of this definition the following Local Stable Manifold theorem can be formulated.

Theorem 1 (Local Stable Manifold)**.** *Let x_0 be a hyperbolic fixed point. Then the intersections of $W^s(x_0)$ and $W^u(x_0)$ with a sufficiently small neighborhood of x_0 contain smooth submanifolds $W^s_{loc}(x_0)$ and $W^u_{loc}(x_0)$ of dimension equal to the number of stable/unstable eigenvectors, respectively. Moreover, $W^s_{loc}(x_0)$ $(W^u_{loc}(x_0))$ is tangent at x_0 to $T^s(T^0)$, where $T^s(T^0)$ is the generalized eigenspace corresponding to the union of all eigenvalues of J (Jacobian matrix) with eigenvalue $|\lambda| < 1 (|\lambda| > 1)$.*

The proof can be found in Kuznetsov (1998: 50). The theorem claims that whenever the system is close enough to a steady state the stable and unstable manifolds are tangent to the respective stable and unstable eigenvectors of the linearized system. On the basis of this theorem, we can easily show the following Lemma.

Lemma 2. *For every two different Ω and $\Omega^{'}$ there exists $\Delta\theta^{'} = \overline{\theta}^{'} - \underline{\theta}^{'}$, such that the system with Ω and $\Delta\theta$ is topological equivalent near the steady state to the system with $\Omega^{'}$ and $\Delta\theta^{'}$.*

Proof. Given the Jacobians in the steady state, the two eigenvalues are δ and $1 \pm a\Omega k$. Since $\delta < 0$, the corresponding eigenvector is the stable one, it is invariant and corresponds to the vertical line in $\alpha = 0$ $(\alpha = 1)$. The opposite is true for $1 \pm a\Omega k$. Then, for an Ω' it is sufficient to define a rescaling of $\Delta\theta$ such that $a\Omega k = a'\Omega' k'$. The rest follows from the Local Stable Manifold theorem. ∎

The result in Lemma 2 implies that the normalization of the parameter Ω is not problematic for qualitative analyses of the dynamical system. In the numerical simulations that follow all parameters have economically meaningful values.

3.3 Global behaviour

Next, we move from a local analysis to a global analysis of the dynamical system. We will describe possible endogenous fluctuations of the quantity and price, which are not driven by a change in the fundamentals.

Theorem 2. *For δ close enough to -1 and $|\rho - \rho_c|$ small enough, the map $\Gamma(x)$ admits a limit-two cycle:*

$$\Gamma(x') = x'' \text{ and } \Gamma(x'')) = x' \qquad \text{for } x' \neq x'', \Gamma(x') \neq \Gamma(x'')$$

While the proof of the theorem can be found in the appendix where we also give a system of nonlinear equations with explicit solution from which one can obtain the limit cycle, here we give a more intuitive explanation. It is clear that for $\delta = -1$ the steady states undergo a flip bifurcation and then a limit-two cycle is born. At the same time, $\rho = \rho_c$ implies, as seen above, the fold bifurcation. Then the steady states could, in principle, have a *strong resonance* with multipliers $-1, 1$. Intuition would suggest that for values close to the strong resonance a nontrivial solution of the map can arise by the presence of a saddle-node and of a possible limit-two cycle. The presence of this last one can affect (see for example Kuznetsov (1998) or Wiggins (2003) for detailed explanations) the global behaviour whenever it undergoes a *Neimark Sacker bifurcation*, i.e. two complex eigenvalues which have modulus equal to unity[6]. While, as said, the proof of the existence of a real solution of the second iterate $\Gamma^{(2)}(x)$ can be written in closed form, unfortunately the Jacobian of the second iterate has a nontrivial expression and only a numerical solution can be obtained.

Then, in what follows, we will use the E&F Chaos program as described in Diks et al. (2008) to produce bifurcation diagrams associated with the relevant values of ρ and δ for a *ceteris paribus* analysis of the bifurcation. In the bifurcation analysis the following parametrization is used: $\beta = 100, \overline{\theta} = 24, \underline{\theta} = 22, v = 0.6, \Omega = 0.01$. Figure 2 and 3 are the bifurcation diagrams associated with δ and ρ. The diagrams plot the long-run behaviour of the system as a function of some parameters.

