From Dynamics to Novelty: An Agent-Based Model of the Economic System

Gustavo Recio
Universidad Rey Juan Carlos
Department of Computer Science
gustavo.recio@uc3m.es

Wolfgang Banzhaf*
Michigan State University
Department of Computer Science and Engineering
banzhafw@msu.edu

Roger White
Memorial University of Newfoundland and Labrador
Department of Geography
roger@mun.ca

Abstract The modern economy is both a complex self-organizing system and an innovative, evolving one. Contemporary theory, however, treats it essentially as a static equilibrium system. Here we propose a formal framework to capture its complex, evolving nature. We develop an agent-based model of an economic system in which firms interact with each other and with consumers through market transactions. Production functions are represented by a pair of von Neumann technology matrices, and firms implement production plans taking into account current price levels for their inputs and output. Prices are determined by the relation between aggregate demand and supply. In the absence of exogenous perturbations the system fluctuates around its equilibrium state. New firms are introduced when profits are above normal, and are ultimately eliminated when losses persist. The varying number of firms represents a recurrent perturbation. The system thus exhibits dynamics at two levels: the dynamics of prices and output, and the dynamics of system size. The model aims to be realistic in its fundamental structure, but is kept simple in order to be computationally efficient. The ultimate aim is to use it as a platform for modeling the structural evolution of an economic system. Currently the model includes one form of structural evolution, the ability to generate new technologies and new products.

Keywords
Artificial chemistries, economic simulation, evolutionary economy, agent-based systems, structural evolution, complexity

1 Introduction

The global economy is both highly integrated and widely decentralized. It is complex, dynamic, and innovative. It is a system that is constantly evolving to become a somewhat different system. Here we introduce an agent-based model that captures the fundamental characteristics of economic systems: their complexity, their dynamics, and their continual structural transformation. Our ultimate objective is to develop a computational theoretical framework that can be used to investigate general questions related to equilibrium, dynamics, and structural change resulting from innovation. To this end we show that the model can replicate basic features of the real economy such as attaining equilibrium following a perturbation, as well as long-term growth of productivity, size, and complexity. The model is designed so that it can be easily modified or extended in order to investigate a wide variety of questions.

Our model can be seen as an alternative to neoclassical microeconomic theory in that it represents the behavior of economic agents (consumers and firms) in order to show how a well-structured

* Corresponding author.

economic system results from that behavior. Neoclassical theory is essentially static and characterizes a system by its equilibrium state. Change is handled not in terms of dynamics but as a problem in comparative statics. Perhaps in part because of the rigor and foundational importance of the neoclassical tradition, as Krugman (1995) pointed out, dynamical approaches have never become central to economic theory, although there have been numerous developments over the years.

1.1 Related Approaches

One approach consists of the dynamic stochastic equilibrium models (e.g., see Blatt, 1983; Colander, 2006; Sbordone et al., 2010). These models have the advantage of representing the dynamics explicitly, and purport to show macroscale phenomena of an economy as emergent from microscale behavior as represented in microeconomic theory. However, microeconomic theory brings with it serious limitations to the realism of these models, as de Grauwe (2010) has pointed out. One of these limitations is that the microscale behavior in these models is simply a behavioral rule (utility maximization under rational expectations) that applies to all microscale agents, so in fact there is no real representation of the microscale. Another is that the focus is still on the equilibrium state rather than the dynamics as such. Recently Dosi and Roventini (2019) called for moving beyond these restrictions, but that has not been done yet. These limitations mean that dynamic stochastic equilibrium models cannot in general be used to investigate questions concerning the dynamics of an economic system at the micro level, nor can they be used to explore questions of structural evolution, since the internal structure of the system is not explicitly represented, and in fact is implicitly assumed to be unchanging. However, a dynamic general equilibrium model proposed by Yang and Borland (Borland & Yang, 1992; Yang & Borland, 1991) takes a step beyond this limitation. In their model, in which (1) producers become more efficient at production that they undertake more often, and (2) transaction costs are positive, if efficiency gains from specialization are greater than transaction costs for buying in goods not produced, then producers will become more specialized. In this case, because of the increased efficiency, per capita output of the system will increase while the number of producers of each good will decrease. Thus both growth and the structure of the economy are endogenously determined. In this respect the model is a departure from other general equilibrium models and closer to the model we propose here.

Another approach, much closer to ours, is embodied in the EURACE project, which developed an agent-based model of the European economy (Dawid et al., 2012; Deissenberg et al., 2008; Teglio et al., 2010). This model has a detailed representation of the various types of economic agents (consumers, firms, financial institutions, etc.) and their interactions. As in our model, there is a heterogeneous population of agents of each type, and interactions among agents are asynchronous. The model is inherently dynamic, and in fact the point of the model is to yield insights into the dynamics of the European economy. It does not, however, include the possibility of structural evolution. In principle it could be used as a foundation on which to build such an evolutionary model, although its complexity would make this difficult to achieve in practice. In contrast, our approach is focused primarily on the problem of long-run structural evolution and only secondarily on microscale dynamics as the underpinning of the structural evolution.

1.2 Models for Structural Evolution and Innovation

Structural evolution is the transformation of a given system into another system, one with a different structure and hence different behaviors. The process is not straightforward to represent formally. In the case of a system represented by a set of equations, structural evolution would mean additional variables, together with additional equations to integrate those variables into the system. However, systems of equations do not generate new variables and equations, and so the initial system would need to be embedded in an algorithm that could do that. For a system described algorithmically, structural change would require the algorithm to change itself as it executes. This is possible, and is the approach taken in this article, but such algorithms have not been widely used. In economic modeling, the problem of structural evolution has essentially been avoided by modeling at a more abstract level, where changes in system structure can be parameterized or treated as aggregate variables—for example, total output is used in this way in dynamic stochastic equilibrium models. This approach translates structural change into dynamics, but at the cost of making the process of structural transformation itself invisible. However, structural evolution is perhaps the most striking characteristic of the economic system, and the process needs to be treated explicitly in order to be understood.

Innovation is at the heart of structural transformation. Historically, it has meant the continual introduction of new products, new processes, and new types of agents into the system, and this transformation has resulted in increasing productivity and growing complexity. This process is similar in many ways to biological evolution. However, there are several important differences, and although it is appropriate to characterize structural change in the economic system as an evolutionary process, it is a different process than that found in biological systems. For example, while in both cases evolution depends on information flows, in biological systems the information is largely encoded in DNA and passed from parents to offspring, whereas in economic systems the flow of information is much less constrained, and in general it makes little sense to speak of parents or inheritance. In this article we use the terms *evolution* and *evolutionary model* to refer to processes that transform a system into another, different (but related) system. We do not thereby imply that the evolution of an economic system necessarily shares any particular characteristics with evolution in a biological system, and in fact our evolutionary model of the economic system embodies quite different processes than do models of biological evolution. Fundamentally, however, both economic and biological evolution represent structural change in the system.

Over the long run, structural evolution—the generation by the system of novel products, technologies, and organizational structures—is the most important characteristic of the economy, and it is the ultimate focus of our interest. However, it as an emergent property, and to understand it we must first understand the structure and dynamics of the economy, and this understanding must reside at the level where invention and innovation occur, i.e., the level of individual agents, since these are the drivers of structural evolution.

More than 80 years ago Schumpeter (1961) had this insight, and although his ideas continue to motivate some lines of inquiry (e.g., Festré et al., 2017), to date little progress has been made in developing formal theory. Nelson and Winter (1982) proposed an evolutionary model of the economy, but in fact the model was not truly evolutionary because possible innovations were pre-defined and selected from a list. Arthur (2009) and Arthur and Polak (2006) have proposed formal models of technological evolution, using logic circuits as the elements that combine to form more complex structures. Similarly, Angus and Newnham (2013) developed a generational model of technological evolution using finite state automata applied to bit strings. The model is fully evolutionary, but its highly abstract nature means that it cannot be made more realistic and hence it cannot give much in the way of detailed insights into the nature of structural transformation in actual economic systems. These combinatorial approaches are similar to our approach to generating new technologies; however, they are essentially stand-alone models and are not embedded in a model of the full economy. Solé et al. (2013) have proposed a model of technological innovations in a system of agents in order to show both parallels and differences with biological evolution as it unfolds in an ecological system. However, this work is relatively abstract and does not fully address the problem of innovation in an economic system. Dosi et al. (2010, 2013) have also proposed a formalization of Schumpeter's insight with a model of endogenous growth. This includes a process of improvement in capital productivity, and so takes a step toward modeling structural change; but the approach is essentially dynamical, and the model includes only one capital good and one consumer good. Dosi and Marengo (2014) examined another aspect of structural change by modeling the dynamics of organizational structure; this interesting work informs our current project to introduce organizational evolution into our model.

More generally, innovation comes in many forms, and is frequently treated as an emergent property of complex system dynamics (Krugman, 1995). Coenen and Díaz López (2010) compare theoretical frameworks for technological, socio-technical, and sectoral systems of innovation, while novelty as a driver of structural change in economic systems is addressed by Encinar and Muñoz (2006) and Gualdi and Mandel (2019). Some writers have focused on the measurement and statistical estimation of innovation (e.g., see Korres & Drakopoulos, 2009). Much of this modeling work on the innovation process is not at present directly relevant to our approach because the phenomena it addresses explicitly are subsumed under a simple stochastic process in our model; however, since our model is also a modeling framework, in future we will draw on this work to replace the stochastic process with a more detailed sub-model of innovation. Nevertheless, empirical work such as Korres and Drakopoulos (2009) is currently useful in providing a basis for the stylized facts we use to judge model results, and in future it may also help in model calibration.

While these approaches provide formal insights into several aspects of technological change in economic systems, they do not address the full problem of structural evolution directly. There is a field of evolutionary economics, but it consists largely of empirical and historical studies (Hodgson, 1993). Although, as Foster (2015) points out, some researchers use historical studies as the basis for simple simulation models, for the most part historical and theoretical approaches are seen as mutually exclusive. Nelson et al. (2018) provide an overview of recent work in the area of evolutionary economics.

The work of Potts (2000) and Dopfer and Potts (2008) is exceptional in that it provides an extensive formal treatment of structural evolution in economic systems. However, it is essentially descriptive, because it does not include an executable model of endogenously generated structural change, and therefore does not demonstrate evolution. Consequently it cannot be used as a platform for investigating various aspects of structural evolution and cannot be used to perform experiments. In contrast, an algorithmic model of the economy developed by Straatman (2008) and Straatman et al. (2008), a precursor to the model presented here, demonstrates open-ended structural evolution. It uses an agent-based model, in which a population of agents represents individual firms and consumers, and is able to show the emergence of innovations in the economy which thus transform it structurally. Agent-based techniques such as the one used in the Straatman et al. model would seem to be the most appropriate approach to modeling emergent phenomena, or even more limited aspects of the economy, as clearly demonstrated by Tesfatsion's labor market models (Tesfatsion, 1998, 2001, 2002). Gintis (2007) used an agent-based model to explore the dynamic properties of the Walrasian general equilibrium model (van Daal & Jolink, 1993), the well-known foundational model of modern economic theory. However, this work, which showed the power of introducing dynamics and heterogeneity to the conventional, static theory, had little effect. Farmer and Foley (2009), in a commentary in *Nature*, subsequently called for an agent-based approach for providing a new and more powerful foundation for economic theory, but to date their call has not, for the most part, been heeded.

This article is organized as follows: Section 2, gives an overview of the model approach taken here. Section 3 provides a description of the von Neumann technology matrices and their use. Section 4 presents our agent-based model of the economy, including a description of production functions, the price mechanism, agent properties and behaviors, and procedures for introducing and eliminating agents. In addition, the operation of the algorithm is described. Section 5 presents results from model runs, including the response to perturbations, dynamical phenomena, and the emergence of new products and technologies—i.e., structural evolution. The final section discusses possible uses of the model for detailed investigations of the dynamics of an economic system as well as its use as a platform for a richer treatment of innovation and structural evolution.

2 Overview of the Model

Our ultimate aim is to model the transformation of an economic system by means of the emergence of novel products, technologies, and organizational structures. However, since such evolution represents the continual structural evolution of a dynamical system, we must have both a suitable dynamical system and a process by which it can evolve. To this end, we develop a fully dynamic

$$\begin{bmatrix} p_{1,1} & p_{1,2} & \cdots & p_{1,n} \\ p_{2,1} & p_{2,2} & \cdots & p_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ p_{n,1} & p_{n,2} & \cdots & p_{n,n} \end{bmatrix} \rightarrow \begin{bmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,k} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,k} \end{bmatrix}, \begin{bmatrix} b_{1,1} & b_{1,2} & \cdots & b_{1,k} \\ b_{2,1} & b_{2,2} & \cdots & b_{2,k} \\ \vdots & \vdots & \ddots & \vdots \\ b_{n,1} & b_{n,2} & \cdots & b_{n,k} \end{bmatrix}$$

Figure 1. From Leontieff's quadratic input-output matrix (left side) to von Neumann technology matrices. *Left*: The (endogenous) Leontieff matrix has n inputs and n outputs. $p_{1,2}$ is the amount of input product 1 needed to produce product 2. *Right*: For the von Neumann technology matrices there are m inputs and n outputs. $a_{1,2}$ is the amount of product 1 used with technology 2. $b_{3,2}$ is the amount of product 3 produced using technology 2. There are m inputs, n outputs and k technologies in that model.

agent-based model of an economic system in which firms interact both with each other and with consumer-employees through market transactions.[1] Since the dynamics arise directly from the individual behavior of these agents and from their interactions with each other, as it does in the real economy, the model is able to mirror the real economy. We show first that the model has appropriate dynamical behavior, specifically that it is capable of attaining a steady state of prices, profits, outputs, and number of producers following a perturbation—though it is the transient disequilibria that drive the dynamics and provide openings for innovation. Having thus established that the dynamic model is a suitable platform for modeling structural evolution, we introduce a very simple evolutionary process by which new technologies and products appear in the system and demonstrate typical patterns of evolution that result. Our model for an economy is very general, because its purpose is to provide a conceptual and computational framework for investigating general issues of growth and change in economic systems. From an epistemological point of view, simulation is thus a tool for theoretical investigation, and for theory development. This is no different than what has been done in the natural sciences over many years (Winsberg, 2001), where it is now recognized that simulations (and the assumptions going into implementing the models that produce them) are another tool in the set of techniques used to uncover reality. In particular, when attempting to treat systems where *time* is not just a parameter but a constituting part of the model, as when innovations emerge (White & Banzhaf, 2020), step-by-step procedures embodied in simulating agent-based systems are powerful exploratory devices. Moreover, given the nature of agent-based simulations, our model framework can be extended or made more detailed as required to investigate particular phenomena of interest. In its evolutionary form it can even, to some extent, do this itself. In other words, as a framework, it has the potential to be filled with increasingly detailed substance.

2.1 Our Modeling Approach

The approach is based on the use of von Neumann technology matrices (von Neumann, 1945, see Figure 1). These are analogous to a Leontieff input-output matrix that is widely used to model the empirical structure of national economies. The Leontieff flow matrix disaggregates the economy into individual sectors. Each row and corresponding column of the matrix represents a particular sector: Each row shows the sales of that sector to every sector of the economy, while the corresponding column shows the purchase of inputs by that sector from other sectors. A technical coefficients matrix is derived from the flow matrix by dividing the values in each column by the total output of the sector corresponding to that column; the values for each column then sum to unity.

The Leontieff input-output approach has several disadvantages, however. The most serious for our purposes is that it is highly aggregated, even when hundreds of sectors are used, as is typically the case. In reality, within each sector different firms will be producing slightly different products, and even firms producing the same product may be using quite different technologies, and thus require different inputs.

