Exploitation, exploration and innovation in a model of endogenous growth with locally interacting agents

Giorgio Fagiolo *, Giovanni Dosi [1]

Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, P.zza Martiri della Libertà, 33, I-56127 Pisa, Italy

Abstract

The paper presents a model of endogenous growth in which firms are modeled as boundedly-rational, locally interacting, agents. Firms produce a homogeneous good employing technologies located in an open-ended technological space and are allowed to either imitate existing similar practices or to locally explore the technological space to find new, more productive, techniques. We first identify sufficient conditions for the emergence of empirically plausible GDP time-series characterized by self-sustained growth. Then, we study the trade-off between individual rationality and collective outcomes by providing an example in which more rational agents systematically perform worse than less rational ones.
© 2003 Elsevier B.V. All rights reserved.

* Corresponding author.
E-mail addresses: fagiolo@sssup.it (G. Fagiolo), gdosi@sssup.it (G. Dosi).
[1] Tel.: +39-050-883-343; fax: +39-050-883-344.

1. Introduction

The analysis of the determinants of self-sustained processes of economic growth fueled by technological advances has received an increasing attention in the past few years.

On the theoretical side, 'Endogenous Growth' and 'Evolutionary' models have been trying to explain how positive feedbacks in knowledge accumulation affect per-capita income growth (Romer, 1990; Grossman and Helpman, 1991; Nelson and Winter, 1982; Verspagen, 1993; Silverberg and Verspagen, 1994). On the empirical side, a rapidly expanding literature on the economics of technological change has been instead exploring the drivers of innovation and diffusion at the levels of firms, sectors and whole Countries (see, among others, Freeman, 1994; Rosenberg, 1994; Nelson, 1995; Stoneman, 1995).

Notwithstanding this great effort, many scholars have recently spelled out a negative assessment on the extent to which 'neoclassical', 'endogenous' and 'evolutionary' growth theories have been able to match 'old' and 'new' growth 'stylized facts' and to provide 'fresh' testable implications (cf. Durlauf and Quah, 1998; McGrattan and Schmitz, 1998; Silverberg and Verspagen, 1996). As argued by a large body of literature (cf. e.g. Nelson, 1998; Dosi et al., 1994), these difficulties are mainly due to the large gap still existing between what we historically know about the microeconomics of technical change, innovation and technological diffusion, and the ways we represent that knowledge in formal models.

For example, economic growth models do not usually account for both systematic heterogeneity observed in technological competencies and the fine details of the mechanisms governing the dynamics of interactions among economic agents. However, microeconomic diversity and institutional settings have been shown to affect in non trivial ways the properties of aggregate dynamics. Hence, any 'representative agent' reduction employed by a good deal of contemporary literature might turn out to be misleading whenever heterogeneity and interactions are important factors in explaining economic growth (see Kirman, 1992, 1998).

Furthermore, technological advances typically involve business firms whose R&D activity is characterized by routinized decisions, trial-and-error, mistakes and unexpected discoveries (cf. Dosi and Lovallo, 1998). Consequently, forward-looking rationality typically imputed to agents in standard models of growth might not be a good proxy, especially when firms face complicated environments where novelty endogenously emerges as the outcome of others' behaviors (cf. Conlisk, 1996; Dosi et al., 2003).

In economies populated by heterogeneous agents (e.g. firms) who repeatedly interact, innovate and adaptively learn about the world where they live in, observed aggregate regularities can hardly be understood as equilibrium paths (Silverberg and Verspagen, 1997).

Empirically observed properties of macroeconomic time-series might be instead more fruitfully interpreted as 'metastable' regularities emerging in a complex evolving system. For example, the observed regularities displayed by the patterns of self-sustained GDP aggregate growth may be described as emergent properties of

an economy composed of many heterogeneous simple firms interacting in some properly defined technological space (cf. Lane, 1993a,b).

Following this intuition, we present a computer-simulated model of endogenous growth in which simple, boundedly-rational firms produce a homogeneous good employing technologies located in an open-ended (i.e. without boundaries) productivity space. Technologies located close to each other have similar productivities, while more distant technologies perform better on average. Entrepreneurs can either imitate existing practices (similar to the one they currently master) or locally explore the technological space to find new and more productive techniques (i.e. innovate). We ask whether (and, if so, in which technological regimes) such an economy is able to generate self-sustaining patterns of aggregate growth with statistical properties similar to those displayed by empirically observed time-series.

The paper is organized as follows. In Section 2, we outline in more detail the building blocks and theoretical conjectures supporting the model described in Section 3. Next, in Sections 4–6, we present an extensive analysis of computer simulations. Section 7 discusses some econometric properties of the simulated time-series. The tension arising in the model between individual rationality and collective performance is illustrated with a simple example in Section 8. Finally, Section 9 draws some conclusions and flags research developments ahead.

2. Decentralized knowledge accumulation, interactions and collective outcomes

A large body of empirically-grounded contributions has recently investigated the main properties of the processes underlying the emergence of self-sustaining growth patterns. In a nutshell, two key sets of insights emerge from this literature (cf. Rosenberg, 1982; Freeman, 1982; Dosi, 1982).

First, technological search and knowledge diffusion in presence of dynamic increasing returns seem to play a primary role among the engines of growth. Technological advances are *endogenously* generated through resource-expensive search undertaken by a multiplicity of agents. Search is generally characterized by radical uncertainty and innovative entrepreneurs are driven by the belief that ''there might be something profitable out there''. As agents are generally unable to form probability distributions on the outcomes of their search efforts, systematic mistakes in innovative search and adoption are very likely.

Second, the process of technical change appears to be driven not only by innovation but also by time-consuming diffusion (see also Jovanovic and Rob, 1989; Jovanovic, 1997). Innovations are indeed not entirely appropriable and knowledge progressively spreads (with some time lags involved) to other agents who might catch-up by investing in imitation. Knowledge accumulation generally entails dynamic increasing returns at the level of individual agents. However, radically new technologies typically involve discontinuities and only part of the old knowledge might be useful in the exploitation of subsequent technologies.

In order to embody the foregoing properties in the present analysis, we will start by describing search and innovation activities, technological diffusion and knowl-

edge accumulation as *direct* interaction processes taking place in some (high-dimensional) technological space (see Kirman, 1998; Chiaromonte and Dosi, 1993). Suppose indeed that the technologies currently adopted by all firms in the economy, as well as those still to be discovered, are associated to points of a metric space (e.g. a 2-dimensional regular lattice). Any metrics the space is endowed with will then metaphorically represent technological dissimilarity: similar technologies will lie close to each other, while more productive technologies will be situated far from existing ones. Both imitative and innovative activity might be, therefore, represented as an interaction process in which the sets of interacting units are firms and technologies. Any firm will directly affect the behaviors of other firms employing similar technologies. Since adopted technologies will typically change in time, interaction structures (e.g. who interacts with whom) are likely to endogenously change over time as well.

More precisely, we will employ the following metaphor. Think of the technological space as an empty, unbounded sea. The notional production possibility set is composed of a discrete, countable set of production activities (technologies, paradigms, etc.), each of which can be thought as an island randomly placed in a point of the sea and endowed with a mine. The economy is populated by a discrete, finite population of firms (or 'Schumpeterian' entrepreneurs) initially distributed across a small subset of islands (i.e. the set of fundamentals). We assume that an island can be at any point in time exploited by more than one agent, while each agent cannot exploit more than an island at the same time.

Every agent currently living on an island represents one of the adopters of technology embodied in it (or, in our metaphor, a 'miner') and extracts a homogeneous good (i.e. GDP). Mining is possibly characterized by increasing returns to scale in the number of current adopters due to knowledge-accumulation. Since distances between islands in the sea are a proxy of technological (productivity) differences and the sea is unbounded, notionally unlimited opportunities do exist in the economy, albeit at each point in time only a small subset of mines are known and exploitable, i.e. those which have been operated by any one firm so far.

We suppose that miners might become explorers by leaving the island they are working on and traveling around to find still unknown, possibly better, islands. The set of fundamentals can be, therefore, enlarged through endogenous innovations. Alternatively, miners might try to capture informative signals about the location of already known, better, islands and decide whether to imitate firms currently working on them. Of course, this representation of the space of technological opportunities and of the related innovation processes is much more abstract than any particular empirical example. However, we believe that it captures some of their general features, including the intrinsic uncertainty of search, the idiosyncratic and cumulative nature of technological learning, and the painstaking patterns of melioration and diffusion of specific bodies of knowledge (i.e. of *technological paradigms*).

Within this framework, we will address the following issues. First, we will attempt to identify conditions under which the economy is able to tackle the trade-off between 'exploitation' of existing technologies and 'exploration' of potentially

superior ones and to generate patterns of self-sustaining growth[2]. Second, extensive Montecarlo (MC) simulations will be performed so as to map technological regimes (i.e. strength of path-dependency in learning achievements, levels of individual willingness to explore, etc.) into aggregate growth patterns. Third, as a plausibility check, we shall investigate whether the model is able to generate GDP time-series displaying statistical properties similar to the empirically observed ones. Finally, we will discuss the conflict arising in the model between individual rationality and collective economic performance. In particular, we shall investigate a simple situation wherein boundedly rational firms are replaced by a representative agent with unbounded computational skills and complete information about the structure of the economy.

3. The model

Consider a technological space represented as a 2-dimensional, infinite, regular lattice endowed with the 'Manhattan' metrics d_1[3]. Time is discrete and the generic time-period is denoted by $t = 0, 1, 2, \ldots$. The economy is populated by a finite set of agents $I = \{1, 2, \ldots, N\}$, $N \ll \infty$, and a countable infinite number of islands, indexed by $j \in \{1, 2, \ldots\}$. There is only one good (GDP), which can be 'extracted' from any island.

Each node (x, y) in the lattice can be either an island or not and each island has a size of one node. Let $p(x, y)$ be the probability that the node (x, y) is an island. We will assume that $p(x, y) = \pi$, all (x, y), where $\pi \in (0, 1)$. Each island j is completely characterized by its coordinates (x_j, y_j) in the lattice and by a 'productivity' coefficient $s_j = s(x_j, y_j) \in \Re_+$ (i.e. the amount of good which can be extracted if there is only one agent on j).

Each agent $i \in I$ is in turn characterized, at each t, by her state $a_{i,t}$ and her position in the lattice $(x_{i,t}, y_{i,t})$. The state of an agent $a_{i,t}$ can be: 'miner', 'explorer' or 'imitator', i.e. $a_{i,t} \in \{'mi', 'ex', 'im'\}$.

Denote by $m_t(x_j, y_j)$ the number of agents currently working on island j and define an island j to be currently 'known' if $m_\tau(x_j, y_j) > 0$ for at least a $\tau : 0 \le \tau \le t$, i.e. if it currently hosts some agents or if it did host some miners in the past. Accordingly, let the set of currently 'known' islands be defined as:

$$L_t = \{j = 1, 2, \ldots : \exists\, 0 \le \tau \le t : m_\tau(x_j, y_j) > 0\} \tag{1}$$

Let us call 'colonized' a known island which is currently exploited at t, i.e. an island $j \in L_t : m_t(x_j, y_j) > 0$. Conversely, all islands which are not in L_t will be 'unknown',

---

[2] The exploitation–exploration trade-off in individual decisions (as well as its consequences for aggregate efficiency) is studied in March (1991).

[3] The distance between any two nodes (x, y) and (x', y') in the lattice is thus: $d_1 = |x - x'| + |y - y'|$. The choices of the lattice (and its dimension), as well as that of the metrics, do not crucially affect our results.

since no agent has previously exploited them. Finally, denote the cardinality of L_t by ℓ_t. Let us turn now to describe how the economy evolves.

3.1. Production

Suppose that at time t agent $i \in I$ is a 'miner' currently located on island $j \in L_t$ with coordinates (x_j, y_j). We assume that i will extract, at no cost, an output $Q_{i,t}$ given by:

$$Q_{i,t} = s(x_j, y_j)[m_t(x_j, y_j)]^{\alpha - 1} \tag{2}$$

where $\alpha \geq 1$. Hence, the current total output of island $j \in L_t$ will be:

$$Q_t(x_j, y_j) = s(x_j, y_j)[m_t(x_j, y_j)]^{\alpha}. \tag{3}$$

Total output (GDP) will obviously read: $Q_t = \Sigma_{j \in L_t} Q_t(x_j, y_j)$.

3.2. Exploration and innovation

At time t, each miner currently working on island $j \in L_t$ decides to become explorer (i.e. $a_{i,t+1} = \text{'}ex\text{'}$) with probability $\epsilon \geq 0$, where ϵ is taken to measure the willingness to explore of agent i (which in this first approximation is the same for all agents).

If i decides to become explorer, she leaves island j and 'sails' around until another—possibly still unknown—island is discovered. During the search, explorer i is not able to extract any output and moves through the lattice following the 'naïve' stochastic rule:

$$Prob\{(x_{i,t+1}, y_{i,t+1}) = (x, y)|(x_{i,t}, y_{i,t})\} = \frac{1}{4} \Leftrightarrow |x - x_{i,t}| + |y - y_{i,t}| = 1, \text{ all } (x,y) \tag{4}$$

While exploring, each agent carries the 'memory' of the last quantity of output produced in the state of 'miner', i.e. $Q_{i,\tau}$, where τ is the last period of mining before leaving.

The new location of the explorer $(x_{i,t+1}, y_{i,t+1})$ might obviously be: (i) 'sea'; (ii) a 'known' island $j \in L_t$; (iii) a 'new' island $j \in \{1, 2, \ldots\} \backslash L_t$. Let us focus on the third case[4]. If the node inspected by explorer i at time $t+1$ is a 'new' island (which happens with probability π), we assume that the new island j^* with coordinates $(x_{j^*}, y_{j^*}) = (x_{i,t+1}, y_{i,t+1})$ is added to the set of 'known' islands, i.e. $L_{t+1} = L_t \bigcup \{j^*\}$ and $\ell_{t+1} = \ell_t + 1$.

