CHAPTER 5

IMPACT OF ENVIRONMENTAL DYNAMICS ON ECONOMIC EVOLUTION: UNCERTAINTY, RISK AVERSION, AND POLICY

Abstract

The general question of how environmental dynamics affect the behavioral interaction in an evolutionary economy is considered. To this end, a basic model of a dynamic multi-sector economy is developed where the evolution of investment strategies depends on the diversity of investment strategies, social connectivity and relative contribution of sector specific investments to production. Four types of environmental dynamics are examined that differ in how gradual and how frequent environmental change occurs. Numerical analysis shows how the socially optimal level of diversity increases with the frequency and rapidity of the changes. When there is uncertainty about which type of environmental dynamics will prevail, the socially optimal level of diversity increases with the degree of risk aversion of the policy maker or the society.

5.1 Introduction

Evolutionary reasoning and agent-based modeling are standard practice in various disciplines, including social sciences (e.g., Binmore, 1994; Galor and Moav, 2002; Tesfatsion, 2006; Mirowski, 2007). A typical evolutionary model uses a population of entities that undergo selection and variation. Although specific domains ask for the development of particular types of model, several common, general questions arise. Here we

A revised version has been published as Nannen et al. (2013).

aim to address one such question, namely how does a dynamic environment influence the behavior of an evolutionary economic system consisting of multiple agents employing different behavioral strategies. The relevance of this question is evident: few economic environments are static.

Generally, one cannot expect evolution in a changing environment to approach a steady state. What matters is not so much how well the agents adapt if given enough time, but how fast they adapt to a new challenge. In a socio-economic context a wide range of environmental variables can be identified: macroeconomic conditions, technological opportunities, policies and institutions, and natural resources. Most studies of social behavior through evolutionary methods have been limited to constant environments, letting selection pressure depend on the population distribution. In part, this allows for analytical treatments, as has been the common approach in evolutionary game theory (both in biology and the social sciences). The addition of a dynamic environment requires a numerical or computational approach. As environmental economics deals with the economic analysis of exploitation of natural resources, abatement of environmental pollution, and human-induced climate change, dynamic environments are prevalent. The evolution of strategies is important when heterogeneous groups of users, polluters, or harvesting strategies are involved (Ostrom, 2000; van den Bergh, 2007). Dynamic environments may cause certain strategies to become evolutionary stable and others to become unstable. We will not only draw upon the social sciences but also make use of certain insights from evolutionary biology. Evidently, many explicit and implicit insights on the influence of environment on evolution are available here.

For our purpose a relevant distinction is between exogenous and endogenous environments. Whereas systems with only exogenous variables are relatively simple, endogenous variables generate complex feedback systems. Unfortunately, most real-world systems studied by biologists and social scientists are of the latter type. Resource dynamics (e.g., Sethi and Somanathan, 1996; Noailly et al., 2003) and dynamic control of a pest population that evolves resistance to pesticides (Munro, 1997) are policy-relevant examples. Another, general example is a coevolutionary system in which two heterogeneous populations cause selection pressure on one another (Epstein and Axtell, 1996). This leads to very complex coevolutionary interactions because the environment of each evolutionary (sub)system is evolving as well. Coevolution thus implies a particular type of dynamic and endogenous environment (Noailly, 2008).

With regard to the evolutionary system, there is a range of theoretical starting points and modeling approaches (Eiben and Smith, 2003; van den Bergh, 2004). First of all, one can choose to use very theoretical, abstract models of the evolutionary game type. However, adding dynamic environments here will lead to systems that are no longer amenable to analytic solutions. Numeric simulations of multi-agent systems form an alternative to the analytic approach that offer much more flexibility in examining system behavior. They allow a distinction between local and global environments, and between stationary and mobile agents. They further allow to study the influence of population size, and the effects of dynamic environments on group and network formation (Bergstrom, 2002; Henrich, 2004). In addition, different assumptions can be made regarding selection factors and innovation mechanisms (random mutations, deterministic trends, recombination) and bounded rationality of agents (habits, imitation).

In this chapter we investigate the impact of various types of general environmental dynamics on the socially optimal type of behavioral interactions among the agents in the population. We consider the general structure of a dynamic non-aggregate multi-sector economy. Agents have individual investment strategies that specify how they invest their respective income. Their objective is to maximize their individual welfare. They prefer investment strategies which give high welfare. Their rational capabilities are bounded and their information is limited. The only information available to the agents is the investment strategies and the welfare of their fellow agents. The behavioral interactions influence how the agents use this information to evolve their investment strategies through imitation. Our framework postulates that the environmental dynamics are beyond the control of the policy maker, while he or she can regulate (some aspects of) the agent interactions. Various types of government regulation, information and education affect the search for and effectiveness of innovation by economic agents. In particular, by regulating how accurately agents can imitate each other, a policy maker can control the diversity of strategies within the population. Examples of policies that influence diversity are patent and copyright laws, conditions for competition for public R&D funds and subsidies, and the support or enforcement of industry standards.

We will study the effect of diversity on welfare numerically through computer simulations. We will address two research questions. The first is whether it is true that different environmental dynamics require different degrees of diversity for the agents to achieve a high welfare. The second question follows from the fact that environmental dynamics are not only beyond the control of the policy maker, but that they are also uncertain to him. This raises the issue of adequate policies under uncertainty: how do agent interactions that work well for one type of environmental dynamics perform under another environmental dynamics? Depending on the degree of risk aversion of the policy maker or the society, different policies can be recommended.

As for the environmental dynamics, we focus on two general aspects of environmental change: how gradually it occurs, and how often. Gradualness and frequency of change are two aspects of an environmental dynamics that can relatively easily be observed and recognized. Depletion of a mineral resource, for example, typically manifests itself over an extended period of time, while a biotic resource like fish can disappear literally overnight. Or a remote agricultural community is normally exposed to environmental hazards less frequently than one surrounded by a heavily industrialized region. If a policy maker can anticipate these aspects of environmental change, he or she might want to steer behavioral interaction such that economic agents can adapt well.

The remainder of this chapter is organized as follows. Section 5.2 describes production and growth in an economy with a very general structure and presents the evolutionary mechanism of behavioral interaction. In Section 5.3 the relation between an investment strategy and the income growth rate is studied. Section 5.4 describes the experimental setup. Section 5.5 provides simulation results and interpretations. Section 5.6 concludes.

5.2 The economic model

5.2.1 General features of the model

Consider a population of agents with the objective to reach a high level of individual welfare, which can only be achieved by a sustained high income growth rate. Each agent can invest its respective income in a finite number of capital sectors. How it allocates its investment over these sectors is expressed by its individual investment strategy. Invested capital is non-malleable: once invested it cannot be transferred between sectors. Standard economic growth and production functions describe how the invested capital accumulates in each sector and contributes to income. These functions are not aggregated: growth and returns are calculated independently for each agent. Two agents with different investment strategies can experience different income growth rates and income levels.

The agents understand that there is a causal link between an investment strategy and economic performance as expressed by the income growth rate, but they cannot use calculus to find an investment strategy that maximizes the income growth rate. Instead, the agents employ the smartest search method that nature has in store, evolution, and they evolve their investment strategies by imitation with variation. Since they prefer a high income growth rate over a low income growth rate, they imitate the investment strategy of a fellow agent when that fellow agent realizes an income growth rate that is high relative to their own income growth rate and that of their other fellow agents. Imitation is not perfect. Changes that are introduced during imitation guarantee diversity in the pool of strategies and keep the evolutionary search alive. In the terminology of evolutionary theory an agent *selects* another agent based on a property (the *phenotype*) that is indicative of its current economic performance and imitates its investment strategy (the *genotype*) with *variation*.

5.2.2 Strategies, investment, and production

All variables and parameters of the economic model are summarized in Table 5.1. The population approach means that accounting of capital investment, production, and income takes place at the level of individual agents. Let $Y_a(t)$ be the income of agent a at time t and let n be the number of available investment sectors. Formally, the investment strategy $s_a(t)$ of an agent can be defined as an n-dimensional vector

$$s_a(t) = [0,1]^n, \quad \sum_i s_{ia}(t) = 1. \tag{5.1}$$

The partial strategy $s_{ia}(t)$—which is the i^{th} element of a strategy—determines the fraction $s_{ia}(t)Y_a(t-1)$ of income that agent a invests in sector i at time t. Each agent must invest its total income in one sector or another, so the partial strategies must be non-negative and sum to one. The set of all possible investment strategies is an $n-1$ dimensional simplex that is embedded in n-dimensional Euclidean space. We call this simplex the *strategy space*.

Capital accumulation in each sector depends on the sector specific investment of each agent and on the global deprecation rate δ. Deprecation is assumed to be equal for

Table 5.1: Variables and parameters of the model

$\|P\|$	population size
k	average number of neighbors per agent
N	neighbors of an agent
$\mathscr{C}$	clustering coefficient of the network
n	number of investment sectors
β	scaling factor of production
δ	discount rate
σ	diversity control parameter
K_{ia}	capital that agent a has accumulated in investment sector i
π_i	production coefficient of investment sector i
s_{ia}	fraction of income that agent a allocates to investment sector i
Y_a	net domestic income of agent a
γ_a	income growth rate of agent a

all sectors and all agents. The dynamic equation for non-aggregate growth per sector is

$$K_{ia}(t) = s_{ia}(t)Y_a(t-1) + (1-\delta)K_{ia}(t-1). \tag{5.2}$$

An extended version of this equation that accounts for dynamic prices can be found in the appendix. To calculate the income $Y_a(t)$ from the capital that agent a has accumulated per sector, we use an n-factor Cobb-Douglas production function with a constant elasticity of substitution,

$$Y_a(t) = \beta \prod_i K_{ia}(t)^{\pi_i(t)}, \tag{5.3}$$

where β is a scaling factor that limits the maximum possible income growth rate. The relative contribution of each sector to production is expressed by a dynamic vector of non-negative production coefficients $\pi(t) = \langle \pi_1(t) \ldots \pi_n(t) \rangle$. To enforce constant returns to scale, all production coefficients are constraint to add up to one,

$$\pi(t) = [0,1]^n, \quad \sum_i \pi_i(t) = 1. \tag{5.4}$$

Similar to the strategy space, the set of all possible vectors of production coefficients is an $n-1$ dimensional simplex that is embedded in n-dimensional Euclidean space.