While we report on the vertical axis the quantity, we underline that a convergence to a specific value $(\widetilde{q}_{low}, \widetilde{q}_{high})$ (not reported in the figures) corresponds to a convergence towards the respective monomorphic state of the population. The two sets of figures support the result of Theorem 2 and they essentially show the different configurations needed to observe the presence of a limit-two cycle. On the one hand it is not surprising to observe a cycle for values close to ρ_c and $\delta = -1$. On the other hand, possible bifurcations that the limit cycle can undergo and the impact on the global behaviour caused by its presence in conjunction with a saddle-node could be less intuitive. In fact, what follows is devoted to explain, with the help of a graphical and numerical analysis, how a simple deterministic model with heterogeneous agents could lead to endogenous fluctuations and cyclical behaviour[7]. Then, at this point, it is worth to be clear about the methodology used and

[6]The Neimark Sacker bifurcation is the discrete version of a *Hopf* bifurcation.

[7]The numerical simulations used in this paper are based on our Mathematica, Matlab and E&F Chaos codes. The codes are available upon request.

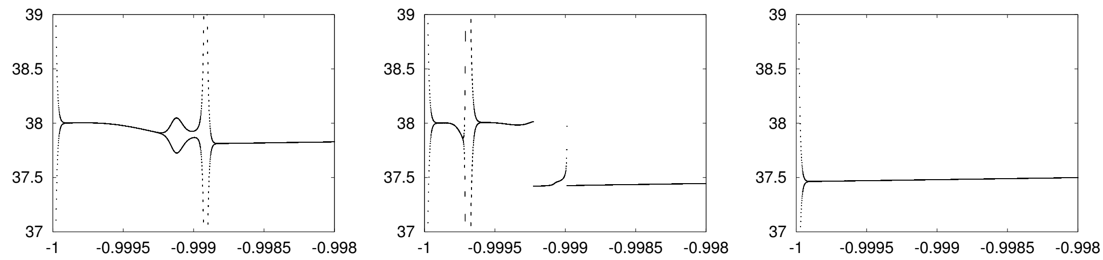

(a) $\rho = 0.7499$ (b) $\rho = 0.749$ (c) $\rho = 0.74$

Figure 2: Bifurcation diagrams of δ for different values of ρ

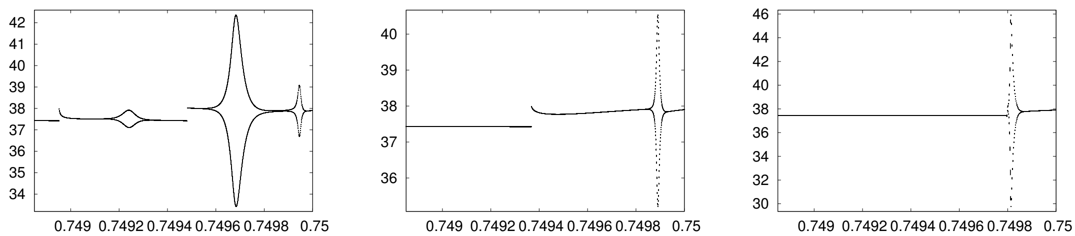

(a) $\delta = -0.999$ (b) $\delta = -0.9989$ (c) $\delta = -0.9988$

Figure 3: Bifurcation diagrams of ρ for different values of δ

the obtained results. Till now, we have performed a simple local bifurcation analysis of the steady states and we have found the presence of a limit two cycle by inspection of the second iterate. The following step will be to detect the Neimark Sacker bifurcation for the limit-two cycle and showing the existence of an attracting orbit generated by the presence of a *closed invariant curve* close to the stable and unstable manifolds of the saddle-node.