In order to capture the crucial distinction between innovations within existing knowledge bases and introduction of radically new 'technological paradigms' (cf. Dosi, 1982), we let the 'intrinsic' productivity coefficient of a 'new' island j^*

---

[4] In the first case $(x_{i,t+1}, y_{i,t+1}) \neq (x_j, y_j)$ for all j, and $a_{i,t+1} = \text{'}ex\text{'}$ (i.e. the exploration goes on), while in the second case, there will be a $j \in L_t$ such that $(x_{i,t+1}, y_{i,t+1}) = (x_j, y_j)$ and hence the explorer i becomes miner on $j \in L_t$, i.e. $a_{i,t+1} = \text{'}mi\text{'}$.

discovered by an explorer carrying the output memory $Q_{i,\tau}$ to be:

$$s_{j*} = s(x_{j*}, \ y_{j*}) = (1 + W) \cdot \{[|x_{j*}| + |y_{j*}|] + \varphi \cdot Q_{i,\tau} + \varpi\} \tag{5}$$

where W is a Poisson-distributed r.v. with mean $\lambda > 0$, ϖ is a uniformly-distributed r.v., independent of W, with mean zero and variance 1; and, finally, $\varphi \in [0, 1]$. The interpretation of Eq. (5) is straightforward. The initial productivity of a 'new' island depends on four factors, namely: (i) its distance from the origin; (ii) past 'skills' of the discoverer, i.e. $\varphi \cdot Q_{i,\tau}$ (i.e. a cumulative learning effect); (iii) a random variable W which allows for low probability high 'jumps' (i.e. changes in technological paradigms); (iv) a stochastic i.i.d. zero-mean noise ϖ controlling for high-probability low-jumps (i.e. incremental innovations).

3.3. Interactions, diffusion of knowledge and imitation

Exploitation of existing technologies is not associated to production only. Indeed, miners might also decide to imitate currently known technologies by taking advantage of informational spill-overs emanated by more productive islands located in their technological neighborhood.

More formally, the process of knowledge diffusion and imitation works as follows. Let m_t be the number of 'miners' currently present in the economy. At time t, agents mining on any 'colonized' island j deliver a signal, which is instantaneously spread in the system. A signal delivered from (x_j, y_j) is *received* by a miner currently located at $(x, y) \neq (x_j, y_j)$, independently of all other delivered signals, with probability:

$$w_t(x_j, \ y_j; \ x, \ y) = \frac{m_t(x_j, \ y_j)}{m_t} \ \exp\{-\rho[|x - x_j| + |y - y_j|]\}, \quad \rho \geq 0. \tag{6}$$

We call $w_t(x_j, \ y_j; \ x, \ y)$ the *intensity* of the signal. Notice that $w_t(x_j, \ y_j; \ x, \ y)$ is increasing in the share of miners working on j and decays exponentially with the distance between source and receiver.

Furthermore, each signal has a *content* $c_t(x_j, \ y_j)$ equal to actual productivity of the island from which it is emitted, given by:

$$c_t(x_j, \ y_j) = \frac{Q_t(x_j, \ y_j)}{m_t(x_j, \ y_j)}. \tag{7}$$

Agent i will simply choose the signal associated to the largest *content* among all signals she has received (and randomizing if ties occur).

Let us suppose that the receiver i is a miner on j. If the selected technology h is not the one she is currently mastering (i.e. $h \neq j$), she will become an 'imitator' ($a_{i,t+1} =$ '*im*'). She will then move toward the imitated island (one step per period) and following the shortest path leading to h. Therefore, she will adopt h after $k = |x_h - x_j| + |y_h - y_j|$ time periods. This allows us to embody in the model the time-consuming nature exhibited by many processes of technological adoption and diffusion. Finally, once the imitated island is reached, she will turn again her state

into 'miner', i.e. $a_{i,t+k+1} = 'mi'$. If on the contrary she chooses to stay on her current island, nothing happens and she will keep working on j at time $t+1$.

3.4. Discussion

Before describing the implementation of the model and discussing the results, some considerations are in order. First, in tune with the philosophy of agent-based and evolutionary modeling, we start by analyzing a very simple economy populated by naïve agents behaving according to routinary, myopic rules. For instance, the exploration rule (4) implies that agents are not aware of (and cannot learn) the fact that islands are on average more and more productive the further away one goes from the origin of the lattice, as the expected location after k periods is simply the starting node:

$$E[(x_{i,t+k}, \ y_{i,t+k})|(x_{i,t}, \ y_{i,t})] = (x_{i,t}, \ y_{i,t}).$$

Moreover, we make the extreme assumption that the activities of exploration, imitation and production are costless and mutually exclusive. In fact, miners cease to produce while imitating and exploring. This can be interpreted as a sort of opportunity cost agents must bear in order to divert resources from production to R&D or imitation. In a more realistic picture, firms should have been endowed with additional decision rules governing allocation of resources among exploitation and exploration activities (see e.g. Nelson and Winter, 1982). The strategy of keeping as small as possible the microeconomics of firm behaviors allows us to focus on the effects of the purported engines of growth only (i.e. innovation, diffusion, etc.).

Second, the parameters governing production, exploration, innovation and imitation define easily identifiable technological regimes. In particular, whether $\epsilon \in (0, 1]$ or $\epsilon = 0$ allows us to discriminate between economies in which endogenous innovation is permitted or not. Furthermore, α tunes the regime of returns to scale in production, with $\alpha > 1$ meaning increasing returns to scale, e.g. due to learning by doing or economies of agglomeration. In addition, λ and π tune the degree of notional 'opportunities' in the economy. Indeed, a large λ lets average productivity of a newly discovered technology to be sensibly larger than that associated to currently 'known' islands. Conversely, a smaller λ implies search processes characterized by small improvements upon currently mastered practices (i.e. incremental innovations). Likewise, a larger π induces a larger average number of per-period discoveries and thus is associated to economies where technological opportunities are very likely.

Third, the strength of path-dependency in innovation depends positively on φ. Large φ's mean that more skilled 'explorers' (i.e. 'miners' who have been more efficient in the past) are likely to discover more productive islands today and, therefore, to produce more in the future, thanks to a sort of 'learning-to-learn' mechanism *à la* Stiglitz (1987).

Fourth, the process of knowledge diffusion governs the interaction regime in the model (see Fagiolo, 1998; Kirman, 1998). Indeed, the behaviors of any firm is directly affected by the information signals emanated by agents employing similar

technologies. The parameter $\rho \geq 0$ tunes the 'degree of locality' of the interactions: the larger ρ, the more the process of diffusion of knowledge is local, since signals will tend to reach, in probability, only 'nearest neighbors'. Two extreme cases are: (i) $\rho = 0$, i.e. interactions are global, as information diffusion does not depend on the distance between source and receiver; and (ii) $\rho = \infty$, i.e. no signals are spread and interactions are shut down.

3.5. Initial conditions, timing and implementation

Suppose that at time $t = 0$ a set of initial islands L_0 (together with their coordinates in the lattice) is given and that all agents are randomly distributed across the ℓ_0 'mines'. Assume also that the intrinsic productivity coefficients of any initial island $j \in L_0$ is simply $s(x_j, y_j) = |x_j| + |y_j|$.

In each $t = 1, 2, \ldots$, given current agents' coordinates and states, the timing of decisions and events occurring in a generic iteration (i.e. in the time interval $(t-1, t]$) runs as follows. First, agents take their decisions: miners update output and choose whether to start searching; explorers select the next portion of the lattice to explore (and, possibly, they find a new island); imitators keep approaching the technologies they have chosen to adopt. Second, interactions take place through information diffusion. Finally, all time-t system variables are accordingly updated and the next iteration starts.

The model is an example of a so-called 'artificial economy' (cf. Lane, 1993a; Epstein and Axtell, 1996). Unless the focus is not on particular stationary cases (e.g. $\epsilon = 0$), one is bound to analyze its main properties by resorting to computer simulations. Analytical solutions are not indeed achievable for the full-fledged form, because of the underlying complication of the stochastic processes updating micro- and accordingly macro-system variables.

In the next sections we will present an overview of simulation results[5], with particular emphasis on the aggregate properties of the simulated time-series of the log of GDP, i.e. $q(\omega) = \{\log Q_t, \ t = 1, \ \ldots, \ T; \ \omega\}$, where ω is a point in the parameter space Ω, that is:

$$\omega \in \Omega \equiv \{(\rho, \ \lambda, \ \alpha, \ \varphi, \ \pi, \ \epsilon, \ N, \ T) \in \mathfrak{R}_+^2 \times [1, \ \infty) \times [0, \ 1]^3 \times \{1, \ 2, \ \ldots\}^2\} \qquad (8)$$

To begin with, we will analyze how the model behaves in some 'benchmark' parametrizations, in order to assess the role played by knowledge-specific increasing returns, imitation and exploration in the dynamics of the economy. In particular, we will start by addressing the question whether the model is able to display patterns of persistent growth and, if so, under which behavioral and system parametrizations (especially concerning the degree of 'open-endedness' of the economy, as well as innovation and diffusion rates).

---

[5] For a thorough discussion of the results presented in the following Sections and for extensions of the model, cf. Fagiolo (2000).

4. The emergence of self-sustained growth: bounded versus open-ended economies

A key feature of the model resides in its ability to allow for an endogenous evolution of the set of fundamentals of the economy. But, in the first place, what happens if one bounds, to some extent, the dynamics governing the progressive enlargement of the technological frontier? Put it differently, is the economy able to generate patterns of self-sustaining aggregate growth if one considers stationary environments where agents behave on the grounds of a fixed set of fundamentals?

The answer to this question is no. To see this, let us first analyze the benchmark case of a 'bounded' economy (i.e. one in which $L_t \equiv L$, $\forall t$) and considering two distinct setups.

4.1. Bounded economies without exploration

Assume first no possibility of exploration whatsoever, i.e. $\epsilon = 0$. In this setup, agents can exchange information about a fixed set of technologies, but they cannot endogenously introduce innovations in the system. To study the behavior of the system, we can focus, without loss of generality, on economies composed of only two islands, i.e. $|L_0| = \ell_0 = 2$.

In this case, we may neglect any spatial consideration and suppose that the productivity coefficients $(s_1, s_2) \in \Re_+^2$ also represent the technological distance between islands. More precisely, let $(s_1, s_2) = (1, s)$, $s = 1, 2, \ldots$, and suppose that if a miner working on island $j \in \{1, 2\}$ at the beginning of time $t-s$ decides to imitate island $j' \neq j$, then she will reach j' at the end of time $t-1$ and start producing at time t. Island 2 plays here the role of the 'best practice' for $s \geq 2$, while the case $s = 1$ depicts the benchmark case of homogeneous technologies.

In either case, the dynamics of the economy is entirely driven by the process of information diffusion (cf. Section 3.3), until one out of the two technologies, say j, manages to capture all N agents. In that case, no signal can be emitted by the other island and, therefore, the economy locks-in at the steady state where total output is $Q^* = s_j N^\alpha$. An example of the behavior of the time-series q_t is shown in Fig. 1(a).

As intuition suggests, however, path-dependency entailed by increasing returns will tend to drive all agents, through waves of imitation, toward the island characterized by the *actual* (not *initial*) highest productivity. This in turn implies non-ergodicity in the stochastic process governing output evolution and, consequently, potential inefficiency.

More formally, define M_{jt} as the random variable: "number of agents 'mining' on island j at time t", $j \in \{1, 2\}$. It can be easily shown (see Fagiolo, 2000 for details), that, if $s \geq 2$, $\underline{M}_t = \{(M_{1t}, M_{2t}), t \geq 1\}$ is a non-stationary, aperiodic Markov process with two absorbing states $\underline{m}_+ = (0, N)$ and $\underline{m}_- = (N, 0)$.

Moreover, let $p^s_-(m_0; \alpha, \rho)$ be the absorption probability in island 1, i.e. the probability of being absorbed in the inefficient limit state if $s \geq 2$ given $m_0 = m_{10} \in \{1, \ldots, N-1\}$ and system parameters. Simulations indicate that $p^s_-(m_0; \alpha, \rho)$ is non-increasing in s and ρ and non-decreasing in m_0 and α. Fig. 1(b and c) show examples of the estimation of $p^s_-(m_0; \alpha, \rho)$ for $s = 1, 2$, as $(m_0; \alpha, \rho)$ vary in the relevant

parameter space. Notice that when the initial number of 'inefficient' adopters is below a certain threshold (which itself increases with the strength of returns to scale α and the technological gap s), the system will inevitably converge to the efficient outcome no matter how large are the incentives to stick to the initial choice. However, when m_0 goes through that threshold, the probability of ending up in the inefficient state becomes strictly positive and grows as the incentives to knowledge accumulation increase. In the limit, when only a few miners are initially aware of the superior technology and returns to scale are increasing ($\alpha > 1$), the probability that the system is absorbed by island 2 converges to zero. Finally, when $s \geq 2$, the more information diffusion is local (i.e. the greater ρ), the smaller the average number of miners which leave their islands and, consequently, the less likely the event that waves of imitation triggers a migration from the efficient technique toward the inefficient one, see Fig. 1(d). Therefore, for a given (m_0, α), the probability of being absorbed in island 1 will decrease with ρ (increasingly fast as s grows).

4.2. Bounded economies with exploration

In a setup without exploration, non-ergodicity of the stochastic process $\underline{M}_t$ implies that the long-run steady-state GDP level is determined by unpredictable, early waves of imitation (cf. David, 1992). As it happens in Polya urn schemes (cf. Arthur, 1994), the system locks-in in the long run. However, in Arthur's model lock-in occurs because population size increases without bound. This implies that the perturbations introduced by individual choices become irrelevant in the long-run. Conversely, in the reduced form of the model presented here the population size is constant and perturbations die away as soon as an island manages to capture all miners.