Production coefficients can depend on an array of economic dynamics, like technological development and environmental dynamics. When the technology or the environment changes, the production coefficients can change as well. Progressive desertification of farm land for example increases the dependency of farmers on irrigation. This can be interpreted as an increase of the production coefficient of irrigation, while some or all of the other production coefficients of the agricultural production process would decrease to compensate. Evolutionary economics raises the question of what happens if the production coefficients change. For this reason we model the environmental dynamics as exogenously defined changes in $\pi(t)$. This is a general approach that can also be applied to other economic dynamics such as technological development. Section 5.4

describes how we implement these changes and how we test whether or not they have an impact on behavioral interactions.

To measure how well a population of agents is adapted to a certain economic environment, we use the expected log income $\mathrm{E}[\log Y(t)]$ of all agents at time t. Expected log income emphasizes an egalitarian distribution of income. Technically speaking, an economic agent with constant relative risk aversion prefers a society with high expected log growth and high expected log income. Let $|P|$ be the total number of agents in the population P. We calculate expected log growth as

$$\mathrm{E}[\log Y(t)] = \frac{1}{|P|} \sum_a \log Y_a(t). \tag{5.5}$$

The individual income growth rate $\gamma_a(t)$ is

$$\gamma_a(t) = \frac{Y_a(t)}{Y_a(t-1)} - 1. \tag{5.6}$$

Expected log income relates to expected log income growth as

$$\mathrm{E}[\log Y(t)] = \sum_{i=1}^{t} \mathrm{E}[\log(\gamma(i)+1)] + \mathrm{E}[\log Y(0)], \tag{5.7}$$

where expected log income growth is defined as

$$\mathrm{E}[\log(\gamma(t)+1)] = \frac{1}{|P|} \sum_a \log(\gamma_a(t)+1). \tag{5.8}$$

5.2.3 The evolutionary mechanism of behavioral interactions

From the point of view of evolutionary modeling, agents and investment strategies are not the same: an agent carries or maintains a strategy, but it can change its strategy and we still consider it to be the same agent (Nowak, 2006). Because every agent has exactly one strategy at a time, the number of active strategies is the same as the number of agents.

To model which agents an agent can imitate we use a generic class of social networks that has been well studied and validated in network theory, namely those that can be generated by a random process with preferential attachment and that have a high clustering coefficient, see Section 4.2.3 on page 71 for details. Before the start of each simulation a stochastic process assigns to each agent a a set of peers N_a that does not change during the course of the simulation. If agent a is a peer of agent b, then a will consider the income growth rate and the investment strategy of b when choosing an agent for imitation, while b will consider the income growth rate and the investment strategy of a. On the other hand, if a and b are not peers, they will not consider each other for the purpose of imitation.

At each time step t an agent may select one of its peers in the social network and imitate its strategy. If that happens, the strategy of the imitating agent changes, while the strategy of the agent that is imitated does not. The choice of which agent to imitate is based on relative welfare as indicated by the current growth rate of income. The imitating agent always selects the peer with the highest current income growth rate. Only if an agent has no peer with an income growth rate higher than its own, the agent does not revise its strategy.

If imitation was the only mechanism by which agents change their strategies, the strategies of agents that form a connected network must converge on a strategy that was present during the initial setup. However, real imitation is never without errors. Errors are called mutations in evolutionary theory. They are fundamental to an evolutionary process because they create and maintain the diversity on which selection can work. In this model we implement mutation by adding some Gaussian noise to the imitation process. That is, when an agent imitates a strategy, it adds some random noise drawn from a Gaussian distribution with zero mean. This causes small mutations along each partial strategy to be more likely than large ones. The exact formula by which agent a imitates and then mutates the strategy of agent b is

$$s_a(t) = s_b(t-1) + \mathcal{N}(0,\sigma), \tag{5.9}$$

where $\mathcal{N}(0,\sigma)$ denotes a normally distributed n-dimensional random vector with zero mean and standard deviation σ per dimension. Because partial investment strategies have to sum to one, we have to enforce $\mathcal{N}(0,\sigma) = 0$, for example by orthogonal projection of the Gaussian noise term onto the simplex, resulting in the loss of one degree of freedom. The error term is further constraint to leave all partial strategies positive. Needless to say that we do not imply that our boundedly rational agents engage consciously in such mathematical exercise. Subjectively they merely allocate their income such that none is left.

The sum of squares of the n partial errors, i.e., the square of the Euclidean distance covered by the error, follows a chi-square distribution with $n-1$ degrees of freedom and mean $(n-1)*\sigma^2$. In equilibrium, when all agents try to imitate the same perfect strategy, the expected standard deviation of the partial strategies will in fact be σ. Since the parameter σ controls the diversity of the investment strategies, we will call it the diversity control parameter, or simply *diversity*. It is the only free parameter of this evolutionary mechanism and has potential policy implications.

5.3 The evolutionary dynamics

5.3.1 The growth rate of a strategy

If we want to understand the impact of environmental dynamics on how agents evolve their strategies we need to understand if and how these environmental dynamics affect which agents are imitated. Whether the strategy of an agent is imitated depends on whether the agent has a higher income growth rate than those agents it is compared with. We call the mapping from investment strategies to income growth rate the *growth function.* The growth function calculates the equilibrium growth rate that an imitating agent realizes if it holds on to a particular investment strategy. If the growth function maps one strategy to a higher equilibrium growth rate than another strategy, then our evolutionary agents will prefer this strategy over the other strategy and imitate it. In this way the growth function indicates which of any two strategies will survive and propagate. Since it depends only on the order of income growth rates—i.e., which of any two agents has a higher income growth rate—whether an agent is imitated, the evolutionary dynamics is invariant under any strictly increasing transformation of the growth function. Any two growth functions that are strictly increasing (or decreasing) transformations of each other lead to the same evolutionary dynamics.

Let us start the derivation of the growth function with an analysis of the equilibrium ratio of sector specific capital to income, $K_{ia}(t)/Y_a(t)$, that will be achieved if an agent holds on to a particular strategy. The dynamic equation of this ratio is

$$\begin{aligned}\frac{K_{ia}(t)}{Y_a(t)} &= \frac{s_{ia}(t)Y_a(t-1) + (1-\delta)K_{ia}(t-1)}{(\gamma_a(t)+1)\;Y_a(t-1)} \\ &= \frac{s_{ia}(t)}{\gamma_a(t)+1} + \frac{1-\delta}{\gamma_a(t)+1}\,\frac{K_{ia}(t-1)}{Y_a(t-1)}.\end{aligned} \tag{5.10}$$

This equation is of the form

$$x(t) = a + bx(t-1), \tag{5.11}$$

which under the condition $0 \le b < 1$ converges monotonically to its unique stable equilibrium at

$$\lim_{t\to\infty} x(t) = a/(1-b).$$

This condition is fulfilled here: investment is always non-negative and sector specific capital cannot decrease faster than δ. With constant returns to scale, income cannot decline faster than capital deprecation, and we have $\gamma_a \ge -\delta$. For the moment, let us exclude the special case $\gamma_a = -\delta$. Then, considering that $0 < \delta \le 1$, we have the required constraint

$$0 \le \frac{1-\delta}{\gamma_a(t)+1} < 1 \tag{5.12}$$

and we conclude that the ratio of capital to income converges to

$$\begin{aligned}\lim_{t\to\infty}\frac{K_{ia}(t)}{Y_a(t)} &= \lim_{t\to\infty}\frac{s_{ia}(t)}{\gamma_a(t)+1}\Big/\left(1-\frac{1-\delta}{\gamma_a(t)+1}\right) \\ &= \lim_{t\to\infty}\frac{s_{ia}(t)}{\gamma_a(t)+\delta}.\end{aligned} \tag{5.13}$$

Equation 5.13 describes a unique stable equilibrium to which the economy of an agent converges monotonically. We ignore the limit notation and combine equation 5.13 with equation 5.3 to calculate income at equilibrium as

$$\begin{aligned}Y_a(t) &= \beta\prod_i\left(\frac{s_{ia}(t)\;Y_a(t)}{\gamma_a(t)+\delta}\right)^{\pi_i(t)} \\ &= \beta\,\frac{Y_a(t)}{\gamma_a(t)+\delta}\prod_i s_{ia}(t)^{\pi_i(t)}.\end{aligned} \tag{5.14}$$

We can now solve for $\gamma_a(t)$ to derive the growth function

$$\gamma_a(t) = \beta\prod_i s_{ia}(t)^{\pi_i(t)} - \delta. \tag{5.15}$$

Let us return to the special case $\gamma_a = -\delta$. According to equation 5.2, capital per sector decreases at the deprecation rate δ only when it receives zero investment, and it cannot decrease faster. This implies that with constant elasticity of substitution, a growth of $\gamma_a = -\delta$ is only possible if every sector with a positive production coefficient receives zero investment. This implies $s_{ia}(t) = 0$ for at least one partial strategy, and so equation 5.15 holds also for the special case $\gamma_a = -\delta$.

5.3.2 Efficiency and level sets of investment strategies

How does the income growth rate of an imitating agent compare to the income growth rate of a rational agent with perfect information? The term $\prod_i s_{ia}(t)^{\pi_i(t)}$ has a single optimum at $s_a(t) = \pi(t)$, allowing a maximum growth of $\gamma^{opt}(t) = \beta \prod_i \pi_i(t)^{\pi_i(t)} - \delta$. This is the income growth rate that a rational agent with perfect information would expect to achieve. Its exact value depends on the location of the production coefficients in the simplex. In an n-factor economy the term $\prod_i \pi_i(t)^{\pi_i(t)}$ varies between a value of $1/n$ in the center of the simplex where all production coefficients are equal, and a value of one in the corners of the simplex where one sector dominates. In order to remove this variability from the growth function, and to allow an easy comparison with the income growth rate of a rational agent with perfect information, we define the efficiency $\mathscr{E}(s,t)$ of a strategy $s(t)$,

$$\mathscr{E}(s,t) = \prod_i \left(\frac{s_i(t)}{\pi_i(t)} \right)^{\pi_i(t)}. \tag{5.16}$$

The efficiency of a strategy measures the fraction $\gamma_a(t)/\gamma^{opt}(t)$ of optimal growth that an agent achieves with this strategy on given production coefficients, assuming that $\delta = 0$. If one strategy leads to a higher equilibrium growth rate than another, it is also more efficient. Efficiency is therefore a monotonic transformation of the growth function that preserves all information on which agent imitates which other agent, removes the variability due to the location of the optimum on the simplex, and allows us to measures growth in terms of what a rational agent with perfect information would achieve.

