*Asian Journal of Economics and Finance*. **2022**, 4, 4 : 381-409
ISSN: 2582-340X
*https://DOI: 10.47509 /AJEF.2022.v04i04.03*

**ARF INDIA**
Academic Open Access Publishing
*www.arfjournals.com*

# Learning to Grow: A Learning-Based Multi-Agent Endogenous Growth Model

**Orlando Gomes**

*Lisbon Accounting and Business School – Lisbon Polytechnic institute (ISCAL-IPL) and CEFAGE (Univ. Évora - ISCAL) research center. E-mail: omgomes@iscal.ipl.pt*

ARTICLEINFO

**Received: 27 July 2022**

**Revised: 20 August 2022**

**Accepted: 26 August 2022**

**Online: 10 November 2022**

***To cite this paper:***

**Orlando Gomes (2022). Learning to Grow: A Learning-Based Multi-Agent Endogenous Growth Model. *Asian Journal of Economics and Finance.* 4(4), 381-409. https://DOI: 10.47509/ AJEF.2022.v04i04.03**

***Abstract:*** To explain economic growth, economists typically resort to a simple and stylized optimal control problem. This problem is solved by a representative agent, who plans, at a pre-specified initial date, how to allocate resources over time, between consumption and savings, with the objective of maximizing intertemporal utility. It is assumed that the agent is rational, well informed, and capable of planning over long horizons (eventually an infinite horizon). In this study, the benchmark economic growth model is revisited and reinterpreted after relaxing some of its key assumptions. In particular, the individual agent in the adapted growth model will not address a single long-term problem; instead, a sequence of short-run decisions has to be pondered. Moreover, the recurrent decisions on how much to save, that the agent will have to make, are shaped by a learning process emerging from the systematic comparison between the agent's own utility and the utility accomplished by other agents in the system. Additionally, it is assumed that the individual agent has imperfect knowledge about the choices of others, what introduces an element of stochasticity into the model. Simulations reveal that approaching standard growth theory through the lens of a learning-based multi-agent system offers important new insights about the growth process, including an opportunity to integrate in a single framework long-term growth and short-term irregular and unpredictable business fluctuations.

***Keywords:*** Economic growth; Savings; Intertemporal utility; Multi-agent system; Learning; Aggregate fluctuations.

***JEL classification:*** O41; E71.

*"Complexity is not a theory but a movement in the sciences that studies how the interacting elements in a system create overall patterns, and how these overall patterns in turn cause the interacting elements to change or adapt. It might study how individual cars together act to form patterns in traffic, and how these patterns in turn cause the cars to alter their position. Complexity is about formation -- formation of structures -- and how this formation affects the objects causing it. (…) This is often a difficult question; we are asking how a process is created from the purposed actions of multiple agents."*

*W.B. Arthur, 2015, p. 3.*

## 1. Introduction

The definition of complexity in the above opening citation is comprehensive enough to cover most of the collective entities and collective enterprises that one finds in nature and society. Any setting involving relatively small particles that freely interact in a decentralized way (i.e., without the need for any central coordination), and whose interaction generates a unique aggregate outcome or unique aggregate pattern, might fit in such definition.

There is, though, an important distinguishing feature between complex natural or engineering systems and socio-economic systems, as the one to be addressed in this article. Differently from the former, the latter might be interpreted as containing a second layer of complexity, because the particles underlying the behavior of the system are not, in this case, innate objects, relatively elementary living organisms, or automata. As pointed out by Hommes (2021), they are people, who are endowed with the ability to reason, to plan for the future, and to establish elaborate forms of communication and interaction.

The degree of complexity and sophistication one should take into account to approach human choices and human interaction is, nevertheless, an unresolved and controversial theme in Economics. Most of the economic theory, built upon the auspices of orthodox neoclassical thinking, evolved under the premise of hyper-rationality, a premise that neglects most of the elements typically associated with a complex system (heterogeneity, interaction, learning, path-dependence, out-of-equilibrium dynamics, and emergence).

To a large extent, a significant part of the scientific breakthroughs in Economics over the last decades are grounded on the notion of optimal rational behavior, a convenient paradoxical assumption that imposes a pronounced contrast between the alleged inexorable capabilities of the human mind (interpreted as a flawless machine) and the simplicity and straightforwardness of the aggregate outcome they typically promote.

Notwithstanding, despite the conceptual usefulness and analytical practicability enclosed in the optimal rationality paradigm, one must recognize that decisions made by flesh-and-bone economic agents are complex, multifaceted, and constrained at many levels. Recent literature in Economics, which advocates the transition from the neoclassical thinking paradigm to the Economics of complexity (Delli Gatti *et al.*, 2010; Holt *et al.*, 2011; Stiglitz and Gallegati, 2011; Bezemer, 2012; Fagiolo and Roventini, 2017), highlights precisely this point: the volume of information and attentiveness required to approach particular choices and particular interaction processes is so overwhelming that it is not reasonable to admit that agents are capable of pursuing optimal courses of action. Hence, hyper-rationality is not a viable and credible simplifying assumption; one can, at most, assume some mild form of rational thinking, which is limited and constrained both in time and in space.

Economists have always been aware of the cognitive and environmental constraints affecting deliberative processes. Despite this, in fact, they have chosen, for many decades, to stick with the rationality paradigm. Although failing to reflect, in many circumstances, the actual behavior of economic agents, supporters of the mainstream view advocate that it leads to rigorous science, namely by avoiding the so-called wilderness of bounded rationality that underlies the complexity view and which leaves to the modeler too many degrees of freedom in setting the foundational principles required for the analysis of economic phenomena (Sims, 1980; Hommes, 2006, Lengnick, 2013).

Nevertheless, rigorous science that explains no actual real-world facts and events is of little practical use. This is why the complexity approach is gaining influence and researchers are progressively starting to consider, in their frameworks of analysis, a series of less conventional elements, from multi-dimensional agent heterogeneity (Kirman, 2006; Chen, 2012) to decentralized network-based interaction rules (Birke, 2009; Bargigli and Tedeschi, 2014; Bramoullé *et al.*, 2014), and also a wide variety of learning mechanisms (Athey, 2018; Mosavi *et al.*, 2020; Babenko *et al.*, 2021).

In this study, the standard economic growth model, one of the models that best symbolizes the orthodoxy of economic thought (Solow, 1956; Cass, 1965; Koopmans, 1965; Lucas, 1988; Romer, 1990), is adapted to a complexity scenario, through the introduction of a few non-conventional assumptions. Specifically, the growth problem is converted into a learning-based multi-agent system, where aggregate results are not the outcome of the optimal decisions of a rational representative agent, emerging instead from a process of systematic interaction and mutual learning across a potentially large population of agents.

A learning-based multi-agent system is defined, in the scientific literature, as a complex modelling structure or modelling device within which a possibly large number of independent and autonomous agents interact in a decentralized fashion, i.e., without the need for any central planner to intervene and coordinate actions (Khalil *et al.*, 2015; Anandakumar and Arulmurugan, 2019). The contact established among agents in this decentralized setup leads to the emergence of an aggregate outcome that might or might not coincide with a given social goal (e.g., the maximization of a social welfare measure).

Typically, in the learning-based multi-agent system, individual decision-makers share a common environment, which is dynamic in the sense that it is reshaped whenever agents modify their actions. The behavior of the agents, in turn, is an adaptive behavior: it evolves through the exploration of the environment, the observation of the actions of others and, ultimately, the incorporation of a learning algorithm in their decision process.

The notion of learning typically associated with this type of system design is the notion of reinforcement learning. Reinforcement learning is a process of trial-and-error through which the individual agent gradually discovers, within a dynamic environment, which is the best course of action to take. This process of discovery is associated with the observation and mimetization of the behavior of others, i.e., with a mechanism of learning (Kaelbling *et al.*, 1996; Szepesvári, 2010; Wiering and Otterlo, 2012; Sutton and Barto, 2018).

Although learning-based multi-agent systems might be thought and conceived especially in the context of purely computational or engineering problems (Borah and Talukdar, 2019; Hou *et al.*, 2021), it should be evident how they associate well with social contexts, and particularly with economic settings. As emphasized by Wolpert *et al.* (1999) and Ellowitz (2008), an economy is nothing more than a system populated by intelligent agents who interact with one another and, most importantly, learn with one another.

To adapt the standard growth framework to a learning-based multi-agent environment, the following non-standard assumptions will be taken into consideration:

i) The representative agent of the standard growth model is replaced by a group of individual decision-makers, each one solving her own consumption utility maximization problem and each one taking her own savings decisions;

ii) Individual agents are unable to make long-term plans. Instead, they solve a sequence of short-run decision problems. In other words, agents are short-sighted, systematically re-evaluate their circumstances, and eventually revise plans;

iii) At the beginning of each planning date, the agent has to make a decision about how much to save from the current planning period to the next one. The decision on the amount of savings at the terminal date is made by comparing the agent's own utility with the level of utility obtained by the other agents. The individual learns by observing the performance of others and the way in which they choose and behave, adapting her behavior accordingly;

iv) The individual agent is unable to observe with certainty the savings rates of other agents, namely those performing better in what concerns utility outcomes. Therefore, although the agent may desire to approximate her own savings rate to the ones of the individuals obtaining the highest levels of utility, such approximation might fail due to incomplete knowledge or incomplete information.

The growth problem, reshaped to include the above assumptions, is formulated in simple terms, taking into consideration a constant marginal returns AK production function (Rebelo, 1991). Therefore, the model can be classified as an endogenous growth model, a class of growth models thoroughly discussed in the growth literature and presented in detail in most of the textbooks in the field (Barro and Sala-i-Martin, 2004; Acemoglu, 2009; Alogoskoufis, 2019; Romer, 2021).

The simple decision problem of the individual agent turns, under the learning-based multi-agent system interpretation, into a complex problem for the whole economy, because of the interaction and learning processes that emerge. Again, this is an approach that goes in the direction of a complexity view of the economy (simple agents acting and interacting in a sophisticated environment) and away from the orthodox neoclassical view (hyper-rational agents acting in a plain and easy to process world).