In order to explore what happens when the perturbation rate does not vanish, we study a second benchmark setup where: (i) the probability of finding new islands (that is the innovation probability) is $\pi = 0$ as before; (ii) exploration (as well as information diffusion) is permitted ($\epsilon > 0$), but only inside the initial set of 'knowledge bases'.

In this economy miners can become explorers with some probability $\epsilon > 0$, but they will only be able to 'sail' within the box containing all initial islands (or, equivalently, on a finite regular lattice with periodic boundaries). This implies that, for a given population size, the lock-in of the system will not generally occur, since there is always a positive probability that 'non conformist' decisions will induce phase-transitions in the system. Notice that here we allow for a high potential source of 'irrationality' and 'idiosyncrasy' in individual behaviors, because agents could always decide to leave the island they are working on, even though all agents are mining on it.

In a two-islands setup, the economy is characterized as before by the Markov process $\underline{M}_t$, together with the stochastic process describing the current number of explorers. However, unlike the previous case, transition probabilities are not only influenced by the propensities to imitate technologies with higher revealed productivity, but also involve a certain probability of 'exploring'. Islands represent
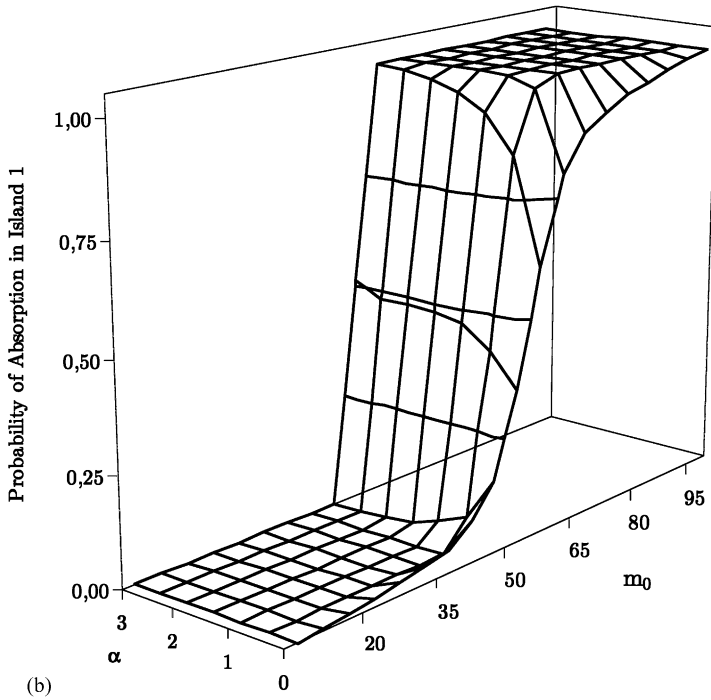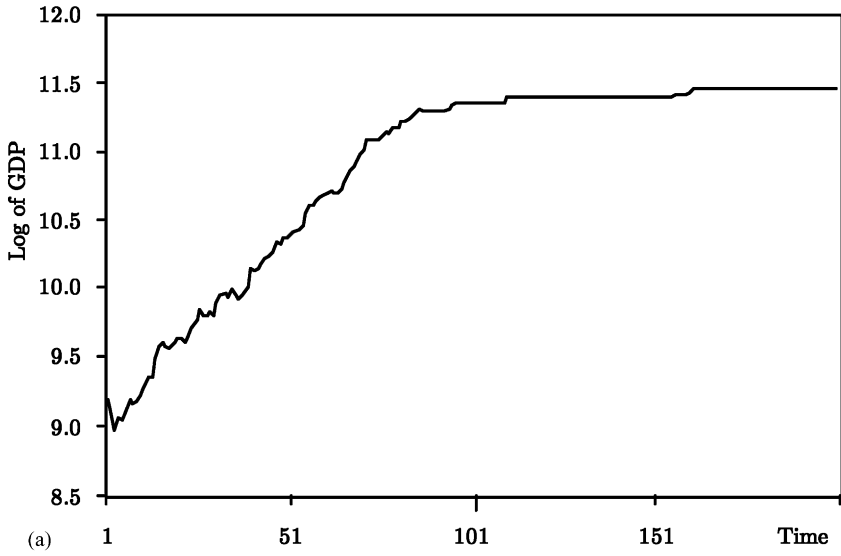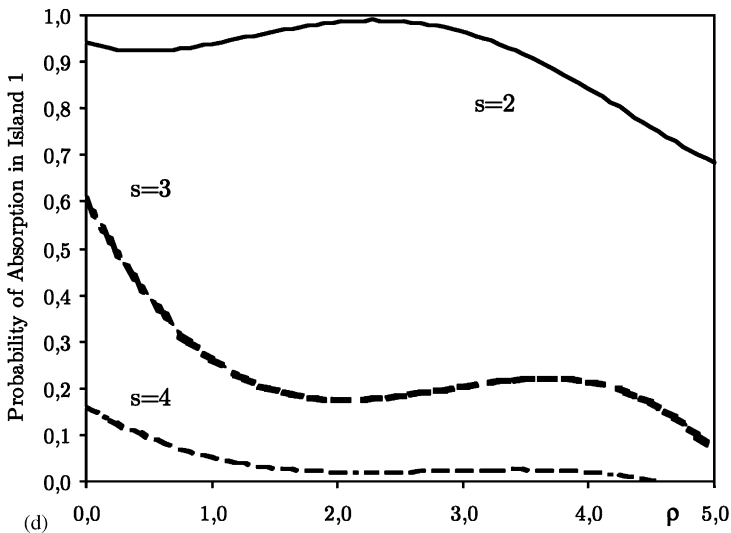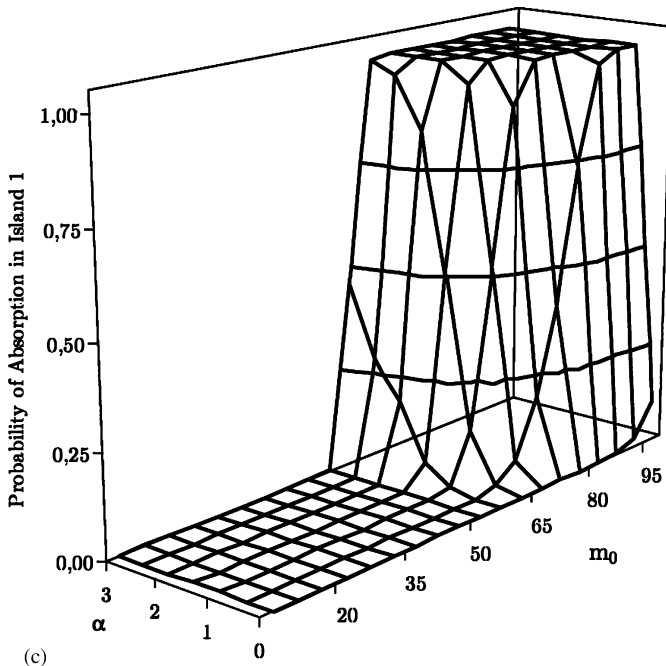
Fig. 1.

Fig. 1. (a) A bounded economy without exploration. A typical log(GDP) time series. Par. Setup: $N = 100$, $\pi = 0.1$, $\rho = 0.1$, $\alpha = 1.5$. (b) A bounded economy without exploration. Estimated probability of absorption in island 1 as a function of α and m_0 in the case $s = 1$. (c) A bounded economy without exploration. Estimated probability of absorption in island 1 as a function of α and m_0 in the case $s = 2$ ($\rho = 0.01$). (d) A bounded economy without exploration. Estimated probability of absorption in island 1 as a function of ρ and s. Par. Setup: $m_0 = 65$, $\alpha = 1.4$.

here 'basins of attraction' among which the system continually oscillates[6]. The stochastic process of exploration/imitation yields persistent output fluctuations but only transitory growth. Over finite time periods, increasing returns and knowledge diffusion induce agents (on average) to move toward currently more efficient islands, cf. Fig. 2(a and b) for the two cases $s_1 = s_2$ and $s_1 < s_2$. However, exploration allows with positive probability 'de-locking' bursts, also toward notionally less efficient islands. In a sense, persistent fluctuations are in this case generated by a problem of *imperfect Schumpeterian coordination* in presence of dynamic increasing returns to learning.

4.3. *Exploring in an open-ended economy: some qualitative results*

In both stationary environments analyzed so far, self-sustaining growth emerges only if one superimposes an exogenous Solow-like drift on the best-practice production function. Otherwise, as long as agents behave on the grounds of fixed fundamentals, economic growth is a transient phenomenon.

Consider now the more general case where $\epsilon > 0$ and the economy is open-ended (i.e. agents explore in a technological space without boundaries). Since firms are able to endogenously induce a drift in the technological frontier, the economy exhibits, for a wide range of parameters, patterns of self-sustaining (exponential) growth, cf. Fig. 3[7].

In all these cases, many other interesting regularities do actually arise. Suppose to start from a fairly uniform distribution of N agents working on the initial set of 'known' islands L_0. First, the number of currently 'known' islands linearly increases in time. However, both the percentage of 'known' islands and the number of 'colonized' technologies fall quickly and then follow a stationary process. This suggests that a typical evolution of the system runs as follows. In the first time periods, diffusion of knowledge drives agents to concentrate on a relatively small cluster of 'known' islands which, thanks to dynamic increasing returns, tend to be the most efficient ones. Relatively ordered spatial patterns of colonized islands are then likely to emerge, due to the local nature of both the exploration and imitation processes. In Fig. 4(a), the path of expansion of a 'best practice' proxy b_t^* is plotted[8], together with four 'snapshots' showing the locations of currently 'colonized' islands in the positive orthant of the technological space for different time periods $t = 0, 500, 1000, 1500$. While in the early time periods of the simulation small (stochastic) events

---

[6] The properties of the stochastic process governing the evolution of the system are qualitatively similar to those discussed in Kirman (1993). For instance, when the ratio between willingness to explore and the size of the population (ϵ/N) decreases, the system tends to spend an increasing number of time periods close to the absorbing states of an $\epsilon = 0$ economy.

[7] All results reported in this Section refer to the parametrization: $\pi = 0.1$, $\rho = 0.1$, $\epsilon = 0.1$, $\lambda = 1$, $\varphi = 0.5$, $N = 100$, $\alpha = 1.5$, $T = 1000$. Cf. Section 5 for an extensive MC investigation of the parameter space.

[8] We define $b_t^* = (x_t^*, y_t^*)$, where $x_t^* = \max\{|x_{jt}|, j \in L_t\}$ and $y_t^* = \max\{|y_{jt}|, j \in L_t\}$, i.e. the vertex of the smallest rectangle containing all currently known islands whose distance from the origin is the maximum one. Notice that b_t^* does not necessarily coincide with a known island.
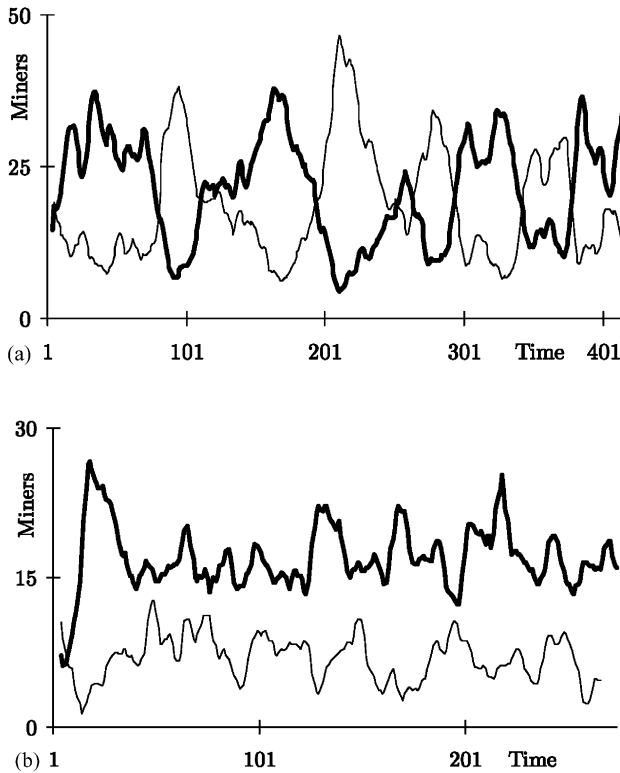
Fig. 2. (a) A bounded economy with exploration: number of miners in island $j = 1, 2$ when $s_1 = s_2$. Thick line: island 2. Par. Setup: $N = 100$, $\rho = 0.1$, $\epsilon = 0.1$, $\alpha = 1.4$, $m_0 = 50$. (b) A bounded economy with exploration: number of miners in island $j = 1, 2$ when $s_1 < s_2$. Thick line: island 2. Par. Setup: $N = 100$, $\rho = 0.1$, $\epsilon = 0.1$, $\alpha = 1.4$, $m_0 = 50$.

select the region of the lattice where exploration will be initially carried through, the path-dependent nature of the overall process tends to keep the economy inside that region. Therefore, 'rare' events (i.e. exceptional discoveries), feeding path-dependently upon diffusion and incremental innovations thereafter, might be able to trigger a self-reinforcing process whose ultimate outcome is a pattern of exponential growth. Indeed, some 'lucky' explorers are likely to find intrinsically superior islands outside the 'realized economy'. Although they might not be able to adequately exploit the opportunities of the 'new' island by themselves, the 'extraordinary' nature of their discovery might nevertheless induce other agents to move there in the future and, consequently, increase its actual productivity. This allows the system to avoid lock-in, provided that $\epsilon > 0$ and the technological regime is characterized by sufficiently strong opportunities (see Section 5).