Efficiency, like the equilibrium growth rate, is a monotonically decreasing function of the Euclidean distance between the strategy and the production coefficients, $|s_a(t) - \pi(t)|$. This function is not symmetric about the optimum but has different slopes in different directions from the optimum. Figure 5.1 shows how the average efficiency of a

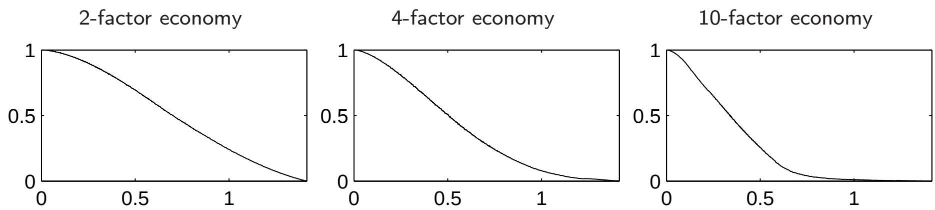

Figure 5.1: Average efficiency as a function of Euclidean distance to the optimum. The x-axis shows the Euclidean distance, the y-axis the corresponding average efficiency. Note the convex shape around the optima.

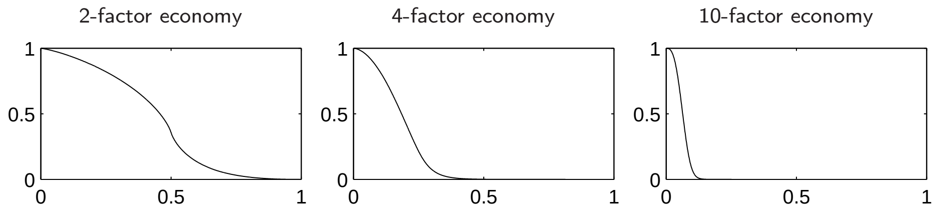

Figure 5.2: Probability of positive growth when the investment strategy and the production coefficients are chosen independently and at random. The x-axis shows the ratio δ/β. The y-axis shows the corresponding probability that the equilibrium growth rate is positive.

strategy decreases as its Euclidean distance to the optimum increases, when strategies and production coefficients are chosen at random from the simplex. Note the inverse S-shape of the graphs. As the Euclidean distance tends to zero, the gradient approaches zero. This implies that when an evolutionary population of agents converges on the optimum, the differences in growth caused by a small amount of diversity σ around the optimal strategy are negligible.

A set of strategies each with identical equilibrium growth rate, say γ', is called a level set and forms a contour hypersurface in the strategy simplex. All strategies that are enveloped by this hypersurface have an equilibrium growth rate that is higher than γ'. This inner set is convex (for a related proof see Beer, 1980) and so from equation 5.15 satisfies

$$\prod_i s_{ia}(t)^{\pi_i(t)} \geq \frac{\gamma' + \delta}{\beta}. \tag{5.17}$$

An important level set is $\prod_i s_{ia}(t)^{\pi_i(t)} > \delta/\beta$. This is the set of all strategies that have a positive equilibrium growth rate. Its size is proportional to $P[\gamma > 0 \mid \pi(t)]$, the probability that a random strategy has a positive equilibrium growth rate with given production coefficients. Let $P[\gamma > 0]$ denote the probability that the equilibrium growth rate is positive if both the strategy and the production coefficients are chosen independently at random from the simplex. Figure 5.2 shows how $P[\gamma > 0]$ decreases as δ/β increases, for economies with respectively 2, 4, and 10 investment sectors. The probability tends to zero as δ/β approaches 1. For given δ/β, the probability that the equilibrium growth rate of a random strategy is positive decreases as the number of investment sectors increases.

The parameters δ and β determine the equilibrium growth rate associated with a given hypersurface, as well as the minimum and maximum equilibrium growth rate that can be achieved with given production coefficients. They do not affect the location of the optimum nor the shape of level sets, both of which depend exclusively on the production coefficients. In other words, δ and β define monotonic transformations of the growth function that are irrelevant to the order of equilibrium growth rates and to the understanding of the evolutionary dynamics as a whole. Also, the rate of convergence in equation 5.13 does not depend on the scaling factor β. We will make use of this fact later on in the experimental design where we use a dynamic β for normalization, significantly reducing the variability of the numeric results.

5.4 Experimental setup

5.4.1 The environmental dynamics

The growth function, up to a monotonic transformation, depends on the coefficients of a Cobb-Douglas type production function. When the production coefficients change with the environmental dynamics, strategies that have previously generated a positive income growth rate can now generate a negative income growth rate. Agents that have converged on a strategy that has previously resulted in a high income growth rate can see their income decline and need to adapt their strategies to the new production coefficients. How does the magnitude and duration of this decline depend on the type of environmental dynamics and on the behavioral interactions among the agents? To answer these questions we model the environmental dynamics as exogenously defined changes in $\pi(t)$. That is, the environmental dynamics that change the production coefficients are the independent variable that the policy maker responds to. The parameters of the imitation mechanism are the dependent variables that the policy maker aims to regulate.

We focus on two aspects of environmental dynamics: how gradual the environment changes, and how frequently. In combination they define four types of environmental dynamics: the production coefficients change gradually and with low frequency, gradually and with high frequency, suddenly and with low frequency, and suddenly and with high frequency. We compare these with two control systems: one without imitation and one with imitation and a static environment. Without imitation, with strategies that are randomly distributed over the strategy space and that stay constant throughout the simulation, the income growth rate of most agents is most likely negative, irrespective of the environmental dynamics. Expected log income will decline and welfare at the population level will be at its lowest. On the other hand, in a static environment where strategies evolve they are expected to converge on the optimum strategy and welfare at the population level will be at its highest.

We consider the general case where a change in the production coefficients is defined as the replacement of one vector of production coefficients by another, with each vector drawn independently and at random from the uniform distribution over the simplex $\sum_i \pi_i(t) = 1$. Replacement is instant for a sudden change and by linear transition for a slow change. A sudden change can be modeled by setting the production coefficients of a 2-factor economy to $\pi = \langle .1, .9\rangle$ up until time t, and to $\pi = \langle .4, .6\rangle$ from $t+1$ onwards. Such extreme changes are characteristic of industries that depend on unreliable resources, e.g., a biotic resource susceptible to climate change like forests or fish. A gradual change can be modeled by changing π from $\langle .1, , 9\rangle$ at time t to $\langle .4, .6\rangle$ at time $t+x$ linearly over x steps, such that

$$\pi(t+j) = \frac{(x-j)\pi(t) + j\pi(t+x)}{x}, \quad 0 \le j \le x, \tag{5.18}$$

where the conditions $\sum_i \pi_i = 1$ and $\pi_i \ge 0$ for all i are fulfilled at all times. Such gradual changes are characteristic of industries that depend on reliable resources, e.g., a mineral resource like iron or coal, where known reserves will typically last for decades if not centuries.

Table 5.2: The environmental dynamics

Environmental dynamics	Observable	Example
gradual, low frequency	reliable resource, Kondrat. wave	oil/gas reserves
sudden, low frequency	unreliable resource, Kondrat. wave	climate change
gradual, high frequency	reliable resource, Juglar's cycle	tech. innovations
sudden, high frequency	unreliable resource, Juglar's cycle	biotic resource

We model low frequency changes by starting the transition from one vector of production coefficients to another vector every 50 years, reflecting a Kondratiev type of wave (Kondratiev, 1935), characteristic of industries that are not a driving force of innovation and change only with the general shift in production methods, e.g., forestry. To model high frequency changes the transition starts every 10 years, corresponding to the fast business cycles observed by Clément Juglar (1863), characteristic of industries that invest heavily in research and development. That is, while we acknowledge that technological innovations are driven by research and development, we treat their effect on the production coefficients as exogenous environmental dynamics that the agents of an industry have to adapt to. We do not claim that the cycles observed by Kondratiev and Juglar are caused by this type of exogenous dynamics. We merely use their observations as examples of frequency patterns that can indeed be detected when present.

We consider different sequences of production coefficients as different instances of the same environmental dynamics as long as the individual vectors of production coefficients are replaced with the same gradualness and frequency. Table 5.2 summarizes the environmental dynamics used for the experiments. Figure 5.3 gives graphic examples of production coefficients that are drawn at random according to the specification of each environmental dynamics. Each row of this figure shows five graphs: one each for the time evolution of the four production coefficients of a 4-factor economy, and one area plot that combines the other four graphs into a single graph, stacking the four individual curves one on top of the other (the upper curve has constant value one), with a different shade of gray for the area under each curve.

With regard to the dependent variable under control of the policy maker, the imitation mechanism has one free parameter, diversity σ, and we specify the optimal behavioral interactions as the diversity $\sigma^{opt}(d)$ that maximizes the expected log income of each agent under given environmental dynamics d,

$$\sigma^{opt}(d) = \underset{\sigma}{\operatorname{argmax}}\, \mathrm{E}(\log Y(t) \,|\, \sigma, d). \tag{5.19}$$

In order to find this optimal value for different environmental dynamics we use repeated numerical simulations with different values of σ and measure the expected log income at the end of each simulation, using standard statistical methods to reduce variance. Having identified the value $\sigma^{opt}(d)$ at which the expected log income is highest under a given environmental dynamics, we proceed to formulate policy advice on the socially optimal level of σ when there is uncertainty over the type of environmental dynamics. To do so we measure the expected log income that an optimal value $\sigma^{opt}(d)$ generates on those environmental dynamics $d' \neq d$ where it is not optimal. We then calculate the value that policy makers with different degrees of risk aversion assign to each $\sigma^{opt}(d)$.