The remainder of the paper is organized as follows. Section 2 revisits the standard AK endogenous growth model and characterizes its main features. Section 3 presents the solution of the model for each planning period and introduces the end-of-period savings rate as the element of heterogeneity that allows to distinguish agents from one another. In section 4, the sequence of periods is linked to one another by establishing a connection between inter-period savings and the initial stock of capital at each period, and by presenting the rule through which learning takes place. Section 5 proceeds with the simulation of the model, allowing to observe how the introduced changes imply the emergence of endogenous fluctuations for aggregate variables that did not exist in the original setting. Section 6 discusses a couple of additional new features that might be integrated into the model. Section 7 concludes.

## 2. The Utility Maximization Problem of the Individual Agent

Conceive an economy populated by a finite number of decision-makers indexed by j. Each of these decision-makers solves a standard optimal control problem, in which utility maximization is subject to a dynamic capital accumulation constraint. Every agent faces an identical problem, except for a single source of heterogeneity, to be characterized further below, which is associated with the savings decisions that, periodically, agents will have to take.

The formulation of the decision problem is standard in growth theory and follows the typical textbook presentation (e.g, Alogoskoufis, 2019; or Romer, 2021). The endogenous variables of the model are the stock of capital per efficiency unit of labor, $k_j(t) \geq 0$, and consumption per efficiency unit of labor, $c_j(t) \geq 0$. Income is generated through a constant marginal returns production function, $y_j(t) = Ak_j(t)$, with $A \geq 0$ a productivity index. Population and labor efficiency both grow at constant rates, respectively $n \geq 0$ and $g \geq 0$. Capital depreciates linearly, at rate $\delta \in (0,1)$. Therefore, the capital accumulation equation, which corresponds to the constraint of the maximization problem, will take the form of the following ordinary differential equation,

$$\dot{k}_j(t) = Ak_j(t) - c_j(t) - (n + g + \delta)k_j(t), \quad k_j(0)\text{given} \tag{1}$$

The objective of the agent consists in maximizing the flow of utility levels from the initial date of the planning problem, $t = 0$, to some future horizon $t = T$. In standard growth models, T tends to infinity; in the model to explore in this study, although the agent is confronted with an infinite horizon, it is assumed that the decision-maker is boundedly rational and short-sighted, thus splitting the whole planning problem in a sequence of short planning problems of length 1. The intertemporal utility expression is as follows (with $\rho \geq 0$ the agent's rate of time preference),

$$U_j(0) = \int_0^T e^{-\rho t} u\left[\frac{C_j(t)}{L_j(t)}\right] L_j(t)dt \tag{2}$$

In equation (2), agent $j$ is interpreted as a household of dimension $L_j(t) > 0$. Hence, the utility of the agent corresponds to per capita consumption utility multiplied by the respective dimension (in this problem, $C_j(t)$ is the overall consumption of the household). Each household grows at the population growth rate n. Recovering the above-mentioned rate of growth of labor efficiency, expression (2) is equivalent to:

$$U_j(0) = \int_0^T e^{-(\rho-n)t} u\left[e^{gt} c_j(t)\right]dt \tag{3}$$

To display equation (3), the initial levels of population (household dimension) and labor efficiency were normalized to 1, in order to simplify notation. The instantaneous utility function in (3) is defined as a standard constant elasticity of intertemporal substitution (CEIS) utility function of the form,

$$u\left[e^{gt} c_j(t)\right] = \frac{\left[e^{gt} c_j(t)\right]^{1-\theta} - C}{1 - \theta} \tag{4}$$

In equation (4), $\theta \in (0, +\infty)\backslash\{1\}$ expresses the degree of concavity of the utility function (it is the inverse of the elasticity of intertemporal substitution). Parameter *C* is a positive parameter that can be calibrated to assure that the obtained values of utility are kept above zero, for certain combinations of parameter values. Rewriting (3) given (4),

$$U_j(0) = \int_0^T e^{-[\rho-n-(1-\theta)g]t} \frac{c_j(t)^{1-\theta}}{1-\theta} dt - \frac{C}{1-\theta} \int_0^T e^{-(\rho-n)t} dt \tag{5}$$

Condition $\rho - n - (1 - \theta)g > 0$ must be met in order to guarantee that intertemporal utility converges to a finite value.

As it is well known in growth literature, the maximization of consumption utility, as displayed in expression (5), subject to capital accumulation constraint (1), allows to derive an optimal consumption trajectory that corresponds to a constant growth rate for this variable. In particular, the optimal result is such that:

$$\dot{c}_j(t) = \frac{1}{\theta}[A - (\rho + \theta g + \delta)]c_j(t) \tag{6}$$

When the agent hypothetically solves the infinite horizon problem, growth rates of capital and consumption will coincide in the long-term balanced growth path (BGP), and will be equal to:

$$\gamma = \frac{1}{\theta}[A - (\rho + \theta g + \delta)] \tag{7}$$

If both endogenous variables grow, in the BGP, at the same rate, then the ratio between the two variables must be constant. Proceeding with the computation, the consumption-capital ratio in the BGP, defined by $\psi$, will be the following constant value,

$$\psi = \frac{\theta - 1}{\theta}(A - \delta) + \frac{1}{\theta}\rho - n \tag{8}$$

Expressions (7) and (8) allow for a considerable simplification on the presentation of the two key dynamic equations of the model, i.e., equations (1) and (6). These are presentable as:

$$\dot{k}_j(t) = (\psi + \gamma)k_j(t) - c_j(t) \tag{9}$$

$$\dot{c}_j(t) = \gamma c_j(t) \tag{10}$$

It is system (9)-(10) that will be approached in the next section, for an agent that has a short-term horizon, i.e., that solves the optimization problem one period of length 1 at a time.

## 3. The Solution of the One-Period Problem and the Savings Transversality Condition

Equations (9) and (10) are linear ordinary differential equations for which it is straightforward to derive explicit expressions for their two variables, capital and consumption, as functions of time. The result of solving the system is:

$$k_j(t) = e^{(\psi+\gamma)t}k_j(0) + \frac{e^{\gamma t}}{\psi}\left(1 - e^{\psi t}\right)c_j(0) \tag{11}$$

$$c_j(t) = e^{\gamma t}c_j(0) \tag{12}$$

The relation between the initial levels of capital and consumption is relevant to proceed with the analysis, and this relation can only be derived by looking at the system's solution at the selected terminal date. At this stage of the analysis, one considers that the agent can only solve a problem of length 1 at a time, and therefore a transversality condition for $t = 1$ must be established. It is assumed that, at such date, the decision-maker selects some savings rate $s_j(1) \in (0,1)$. As it will be clarified later, the choice of the savings rate by the individual agent will be made through a learning process, by observing the gains that individuals eventually choosing different savings rates achieve regarding utility levels.

By selecting savings rate $s_j(1)$, agent $j$ is faced with the transversality condition $c_j(1) = [1–s_j(1)]\, y_j(1)$, which is equivalent to $c_j(1) = [1–s_j(1)]Ak_j(1)$. Observing that expressions (11) and (12) correspond, for $t = 1$, to

$$k_j(1) = e^{\psi+\gamma}k_j(0) + \frac{e^{\gamma}}{\psi}\left(1 - e^{\psi}\right)c_j(0) \tag{13}$$

$$c_j(1) = e^{\gamma}c_j(0) \tag{14}$$

one can apply the transversality condition to these expressions and obtain:

$$e^{\gamma}c_j(0) = \left[1 - s_j(1)\right]A\left\{e^{\psi+\gamma}k_j(0) + \frac{e^{\gamma}}{\psi}\left(1 - e^{\psi}\right)c_j(0)\right\} \tag{15}$$

Rearranging,

$$c_j(0) = \frac{\psi e^{\psi}\left[1 - s_j(1)\right]}{e^{\psi} - s - (e^{\psi} - 1)s_j(1)}k_j(0) \tag{16}$$

In equation (16), $s \equiv 1 - \psi/A$, is the BGP savings rate, which is common across agents. The replacement of (16) in (11) and (12), makes it possible to present both variables with reference to the initial level of capital. Observe that consumption grows at a constant rate, while capital depends on time on a more sophisticated way.

$$k_j(t) = \frac{e^{\psi(1-t)} - s - \left[e^{\psi(1-t)} - 1\right]s_j(1)}{e^{\psi} - s - (e^{\psi} - 1)s_j(1)} e^{(\psi+\gamma)t} k_j(0) \tag{17}$$

$$c_j(t) = \frac{\psi e^{\psi}\left[1 - s_j(1)\right]}{e^{\psi} - s - (e^{\psi} - 1)s_j(1)} e^{\gamma t} k_j(0) \tag{18}$$

Concerning utility, the substitution of consumption in equation (18) into expression (5) yields (observe that the discount factor is rewritten given the definitions of $\psi$ and $\gamma$):

$$\begin{aligned} U_j(0) = \int_0^1 e^{-[\psi+(1-\theta)\gamma]t} \frac{\left[\frac{\psi e^{\psi}[1-s_j(1)]}{e^{\psi}-s-(e^{\psi}-1)s_j(1)} e^{\gamma t} k_j(0)\right]^{1-\theta}}{1-\theta} dt \\ - \frac{C}{1-\theta}\int_0^1 e^{-(\rho-n)t} dt \end{aligned} \tag{19}$$

which is equivalent to:

$$U_j(0) = \frac{\left[\frac{\psi e^{\psi}[1-s_j(1)]}{e^{\psi}-s-(e^{\psi}-1)s_j(1)} k_j(0)\right]^{1-\theta}}{1-\theta} \int_0^1 e^{-\psi t} dt - \frac{C}{1-\theta}\int_0^1 e^{-(\rho-n)t} dt \tag{19}$$

Solving the integrals,

$$U_j(0) = \frac{\left[\frac{\psi e^{\psi}[1-s_j(1)]}{e^{\psi}-s-(e^{\psi}-1)s_j(1)} k_j(0)\right]^{1-\theta}}{1-\theta} \frac{e^{\psi} - 1}{\psi e^{\psi}} - \frac{C}{1-\theta} \frac{e^{\rho-n} - 1}{(\rho - n)e^{\rho-n}} \tag{20}$$

Expression (20) represents the present-value of utility, at the initial date, for a given period of length 1 in which the agent solves the optimization problem. Observe that this utility value is explicitly presented as a function of the various parameters of the model, and without any dependence relatively to variable time. In possession of the values for every parameter of the model, it is possible to explicitly quantify utility.

