

Impact of environmental dynamics on economic evolution: A stylized agent-based policy analysis

Volker Nannen*

*IFISC, Institute for Cross-Disciplinary Physics and Complex Systems, Universitat de les Illes Balears, 07122 Palma, Palma de Mallorca, Spain
Departament d'Arquitectura i Tecnologia de Computadors, Universitat de Girona, Edifici P-IV, Campus de Montilivi, 17071 Girona, Spain*

Jeroen C.J.M. van den Bergh

*ICREA, Barcelona, Spain
Institute of Environmental Science and Technology, Universitat Autònoma de Barcelona, Edifici Cn, 08193, Bellaterra, Spain
Dept. of Economics and Economic History, Universitat Autònoma de Barcelona, Edifici Cn, 08193, Bellaterra, Spain
Faculty of Economics and Business Administration, VU University Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, the Netherlands
Institute for Environmental Studies, VU University Amsterdam, the Netherlands*

A.E. Eiben

Department of Computer Science, Faculty of Sciences, VU University Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam, the Netherlands

Abstract

The general problem 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 these 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 frequently the environment changes. Numerical analysis shows how the socially optimal level of diversity increases with the frequency and speed of environmental change. When there is uncertainty about the specific type of environmental dynamics—whether for lack of data or because it is not constant—the socially optimal level of diversity increases with the degree of risk aversion of the policy maker or the society.

Keywords: agent-based economics, dynamic environments, evolutionary economics, optimal diversity, public policies, risk aversion

1. Introduction

1.1. Economic evolution in dynamic environments

Evolutionary reasoning and agent-based modeling have become standard practice in various disciplines, including social sciences [e.g., 1, 2]. An 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 issues arise. Here we aim to address one such issue, namely, how environmental dynamics bears upon the behavioral interaction of an evolutionary economy consisting of multiple agents with heterogeneous economic

* Corresponding author at: Departament d'Arquitectura i Tecnologia de Computadors, Universitat de Girona, Edifici P-IV, Campus de Montilivi, 17071 Girona, Spain.

Email addresses: vnannen@gmail.com (Volker Nannen), jeroen.bergh@uab.es (Jeroen C.J.M. van den Bergh), gusz@few.vu.nl (A.E. Eiben)

strategies. Specifically, we will study whether different types of environmental dynamics call for distinct behavioral interaction, and how a policy maker can make use of this insight even when the current type of environmental dynamics is unknown. The relevance of these questions is evident: few economic environments are static, and their dynamics are often irregular and unpredictable.

Generally, one cannot expect evolution in a changing environment to approach a steady state. Often, what matters is not 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. When heterogeneous groups of users, polluters, or harvesting strategies are involved, a reduction to a representative agent model is often not possible. Instead, the heterogeneity and interaction of agents and strategies is better treated by an evolutionary agent-based model [3, 4].

With regard to the evolutionary system, there is a range of theoretical starting points and modeling approaches [5, 6]. First of all, one can choose to use very theoretical, abstract models of the evolutionary game type. In evolutionary game theory, studies of social behavior have been mostly limited to constant environments, letting selection pressure depend on the population distribution. This simplification allows for analytical treatment. However, addition of a dynamic environment typically leads to a system that is no longer amenable to analytic solutions. Numeric simulations of multi-agent systems form an alternative to the analytic approach that offers 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 [7, 8]. 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).

For our purpose a relevant distinction is between external and internal environmental dynamics. 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. Yaari et al. [9] use local positive feedback loops to model growth in post-liberalized Eastern Europe. Implications of such a model to policy control are discussed by Challet et al. [10]. Ietri and Lamieri [11] and Windrum and Birchenhall [12] model coevolution of consumers and firms to study product innovation, technological successions, and the emergence of business cycles. Resource dynamics [e.g., 13, 14] and dynamic control of a pest population that evolves resistance to pesticides [15] are policy-relevant examples. Another, general example is a coevolutionary system in which two heterogeneous populations cause selection pressure on one another [16]. 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 internal environmental dynamics [17]. However, coevolution of multiple, interacting evolving populations easily leads to intractable models, and therefore is not a logical starting point for the present analysis (it may be addressed in a subsequent paper). We want to first know how a single population evolves under various types of environmental dynamics.

In this article we investigate the impact of different types of general external environmental dynamics on the socially optimal type of behavioral interactions among the agents in the population. For this purpose we develop a stylized, abstract evolutionary agent-based model that would best describe a community of small and medium-sized enterprises which serve similar markets with similar products, and which are characterized more by cooperation and an open exchange of information than by competition and secrecy. We identify the key control variable and establish its optimal level under different environmental dynamics and for different degrees of risk aversion of the policy maker or the society. While our results are solid insofar as they are robust over a wide range of implementation choices and parameter values, specific application would require adaptation of the model, for example, by including competition, niches, savings behavior, or coevolution of multiple populations.

Our use of evolutionary agent-based simulations for policy analysis in a dynamic environment departs somewhat from the main body of literature in evolutionary economics. This literature, from Schumpeter [18] and Nelson and Winter [19] to Dosi [20] and Silverberg et al. [21], develops an evolutionary theory that explains the emergence of complex empirical phenomena like technological progress and innovation diffusion from the interaction of simple agents. By contrast, the present approach was inspired by a quantitative approach to evolutionary theory that focuses on system control, as developed for example in the vast literature on the evolution of drug resistance in pathological agents and how to prevent it, or the mathematical models on cancer and HIV of Nowak [22].

Within economics, a quantitative approach with focus on control is more typical of the management literature. Here, the problem of how firms have to diversify and reinvent themselves in the face of various types of environmental dynamics is ubiquitous. While normally not phrased in evolutionary terms, the reader will agree that the essential elements of an evolutionary process—a population that undergoes variation and selection—are usually well defined. For example, Ansoff [23] and Ansoff et al. [24] conclude that a firm needs to diversify its portfolio of methods and products in order to stay profitable in a changing environment, and that the level of diversification—ranging in qualitative terms from “incremental” to “creative”—needs to match the degree of changeability and predictability of the environmental dynamics. Thore et al. [25, 26] use data envelope analysis (DEA) to track the time path of the empirical production frontier for the U.S. computer industry. They observe that individual firms left the efficiency frontier and returned to it in accordance with the life cycles of their products and technologies. Building on this, Phillips and Tuladhar [27] show that the number of years a computer manufacturer stayed at the efficiency frontier corresponds to the amount by which it could adjust output levels to environmental stimuli. The authors proceed to link the optimal level of output flexibility to the variety of environmental stimuli, in accordance with the law of requisite variety [28].

1.2. *A model of environmental change*

Economic change in general and business cycles in particular are often characterized by sustained co-movement of several variables and sectors. For example, Hornstein and Praschnik [29] study co-movement between production sectors, Croux et al. [30] study co-movement in the output of economic regions, and Forbes and Rigobon [31] study co-movement in stocks. We take account of this empirical stylized fact by defining economic dynamics as co-movement of two or more production coefficients of a non-aggregate Cobb-Douglas type multi-sector economy. The production coefficients can depend on an array of economic dynamics, like technological development and resource dynamics. When the technology or the environment changes, the production coefficients can change as well.

Water management provides a specific example of a production input that is highly dependent on environmental conditions, available technology and the specific requirements of currently profitable crops. An example of global scale is the progressive desertification of farm land, which increases the dependency of farmers on irrigation. Specific local examples are provided by the Netherlands and South India. Since more than a millennium the already high Dutch water table has been rising even further due to soil compaction and the reduction of swamp areas that buffer excess surface water [32]. In response, the Dutch have learned to control the water table with a complex system of dikes, channels and water pumps of increasing cost and sophistication. South India is largely semi-arid. Most of its large rivers and tributaries are non-perennial. Originally, crops that depend entirely on rain fall and stored moisture in the soil were dominant. However, during the last millennium, and particularly during the Vijayanagara period (14th to 17th century), consumer preferences for wet rice and other irrigated crops stimulated heavy investment in a complex network of reservoirs, ponds, and irrigation canals [33]. These examples can be modeled as a gradual increase, possibly from zero, of the production coefficient of water management. Increased investment in water management necessarily reduces the relative importance of all other inputs, i.e., lowering the absolute value of their coefficient after normalization.