1 Octave code is available at https://github.com/gustarcio/Recio-Economy.git.

The von Neumann approach, in contrast, disaggregates the input-output matrix into a pair of matrices, and represents production technologies explicitly, with rows representing products and columns, technologies. Each column of the output matrix shows the quantity of each of the outputs produced by that technology. One key advantage of representing technologies explicitly is that it is possible to include several technologies producing the same product. This is a common situation in real economies, almost always present when a new technology is introduced to produce an existing product. The advantage for modeling emergent technologies is evident. It is also possible to have one technology produce multiple products, and at a minimum there will always be at least two—an intended product and waste. This feature provides a potential link with environmental issues, which promise to become a focus of work in this area (e.g., see Farmer et al., 2015).

While the technology matrices, which are in effect a representation of two systems of equations, define the possibilities inherent in the economy, these possibilities are realized in our model by a population of agents representing both firms and consumers. In general, each product is produced by a number of agents, but these agents are to some degree heterogeneous, both in terms of their characteristics (e.g., location) and their current state (e.g., retained earnings or production rate). Consumer agents are also heterogeneous, having different patterns of desired consumption reflecting different utility functions, and different wealth. It is the consumer agents who provide the labor required by production technologies. Agents buy from each other in order to acquire the inputs needed to produce, or in the case of consumers, to acquire the consumer goods they desire; the prices at which they do so are determined endogenously by a process that tends toward market clearing. The process is dynamic, with each agent adjusting purchase and production levels at each time step. Furthermore, the size of the system is dynamic, with the number of agents varying during the course of a model run. If shortages of a product persist for a time, in spite of adjustments of output levels by existing firms, new agents producing that product enter the system; if surpluses persist, some agents fail and disappear. It is the variable size of the system that provides a window for innovation: New firms entering to supply a product in short supply can come with a new technology. In this case, the technology matrices acquire an additional column each. Similarly, the introduction of a new product adds a new row to the matrices.

The model is somewhat similar to that developed by Straatman and colleagues (Straatman, 2008; Straatman et al., 2008, 2013) in making use of von Neumann technology matrices. However, there are a number of significant differences. The market mechanism for price determination is more robust than that in Straatman et al. (2008) and much more efficient than those in Kephart et al. (2000) and Straatman et al. (2013). The characteristics of agents are quite different from those specified in the Straatman model. Firm agents have different decision rules and are restricted to using a single technology, while consumer agents are heterogeneous and more complex in their behavior, including in their labor market decisions. Finally, each technology involves the production of waste. This reflects the fact that any economic system must increase entropy. It will also, in future, permit the model to incorporate elements of the circular economy, in which wastes are recycled.

Ponzi et al. (2003, 2006) have also adapted the von Neumann approach to model the dynamics of an economic production network. However, their model, not being agent-based and having no stochastic elements, is analytical and thus gives very "clean" output. One striking result is that nested, multi-scale oscillations in prices appear. This is the result of non-linear dynamics in the production and price setting process when supply chains have several links. This contrasts with the dynamics of the evolution-free runs of our model (see section 6) which seem simple by comparison. This is due primarily to parameter values being chosen so that the dynamics stabilize relatively quickly, because our focus is on the nature of structural evolution of the system, which can be seen more clearly when the baseline dynamic behavior is not highly complex. Also, these runs are with small models with relatively simple supply chain structures. In runs where structural evolution is present (section 6), the models become much larger and the supply chain structures more complex. Consequently in those runs the dynamics of most variables are generally complex and multi-scale.

2.2 Important Considerations

In designing the model, the aim has been to make it realistic in its fundamentals while keeping it abstract and as simple as possible. The motivation for simplicity derives from the need to minimize run time. When the model is used for its ultimate intended purpose of representing the structural evolution of an economic system, it will be necessary to have a very large number of very long runs (tens or hundreds of thousands of iterations). This will only be possible with a relatively simple model. Of course reality is not simple, and we have also taken care in model design to make it possible to add realistic complications without having to alter its fundamental structure. Thus the model can easily be adapted to investigate more detailed and specific situations. This sort of easy adaptability to enhance the realism of a model is one of the advantages of an agent-based approach.

3 Von Neumann Matrices and Their Limitations

Technology matrices were introduced in 1945 by John von Neumann to show the conditions for the existence of a general equilibrium in an economic system undergoing balanced growth. Also called *von Neumann matrices*, they have occasionally been used to model growth in economic systems, most notably by Abramov (2014), as well as to model structural evolution (Straatman et al., 2013). Here we use them to represent the production processes by which inputs are transformed into outputs. They thus serve also to define basic economic sectors, where each sector can be identified with either a particular product or a particular technology. Both possibilities exist because in general a given technology may have more than one product as its output (although here, in our discussion of Von Neumann technology matrices, for simplicity we restrict technologies to one output), and several different technologies may exist to produce a given product. The latter situation is common in an economy undergoing structural evolution as new, more efficient processes appear to produce existing products.

Production processes are normally represented by equations that describe the transformation of inputs into outputs. For example, consider the simple economic system described as consisting of seven products and five manufacturing technologies. For brevity, we shall call this example henceforth the *Straatman Economy*. Products are labeled $P1$ to $P7$ and technologies, $T1$ to $T5$. By convention, product $P1$ represents labor and product $P2$ represents money. The production processes of the economy can then be described, for example, as

$$\begin{aligned}
&T1: 2 \times P1 && \rightarrow 2 \times P4 \\
&T2: 2 \times P1 + 2 \times P3 + 2 \times P4 && \rightarrow 3 \times P5 \\
&T3: 3 \times P1 + 3 \times P5 && \rightarrow 6 \times P6 \\
&T4: 6 \times P1 + 6 \times P6 && \rightarrow 6 \times P7 \\
&T5: 6 \times P7 && \rightarrow 13 \times P1
\end{aligned} \tag{1}$$

where each equation represents a given technology (see Table 1 for the corresponding von Neumann matrices and Figure 2 for a graphical depiction of the economy). Thus, the first technology, $T1$, uses two units of product $P1$, labor, to generate two units of product $P4$. Those products requiring only labor for their production, such as product $P4$, represent raw materials, e.g., those coming from the mining industry. Although it is being consumed, there is no production of product $P3$ and consequently it represents a free good, such as sunlight. The last equation of the above system represents consumption as it generates labor using a manufactured product, $P7$, which can thus be identified as a consumption good. In economic terms, the system described in Equation 1 is in balance as the total labor used equals the labor generated by consumption. The above equations can be converted into their equivalent von Neumann technology matrices to describe the production process.

In von Neumann matrices the columns represent technologies whereas the rows are identified with products. These equivalent matrices are shown in Table 1. Notice that the output matrix reflects

Table 1. Von Neumann technology matrices for the simple economy of Equation 1.

	Input matrix					Output matrix				
	*T*1	*T*2	*T*3	*T*4	*T*5	*T*1	*T*2	*T*3	*T*4	*T*5
*P*1	1	$\frac{2}{3}$	$\frac{1}{2}$	1	0	0	0	0	0	1
*P*2	0	0	0	0	0	0	0	0	0	0
*P*3	0	$\frac{2}{3}$	0	0	0	0	0	0	0	0
*P*4	0	$\frac{2}{3}$	0	0	0	1	0	0	0	0
*P*5	0	0	$\frac{1}{2}$	0	0	0	1	0	0	0
*P*6	0	0	0	1	0	0	0	1	0	0
*P*7	0	0	0	0	$\frac{6}{13}$	0	0	0	1	0

Note. Production processes or technologies are represented as columns *T*1, ..., *T*5, whereas products are represented as rows *P*1, ..., *P*7. *Left*: Input matrix; *Right*: Output matrix. Entries are normalized for the production of one piece of output each. *P*2: Money, not produced; *P*3: Free good, not produced; Raw materials, consumer products etc. are further described in the text.

our assumption that each technology produces only a single product, an assumption that can easily be relaxed. The first four columns of the input and output matrices, *T*1 to *T*4, represent production whereas the last one, *T*5, represents consumption. Note that product *P*2, representing money, is not included in the system of equations described in Equation 1, as it is only used to exchange products at their current price. In the matrix representation of the system, shown in Table 1, this situation takes the form of a zero valued row vector for product *P*2. The price for each product needs to be set in order to simulate the dynamics of the economy. Assuming that no profit is generated by any of the production processes involved (i.e., the price of the output equals the cost of the inputs), product prices can be described by the following set of equations:

$$\begin{aligned} T1 &: 2 \times p(P1) &&= 2 \times p(P4) \\ T2 &: 2 \times p(P1) + 2 \times p(P3) + 2 \times p(P4) &&= 3 \times p(P5) \\ T3 &: 3 \times p(P1) + 3 \times p(P5) &&= 6 \times p(P6) \\ T4 &: 6 \times p(P1) + 6 \times p(P6) &&= 6 \times p(P7) \\ T5 &: 6 \times p(P7) &&= 13 \times p(P1) \end{aligned} \tag{2}$$

The above system contains four independent equations (note that the last equation accounts for balanced growth and is computed as the combination of the rest) and six unknowns; therefore the system has 2 *df* and an infinite number of solutions. A definite solution is obtained if two of the prices are set exogenously. Thus we set the prices for labor and the free good to one and zero respectively. Then the system has the following set of prices as the unique solution:

$$p(Pi) = \{p(P1), p(P3), p(P4), p(P5), p(P6), p(P7)\} = \left\{1, 0, 1, \tfrac{4}{3}, \tfrac{7}{6}, \tfrac{13}{6}\right\}$$

These are the equilibrium prices, i.e., the prices that result if we assume equilibrium. The price of commodity *P*2, money, which is not included explicitly in Equation 2, can be set arbitrarily,

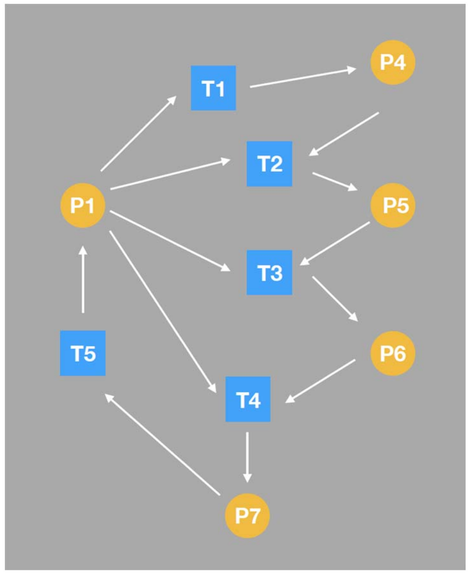

Figure 2. Simplified bipartite graph illustrating the von Neumann production process of Equation 1. *P*2 (money) and *P*3 (free good) are absent from the graph.

here to one for computational purposes involving the von Neumann matrices, giving the following complete set of prices for all commodities involved in the economy:

$$p(Pi) = \{p(P1), p(P2), p(P3), p(P4), p(P5), p(P6), p(P7)\} = \left\{1, 1, 0, 1, \tfrac{4}{3}, \tfrac{7}{6}, \tfrac{13}{6}\right\}$$

As mentioned, these prices represent an equilibrium situation involving the production of products at zero profit. Equilibrium conditions can also be calculated for positive profit levels, e.g., a 5% return rate for all technologies. Positive levels of equilibrium profit are useful for representing a system undergoing balanced growth. In this article we work only with systems where equilibrium profits are zero. However, since we are interested in the dynamics of the system rather than its equilibrium state, we treat sectoral profits as being in general non-zero, so that the system can explore non-equilibrium states, even though the equilibrium profit levels remain zero.

If we allow individual prices to deviate from their equilibrium, values corresponding to a given global profit level will permit individual technologies to exhibit different profits. We modify the system described in Equation 2 to include profits for each sector as:

$$\begin{aligned}
p(P4) &= p(P1) + \pi(T1) \\
p(P5) &= \tfrac{2}{3}p(P1) + \tfrac{2}{3}p(P3) + \tfrac{2}{3}p(P4) + \pi(T2) \\
p(P6) &= \tfrac{1}{2}p(P1) + \tfrac{1}{2}p(P5) + \pi(T3) \\
p(P7) &= p(P1) + p(P6) + \pi(T4) \\
p(P1) &= \tfrac{6}{13}p(P7) + \pi(T5)
\end{aligned} \tag{3}$$

where $p(Pi)$ stands for product prices as before and $\pi(Tj)$ represents the profit of the transformation described by the j-th column of the input and output matrices for each of the five technologies. With both profits and prices treated as variable in a dynamic system, they are no longer fully determined by Equations 2 and 3. Thus, additional mechanisms are required. Briefly, prices are determined by the gap between supply and demand (whether positive or negative), while supply and demand are determined by the aggregated production decisions made by individual agents (firms) based on anticipated profits. The latter are calculated using current prices and the coefficients describing the relevant technology. These are key aspects of the model developed in this article and will be described in more detail in the next section.

Returning to Equation 2, this system will become singular if the profit of the last equation, representing the transformation of final consumption into labor, does not satisfy

$$\pi(T5) = -\left(\frac{1}{3}\pi(T1) + \frac{1}{2}\pi(T2) + \pi(T3) + \pi(T4)\right) \tag{4}$$

and thus it will have no solution under these conditions. Note that in general, inconsistencies occur if the left-hand sides of the equations in the system are linearly dependent, and the constant terms do not satisfy the dependence relation. This dependent equation, representing the production of labor using consumer goods as inputs, is perhaps the main source of limitations of von Neumann technology matrices when modeling dynamics.

Treating the transformation of consumption to labor as just another technology has the advantage of introducing balance into the system, so that it generates the same total amount of labor as is required to manufacture the consumable product(s). This procedure works as long as the consumable product can only be manufactured in a single manner. For an economy where there are multiple different ways of creating the same consumable and, by extension, any other intermediate product, it becomes impossible to obtain such a balance from a single transformation equation representing consumption. This implies that consumption can no longer be considered an additional production process. Hence, an alternative way of treating the generation of labor in the system, capable of overcoming this limitation, is introduced and will be described in the following section. In brief, labor becomes available in proportion to the number of consumers in the system. All the available labor, or only part of it, may be used for executing the production processes. Consumers obtain money in return for their labor and that money can be saved or spent on consumption. This does not involve any profit for consumption and more closely resembles the consumer behavior of real economies.

In general, then, von Neumann technology matrices provide a simple but powerful description of production processes and have thus proven to be of value as the foundation of more than one model of an economic system. However, they have several important limitations, and these must be accepted or worked around. The most striking limitation from the standpoint of standard microeconomic theory is that, like the Leontieff framework, they assume constant returns to scale. This is a limitation we accept for the sake of computational efficiency and because it does not seriously affect our ultimate goal of modeling innovation and transformation in an economic system. For more realistic models, however, it would need to be relaxed. Other limitations have to do with the assumption of a fixed rate of profit and the use of a consumer technology to generate labor, both required to ensure analytic solutions. These we have relaxed, since we are interested in dynamics, by embedding the matrices in an agent-based simulation model which updates prices and production levels at each iteration, so that the path toward (or away from) equilibrium can be followed explicitly. The model also benefits from the ease with which the matrices can be expanded or contracted to represent the appearance of new products and technologies, or their respective disappearance.

4 The Dynamic Agent-Based Model of an Economy

The modeled economy consists of

1. Production processes, as described by the von Neumann technology matrices;
2. A population of agents representing firms that adjust production levels on the basis of expected profit;

3. Another population of agents representing consumers that earn income from labor and allocate their income to necessities, discretionary goods, and savings; and
4. A price adjustment mechanism permitting the system to move toward equilibrium in the absence of exogenous shocks.

Unlike the model described in the previous section, consumer agents appear in the von Neumann matrices only as labor inputs. The technology to produce labor has been removed from the matrices. In other words, consumers are now treated as a qualitatively different kind of agent compared to producer (firm) agents.