As mentioned, the study of the Jacobian of the second iterate can only be done by numerical simulations given the iterative complexity deriving from the population dynamic. However, the bifurcation diagram in figure 2 and 3 suggests some values; namely we obtain for $\rho = 0.7499$ and $\delta = -0.999$ two complex eigenvalues whose modulus is equal to $|\lambda| \approx 1.000007782$, while for $\rho = 0.745$ and $\delta = -0.999$ the eigenvalues are real. Therefore, by a sort of continuity argument we are confident to make the claim that there exists a point for which the limit-two cycle undergoes a Neimark Sacker bifurcation. To complete the proof, we refer to a common result in dynamical systems, which says that the fixed point undergoing a Neimark Sacker bifurcation is surrounded by an isolated closed invariant curve for values "after" or "before" the critical value of the parameter (Kuznetsov (1998):125-137).

Having identified the bifurcation of the cycle, the next step is to study a possible homoclinic orbit which belongs to the intersection of the stable and unstable manifolds which would lead to the existence of a *strange attractor*. An intersection of the invariant manifolds would imply the presence of infinitely many intersections and a nontrivial global behaviour of the map

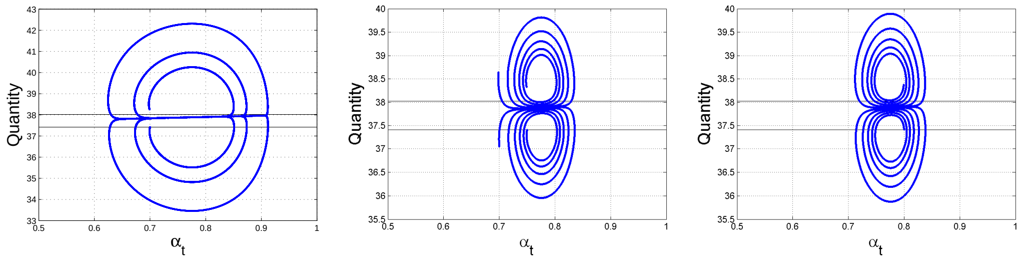

(a) $\alpha_0 = 0.7$, $\widetilde{q}_0 = \widetilde{q}_{low}$ (b) $\alpha_0 = 0.75$, $\widetilde{q}_0 = \widetilde{q}_{low}$ (c) $\alpha_0 = 0.8$, $\widetilde{q}_0 = \widetilde{q}_{low}$

Figure 4: Phase Diagram

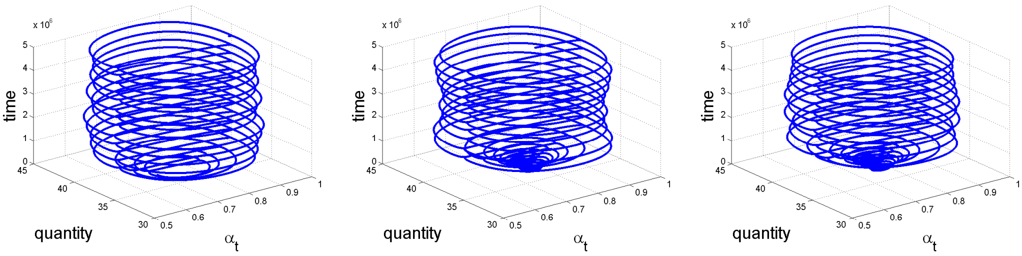

(a) $\alpha_0 = 0.7$, $\widetilde{q}_0 = \widetilde{q}_{low}$ (b) $\alpha_0 = 0.75$, $\widetilde{q}_0 = \widetilde{q}_{low}$ (c) $\alpha_0 = 0.8$, $\widetilde{q}_0 = \widetilde{q}_{low}$