We are not aware of any empirical work that systematically studies how the Cobb-Douglas production coefficients of water management and other agricultural inputs change over time, and how well this change is reflected in the investment behavior of the producers. Empirical studies that track production frontiers and producer efficiency over time typically use non-parametric methods, not the least because a Cobb-Douglas type model allows for only one output. However, a number of studies exist that show significant regional differences both in the Cobb-Douglas production coefficients of irrigation and other agricultural inputs like seeds and fertilizers, and in how efficiently producers allocate these resources. In the case of maize production, Hasan [34] reports that the Cobb-Douglas production coefficient for irrigation is 0.13 in the region of Dinajpur, Bangladesh, and 0.11 in the region of Panchagarh, Bangladesh. Gani and Omonona [35] report that the same coefficient is 0.03 in northern Taraba state, Nigeria. The authors show that there is considerable diversity in the resource allocation among the producers of each region. They further show that the maize producers of each region as a group either significantly over- or under-utilize irrigation and other inputs. In other words, the distribution of how individual farmers allocate their resources is not an unbiased distribution around the optimal allocation strategy for that crop in that region, but rather around some other dominant strategy, distinct in each region, which the authors conclude to

be rather inefficient. This inefficiency seems to be more pronounced in Bangladesh, where maize was introduced comparatively recently. An interpretation offered by the model developed here is that the maize producers of each region had previously optimized their investment behavior under different technological or environmental conditions, or for different crops, and that current dominant investment strategies are transitional, reflecting the fact that each group of producers is in the process of evolving their investment strategy from the previous optimal to the current optimal strategy.

1.3. *Public policy and the evolution of economic behavior*

In the present model agents have individual investment strategies that specify how they invest their respective income. Their objective is to maximize their individual welfare. Agents 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.

We assume that when an agent imitates, it tries to do so faithfully. Here, agents imitate investment portfolios, and how closely the original is matched depends on the accuracy of information on the original. Control over the accuracy of this information is a powerful policy instrument. For example, laws regarding the quality, frequency, and accessibility of balance sheet information of public companies were one of the major responses to the crisis of 1929, in this case to minimize the error made by investors. Yet for the vast majority of firms and countries, balance sheets are either not public, or are not of a quality that allows ready identification of the investment strategy. So the output (product) may be known but not the input. Other examples of policy tools that facilitate knowledge transfer between firms are sponsorship of trade fairs, scientific venues, scholarships, and educative publications.

By regulating how accurately agents can imitate each other, a policy maker can control the diversity of strategies within the population. We will study the effect of diversity on population welfare numerically through computer simulations. We will address two research questions. The first is whether different types of environmental dynamics call for distinct levels of diversity for the agents to achieve a high welfare. If this is the case, then it would be advantageous for a policy maker to first identify the specific type of environmental dynamics and, if identification is possible, to steer the level of diversity of strategies such that it would be optimal for that type of dynamics. The second question follows from the fact that the policy maker is often uncertain about the environmental dynamics, either for lack of data, or because the dynamics is inherently irregular and unpredictable. 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 fast it occurs, and how often, similar to the work of Ciarli et al. [36]. Speed 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. And 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.

As has been extensively discussed by Wilhite [37], a study of economic interaction between social agents needs to give proper attention to the structure of the social network. Communication links between economic agents, whether individuals or organizations, are neither regular nor random. They are the result of a development process that is steered by, among others, geographic proximity, shared history, ethnic and religious affiliation, and common economic interests. Our simulations use a generic class of social networks that reproduce a number of stylized facts commonly found in real social networks. These are the small world property of Erdős and Renyi [38], the scale-free degree distribution of Barabási and Albert [39] and the high clustering coefficient of Watts and Strogatz [40].

The remainder of this article is organized as follows. Section 2 describes production and growth in an economy with a very general structure and presents the evolutionary mechanism of behavioral interaction. In Section 3 the

relationship between an investment strategy and the generated income growth rate is studied. On the basis of this relation the environmental dynamics and the experimental setup are formulated in Section 4. Section 5 provides simulation results and interpretations. Section 6 concludes.

2. The economic model

2.1. General features of the model

Consider a population of agents, each of which has the objective to reach a high level of individual welfare, something that can only be achieved through 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 an agent prefers a high income growth rate over a low income growth rate, it imitates the investment strategy of a fellow agent when that fellow agent realizes an income growth rate that is high relative to its own income growth rate and the income growth rate of its 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*.

One can argue that the criterion for selection should not be the income growth rate, since income depends not only on the potential of the investment strategy, but also on savings behavior. More sophisticated criteria like the growth of returns on investments would indicate the potential of an investment strategy more accurately. Since the rationale for such criteria would be to avoid dependency of the evolution of investment portfolios on the savings behavior, we decided for a simple model without savings, so that the income growth rate is equal to the growth of returns to investments. Furthermore, in this model there is no cost of shifting strategy. If there was, it would introduce a time lag between imitation of a strategy (the ‘genotype’) and its expression as income growth rate (the ‘phenotype’). This slows down the evolutionary process without changing the general results. A cost of adoption introduces at least one free parameter, which is why we decided not to include it in the present analysis.

The evolutionary model was deliberately designed to be neutral to as many economic implementation choices as feasible, among them dependency on returns on scale. In this model, dependency of returns on scale can only affect the evolutionary process if it influences how income growth rate changes with a change in strategy. If it does, either large or small investors can express an improved strategy with a higher income growth rate and are more likely to be imitated. Such a bias does not depend in any way on the type of strategy, and is entirely absorbed in the noise terms of the evolutionary process. We choose constant returns for mathematical elegance, not the least because it allows us to normalize the production coefficients.

All variables and parameters of the economic model are summarized in Table 1. For calibrated values of the free parameters, see Table 3.

2.2. Strategies, investment, and production

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

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

Table 1: Variables and parameters of the model

$ P $	population size
n	number of investment sectors
β	scaling factor of production
δ	capital depreciation
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

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 depreciation rate δ . Depreciation is assumed to be equal for all sectors and all agents. The dynamic equation for non-aggregate growth per sector is

$$K_{ia}(t) = s_{ia}(t)Y_a(t-1) + (1 - \delta)K_{ia}(t-1). \quad (2)$$

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 type production function with constant elasticity of substitution,

$$Y_a(t) = \beta \prod_i K_{ia}(t)^{\pi_i(t)}, \quad (3)$$

where β is a scaling factor that controls 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 $\pi(t) = \langle \pi_1(t) \dots \pi_n(t) \rangle$. To enforce constant returns to scale, all production coefficients are constraint to add up to one,

$$\pi(t) = [0, 1]^n, \quad \sum_i \pi_i(t) = 1. \quad (4)$$

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. As explained in the introduction and because of the conclusions reached in Section 3, we model the environmental dynamics as externally defined co-movement of two or more production coefficients $\pi(t)$. This is a general approach that can also be applied to other economic dynamics such as technological development. Section 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 $E[\log Y(t)]$ 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

$$E[\log Y(t)] = \frac{1}{|P|} \sum_a \log Y_a(t). \quad (5)$$

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

$$\gamma_a(t) = \frac{Y_a(t)}{Y_a(t-1)} - 1. \quad (6)$$

Expected log income relates to expected log income growth as

$$E[\log Y(t)] = \sum_{i=1}^t E[\log(\gamma(i) + 1)] + E[\log Y(0)], \quad (7)$$

where expected log income growth is defined as

$$E[\log(\gamma(t) + 1)] = \frac{1}{|P|} \sum_a \log(\gamma_a(t) + 1). \quad (8)$$

2.3. The evolutionary mechanism of behavioral interactions

From the point of view of evolutionary theory [22], 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. Because every agent has exactly one strategy at a time, the number of active strategies is the same as the number of agents.

The social structure of the agent population is modeled as a bi-directional network where the nodes are agents and the edges are communication links. Before the start of every simulation, a stochastic process assigns to each agent a a random set of peers that does not change during the course of the simulation. At each time step t of the simulation an agent may select one of its peers and imitate its strategy. That is, if agent a is a peer of agent b , then a can imitate b , while b can imitate a . On the other hand, if a and b are not peers, they can not imitate each other. For a detailed discussion of the network properties and its generation see [41].