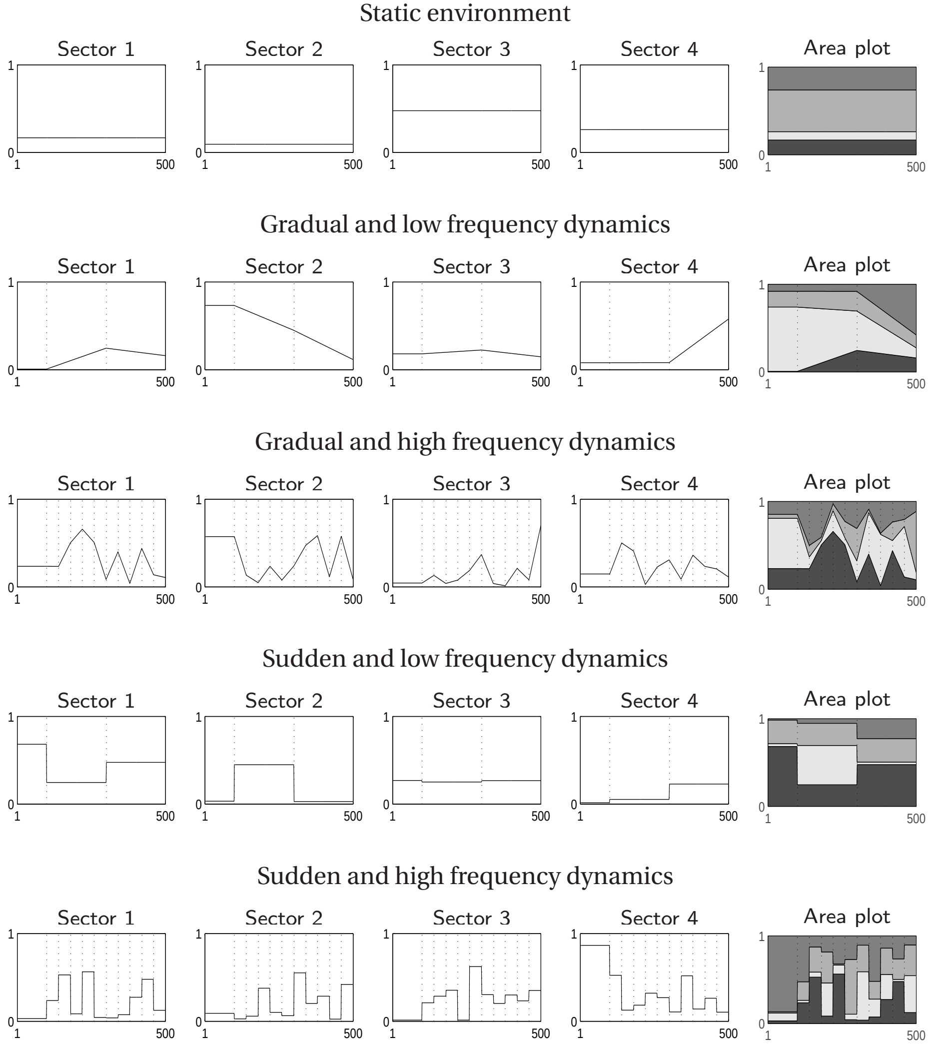

Figure 5.3: Environmental dynamics: changes in the production coefficients. Each row illustrates a different type of environmental dynamics. The sequences of production coefficients are chosen at random. The x-axis shows the 500 time steps (initialization and main experimental phase). The y-axis of the four graphs positioned at the left of each row shows the value of one particular production coefficient in a 4-factor economy. The single graphs at the right are area plots that stack the values of the same four production coefficients one on top of the other, with a different shade of grey under each curve.

5.4.2 Implementation details, model calibration, and scaling

The numerical simulations are based on a discrete synchronous time model where the income and strategy of each agent is updated in parallel at fixed time intervals. We consider each time step t to simulate one financial quarter. As no significant financial market requires a publicly traded company to publish financial results more than 4 times a year, we consider it the limit of feasibility to account for growth and to review an economic strategy as often as 4 times a year. Most economic agents will alter their strategy less often. Each simulation step is divided into two separate update operations: *updating the economy*—each agent invests its income according to its own investment strategy and the individual incomes and growth are calculated by the non-aggregate growth and production and growth functions—and *updating the strategies*, when all agents compare their income growth rate with that of their peer group, and when those agents that decide to imitate change their respective strategies simultaneously.

Each computer simulation spans 500 time steps, simulating 500 financial quarters or 125 years. These are divided into an initialization phase of 100 time steps or 25 years, and a main experimental phase of 400 time steps or 100 years. An initialization phase is needed to avoid influencing the simulation results by an arbitrary choice of initial values. During initialization the simulated economy stabilizes and a "natural" distribution of strategies and growth emerges. Initial conditions are always defined in the same way: all strategies and the initial production coefficients are drawn independently at random from the simplex. The production coefficients are kept static throughout the initialization phase but the agents can imitate in the same way as they do during the main experimental phase, with the same σ. During the 400 time steps (100 years) of the main experimental phase the agents have to adapt to the dynamic changes in the production coefficients. To avoid any initialization effect, the increase in log income is measured from the beginning of the main experimental phase.

Numerical methods are inherently constraint by the availability of computational resources. The computational complexity of multi-agent systems typically scales at least polynomially with system size. The accepted method is to extensively study a system that is large enough to incorporate all the essential ingredients of the model, and to only increase the system size to test whether the obtained results are scalable. Here the main experiments are based on an economy of 200 agents and a 4-factor economy. Sensitivity and scalability are tested with 1,000 agents and with a 10-factor economy. To understand whether the results depend on the specific implementation of the evolutionary mechanism we also test more sophisticated implementations: one version where each agent imitates with probability .1 at every step—as opposed to probability one in the main experiment—and one version where imitation is partial, such that a new strategy is a linear combination of the imitated strategy (with weight .1) and the strategy of the imitating agent (with weight .9). As before, σ controls the standard deviation of the normally distributed errors per partial strategy.

Recall that the rate of capital deprecation δ (equation 5.2) and the scaling factor β (equation 5.3) of the production function have no effect on the evolutionary dynamics and the adaptive behavior of the agents. For the present model we set $\delta = .01$ per time step—about 4% per year—for all sectors. We use a dynamic β to reduce variability of the numeric results. As seen in Section 5.3, different vectors of production coefficients

Table 5.3: Values of economic parameters

population size $\|P\|$	200
investment sectors n	4
duration of the initialization phase	100 time steps
duration of the main experimental phase	400 time steps
capital deprecation δ	.01
dynamically normalized scaling factor $\beta(t)$	$.015\prod_i \pi_i(t)^{-\pi_i(t)}$
average network connectivity	10

have different optimal equilibrium growth rates. To study how efficient the agents adapt to different vectors of production coefficients, we correct for this variability in optimal growth by dynamically normalizing the scaling factor β. To keep the optimal income growth rate at a value of .005 (i.e., an income growth rate of about 2% per year), we let the scaling factor $\beta(t)$ depend on the vector of production coefficients,

$$\beta(t) = .015\prod_i \pi_i(t)^{-\pi_i(t)}. \tag{5.20}$$

With this normalization the equilibrium growth rate of all strategies is constraint to the range $[-.01, .005]$, where the minimum of $-.01$ is realized when $s_{ia}(t)=0$ for some positive π_i and where the maximum of .005 is realized when $s_a(t)=\pi(t)$. Numerical tests show that with these parameter values the probability that a random strategy has a negative equilibrium growth rate on random production coefficients is about .65.

To model which agents an agent can imitate we use a generic class of social networks that has been well studied and validated in network theory, namely those that can be generated by a random process with preferential attachment and that have a high clustering coefficient, see Section 4.2.3 on page 71 for details. Here we use an average connectivity of $k = 10$. In a population of 200 agents this value results in a highly connected network—the average distance between any two agents in the network is 2.7—while maintaining the overall qualities of a complex network.

To improve the general validity of our results we use large number of numerical simulations where—rather than closely calibrating those factors that affect the evolutionary dynamics on a specific economy—we define broad parameter ranges and collect statistical information over a representative sample of different possible economies that fall within these ranges. For example, in order to obtain results that are valid for the general class of scale-free social networks with a high cluster coefficient, each simulation is based on an independent random instance of the social network. Likewise, in order to obtain general results for specific environmental dynamics, each simulation uses an independent random sequences of production coefficients, which are replaced according to the gradualness and frequency of the respective environmental dynamics. The number of simulations needed to obtain reliable statistical results is determined by standard methods of variance reduction. The values of all economic parameters are listed in Table 5.3.

5.5 Results

5.5.1 Economic significance of diversity

Figure 5.4 shows the expected log income of an agent for different values of the diversity control parameter σ, for the two control systems and the four types of environmental dynamics. 50,000 simulations are used for each graph. 500 different values of σ from the range $[0, .5]$ are evaluated, and results are averaged over 100 simulations per value. The plots are smoothed with a moving average with a window size of ten values.

In the first graph of Figure 5.4—the control system without imitation—expected log income is uniformly negative for all levels of σ. This graph is based on a static environment, but the same is observed for any environmental dynamics. All other graphs of Figure 5.4 show systems with imitation and there is a clear functional relation between the value of σ and log income. For each system there is a single optimum $\sigma^{opt}(d)$ that maximizes log income under the given environmental dynamics d, exact values are given in Table 5.4. The value of $\sigma^{opt}(d)$ is higher for more frequent changes than for less frequent changes, and higher for sudden changes than for gradual changes. Its value is lowest in the static environment. Further to this, the graphs show a clear pattern in the relationship between σ and expected log income: the slope to the left of the optima, i.e., for small values of σ, is much steeper than to the right, where σ is large. We will revisit this fact in our discussion of policy advice under uncertainty.