Figure 5: Phase Diagram in 3D

with sensitive dependence on the initial conditions[8]. In our model, the stable manifold is simply the vertical axis passing through the α of the steady state, while the unstable manifold as usual is hard to identify and normally it is approximated by using iterated points on the corresponding unstable eigenvector. Even if we are not able to show a possible intersection, we give an alternative explanation of the presence of a nontrivial appearance of our attractor. As will become evident, the numerical analysis that we perform supports an interaction between the invariant closed curve born from the Neimark Sacker bifurcation and the unstable manifold of the saddle-node. In Figure 4 the phase diagrams for three different initial conditions can be seen, while Figure 5 shows the evolution over time for five million iterations. The diagrams clearly show a non-convergence and an irregular behaviour which is not repetitive even for periods of length greater than the one reported here. The limit-two cycle is, for the parametrization used until now, given by $x' = \left[\alpha'^{(2)} = 0.7758561993315862, q'^{(2)} = 37.333508591139584\right]^T$, $x'' = \left[\alpha''^{(2)} = 0.7749371859015566, q''^{(2)} = 38.435143334811066\right]^T$. The corresponding time series for the quantity and for the proportion of types are reported in figure 6.

The discussion above clearly evidences that any asymptotic convergences and the global behaviour are fundamentally affected by the presence of a limit-two cycle. From an economic point of view the obtained results are relevant in terms of time series, but more importantly in terms of

[8]An intersection of a stable and unstable manifold leads to an infinitely many intersections just observing that by definition each point on a manifold must be a preimage of the next iteration.

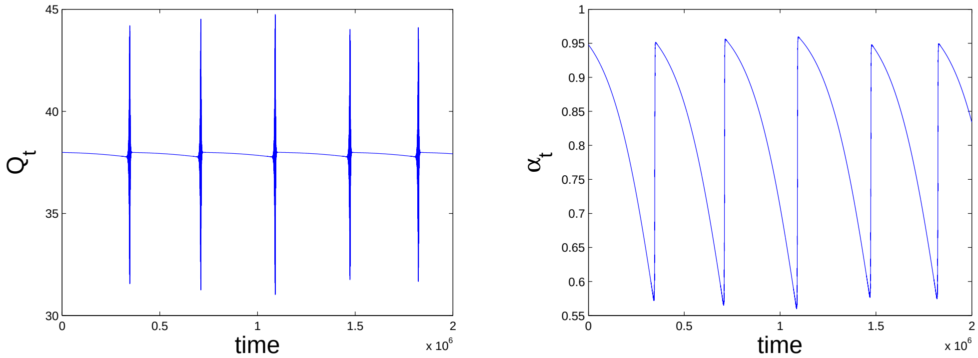

(a) Time Series for Quantity

(b) Time Series for α

Figure 6: Time series

possible under/over production in which long periods of almost stable price are replaced by periods of oscillations where no change in the fundamentals is observed. In addition, besides this we can then, at this point, claim more about the convergence towards a unique monomorphic state where all the principals adopt the same "rule" of writing contracts.

Theorem 3. *Whenever the conditions for the existence of the limit-two cycle are not satisfied (δ and ρ not too close to their critical values) and if some principals are optimistic then the following holds:*

1. *for a low degree of optimism ($\rho < \rho_c$), the population will converge to a situation where all the principals are optimistic.*

2. *for a high degree of optimism ($\rho > \rho_c$), the population will converge to a situation where all the principals are unbiased.*

Proof. The proof is immediate in the absence of heteroclinic/homoclinic tangles which we just detected by the bifurcation diagram, but for which we do not have a formal proof. In fact, in the absence of a limit-two cycle, which acts as an attractor there are only two fixed points, one stable and the other unstable. Then since the only source of uncertainty (the cycle) is not present, it is sufficient to observe that as ρ increases and thereby surpasses ρ_c (but far enough from it), then the two fixed points exchange stability. This proofs the first part of the theorem. Due to the symmetry of the system the same result holds for the second part of the theorem. ∎

4 Discussion and concluding remarks

Our paper introduces an evolutionary learning model with beliefs into an adverse selection problem. We assume that principals have one of two possible beliefs about the distribution of the ability of agents. Either principals are biased in the direction of optimism or they are unbiased in their belief. In our model the evolutionary learning takes place in the form of a linear profit comparison.