Observe, as well, that the only potential sources of heterogeneity in the values of utility across individual agents are the end-of-period selected savings rate and the initial level of capital. In the next section, with the discussion about sequential planning problems, it will become evident that the values of $s_j$ (1) and $k_j$ (0) are intertwined: the more the agent saves at the end of one period, the larger will be the stock of capital available to the agent at the beginning of the following period.

To illustrate the relation between values $s_j$ (1) and $U_j$ (0), for a given planning horizon, consider a numerical example. Let $A = 0.2$; $\rho = 0.04$; $g =$

0.05; $\delta = 0.025$; $n = 0.02$; $\theta = 1.5$; $k_j(0)=1$; $C = 5$. With this array of parameter values, condition $\rho - n - (1-\theta)g = 0.045 > 0$ is satisfied, and the consumption-capital ratio and the BGP savings rate amount, respectively, to: $\psi = 0.065$ and $s = 0.675$. The BGP growth rate is: $\gamma = 0.04$.

Under the selected values, equation (20) is equivalent to the following nonlinear relation between the terminal savings rate and period utility:

$$U_j(0) = 9.9007 - 7.3523\sqrt{\frac{0.3922 - 0.0672s_j(1)}{1 - s_j(1)}} \tag{21}$$

Fig. 1 displays the dependence of $U_j(0)$ on $s_j(1)$. The relation is of opposite sign, meaning that lower savings at the end of the period will imply higher utility over the planning horizon. This is an intuitive result, because saving less at the terminal date frees resources to consume more over the time period under consideration. However, this is also a result with a twist: by saving more at the current period, the initial capital in the following period will be higher, contributing to additional production of wealth, consumption and, ultimately, utility; as one regards by observing expression (20), the level of $k_j(0)$ is decisive for determining $U_j(0)$. This point will be further discussed in the following section.

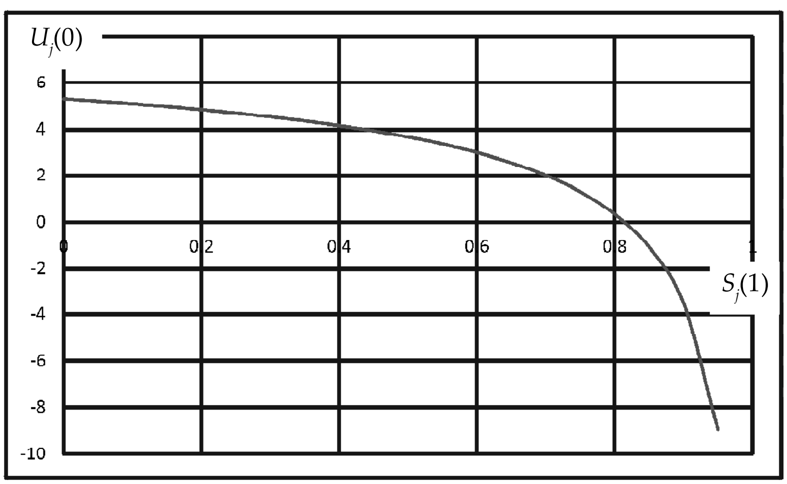


**Figure 1: Relation between end of period savings and period utility**

## 4. Multiple Planning Periods and the Choice of the End-of-Period Savings Rate

As in conventional growth models, the devised structure of analysis assumes that agents face an infinite horizon. Unlike conventional growth models, though, it is considered that decision-makers are uncapable of solving the optimal control problem at once and, thus, they split it in smaller problems of length 1. The endeavor of the agents that populate the economy will then be to repeatedly solve a series of endogenous growth problems with a similar structure.

However, besides solving the never-ending sequence of optimality problems, the individual agent has one more task to perform: she will have to choose, at the end of every time period, how much to save for the following period, in order to set the most convenient initial level of capital which, in turn, will lead to the best performance possible concerning the period utility level to achieve. The savings rate at the terminal dates will be the source of heterogeneity among agents, because agents potentially choose different end-of-period savings. This section explains how this savings choice is performed.

Let $k_j(t, \tau)$ be the stock of capital of agent $j$ at date $t$ for the problem $\tau = 1,2,\ldots$ solved by the agent; $\tau$ represents each of the sequential periods for which agent $j$ solves a maximization utility problem. Savings at the end of period $\tau$, for agent $j$, are defined as $s_j(1, \tau)\, y_j(1, \tau)$, which is equivalent to $As_j(1, \tau)\, k_j(1, \tau)$. The fundamental assumption about savings is that the savings rate at the terminal date will determine available capital at the beginning of the new period. Analytically,

$$k_j(0, \tau+1) = k_j(1, \tau) + s_j(1, \tau) y_j(1, \tau) = k_j(1, \tau)\left[1 + As_j(1, \tau)\right] \qquad (22)$$

Equation (22) indicates that the amount of capital available for the agent at the beginning of planning period $\tau + 1$ is the amount of capital with which the agent ends the previous period plus savings at the end of such period, given the selected terminal savings rate.

Therefore, one finds two countervailing forces in the decision problem of the agent, given her ultimate goal, which is to maximize utility for the period under consideration. On one hand, to achieve high utility levels, the agent must consume more and, thus, save less. However, higher savings will lead to increased levels of capital which allow to generate additional income and, hence, free resources for consumption, leading to higher utility. There is, in this case, a conflict between the short-run and the long-run. Because the agent solves a sequence of short-term problems, one might jump to the conclusion that the agent will privilege current consumption over savings. However, the agent will compare the performance of the various decision-makers and eventually conclude that those who saved more in the past are obtaining better utility outcomes today. Since each agent learns with the behavior of others, and by observing the performance of others, agents may eventually adjust their behavior in the direction of increased savings.

In certain conditions, namely those associated with the learning mechanism that is characterized below, the tendency for a corner solution (i.e., the tendency for a progressive increase in savings over time or, the opposite, a progressive concentration in the short-run goal of increasing consumption at the expenses of savings) might not prevail. Instead, endogenous fluctuations eventually emerge. Endogenous fluctuations will signify that, by observing the behavior of others, the individual agent will at some periods of time prefer to raise savings, while in other periods the best option to maximize utility is to lower savings. On the aggregate, if agents choose boundedly unstable trajectories for savings, the economy will exhibit business fluctuations that accompany the process of long-term growth.

Once a planning period is over, the agent has the possibility of observing the utility accomplished by all the other players. As a result, the agent will adapt her behavior, trying to place her savings rate closer to the one of the best performing agent. The eventual obstacle faced by the agent is the possible inability or incapability to fully observe or understand how much other agents are saving. The individual agent will make her best guess about the savings rate chosen by the other agents, but this estimation is not exempt of error. Savings rates are observed with noise, and the agent may fail in accurately predicting their true values.

Given the above reasoning, the decision-maker will adopt the following leaning rule for establishing her own terminal savings rate in the period that follows:

$$s_j(1,\tau+1) = \zeta s_j(1,\tau) + (1-\zeta)\left[s_m(1,\tau) + \varepsilon_j\right] \tag{23}$$

In equation (23), $\zeta \in (0,1)$ represents the measure of the adjustment; $s_m$ $(1, \tau)$ is the end-of-period savings rate of the agent with the highest utility; and $\varepsilon_j \sim iid(0, \sigma)$ is a noise variable with zero mean and standard-deviation equal to ?.

The equation indicates that the end-of-period savings rate at period $\tau + 1$ is approximated to the end-of-period savings rate (observed with noise) of the best performing agent, at $\tau$. Parameter $\zeta$ measures how fast is the agent willing to abandon her original choice in favor of the choice of others; it translates, in a sense, how fast the agent is willing to learn (the lower the value of the parameter, the higher is the willingness to learn). Note that if the maximum utility is achieved by the agent taking the decision, then no uncertainty is involved, and the agent keeps her previous decision: $s_j(1, \tau + 1) = s_j(1, \tau)$.

At this stage, all the ingredients of the model are duly formalized and might be put together to discuss growth in the context of a learning-based multi-agent system. This is done in the next section, via simulation of the model. Recapping, the analysis so far has achieved the following:

i) The short-run planning problem of an individual decision-maker has been presented and solved;

ii) The intertemporal utility expression for the individual decision-maker has been derived;

iii) The source of heterogeneity in the model has been highlighted. This is attached to the savings rate that the agents select at the terminal date;

iv) It has been emphasized that the end-of-period savings rate is chosen with the objective of approaching maximum utility in the period;

v) It has been noted that the savings rate at the end date in one period will determine the initial stock of capital in the following period;

vi) It has been underlined that the savings rate of the individual is chosen through a learning process: she compares her own savings rate with the one of the agent with the highest utility (this savings rate is observed with noise), and approaches the second at a given rate.

The next section will evaluate the dynamics of the formalized growth model, by taking a prototypical numerical example and, from the example, by simulating some of the relevant trajectories for the individual agents and for the economy as a whole. The outcome will be a result of systematic fluctuations, implying that no convergence to a BGP will ever take place.

## 5. Simulation of the Multi-Period Growth Model

In the simulation of the model to undertake in this section, the focus will be placed on the graphical representation of time trajectories for the following variables:

i) Individual utility per planning period;

ii) End-of-period savings rates selected by individual agents, after learning;

iii) Aggregate income growth, at the end of planning dates;

iv) Growth of aggregate consumption and aggregate investment, at the end of planning dates.

The parameter values to use in the simulation are those already proposed in section 3. To these, it is added the value $\zeta = 0.75$. In order to keep the graphical analysis tractable, only three agents (or three identical groups of agents) are assumed, although many more could be considered under the framework explored in the previous sections.