Our first research question can now be answered: almost any level of diversity σ will allow the evolutionary agents to reach a positive log income under any environmental dynamics, yet a unique optimum where log income is highest can be identified for each environmental dynamics. So while it is not mandatory to define policies that effect σ, in the sense that imitating agents can almost always return to positive growth, it is optimal in the sense that there can be a significant gain in expected log income.

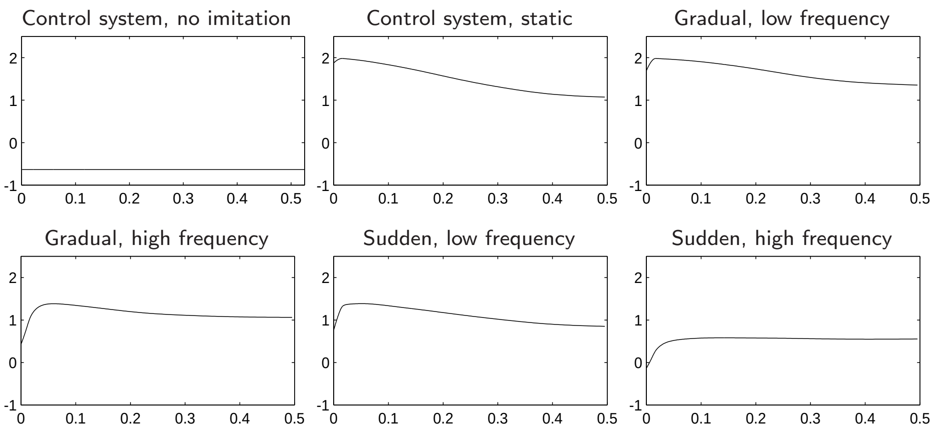

Figure 5.4: Expected log income as a function of the diversity parameter σ. The x-axes show the diversity σ, the y-axes the resulting log income.

Table 5.4: Optimal level of diversity σ for each environmental dynamics

Environmental dynamics	Optimal σ
static	.005
gradual, low frequency	.008
sudden, low frequency	.046
gradual, high frequency	.057
sudden, high frequency	.123

5.5.2 Policy advice under uncertainty

When the type of environmental dynamics that affects an economy is unknown, the optimal diversity σ depends on the risk preference of the policy maker. We consider four types of risk preference: extreme risk seeking (maxmax), modest risk seeking (max average), modest risk aversion (minimax regret), and extreme risk aversion (maxmin). We calculate the optimal value for each preference from a generalization table where each σ that is optimal under one environmental dynamics is applied to the other tested environmental dynamics, including the static environment. The result is shown in Table 5.5. Each row shows the expected log income for the same environmental dynamics but different σ, each column shows results for the same σ but different environmental dynamics. Each entry is averaged over 10,000 simulations (with different instances of the social network and different random sequences of production coefficients). The values in the diagonal are highest for each row, confirming that the optimal value is indeed the best choice for a given environmental dynamics.

For each risk preference, each tested σ can now be associated with an expected value, and the σ with the best such value is considered optimal for the type of risk preference. This is shown in Table 5.6. Each row shows the value associated with each σ under a given risk adversity, with the optimal value in Italic type. A risk seeker looks at the highest expected log income that each tested σ has achieved under the different environmental dynamics, and chooses the highest of these. Risk neutrality means choosing the σ that maximizes the average expected log income over all environmental dynamics. Under minimal regret the σ is chosen that minimizes the greatest possible difference between actual log income and the best log income that could have been achieved. Minimal regret first calculates the maximum possible regret for each σ and all environments, and then chooses the σ that minimizes this maximum. Risk aversion means choosing the σ that promises the highest minimum log income under any environmental dynamics.

The numerical results clearly show that under uncertainty the optimal value of σ rises with the degree of risk aversion. This is in line with our earlier observation about the functional relationship between σ and expected log income: the gradient is steeper for lower values of σ than for higher values, which makes higher values of σ the safer bet. These observations are confirmed by the control experiments that test for sensitivity and scalability and that use alternative implementations of the imitation mechanism.

Table 5.5: Expected log income when a value of σ that is optimal under one environmental dynamics is applied to other environmental dynamics

Environmental dynamics that σ is applied to	Environmental dynamics that σ was optimized for: static	gradual, low freq.	sudden, low freq.	gradual, high freq.	sudden, high freq.
Optimal σ	.005	.008	.046	.057	.123
static	1.996	1.995	1.952	1.934	1.794
gradual, low frequency	1.988	1.991	1.969	1.959	1.880
sudden, low frequency	1.117	1.250	1.412	1.406	1.331
gradual, high frequency	0.550	0.721	1.377	1.385	1.308
sudden, high frequency	-0.042	0.092	0.529	0.548	0.596

Notes. Each column shows the expected log income when the diversity σ that is optimal for one type of environmental dynamics is applied to another type of environmental dynamics. Each row shows the expected log income when agents adapt to a specific environmental dynamics with a value σ that is optimal for another dynamics.

Table 5.6: Optimal policy advice under uncertainty and different degrees of risk aversion

Type of policy maker or society	Environmental dynamics that σ was optimized for: static	gradual, low freq.	sudden, low freq.	gradual, high freq.	sudden, high freq.
Optimal σ	.005	.008	.046	.057	.123
risk seeking	*1.996*	1.995	1.969	1.959	1.880
risk neutral	1.122	1.210	*1.448*	1.446	1.382
minimal regret	0.835	0.664	0.067	*0.062*	0.202
risk averse	-0.042	0.092	0.529	0.548	*0.596*

Notes. Each entry is calculated from a column of Table 5.5 and shows the value that a risk preference assigns to a particular diversity σ. Each row shows the value that is optimal under that risk preference in Italic type (minimum for minimal regret, maximum for the others).

We also tested more sophisticated imitation mechanisms where either only a random selection of 10% of all agents would imitate per step, or where imitation was partial, such that a new strategy is a linear combination of the imitated strategy and the strategy of the imitating agent (again with normally distributed errors per partial strategy). We also made the selection process—the choice of which agent to imitate—dependend on income instead of growth. The arrangement of optimal values $\sigma^{opt}(d)$ is similar for each environmental dynamics. The gradient is always steeper for small values of σ than for large ones. Under uncertainty the optimal value of σ always increases with the degree of risk aversion.

5.5.3 Evolutionary dynamics

Figure 5.5–5.10 show the evolution of some key statistics during the 500 times steps of the simulation. Figure 5.5 discusses the control systems without imitation (a static environment is used). Figure 5.6 discusses the control system with imitation and a static environment. The remaining four figures show the four types of environmental dynamics where agents imitate and where changes occur gradually and with low frequency (Figure 5.7), gradually and with high frequency (Figure 5.8), suddenly and with low frequency (Figure 5.9), and suddenly and with high frequency (Figure 5.10).

Each figure contains six statistics that describe the economic performance of the agent population, the heterogeneity of their strategies, and the relevance of connectivity in the social network at each of the 500 time steps of the simulation. These statistics are averaged over 10,000 simulations. Each simulation uses a different instance of the social network and a different sequence of vectors of production coefficients. Two area plots on top of each group of six illustrate how the production coefficients change under the respective environmental dynamics. Each area plot shows a single different random sequence of production coefficients. To further ease the analysis, in Figure 5.7–5.10 dotted vertical lines are inserted into each area plot and each of the six statistics to show the points in time where the transition to a new set of production coefficients starts.

All statistics react visibly to any change in the production coefficients. This is particularly interesting for those environmental dynamics where change occurs gradually, because when a new vector of random production coefficients is introduced only the momentum changes, not the rate of change. And yet there is a clear and strong economic response to this change in momentum. Note that many statistics have reached some sort of equilibrium after the 100 steps of the initialization phase.

Of the six statistics, the first four visualize the economic performance of the agents. The first statistic shows the average efficiency (see equation 5.16) over all strategies, allowing a direct comparison to what rational agents with perfect information would achieve. The second statistic shows the behavior of the Gini coefficient, a measure of how egalitarian the accumulated capital is distributed. The third statistic shows average log income, which generally behaves as expected: after each change the income level drops temporarily, only to grow continuously thereafter. The fourth statistic shows average log growth, which falls dramatically immediately after a change, as most strategies become obsolete, but peaks within no less than ten time steps after the change, indicating that the recovery process of our evolutionary economy starts almost immediately after a change. The fifth statistic measures the variance of partial strategies within the population and shows how the heterogeneity of strategies is affected by a change in production coefficients. As discussed in Section 5.2.3, at equilibrium the square root of this variance approaches the value of the diversity control parameter σ. The sixth and final statistic measures the covariance between log income and connectivity, to emphasize the effect of a skewed distribution of connectivity on the evolutionary process. It shows how the correlation between log income and network connectivity rises each time that a new change in the production coefficients is initiated. Evidently the highly connected agents are among the first to learn and profit from the improved strategies.

When read in combination with statistic one through four on economic performance, the fifth statistic on the variance of partial strategies allows us to identify the different phases of the adaptive evolutionary search process. After each onset of a new environmental change the standard deviation of partial strategies peaks for a brief period, then drops rapidly, and finally returns slowly to its former state. This shows that in the immediate aftermath of a change those agents that were previously most successful loose their attractive power—average efficiency is at its lowest—and that the diversity of the pool of strategies increases significantly. This is the early phase of unstructured exploration. As the agents evaluate new strategies, some agents are more successful than others—average efficiency rises again—and get heavily imitated, leading to a rapid decline in diversity. This is the second phase of structured, directed search. What is remarkable is that during this second phase the average efficiency declines for a second time and reaches a low point between ten to twenty time steps after the onset of the environmental change. Finally, during the last phase of exploration, the agents seem to finally settle into the new order. Average efficiency increases again, and as more and more agents approach the (moving) optimum, they diversify around it.

5.6 Conclusions

We have studied the general question of how different types of environmental dynamics affect behavioral interaction in an evolutionary economy. For this purpose a simple model of evolutionary formation of investment strategies through variation and selection was presented. Variation occurs when an agent replaces its own strategy by that of another agent (imitation) in an imperfect way. Selection occurs when an agent bases its choice to imitate another agent on some property of the other agent, here individual income growth. The evolutionary mechanism has one free parameter that controls diversity by defining how closely agents imitate each other. This parameter has a clear policy dimension as there are various laws and regulations that regulate how closely agents imitate each other.