Second, in accordance with empirically observed patterns of innovation, diffusion and adoption (see e.g. Dosi, 1982), the model generates *s-shaped* diffusion curves in the number of agents currently mastering a given technology. Moreover, because
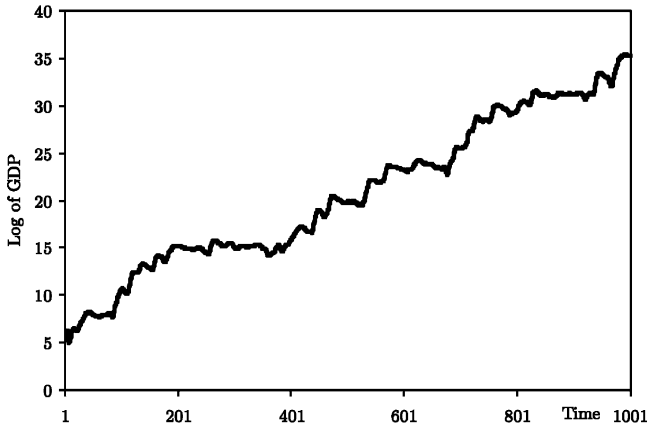
Fig. 3. Patterns of exponential growth in an open-ended economy with exploration. Par. Setup: $N = 100$, $\pi = 0.1$, $\rho = 0.1$, $\alpha = 1.5$, $\epsilon = 0.1$, $\lambda = 1$, $\varphi = 0.5$, $T = 1000$.

many techniques are allowed to coexist over the same time intervals (if they exhibit sufficiently similar realized productivities), one usually detects overlapping diffusion patterns as those depicted in Fig. 4(b). As the set of current available technologies keeps enlarging due to the unceasing process of exploration and innovation, firms migrate toward more productive islands, entailing processes of diffusion, which occur at different rates. These rates typically depend on the characteristics of the technologies involved in the process, the incentives provided by the economic environment and the features of the adopters themselves. In very general terms, the speed at which innovations are adopted (and replaced) is increasing in both their absolute *initial* productivity distance and the extent to which interactions are global. Also, if information is diffused not too locally, radical innovations tend to retain their leadership much longer than incremental ones. Yet, the rate at which innovations are substituted is decreasing with the average willingness to explore of the agents in the system.

5. The sources of self-sustaining growth

The basic conclusion stemming from the analyses presented so far is that patterns of exponential growth might be endogenously generated in the system only if firms are able to explore in an open-ended technological space. In this Section, we study by means of extensive MC exercises how system parameters affect the distribution of long-run average grow rates (AGR):
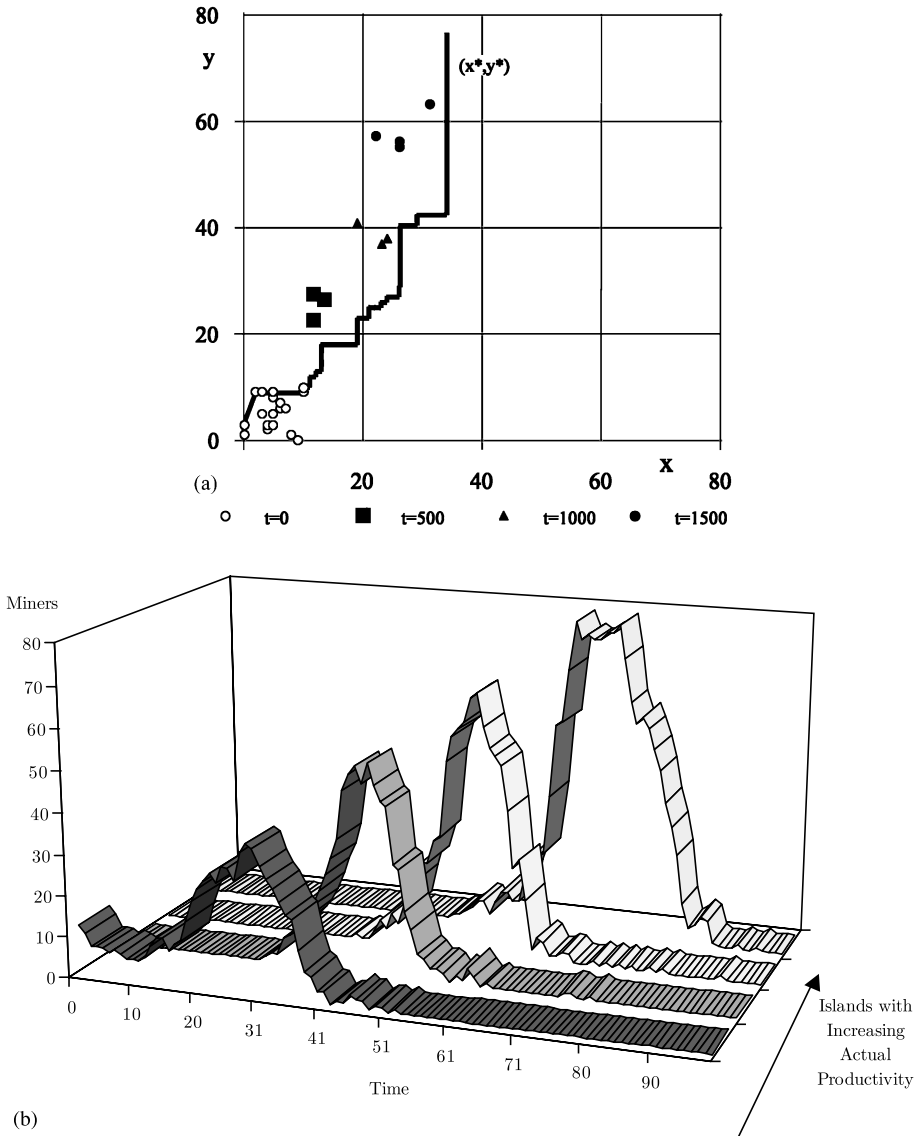
$$g_m(\omega) = \frac{q_{m,T} - q_{m,0}}{T}, \tag{9}$$

Fig. 4. (a) Spatial diffusion of colonized islands and 'Best Practice' proxy (x_t^*, y_t^*). Par. Setup: $N = 100$, $\pi = 0.1$, $\rho = 0.1$, $\alpha = 1.5$, $\epsilon = 0.1$, $\lambda = 1$, $\varphi = 0.5$. (b) Diffusion of technological innovations. An example of overlapping S-shaped patterns of adoption. Par. Setup: $N = 100$, $\pi = 0.1$, $\rho = 0.1$, $\alpha = 1.5$, $\epsilon = 0.1$, $\lambda = 1$, $\varphi = 0.5$, $T = 1000$.

where $m = 1, \ldots, M$ is the MC run, T is the econometric sample-size, $\omega \epsilon \Omega$ is defined in Eq. (8) and $q_{m,t}$ is the log of aggregate GDP at time t. In particular, we will ask how the overall performance of the economy, as measured by the mean of AGR:

$$\bar{g}_M(\omega) = M^{-1} \sum_{m=1}^{M} g_m(\omega), \tag{10}$$

changes in different technological and learning regimes (i.e. in different regions of the parameter space)[9].

A first clear-cut result that MC simulations point out is that—everything else being constant—$\bar{g}_M(\omega)$ appears to be positively influenced by: (a) the extent to which the system is fueled with innovation 'opportunities' (i.e. larger λ and π); (b) the magnitude of path-dependency affecting the innovation process (i.e. larger φ); (c) the degree of globality of the information diffusion in the interaction process (i.e. smaller ρ).

This claim is supported by the surfaces in Fig. 5(a and b) where, for a given choice of α and ϵ, we plot MC mean of AGRs against $(\log_{10}\rho, \varphi)$ in two distinct opportunity setups (i.e. for different values of π and λ). Notice that, as typically happens in evolving complex systems (see Batten, 2000), the causal relationships between system parameters and aggregate variables are characterized by threshold effects and non-linearities (see Fig. 5c and d). On the one hand, path dependence linearly affects the mean of AGRs. On the other hand, as one gradually increases the rate of information diffusion, an abrupt change in AGRs usually arises around $\rho^*(\omega) \cong 1.0$. If $\rho < \rho^*(\omega)$, the performance of the system is barely influenced by ρ. When $\rho > \rho^*(\omega)$, small changes in the degree of locality of interactions bring about dramatic consequences in the mean of AGRs.

Let us turn now to study how the willingness to explore of the system (ϵ) affects AGRs. As intuition suggests, larger AGRs could be attained if the economy somehow manages to optimally solve the trade-off between exploitation and exploration (cf. March, 1991; Allen and McGlade, 1986). However, it turns out that the levels of willingness to explore required to optimally balance between exploitation and exploration strongly depend on the technological and learning regime which characterize the economy. As illustrated in Fig. 6(a–d), four distinct regimes emerge in setups where returns to scale are increasing ($\alpha > 1$). When no interaction takes place ($\rho = \infty$) and opportunities are low, higher exploration rates are totally harmful because agents hardly find radically new practices and, if they do,

---

[9] All results presented below are not affected by the particular choice of the AGRs. Indeed, employing alternative specifications as $g'_m = [(q_{m,T}/q_{m,0})^{1/(T+1)} - 1]$ or $g''_m = [(Q_{m,T}/Q_{m,0})^{1/(T+1)} - 1]$, will only change the scale of attainable growth rates. Moreover, we have chosen values of T in such a way that recursive MC mean and variance of AGR converge. Therefore, properties about $\bar{g}_M(\omega)$ are not influenced by the econometric sample size. Finally, in the chosen range for T and M, the MC variance of AGR is typically negligible. This allows us to avoid reporting confidence intervals for $\bar{g}_M(\omega)$.

Fig. 5. a Mean of MC AGR distributions as a function of , . Low opportunity regime: 1, 0.1. Par. Setup: N100, 1.5, 0.1, T1000, M10000. b Mean of MC AGR distributions as a function of , . High opportunity regime: 5, 0.4. Par. Setup: N100, 1.5, 0.1, T1000, M10000. c Non linear behaior of mean of MC AGR distributions as a function of . Par. Setup: N100, 0.4, 1.5, 0.1, 5, 0.5, T1000, M10000. d Linear behaior of mean of MC AGR distributions as a function of . Par. Setup: N100, 0.4, 1.5, 0.1, 5, 0.1, T1000, M10000.
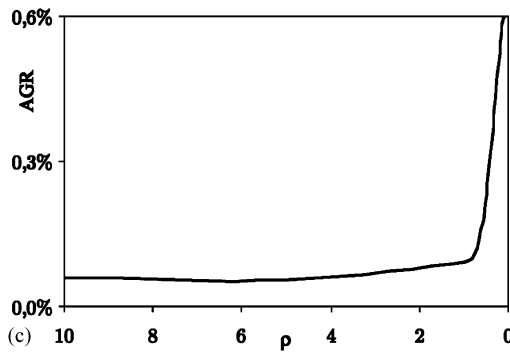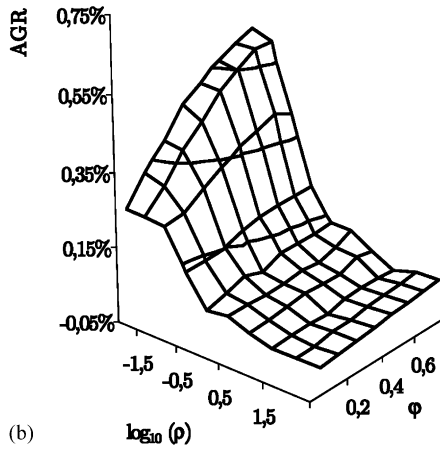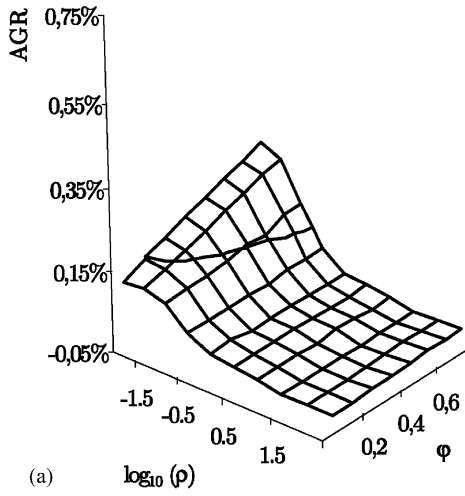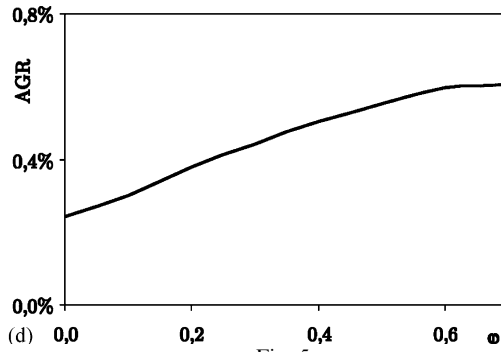
Fig. 5.

(d)

Fig. 5

they cannot benefit from increasing returns to scale. Hence $\bar{g}_M(\omega)$ monotonically decreases with ϵ —no matter the degrees of path-dependence (cf. Fig. 6(a)). Conversely, economies in which information is globally diffused ($\rho = 0$) and innovators strongly benefit from learning by doing (high φ) typically maximize their AGR when all agents commit themselves to exploration and production on new islands only lasts one period (see Fig. 6(b)). Moreover, if information is spread locally—i.e. $0 \ll \rho \ll \infty$ as in Fig. 6(c)—the overall performance of the economy increases either if few explorers are around or if there are very many: in the first case, a large population of miners can continually exploit both increasing returns to scale and incremental, path-dependent, innovations through small-scale migrations driven by local imitation. In the second case, thanks to local information diffusion, small clusters of 'colonized islands' can immediately benefit from the large-scale introduction of innovations.