Individual agents start by choosing relatively close but distinct terminal savings rates for the first planning period: $s_1$ (1,1) = 0.25, $s_2$ (1,1) = 0.3, $s_3$ (1,1) = 1/3. Given parameter values, the expressions for the several variables involved (e.g., capital, income, consumption, and utility), and the adopted rules of motion (namely for the evolution of the savings rate), it is possible to present time trajectories for the relevant variables of the model. In the representation of these trajectories, one hundred points in time are considered (which, in this case, have correspondence in one hundred planning periods for each individual).

Let us begin by addressing utility. To turn results comparable, the utility obtained by each agent at each period, $U_j$ (0, τ), is divided by the average value of the utilities of the three agents. Hence, in this setting, utility measures with a value above 1 are utility levels above the average and those below 1 are utility levels below the average. The first graphic (Fig. 2) is drawn for the case of certainty, i.e., when there is no noise associated with the perception of the savings rate of the player with the highest utility.

In the case of complete knowledge and no uncertainty, the highest utility begins, in the first time steps, by being the one associated with the agent that saves less. However, after a few periods, given the accumulation of capital that savings provokes, this result is modified in favor of the agent that saves more. Note that the utility levels converge over time; this occurs because, under the learning mechanism that is considered, end-of-period savings rates converge, as well, to a unique value, what is represented in Fig.3. As this figure reveals, the convergence process is relatively fast. Although in a first phase (until period 4) the convergence goes in the direction of the lower savings, the movement is rapidly inverted and stabilizes at an intermediate value of $s_j$ (1).

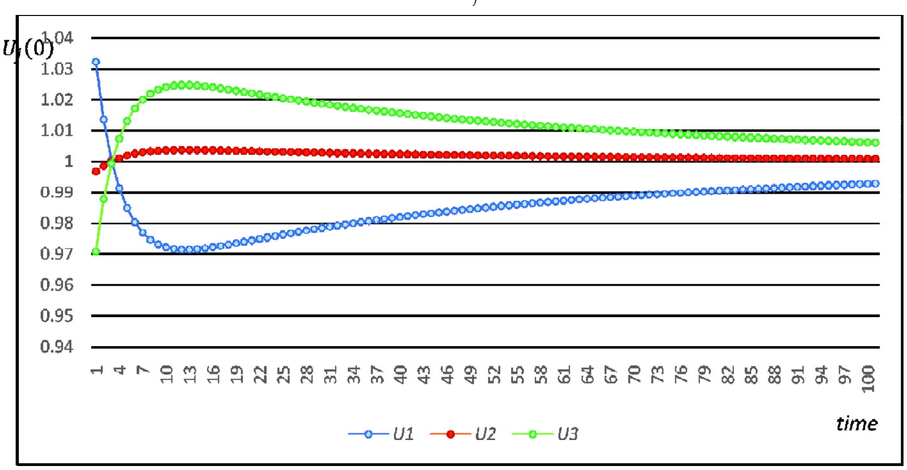


**Figure 2: Evolution of period utility under perfect knowledge**

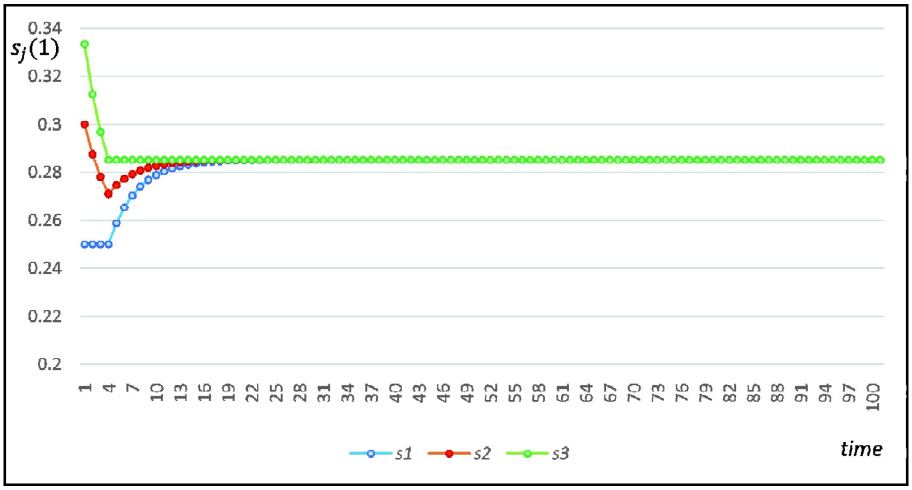


**Figure 3: Convergence in end-of-period savings rates under perfect knowledge**

The results displayed in Figs. 2 and 3 radically change when introducing noise in the perception of the savings rates selected by other agents. Let ?=0.05 be the standard-deviation associated with the noise term. In this case, results will change every time the model is run, but there is a pattern: there is not an unequivocal prevalent utility result, i.e., different agents have the best utility outcome at different time periods each time the model is run. Associated with this outcome, there is also a bounded instability process attached to the way in which each the savings rate evolves over time. Fig. 4 and Fig. 5 display three possible examples of the trajectories followed by the utility of each agent and by each agent's end-of-period savings rate.