The strategy of the imitating agent changes, while the strategy of the imitated agent 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

$$s_a(t) = s_b(t-1) + \mathcal{N}(0, \sigma), \quad (9)$$

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 $(n - 1) * \sigma^2$. If all agents try to imitate the same optimal investment strategy of a static environment, the expected standard deviation of the partial strategies will in fact converge on 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. As has been shown by Nannen and Eiben [42], of the many parameters that can define an evolutionary process of imitation, calibration of diversity has the highest impact on agent welfare.

3. The evolutionary dynamics

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 the choice of agent to be imitated. Here,

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 to. We call the mapping from investment strategies to income growth rate the *growth function*. The growth function calculates the asymptotic growth rate that an imitating agent realizes if it holds on to a particular investment strategy. In the out-of-equilibrium system the actual income growth rate of an agent depends also on the investment strategy it was using before it adopted its current strategy. Dependency on the previous strategy constitutes a random factor to the income growth rate that is independent of the current strategy and that is therefore ignored in this analytical treatment. It is of course present in our out-of-equilibrium simulations.

If the growth function maps one strategy to a higher asymptotic 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 asymptotic 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

$$\begin{aligned} \frac{K_{ia}(t)}{Y_a(t)} &= \frac{s_{ia}(t)Y_a(t-1) + (1-\delta)K_{ia}(t-1)}{(\gamma_a(t)+1)Y_a(t-1)} \\ &= \frac{s_{ia}(t)}{\gamma_a(t)+1} + \frac{1-\delta}{\gamma_a(t)+1} \frac{K_{ia}(t-1)}{Y_a(t-1)}. \end{aligned} \tag{10}$$

This equation is of the form

$$x(t) = a + bx(t-1), \tag{11}$$

which under the condition $0 \leq b < 1$ converges monotonically to a limit at

$$\lim_{t \rightarrow \infty} x(t) = a/(1-b). \tag{12}$$

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 depreciation, and we have $\gamma_a \geq -\delta$. For the moment, let us exclude the special case $\gamma_a = -\delta$. Then, considering that $0 < \delta \leq 1$, we have the required constraint

$$0 \leq \frac{1-\delta}{\gamma_a(t)+1} < 1 \tag{13}$$

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

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{K_{ia}(t)}{Y_a(t)} &= \lim_{t \rightarrow \infty} \frac{s_{ia}(t)}{\gamma_a(t)+1} / \left(1 - \frac{1-\delta}{\gamma_a(t)+1}\right) \\ &= \lim_{t \rightarrow \infty} \frac{s_{ia}(t)}{\gamma_a(t)+\delta}. \end{aligned} \tag{14}$$

Equation 14 describes the limit ratio of sector specific capital to income to which the economy of an agent converges monotonically. We ignore the limit notation and combine equation 14 with equation 3 to calculate asymptotic income as

$$\begin{aligned} Y_a(t) &= \beta \prod_i \left(\frac{s_{ia}(t) Y_a(t)}{\gamma_a(t) + \delta} \right)^{\pi_i(t)} \\ &= \beta \frac{Y_a(t)}{\gamma_a(t) + \delta} \prod_i s_{ia}(t)^{\pi_i(t)}. \end{aligned} \tag{15}$$

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

$$\gamma_a(t) = \beta \prod_i s_{ia}(t)^{\pi_i(t)} - \delta. \tag{16}$$

Let us return to the special case $\gamma_a = -\delta$. According to equation 2, capital per sector decreases at the depreciation 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 16 holds also for the special case $\gamma_a = -\delta$.

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 for a maximum growth rate of

$$\gamma^{opt}(t) = \beta \prod_i \pi_i(t)^{\pi_i(t)} - \delta. \tag{17}$$

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 $E(s, t)$ of a strategy $s(t)$,

$$E(s, t) = \prod_i \left(\frac{s_i(t)}{\pi_i(t)} \right)^{\pi_i(t)}. \tag{18}$$

The efficiency of a strategy measures the fraction $\gamma_a(t)/\gamma^{opt}(t)$ of optimal growth that an agent achieves with its current strategy on the current production coefficients, assuming that $\delta = 0$. If one strategy leads to a higher asymptotic growth rate than another, it is also more efficient. This measure of efficiency is a monotonous 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 measure growth in terms of what a rational agent with perfect information would achieve.

Efficiency, like the asymptotic growth rate, is a monotonically decreasing function of the Euclidean distance between a strategy and the current production coefficients, $|s_a(t) - \pi(t)|$. This function is not symmetric about the optimum; its slope varies with the direction from the optimum. Figure 1 shows how the average efficiency of a 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 asymptotic 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 asymptotic growth rate that is higher than γ' . This inner set is convex [for a related proof see 43] and according to equation 16 satisfies

$$\prod_i s_{ia}(t)^{\pi_i(t)} \geq \frac{\gamma' + \delta}{\beta}. \tag{19}$$

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 asymptotic growth rate. Its size is proportional to $P[\gamma > 0 | \pi(t)]$, the probability that a random strategy has a positive asymptotic growth rate with given production coefficients. Let $P[\gamma > 0]$ denote the probability that the asymptotic growth rate is positive if both the strategy and the production coefficients are chosen independently at random from the simplex. Figure 2 shows how $P[\gamma > 0]$ decreases as δ/β increases, for economies with respectively two, four, and




Figure 1: Average efficiency as a function of Euclidean distance to the optimum. *Notes:* The x -axis shows the Euclidean distance, the y -axis the corresponding average efficiency. Note the convex shape around the optima.

ten investment sectors. The probability tends to zero as δ/β approaches one. For given δ/β , the probability that the asymptotic growth rate of a random strategy is positive decreases as the number of investment sectors increases. What is true for a single optimum is also true for multiple optima: the larger the number of sectors over which agents have to spread their investments, the higher the pressure to be close to any given optimum, whether niche or global optimum.

The parameters δ and β determine the asymptotic growth rate associated with a given hypersurface, as well as the minimum and maximum asymptotic growth rate that can be achieved with given production coefficients. δ and β define monotonous transformations of the growth function that are irrelevant to the order of asymptotic growth rates, and which do not affect the location of the optimum nor the shape of level sets. In other words, in this type of model, environmental dynamics will affect the evolution of strategies only in so far as they are expressed as co-movement of two or more production coefficients. Moreover, the rate of convergence in equation 14 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.

4. Experimental setup

4.1. The environmental dynamics

The previous section has shown how the growth function, up to a monotonous transformation, depends on the co-movement of two or more production coefficients of a Cobb-Douglas type multi-sector economy. 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 this question we model the environmental dynamics as externally defined changes in $\pi(t)$. That is, the environmental




Figure 2: Probability of positive growth when the investment strategy and the production coefficients are chosen independently and at random. *Notes:* The x -axis shows the ratio δ/β . The y -axis shows the corresponding probability that the asymptotic growth rate is positive.

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 where 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 a static environment and a control systems without imitation. As discussed in Section 3.2, when strategies evolve in a static environment, they are expected to converge on the optimum strategy, and welfare at the population level will be at its highest. On the other hand, without imitation and 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.

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. To give an example, a sudden change can be modeled by setting the production coefficients of a two-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. Gradual change in this model reaches its final state at the end of a cycle. It is 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

$$\pi(t + j) = \frac{(x - j)\pi(t) + j\pi(t + x)}{x}, \quad 0 \leq j \leq x, \quad (20)$$

where the conditions $\sum_i \pi_i = 1$ and $\pi_i \geq 0$ 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. They are also characteristic of technical change, or changes in human capital.

As for the frequency of change, we model low frequency changes by starting the transition from one vector of production coefficients to another vector every fifty years, reflecting a Kondratiev type of wave [44] that is 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 ten years, corresponding to the fast business cycles observed by Juglar [45]. Such frequent changes are characteristic of industries that invest heavily in research and development, e.g., telecommunications and biotechnology. Both the Kondratiev wave and the Juglar cycle are generally understood to be related to capital investment dynamics and technological restructuring, as opposed to a mere change in the employment of existing inventory, which is generally understood to be represented by the even shorter Kitchin cycle [46].

