

The microfoundations of business cycles: an evolutionary, multi-agent model

Giovanni Dosi · Giorgio Fagiolo ·
Andrea Roventini

© Springer-Verlag 2008

Abstract This work presents an evolutionary model of output and investment dynamics yielding endogenous business cycles. The model describes an economy composed of firms and consumers/workers. Firms belong to two industries. The first one performs R&D and produces heterogeneous machine tools. Firms in the second industry invest in new machines and produce a homogenous consumption good. Consumers sell their labor and fully consume their income. In line with the empirical literature on investment patterns, we assume that firms' investment decisions are lumpy and constrained by their financial structure. Simulation results show that the model is able to deliver self-sustaining patterns of growth characterized by the presence of endogenous business cycles. The model can also replicate the most important stylized facts concerning micro- and macro-economic dynamics.

Keywords Evolutionary dynamics · Agent-based computational economics · Lumpy investment · Output fluctuations · Endogenous business cycles

JEL Classification C15 · C22 · C49 · E17 · E22 · E32

Thanks to Giulio Bottazzi, Andrea Ginzburg, Marco Lippi, Marco Mazzoli, Mauro Napoletano and Sandro Sapio. All usual disclaimers apply.

G. Dosi · A. Roventini
Sant'Anna School of Advanced Studies, Pisa, Italy

G. Fagiolo (✉)
Laboratory of Economics and Management, Sant'Anna School of Advanced Studies,
Piazza Martiri della Libertà, 56127 Pisa, Italy
e-mail: giorgio.fagiolo@sssup.it

A. Roventini
University of Modena and Reggio Emilia, Modena, Italy

1 Introduction

One of the most important features that characterize the overall activity of modern economies is the existence of widespread and persistent fluctuations. However, despite the long-standing importance of the issue for economic analysis, one still lacks a generally accepted explanation for business fluctuations. In the last resort, contemporary business cycles theories can be traced back to two alternative generating mechanisms: the Real Business Cycle (RBC) perspective (see e.g. King and Rebelo 1999) and the New-Keynesian (NK) paradigm (cf. e.g. Mankiw and Romer 1991).

In our view, both RBC and NK theories share two serious weaknesses. First, both streams of literature dramatically underestimate the role of endogenous technological shocks occurring at the microeconomic level. Indeed, one finds it very hard to believe the existence of *macroscopic*, possibly negative, technological shocks necessary for the RBC story to hold. Conversely, NK models put an almost exclusive emphasis upon monetary and price shocks as drivers of the fluctuations, thus neglecting all technological factors. Second, both RBC and NK explanations show a persistent clash between the microeconomics that one assumes in the models and the regularities in microeconomic behaviors that one empirically observes.¹

At the end of the day, almost no attempt has been made in the literature to explain the properties of business cycles on the grounds of multiple individual agents embodying the observed microeconomic regularities about firms' investment and pricing behaviors.

In this paper, we try to bridge such gaps by proposing an evolutionary model where both output and investment dynamics are grounded upon lumpy investment decisions undertaken by boundedly-rational firms. The latter are constrained by their financial structure, but, at the same time, are always able to discover new production technologies (well in tune with the seminal Nelson and Winter (1982) contribution). Building on earlier, admittedly more rudimentary, attempts to explore the properties of evolutionary models with "Keynesian" demand propagation effects (cf. Chiaromonte and Dosi 1993; Dosi et al. 1994), we present a modified version of the model discussed in Dosi et al. (2005, 2006) to study the statistical properties of micro and macro dynamics. More precisely, in line with the recommendations advanced by the recent literature on the empirical validation of ACE models (Fagiolo et al. 2007; Brenner and Werker 2007), the model that we present below tries to capture explicitly in its behavioral assumptions some of the micro regularities empirically detected by both micro-econometric analysis and behavioral economics.

The model depicts an economy composed of boundedly-rational firms (operating in two vertically-linked industries) and consumers/workers. Firms in the "upstream" industry perform R&D and produce technologically heterogeneous machines. The latter are used in the "downstream" industry to produce

¹See, e.g., the sharp critique to the "representative agent fallacies" in Kirman (1989).

a consumption good bought by workers with their wages. Microeconomic investment lumpiness is at the root of the business cycle fluctuations characterizing the economy. Investment can be either employed to increase the capital stock or to replace existing capital goods. Consumption-good firms plan their expansion investment according to an (S,s) pattern.² Firms will then decide to expand their stock of capital only if they (adaptively) expect a significant demand growth. Similarly to what happens for expansion investment, firms employ routinized behaviors to decide their replacement investment through a payback-period routine. Finally, the financial profile of each firm affects its investment policies, due to the presence of financial constraints.

Simulation results show that the model is able to deliver both self-sustaining growth patterns and endogenous business cycles. Moreover, we show that the model is able to replicate many business-cycle stylized facts concerning, e.g., volatility, auto- and cross-correlation patterns. Finally, the micro-structure of the simulated economy is quite in tune with the evidence on, e.g., persistent heterogeneity in firm efficiencies, size and growth rate distributions.

The work is organized as follows. Section 2 provides a short overview of micro and macro empirical evidence. In Section 3, the model is formally presented. Section 4 discusses qualitative and quantitative results of simulation exercises. Finally, Section 5 concludes.