The model exhibits two kinds of dynamics. The first is what we might call the normal system dynamics—the continual adjustments of prices, production, and consumption levels that represent a movement of the system toward equilibrium. The second is the system size dynamics—a variation in system size as new firms and consumers appear and disappear in response to the profit opportunities and competitive failures created by the temporary imbalances characterizing the normal system dynamics. The system size dynamics represent a continual source of exogenous shocks that perturb the normal system dynamics. Above these two levels of dynamics is the level of structural evolution of the system, represented by the appearance of novel technologies, products, and types of agents.

4.1 Production Processes

The production processes are described by the von Neumann technology matrices (see Table 2 for an example). The column vectors of the input matrix are normalized to show required inputs per unit of output. For simplicity, each technology produces only a single target output, but it may also "produce" one or more of its inputs. This situation represents the use of capital goods, and the amount by which the output quantity is smaller than the input quantity represents depreciation; for example, in Table 2, *P*5, represents a capital good used in technology *T*5 with a depreciation

Table 2. Modified technology matrices where consumption is no longer treated as a production process (*T*5 becomes superfluous).

	Input matrix					Output matrix				
	*T*1	*T*2	*T*3	*T*4	*T*5	*T*1	*T*2	*T*3	*T*4	*T*5
*P*1	1	$\frac{2}{3}$	$\frac{1}{2}$	1	0.25	0	0	0	0	0
*P*2	0	0	0	0	0	0	0	0	0	0
*P*3	0	0	0	0	0	0.01	0.01	0.01	0.01	0.01
*P*4	0	$\frac{2}{3}$	0	0	0.9	1	0	0	0	0
*P*5	0	0	$\frac{1}{2}$	0	0.07818	0	1	0	0	0.07036
*P*6	0	0	0	1	0	0	0	1	0	0
*P*7	0	0	0	0	0	0	0	0	1	0
*P*8	0	0	0	0	0	0	0	0	0	1

Note. Additionally, *P*3 is no longer considered a free good and now represents the waste product. This model includes an extra manufacturing technology (*T*5) that generates an additional consumable product, *P*8. The economy being modeled here consists of eight products and five manufacturing technologies.

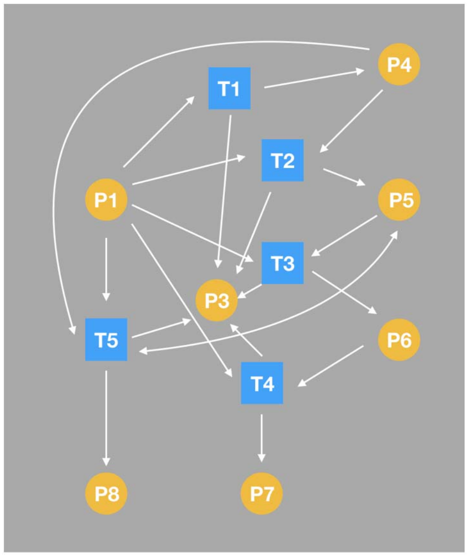

Figure 3. Simplified bipartite graph illustrating our economy of Table 2. *P*2 (money) is absent from the graph; *P*1 is labor, *P*3 is waste, *P*5 is a capital good, *P*7 and *P*8 are consumables.

rate of 10%—i.e., the output capital is 10% less than the required input (depreciation $= (0.07818 - 0.07036)/0.07818 \simeq 0.10$). In addition, each technology also produces waste. In the current version of the model, as shown in Table 2 and Figure 3, we work with a single, generalized, waste product for the sake of computational efficiency; since we are no longer using the free good shown in row 3 of Table 1, we appropriate that row for the waste product. We use the single waste product in row 3 as a placeholder for a more extensive representation in which each technology may produce different wastes. Waste is included explicitly since it is an inevitable by-product of any process, and waste products may become inputs for other production technologies at a later time. In principle, the full technology matrices are thus capable of modeling a circular economy in which wastes are recycled, and an evolving economy in which new technologies appear to make use of wastes.

4.2 Agents and Their Interaction Rules

An agent represents an autonomous entity that is able to observe certain aspects of the system being modeled and reacts to or acts upon its environment in response to those observations, on the basis of specified behavioral rules. In particular, an agent has the ability to interact with other agents by buying and selling goods. Agents are represented as vectors subdivided into blocks, where each block corresponds to a group of defining features. Four types of blocks are used: identity, possessions, production skills, and consumption pattern:

$$Agent = \{[\,identity], [\,possessions], [\,technology], [\,consumption]\,\} \tag{5}$$

The first block consists of three elements: the agent's geographical location in a 2D world, given by its x and y coordinates, and the agent number or ID. The next block represents the agent's possessions, i.e., how much of each product an agent has in stock (whether purchased or produced) and money; each agent is initialized with a stock of both money and the good it produces, so that inputs will be available for purchase by others. The length of this block corresponds to the number of existing products, but in general most elements will be zero. The technology block shows what production processes the agent is able to perform. It is Boolean valued, meaning that the agent has or does not have the ability to perform the production process described by the technology column of the von Neumann matrices; the length of the technology block corresponds to the number of technologies. Currently an agent is limited to implementing only one of the existing technologies. The fourth block defines an agent's consumption pattern, that is, the proportion of its income that is spent on each consumer good, including savings. It is randomly generated for each consumer, so that individual consumers differ in their consumption preferences. This last block has the same length as the number of consumable products in the economy and is generally normalized to one.

In the current implementation, agents are either producers or consumers, never both. A producer is identified by having at least one technology skill and all zeros in its consumption pattern. A consumer agent has all zeros in its technology block and some non-zeros in its consumption pattern. All agents are initialized with a standard quantity of relevant possessions: Producers are endowed with a stock of the product their technology produces as well as a stock of money; consumers are endowed only with money to afford consumption. The initialized stocks of product ensure that inputs are available in the system for the first round of production, and the initial stocks of money ensure that available inputs can be purchased.

In addition to defining agent structure, it is necessary to specify agent behavior. For producer agents, behavior is determined by the rate of profit, defined as the ratio of anticipated revenues to anticipated costs; thus a ratio of one represents the break-even or zero profit point. Anticipated profit is calculated using current prices; in other words, agents assume that current prices will continue into the next iteration. The underlying behavioral assumption is that producers will attempt to maximize their output when the rate of profit is high, but produce less at lower rates of profit. However, one should keep in mind that the current observed rate of profit is not necessarily the one that will be valid when the production is effected and sold in the next time period (iteration). Thus the lower the currently observed profit rate, the greater the risk that production undertaken now will be unprofitable by the time it is sold. On the other hand, in order to keep the business in operation, it is better to undertake some production even if at current profit rates it would be unprofitable. More generally, the current profit rate is an indicator of the supply and demand difference in the market for that product.

The current profit rate, defined as the ratio of the value of output at current prices to the cost of inputs, is easily calculated from the technology matrices and the vector of product prices:

$$\pi = \frac{p \times O}{p \times I} \tag{6}$$

A profit rate of 1 represents break-even, and rates <1 represent losses.

The planned production level Q of an agent is chosen to follow a sigmoid function:

$$Q = Q_{min} + (Q_{max} - Q_{min}) \left(\frac{1}{1 + e^{-\alpha(\pi - \pi_0)}} \right) \tag{7}$$

where Q_{min} and Q_{max} represent the lower and upper bounds of an agent's production, the parameter α determines the sensitivity of an agent to the rate of profit, and π_0 represents the level at which the agent makes the greatest increase or decrease in output in response to a change in profit (see also Figure 4). Q_{max} in effect defines the size of the firm, and π_0 can be set either to $\pi_0 = 1$ or to some other, probably higher, level like $\pi_0 = 1.05$ to represent a normal expected rate of return. For the examples shown in this article, the sigmoid response described in Equation 7 has been evaluated with parameters $Q_{min} = 0$, $Q_{max} = 1.5$, $\alpha = 10$, and $\pi_0 = 1$, and is shown in Figure 4. If a production agent operates for at least five iterations with low output levels ($Q < 0.3$, i.e., 20% of its maximal capacity in the current example), which is to say it is operating with significant losses, it runs the risk of being removed from the system as a failed firm as described below.

Equation 7 quantifies the agent's desired level of output, but it will not necessarily be able to achieve that level. Its plans are constrained both by the amount of money it has available to buy required inputs and by the global stock of available products. The agent thus produces only as much as these constraints allow, up to the desired production rate, Q. Note that for input, production, and

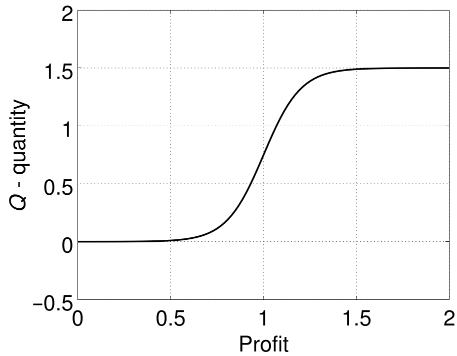

Figure 4. Production function $Q(\pi)$ (amount of production Q over profit π) used to simulate producer agent behavior. A profit level of 1 represents the break-even point, where costs are equal to revenue. $Q_{min} = 0$, $Q_{max} = 1.5$, $\alpha = 10$, and $\pi_0 = 1$.

consumption, the terms *quantity* (or *level*) and *rate* may be used interchangeably, since quantities here always refer to the same length of time—one iteration—and are thus effectively also a rate.

In buying inputs, a producer agent deals directly with other producer agents, approaching particular agents that can supply it with the required goods. The model has a spatial bias in this process, with the producer first approaching the nearest agents that could in principle supply the goods. If they are not able to supply them in the required quantities, other, more distant, producers are approached. This is currently the only use made of the spatial dimension of the model, and the selection of supply agents could just as well be random. The spatial element is included in the model in order to make possible future investigations of the evolution of spatial structures in the economy (e.g., business district, department store, mall, power center) but that is not done here.

Consumer agents behave in a rather different manner. In their role as employees they supply labor to producer agents, on demand and up to an upper limit representing the standard working hours per iteration. In return, they are paid for their labor at the current market price (i.e., the current wage rate), and that income permits them to buy consumer goods. Each producer hires individual consumers to supply labor—in other words, labor is not treated as an aggregated pool. Consequently, some consumers may be fully employed and others underemployed or not employed at all. Unlike the supplier-customer relationships among producer agents, however, distance is not a factor: Consumers are selected at random to provide labor to producer agents. This simplification was introduced for the sake of efficiency. It could easily be relaxed if a greater degree of realism were required, or if, for example, it were desired to model the dynamics of a spatially distributed population; but since the ultimate purpose of the current implementation is to model technological evolution this simplification is not a problem.

In their role as consumers, agents must consume a minimum amount s of consumable products at each iteration, allowing survival, here with a rate of $s = 0.15$. This is analogous to the metabolic rate included in Straatman et al. (2008). In the case of a single consumption good, each consumer must purchase at least 0.15 units of the consumable product in each iteration. In the case of multiple consumer goods, units are distributed among the different products in the proportions shown in the consumption block of an agent's definition vector. In addition, it is assumed that a consumer agent attempts to maintain a savings buffer that will cover the minimum consumption for five iterations at current price levels. If savings (the stock of product $P2$—money—in the possessions block of the agent definition vector) are less than this minimum level, any income above the amount required for the survival consumption is diverted to savings. Once savings are sufficient, additional income is divided among the various consumables, including savings, in the proportions shown in the consumption block of the agent definition vector. This is discretionary consumption. Note that the target minimum level of the savings depends on current price levels, so that if relevant prices rise, so does the target level of savings, thus forcing a decrease in discretionary consumption. If a

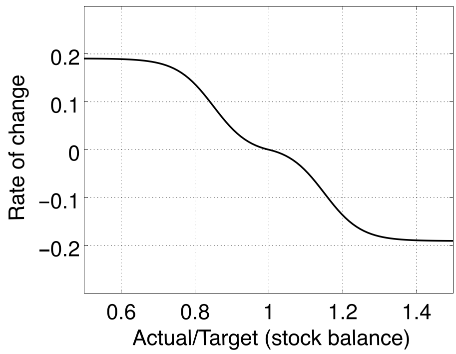

Figure 5. Price adjustment function $\Delta p(x)$ used to adjust prices as a function of the balance ratio x of actual and target stocks. Note that in the neighborhood of the equilibrium point at (1,0) price adjustments vanish.

consumer agent has exhausted its savings buffer and is unable to purchase the minimum amount of consumption goods, it is removed from the system.

4.3 The Price Mechanism

In previous agent-based economic models structured by von Neumann matrices (Straatman et al., 2008, 2013) we have used agent-based mechanisms to establish prices, and we have shown in both cases that the price mechanism gave prices that approximated the equilibrium prices for the system as calculated by the Leontieff inverse. Farmer (2002) has also proposed an agent-based price mechanism and shows that depending on agent strategies, it can yield different sets of prices with different efficiency properties. Because in this article we are not interested in prices as such, but rather prices as part of a production system that is able to equilibrate, we substitute a price adjustment calculation for agent-based price determination. We show that the calculated prices are able to return the system to a steady state following a perturbation, though they are not, strictly speaking, equilibrium prices since they continue to fluctuate within narrow limits. This price mechanism significantly reduces run time when the modeled systems become large, as in the case of evolving systems (section 6), which after several thousand iterations may have over 500 agents and 20 prices, and may continue to expand indefinitely.

Prices are adjusted in response to changes in the level of stocks of the various products in the system. The desired (target) level of stocks for a given product is defined as the average of stock levels over the past five iterations and determined by the number of agents producing the product times the standard level of stock with which producers of this product are initialized. If current stock levels are higher than the desired level, then the product is in surplus and prices are adjusted downward. If actual stock levels are lower than desired, then the product is in short supply, and prices are adjusted upward. This adjustment for the number of producers is necessary to avoid a situation where a product seems to be in short supply because of the disappearance of a producer. The product would appear to be in short supply because the total supply represents the output of remaining producers, while desired supply, representing an average over the last five iterations, would mostly reflect the larger output due to the presence of the lost producer's output in the totals for the first four iterations.

The price adjustment function again takes a sigmoidal form, symmetrical around the point at which the actual stock level equals the desired level, i.e., at $x = 1$ (see Figure 5). Deviations from the desired level are represented as a proportion of actual to desired level. Specifically, prices are adjusted according to

$$\Delta p = \begin{cases} -\left(\frac{\Delta p_{max}}{1+e^{-\alpha(x-(1+x_{off}))}}\right) + \Delta p_{max} - \Delta p_{adj} & \text{if } x \leq 1 \\ -\left(\frac{\Delta p_{max}}{1+e^{-\alpha(x-(1-x_{off}))}}\right) + \Delta p_{adj} & \text{if } x > 1 \end{cases} \tag{8}$$

where

$x =$ the stock balance (actual stock relative to desired stock);

$x_{off} =$ an offset with respect to 1, at which the inflection in the price adjustment curve occurs; $x_{off} = 0.15$ in Figure 5;

$\Delta p_{max} =$ the limiting value of the sigmoid function, here set equal to 0.2;

$\Delta p_{adj} =$ an adjustment to both tails of the piecewise sigmoidal function to give the same value when evaluated at $x = 1$; it is calculated as

$$\Delta p_{adj} = -\left(\frac{\Delta p_{max}}{1 + e^{-\alpha x_{off}}}\right) + \Delta p_{max} \tag{9}$$

Current prices p_t are then modified to

$$pt + 1 = p_t(1 + \Delta p) \tag{10}$$

4.4 Fluctuations in the Number of Agents

Consider first the production system proposed by Straatman et al. (2008) described in Table 1, where the dynamics is generated by agents that cannot vary their output. For simplicity, let the output of each agent be equal to one unit. Then the production system will be balanced with the following number of agents per technology (or any multiple of these): 2 agents with technology $T1$, 3 agents with $T2$, 6 with $T3$, 6 with $T4$, and 13 consumer agents. This is in fact the eigenvector-eigenvalue solution to the zero profit vector version of von Neumann matrices expressed as

$$O - (1 + \pi)I = O - I = 0 \tag{11}$$

Note that the value of 13 consumer agents is a requirement for the system to be balanced.