We do not claim that the waves and cycles observed by Kondratiev and Juglar are caused by external environmental dynamics. Indeed, Dosi et al. [47] present an evolutionary multi-agent model where business cycles are endogenous, driven by investment dynamics. We merely use the observations by Kondratiev and Juglar to define the longest and the shortest frequency pattern of investment dynamics that have actually been observed with a reasonable amount of certainty. That is, while we acknowledge that technological innovations are driven by research and development, we treat their effect on the production coefficients as external environmental dynamics that producers have to adapt to. Since the empirical evidence for our different types of environmental dynamics displays a significant amount of variability, we assume that the policy maker is often uncertain about the current type of environmental dynamics.

We consider different sequences of production coefficients to be instances of the same environmental dynamics as long as the individual vectors of production coefficients are replaced with the same speed and frequency. Table 2 summarizes the environmental dynamics used for the numerical analysis. Figure 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 four-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.

The imitation mechanism has one free parameter, the diversity σ , which we assume to be under control of the policy maker. From the point of view of the policy maker, the level of diversity is optimal in the presence of environmental dynamics d when it maximizes the expected log income of each agent under this dynamics,

$$\sigma^{opt}(d) = \underset{\sigma}{\operatorname{argmax}} E(\log Y(t) | \sigma, d). \tag{21}$$

In order to find this optimal level for different environmental dynamics we use repeated numerical simulations with varying levels of σ and measure the expected log income at the end of each simulation, using standard statistical methods to reduce variance. Having identified the level $\sigma^{opt}(d)$ 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 about the type of environmental dynamics. To do so we measure the expected log income that an optimal level $\sigma^{opt}(d)$ generates on those environmental dynamics $d' \neq d$ where it is not optimal. We then calculate the level that policy makers with various degrees of risk aversion assign to each $\sigma^{opt}(d)$.

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 four times a year, we consider it the limit of feasibility to account for growth and to review an economic strategy as often as four 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 so that simulation results do not depend on the 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 diversity 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 initialization effects, 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 four-factor economy. Sensitivity and scalability are tested with 1,000 agents and with a ten-factor economy. In order 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

Table 2: The environmental dynamics used for numerical analysis

Environmental dynamics	Example
gradual, low frequency	oil/gas reserves
sudden, low frequency	climate change
gradual, high frequency	technological innovation
sudden, high frequency	biotic resource




Figure 3: Environmental dynamics: changes in the production coefficients. *Notes:* 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 four-factor economy. The single graphs at the right are area plots that visually combine the values of the same four production coefficients by stacking them one on top of the other, with a different shade of gray under each curve.

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 depreciation δ (equation 2) and the scaling factor β (equation 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 3, different vectors of production coefficients have distinct optimal asymptotic 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

Table 3: Values of free simulation parameters

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

factor $\beta(t)$ depend on the vector of production coefficients,

$$\beta(t) = .015 \prod_i \pi_i(t)^{-\pi_i(t)}. \tag{22}$$

With this normalization the asymptotic growth rate of all strategies is constraint to the range $[-.01, .005]$, where the minimum of $-.01$ is realized when $s_{i\alpha}(t) = 0$ for some positive π_i and where the maximum of $.005$ is realized when $s_{\alpha}(t) = \pi(t)$. Numerical tests show that with these parameter values the probability that a random strategy has a negative asymptotic growth rate on random production coefficients is about $.65$. In other words, model parameters are calibrated such that two thirds of the strategy space has negative income growth. In the long run, agents can only experience positive income growth if they adapt.

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 possible economies that fall within these ranges. For example, 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 speed and frequency of the respective environmental dynamics. Likewise, 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. The generated networks have an average clustering coefficient¹ of $.66$ and an average connectivity of ten, meaning that the overall network can be subdivided into a few local clusters that maintain a high degree of internal connectivity, and a somewhat lower degree of connectivity to agents outside the respective cluster. The number of simulations needed to obtain reliable statistical results is determined by standard methods of variance reduction. The values of the free parameters used in the simulations are listed in Table 3.

5. Results

5.1. Economic significance of diversity

Figure 4 shows the expected log income of an agent for varying levels of the diversity control parameter σ : for the control system without imitation, the static environment, and the four types of environmental dynamics. 50,000 simulations are used for each graph. 500 different levels of σ from the range $[0, .5]$ are evaluated, and results are averaged over 100 simulations per level. The plots are smoothed with a moving average with a window size of ten levels.

In the first graph of Figure 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 4 show systems with imitation and there is a clear functional relation between the level of σ and log income. For each system there is a single optimum $\sigma^{opt}(d)$ that maximizes log income under

¹In their seminal article Watts and Strogatz [40] define the clustering coefficient of a node as the number of all direct links between the immediate neighbors of that node divided by the maximum number of links that could possibly exist between them. They define the clustering coefficient of the entire network as the average clustering coefficient of the nodes of the network.

the given environmental dynamics d . The exact values of the optima are given in Table 4. The level 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 level 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 levels of σ , is much steeper than to the right, where σ is large. In other words, if agents tend to imitate faithfully, providing them with perfect information will not create the variation that is needed to adapt to change. We will revisit this fact in our discussion of policy advice under uncertainty.




Figure 4: Expected log income as a function of the diversity parameter σ .

Our first research question can now be answered: almost any level of diversity σ will allow the evolutionary agents to reach a positive income under any environmental dynamics, yet a unique optimum where 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 income.

5.2. Policy advice under uncertainty

The environmental dynamics does not need to be constant. And even when it is, the policy maker will often be uncertain about it for lack of data. In such cases of uncertainty 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 level 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

Table 4: Optimal level of diversity σ for a static environment and for each type of environmental dynamics

Environmental dynamics	Optimal diversity σ
static	0.005
gradual, low frequency	0.008
sudden, low frequency	0.046
gradual, high frequency	0.057
sudden, high frequency	0.123

Table 5: Expected log income when a level of diversity 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.
Static	1.996	1.995	1.95	1.93	1.79
Gradual, low freq.	1.988	1.991	1.97	1.96	1.88
Sudden, low freq.	1.12	1.25	1.412	1.406	1.33
Gradual, high freq.	0.55	0.72	1.38	1.39	1.31
Sudden, high freq.	-0.04	0.09	0.53	0.55	0.60

Notes: Each column shows the expected log income when the level of diversity 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 level of diversity σ that is optimal for another dynamics.

result is shown in Table 5. Each entry shows the expected log income when a level of diversity σ is applied to environmental dynamics different from the one where it is optimal. That is, each row shows the expected log income for the same environmental dynamics but different σ , and each column shows results for the same σ but different environmental dynamics. The log income of each entry is averaged over 10,000 simulations, where each simulation is based on a distinct instance of the social network and a distinct sequences of production coefficients. The values in the diagonal entries are highest for each row, confirming that the optimal level is indeed the best choice for a given environmental dynamics.

For each risk preference, each tested σ can now be associated with an expected log income value, and the σ with the best such value is considered optimal for the type of risk preference. This is shown in Table 6. Each row shows the expected log income associated with each σ under a given risk adversity, with the optimal expected log income 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.

Often, policy makers are judged based on posterior knowledge. A minimax regret policy maker does not know the nature of the change, but knows that his critics will hold him accountable for the discrepancy between what could have been achieved if he had known, and what he actually achieved. So he minimizes this discrepancy [48]. Under minimax 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. Minimax 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 level of σ rises with the degree of risk aversion. This is in line with our earlier observation about the functional relationship between diversity σ and

Table 6: Optimal policy advice under uncertainty and different degrees of risk aversion

Type of policy maker/society	Environmental dynamics that σ was optimized for				
	Static	gradual, low freq.	Sudden, low freq.	Gradual, high freq.	Sudden, high freq.
Risk seeking	<i>1.996</i>	1.995	1.97	1.96	1.88
Risk neutral	1.12	1.21	<i>1.448</i>	1.446	1.38
Minimax regret	0.84	0.66	0.07	<i>0.06</i>	0.20
Risk averse	-0.04	0.09	0.53	0.55	<i>0.60</i>

Notes: Each entry is calculated from a column of Table 5 and shows the expected log income that a risk preference associates with a particular level of diversity σ . Each row shows the expected log income that is optimal under that risk preference in *Italic type* (minimum for minimax regret, maximum for the others).

expected log income: the gradient is steeper for lower levels of σ than for higher levels, which makes higher levels of σ the safer bet. These observations are confirmed by control experiments that test for sensitivity and scalability and that use alternative implementations of the imitation mechanism.