If agents in an economy with a Cobb-Douglas type production function use relative income growth rate to determine which agent to imitate, the evolutionary dynamics are governed by the equilibrium growth rate of a strategy. This equilibrium growth rate is uniquely determined by the production coefficients of the Cobb-Douglas function. Modeling environmental dynamics as dynamic changes in these production coefficients enables us to study the impact of such environmental dynamics on the optimal behavioral interactions. This is a general approach that can be applied to model technological or macroeconomic dynamics as well as environmental hazards like (climate change induced) desertification and diseases.

We specified four different types of environmental dynamics that differ in the gradualness and frequency of change. We further specified one control system without imitation and one control system with imitation and a static environment. To achieve general results that are valid for a broad class of economies, all numerical results were based on large number of computer simulations, each with different instances of those factors that affect the evolutionary dynamic.

Our first research question was whether or not different values for the diversity control parameter are optimal under different environmental dynamics. We established that for almost all tested values of this parameter and all tested environments the agents quickly adapt and find strategies with a positive equilibrium growth rate. We found further that each environmental dynamics has a unique optimal diversity that maximizes log income. This optimum increases with the frequency and rapidity of the changes.

Our second research question was whether or not different policies can be defined for different degrees of risk aversion when the type of environmental dynamics is unknown. Here we found that if there is uncertainty about the environmental dynamics, the optimal value for σ increases with the degree of risk aversion. The generality of our findings were confirmed by control experiments that tested for scalability and sensitivity of the economic parameters and that used alternative implementations of the imitation mechanism.

Various types of public policies directly and indirectly affect the imitation behavior of economic agents, the diversity of their investment strategies, and their ability to adapt to a changing environment. Numeric simulations of stochastic multi-agent systems can be used to evaluate such policies even when there is uncertainty on the specific nature of the environmental dynamics. Despite, or rather because of, their stochastic nature they can identify the preferred policy under a particular degree of risk aversion.

Appendix 5.A Evolution with variable prices

The interested reader will be curious to know how variable prices affect the evolutionary process. Regardless of the market structure and price formation, equation 5.2 for non-aggregate growth per investment sector i can be extended to include a dynamic price $p_i(t)$,

$$K_{ia}(t) = \frac{s_{ia}(t)Y_a(t-1)}{p_i(t)} + (1-\delta)K_{ia}(t-1). \tag{5.21}$$

The ratio of capital to income (equation 5.13) now converges to

$$\lim_{t\to\infty} \frac{K_{ia}(t)}{Y_a(t)} = \lim_{t\to\infty} \frac{s_{ia}(t)/p_i(t)}{\gamma_a(t)+\delta}. \tag{5.22}$$

The existence of this limit and the speed of convergence depend on the behavior of $p_i(t)$. If the price converges, the ratio of capital to income will converge as well. In that case the growth rate at equilibrium is

$$\gamma_a(t) = \beta \prod_i p_i(t)^{-\pi_i(t)} \prod_i s_{ia}(t)^{\pi_i(t)} - \delta. \tag{5.23}$$

That is, as long as the market structure does not prevent the capital-income ratio to converge in reasonable time, variable prices have a similar effect on the evolutionary process as the scaling factor. Both are monotonic transformations of the growth function that do not affect the evolutionary dynamics.

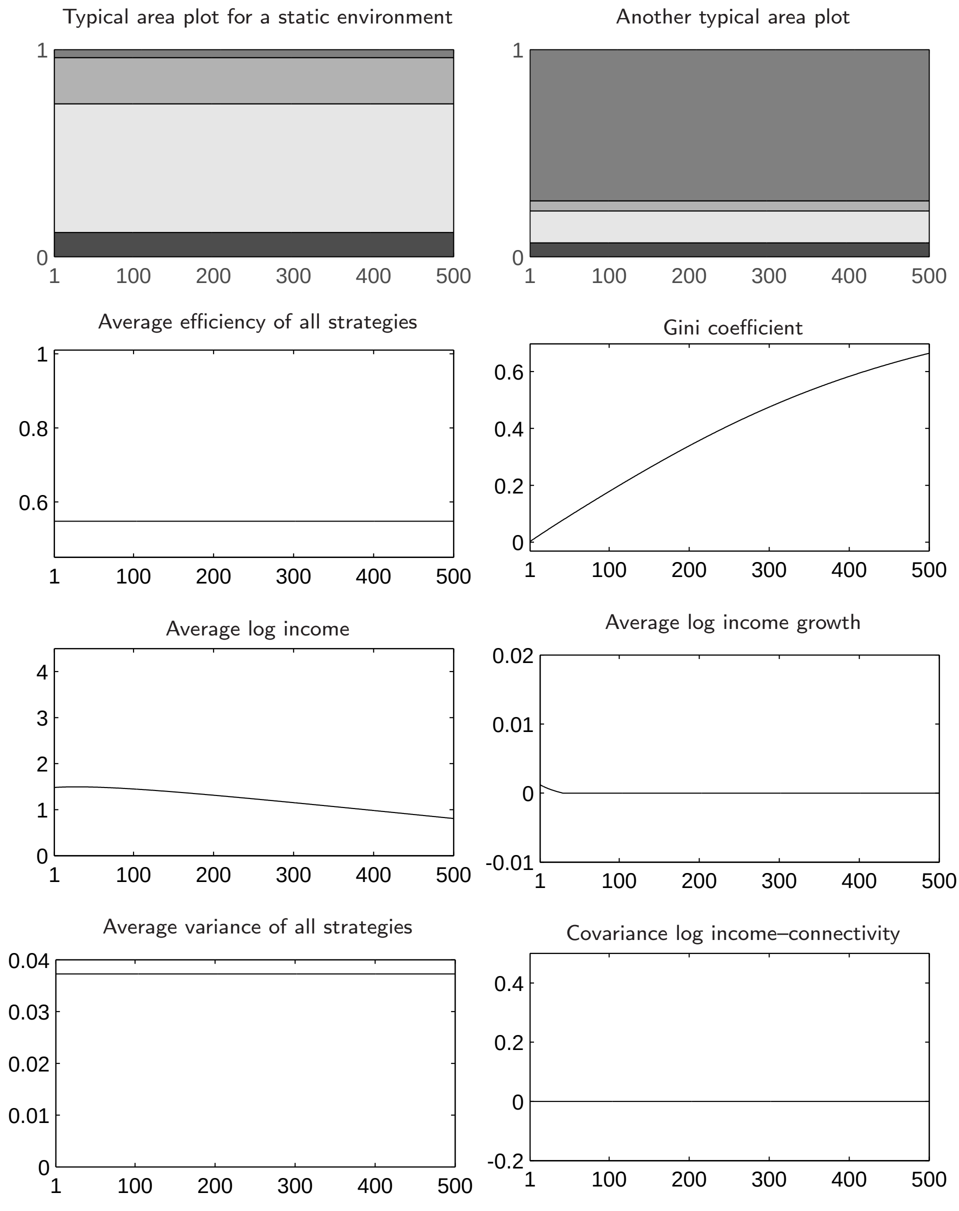

Figure 5.5: Average time evolution of an economy *without imitation* (the environment is static).

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x-axis shows the 500 time steps while the y-axis shows the respective statistics.

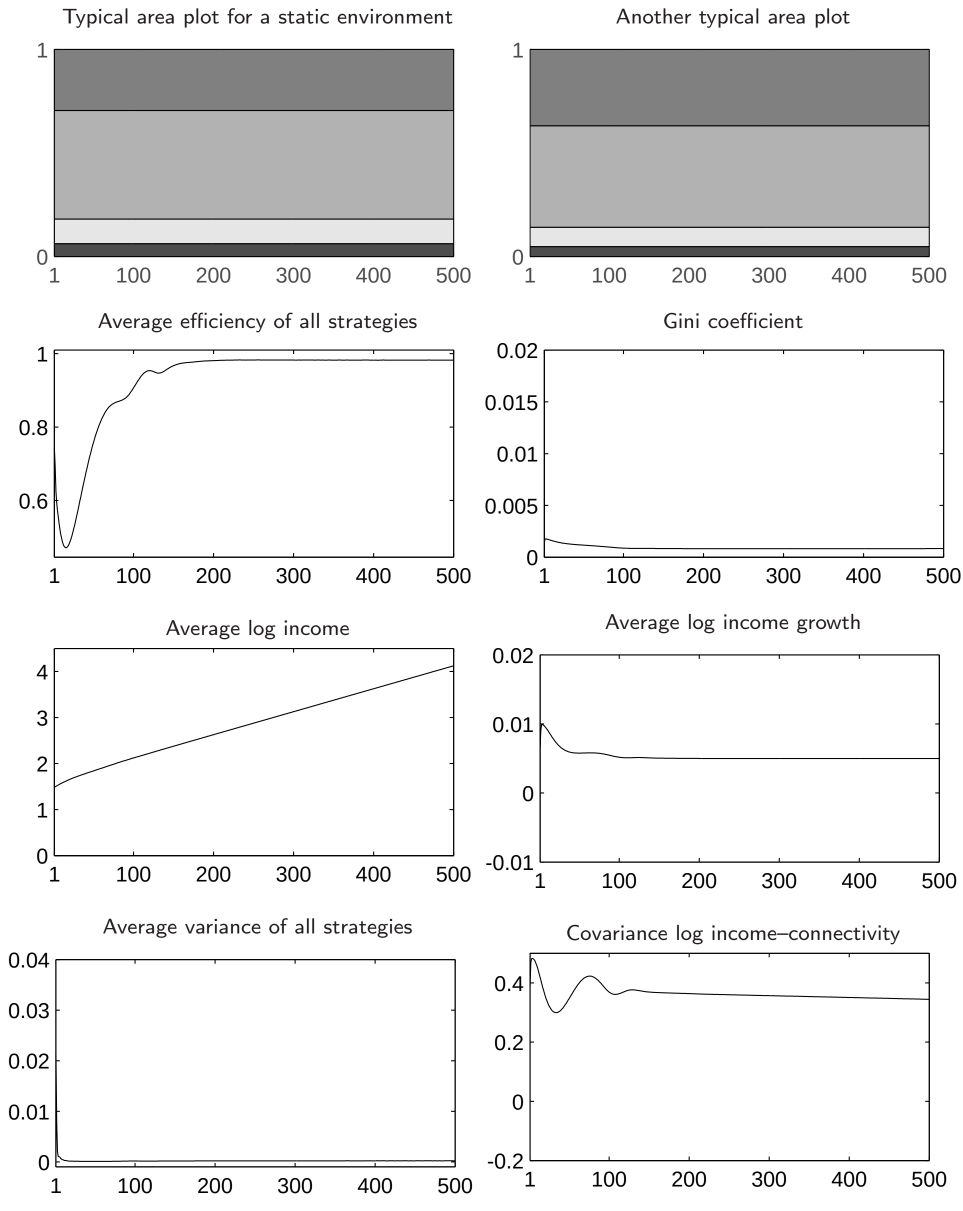

Figure 5.6: Average time evolution of an economy with imitation and a *static environment.*

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x-axis shows the 500 time steps while the y-axis shows the respective statistics. Note how most statistics have stabilized during the first 100 steps of the initialization phase.