Now suppose that the agents are able to vary their output and consumption levels. With stationary prices, the above concentration of producers and technologies will still provide a balanced system. But now let a new producer agent of type $T4$ be introduced into the system. At current prices and profit levels, there will now be an overproduction of product $P7$, the consumable. There will also be an increase in the demand for the products used in the production of $P7$. The stationary state of the system has been perturbed, and as a consequence, prices will change. Eventually, the system will settle into a new stationary state, one with a different set of prices and in particular, a lower price for the consumable $P7$, which enables the unchanged number of consumer agents to consume more and thus produce the extra labor required by the new agent using $T4$.

In contrast, in the approach adopted in our model, where the consumption good ($P7$) is not treated as an input that a consumer agent uses to produce labor, i.e., where this process is removed from the von Neumann technology matrices (as shown in Table 2) and treated separately, the number of consumer agents no longer represents a rigid constraint for the balance of the system. A stationary balanced state can exist for a wide range of consumer agent numbers as long as some demand for the consumable product is generated.

Production, consumption, and prices fluctuate as the system moves toward equilibrium values for these quantities. However, the system is now continually perturbed by the introduction of new agents and by the elimination of unsuccessful ones. As long as new agents do not introduce new technologies, the system will still approach an approximate steady state in which production, consumption, prices, and numbers of agents of various types—i.e., system size—all fluctuate around their equilibrium values. When new technologies are able to appear, however, even equilibrium system size becomes indeterminate. In general, in the case of systems undergoing structural evolution, equilibrium becomes a local and temporary phenomenon.

New agents are introduced continually in our simulation, and are removed in response to the system dynamics. Introductions are of three types: introduction of a producer agent, introduction of a producer agent to recover a technology not currently being implemented, and introduction of a consumer agent. If a new producer agent is introduced, the technology assigned to the new agent is selected from the existing technologies under the criterion of possessing the greatest mean profit over the past five iterations. In the case of a new agent whose role is the re-introduction of a technology that is currently not implemented by any agent, special treatment is required, because in the absence of production, the mean profit over the last 5 iterations might have fallen to zero, so the technology is unlikely to meet the maximum profit criterion for introduction of an agent. Note that the absence of a single agent to implement a technology can destabilize the system, because ultimately it will become impossible to implement other technologies that require its product as an input.

If a new consumer agent is introduced and there is more than one consumer good, its pattern of consumption as defined by the fourth vector block of Equation 5 is determined randomly. New producer agents are endowed with a stock of both the product they will produce and money. New consumers are endowed with a stock of money. All of these stock endowments are equal to the initial endowments of the original agents in the system.

Producer agents risk being eliminated if their output is less than a certain fraction (0.3 in the runs discussed in this article) of their maximum production capacity for five consecutive iterations. Their production may fall below this threshold either because the current profit rate is so low that by applying Equation 6 they decide to under-produce, or because they are unable to purchase sufficient inputs due to the unavailability of those inputs, or because they lack the money to buy them. Consumer agents risk removal if they consume only the minimum or survival level for five consecutive iterations.

Whether an agent is introduced or removed, and if so, what kind of agent, are determined stochastically. Using a uniform random distribution ($0 < r < 1$), for simulations discussed in this article, the probabilities of introducing a new agent or removing an existing one are both equal to 20%. All agents satisfying the criteria for removal are in a single pool, and one is chosen randomly. Since normally there are many more producers than consumers in this pool, most removals are of producer agents. With an additional probability of 10% an agent is introduced that possesses a technology that was previously active but is currently not used by any agent. Finally, consumer agents are introduced with a 2% probability. Thus in each iteration, either there is no change in the number of agents, or one or more agents are added or removed. In the evolving systems described below (section 6) the maximum number of agents generated by the system is typically around 300 producers and 200 consumers.

4.5 Execution

The general flowchart of the algorithm is shown in Figure 6. The algorithm starts by determining all the initial parameters, including those defining products and technologies, the von Neumann input and output matrices, the initial prices for the products, and a population of consumer and producer agents with their initial stocks of commodities. Execution continues into the main loop. Here, the variables are updated, including those that will be used to control the dynamics of the system, such as the vector of average profits over five iterations, and the random reordering of the list of agents that determines the order in which they will execute their production or consumption plans. Next, each agent takes its turn to interact with the system by executing its plan, purchasing inputs, and producing outputs to contribute to the stock of goods in the case of producers, or by diminishing the stock in the case of consumers. Once all agents have taken their turn, the global stock levels are updated and prices are modified in response. Finally, the number of agents is adjusted as described in section 4.4. The process is repeated for a specified number of iterations.

Since agents execute their production and consumption plans (Figure 7) by trading directly with each other, rather than buying from and selling into global stocks, they cannot be updated

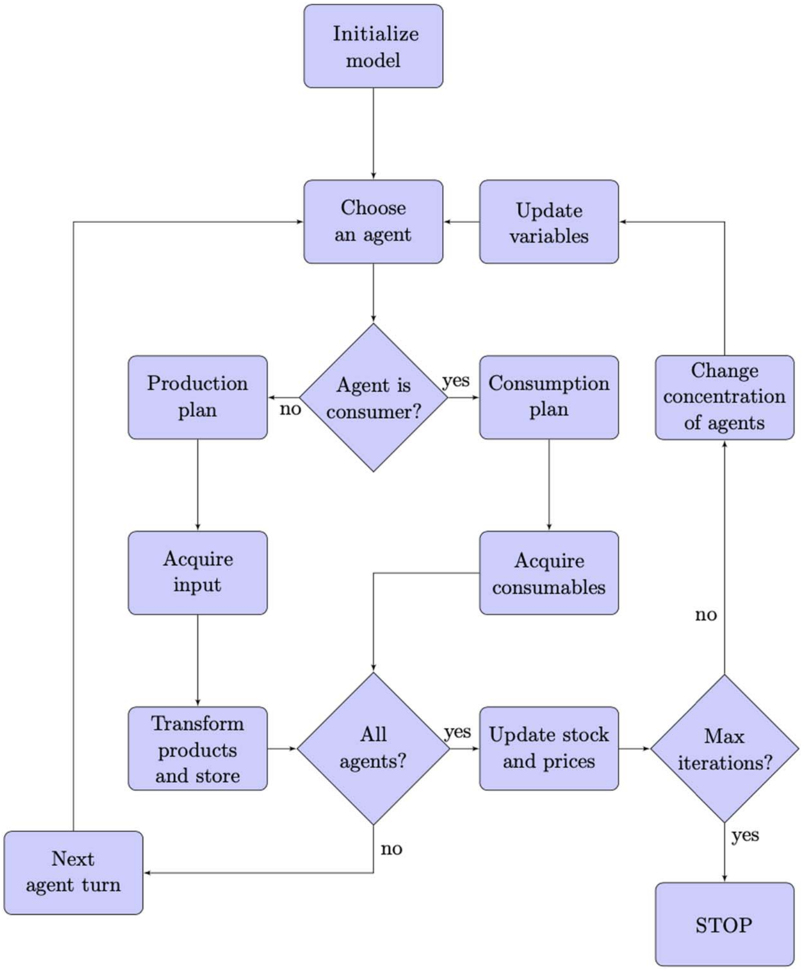

Figure 6. Flowchart of the algorithm. All agents present in an iteration operate in turn to implement their production plan. By executing their production or consumption plans, agents modify the global stock. When all agents have taken their turn, prices are updated according to current stock (demand and supply) and a new iteration begins.

simultaneously. However, this means that the environment in which an agent executes its plan depends on its position in the update queue. For example, if an agent is too far down in the update queue it might not be able to fully implement its production plan, because sufficient inputs are not available, whereas earlier they would have been. New production is not available for sale to other producers or consumers until the next iteration; thus all input purchases are of goods produced in the previous timestep. In order not to systematically favor some agents, the update queue is shuffled randomly for each iteration. Once all agents have taken their turn, global stock levels of each product are updated, and these values are used to establish new prices according to Equations 8 to 10. Also updated are the five-iteration averages of profit levels for each technology, and these, together with the other criteria mentioned in the previous section, are used in the process that introduces or removes an agent.

5 Results

A preliminary study of highly non-equilibrium situations of the system was already studied in an earlier conference contribution (Recio et al., 2020). Here we want to examine the system and its

```
PRODUCTION PLAN
---------------------------------------------------------------------------
  initialize production plan;
  %work out the maximum production limited by agent's money
  maxProductionA;
  %work out the maximum production limited by available products in the market
  maxProductionB;
  %expected production given by the sigmoid function (profit driven)
  maxProductionC;
  maxProduction=min([maxProductionA maxProductionB maxProductionC]);
  productionPlan=maxProduction;
  return productionPlan;

CONSUMPTION PLAN
---------------------------------------------------------------------------
  initialize consumptionPlan;
  %work out minimum consumption plan subject to pattern
  minimumPlan;
  minimumCost;
  %consume minimum if savings are not enough to stay in
  %the system for 5 iterations, otherwise discretionary consumption
  survivalAmount = 5 x minimumCost;
  if agentBudget less than survival then
     consumptionPlan=minimumPlan;
  else
     %money being not essential for survival
     excessMoney;
     %spend some of the excess money in discretionary consumption
     discretionaryPlan;
     consumptionPlan=minimumPlan+discretionaryPlan;
  end
  return productionPlan;
```

Figure 7. Pseudocode for production and consumption plans of agents. Production is limited by the minimum of the three possible measurements: the maximum production allowed with the current agent's budget, the maximum production allowed with the current market stock available, and an expected production function (profit-driven equilibrium). Consumers substantially reduce their consumption to a minimum when their budget is smaller than a threshold, otherwise they consume a random amount of their savings in discretionary consumption.

potential more systematically. With that intention, a number of experiments were carried out to assess the ability of the model to simulate fundamental aspects of the dynamics of an economic system. In particular we focus on the stability characteristics of prices, production, consumption, profits, and the number of agents of various types. All experiments, unless otherwise specified, implement the system described in Table 2. However, several experiments were run excluding a capital good (product *P*8) and its associated technology *T*5. Which model variant is being used can be determined from the legend box on the graphs. In the models described in the next section, when there are two different technologies producing the same good, the more efficient one generally eliminates the less efficient one. However, competing technologies may persist for long periods, or even indefinitely, given the continual system perturbation by new agents, products, and technologies. Also, new products, technologies, or consumers with different consumption patterns may induce price changes that reverse the relative efficiency of two technologies.

5.1 Operation of the Price Mechanism

The first experiment demonstrates the ability of the price mechanism to stabilize the system in the absence of exogenous events. Figures 8(a) to 8(c) show, for a typical run, the available stock levels of products, the product prices, and the profit rates for each of the four technologies. The number of agents was kept constant and the model was run for 500 iterations. The high frequency variations in the stock of available products (Figure 8(a)) are the result of the stochastic behavior of consumers, which generates minor fluctuations in the demand for the consumable product, *P*7. The target stock for this product was set to 60, and in fact the plot oscillates around this value for the entire run. Variations in the demand for the consumable product cause variations in the demand for the inputs to produce it; thus the effect ripples through all sectors of the economy, as seen in the figure. Prices (Figure 8(b)) also show high frequency fluctuations as they respond to a changing

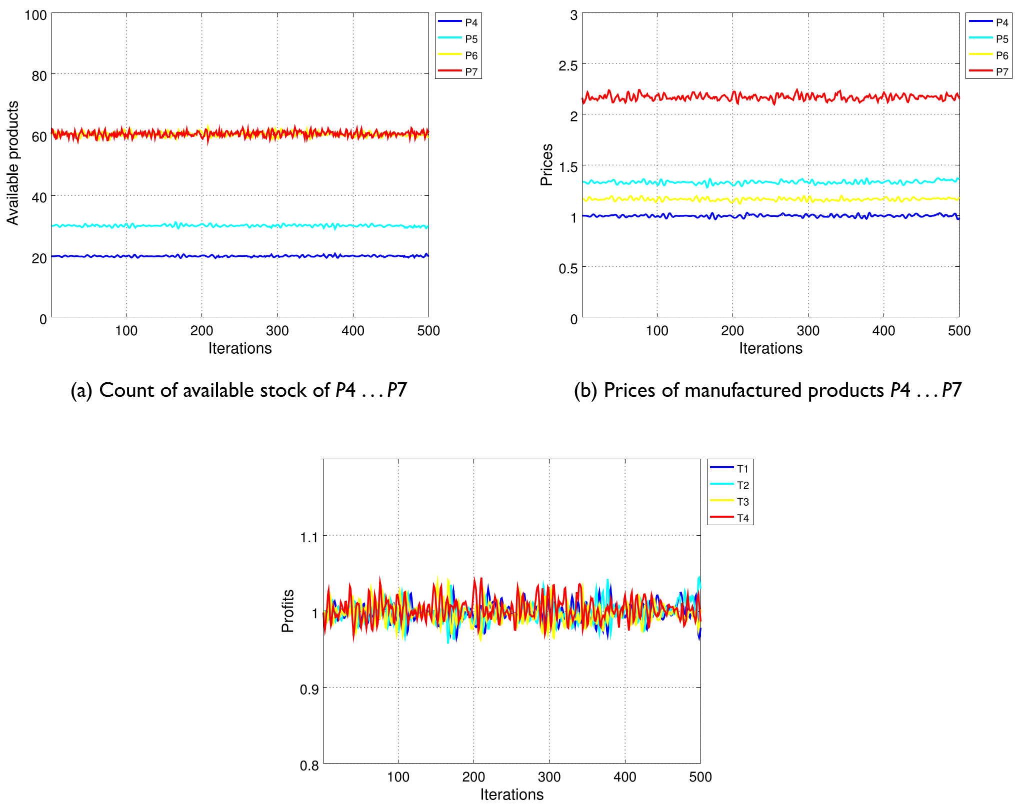

(a) Count of available stock of *P4 . . . P7*

(b) Prices of manufactured products *P4 . . . P7*

(c) Profit ratios of technologies *T1 . . . T4*

Figure 8. (a) Stochastic behavior of consumers is directly related to minor fluctuations in the demand for the consumable product, *P7*, which at the same time, causes variations in the inputs to produce it. These demand-driven variations create the ripple effect seen in all sectors shown in the figure. (b) Prices show similar fluctuations in response to agents' changing supply and demand of products. All prices fluctuate around their equilibrium values, a sign of an effective price mechanism. (c) A profit ratio of 1.0 represents zero profit, since the profit ratio is the ratio of revenue to cost. Note that all profit ratios fluctuate around their equilibrium value (zero profit), which gives an indication of the effectiveness of the price mechanism.