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—dependent on income instead of growth. In all cases the arrangement of optimal levels $\sigma^{opt}(d)$ is similar for each environmental dynamics. The gradient is always steeper for small levels of σ than for large ones. Under uncertainty the optimal level of σ always increases with the degree of risk aversion.

5.3. Evolutionary dynamics

Figure 5–10 show the evolution of some key statistics during the 500 time steps of the simulation. Figure 5 discusses the control systems without imitation (a static environment is used). Figure 6 discusses imitation in a static environment. The remaining four figures discuss the four types of environmental dynamics where changes occur gradually and with low frequency (Figure 7), gradually and with high frequency (Figure 8), suddenly and with low frequency (Figure 9), and suddenly and with high frequency (Figure 10). Each figure uses the optimal level of σ for the respective type of environmental dynamics (Table 4). Figure 4 shows that there is a fair degree of robustness for the optimal level of σ . We evaluated the each type of environmental dynamics for all levels of σ between zero and 0.5. The patterns shown in Figure 5–10 are representative for a wide range (± 0.1) of sigma.

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 distinct instance of the social network and a distinct 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 distinct random sequence of production coefficients. To further ease the analysis, in Figure 7–10 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 while we study an out-of-equilibrium system, 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 18) 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 2.3, in a static environment the square root of this variance approaches the level 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




Figure 5: Average time evolution of an economy without imitation (the environment is static). *Notes:* 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.

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 lose 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 the 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

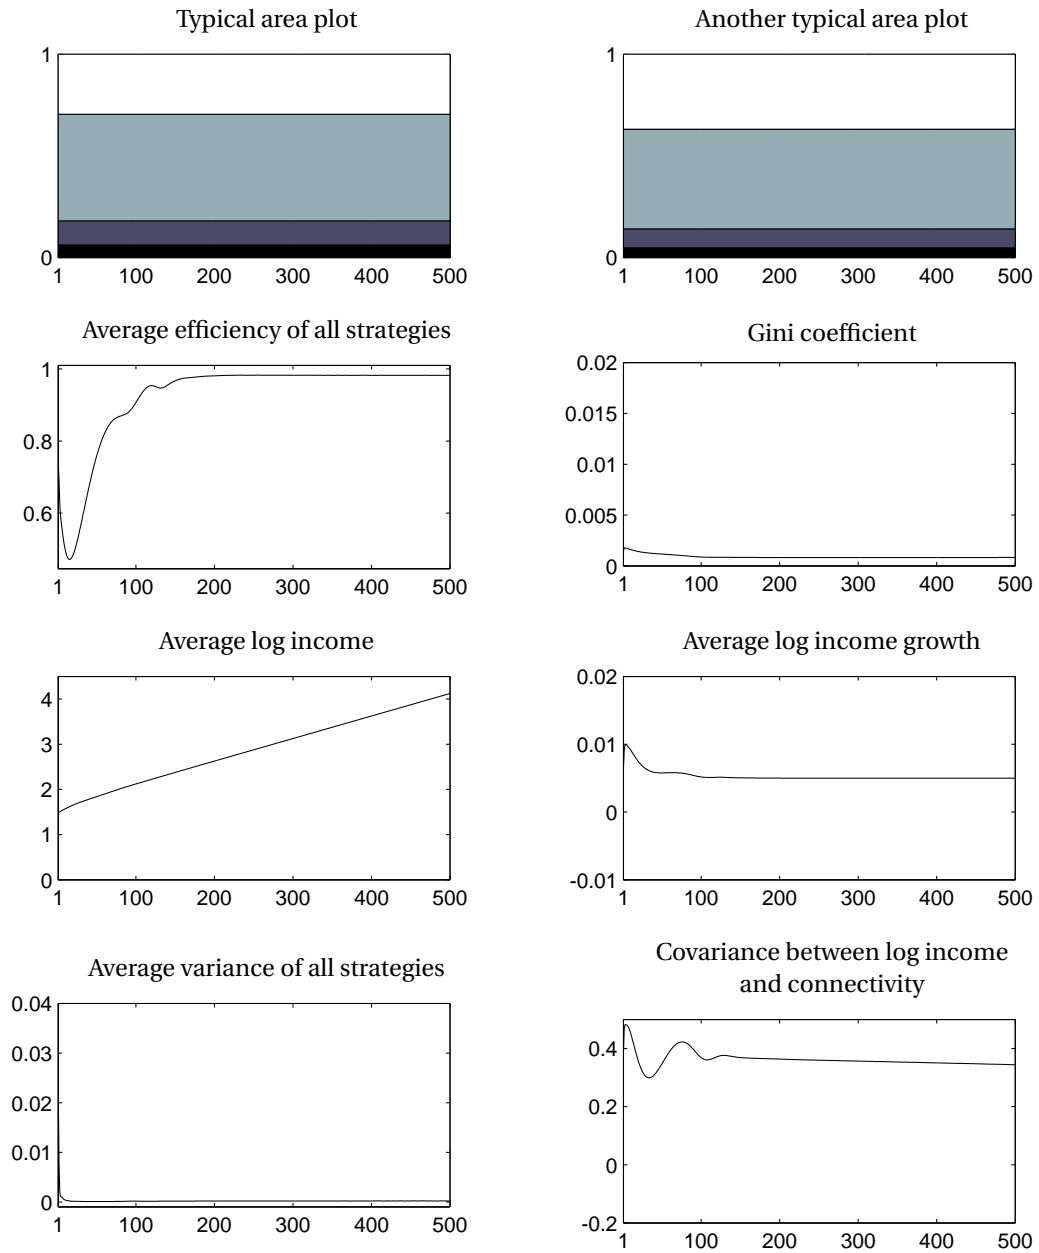


Figure 6: Average time evolution of an economy and a static environment. *Notes:* The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x -axis shows the 500 time steps while the y -axis shows the respective statistics. Note how most statistics have stabilized during the first 100 steps of the initialization phase.

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.

The keen observer will note that in the gradual scenarios of Figure 7 and 8, growth rate suddenly increases at the onset of a cycle, after which it declines with increasing speed. However, in the sudden scenarios (Figure 9 and 10), it first declines sharply, then rises, after which it declines with decreasing speed. This requires some explanation. Two effects are important here: the way the growth rate of an agent stabilizes after the agent has changed




Figure 7: Average time evolution of an economy with imitation and a dynamic environment characterized by gradual, low frequency changes. *Notes:* 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.

its investment portfolio, and the way a community of agents organizes itself around a dynamic optimum. When an agent changes its investment strategy, it will take some time until its growth rate has reached the corresponding income growth rate. The growth rate will approach the stable state always from above. This means that the income growth rate of an agent is highest immediately after the change, and then slowly diminishes until it has converged at the stable state that characterizes the strategy [49].