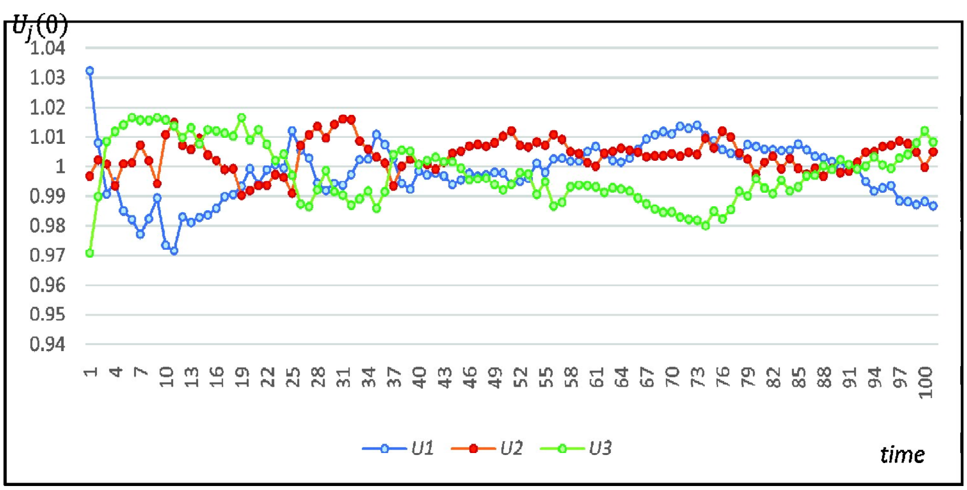

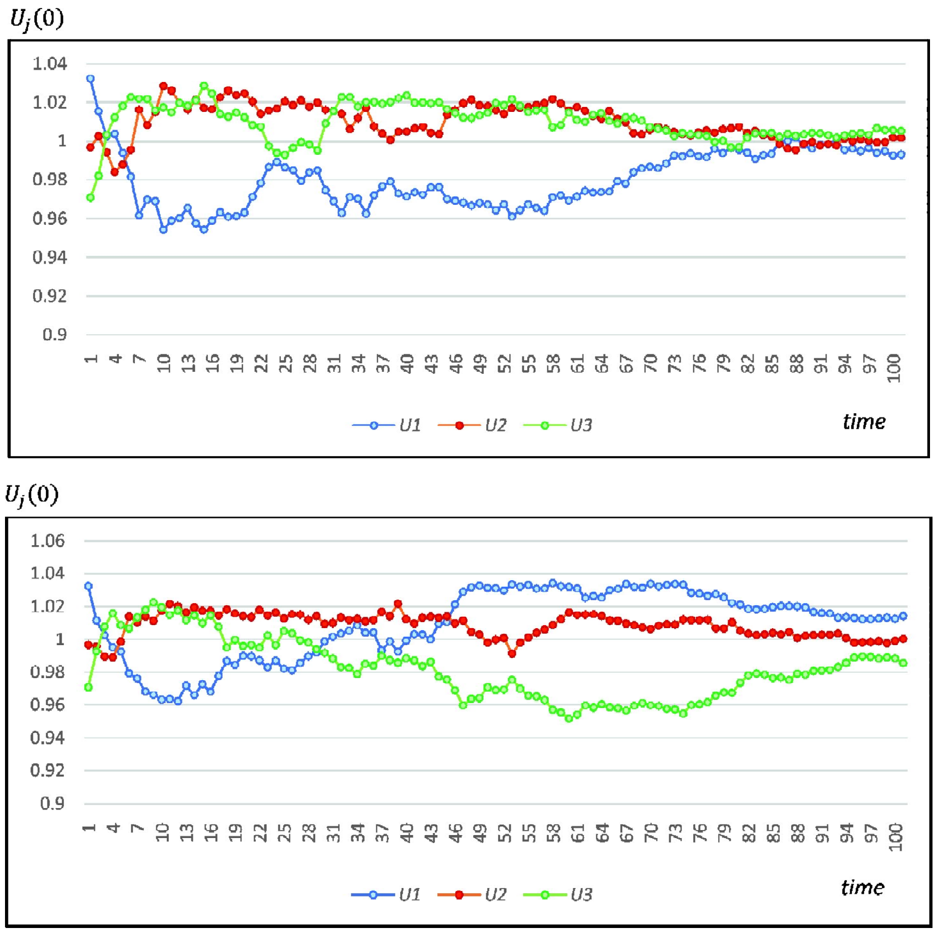


**Figure 4: Evolution of period utility with noise**

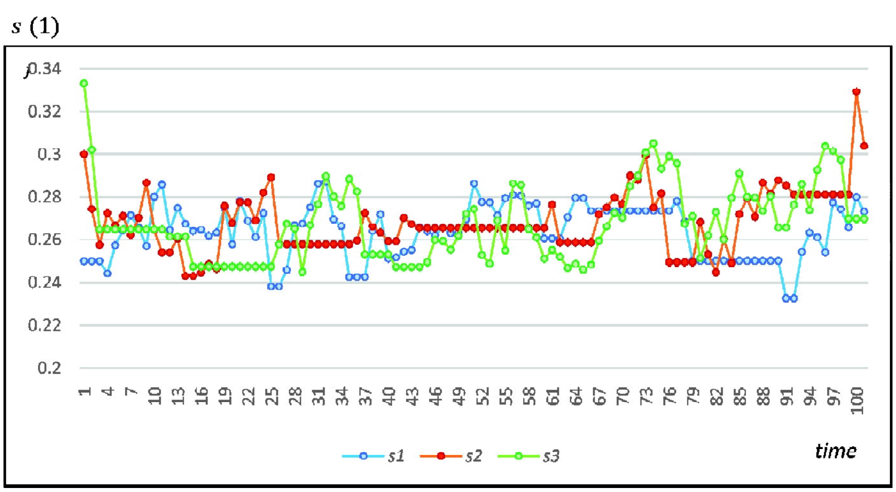

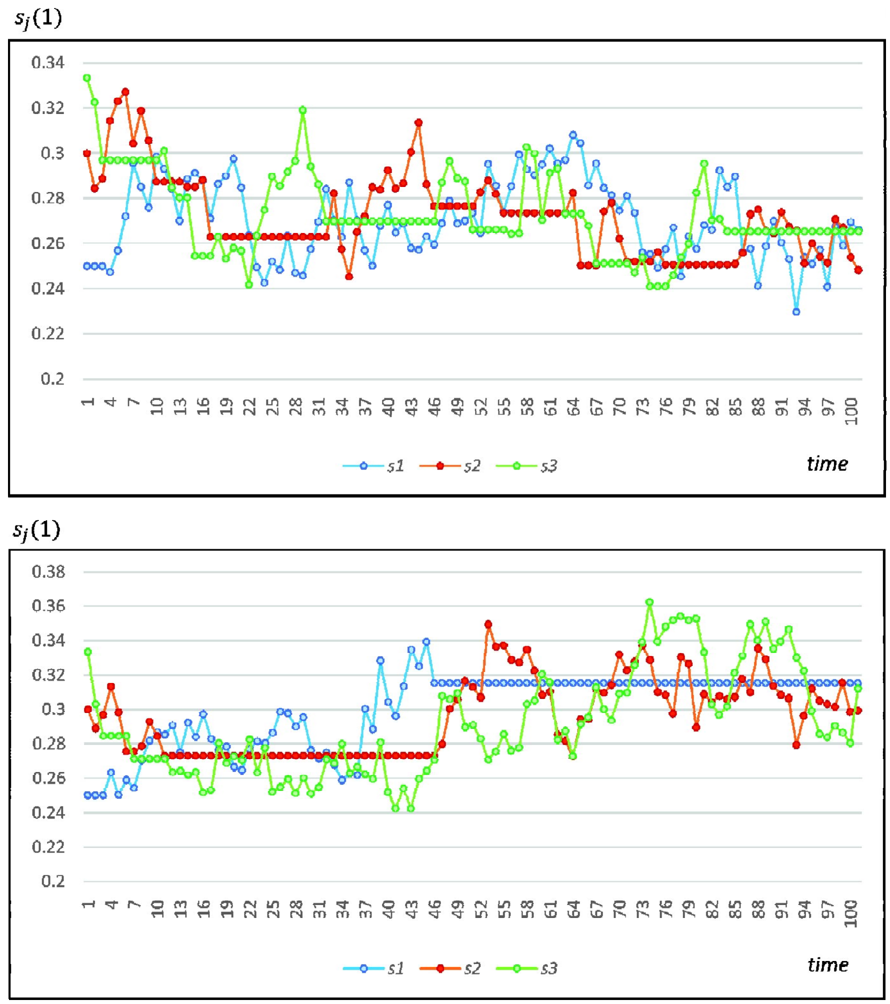


**Figure 5: Time trajectories of savings rates with noise**

Observe, from the figures, that for the agent with the highest utility over some length of time, the corresponding trajectory of the terminal savings rate is linear. This occurs because the agent with the highest utility does not have an incentive to change the savings rate and therefore faces no uncertainty: she just keeps selecting the same rate as before. When the utility level of the individual agent is hypothetically overtaken by the utility level of another agent, the first individual abandons her savings rate and tries to guess what the savings rate of the other agent is (what is done with error, thus justifying the oscillating irregular path followed by the savings trajectories, as displayed in Fig. 5).

Because individual results, contingent on the interaction among agents in a scenario of imperfect information, lead to bounded instability and irregular fluctuations at the individual (micro) level, one should expect these fluctuations to be passed on to the aggregate (macro) level. The figures to be presented below confirm this supposition. Particularly relevant is the result that the growth model involving the subtleties of learning and interaction is capable of explaining not only long-term growth but also of addressing potential sources of business cycles in the economy. In this case, aggregate fluctuations have, as candidate underlying forces, a mix of elements including decentralized interaction, learning and imperfect information.

Let us return to the case of no noise, for a first approach to the outcomes of the aggregate economy. Define aggregate income, aggregate consumption, and aggregate investment, at the end of each planning date, in the following terms:

$$Y(1,\tau)=\sum_{j=1}^{3} y_j(1,\tau)\,;\; C(1,\tau)=\sum_{j=1}^{3} c_j(1,\tau)\,;\; I(1,\tau)=\sum_{j=1}^{3}\left[y_j(1,\tau)-c_j(1,\tau)\right] \quad (24)$$

The evolution of the growth rates of these three variables, with perfect knowledge and the absence of noise, is represented in Fig. 6. It is evident that, after a brief transient phase, the growth rates converge to a same positive value. Because the equilibrium growth rate is positive and derived within the model, this has the nature of an endogenous growth model. Sustained growth is guaranteed by the fact that in the beginning of each planning date, capital accumulated from the previous period plus savings is reinvested in production.

A possible comparison one might undertake at this point is the comparison between the infinite horizon problem of the representative agent in the orthodox formulation of the model and the outcome in Fig.6. As mentioned in section 3, for the adopted parameter values, the infinite horizon problem yields a BGP growth rate equal to $\gamma = 0.04$. In the multi-agent system, the observation of the figure points to a lower growth rate for the aggregate economy, $\gamma = 0.023$. This is an expected result: it is certainly more efficient to solve the infinite-horizon problem at once, instead of solving it period per period and choosing a savings rate every time the optimization problem is addressed. The problem is that solving the infinite horizon problem with no regard for the choices of others involves assuming a notion of rationality that goes beyond the capabilities of the average agent: the economy grows less than it could eventually grow because decision-makers are short-sighted, suffer the influence of others, and have to learn to eventually arrive to a BGP outcome (i.e., agents are not hyper-rational).

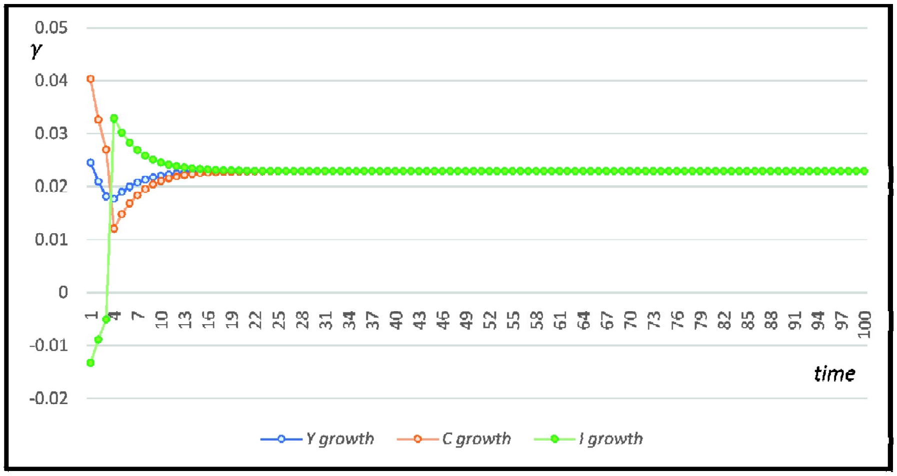


**Figure 6: Growth paths of income, consumption, and investment under perfect knowledge**

Once the assumption of perfect knowledge is relaxed, fluctuations emerge in the paths of the three variables - income, consumption, and investment - as Fig. 7 and Fig. 8 allow to verify. In each case, three panels are drawn, representing three possible realizations of the model. Income is isolated from the other two variables, to highlight the growth of the economy, which is represented by the evolution of income. Observe that, in each of the examples, the growth rate of income fluctuates at values close to the BGP $\gamma = 0.023$ but never rests in this value. Fig. 8 reveals the result that investment is, in every example, more volatile than consumption, an observation that has empirical support.