The most interesting regime, however, arises in all other 'intermediate' settings where MC mean of AGRs are maximized by an interior value of ϵ, cf. Fig. 6(d). The intuition here corresponds to that suggested in March (1991), p. 71. As he points out, systems that engage in exploration to the exclusion of exploitation "exhibit too many undeveloped new ideas and too little distinctive competencies", while, at the opposite extreme, they "are likely to find themselves trapped in sub-optimal stable equilibria". In our model, this condition applies in two setups, namely: (a) agents face very large opportunities but they are unable to completely exploit dynamic increasing returns because information is not spread around; (b) interactions are global but knowledge does not accumulate as the economy evolves. In both situations, higher economic performances cannot be attained by entirely committing either to technological search or to production. As a result, losses stemming from the exploitation–exploration trade-off are reduced by an appropriate balance between the two forces[10].

_____

[10] Simpler patterns arise when one analyzes how different regimes of returns to scale in production affect economic performance. When interactions are shut down ($\rho = \infty$), MC means of AGRs tend to be decreasing with α when opportunities are low (and only mildly increasing for large α's when they are high). Conversely, if information is globally diffused, AGRs are monotonically increasing in α for any (λ, π) and φ.
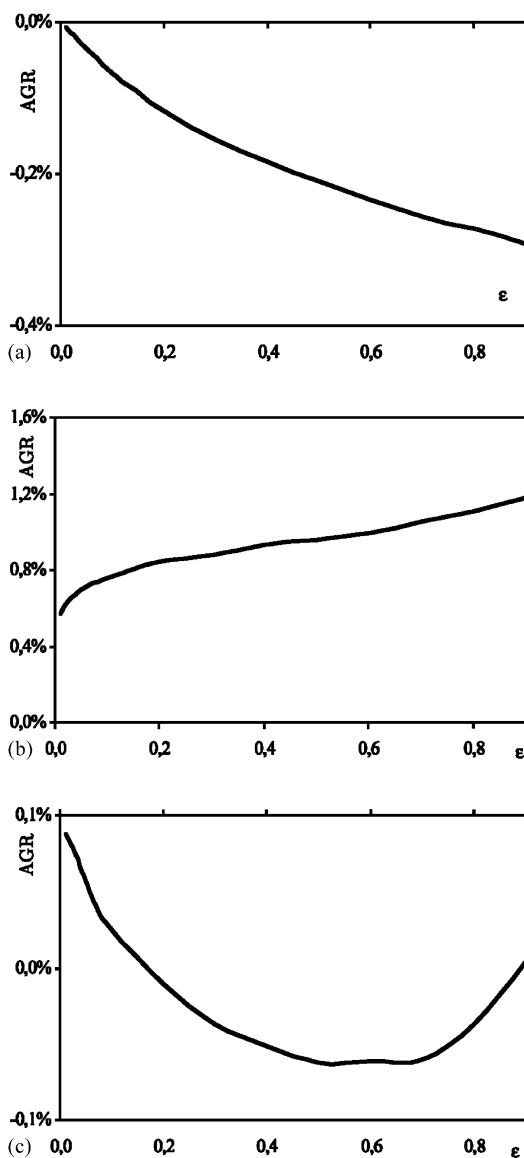
Fig. 6. (a) Mean of MC AGR distributions as a function of the willingness to explore ϵ. Technological regime: $\lambda = 1$, $\pi = 0.1$, $\rho = \infty$, $\varphi = 0.1$. Other parameters: $\alpha = 1.4$, $N = 100$, $M = 10\,000$. (b) Mean of MC AGR distributions as a function of the willingness to explore ϵ. Technological regime: $\lambda = 5$, $\pi = 0.4$, $\rho = 0$, $\varphi = 0.5$. Other parameters: $\alpha = 1.4$, $N = 100$, $M = 10\,000$. (c) Mean of MC AGR distributions as a function of the willingness to explore ϵ. Technological regime: $\lambda = 1$, $\pi = 0.1$, $\rho = 0.1$, $\varphi = 0.5$. Other parameters: $\alpha = 1.4$, $N = 100$, $M = 10\,000$. (d) Mean of MC AGR distributions as a function of the willingness to explore ϵ. Technological regime: $\lambda = 5$, $\pi = 0.4$, $\rho = \infty$, $\varphi = 0.5$. Other parameters: $\alpha = 1.4$, $N = 100$, $M = 10\,000$.
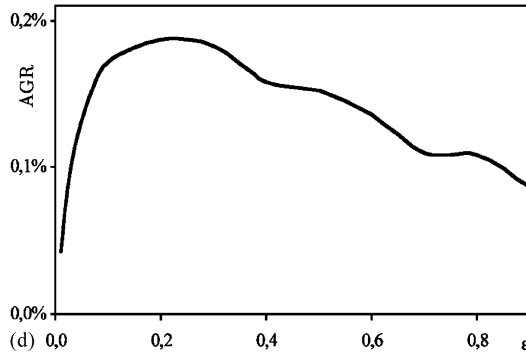
Fig. 6 (*Continued*)

This point arises even more strongly when one allows for heterogeneity in agents' willingness to explore. Consider for instance an economy in which an initial distribution $E = (\epsilon_1, \epsilon_2, \ldots, \epsilon_N)$, $\epsilon_i \in [0, 1]$ and $\epsilon_{i'} \neq \epsilon_{i''}$ for some $i' \neq i''$ is given. To keep things simple, let us suppose that E is such that $\epsilon_i = 0$, $i = 1, 2, \ldots, \lfloor \mu N \rfloor$ and $\epsilon_i = \epsilon_0$, $i = \lfloor \mu N \rfloor + 1$, $\ldots$, N, where $\mu \in [0, 1]$ and $\epsilon_0 \in (0, 1]$. The aggregate consequences of increasing μ's (in terms of economy's AGRs) are once again strictly related to the handling of the exploitation–exploration trade-off, which in turn depends on the prevailing technological and institutional regimes, cf. Fig. 7(a–c). Again, in all 'intermediate' setups described above, AGRs are maximized by some $0 < \mu^*(\omega) < 1$, with poor performances when the economy commits either small or too many resources in the exploration of unknown knowledge bases.

6. Growth rates volatility and system performance

Higher average economic performances are generated in the model if the economy is gradually injected by increasingly powerful sources of growth (i.e. stronger increasing returns to scale, more global knowledge diffusion, higher path-dependency and technological opportunities). It is then of interest to assess how the volatility of aggregate performances (both across MC samples and within time-series realizations) is affected by system parameters governing these forces.

Despite what one could have expected, patterns of self-sustaining growth characterized by higher AGRs are *not* generally associated with overly increasing levels of growth rates volatility. On the one hand, a strong positive correlation emerges between $\bar{g}_M(\omega)$ and MC sample standard deviations:

Fig. 7. a Mean of MC AGR distributions as a function of the share of sedentary agents in the population . Technological regime: low opportunities 1, 0.1, no info diffusion , low path dependency 0.1. Par. Setup: 1.4, 0.1, N100, M10000, T500. b Mean of MC AGR distributions as a function of the share of sedentary agents in the population . Technological regime: mild opportunities 2, 0.2, global info diffusion 0, no path dependency 0. Par. Setup: 1.4, 0.1, N100, M10000, T500. c Mean of MC AGR distributions as a function of the share of sedentary agents in the population . Technological regime: high opportunities 1, 0.1. Par. Setup: 1.4, 0.1, N100, M10000, T500.
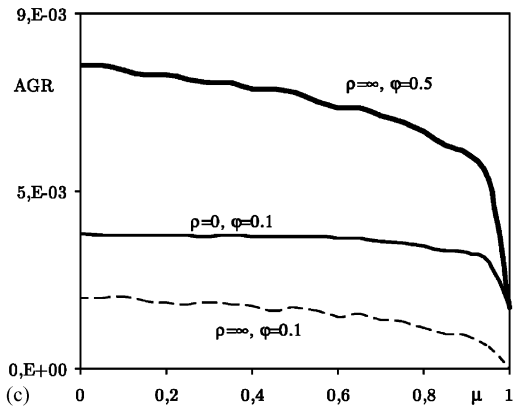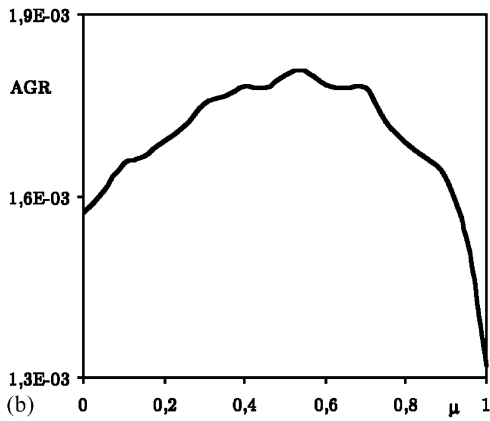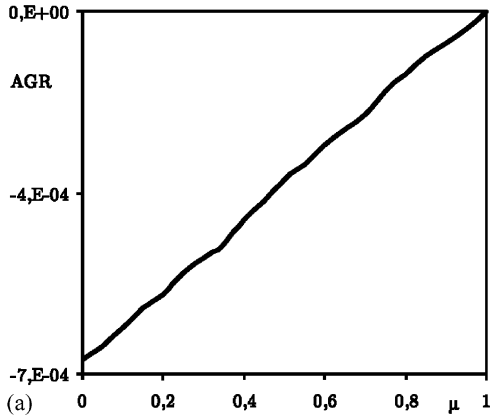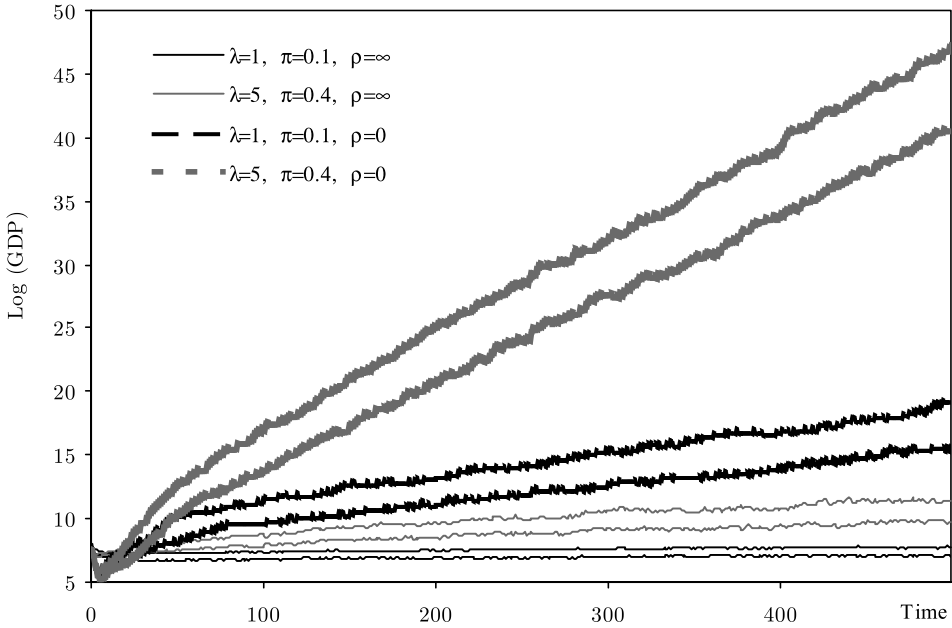
Fig. 7

Fig. 8. Time evolution of 5- and 95-percentile of the MC distribution of log(GDP) time-series observations over $M = 10\,000$ MC simulations in different technological regimes. Par. Setup: $\alpha = 1.5$, $\varphi = 0.4$, $\epsilon = 0.1$, $N = 100$.

$$\sigma_{g_M}(\omega) = [M^{-1} \sum_{m=1}^{M} g_m^2(\omega) - \bar{g}_M^2(\omega)]^{\frac{1}{2}}, \qquad (11)$$

so that the latter appear to be monotonically increasing with λ, π, φ and $-\rho$, everything else being constant[11]. On the other hand, MC sample standard deviations never 'explode' as one increases the strength of the sources of growth. Therefore, despite the self-reinforcing nature of the mechanisms triggering economic growth in the system (i.e. exploration, innovation and more efficient production), the model yields sufficiently ordered growth paths, which turn out neither to overlap nor to converge as long as one considers sets of GDP time-series generated by points in the parameter space far enough from each other. To illustrate this property, Fig. 8 plots time-series describing the 5 and the 95% percentiles of the MC distributions $q_t(\omega) = \{q_{m,t}(\omega), m = 1, \ldots, M\}$, as $t = 1, \ldots, T$, in four different parameter setups ($M = 10\,000$). Notice that even in a global information/high opportunities setup, the band including the 90% of MC observations does not widen as T grows. Moreover, 90% confidence intervals do not overlap even for very small econometric sample sizes.

Let us turn now to the properties of the *within-sample* volatility of growth rates time-series (GRTS) $\underline{h}_m(\omega) = \{h_{m,t}(\omega), t = 1, \ldots, T\}$, where:

---

[11] For a similar property displayed by actual time-series in a cross-section of countries (cf. Fatas, 2000).

$$h_{m,t}(\omega) = \frac{q_{m,t} - q_{m,t-1}}{q_{m,t-1}}. \tag{12}$$

Here a first important result is that, unlike MC sample standard deviations, self-sustaining growth does not always imply a larger volatility in GRTS (for a given econometric sample size T), as measured by the MC mean of its standard deviation:

$$\sigma(\underline{h}_m(\omega)) = \sqrt{T^{-1} \sum_{t=1}^{T} h_{m,t}^2(\omega) - \left[T^{-1} \sum_{t=1}^{T} h_{m,t}(\omega) \right]^2}. \tag{13}$$

In particular, when radical innovations are very likely, setups typically yielding self-sustaining growth (e.g. small ρ's, large φ's) are characterized by a lower magnitude of average volatility, whereas economies usually generating stationary GDP time-series or very mild growth display a higher CRTS variation, cf. Fig. 9(a).