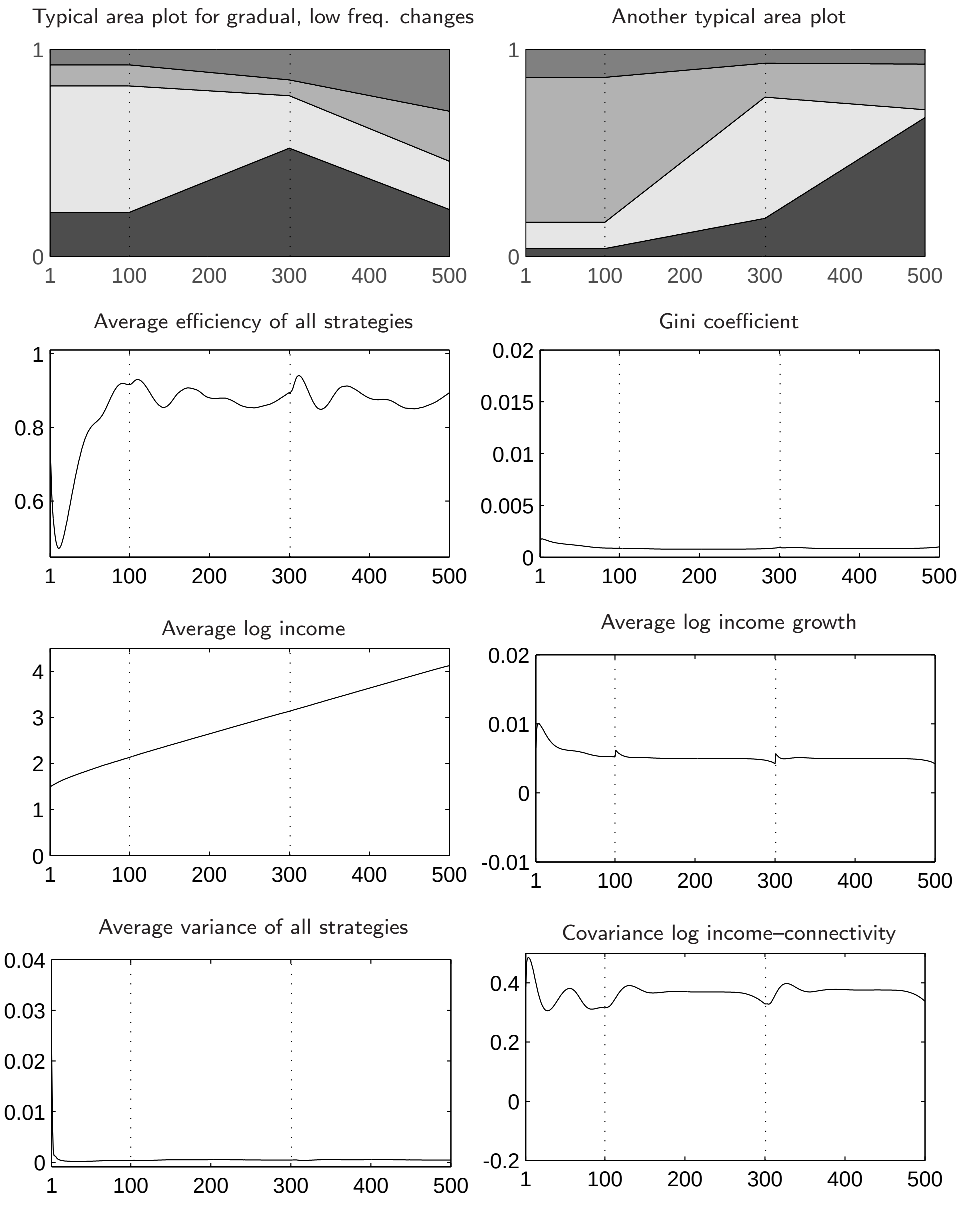

Figure 5.7: Average time evolution of an economy with imitation and a dynamic environment characterized by *gradual, low frequency* changes.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x-axis shows the 500 time steps while the y-axis shows the respective statistics.

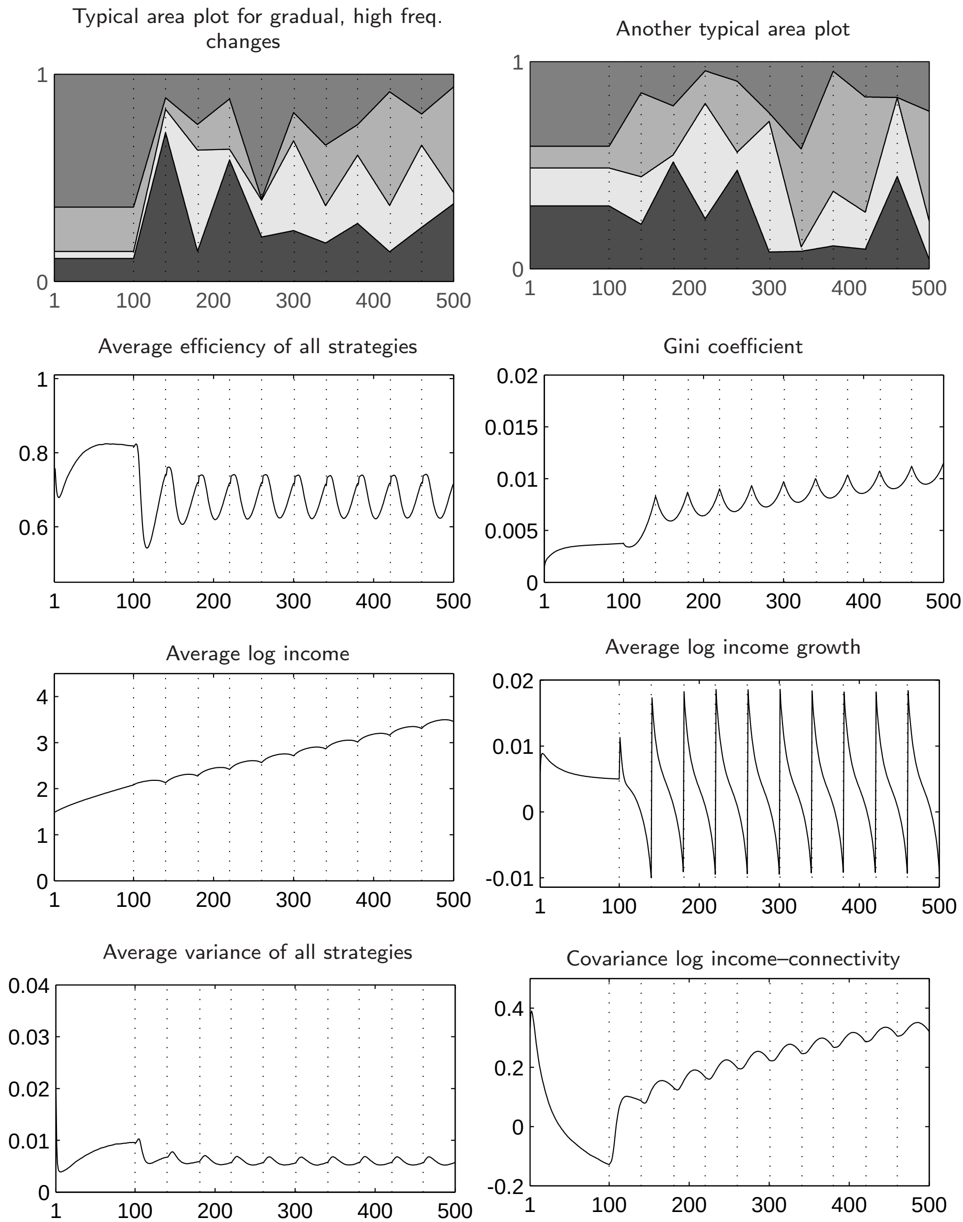

Figure 5.8: Average time evolution of an economy with imitation and a dynamic environment characterized by *gradual, high frequency* changes.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x-axis shows the 500 time steps while the y-axis shows the respective statistics.