In the scenarios with sudden change, the community of agents forms a symmetric cloud around the current optimum. Change typically involves a large distance between the old and new optimum, with the effect that the

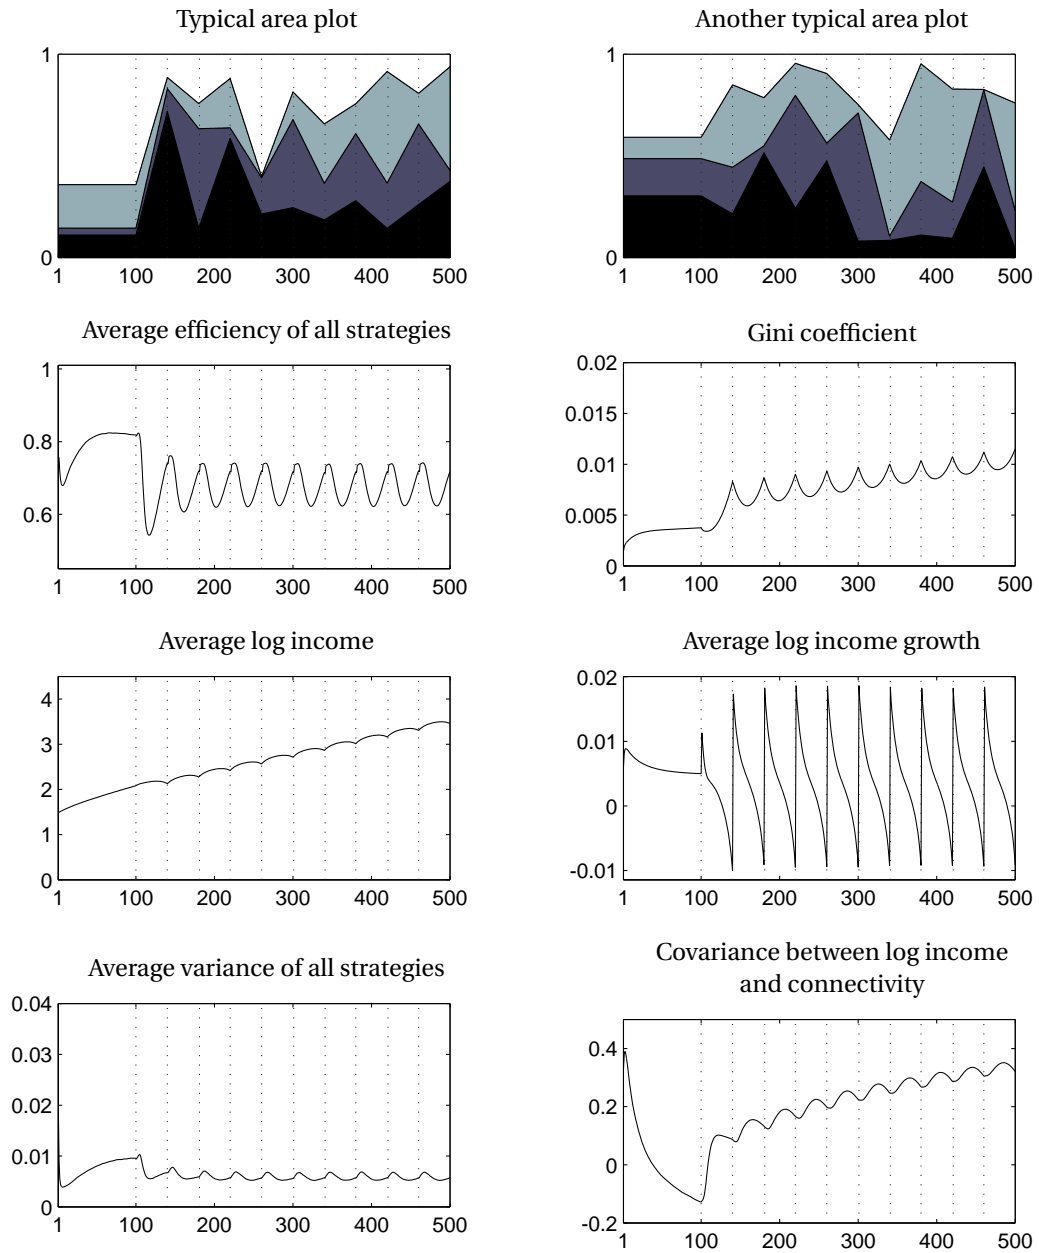


Figure 8: Average time evolution of an economy with imitation and a dynamic environment characterized by gradual, high frequency changes. *Notes:* The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x -axis shows the 500 time steps while the y -axis shows the respective statistics.

agents find themselves far from the new optimum, creating the observed depression in growth rates. Soon (if the diversity parameter is large enough) they have localized the new optimum and converge on it. As described above, the agents will now approach the stable growth rate from above. In the simulations where change is gradual, the direction of change is constant for the duration of a cycle, and changes only between cycles. During a cycle, the agents trace the moving optimum like the tail of a comet. The length of the tail grows slowly during the cycle, and the average growth rate decreases. When the optimum changes direction, possibly crossing its tail, for a shot moment the agents can form a dense cloud around the optimum and increase their growth rate.




Figure 9: Average time evolution of an economy with imitation and a dynamic environment characterized by sudden, low frequency changes. *Notes:* 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.

6. Summary

We have studied the general problem 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

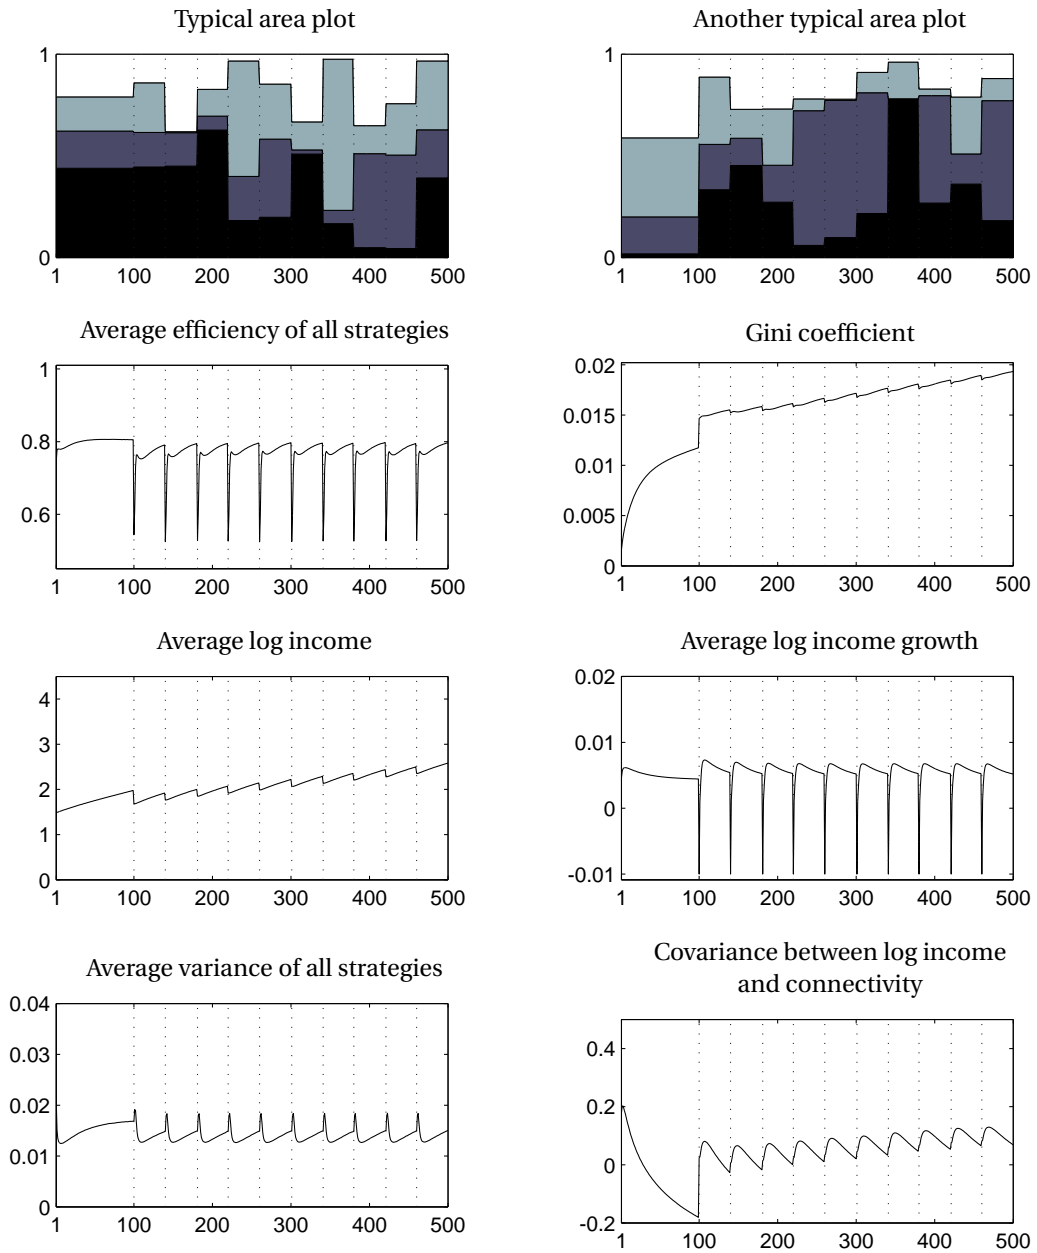


Figure 10: Average time evolution of an economy with imitation and a dynamic environment characterized by sudden, high frequency changes. *Notes:* The two area plots on top show single random sequences of production coefficients. The other statistics are averaged over 10,000 simulations. The x -axis shows the 500 time steps while the y -axis shows the respective statistics.

has a clear policy dimension and policy makers might want to control the amount of information that agents have about the investment strategy of their peers. They can achieve this, for example, by giving access to balance sheets, or by extracting and publishing the investment strategy and economic performance.