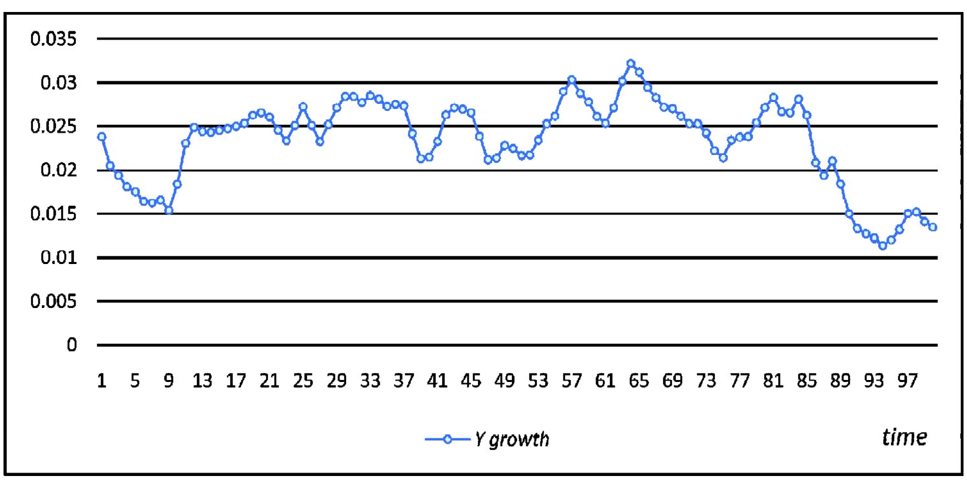

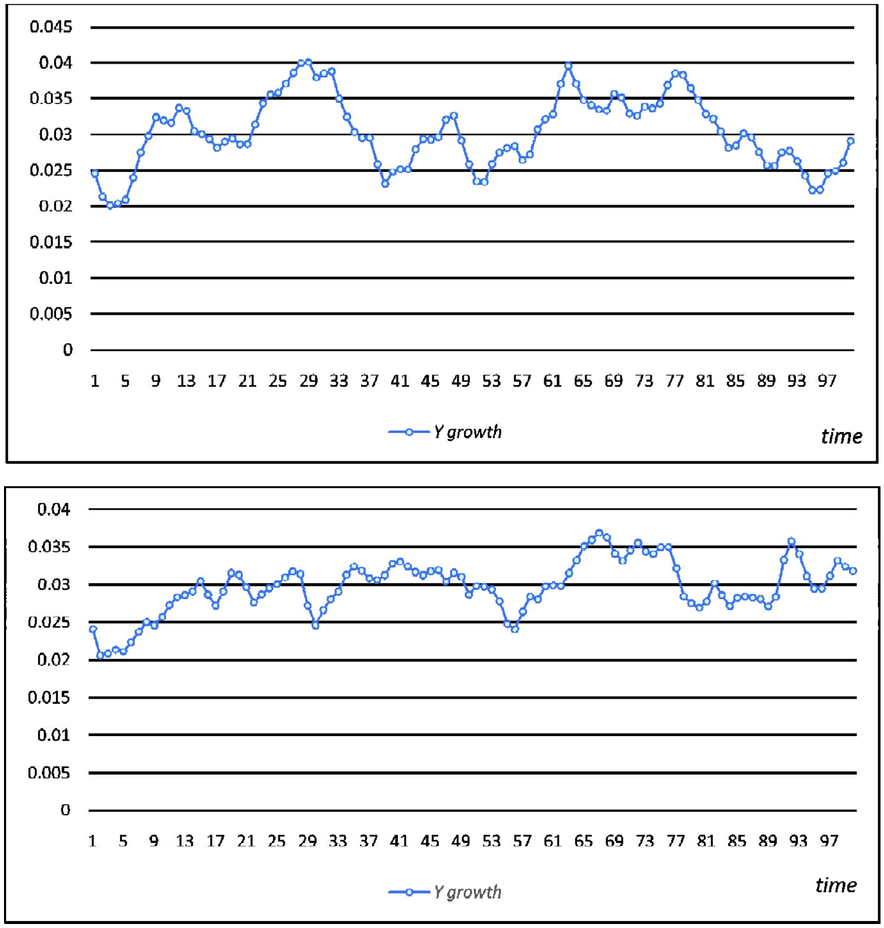


**Figure 7: Growth path of income, in the model with noise**

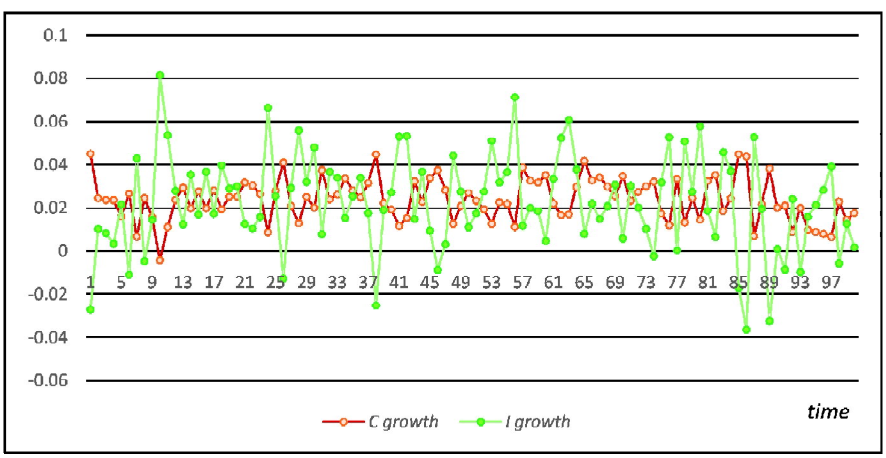

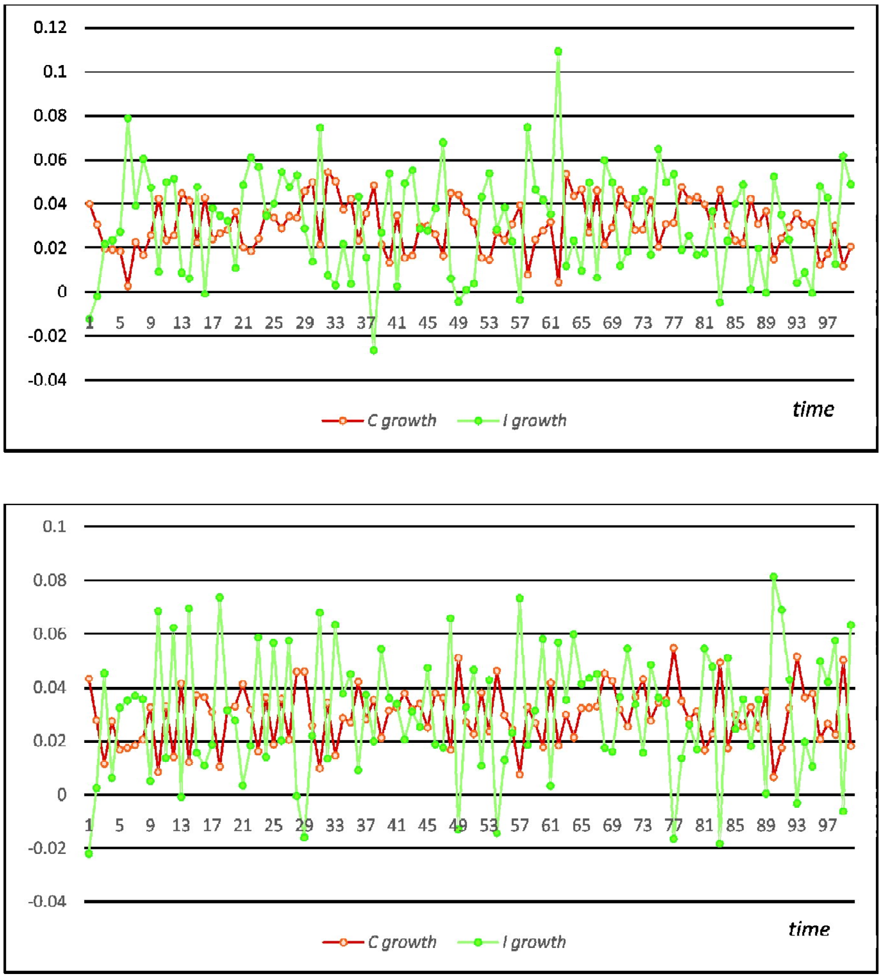


**Figure 8: Growth paths of consumption and investment, in the model with noise**

To complement the above graphical analysis, Fig. 9 displays the relation between the average end-of-period savings rate (i.e., the average between the three individual savings rates) and the growth rate of income. Again, three panels are represented for three different realizations of the model in the imperfect knowledge case. As expected, the periods in which savings are higher are also those in which the economy grows more (because these are also the periods in which the economy accumulates capital at a faster pace). In each graphic, a trend line is added to emphasize the predominance of the characterized result and the found regularity.