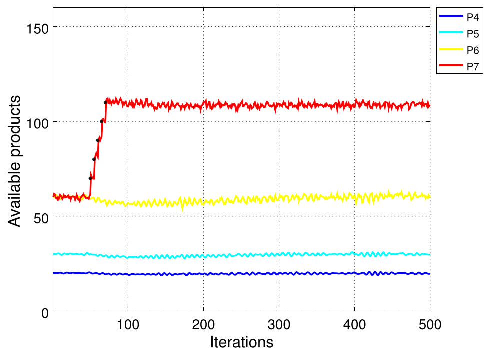

(a) Increasing number of producers for producing *P7*

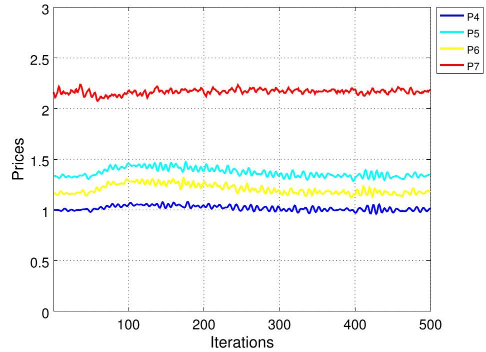

(b) Price variation when new producers of *P7* appear

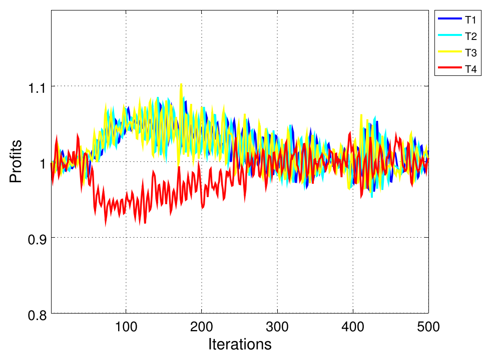

(c) Profit ratios similarly need some time to get back into balance

Figure 9. (a) Black dots indicate the increase in stock due to the introduction and endowment of new agents. Count of other products temporarily dips. (b) Note an increase in the price of input products (*P4*, *P5*, *P6*) for a period of time until the system returns to balance again. This is due to gradual adjustments in the economy. (c) In this case, the current demand for product *P7* must be shared among more producers every time a new agent is introduced and hence their profit gets substantially reduced. This turns out larger profits for required inputs due to changes taking place gradually.

supply and demand situation. Prices are used in Equation 6 to calculate the profit ratios (Figure 8(c)), which fluctuate within a 5% band around their equilibrium ratio of 1.0. The profit ratios are in turn used in Equation 7 to determine production levels of the individual agents, and these in aggregate constitute the available stocks of the products shown in Figure 8(a). Since all of these quantities fluctuate around their equilibrium values, the price mechanism is working efficiently.

5.2 Varying the Number of Producer Agents

The next set of experiments is aimed at investigating the dynamic response of the system to variations in the number of producer agents. In order to clearly single out these effects, rather than introducing new agents continually in the manner described above, in the first experiment we add a total of five new agents using technology *T4* to produce the consumer good *P7*. These agents are introduced at iterations 50, 55, 60, 65, and 70 respectively. Each new agent is initially endowed with the same stocks of money and product *P7* as the other agents producing *P7* had at initialization. The target level of available stock for the system is increased by the same amount, thus growing from 60 at iteration 50 to 110 at iteration 70. Results are shown in Figures 9(a) to 9(c). The introduction of a new agent means that the current demand for an input product must be shared among more producers. Since supply does not initially change, all producers will, on average, experience reduced profit levels and will, as a consequence, reduce output. However, the reduction happens gradually,

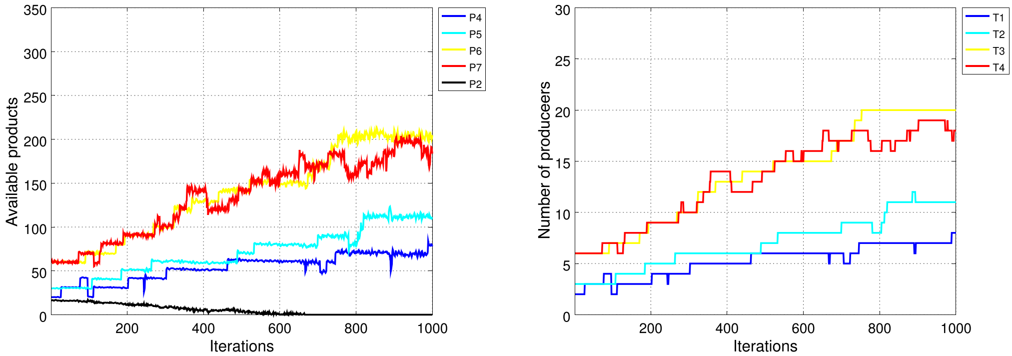

(a) Available stock of products *P*4, . . ., *P*7

(b) Producers for technologies *T*1, . . ., *T*4

Figure 10. Continuous variation in the number of agents according to the method described in Section 4.4 for 1,000 iterations. (a) It is worth mentioning that the labor shown in this plot (the black line, *P*2) corresponds to labor available but not currently being used, when this value falls to zero the labor supply of the system is fully used. (b) The system shows an initial growth followed by a slowdown as it tends to get into balance once the labor supply is fully used.

and in the meantime the production of input products (*P*4, *P*5, *P*6) must increase, with the increased production elicited by higher prices (Figure 9(b)) and hence higher profits (Figure 9(c)).

The adjustment process takes on the order of 200 iterations, as can be seen most clearly in Figure 9(c), where by iteration 250 profits for all technologies are again oscillating around the neutral level 1.0. Since no agent is removed from the system in this experiment, profits are only able to return to the neutral level for all technologies (and in particular *T*4, with the increased number of agents) because initially there is excess labor in the system. In other words, at the outset consumer agents were not able to sell all the labor they were willing to make available. As the price for *P*7 fell initially, they were able to consume more of it, thus increasing demand for that product, and ultimately, through the supply chain, other products as well. The extra production meant that they were collectively able to sell more labor, thus increasing their incomes and, through their purchases of *P*7, absorb the additional production of this good due to the five new agents.

The next experiment shows the adjustment process operating when agents are introduced and removed continually, as described in section 4.4. However, in order to increase the removal rate for illustration purposes, the threshold criterion to determine when an agent should become a candidate for removal is increased from 0.3 to 0.5 of the maximum production level. Figure 10(a) shows the stock of available products. In the case of labor (the black line, *P*2), the graph shows labor available for use but not currently being used. Once it falls to zero, all labor available to the system is being used. As can be seen, the system is able to grow in terms of both stocks of available products (Figure 10(a)) and number of producer agents for each product (Figure 10(b)). The labor supply is fully used around iteration 300, but both the number of producers and available stocks continue to grow for another 200 or so iterations. This is possible because, while the number of producers grows beyond the point at which the labor supply is fully used, not all are able to carry out their production plans. They thus become candidates for removal. However, the removal rate is relatively low, so the number of producers in the candidate list grows for approximately 200 iterations until the number removed at each iteration equals, on average, the number introduced. This lag is unrealistic, and can easily be corrected either by modifying the criteria for introducing new producers or by increasing the removal rate. Prices (Figure 11(a)) and profits (Figure 11(b)) fluctuate within relatively constant limits, although in both cases the fluctuations are greater than in the absence of entry and exit of agents, and in the case of profits the fluctuations are somewhat greater while the system still has excess labor available.

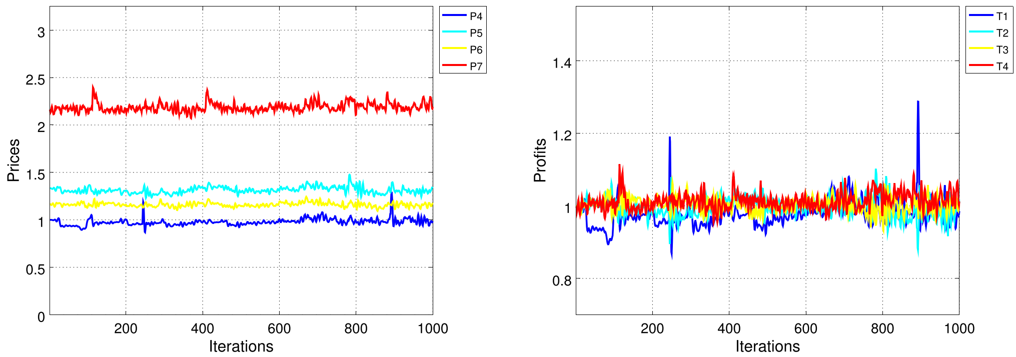

(a) Price fluctuations (b) Profit fluctuations

Figure 11. Fluctuations with a continuous variation in the number of producer agents. (a) Though within limits, price fluctuations are noticeably greater than in the absence of variation in the number of agents. (b) Profits, as directly related to prices, show similar fluctuations that are also greater than in the absence of variation in the number of agents. The peaks of these fluctuations tend to gradually decrease as the system labor becomes fully used.

5.3 Multiple Consumer Goods

Figures 12 and 13 show the behavior of the system when a new consumer good, *P*8, together with new technology required to produce it, *T*5, is introduced at iteration 100. After an initial shock, prices (Figure 12(a)) and profits (Figure 12(b)) quickly stabilize, but the available stock of goods (Figure 13(a)) and the number of producers (Figure 13(b)) both continue to grow until about iteration 1,000. This long adjustment period is largely due to the fact that labor supply is not yet fully used when the new product is introduced, so the system adjusts as the amount of labor used continues to grow until about iteration 600, and the adjustment continues for another 400 or so iterations after the labor supply is fully used, until the number of agents entering the system is balanced by the number of those being removed, as described above. The system finally reaches a stationary state at about iteration 1,000. The introduction of the new product, *P*8, seems in itself to have little effect on the system. Figure 13(a) shows the availability of the waste product (*P*3 in Table 2). Note that its supply is cumulative and thus continually growing, since waste removal is not implemented in this system and, in particular, there is no technology making use of the waste as an input.

If more than one consumer good is available, consumers display different consumption patterns. Figure 14 shows consumption patterns of two different consumers when the second consumption

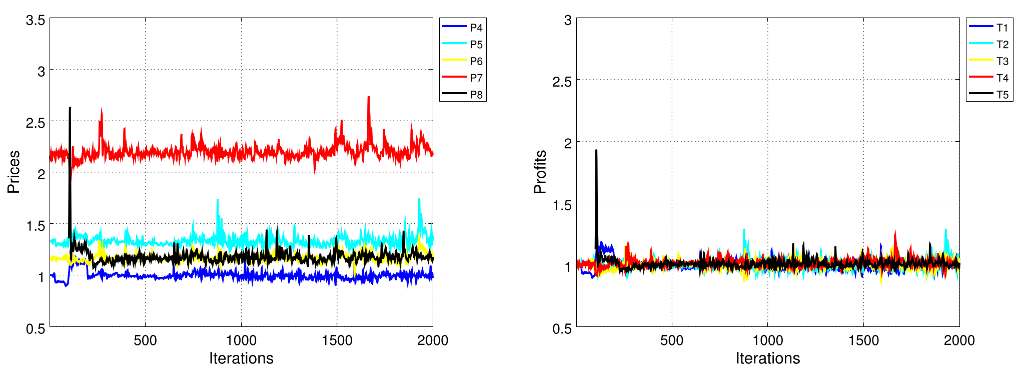

(a) Prices for five products *P*4, . . ., *P*8 (b) Profits using the five technologies *T*1, . . ., *T*5

Figure 12. Initial shock after the introduction of new consumer product *P*8 with technology *T*5. (a) The new product quickly draws a large demand, which causes a price increase. The system settles down after a short period of time. (b) Profit quickly settles down as the supply and demand of the new product are adjusted through the prices.

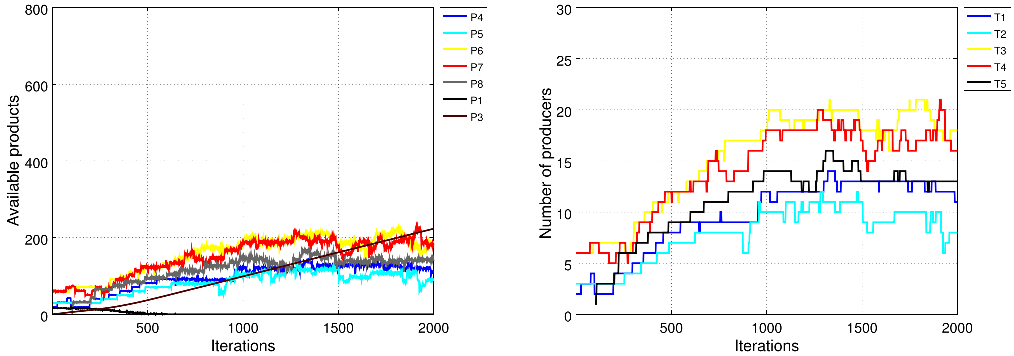

(a) Available products $P4, \ldots, P8$ plus $P1, P3$ (b) Producers using technologies $T1, \ldots, T5$

Figure 13. (a) There is no shock after introducing new consumable *P*8, but a gradual adjustment while the labor supply is becoming fully used and for some 400 iterations afterwards. Waste product *P*3 continues to be produced through the run. (b) Note gradual adjustment during and after the saturation of the labor supply until a steady state is achieved at about iteration 1,000.

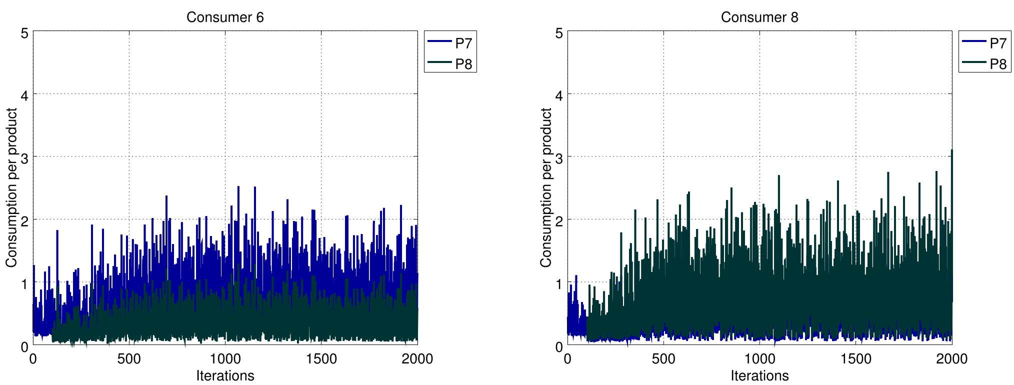

(a) (b)

Figure 14. Consumption patterns of two goods by consumers 6(a) and 8(b). At every iteration the amount of consumption for each consumable is proportional to an encoded pattern that can be found in the agent's definition (consumption share parameters).

good (*P*8) is introduced at iteration 100. All consumers alter their consumption iteration by iteration in response to variations in their labor income and the need to maintain a buffer amount of savings. In terms of preferences, Consumer 6 consumes the two products in approximately equal amounts (the actual share parameters are 0.507134 and 0.492866), though the volatility of consumption of *P*8 seems to be greater, while Consumer 8 has a clear preference for product *P*7 (share parameters: 0.601994 and 0.398006). Currently the consumers do not adjust the product mix of their consumption as the relative prices change. The share parameters for all consumers are shown in Table 3.

5.4 Multiple Technologies to Produce a Single Good

We now consider an economy with two ways of producing a single good. The technology matrices of previous examples have been expanded with a new column, *T*6, representing the additional technology (see Table 4 and Figure 15).

The simulation starts using two agents with technology *T*1, three agents with technology *T*2, six agents with technology *T*3, three agents with technology *T*4, another six with *T*6, as well as 13 consumer agents. With respect to the previous setup, rather than six agents using *T*4, the production

Table 3. Consumption share parameters for products *P*7 and *P*8 for all 13 consumers.