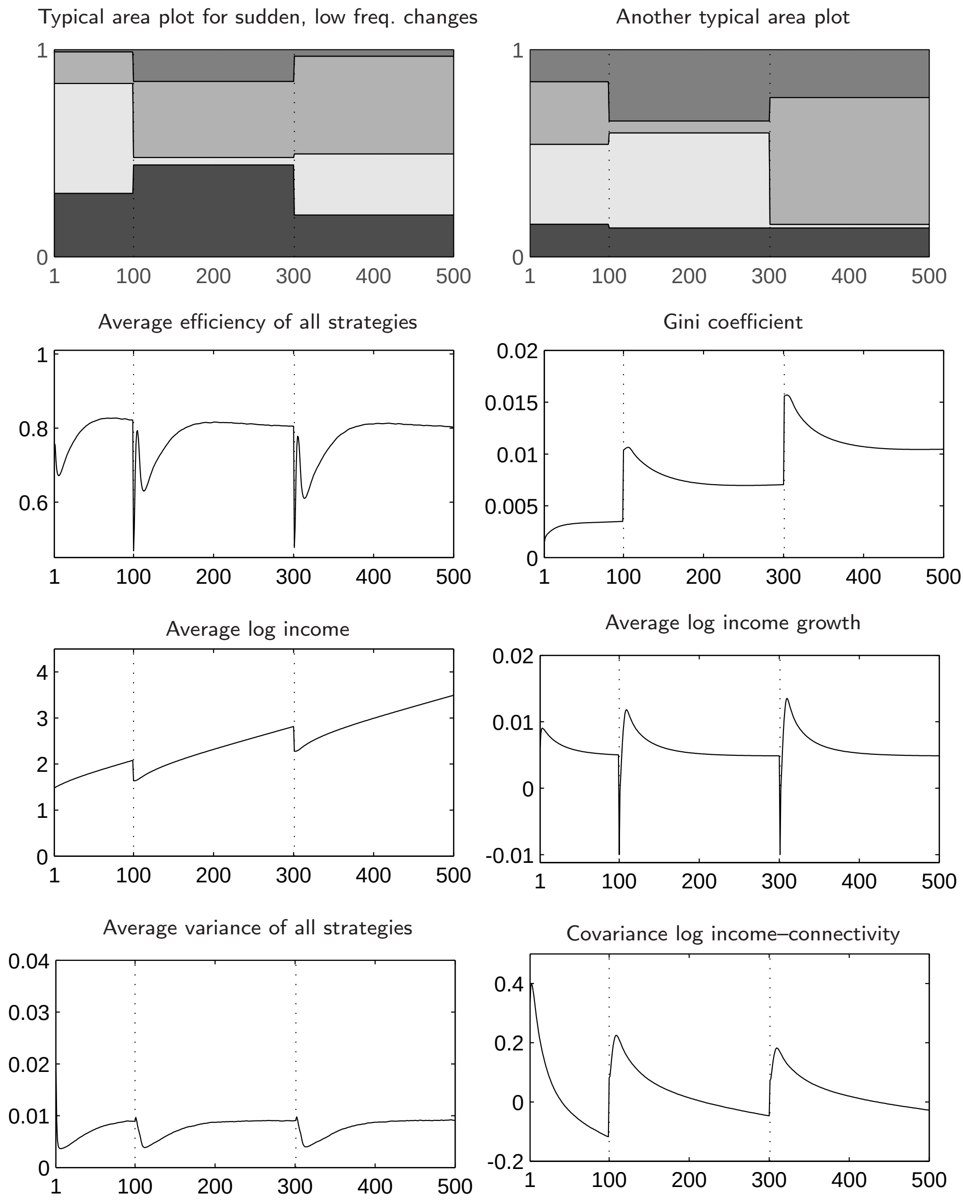

Figure 5.9: Average time evolution of an economy with imitation and a dynamic environment characterized by *sudden, low frequency* changes.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x-axis shows the 500 time steps while the y-axis shows the respective statistics.

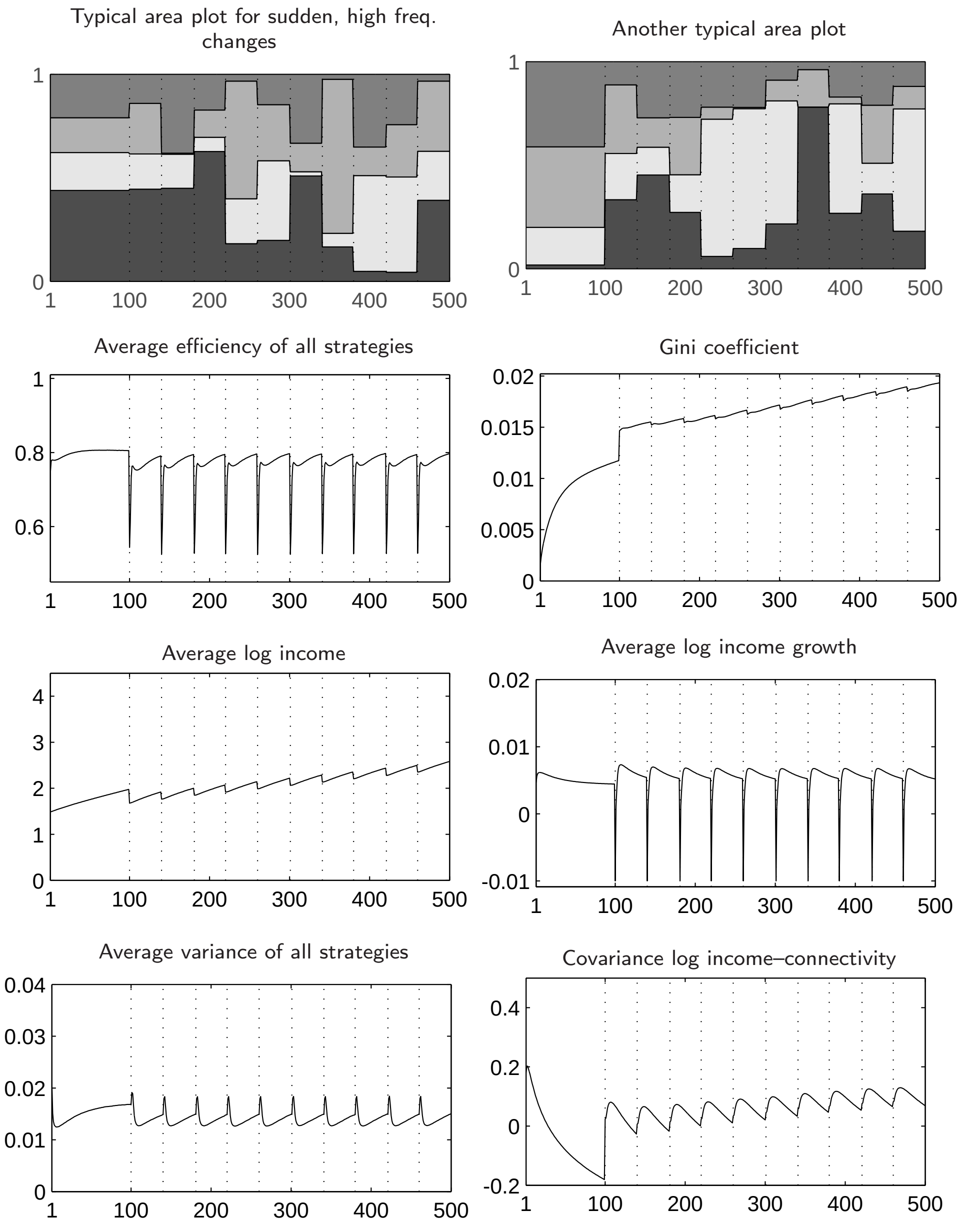

Figure 5.10: Average time evolution of an economy with imitation and a dynamic environment characterized by *sudden, high frequency* changes.

The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x-axis shows the 500 time steps while the y-axis shows the respective statistics.

REFERENCES

Beer, G., 1980. The Cobb-Douglas Production Function. Math. Mag. 53 (1), 44–48.

Bergstrom, T. C., 2002. Evolution of Social Behavior: Individual and Group Selection. J. Econ. Perspect. 16 (2), 67–88.

Binmore, K., 1994. Game Theory and the Social Contract. Vol.1 - Playing Fair, and Vol.2 - Just Playing. MIT Press, Cambridge, MA.

Eiben, A. E., Smith, J. E., 2003. Introduction to Evolutionary Computing. Springer, Berlin / Heidelberg.

Epstein, J. M., Axtell, R., 1996. Growing Artificial Societies: Social Science from the Bottom Up. MIT Press, Cambridge, MA.

Galor, O., Moav, O., 2002. Natural selection and the origin of economic growth. Q. J. Econ. 117 (4), 1133–1191.

Henrich, J., 2004. Cultural group selection, coevolutionary processes and large-scale cooperation. J. Econ. Behav. Organ. 53 (1), 3–35.

Juglar, C., 1863. Crises Commerciales. Imprimerie de Veuve Berger-Levrault, Strasbourg.

Kondratiev, N. D., 1935. The Long Waves in Economic Life. Rev. Econ. Stat. 17 (6), 105–115.

Mirowski, P., 2007. Markets come to bits: Evolution, computation and markomata in economic science. J. Econ. Behav. Organ. 63 (2), 209–242.

Munro, A., 1997. Economics and biological evolution. Environ. Resour. Econ. 9 (4), 429–449.

Nannen, V., van den Bergh, J. C. J. M., Eiben, A. E., 2013. Impact of environmental dynamics on economic evolution: A stylized agent-based policy analysis. Technol. Forecast. Soc. Change 80 (2), 329–350.

Noailly, J., 2008. Coevolution of economic and ecological systems. J. Evol. Econ. 18 (1), 1–29.

Noailly, J., van den Bergh, J. C. J. M., Withagen, C. A., 2003. Evolution of harvesting strategies: replicator and resource dynamics. J. Evol. Econ. 13 (2), 183–200.

Nowak, M. A., 2006. Evolutionary Dynamics: Exploring the Equations of Life. Harvard University Press, Cambridge, MA.

Ostrom, E., 2000. Collective Action and the Evolution of Social Norms. J. Econ. Perspect. 14 (3), 137–158.

Sethi, R., Somanathan, E., 1996. The Evolution of Social Norms in Common Property Resource Use. Amer. Econ. Rev. 86 (4), 766–788.

Tesfatsion, L., 2006. Agent-based computational economics: A constructive approach to economic theory. In: Judd, K. L. (Ed.), Handbook of Computational Economics. Elsevier, pp. 831–880.

van den Bergh, J. C. J. M., 2004. Evolutionary modelling. In: Proops, J., Safonov, P. (Eds.), Modelling in Ecological Economics. Edward Elgar Publishing, Cheltenham, pp. 9–35.

van den Bergh, J. C. J. M., 2007. Evolutionary thinking in environmental economics. J. Evol. Econ. 17 (5), 521–549.