If agents in a Cobb-Douglas type multi-sector economy use relative income growth rate to determine which agent to imitate, the evolutionary dynamics are governed by the asymptotic growth rate of a strategy. This asymptotic growth rate is uniquely determined by the production coefficients of the Cobb-Douglas function. Modeling environmental dynamics as co-movement of two or more 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 co-movement of technological or macroeconomic variables as well as environmental hazards like (climate change induced) desertification and diseases. We specified four different types of environmental dynamics that differ in the speed and frequency of change, and compared them to a static environment. We further specified a static environment and a control system without imitation. To achieve general results that are valid for a broad class of economies, all numerical results were based on a large number of computer simulations, each with different instances of those factors that affect the evolutionary dynamics.

Our first research question was whether or not distinct levels for the diversity control parameter are optimal under different environmental dynamics. We established that for each environmental dynamics a distinct and unique level of optimal diversity can be found that maximizes income. This optimum increases with the speed and frequency of change. We found further that while optimization can significantly increase social welfare, it is not necessary for economic survival, since the agents were able to quickly adapt and find strategies with a positive asymptotic growth rate for almost all tested levels of diversity and in all tested environments.

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

That a higher level of diversity is better in more volatile environments does not come as a surprise. Diversity is directly related to the variance of the probability density of new genotypes. As the distance of a new optimum from the current generation of genotypes increases, so does the diversity that maximizes the probability that a new genotype discovers the new optimum. What may surprise is that the level of diversity that can generate the highest social welfare is not necessarily the safest bet under uncertainty. Perhaps our most counter-intuitive result is that it is better to err on the side of diversity than to err on the side of accuracy (in imitation). In engineering, similar studies and results are known as Taguchi methods [50]. We are not aware of any equivalent study on socially optimal behavior or policy control. We find however that our policy recommendations on optimal diversity *between* enterprises agree with established recommendations for strategic management, albeit *within* the enterprise, and provided that the enterprise is large enough to sustain a portfolio of products and methods of adequate diversity.

The presented model would best describe a community of small and medium-sized cooperating enterprises that serve similar markets with similar products. Various types of public policies directly and indirectly affect the imitation behavior of such economic agents, the diversity of their investment strategies, and their ability to adapt to a changing environment. Numeric simulations of randomized multi-agent systems can be used to evaluate such policies even when there is uncertainty about the specific nature of the environmental dynamics. Despite, or rather because of, their random nature they can identify the preferred policy under a particular degree of risk aversion.

Appendix: 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 2 for non-aggregate growth per investment sector i can be extended to include a dynamic price $p_i(t)$,

$$K_{ia}(t) = \frac{s_{ia}(t)Y_a(t-1)}{p_i(t)} + (1-\delta)K_{ia}(t-1). \quad (23)$$

The ratio of capital to income (equation 14) now converges to

$$\lim_{t \rightarrow \infty} \frac{K_{ia}(t)}{Y_a(t)} = \lim_{t \rightarrow \infty} \frac{s_{ia}(t)/p_i(t)}{\gamma_a(t) + \delta}. \quad (24)$$

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 asymptotic growth rate is

$$\gamma_a(t) = \beta \prod_i p_i(t)^{-\pi_i(t)} \prod_i s_{ia}(t)^{\pi_i(t)} - \delta. \quad (25)$$

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 monotonous transformations of the growth function that do not affect the evolutionary dynamics.

Acknowledgements

We thank David Bree for valuable comments.

References

- [1] L. Tesfatsion, Agent-based computational economics: A constructive approach to economic theory, in: K. L. Judd (Ed.), *Handbook of Computational Economics*, Elsevier, 831–880, doi:10.1016/S1574-0021(05)02016-2, 2006.
- [2] P. Mirowski, Markets come to bits: Evolution, computation and markomata in economic science, *Journal of Economic Behavior & Organization* 63 (2) (2007) 209–242, doi:10.1016/j.jebo.2005.03.015.
- [3] E. Ostrom, Collective Action and the Evolution of Social Norms, *Journal of Economic Perspectives* 14 (3) (2000) 137–158, doi:10.1257/jep.14.3.137.
- [4] J. C. J. M. van den Bergh, Evolutionary thinking in environmental economics, *Journal of Evolutionary Economics* 17 (5) (2007) 521–549, doi:10.1007/s00191-006-0054-0.
- [5] A. E. Eiben, J. E. Smith, *Introduction to Evolutionary Computing*, Springer, Berlin / Heidelberg, 2003.
- [6] K. Safarzyńska, J. C. J. M. van den Bergh, Evolutionary models in economics: a survey of methods and building blocks, *Journal of Evolutionary Economics* 20 (3) (2010) 329–373, doi:10.1007/s00191-009-0153-9.
- [7] T. C. Bergstrom, Evolution of Social Behavior: Individual and Group Selection, *Journal of Economic Perspectives* 16 (2) (2002) 67–88, doi:10.1257/0895330027265.
- [8] J. Henrich, Cultural group selection, coevolutionary processes and large-scale cooperation, *Journal of Economic Behavior & Organization* 53 (1) (2004) 3–35, doi:10.1016/S0167-2681(03)00094-5.
- [9] G. Yaari, A. Nowak, K. Rakocy, S. Solomon, Microscopic study reveals the singular origins of growth, *The European Physical Journal B. Condensed Matter and Complex Systems* 62 (4) (2008) 505–513, doi:10.1140/epjb/e2008-00189-6.
- [10] D. Challet, S. Solomon, G. Yaari, The universal shape of economic recession and recovery after a shock, *Economics* 3 (2009) 2009–2036, doi:10.2139/ssrn.1726761.
- [11] D. Ietri, M. Lamieri, Innovation creation and diffusion in a social network: an agent based approach, in: R. Leombruni, M. Richiardi (Eds.), *Wild@ace 2004 Conference*, Moncalieri, Italy, 2004.
- [12] P. Windrum, C. Birchenhall, Structural change in the presence of network externalities: a co-evolutionary model of technological successions, *Journal of Evolutionary Economics* 15 (2) (2005) 123–148, doi:10.1007/s00191-004-0226-8.
- [13] R. Sethi, E. Somanathan, The Evolution of Social Norms in Common Property Resource Use, *The American Economic Review* 86 (4) (1996) 766–788.
- [14] J. Noailly, J. C. J. M. van den Bergh, C. A. Withagen, Evolution of harvesting strategies: replicator and resource dynamics, *Journal of Evolutionary Economics* 13 (2) (2003) 183–200, doi:10.1007/s00191-003-0146-z.
- [15] A. Munro, Economics and biological evolution, *Environmental and Resource Economics* 9 (4) (1997) 429–449, doi:10.1007/BF02441760.
- [16] J. M. Epstein, R. Axtell, *Growing Artificial Societies: Social Science from the Bottom Up*, MIT Press, Cambridge, MA, 1996.
- [17] J. Noailly, Coevolution of economic and ecological systems, *Journal of Evolutionary Economics* 18 (1) (2008) 1–29, doi:10.1007/s00191-007-0067-3.
- [18] J. A. Schumpeter, *The Theory of Economic Development*, Harvard University Press, Cambridge, MA, 1934.
- [19] R. R. Nelson, S. G. Winter, *An Evolutionary Theory of Economic Change*, Harvard University Press, Cambridge, MA, 1982.
- [20] G. Dosi, Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change, *Research Policy* 11 (3) (1982) 147–162, doi:10.1016/0048-7333(82)90016-6.
- [21] G. Silverberg, G. Dosi, L. Orsenigo, Innovation, diversity and diffusion: a self-organisation model, *The Economic Journal* 98 (393) (1988) 1032–1054, doi:10.2307/2233718.
- [22] M. A. Nowak, *Evolutionary Dynamics: Exploring the Equations of Life*, Harvard University Press, Cambridge, MA, 2006.
- [23] H. I. Ansoff, Strategies for diversification, *Harvard Business Review* 35 (5) (1957) 113–124.
- [24] H. I. Ansoff, P. A. Sullivan, P. Antoniou, H. Chabane, S. Djohar, R. Jaja, P. Wang, Empirical proof of a paradigmatic theory of strategic success behaviors of environment serving organizations, in: D. E. Hussey (Ed.), *International Review of Strategic Management*, John Wiley & Sons Ltd., ISBN 978-0-471-93968-9, 173–203, 1993.
- [25] S. Thore, G. Kozmetsky, F. Phillips, DEA of financial statements data: the US computer industry, *Journal of Productivity Analysis* 5 (3) (1994) 229–248, doi:10.1007/BF01073909.
- [26] S. Thore, F. Phillips, T. W. Ruefli, P. Yue, DEA and the management of the product cycle: The US computer industry, *Computers & Operations Research* 23 (4) (1996) 341–356, doi:10.1016/0305-0548(95)00043-7.
- [27] F. Phillips, S. D. Tuladhar, Measuring organizational flexibility: an exploration and general model, *Technological Forecasting and Social Change* 64 (1) (2000) 23–38, doi:10.1016/S0040-1625(99)00077-3.
- [28] W. R. Ashby, *An introduction to cybernetics*, vol. 80, Chapman & Hall, London, 1956.
- [29] A. Hornstein, J. Praschnik, Intermediate inputs and sectoral comovement in the business cycle, *Journal of Monetary Economics* 40 (3) (1997) 573–595, doi:10.1016/S0304-3932(97)00056-1.
- [30] C. Croux, M. Forni, L. Reichlin, A measure of comovement for economic variables: Theory and empirics, *The Review of Economics and Statistics* 83 (2) (2001) 232–241, doi:10.1162/00346530151143770.