Even more unexpectedly, persistently higher AGRs seem to be attained by the system through a process characterized by GRTS volatility decreasing in time (i.e. across subsequent phases of development). To illustrate this property, consider, as done in Fig. 9(b), four prototypal environments yielding: (a) stationary GDP time-series; (b) levels of GDP evolving around a S-shaped trend; and self-sustaining growth emerging from (c) a low opportunities setup; or (d) a high opportunities setup. As one takes into account the time evolution of MC mean of the distributions of recursive standard deviation of GRTS (i.e. computed over enlarging econometric sub-samples $\{T_0, T_0+1, \ldots, \bar{T}\}$, for $\bar{T} = T_0+20, T_0+21, \ldots, T$ and $T_0 = 50$) a striking pattern arises. Indeed, recursive standard deviation of GRTS appears to behave as $\bar{T}^{\beta-1}$, $\beta > 0$ in each of the above environments. However, while in the stationary GDP case one has $1 \le \beta < 2$, as soon as some evidence of persistent growth emerges in the system, β becomes less than unity and recursive standard deviations turn out to be monotonically decreasing toward some positive constant level. In general, a negative correlation emerges between β and the overall performance of the economy: the more one fuels the system with opportunities and path-dependency, the higher the rate at which GRTS volatility, as measured by average recursive standard deviation, decreases in time.

Therefore, the model seems to account for the appearance, over finite time periods, of distinct 'phases' of development. Under structural conditions above certain thresholds, the economy displays an aggregate dynamics wherein phases of almost steady positive growth rates are punctuated by temporary slowdowns. Exponential growth thus emerges as the outcome of a process leading to 'ordered' GDP time-series characterized by fairly moderate variability both across independent histories and, more importantly, within the sample path.

---

[12] For a critical discussion on trend versus difference stationarity and drawbacks of ADF tests cf. Fagiolo, 2000.
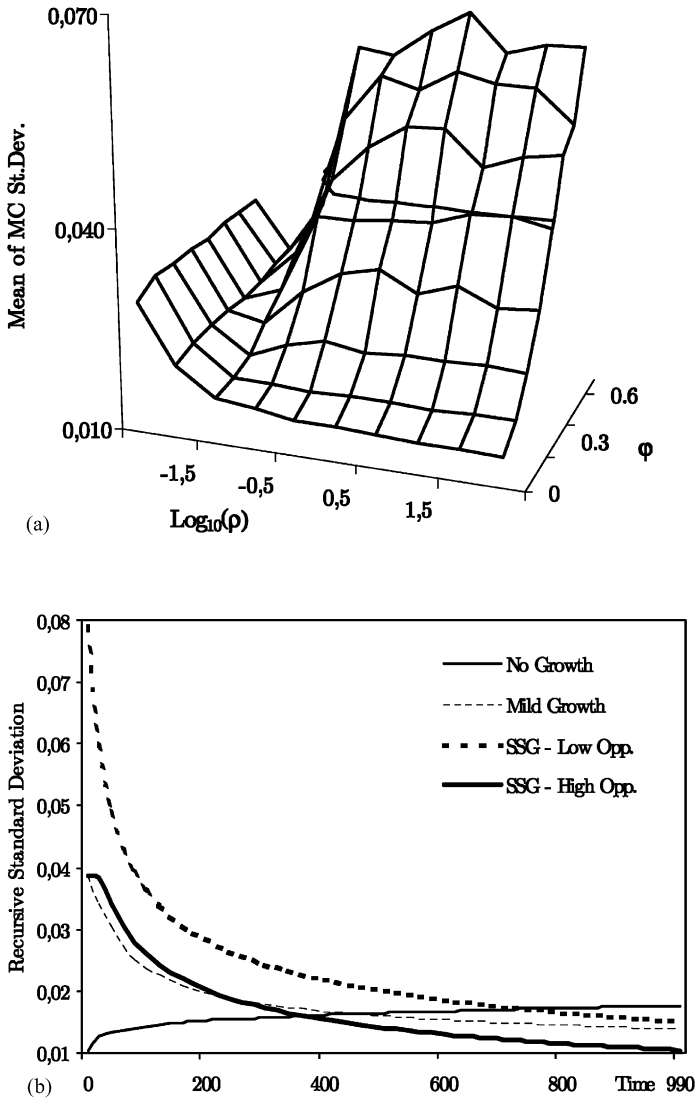
Fig. 9. (a) Mean of MC standard deviations of log (GDP) time-series growth rates as a function of (ρ, φ). High opportunities: $\lambda = 5$, $\pi = 0.4$. Par. Setup: $N = 100$, $\alpha = 1.5$, $\epsilon = 0.1$, $T = 1000$, $M = 10\,000$. (b) Time evolution of GDP time-series growth rates (GRTS) volatility in four paradigmatic growth regimes. Y-axis: MC mean of recursive standard deviations of GRTS (within simulations).

7. Statistical properties of simulated GDP time-series

The foregoing exercises have attempted to shed some light on the mechanisms underlying the emergence of self-sustained growth in the model. In this Section, we will ask whether (and if yes, under which technological regime) the model is able to

generate simulated GDP time-series which display statistically properties similar to those empirically detected in actual output time-series (e.g. non-stationarity, auto-correlation in output growth, persistence of oscillations, etc.).

Let us start to address this 'exercise in plausibility' by noticing that (when they arise) patterns of self-sustaining growth are always associated in the model to 'difference stationary' log(GDP) time-series (as opposed to trend-stationary ones). In fact, according to standard ADF tests—and irrespective of the employed Dickey–Fuller regression specification—one cannot reject the null of a unit root (at 5%), which, on the contrary, is systematically not accepted for both first differences $\Delta q_{m,t}$ and growth rates $h_{m,t} = \Delta Q_{m,t}/Q_{m,t-1}$[12].

Even more interestingly, we find that the ways in which system parameters affect the likelihood of generating $I(1)$ time-series (i.e. patterns of self-sustaining growth) are very similar to the ways in which system parameters affect system performances (i.e. mean of AGRs). Indeed, the behavior of MC mean of ADF(1) test statistics $t_1(q_{m,t}(\omega))$ mimics the one found for MC mean of AGRs as one tunes the relevant system parameters. As Fig. 10(a) shows, $t_1(q_{m,t}(\omega))$ are indeed increasing exponentially with ρ and linearly with φ—for any given levels of the opportunity setup (λ, π), willingness to explore (ϵ) and returns to scale (α). Consequently, the null hypothesis (i.e. presence of a unit-root in the log(GDP) time-series) is accepted with an increasing MC frequency as one tunes up the sources of growth. Accordingly, the portion of the (φ, $\log_{10} \rho$)-plane containing MC frequencies of the ADF(1) test acceptance greater than 90% is larger, the greater the magnitude of opportunities. In fact, a sort of threshold emerges in the (φ, $\log_{10} \rho$)-plane: beyond some given combinations of path-dependency and globality of interactions, the model delivers almost always difference-stationary log(GDP) time-series, cf. Fig. 10(b). Results in line with the analysis in Section 5 also arise when one investigates how the MC acceptance frequency of the ADF tests varies with the magnitude of the willingness to explore (ϵ). Again, in all those cases whereby the economy is characterized either by global interactions and low path-dependency or by no information diffusion and high opportunities, the system seems to be able to generate self-sustaining patterns of growth *only if* a suitable balance between R&D and production is achieved, cf. Fig. 10(c).

As a result, the model is able to deliver quite precise predictions about how technological regimes affect both system performances and the likelihood that a self-sustaining process of growth is triggered in the economy.

In addition to all that, the model is able to robustly generate growth rate time-series whose autocorrelation functions (ACF) are characterized by positive, statistically significant, terms over short horizons and declining toward zero over higher lags[13]. Fig. 11(a and b) shows examples of MC mean (over $M = 1000$ replications) of the ACF of GDP growth rates and the associated estimates for the

---

[13] This is indeed one of the few unquestioned 'stylized facts' in the business cycle literature. This shape for the ACF of quarterly GDP growth rates has been observed in almost all Countries, with some notable exceptions: see Campbell and Mankiw (1989).
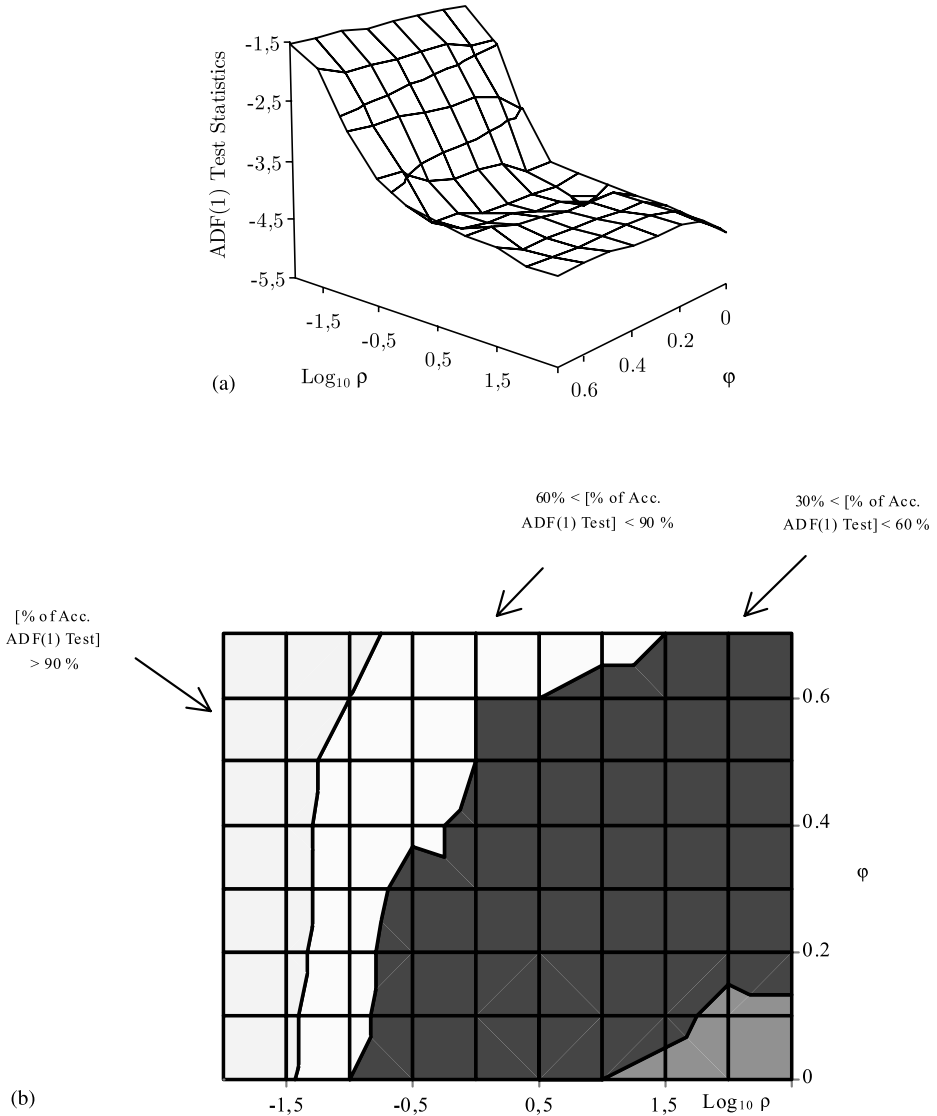
Fig. 10. (a) A MC study of thresholds in the emergence of unit-roots in log(GDP) time-series. Mean of MC ADF(1) test statistics distribution in a high opportunity regime ($\lambda = 5$, $\pi = 0.4$). Critical values: -3.441 (5%); -4.022 (1%). Par. Setup: $\epsilon = 0.1$, $\alpha = 1.5$, $N = 100$, $T = 1500$, $M = 10\,000$. (b) A MC study of thresholds in the emergence of unit-roots in log(GDP) time-series. Frequency of acceptance of the 5%-ADF(1) test in a high opportunity regime ($\lambda = 5$, $\pi = 0.4$). Par. Setup: $\epsilon = 0.1$, $\alpha = 1.5$, $N = 100$, $T = 1500$, $M = 10\,000$. (c) A MC study of thresholds in the emergence of unit-roots in log(GDP) time-series. Frequency of acceptance of the 5%-ADF(1) test as a function of ϵ in a high opportunity, no info diffusion regime with low path dependency ($\lambda = 5$, $\pi = 0.4$, $\rho = \infty$, $\varphi = 0.1$). Par. Setup: $\alpha = 1.5$, $N = 100$, $T = 1500$, $M = 10\,000$.
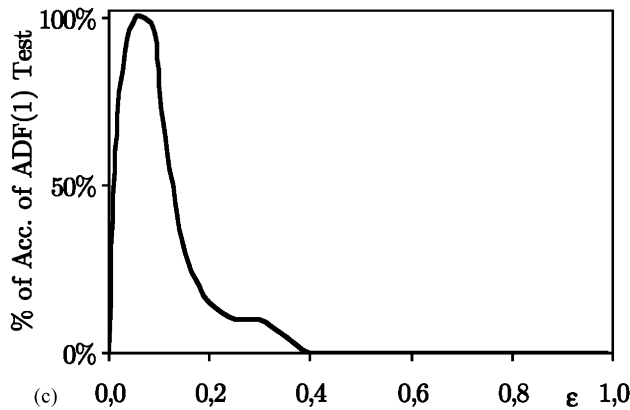
Fig. 10 (*Continued*)

spectral densities in some revealing parameter regions yielding $I(1)$ patterns of log(GDP) time-series. In particular, whenever interactions are global and opportunities are large enough, positive shocks to output growth rates coming from innovations are almost instantaneously spread in the economy, leading to high and positive first- and second-order autocorrelation coefficients, followed by not significant values over longer horizons. Conversely, when interactions are local and opportunities are low, GDP growth does not display ACF coefficients significantly different from zero (as they almost always fall inside 5% Bartlett confidence bands). Accordingly, spectral densities, albeit much smoother than empirical ones, usually display a peak around low frequencies and then tend to decrease as the length of the period becomes small.