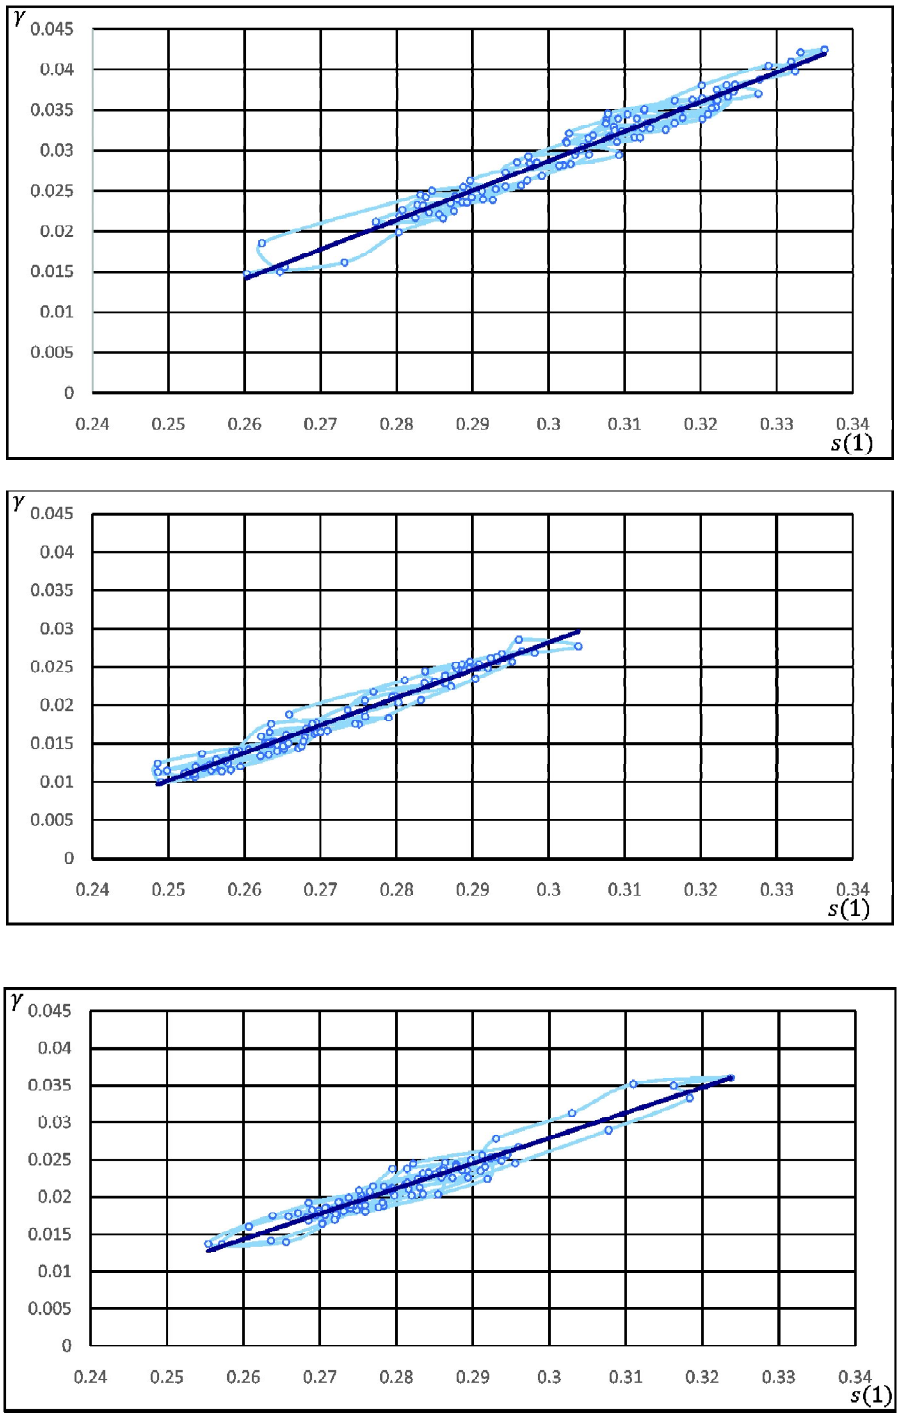


**Figure 9: Correlation between the savings rate and the growth rate of income**

## 6. Possible Extensions

The growth model characterized along the previous sections might be modified and extended in various directions. To exemplify the adaptability of the theoretical framework, two possible changes to the benchmark structure are proposed and briefly discussed below.

The first adaptation concerns the nature of the capital stock. Expression (22) indicates that capital is fully private, i.e., the capital accumulated in one period and the agent's savings in the same period, fully revert to the same individual at the beginning of the subsequent period. A redistributive policy or the public good nature of the capital good may imply that such assumption eventually does not hold.

If one assumes that part of the capital is private and returns to the agent in the next period, while the remaining share is equitably distributed across all agents, then the dynamics observed in the previous section will certainly suffer a change. Let $\omega \in (0,1)$ be the share of private capital. In this scenario, the amount of capital available to agent $j$ at the beginning of period $\tau + 1$ will no longer be the one presented in equation (22); it will be the following:

$$k_j(0, \tau + 1) = \omega k_j(1, \tau)\left[1 + As_j(1, \tau)\right] + (1 - \omega)\frac{\sum_{i=1}^{J} k_i(1, \tau)[1 + As_i(1, \tau)]}{J} \tag{25}$$

In equation (25), the term $\sum_{i=1}^{J} k_i(1, \tau)[1 + As_i(1, \tau)]$ corresponds to the stock of capital accumulated in the economy at the end of period $\tau$. A fraction $1 - \omega$ of this stock is equitably distributed across agents and it will add to the share of private capital to form the available stock of this input at the beginning of the new planning period, i.e., $\tau + 1$.

If a share of the capital is of public use, and therefore can be employed by every agent in the same amount, regardless of who generated it, this will introduce a tragedy of the commons type of problem. Agents will expect others to accumulate capital that can be used by all. Meanwhile, the same agents reinforce consumption and lower savings to increase period utility. For a lower than 1 value of $\omega$, agents have progressively less incentive to save, because their accumulated capital will be distributed by the universe of the agents in the economy. In this case, shares of savings eventually fall to zero and the only capital available to initiate the next period is the one that is accumulated from the previous period.

This result is depicted in Fig. 10, where the trajectories of the savings rates at the terminal date are drawn for the same parameter values used in previous examples. In this case, it is assumed $\omega = 2/3$. Three panels are

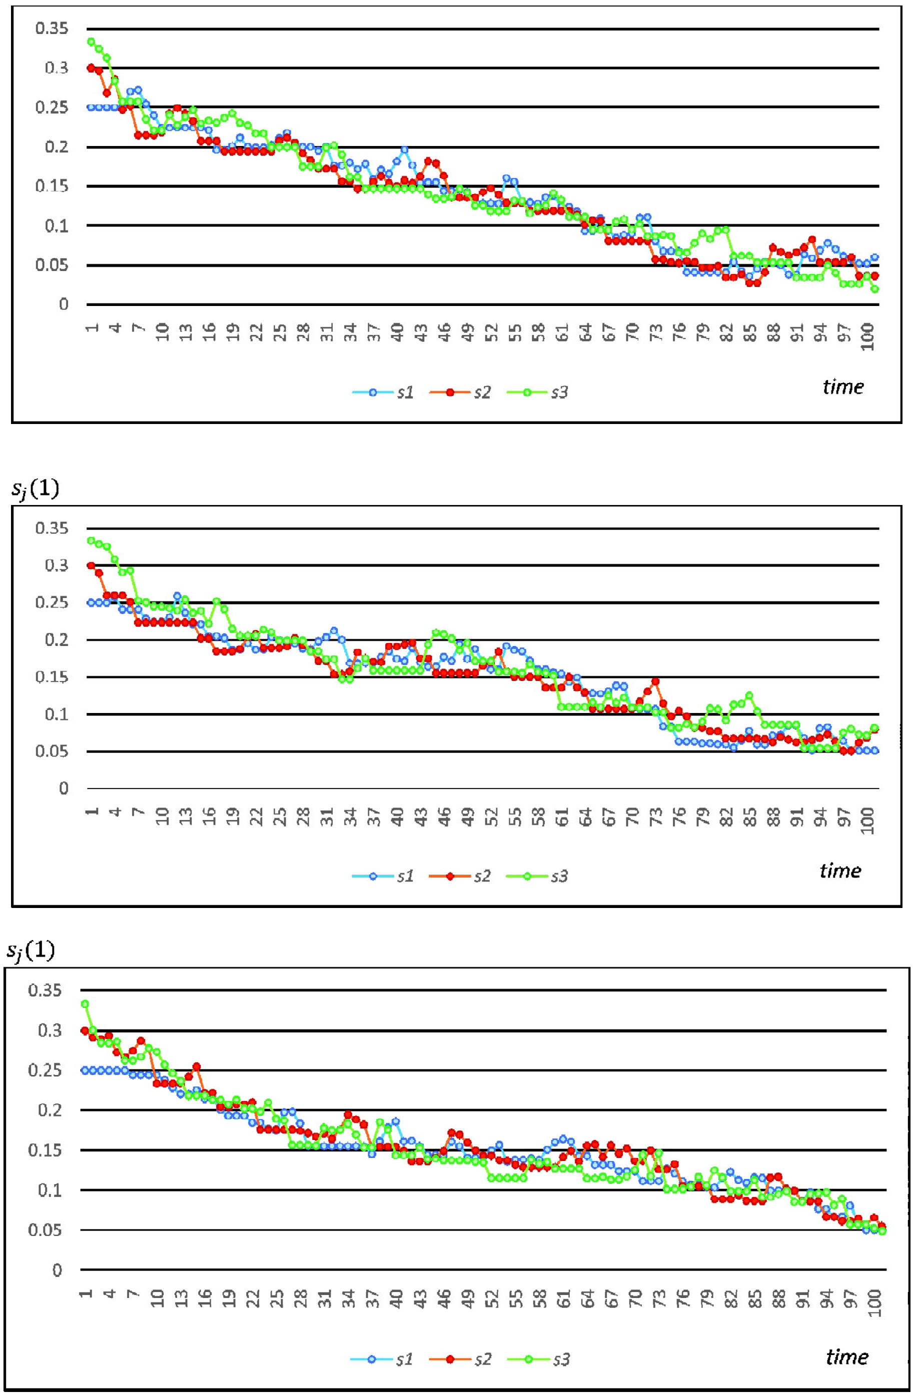


**Figure 10: Time trajectories of savings rates with capital as a partial public good**

drawn, representing three possible realizations of the model given the presence of the noise variable. One observes that if part of the available capital has a public good nature this will progressively lower end-of-period savings, leading to a convergence to zero in the long term. Therefore, in this model agents have no incentive to cooperate and share resources; when such resources are shared, agents save less, accumulate lower amounts of capital and, therefore, will consume less in the long-run, with the associated negative effect over utility.

Another possible extension consists in conceiving that agents are capable of assessing future outcomes of current decisions and, therefore, to evaluate utility over long horizons. In this case, agents will compare utility across periods under the assumption that all players will maintain their current end-of-period savings rate. The utility levels to compare, for each agent $j$, will be, in this case, $\sum_{\tau=1}^{T} \beta^{\tau} U_j(0,\tau)$, with $\beta \in (0,1)$ an intertemporal discount factor. The savings rate will then be chosen through a rule similar to (23), where the only change consists in the measure of utility that is subject to evaluation, and that now is not circumscribed to the current period.

If the discount factor is high enough (i.e., if the future is not too strongly discounted), this new assumption constitutes a factor of inertia, because the rank of utilities across agents is less likely to suffer changes. Consequently, the best ranked player will tend to be perpetuated over time, and the other agents will adjust their choices in order to meet, in the long-term, the savings rate of the agent with the best performance in terms of utility.

## 7. Conclusion

The models devised by economists to explain economic growth have the virtue of setting the foundations of such explanation in the behavior of agents that make choices and plan for the future. These decision-makers are simultaneously households, who choose the optimal paths for consumption and savings, and firms, which accumulate capital given some production technology. If the production technology is such that constant returns in the accumulation of capital prevail, then a process of sustained endogenous growth emerges.

The main issue with the mainstream model of growth is that the behavior of the economic agents is excessively stylized. Relaxing some of the assumptions that shape such behavior may allow to acquire additional insights about the process of growth. First, in the mainstream model, agents decide and act in a fully rational way. This signifies that agents are capable of formulating and solving intertemporal problems in long horizons (even infinite horizons), that they are fully autonomous in the sense that they do not need to learn from the behavior of others, and that they face no obstacle in collecting and processing information about the future, thus acting under perfect foresight. The corollary of all the prior observations is that all agents act alike and, therefore, the intertemporal consumption choice they face can be reduced to the intertemporal decision problem of a single representative agent.