- [31] K. J. Forbes, R. Rigobon, No contagion, only interdependence: measuring stock market comovements, *The Journal of Finance* 57 (5) (2002) 2223–2261, doi:10.1111/0022-1082.00494.
- [32] M. J. van der Meulen, A. J. F. van der Spek, G. de Lange, S. H. L. L. Gruijters, S. F. van Gessel, B.-L. Nguyen, D. Maljers, J. Schokker, J. P. M. Mulder, R. A. A. van der Krogt, Regional Sediment Deficits in the Dutch Lowlands: Implications for Long-Term Land-Use Options (8 pp), *Journal of Soils and Sediments* 7 (1) (2007) 9–16, doi:10.1065/jss2006.12.199.
- [33] K. R. Krishna, *Agroecosystems of South India: Nutrient Dynamics, Ecology and Productivity*, Universal-Publishers, Boca Raton, Florida, USA, 2010.
- [34] M. F. Hasan, Economic efficiency and constraints of maize production in the northern region of Bangladesh, *Journal of Innovation & Development Strategy* 2 (1) (2008) 18–32.
- [35] B. S. Gani, B. T. Omonona, Resource use efficiency among small-scale irrigated maize producers in Northern Taraba State of Nigeria, *Journal of Human Ecology* 28 (2) (2009) 113–119.
- [36] T. Ciarli, R. Leoncini, S. Montresor, M. Valente, Technological change and the vertical organization of industries, *Journal of Evolutionary Economics* 18 (3–4) (2008) 367–387, doi:10.1007/s00191-008-0092-x.
- [37] A. Wilhite, Economic Activity on Fixed Networks, in: L. Tesfatsion, K. L. Judd (Eds.), *Handbook of Computational Economics*, Elsevier, ISBN 9780444512536, 1013–1045, doi:10.1016/S1574-0021(05)02020-4, 2006.
- [38] P. Erdős, A. Renyi, On random graphs I., *Publicationes Mathematicae*, Debrecen 6 (1959) 290–297.
- [39] A.-L. Barabási, R. Albert, Emergence of scaling in random networks, *Science* 286 (5439) (1999) 509–512, doi:10.1126/science.286.5439.509.
- [40] D. J. Watts, S. H. Strogatz, Collective dynamics of 'small-world' networks, *Nature* 393 (6684) (1998) 440–442, doi:10.1038/30918.
- [41] V. Nannen, J. C. J. M. van den Bergh, Policy instruments for evolution of bounded rationality: Application to climate–energy problems, *Technological Forecasting and Social Change* 77 (1) (2010) 76–93, doi:10.1016/j.techfore.2009.05.003.
- [42] V. Nannen, A. E. Eiben, A method for parameter calibration and relevance estimation in evolutionary algorithms, in: *Genetic and Evolutionary Computation, GECCO'06*, ACM, ISBN 1595931864, 183–190, doi:10.1145/1143997.1144029, 2006.
- [43] G. Beer, The Cobb-Douglas Production Function, *Mathematics Magazine* 53 (1) (1980) 44–48, doi:10.2307/2690031.
- [44] N. D. Kondratiev, W. F. Stolper, The Long Waves in Economic Life, *The Review of Economics and Statistics* 17 (6) (1935) 105–115, doi:10.2307/1928486.
- [45] C. Juglar, *Crises Commerciales*, Imprimerie de Veuve Berger-Levrault, Strasbourg, 1863.
- [46] J. Kitchin, Cycles and trends in economic factors, *The Review of Economics and Statistics* 5 (1) (1923) 10–16, doi:10.2307/1927031.
- [47] G. Dosi, G. Fagiolo, A. Roventini, The microfoundations of business cycles: an evolutionary, multi-agent model, *Journal of Evolutionary Economics* 18 (3–4) (2008) 413–432, doi:10.1007/s00191-008-0094-8.
- [48] G. Loomes, R. Sugden, Regret theory: An alternative theory of rational choice under uncertainty, *The Economic Journal* 92 (368) (1982) 805–824, doi:10.2307/2232669.
- [49] V. Nannen, *Convergence of Income Growth Rates in Evolutionary Agent-Based Economics*, Tech. Rep., Vrije Universiteit, Amsterdam, 2010.
- [50] G. Taguchi, Y. Wu, *Introduction to Off-Line Quality Control*, Central Japan Quality Control Association, Nagoya, Japan, 1980.

Volker Nannen is currently a post-doctoral researcher at the University of Girona. He works on stochastic methods and models to quantify and predict the evolution of complex multi-agent systems.

Jeroen van den Bergh is ICREA Research Professor in Barcelona. He is affiliated with the Institute of Environmental Science and Technology and the Dept. of Economics & Economic History of Universitat Autònoma de Barcelona. In addition, he is professor of Environmental and Resource Economics in the Faculty of Economics & Business Administration and the Institute for Environmental Studies, VU University Amsterdam. He is fellow of NAKI and the Tinbergen Institute. He is editor-in-chief of the Elsevier journal “Environmental Innovation and Societal Transitions”. His research is on the intersection of economics, environmental science and innovation studies.

A.E. Eiben is professor and head of the Computational Intelligence Group at the Department of Computer Science at the Vrije Universiteit Amsterdam. His academic research lies within computational intelligence or natural computing with evolutionary computing as the binding factor. He is doing fundamental research in evolutionary computing concerning parameter settings and applied research in swarm robotics, evolutionary art, socio-economical modelling, etc.