   As a further test of the ability of the model to generate simulated output growth time-series with statistical properties similar to the empirically ones, we have investigated whether GDP fluctuations are characterized by a permanent component and, if so, how big such a component might be. Following Campbell and Mankiw (1987) and Cochrane (1988), we have computed non-parametric measures of persistence of GDP fluctuations based on sample estimates of auto-correlations of output growth (cf. the Appendix A for details).

   As Table 1 shows, both measures decrease as the window size k grows, but they generally stabilize around values exceeding unity in all experimented parametrizations. Despite the well-known drawbacks of this estimation procedure, our results, quite in tune with the findings of Campbell and Mankiw (1989), imply that simulated GDP time-series do not appear to revert toward any smooth exogenous trend and exhibit very persistent fluctuations: a 1% shock to output should indeed change the long-run univariate forecast of GDP levels by far more than 1%. Also, persistence turns out to be *higher* the more interactions are global, the *larger* the likelihood of 'radical' innovations and the *smaller* the density of islands in the economy.

   Finally, the model, notwithstanding increasing returns to learning achievements, *does not* display the well-known 'size-effects' characterizing many endogenous
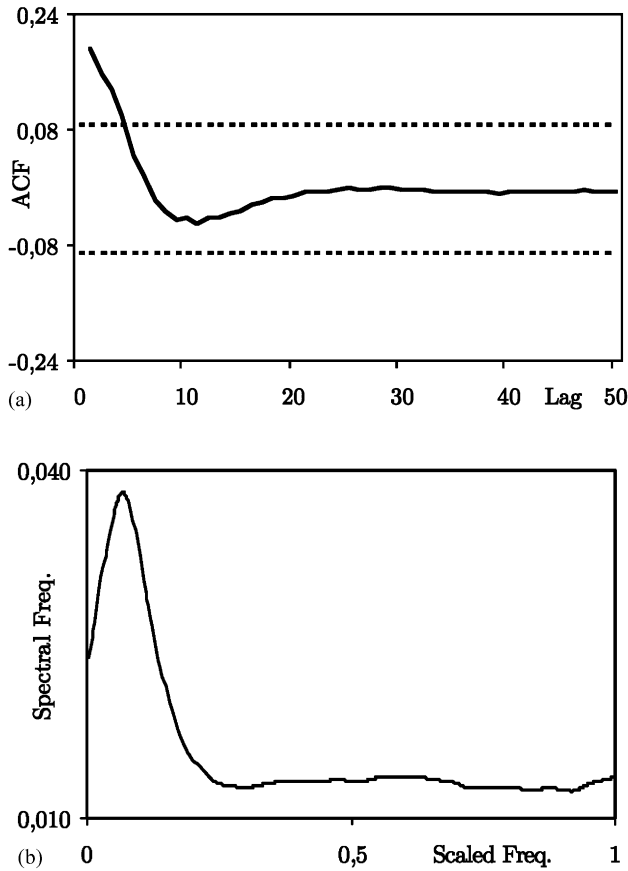
Fig. 11. (a) A MC study of growth rates time-series autocorrelation structure. Mean of MC autocorrelation function. Technological regime: high opportunities ($\lambda = 5$, $\pi = 0.4$), global info diffusion ($\rho = 0$) and high path-dependency ($\varphi = 0.5$). Dotted lines: 95% Bartlett bands. Parameter setup: $\alpha = 0.1$, $\epsilon = 0.1$, $N = 100$, $M = 10\,000$, $T = 1500$. (b) A MC study of growth, rates time-series autocorrelation structure. MC estimate of log(GDP) growth rates spectral density. Technological regime: high opportunities ($\lambda = 5$, $\pi = 0.4$), global info diffusion ($\rho = 0$) and high path-dependency ($\varphi = 0.5$). Frequencies are scaled so as to map the unit interval. Spectra computed by smoothing the periodogram using a Bartlett window with width = 50. Parameter setup: $\alpha = 0.1$, $\epsilon = 0.1$, $N = 100$, $M = 10\,000$, $T = 1500$.

growth models (see Jones, 1995, 1999). Fig. 12 depicts the behavior of MC mean of AGRs as a function of the population size (N) and econometric sample size (T), in a parameter setup usually yielding self-sustaining growth. If any, a weak evidence on *falling* AGRs the larger the size of the economy for a given time-length emerges. Moreover, AGRs do not display any monotone pattern when N and T both increase. The intuition behind this result is that, while ceteris paribus larger economies face potentially higher returns to knowledge exploitation, it also holds that they must cope with higher 'adjustment lags' to new knowledge bases (as

Table 1
MC estimates of persistence in GDP growth

K	$\lambda = 1$		$\lambda = 3$		$\lambda = 5$	
	$\hat{V}^k$	$\hat{A}^k(1)$	$\hat{V}^k$	$\hat{A}^k(1)$	$\hat{V}^k$	$\hat{A}^k(1)$
$\rho = 0,\ \varphi = 0.5$						
$\pi = 0.1$ 10	1.716 (0.294)	1.319	2.157 (0.294)	1.508	2.479 (0.425)	1.645
20	1.637 (0.387)	1.288	1.949 (0.461)	1.433	2.300 (0.544)	1.584
30	1.449 (0.417)	1.212	1.745 (0.502)	1.356	2.141 (0.616)	1.528
40	1.324 (0.438)	1.159	1.625 (0.537)	1.309	2.044 (0.676)	1.493
50	1.262 (0.465)	1.131	1.539 (0.567)	1.274	1.959 (0.722)	1.462
$\pi = 0.4$ 10	1.585 (0.271)	1.263	1.987 (0.340)	1.430	2.285 (0.391)	1.548
20	1.526 (0.361)	1.239	1.776 (0.420)	1.352	2.141 (0.507)	1.498
30	1.371 (0.394)	1.174	1.616 (0.465)	1.289	2.072 (0.596)	1.474
40	1.277 (0.422)	1.133	1.532 (0.507)	1.256	2.009 (0.664)	1.452
50	1.226 (0.452)	1.111	1.467 (0.541)	1.229	1.941 (0.716)	1.427
$\rho = 0.1,\ \varphi = 0.1$						
$\pi = 0.1$ 10	1.149 (0.197)	1.072	1.390 (0.238)	1.180	1.623 (0.278)	1.279
20	1.198 (0.284)	1.095	1.451 (0.343)	1.206	1.641 (0.388)	1.286
30	1.183 (0.340)	1.088	1.387 (0.399)	1.179	1.531 (0.440)	1.242
40	1.155 (0.382)	1.075	1.323 (0.437)	1.151	1.452 (0.480)	1.209
50	1.135 (0.419)	1.065	1.287 (0.475)	1.136	1.410 (0.520)	1.192
$\pi = 0.4$ 10	1.096 (0.188)	1.047	1.334 (0.228)	1.155	1.513 (0.259)	1.232
20	1.171 (0.277)	1.082	1.422 (0.336)	1.193	1.560 (0.369)	1.251
30	1.195 (0.344)	1.093	1.383 (0.398)	1.176	1.478 (0.425)	1.217
40	1.200 (0.397)	1.095	1.337 (0.442)	1.157	1.427 (0.472)	1.196
50	1.203 (0.443)	1.097	1.385 (0.485)	1.147	1.409 (0.520)	1.189

Campbell and Mankiw (1989) statistics (MC standard deviations in parentheses). Par. Setup: $N = 100$, $\alpha = 1.4$, $\epsilon = 0.1$, $M = 10\,000$, $T = 500$. (a) Top: global information ($\rho = 0$) and high path-dependency ($\varphi = 0.5$). (b) Bottom: local information ($\rho = 0.1$) and low path-dependency ($\varphi = 0.1$).

proxied in our model by the time it takes to move a certain fraction of the N agents toward superior islands).

8. Individual rationality and collective outcomes

Given the rather heroic assumptions made on the rational bounds of the agents populating our economy, the model is particularly suited to explore potential conflicts between degrees of individual rationality and collective performances. For instance, what happens to average growth rates if the population of our naïve entrepreneurs is injected with more 'rational' players, which behave on the grounds of some (appropriately defined) expectations about the net returns from exploration?

In order to illustrate this point, consider the following simple example. Assume an economy characterized by: (i) constant returns to scale in production (i.e. $\alpha = 1$); (ii) no knowledge diffusion (i.e. $\rho = \infty$); (iii) no path-dependency in innovation (i.e. $\varphi = $
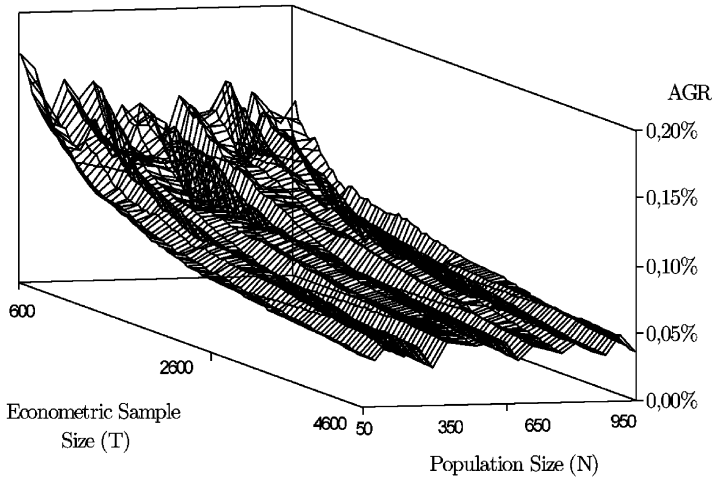
Fig. 12. Mean of MC AGR distributions as a function of econometric sample size (T) and population size (N). Par. Setup: $\lambda = 1$, $\pi = 0.1$, $\alpha = 1.5$, $\epsilon = 0.1$, $\rho = 0.01$, $\varphi = 0.5$, $M = 10\,000$.

0); (iv) all N agents working at time $t = 0$ on a single island ($\ell_0 = 1$) with coordinates (x^*, y^*) and initial productivity $s^* = x^* + y^*$; (v) a constant positive per-period 'transportation' cost $\beta > 0$, which explorers must pay during their search. As far as behavioral assumptions are concerned, let us concentrate on two different settings, namely: (a) the population is composed of N agents behaving according to the rules defined in Section 3; and: (b) the population is made of N copies of a 'representative individual' (RI), with unbounded computational skills and complete information. In particular, assume that the RI knows: (i) the coordinates (x^*, y^*); (ii) the system parameters; (iii) the model of the economy. Although the RI is aware that, on average, the initial productivity of a new island is increasing in its distance from the origin, she does not know where new islands are actually located (i.e. the RI does not have rational technological expectations). Therefore, assume that if she is exploring around the node (x, y), she will make use of a more 'rational' exploration rule putting equal probability on the nodes $(x+1, y)$ and $(x, y+1)$. Finally, suppose for simplicity that the intertemporal discount rate is zero.

At time $t = 1$, the problem for the RI is to decide whether to continue to produce the good at time $t = 2$ or to start exploring. In the first case, she will get a per-period net output from mining equal to $Q_M = s^*$. In the second case, the expected per-period net output from exploration will be: $Q_E = [(1+\lambda)(s^* + \tau) - \beta\tau]/\tau$ where $\tau = 1/\pi$ is the expected length of exploration (or, equivalently, the expected distance between (x^*, y^*) and a new island). Then, the RI will decide to remain on island (x^*, y^*) if and only if $Q_M > Q_E$, i.e. iff:

$$\pi < \frac{1}{1 + \lambda} - \frac{1}{s^*} + \frac{\beta}{(1 + \lambda)s^*} = \pi^*(\beta, \lambda, s^*) \tag{14}$$

As one can easily check, $\pi^*(\beta, \lambda, s^*)$ is increasing in β, decreasing in λ, and

increasing in s^* if $\lambda > \beta - 1$ (i.e. if opportunities are large enough compared with 'transportation' costs). Notice that if $s^* \to \infty$ the *RI* will always stay on (x^*, y^*), while if $\lambda \to \infty$ she will always leave.

Consider now, for any value of s^*, the set of (β, λ) satisfying Eq. (14) for some $\pi \in (0, 1)$. In such a parameter region, the *RI* will decide to continue to work as a 'miner' and get a constant output $Q_M = s^*$. On the contrary, any economy characterized by the same $(\beta, \lambda, \pi, s^*)$ and composed of myopic agents behaving as in setup (a) above (cf. Section 3), will face a rather 'poor' environment, in which there is neither knowledge diffusion, nor path-dependency in innovation, nor increasing returns to scale. Furthermore, let us assume that our 'naïve' agents are characterized by a very low 'willingness to explore' (i.e. $\epsilon = 0.05$). Notwithstanding all that, as Fig. 13 shows, the economy is able to get a per-capita net output persistently greater than Q_M.