Consumer	*P*7	*P*8
1	0.979140	0.020860
2	0.465138	0.534862
3	0.250548	0.749452
4	0.796751	0.203249
5	0.792195	0.207805
6	0.507134	0.492866
7	0.393417	0.606583
8	0.601994	0.398006
9	0.973243	0.026757
10	0.587346	0.412654
11	0.709500	0.290500
12	0.285244	0.714756
13	0.437202	0.562798

Note. As an example, out of the total consumption of consumer agent number 3, roughly 25% is spent on consuming product *P*7, whereas 75% is spent on *P*8.

of the consumable *P*7 has been split between the two technologies, with three agents using each of them (*T*4 and *T*6). Since the new technology uses different inputs to produce the same output, the demand for intermediate products is different, and so prices must adjust from their previous values. In order to facilitate comparison with the previous results, the number of agents of each type was frozen. In the absence of fixed agent numbers, it is likely that the new technology would have ultimately been eliminated, because the price levels prevailing during approximately the first 300 iterations render it a less efficient technology, as can be seen in the relative profit levels for the two technologies (Figure 16(a)). Under those circumstances more producers using the more profitable technology would have entered, and successive producers using the less profitable one would have lost market and been removed. However, in the present circumstances, in order for the system to equilibrate, the prices for products *P*4, *P*5, and *P*6 must rise (Figure 16(b)), and eventually the stock of these products falls slightly (Figure 16(c)).

5.5 Structural Evolution

The major part of the present article has been an evaluation of the dynamical behavior of the agent-based model whose structure is described by a pair of von Neumann matrices. The point was to show that the model has good stability properties and generates plausible dynamical behaviors. However, the ultimate reason for developing the model is to use it to investigate the structural evolution of such an economy. This requires augmenting the dynamic model with an endogenous procedure for continually modifying the von Neumann matrices as the model executes. In this section we describe that procedure and present several examples of the resulting structural

Table 4. Technology matrices modified to have multiple technologies producing a single commodity.

	Input matrix					Output matrix				
	*T*1	*T*2	*T*3	*T*4	*T*6	*T*1	*T*2	*T*3	*T*4	*T*6
*P*1	1	$\frac{2}{3}$	$\frac{1}{2}$	1	1.9	0	0	0	0	0
*P*2	0	0	0	0	0	0	0	0	0	0
*P*3	0	0	0	0	0	0.01	0.01	0.01	0.01	0.01
*P*4	0	$\frac{2}{3}$	0	0	0	1	0	0	0	0
*P*5	0	0	$\frac{1}{2}$	0	0	0	1	0	0	0
*P*6	0	0	0	1	0.2	0	0	1	0	0
*P*7	0	0	0	0	0	0	0	0	1	1

Note. Two technologies, *T*4 and *T*6, are producing the consumable product *P*7 in a different manner, i.e., using differing amounts of resources.

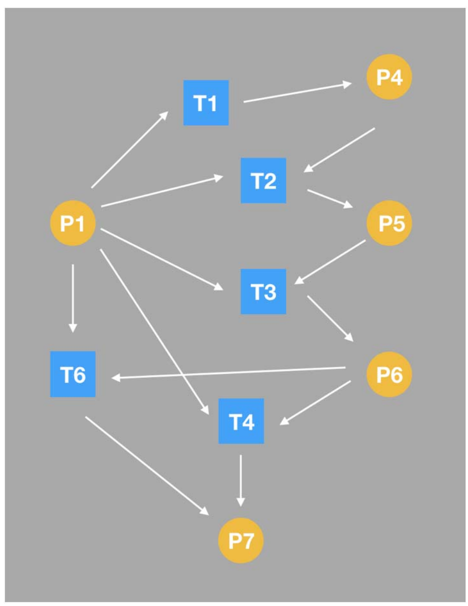

Figure 15. Bipartite graph illustrating the von Neumann production process with two technologies for one product. *P*1 is labor, *P*7 is consumable, *P*2 and *P*3 are not shown.

evolution of the system. The aim here is not to provide a comprehensive analysis of the evolutionary possibilities of the system—that will be the subject of a subsequent article—but to illustrate that the model is capable of generating a rich but realistic variety of true evolutionary behavior.

Before describing the use of the model to explore structural evolution, it is useful to recall that the evolution of an economic system can be modeled in different ways, some close to the paradigm of biological evolution, some more remote and closer to biological learning, intentional design, or other phenomena, like physical phase transitions, etc. However, when firms are introduced using

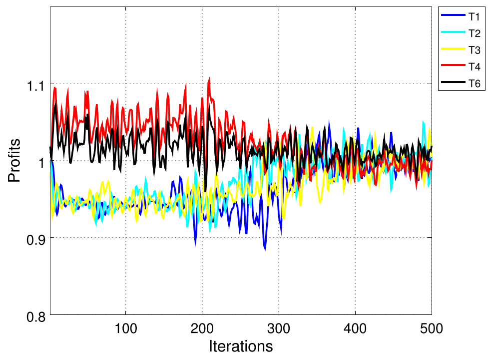

(a) Profits

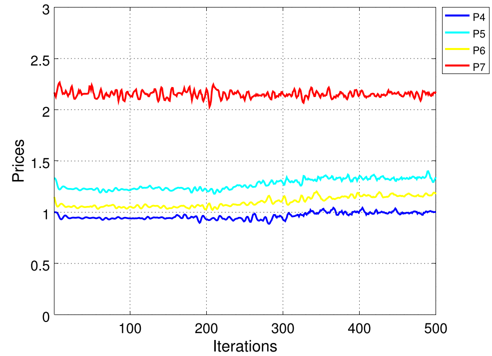

(b) Prices

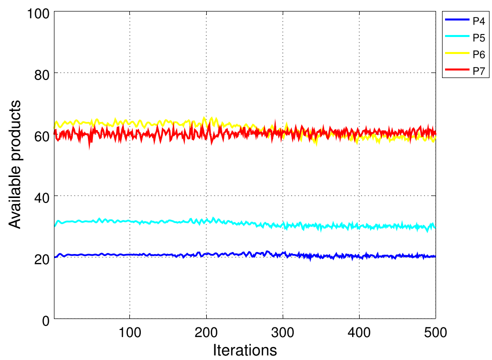

(c) Available products

Figure 16. Situation when there are two ways of producing a single good (*T4* and *T6*). (a) In the absence of a fixation of agent numbers, it is likely that the new technology *T6* would have ultimately been eliminated. This technology is clearly less efficient as can be seen in the profit plot. (b) In order for the system to equilibrate, the prices for products *P4*, *P5*, and *P6* must rise. (c) Eventually the stock of products *P4*, *P5*, and *P6* falls slightly.

a technology or producing a product not previously present in the system, this process can be described as analogous to an innovation process in natural evolution. A new technology uses as inputs products already in the system, but in a novel combination, and may produce either an existing or a new product, i.e., one not previously present in the system. In the case of a new technology, the set of inputs is chosen randomly. The introduction of a new product is also a random event. This stochasticity plays the role of mutations in a biological system. Selection occurs once the firm embodying the new technology or producing the new product enters the system: Either it is adopted by the system or it fails and is eliminated.

Growth and evolution in the economic model are forced by the continual introduction of new agents, technologies, and products. Producer and consumer agents are introduced as described above in section 5.4. Specifically, at each iteration a new agent (either consumer or producer) is introduced with a probability 0.075, an agent on a removal list is removed with a probability of 0.125, and a producer agent for an inactive technology is reintroduced with a probability of 0.05. Thus, on average, an agent is added or removed every fourth iteration. If the coefficients in the technology matrices are such that the economy is sustainable, so that agent additions exceed removals, then in the absence of innovation (i.e., the introduction of new technologies and products) this forcing produces balanced growth in the system—the phenomenon for which von Neumann originally proposed the matrices. Here, however, in order to model structural evolution, a certain proportion of new producer agents are introduced with a novel technology that either produces a new product, or produces an existing product in a new way, using a different set of inputs. In general, a new technology is introduced as a pair of technologies, the first producing a new intermediate product and the second using the new intermediate together with other inputs to produce a new final demand product (i.e., a consumer good). This approach permits the introduction of new producer goods while ensuring that they have a market; the new producer good is then available for use in other new technologies introduced subsequently. However, if a new technology is introduced to produce an existing product then this is done using only a single technology rather than a pair, since the market already exists. Note that since the introduction of new firms, new products, and new technologies are stochastic events, as is the selection of inputs that defines a new technology, innovation is not hand-coded.

In order to see evolutionary trends it is necessary to have long model runs, so that a significant amount of innovation can occur able to restructure the economy in a major way. However, as the size of the evolving system grows (each product or technology introduction increases the size of the von Neumann matrices) the run time per iteration becomes longer, so it is desirable to generate as much evolutionary activity as possible in a limited number of iterations. Of course this evolutionary activity also results in extinctions of products and technologies, and if the rate of introductions is too large, the system tends to become evolutionarily unstable and collapses to one that is small, i.e., one with a limited number of agents, products, and technologies—but which is nevertheless undergoing a continual churn of introduction and elimination of agents, products, and technologies. In other words, this is an instance where an evolutionary process operates to produce change without evolution. The rates of introduction of new products and technologies were therefore adjusted to maximize the potential rate of evolution while ensuring that the majority of systems were in fact able to evolve.

In order to get some idea of the behavior of the model when run in its evolutionary mode, we made 60 runs with parameters controlling the introduction rates set to the highest levels consistent with evolutionary stability. Qualitatively, the results can be characterized as showing one of four patterns: long-term steadily successful evolution ("steady evolution"), long-term successful evolution occasionally interrupted by a crisis ("evolution with crises"), evolution that starts successfully but then experiences a crisis from which the system does not recover ("failed evolution"), and evolution that never gets under way ("no evolution"). There were only a few cases of this last type. In these runs either newly introduced technologies were quickly eliminated, or they were adopted by the system, but at the cost of the loss of a similar number of existing technologies; as a result, while the particular technologies present in the system changed, the system did not become more complex. Table 5 shows the distribution of runs among these four evolutionary classes. Since all parameter values were identical for these runs, the differences in outcome are due solely to the effect of the

Table 5. Run statistics for the 60 runs executed until iteration 10,000.

Type	Number of occurrences
Steady evolution	11
Evolution with crises	20
Failed evolution	24
No evolution	5
Total	60

Note. Number of runs per class. Crises are defined as a drop of more than 25% in the total number of producer agents in the system.

stochastic perturbation. Here we show a typical example of each type. For these runs the parameters were such that the probability at any iteration of introducing a new technology or technology pair was 0.001. Within this global figure the probability of introducing a new consumer good was 0.0002, i.e., one fifth of new technology introductions produced a new consumer good.

A typical successfully evolving system with a crisis is shown in Figure 17. Over the 10,000 iterations of the simulation it exhibits growth in the number of consumers, producers, technologies,

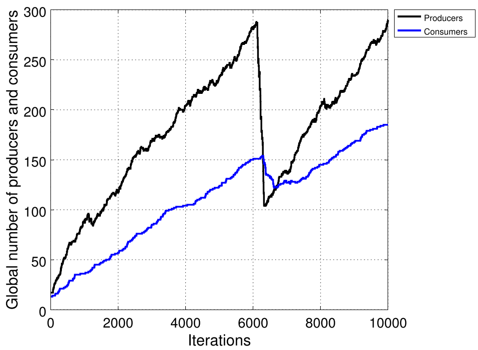

(a) Number of producer and consumer agents

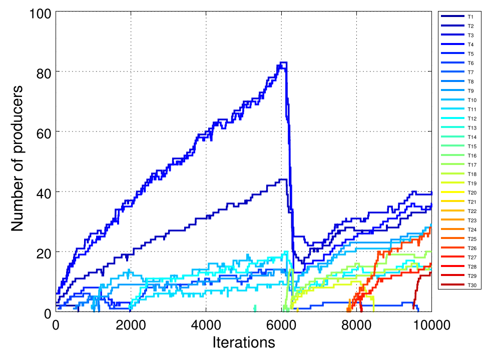

(b) Number of producers for each technology

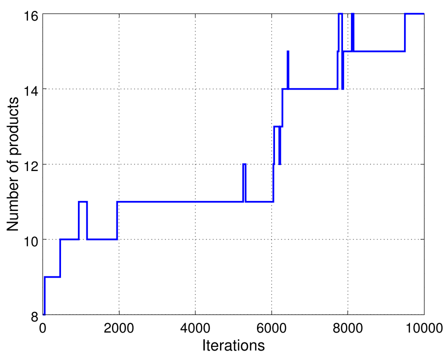

(c) Number of products available in the system

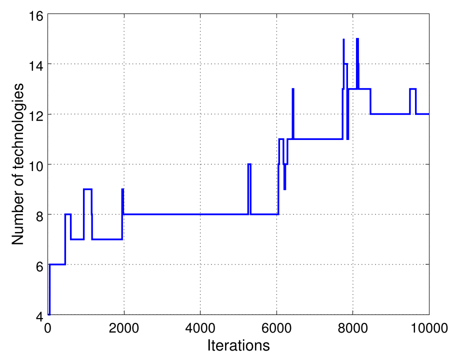

(d) Number of technologies present in the system

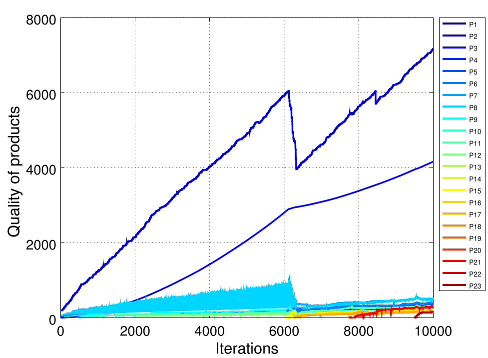

(e) Available products present in the system

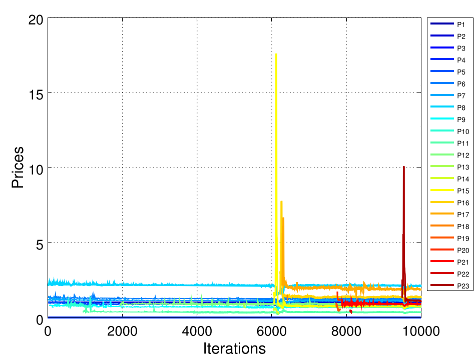

(f) Product prices

Figure 17. Evolution with crisis. In (b) the two nearly coinciding upper curves peaking at about iteration 6,000 correspond to the initial consumer good and the manufactured good used to produce it; the curve below them corresponds to the manufactured good used to produce the manufactured good shown by the upper curve.

and products. However, in addition to minor fluctuations, it experiences one major crisis, at around iteration 6,000, with a crash in the number of producers followed shortly by a rapid decline in the number of consumers (Figure 17(a)). Consumers are eliminated because of a shortage of the goods that feed them, due in turn to a shortage of the goods required to make those products. It should be noted, however, that not all crises result in a reduction in the number of consumers; in about half of the runs exhibiting evolution with crises, the number of consumers never declines. As shown in Figure 17(b), this particular crisis was initiated by the appearance of several new technologies, especially *T*17, but note that the effect was disproportionate to the magnitude of the shock: Two of the new technologies disappeared relatively quickly, and *T*17 peaked at only 14 producer agents before undergoing its own (temporary) crisis due to the entry of even newer technologies. The net result of the crisis was that the structure of the system was transformed, with several new technologies becoming relatively important and the number of products and active technologies increasing significantly (Figures 17(c) to (e)). This is reminiscent of a Schumpeterian gale of creative destruction. It also resembles the behavior of Lindgren's (1991) iterated *N*-person prisoner's dilemma model in which agents may evolve new strategies. That model produced long periods of relative stasis in which certain strategies predominated interrupted by short crises of unstable behavior. Angus and Newnham's (2013) evolutionary model also generates intermittent creation and destruction events. The emergent temporal clustering of innovation or structural evolution thus seems to be a robust feature of both evolutionary models and actual economic systems—as well as, perhaps, biological systems. Another burst of innovation at around iteration 8,000 further transformed the system