The higher the fraction of principals with a particular belief and the higher the payoff difference between two randomly chosen principals from the same reference group the higher the probability to switch to the other belief. On the basis of this set-up we are able to show that under certain conditions bimorphic populations can exist. This means that both beliefs coexist. In addition, we detect cycling patterns and report some strong evidence that a complicated Neimark Sacker bifurcation can occur on the system.

To conclude we would like to return to some of the assumptions made in the model and discuss their implications and possible changes in future research. First, we made the assumption that there are just two types of agents. The mechanism design literature allows also for a continuity of types. On the one hand, from a formal point of view it is possible to include the continuity of abilities. In order to allow comparison of principals one would have to define a norm, which represents the reference group. In other words, a principal is willing to compare profits if a randomly drawn different principal was matched with an agent with an ability close to the own match. Also, numerical simulations can be done in this case. However, we decided against this route simply because of economical considerations. The additional generality of the result is not enough to outweigh the additional mathematical complexity of them model at this point. Surely, though, this is a possible future research. Second, we choose the demand function and the cost functions in such a way that we obtain a linear demand and supply function. This is done so we are able to achieve a comparability with a literature that examines heterogeneity of market participants in the framework of cobweb models (Hommes, 2013). In fact, what we can say is that our adverse selection story adds complexity and fluctuations to a stable cobweb model, while in this research field the complexity comes from the use of an unstable cobweb model. Last, our choice of a linear switching rule for principals might appear to be an ad hoc-decision. Obviously, there are other, non-linear rules that could be used, such as any form of logistic function and we certainly think that it could be added to a model like ours. However, one should keep in mind that by choosing a linear switching rule we stack the deck against us in a certain way. This is so, because using the linear form denies any interior fixed point.

Appendix

Proposition 1

Proof. As standard, we have that the quantities of the efficient and inefficient type are defined implicitly by: $S'_{\overline{q}}(\cdot) = \underline{\theta}$, $S'_{\underline{q}}(\cdot) = \overline{\theta} + \frac{\phi}{1-\phi}\Delta\theta$. Using the specific functional form for $S(\cdot)$, we simply obtain:

$$\underline{q}_t^v = \underline{q}_t^\rho = \beta + \delta\tilde{q}_{t-1} - \underline{\theta} \tag{8}$$

$$\overline{q}_t^\phi = \beta + \delta\tilde{q}_{t-1} - \overline{\theta} - \frac{\phi}{1-\phi}\Delta\theta \tag{9}$$

with $\Delta\theta = \overline{\theta} - \underline{\theta}$.
From $\rho > v$, follows: $\underline{q}_t^v = \underline{q}_t^\rho > \overline{q}_t^v > \overline{q}_t^\rho$ ■

Proposition 2

Proof. Given equation (9), computing the difference in quantities for the inefficient type we obtain: $\overline{q}_{t+1}^v - \overline{q}_{t+1}^\rho = \Delta\theta\left[\frac{\rho}{1-\rho} - \frac{v}{1-v}\right]$. While the differences in payoffs, taking into account that the rent for the inefficient type is zero, are simply obtaining as follows:

$$\underline{\pi}_{t+1}^\rho - \underline{\pi}_{t+1}^v = -\Delta\theta\overline{q}_{t+1}^\rho + \Delta\theta\overline{q}_{t+1}^v = \Delta\theta\left[\overline{q}_{t+1}^v - \overline{q}_{t+1}^\rho\right] = (\Delta\theta)^2\,\frac{\rho - v}{(1-\rho)(1-v)}$$

$$\begin{aligned}\overline{\pi}_{t+1}^v - \overline{\pi}_{t+1}^\rho = \left[\overline{q}_{t+1}^v - \overline{q}_{t+1}^\rho\right]\left\{\beta - \frac{\left[\overline{q}_{t+1}^v + \overline{q}_{t+1}^\rho\right]}{2} + \delta\widetilde{q}_{t+1}\right\} = \\ \Delta\theta\frac{\rho - v}{(1-\rho)(1-v)}\left\{\delta\left[\widetilde{q}_{t+1} - \widetilde{q}_t\right] + \frac{1}{2}\left(\frac{v}{1-v} + \frac{\rho}{1-\rho}\right)\right\}\end{aligned}$$