Thus, even in this very simple setting, higher collective performance are generated despite (or thanks to) 'irrational' individuals. Notice that this property adds to another one shared with 'new growth' literature (cf. Aghion and Howitt, 1992), according to which in presence of externalities or dynamic increasing returns a systematic divergence between endogenously generated growth rates and socially optimal ones is likely to emerge.
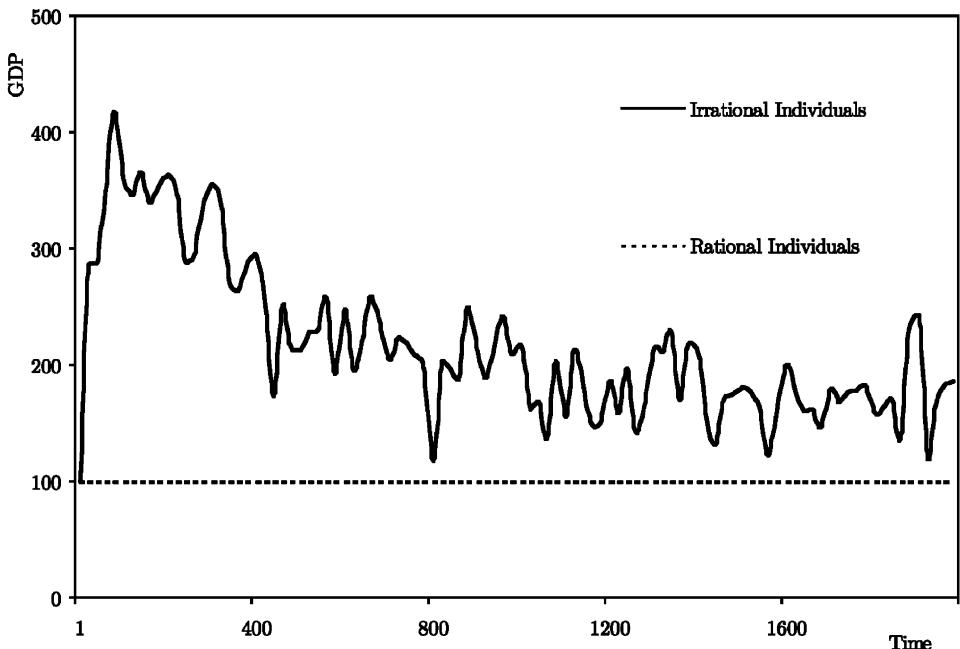

Fig. 13. Individual vs. collective rationality: a simple example. GDP time series generated by irrational and 'more rational' firms. Par. Setup: $s^* = 100$, $N = 100$, $\epsilon = 0.05$, $\varphi = 0$, $\lambda = 5$, $\pi = 0.15$, $\rho = \infty$, $\alpha = 1$.

9. Conclusions

The paper presents a simple model in which self-sustaining growth endogenously emerges, under suitable technological and behavioral conditions, as the result of imperfect coordination among stylized, boundedly-rational, heterogeneous, firms locally interacting in an open-ended technological space.

The model shows that the very possibility of notionally unlimited (albeit unpredictable) technological opportunities is a necessary condition for patterns of persistently fluctuating exponential growth to be generated in the economy.

In that circumstance, exponential growth is attained whenever technological opportunities (as captured by both the density of 'islands' π and the mean of Poisson jumps to radically new paradigms λ), path-dependency (i.e. the fraction of idiosyncratic knowledge, φ, that agents are able to carry over to newly discovered technologies) and globality of interactions in the information diffusion process ($-\rho$), are beyond identifiable thresholds. In that region of the parameter space, the system goes through subsequent phases of development and exhibits ordered GDP time-paths characterized by small growth-rates volatility. Furthermore, the overall performance of the economy appears to be monotonically increasing in any of the latter sources of growth, with the degree of globality of interactions engendering a strong threshold effect in the average performance of the system.

A trade-off between exploitation of the fundamentals and exploration of still unknown technologies clearly emerges, however, when one investigates how growth is affected by the propensity to explore (ϵ). In well-defined technological regimes, the system generates self-sustaining patterns of growth and higher overall performances only if a suitable balance between R&D and production is achieved.

As mentioned, the model could be considered as a sort of 'reduced form' evolutionary model, with an almost exclusive emphasis upon the learning/diffusion aspects of economic evolution, while repressing the competition/selection features of market interactions. Although the limitations stemming from this assumption are quite obvious (for example, the 'microeconomics' is bound to be rather limited), the model is nonetheless able to generate GDP time-series with statistical properties which robustly replicate a few of the stylized facts of macroeconomic dynamics (including GDP growth autocorrelation structure and persistence of fluctuations). Moreover, it does so in ways that are in tune with general microeconomic stylized facts, such as time-consuming diffusion of innovations and persistent asymmetries in efficiency among agents. All this represents indeed a significant interpretative advantage vis-à-vis 'new-growth' theories. Furthermore, unlike most of the latter, the foregoing model is also able to avoid size effects (i.e. the positive influence of the sheer size of an economy upon its growth).

Many extensions of the model can be conceived. First, one should try to explore how the results presented here change with the introduction of a demand side (and, thus, of some 'Keynesian' features). Second, one might likewise study the relevance of adding explicit selection processes affecting the frequency in the population (i.e. the size) of different agents, which are 'carriers' of different technologies. Finally,

given reasonable rules of interaction between economies, one may also investigate convergence–divergence issues across-countries.

Acknowledgements

Appendix A

Assume that the change in log of GDP Δq_t follows a stationary *ARMA* process, so that $\Delta q_t = A(L)v_t$, with v_t a white-noise process. Following Cochrane (1988), Campbell and Mankiw (1987, 1989), we computed estimates of the following persistence measures:

$$V^k = \frac{1}{k+1} \frac{Var(q_{t+k+1} - q_t)}{Var(q_{t+1} - q_t)} = \left[1 + 2\sum_{j=1}^{k}\left(1 - \frac{j}{k+1}\right)\rho_j\right],$$

where ρ_j is the jth autocorrelation coefficient of Δq_t and $A(1) = 1 + \Sigma_{j=1}^{\infty} A_j$.

If $\{q_t\}$ were even more persistent than a random walk, both $A(1)$ and V^k would exceed unity. Moreover, as:

$$A(1) = \sqrt{\frac{V}{1 - R^2}}$$

where $V = \lim_{k \to \infty} V^k \equiv 1 + 2(\rho_1 + \rho_2 + \ldots)$ and $R = 1 - Var(v_t)/Var(\Delta q_t)$, estimation of V^k and $A(1)$ can be done non-parametrically employing sample estimates of the autocorrelation function, i.e. $r_j = \hat{\gamma}(j)/\hat{\gamma}(0)$. An estimate of V^k (consistent for V if k is large) is found simply by replacing population auto-correlations with sample counterparts (after having corrected by a downward bias), i.e.:

$$\hat{V}^k = \frac{T-k}{T}\left[1 + 2\sum_{j=1}^{k}\left(1 - \frac{j}{k+1}\right)r_j\right],$$

while $A(1)$ is estimated (for large k) by:

$$\hat{A}^k(1) = \sqrt{\frac{\hat{V}^k}{1 - r_1^2}}.$$

Notice also that since r_1^2 underestimates R^2, $\hat{A}^k(1)$ tends to underestimate $A(1)$ for large k. Also, the standard error of $\hat{V}^k$ is equal to:

$$S.E.(\hat{V}^k) = V^k \left[\frac{3T}{4(k+1)} \right]^{-\frac{1}{2}}.$$

Finally, both MC standard deviations and theoretical standard errors are increasing with V^k (or with its estimate).

References

Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. Econometrica 60 (2), 323–351.

Allen, P., McGlade, J., 1986. Dynamics of discovery and exploitation: the case of the scotian shelf groundfish fisheries. Canadian Journal of Fishery and Aquatic Sciences 43, 1187–1200.

Arthur, W., 1994. Increasing Returns and Path-Dependency in Economics. University of Michigan Press, Ann Arbor.

Batten, D., 2000. Discovering Artificial Economies. Westview Press, Oxford.

Campbell, J., Mankiw, N., 1987. Are output fluctuations transitory. Quarterly Journal of Economics 102, 857–880.

Campbell, J., Mankiw, N., 1989. International evidences on the persistence of economic fluctuations. Journal of Monetary Economics 23, 319–333.

Chiaromonte, F., Dosi, G., 1993. Heterogeneity, competition and macroeconomic dynamics. Structural Change and Economic Dynamics 4 (1), 39–63.

Cochrane, J., 1988. How big is the random walk in gnp. Journal of Political Economy 96, 893–920.

Conlisk, J., 1996. Why bounded rationality. Journal of Economic Literature 34, 669–700.

David, P., 1992. Path-dependence and predictability in dynamic systems with local externalities: a paradigm for historical economics. In: Foray, D., Freeman, C. (Eds.), Technology and the Wealth of Nations. Pinter, London.

Dosi, G., 1982. Technological paradigms and technological trajectories. Research Policy 11 (2), 147–162.

Dosi, G., Lovallo, D., 1998. Rational entrepreneurs or optimistic martyrs? Some considerations on technological regimes. In: Garud, R., Nayyar, P., Shapiro, Z. (Eds.), Foresights and Oversights in Technological Change. Cambridge University Press, Cambridge.

Dosi, G., Freeman, C., Fabiani, S., 1994. The process of economic development: introducing some stylized facts and theories on technologies, firms and institutions. Industrial and Corporate Change 3, 168–203.

Dosi, G., Marengo, L., Fagiolo, G., 2003. Learning in evolutionary environments. In: Dopfer K. (Ed.), The Evolutionary Principles of Economies. Cambridge University Press, Cambridge, in press.

Durlauf, S., Quah, D., 1998. The new empirics of economic growth. Discussion Paper No. 384. London School of Economics, Centre for Economic Performance, London.

Epstein, J., Axtell, R., 1996. Growing Artificial Societies: Social Science from the Bottom-Up. MIT Press and Brooking Press, Washington, DC.

Fagiolo, G., 1998. Spatial interactions in dynamic decentralized economies: a review. In: Cohendet, P., Llerena, P., Stahn, H., Umbhauer, G. (Eds.), The Economics of Networks. Interaction and Behaviours. Springer, Berlin, Heidelberg.

Fagiolo, G., 2000. Endogenous growth in open-ended economies with locally interacting agents. Working Paper 2000/07. LEM—St. Anna School of Advanced Studies, Pisa, Italy.

Fatas, A., 2000. Endogenous growth and stochastic trends. Journal of Monetary Economics 45 (1), 107–128.

Freeman, C., 1982. The Economics of Industrial Innovation, second ed. Francis Pinter, London.

Freeman, C., 1994. The economics of technical change. Cambridge Journal of Economics 18, 463–514.

Grossman, G., Helpman, E., 1991. Trade, knowledge spillovers and growth. European Economic Review 35, 517–526.

Jones, C., 1995. R and d based models of economic growth. Journal of Political Economy 103 (4), 759–784.

Jones, C., 1999. Growth with or without scale-effects. American Economic Review 89 (2), 139–149.

Jovanovic, B., 1997. Learning and growth. In: Kreps, D., Wallis, K. (Eds.), Advances in Econometrics: Theory and Applications. Cambridge University Press, Cambridge.

Jovanovic, B., Rob, R., 1989. The growth and diffusion of knowledge. Review of Economic Studies 56 (4), 569–582.

Kirman, A., 1992. Whom or what does the representative individual represent. Journal of Economic Perspectives 6 (2), 117–136.

Kirman, A., 1993. Ants, rationality and recruitment. Quarterly Journal of Economics 108, 137–156.

Kirman, A., 1998. Economies with interacting agents. In: Cohendet, P., Llerena, P., Stahn, H., Umbhauer, G. (Eds.), The Economics of Networks. Interaction and Behaviours. Springer, Berlin, Heidelberg.

Lane, D., 1993a. Artificial worlds in economics: part 1. Journal of Evolutionary Economics 3, 89–107.

Lane, D., 1993b. Artificial worlds in economics: part 2. Journal of Evolutionary Economics 3, 177–197.

March, J., 1991. Exploration and exploitation in organizational learning. Organization Science 2, 71–87.

McGrattan, E., Schmitz, J., 1998. Explaining cross-country income differences, Staff Report 250, Federal Reserve Bank of Minneapolis.

Nelson, R., 1995. Recent evolutionary theorizing about economic change. Journal of Economic Literature 33, 48–90.

Nelson, R., 1998. The agenda for growth theory: a different point of view. Cambridge Journal of Economics 22 (4), 497–520.

Nelson, R., Winter, S., 1982. An Evolutionary Theory of Economic Change. The Belknap Press of Harvard University Press, Cambridge.

Romer, P., 1990. Endogenous technological change. Journal of Political Economy 98 (5), 71–102.

Rosenberg, N., 1982. Inside the Blackbox. Cambridge University Press, Cambridge.

Rosenberg, N., 1994. Exploring the Black Box: Technology, Economics and History. Cambridge University Press, Cambridge.

Silverberg, G., Verspagen, B., 1994. Collective learning, innovation and growth in a boundedly rational, evolutionary world. Journal of Evolutionary Economics 4, 207–226.

Silverberg, G., Verspagen, B., 1996. Evolutionary theorizing on economic growth. In: Kuusi, O. (Ed.), Innovation Systems and Competitiveness. ETLA, Helsinki.

Silverberg, G., Verspagen, B., 1997. From the artificial to the endogenous: modeling evolutionary adaptation and economic growth. In: Helmstadter, E., Perlman, M. (Eds.), Behavioral Norms, Technological Progress, and Economic Dynamics: Studies in Schumpeterian Economies. University of Michigan Press, Ann Arbor.

Stiglitz, J., 1987. Learning to learn, localized learning and technological progress. In: Dasgupta, P., Stoneman, P. (Eds.), Economic Policy and Technological Performance. Cambridge University Press, Cambridge.

Stoneman, P. (Ed.), Handbook of the Economics of Innovation and Technological Change. Blackwell, Oxford 1995.

Verspagen, B., 1993. Uneven Growth between Interdependent Economies: An Evolutionary View. Aldershot, Avebury.