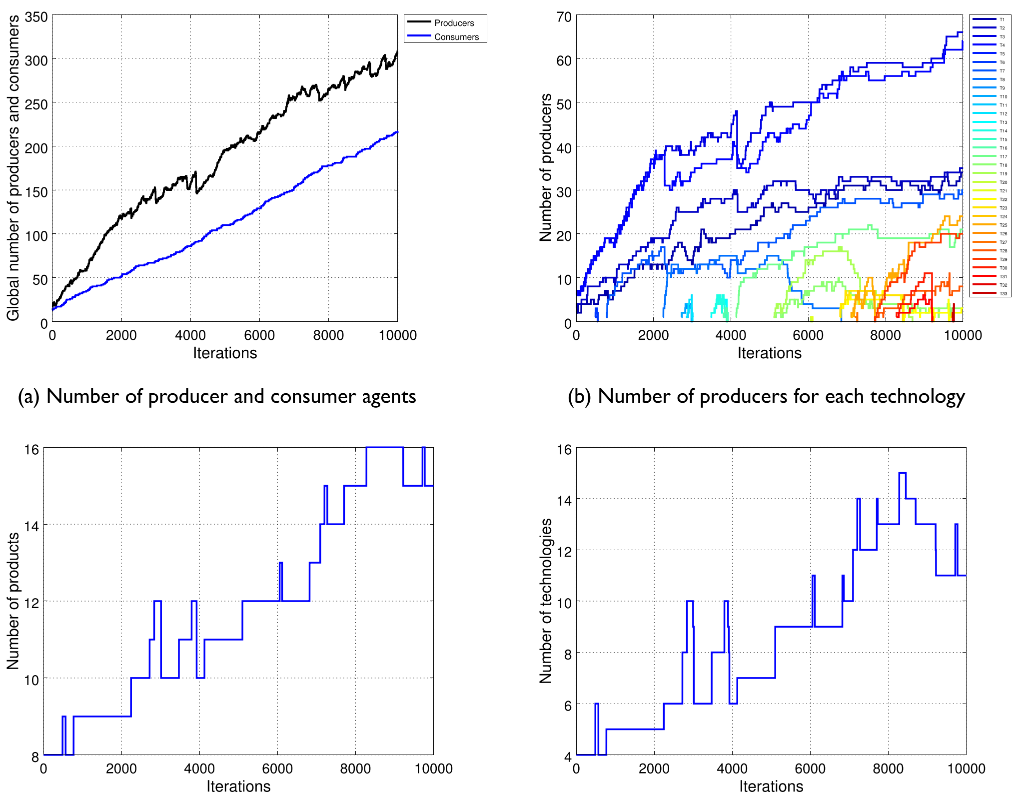

(a) Number of producer and consumer agents　　(b) Number of producers for each technology

(c) Number of products available in the system　　(d) Number of technologies present in the system

Figure 18. Steady evolution. The two highest curves in (b) correspond to the initial consumer good and the manufactured good used to produce it.

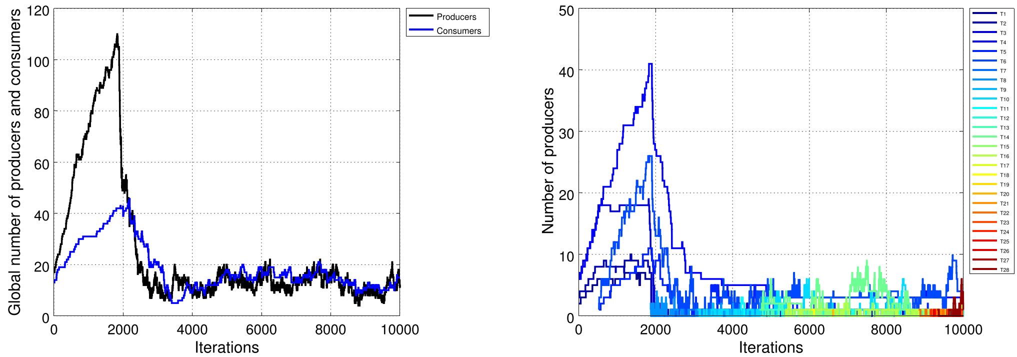

(a) Number of producer and consumer agents

(b) Number of producers for each technology

Figure 19. Failed evolution. In (b), at the final iteration (10,000), there are seven active technologies.

by increasing product diversity—this time, however, without causing a crisis. Figure 22 (see below) shows the structure of the economy at iteration 8101, after this second transformation. Note that the introduction of a new product is frequently accompanied by a spike in its price (Figure 17(f)), as well as an associated spike in profit (not shown).

While about two-thirds of runs exhibiting continuing evolution produced one or more temporary crises like the one just described, the other third reached 10,000 iterations without undergoing a crisis—in other words they exhibited steady evolution. Aside from the crises, however, the results were similar, with increasing numbers of producers, consumers, technologies, and products (see Figure 18). For comparison, we show a system that was unable to recover from an early crash, and was thus subsequently unable to evolve, in Figure 19.

Structural evolution may be characterized as increasing diversity of technologies and products in the system. By this criterion the systems shown in Figures 17 and 18 are clearly evolving, since the number of both technologies and products grows over time. The diversity includes multiple technologies to produce a single product. New technologies to produce a good already present in

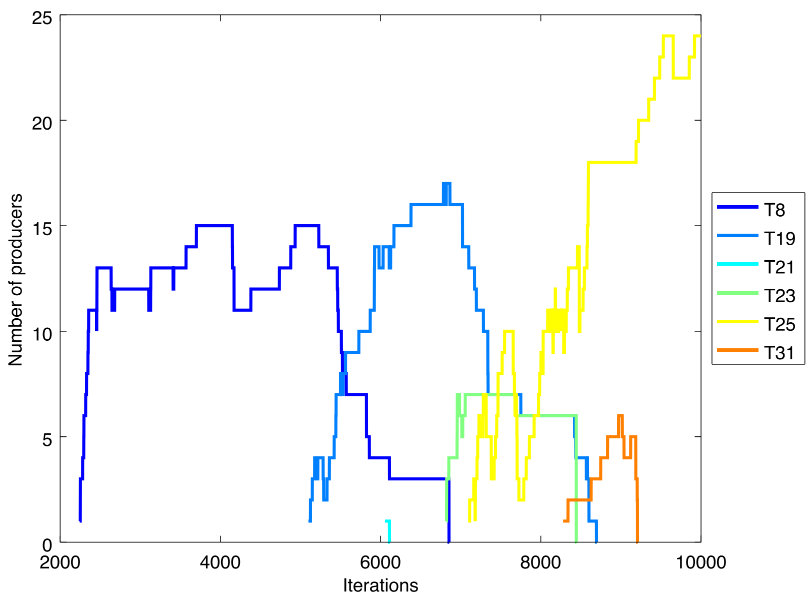

Figure 20. Multiple technologies for producing product *P*11 in the steady evolution system of Figure 18.

the system may be introduced by the stochastic process, but the probability is biased toward goods that are highly profitable or that are facing a supply shortage of one or more inputs. Some of these technologies will be more efficient than others, thus favoring the survival of agents that use them, but efficiency depends on relative prices for the inputs used, so the relative efficiencies of competing technologies may change over time as input prices change. Two or more competing technologies may coexist for long periods, as seen in Figure 20 (from the run shown in Figure 18), where, for example, technologies $T8$ and $T19$ are both present for more than 1,700 iterations. The temporary (and not so temporary) presence of multiple technologies is a means by which the system increases its efficiency, analogous to the role of genetic variation in the evolution of a biological population.

This increasing diversity is made possible by the continuing growth of system size as measured by the number of agents. However, Bedau et al. (1997) have suggested that while increasing diversity is one measure of evolutionary activity, a stronger one is increasing average productivity. In our system the link with the environment is the raw material. Since there is no limit to the amount of raw material that may be produced, it does not constrain the ultimate size of the system. It can, however, be used as the basis of an efficiency measure. We therefore define the average productivity of the system by the number of consumers (i.e., people) who can be supported per unit of raw material—after all, the ultimate purpose of the economic system is to keep people alive. We will call this the *efficiency* of the system.

By this measure we see that the evolving systems are characterized by generally increasing efficiency, as seen in Figure 21(a) and 21(b). In both cases (but much more in the case of steady

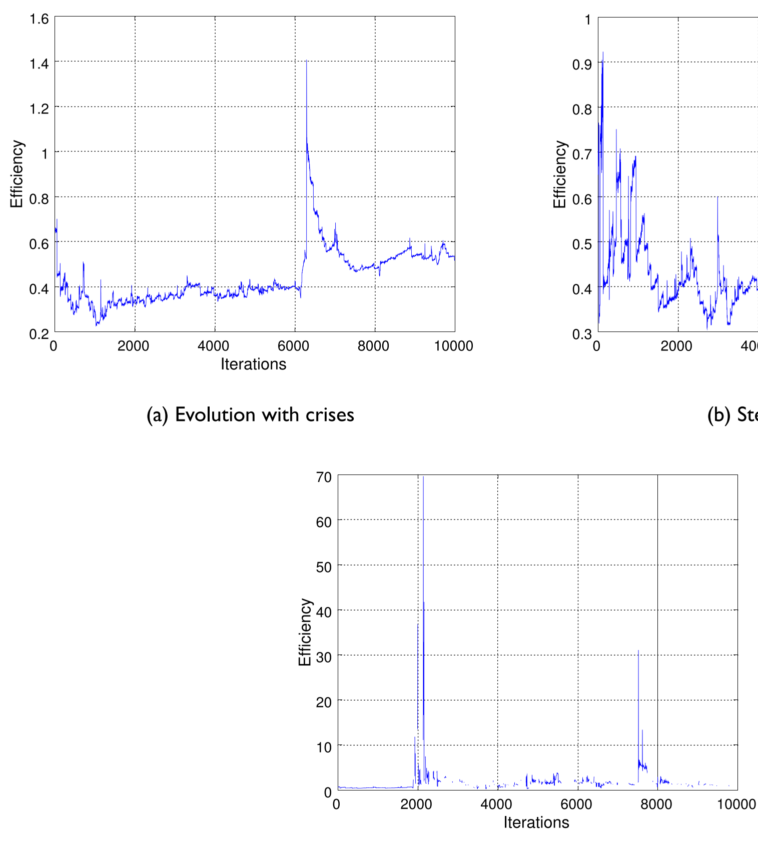

(a) Evolution with crises

(b) Steady evolution

(c) Failed evolution

Figure 21. System efficiency in different classes of runs. Efficiency is defined as the number of consumers per unit of raw material input.

evolution) there is some initial instability in the measure, but then relatively stable growth. The crisis in the evolution with crises case (seen in Figure 17(a), 17(b)), however, produces a temporary spike in this measure (seen in Figure 21(a)) until the consumers who cannot be supported due to the crisis-induced fall of production of consumer goods are removed from the system. In contrast, the failed evolution system shown in Figure 19, which is unable to evolve after its collapse, is characterized by steady efficiency (Figure 21(c)). However, even excluding the peaks, the efficiency is significantly higher than in the cases of successful evolution.

This apparently paradoxical result is explained by the fact that the original consumer good remains dominant, and it is relatively efficient in transforming the raw material. In the cases of successful evolution, many more technologies and goods appear, so productivity is diverted from consumers to goods. With failed evolution, the population remains relatively steady (Figure 19(a)), though at a low level. With successful evolution, the population grows (Figures 17(a) and 18(a)), although the efficiency of the system as measured by the population supported per unit of raw material remains much lower than in the failed evolution case, at least over the simulation period. In other words, systems that evolve successfully divert increasing amounts of resources into the production of a wider range of consumer goods. This is suggestive of modern consumer societies where growth largely takes the form of an increasing variety of goods rather than a population growing at the maximum possible rate. The failed evolution case suggests a subsistence society where production is dominated by a few basic goods that support a relatively large population.

Another characteristic of a successfully evolving system is structural complexity. A convenient measure of this is the average number of inputs per technology. In three runs showing successful evolution (including the runs shown in Figures 17 and 18), the mean number of inputs for the four initial technologies was 2.00. For the subsequently added technologies that were still present in the system at the end of the run the mean number of inputs had grown to 4.52, and the maximum number of inputs for a single technology was 10. For larger systems other indicators of network structure would be useful measures of complexity.

Although the mean number of inputs to a technology shows increasing complexity, by this measure the systems shown here still do not seem very complex. However, a look at the system structure as represented by the von Neumann matrices shows that these systems are not so simple. Figure 22 is a graphical representation of the von Neumann matrices of the evolution with crisis case shown in Figure 17, after the system has recovered from its crisis, and after the secondary wave of innovation at about iteration 8,000 (the diagram excludes flows representing capital goods). Of the four consumer goods (*P*8, *P*15, *P*16, *P*17), *P*16 has the longest supply chain as measured from the raw material, *P*5. It involves eight products: *P*5, *P*6, *P*7, *P*8, *P*9, *P*11, *P*12, and *P*15. Of course it is

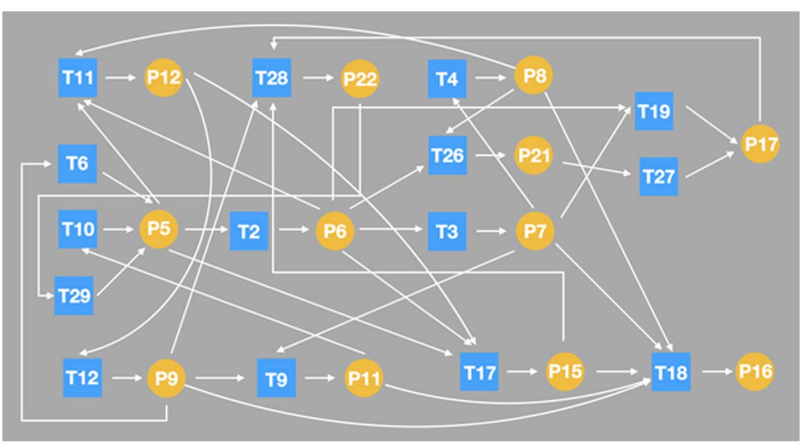

Figure 22. Structure of the von Neumann production process at iteration 8101 of the system shown in Figure 17 (evolution with crisis). Flows representing capital goods are not shown.

not a simple chain, but rather has a branching structure. Another consumer good, *P*17, is produced by two competing technologies and thus has two supply chains. However, note that while the raw material was originally produced by *T*1, which used labor as the sole input, *T*1 was eliminated early in the run by a more efficient technology, and the raw material is now produced only by *T*6, *T*10, and *T*29, so all supply chains in this diagram are in effect cycles.

6 Discussion and Conclusions

The model as described in this article is relatively simple. However, one of its strengths is that many of the simplifications can be relaxed in a straightforward manner, and in ways that are in accord with standard microeconomic theory. For example, consumers who now have individual but fixed consumption patterns could be endowed with price sensitive preference functions so that their consumption patterns could change to reflect changes in relative product prices. On the producer side, agents could be allowed to possess more than one technology. This might represent a case of vertical integration, where an agent possessed several technologies that were chained, so that the output of one was the input of the next. Or the situation might be one of horizontal integration, and in that case, at each iteration, agents would adjust their production plans for the various products to reflect their anticipated relative profitability. Money, which is currently used only for the exchange of products, and hence appears in the technology matrices as a row (*P*2) of zeros, could be given a bigger role, with agents able to buy and sell it at a price (the interest rate). This would permit more flexibility in the production processes, with firms able to borrow money if they did not have enough to implement their production plan, and would also permit the appearance of a financial sector.

On the other hand, in one respect the model is relatively inflexible. Production functions for all technologies are linear, since they are represented by von Neumann matrices. Economies or diseconomies of scale could be introduced, as could input substitutions reflecting price changes, but only at the cost of using a much less efficient representation of production technologies. However, for our present purposes we consider the assumption of linear production functions to be an acceptable limitation.

Structural evolution can occur only in dynamical systems, which is the reason we developed the dynamical model of this article in the first place. It is endowed with the ability to evolve by the spontaneous appearance of new agents, products, and technologies, which are then adopted or rejected by the system. We focused on evaluating the dynamic behavior of the system because the dynamics provide the environment within which evolution takes place: If the dynamical behavior of the model is not reasonable, then any results concerning evolution when the model is run in its evolutionary mode are unlikely to be meaningful.