■

Derivation of the nonlinear map

Proof. The derivation of equation (7) is straightforward. Using it in equation (4) to eliminate its dependence on $\widetilde{q}_{t+1}$ gives:

$$\overline{\pi}_{t+1}^v - \overline{\pi}_{t+1}^\rho = a\Delta\theta\delta\left\{\left[\beta - \overline{\theta} + (\delta - 1)\widetilde{q}_t - (1-\alpha_t)\,c\right] + b\right\}$$

Then using the expressions for the differences in realized payoffs in the replica equation (2), we obtain (6). ■

Theorem 2

Proof. To simplify the algebra, let $\Phi \equiv a\Omega$, $R \equiv \Phi\left(1-v\right)^{2}\frac{\delta}{1+\delta}c$,$S \equiv \Phi k$. From equation (7), the solution of the second iterate for the quantity must solve:

$$q^{(2)} = \frac{\beta - \overline{\theta}}{1-\delta} - \left(1-\alpha_{\tau}\right)\frac{\delta}{1-\delta^{2}}c - \left(1-\alpha_{\tau+1}\right)\frac{1}{1-\delta^{2}}c \tag{10}$$

Equation (6), using (10) and after some algebra can be written for a generic time τ as:

$$\alpha_{\tau+1}\left[1+\alpha_{\tau}\left(1-\alpha_{\tau}\right)R\right] = \alpha_{\tau}\left[1+\alpha_{\tau}\left(1-\alpha_{\tau}\right)R\right] + \alpha_{\tau}\left(1-\alpha_{\tau}\right) \tag{11}$$

The same relationship obviously holds for $\alpha_{\tau+1}$, $\alpha_{\tau+2}$. Using both relationships, rearranging and using an adjunct variable H, we obtain:

$$\left[H+\left(1-\alpha\right)S\right]\left[H-\alpha S\right]\left[2H-1\right]+H^{2}=0 \tag{12}$$

$$H = 1+\alpha\left(1-\alpha\right)R \tag{13}$$

Equations (12) and (13) determine $\alpha^{(2)}$, (11) $\alpha_{\tau+1}$, (10) $q^{(2)}$.

While equation (10) is linear, the solution for $\alpha^{(2)}$ can be in the set of complex number. However, observe that R is a function of δ while S is not. Observe that $\delta < 0$ implies $R < 0$ which converge to $-\infty$ for $\delta \to -1$. Equation (13) describes a parabola with minimum in $\frac{1}{2}$ and equal to $H_{min} = 1+\frac{R}{4}$; therefore $\delta \to -1$, implies $H_{min} \to -\infty$. Then, to conclude the proof it is sufficient to show that it is possible to obtain a real solution for equation (12). Solving (12) for α gives:

$$\alpha = \frac{2H+S}{2S} \pm \frac{1}{2}\sqrt{\frac{2HS^{2}-4H^{2}-S^{2}}{S^{2}(2H-1)}}$$

Therefore it is possible to choose an S small enough without affecting R such that a real solution exists, i.e. the graphs of (12) and (13) in the $\alpha - H$-plane have an intersection. See Figure 7 for an illustration. ■