The orthodox model of growth is at odds with a complexity view of the economy, i.e., with a view in which the emergent outcome is the result of the interaction among a large number of individuals that are unsophisticated, and thus follow simple rules, and that have to manage themselves in a complex world that they cannot fully understand or know. It is with this idea in mind, i.e., that the economy should be perceived and interpreted as a complex system, that the benchmark economic growth model is modified in this paper. A few changes are introduced, making the growth model approach a system that can be classified as a learning-based multi-agent system.

As the designation indicates, two of the most significant changes to the model are the introduction of heterogeneity (thus transforming the economy in a multi-agent system) and the inclusion of a process through which agents modify their behavior by observing the behavior of others and, hence, potentially approaching the behavior of those who perform the best (this is a learning mechanism). Besides these, other changes are required, specifically the idea that the capacity of individuals to solve intertemporal plans is limited. Therefore, instead of solving a utility maximization problem for the whole horizon of life, the agent splits this problem in a sequence of many problems of equal length. By solving these smaller scale problems, the individual is capable of adjusting her behavior in the moment between the ending date of a problem and the beginning date of the next one. This adjustment takes the form of a choice of how much to save from one period to the next. Savings have the advantage of increasing the capital stock, although they have the disadvantage of provoking a substitution effect that lowers consumption at the expenses of savings.

The agent will not be fully informed and, therefore, she will not know how much to save to attain her utility maximization goal. Hence, she will compare her utility with the utility of those who adopt different savings rates, and adjust the respective choice whenever the choices of others perform better. If there were no informational flaws, and the agents were capable of exactly observing how much the others save, the dynamics underlying the model would converge to a BGP, where the main economic aggregates would all grow at the same rate, and the savings rate selected by each individual would rapidly become the same. In this case, the differences relatively to the benchmark infinite-horizon representative-agent model are not much: the main result is, in each scenario, the formation of a BGP where the behavior of every agent is the same and, therefore, there is a coincidence between individual outcomes and aggregate outcomes.

The previous result is radically modified once one considers that the learning process is subject to some sort of information or knowledge imperfection. If the agent is not capable of knowing exactly how much the agent that performs best in terms of utility saves, the incurred potential errors might lead to a systematic change in the selected savings rate at the end of each time period. As a result, capital, consumption and income will not converge to the BGP, i.e., they will not grow, after the transient phase has faded out, at a constant rate. The time trajectories of the main aggregate variables will display endogenous fluctuations under the form of a bounded instability evolution process, in which the corresponding growth rates oscillate around a constant value but never converge towards it.

The main message is that adapting conventional growth models in order to interpret them as complex systems, more specifically learning-based multi-agent systems, is a fundamental step to acquire additional insights on the growth process. Although the obtained results might not appear as straightforward as in the infinite-horizon problem, much of the introduced assumptions add realism to the model: the cognitive limitations in planning ahead too far in the future, the information processing limitations associated with knowing how much others save, and the decision-making limitations, that make agents learn with others instead of taking decisions on their own, are all elements that place us closer to reality.

## References

Acemoglu, D. (2009). Introduction to Modern Economic Growth. Princeton. NJ: Princeton University Press.

Alogoskoufis, G. (2019). Dynamic Macroeconomics, Cambridge, MA: MIT Press.

Anandakumar, H.; and R. Arulmurugan (2019). “Machine Learning Based Multi Agent Systems in Complex Networks.” Third International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC), pp. 300-304.

Arthur, W.B. (2015). Complexity and the Economy, New York: Oxford University Press.

Athey, S. (2018). "The Impact of Machine Learning on Economics." The Economics of Artificial Intelligence: an Agenda, pp. 507-547.

Babenko, V.; A. Panchyshyn; L. Zomchak; M. Nehrey; Z. Artym-Drohomyretska; and T. Lahotskyi (2021). “Classical Machine Learning Methods in Economics Research: Macro and Micro Level Example.” WSEAS Transactions on Business and Economics, vol. 18, pp. 209-217.

Bargigli, L.; and G. Tedeschi (2014). "Interaction in Agent-based Economics: a Survey on the Network Approach." Physica A: Statistical Mechanics and its Applications, vol.399, pp.1-15.

Barro, R.J. and X. Sala-i-Martin (2004). Economic Growth, 2nd ed., New York: McGraw-Hill.

Bezemer, D.J. (2012). "The Economy as a Complex System: the Balance Sheet Dimension." *Advances in Complex Systems*, vol. 15(5), pp. 1-22.

Birke, D. (2009). "The Economics of Networks: a Survey of the Empirical Literature." *Journal of Economic Surveys,* vol. 23(4), pp. 762-793.

Borah, K.J.; and R. Talukdar (2019). "Revisited: Machine Intelligence in Heterogeneous Multi-Agent Systems." International Conference on Aerospace System Science and Engineering, Singapore: Springer, pp. 219-229.

Bramoullé, Y.; R. Kranton; and M. D'Amours (2014). "Strategic Interaction and Networks." *American Economic Review,* vol. 104(3), pp. 898-930.

Cass, D. (1965). "Optimum Growth in an Aggregative Model of Capital Accumulation." *Review of Economic Studies,* vol. 32(3), pp. 233-240.

Chen, S.H. (2012). "Varieties of Agents in Agent-based Computational Economics: a Historical and an Interdisciplinary Perspective." *Journal of Economic Dynamics and Control,* vol. 36(1), pp. 1-25.

Delli Gatti, D.; E. Gaffeo; and M. Gallegati (2010). "Complex Agent-based Macroeconomics: a Research Agenda for a New Paradigm." *Journal of Economic Interaction and Coordination,* vol. 5(2), pp. 111-135.

Ellowitz, J. (2008). "Multi-agent Reinforcement Learning and Genetic Policy Sharing." arXiv preprint arXiv:0812.1599.

Fagiolo, G.; and A. Roventini (2017). "Macroeconomic Policy in DSGE and Agent-based Models Redux: New Developments and Challenges Ahead." *Journal of Artificial Societies & Social Simulation,* vol. 20(1).

Holt, R.P.F.; J.B. Rosser, J.B.; and D. Colander (2011). "The Complexity Era in Economics." *Review of Political Economy,* vol. 23(3), pp. 357-369.

Hommes, C.H. (2006). "Heterogeneous Agent Models in Economics and Finance." in: L. Tesfatsion & K.L. Judd (eds.), Handbook of Computational Economics, edition 1, vol. 2, chapter 23, pp. 1109-1186.

Hommes, C.H. (2021). "Behavioral and Experimental Macroeconomics and Policy Analysis: a Complex Systems Approach." *Journal of Economic Literature,* vol. 59(1), pp. 149-219.

Hou, J.; F. Wang; L. Wang, L.; and Z. Chen (2021). "Reinforcement Learning Based Multi-Agent Resilient Control: From Deep Neural Networks to an Adaptive Law." Proceedings of the AAAI Conference on Artificial Intelligence, vol.35(9), pp. 7737-7745.

Kaelbling, L.P.; M.L. Littman; and A.W. Moore (1996). "Reinforcement Learning: a Survey." *Journal of Artificial Intelligence Research,* vol.4, pp. 237-285.

Khalil, K.M.; M. Abdel-Aziz; T.T. Nazmy; and A.B.M. Salem (2015). "MLIMAS: a Framework for Machine Learning in Interactive Multi-agent Systems." Procedia Computer Science, vol.65, pp. 827-835.

Kirman, A. (2006). "Heterogeneity in Economics." *Journal of Economic Interaction and Coordination,* vol.1(1), pp. 89-117.

Koopmans, T.C. (1965). "On the Concept of Optimal Economic Growth." in The Econometric Approach to Development Planning. Amsterdam: North Holland.

Lengnick, M. (2013). "Agent-based Macroeconomics: a Baseline Model." *Journal of Economic Behavior & Organization,* vol. 86(C), pp. 102-120.

Lucas, R.E. (1988). "On the Mechanics of Economic Development." *Journal of Monetary Economics,* vol. 22(1), pp. 3-42.

Mosavi, A.; Y. Faghan; P. Ghamisi; P. Duan; S.F. Ardabili; E. Salwana; and S.S. Band (2020). "Comprehensive Review of Deep Reinforcement Learning Methods and Applications in Economics." *Mathematics,* vol. 8(10), article 1640, pp. 1-42.

Rebelo, S. (1991). "Long-Run Policy Analysis and Long-Run Growth." *Journal of Political Economy,* vol. 99(3), pp. 500-521.

Romer, D. (2021). Advanced Macroeconomics, 5th ed., New York: McGraw-Hill Education.

Romer, P.M. (1990). "Endogenous Technological Change." *Journal of Political Economy,* vol. 98(5), pp. S71-S102.

Sims, C. A. (1980). "Macroeconomics and Reality." *Econometrica,* vol. 48(1), pp. 1-48.

Solow, R.M. (1956). "A Contribution to the Theory of Economic Growth." *Quarterly Journal of Economics,* vol. 70(1), pp. 65-94.

Stiglitz, J.E. and M. Gallegati (2011). "Heterogeneous Interacting Agent Models for Understanding Monetary Economies." *Eastern Economic Journal,* vol. 37(1), pp. 6-12.

Sutton, R.S.; and A.G. Barto (2018). Reinforcement Learning: an Introduction. Cambridge, MA: MIT press.

Szepesvári, C. (2010). "Algorithms for Reinforcement Learning." Synthesis Lectures on Artificial Intelligence and Machine Learning, vol.4(1), pp. 1-103.

Wiering, M.A.; and M. Van Otterlo (2012). "Reinforcement Learning." Adaptation, Learning, and Optimization, vol.12(3), 729.

Wolpert, D.H.; Wheeler, K.R.; and K. Tumer (1999). "General Principles of Learning-based Multi-agent Systems." Proceedings of the Third Annual Conference on Autonomous Agents, pp. 77-83.