The experiments described in sections 5.1 to 5.4 show that the model does indeed capture the standard dynamics of an economic system: Prices, profits, supply, and demand are all equilibrated, and following a structural perturbation, the system re-equilibrates to the expected new steady state. The results from running the model in its evolutionary mode (section 5.5) show that the model is capable of generating true evolutionary behavior in that the underlying dynamical system becomes larger (it has more equations with more variables—i.e., larger von Neumann matrices); it becomes different (variables representing products change as some are eliminated and others are introduced); it grows more complex (the equations on average involve more variables); and it becomes more efficient, with more consumers (the ultimate output) supported per unit of raw material (the primary input). These results strongly suggest that the model will be a useful tool for a formal, systematic exploration of structural evolution in economic systems.

The sort of evolution handled by the model is essentially technological evolution, since the matrices describe the transformation of products into other products. As a model of technological evolution it is necessarily very abstract, since, except for a few labeled products (labor, money, raw material) we cannot identify the products. Even if we know what all the initial products are, as new products and technologies are introduced they can only be characterized by their inputs. For the first few products this may be sufficient to identify them generically, e.g., if the two inputs are labor and grain and the capital good is a millstone, then the new product can probably be identified as flour; but as the system evolves the identifications will quickly become very imprecise. In reality, both production processes and products depend on the physical and chemical properties of the materials involved, as well as on the research that progressively reveals those properties. For modeling near-term evolution of a specific sector of the economy, it might be possible to be more specific about these properties, and therefore the results would be to some degree identifiable. However, in general it would be almost impossible to include this kind of information in the model, and thus the specific nature of the evolving economy will remain almost completely abstract, i.e., we will know that a new product n appears, requiring inputs j, k, l to manufacture using technology T, but we will not know what the new product is, or what the inputs actually are, or what the nature of the technology is. Nevertheless, the model remains valuable as a way to build a systematic understanding of general structural features of an evolving economy, including a better understanding of the conditions that produce periods of systematic restructuring, with or without crises, or rising inequality of consumer endowments.

While technological evolution is undoubtedly the most important kind of economic evolution, there are others that are also important: specifically organizational, financial, and spatial evolution. These involve much more specific types of innovations (e.g., banks, insurance companies, hedge funds) identifiable by their structure and structural relationships within the system, and could be modeled with modifications or extensions of the current model. For example, organizational evolution concerns the internal structure of an agent. In the model described in this article, all producer agents have the same structure; they differ only in the technology they possess and therefore the products they process. In real economic systems, however, agents have a variety of structures. For example, a conglomerate is an agent that contains other agents that differ from each other, and a chain operation is an agent that contains other agents that are essentially the same as each other. There are other forms of composite agents, and more generally, other, non-producer, organizations such as trade unions that act as agents in real economic systems. We have recently experimented with a prototype model that reads the model described in this article as data and then, using an evolutionary technique called genetic programming (Banzhaf et al., 1998; Koza, 1992), creates new, more complex types of agents composed of individual agents already present in the original model, and then reintroduces them as composite agents to that model. In a first trial, this approach has generated composite agents with the characteristics of franchises, chains, and co-ops. This exercise shows that it is possible to model various types of economic evolution. More importantly, it demonstrates that using genetic programming opens the way to modeling much deeper structural transformations than can be captured in the relatively rigid framework of the von Neumann matrices. This is important because evolution is undetermined and open-ended at all levels.

The model presented in this article is a simple but comprehensive and fully dynamic microeconomic model that has the endogenous capacity for structural evolution. Currently it is the only such model. Other models are dynamic but not evolutionary, or they are evolutionary but do not represent the structural evolution of a comprehensive economic system. This model constitutes an extensible platform for investigating a wide variety of questions regarding the evolution of economic systems, and in that respect it is unique.

We have said that structural evolution can only occur in dynamical systems. Today, 300 years after Newton and Leibniz laid the foundations for the mathematical treatment of dynamics, the theory and techniques for treating dynamical systems are well developed; the field is vast. In contrast, there is no comprehensive formal treatment of structural evolution, most probably because the undetermined, open-ended nature of evolution cannot be encompassed within the analytic framework of mathematics. The way forward lies in simulation, because computation allows open-endedness and indeterminacy; and techniques like genetic programming, which allow models to restructure themselves as they execute, are likely to prove central to a computational treatment of evolution. But the field is in its infancy and progress will initially seem slow.

Acknowledgments
Discussions with Dr. Bas Straatman and Professor Susan Stepney are gratefully acknowledged.

Conflicts of Interest
The authors declare that there is no conflict of interest regarding the publication of this paper.

References
Abramov, A. P. (2014). *Balanced and cyclical growth in models of a decentralized economy*. Springer. https://doi.org/10.1007/978-3-319-07917-2

Angus, S. D., & Newnham, A. (2013). The bit-economy: An artificial model of open ended technology discovery. *Complexity*, *18*(5), 57–67. https://doi.org/10.1002/cplx.21449

Arthur, W. B. (2009). *The nature of technology: What it is and how it evolves*. Simon & Schuster.

Arthur, W. B., & Polak, W. (2006). The evolution of technology within a simple computer model. *Complexity*, *11*(5), 23–31. https://doi.org/10.1002/cplx.20130

Banzhaf, W., Nordin, P., Keller, R. E., & Francone, F. D. (1998). *Genetic programming: An Introduction*. Morgan Kaufmann.

Bedau, M. A., Snyder, E., Brown, C. T., & Packard, N. H. (1997). A comparison of evolutionary activity in artificial evolving systems and in the biosphere. In P. Husbands & I. Harvey (Eds.), *Proceedings of the fourth European conference on artificial life* (pp. 125–134). MIT Press.

Blatt, J. M. (1983). *Dynamic economic systems*. Routledge.

Borland, J., & Yang, X. A. (1992). Specialization and a new approach to economic organization and growth. *The American Economic Review*, *82*(2), 386–391.

Coenen, L., & Díaz López, F. J. D. (2010). Comparing systems approaches to innovation and technological change for sustainable and competitive economies: An explorative study into conceptual commonalities, differences and complementarities. *Journal of Cleaner Production*, *18*(12), 1149–1160. https://doi.org/10.1016/j.jclepro.2010.04.003

Colander, D. (Ed.) (2006). *Post Walrasian macroeconomics: Beyond the dynamic stochastic general equilibrium model*. Cambridge University Press.

Dawid, H., Gemkow, S., Harting, P., van der Hoog, S., & Neugart, M. (2012). The eurace@unibi model: An agent-based macroeconomic model for economic policy analysis. *Bielefeld Working Papers in Economics and Management*, 05-2012, 1–61. Bielefeld University, Department of Business Administration and Economics. https://doi.org/10.2139/ssrn.2408969

de Grauwe, P. (2010). The scientific foundation of dynamic stochastic general equilibrium (DSGE) models. *Public Choice*, *144*(3/4), 413–443. https://www.jstor.org/stable/40835621, https://doi.org/10.1007/s11127-010-9674-x

Deissenberg, C., van der Hoog, S., & Dawid, H. (2008). EURACE: A massively parallel agent-based model of the European economy. *Applied Mathematics and Computation*, *204*(2), 541–552. https://doi.org/10.1016/j.amc.2008.05.116

Dopfer, K., & Potts, J. (2008). *The general theory of economic evolution*. Routledge. https://doi.org/10.4324/9780203507407

Dosi, G., Fagiolo, G., Napoletano, M., & Roventini, A. (2013). Income distribution, credit and fiscal policies in an agent-based Keynesian model. *Journal of Economic Dynamics and Control*, *37*(8), 1598–1625. https://doi.org/10.1016/j.jedc.2012.11.008

Dosi, G., Fagiolo, G., & Roventini, A. (2010). Schumpeter meeting Keynes: A policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control*, *34*(9), 1748–1767. https://doi.org/10.1016/j.jedc.2010.06.018

Dosi, G., & Marengo, L. (2014). The dynamics of organizational structures and performances under diverging distributions of knowledge and different power structures. *Journal of Institutional Economics*, *11*(3), 535–559. https://doi.org/10.1017/S1744137414000204

Dosi, G., & Roventini, A. (2019). More is different ... and complex! The case for agent-based macroeconomics. *Journal of Evolutionary Economics*, *29*(1), 1–37. https://doi.org/10.1007/s00191-019-00609-y

Encinar, M.-I., & Muñoz, F.-F. (2006). On novelty and economics: Schumpeter's paradox. *Journal of Evolutionary Economics*, *16*, 255–277. https://doi.org/10.1007/s00191-006-0018-4

Farmer, J. D. (2002). Market force, ecology and evolution. *Industrial and Corporate Change*, *11*(5), 895–953. https://doi.org/10.1093/icc/11.5.895

Farmer, J. D., & Foley, D. (2009). The economy needs agent-based modelling. *Nature*, *460*, 685–686. https://doi.org/10.1038/460685a

Farmer, J. D., Hepbum, C., Mealy, P., & Teytelboym, A. (2015). A third wave in the economics of climate change. *Environmental and Resource Economics*, *62*(2), 329–357. https://doi.org/10.1007/s10640-015-9965-2

Festré, A., Lakomski-Laguerre, O., & Longuet, S. (2017). Schumpeter and Schumpeterians on economic policy issues: Re-reading Schumpeter through the lens of institutional and behavioral economics. An introduction to the special issue. *Journal of Evolutionary Economics*, *27*(1), 3–24. https://doi.org/10.1007/s00191-015-0442-4

Foster, J. (2015). Joseph Schumpeter and Simon Kuznets: Comparing their evolutionary economic approaches to business cycles and economic growth. *Journal of Evolutionary Economics*, *25*(1), 165–172. https://doi.org/10.1007/s00191-014-0356-6

Gintis, H. (2007). The dynamics of general equilibrium. *The Economic Journal*, *117*(523), 1280–1309. https://doi.org/10.1111/j.1468-0297.2007.02083.x

Gualdi, S., & Mandel, A. (2019). Endogenous growth in production networks. *Journal of Evolutionary Economics*, *29*(1), 91–117. https://doi.org/10.1007/s00191-018-0552-x

Hodgson, G. M. (1993). *Economics and evolution: Bringing life back into economics*. The University of Michigan Press. https://doi.org/10.3998/mpub.14010

Kephart, J. O., Hanson, J. E., & Greenwald, A. R. (2000). Dynamic pricing by software agents. *Computer Networks*, *32*(6), 731–752. https://doi.org/10.1016/S1389-1286(00)00026-8

Korres, G. M., & Drakopoulos, S. (2009). Economics of innovation: A review in theory and models. *European Research Studies*, *XII*(3), 25–38. https://doi.org/10.35808/ersj/229

Koza, J. R. (1992). *Genetic programming: On the programming of computers by means of natural selection*. MIT Press.

Krugman, P. (1995). *Development, geography, and economic theory*. MIT Press. https://doi.org/10.7551/mitpress/2389.001.0001

Lindgren, K. (1991). Evolutionary phenomena in simple dynamics. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmussen (Eds.), *Artificial Life II, SFI studies in the sciences of complexity* (pp. 295–312). Addison-Wesley.

Nelson, R., Dosi, G., Helfat, C., Pyka, A., Saviotti, P., Lee, K., Winter, S., Dopfer, K., & Malerba, F. (2018). *Modern evolutionary economics: An overview*. Cambridge University Press. https://doi.org/10.1017/9781108661928

Nelson, R. R., & Winter, S. G. (1982). *An evolutionary theory of economic change*. Belknap Press.

Ponzi, A., Yasutomi, A., & Kaneko, K. (2003). Multiple timescale dynamics in economic production networks. ArXiv. https://doi.org/10.48550/arXiv.nlin/0309002

Ponzi, A., Yasutomi, A., & Kaneko, K. (2006). Complex dynamical behaviour in economic production networks. *Journal of Economic Behavior & Organization*, *61*(4), 729–749. https://doi.org/10.1016/j.jebo.2004.07.013

Potts, J. (2000). *The new evolutionary microeconomics*. Edward Elgar.

Recio, G., Banzhaf, W., & White, R. (2020). A study of severe disruption in an artificial economy. In *ALIFE 2020: The 2020 Conference on Artificial Life* (pp. 180–187). MIT Press. https://doi.org/10.1162/isal_a_00336

Sbordone, A., Tambalotti, A., Rao, K., & Walsh, K. (2010). Policy analysis using DSGE models: An introduction. Federal Reserve Bank of New York Economic Policy. https://www.newyorkfed.org/research/epr/10v16n2/1010sbor.html. https://doi.org/10.2139/ssrn.1692896

Schumpeter, J. A. (1961). *The theory of economic development: An inquiry into profits, capital, credit and the business cycle*, Harvard University Press.

Solé, R. V., Valverde, S., Casals, M. R., Kauffman, S., Farmer, D., & Eldredge, N. (2013). The evolutionary ecology of technological innovations. *Complexity*, *18*(4), 15–27. https://doi.org/10.1002/cplx.21436

Straatman, B. (2008). *Innovation and evolution in an artificial spatial economy*. Unpublished doctoral dissertation. Memorial University of Newfoundland.

Straatman, B., Marceau, D. J., & White, R. (2013). A generic framework for a combined agent-based market and production model. *Computational Economics*, *41*(4), 425–445. https://doi.org/10.1007/s10614-012-9341-z

Straatman, B., White, R., & Banzhaf, W. (2008). An artificial chemistry-based model of economies. In S. Bullock, J. Noble, R. Watson, & M. Bedau (Eds.), *Artificial Life XI: Proceedings of the eleventh International Conference on the Simulation and Synthesis of Living Systems* (pp. 592–599). MIT Press.

Teglio, A., Raberto, M., & Cincotti, S. (2010). Balance sheet approach to agent-based computational economics: The EURACE project. In C. Borgelt, G. Gonzélez-Rodríguez, W. Trutschnig, M. Asunción Lubiano, M. Ángeles Gil, P. Grzegorzewski, & O. Hryniewicz (Eds.), *Combining soft computing and statistical methods in data analysis* (pp. 603–610). Springer. https://doi.org/10.1007/978-3-642-14746-3_74

Tesfatsion, L. (1998). Preferential partner selection in evolutionary labour markets: A study in agent-based computational economics. In V. W. Porto & N. Saravanan (Eds.), *Evolutionary programming VII: Proceedings of the seventh annual conference on evolutionary programming* (pp. 15–24). Springer-Verlag.

Tesfatsion, L. (2001). Structure, behavior, and market power in an evolutionary labor market with adaptive search. *Journal of Economic Dynamics and Control*, *25*(3/4), 419–457. https://doi.org/10.1016/S0165-1889(00)00032-4

Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, *8*(1), 55–82. https://doi.org/10.1162/106454602753694765

van Daal, J., & Jolink, A. (1993). *The equilibrium economics of Léon Walras*, Routledge. https://doi.org/10.4324/9780203401460

von Neumann, J. (1945). A model of general economic equilibrium. *Review of Economic Studies*, *13*(1), 1–9. https://doi.org/10.2307/2296111

White, R., & Banzhaf, W. (2020). Putting natural time into science. In A. Adamatzky & V. Kendon (Eds.), *From astrophysics to unconventional computing: Essays presented to Susan Stepney on the occasion of her 60th birthday* (pp. 1–21). Springer. https://doi.org/10.1007/978-3-030-15792-0_1

Winsberg, E. (2001). Simulations, models, and theories: Complex physical systems and their representations. *Philosophy of Science*, *68*(3), S442–S454. https://doi.org/10.1086/392927

Yang, X., & Borland, J. (1991). A microeconomic mechanism for economic growth. *Journal of Political Economy*, *99*(3), 460–482. https://doi.org/10.1086/261762