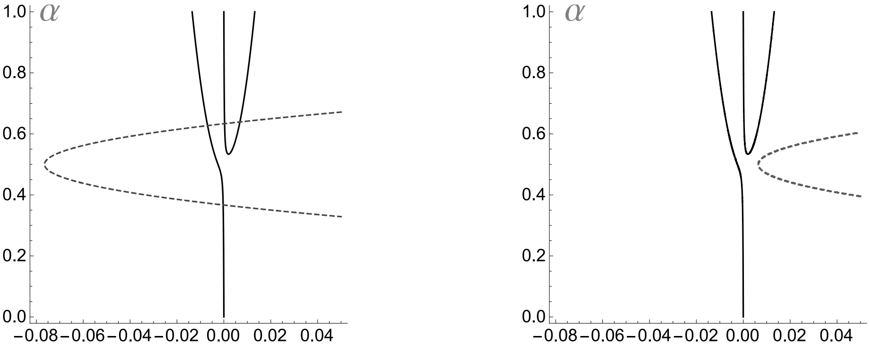

(a) $\delta = -0.9988$ (b) $\delta = -0.9987$

Figure 7: Two examples for the graphs of equations (12) and (13), where the latter is the dashed line. In both the following values were used: $\beta = 100, \overline{\theta} = 24, \underline{\theta} = 22, v = 0.6, \rho = 0.745, \Omega = 0.01$ varying only the value for δ.

References

Alchian, A. A. (1950). Uncertainty, evolution, and economic theory. *Journal of Political Economy*, 58(3):211–221.

Alós-Ferrer, C. (2004). Cournot versus walras in dynamic oligopolies with memory. *International Journal of Industrial Organization*, 22(2):193–217.

Ania, A. B., Tröger, T., and Wambach, A. (2002). An evolutionary analysis of insurance markets with adverse selection. *Games and Economic Behavior*, 40(2):153–184.

Arifovic, J. and Karaivanov, A. (2010). Social learning in a model of adverse selection. In *Industrial Organization, Trade and Social Interaction: Essays in Honour of B. Curtis Eaton*. University of Toronto Press, Toronto.

Branch, W. A. and McGough, B. (2008). Replicator dynamics in a cobweb model with rationally heterogeneous expectations. *Journal of Economic Behavior & Organization*, 65(2):224–244.

Brock, W. A. and Hommes, C. M. (1997). A rational route to randomness. *Econometrica*, 65(5):1059–1095.

Diks, C., Hommes, C., Panchenko, V., and van der Weide, R. (2008). E&f chaos: A user friendly software package for nonlinear economic dynamics. *Computational Economics*, 32(1-2):221–244.

Evans, G. W. and Honkapohja, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton and Oxford.

Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press, Cambridge, MA.

Gintis, H. (2009). *Game Theory Evolving*. Princeton University Press, Princeton.

Guckenheimer, J. and Holmes, P. (1983). *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer, New York.

Guerrieri, V., Shimer, R., and Wright, R. (2010). Adverse selection in competitive search equilibrium. *Econometrica*, 78(6):1823–1862.

Hommes, C. M. (2013). *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press, Cambridge.

Kuznetsov, Y. A. (1998). *Elements of Applied Bifurcation Theory*. Springer, New York.

Laffont, J.-J. and Martimort, D. (2002). *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, Princeton.

Pavan, A. and Calzolari, G. (2010). Truthful revelation mechanisms for simultaneous common agency games. *American Economic Journal: Microeconomics*, 2(2):132–190.

Piccolo, S. and Pagnozzi, M. (2013). Information sharing between vertical hierarchies. *Games and Economic Behavior*, 79:201–222.

Sandholm, W. H. (2011). *Population Games and Evolutionary Dynamics*. MIT Press, Cambridge, Mass.

Santos-Pinto, L. (2010). Positive self-image in tournaments. *International Economic Review*, 51(2):475–496.

Schlag, K. H. (1998). Why imitate, and if so, how? *Journal of Economic Theory*, 78:130–156.

Selten, R. and Ostmann, A. (2001). Imitation equilibrium. *Homo oeconomicus*, 18(1):111–149.

Vega-Redondo (1997). The evolution of walrasian behavior. *Econometrica*, 65(2):375–384.

Vives, X. (2001). *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press, Cambridge, MA.

Weibull, J. W. (1995). *Evolutionary Game Theory*. MIT Press, Cambridge.

Wiggins, S. (2003). *Introduction to Applied Nonlinear Dynamical Systems and Chaos*. Springer, New York.

Young, H. P. (1998). *Individual Strategy and Social Structure*. Princeton University Press, Princeton.

Young, H. P. (2004). *Strategic Learning and its Limits*. Oxford University Press, Oxford.