2 Macro and micro stylized facts

A good check of the robustness of any model claiming to be able to “explain” business cycles resides in its ability to: (i) account simultaneously for more than one macroeconomic “stylized fact” (SF) concerning the dynamics of output, investment, employment, etc.; and (ii) provide a story that is coherent with the observed microeconomics of business decisions and innovation patterns. Let us thus recall some important empirical regularities characterizing micro and macro dynamics (see Dosi et al. 2006, for more details).

A key issue in the empirical business cycle literature concerns the properties of aggregate output and of its main components (i.e. investment, consumption and inventories). At the business cycle frequencies, the series display a typical “roller coaster” shape, implying the repeated interchange of expansions and recessions, which are part of the very definition of the business cycle. A closer inspection of these patterns (see e.g. Stock and Watson 1999; Napoletano et al. 2006) reveals that investment is considerably more volatile than output (SF1); consumption is less volatile than output (SF2); investment, consumption and change in inventories tend to be procyclical and coincident variables (SF3); aggregate employment tends to be procyclical and lagging, whereas the unemployment rate is countercyclical and lagging (SF4).

Furthermore, at a microeconomic level, the empirical literature on industrial dynamics and technological change has singled out, over the last couple

²See Caballero (1999).

of decades, an impressive number of robust statistical regularities concerning the distributions of firm properties and their dynamics. Let us begin here with those stylized facts pertaining to firms' investment decisions. First, empirical works based on plant level data (Caballero 1999) show that lumpiness is an intrinsic feature of firm investment (SF5). Second, the financial structure of firms is not neutral with respect to their investment choices (SF6), see Hubbard (1998).

As far as the drivers of growth are concerned, a growing number of contributions has robustly highlighted the central role of technological learning, innovation and diffusion carried out by business firms (see Dosi et al. (1994) for a critical overview). First of all, firms appear to be the main locus where technological accumulation takes place (SF7). Technological learning and accumulation tend to be mostly local: technical advances typically occur in a neighborhood of currently-mastered technologies, but is seldom "punctuated" by major, low-probability advances. Furthermore, innovations take time to diffuse (SF8). Technological diffusion is slowed down by information asymmetries and, even more important, by the fact that firms require time to learn how to master new technologies and to develop new skills. Finally, most innovations are industry-specific (SF9). Therefore, the overall pattern of business fluctuations cannot be fully explained by economy-wide innovative shocks.

In turn, the foregoing regularities concerning asymmetries in innovation profiles and technological diffusion map onto the intersectoral patterns of realized performances and productivities. Extensive studies on longitudinal micro-level data sets (cf. the surveys in Bartelsman and Doms 2000; Dosi 2005) confirm that productivity dynamics is characterized by a few robust regularities, including: (a) productivity dispersion among firms are considerably large (SF10); and (b) inter-firm productivity differentials are quite persistent over time (SF11).

Together, heterogeneity is also a persistent feature of firm size distributions, both among firms belonging to the same industrial sector and across different industrial sectors (see, Dosi (2005) for a survey). More specifically, firm size distributions tend to be considerably right skewed, with upper-tails made of few large firms (SF12). Moreover, a growing evidence highlights microeconomic processes of growth, entailing some underlying correlation structure and lumpiness. More precisely: firm growth-rate distributions are *not* Gaussian and can be well proxied by fat-tailed, tent-shaped densities (SF13).³

³See Bottazzi and Secchi (2006). Castaldi and Dosi (2004) and Fagiolo et al. (2008) find that country output growth-rate distributions are also fat-tailed, over both cross-section and time-series dimensions.

In the model presented below, we take explicitly on board micro-regularities pertaining to firm investment and innovating behaviors (SF5 to SF9) in the way we design the agents populating our economy. We aim at building a model that, at the same time, is able to replicate and to explain the stylized facts concerning the business cycle (SF1 to SF4) on the basis of micro-dynamics patterns which replicate the statistical regularities displayed by the evolution of firm productivity, size and growth over time (SF10 to SF13).

3 The model

We model an economy populated by F firms and L workers/consumers. Firms belong to two industries: there are F_1 consumption-good firms (labeled by j in what follows) and F_2 machine-tools firms (labeled by i). Of course, $F = F_1 + F_2$. Consumption-good firms invest in machine-tools and produce a homogeneous product for consumers. Machine-tool firms produce heterogeneous capital goods and perform R&D. Workers inelastically sell labor to firms in both sectors and fully consume the income they receive. Investment choices of consumption-good firms determine the level of income, consumption and employment in the economy.

In the next subsection, we shall first describe the dynamics of events in a representative time-period. Next, we shall provide a more detailed account of each event separately.

3.1 The dynamics of microeconomic decisions

In any discrete time period $t = 1, 2, \dots$, the timeline of events runs as follows:

1. Capital-good firms advertise their machines, sending a “brochure” to a subset of consumption-good firms.
2. Consumption-good firms make their production and investment decisions, choose their supplier and order the machines.
3. Both capital- and consumption-good firms hire workers according to their production plans and start producing.
4. Consumption-good market opens.
5. Capital-good firms deliver the machine-tools ordered by consumption-good firms.
6. Exit, technical change and entry take place. Capital-good firms stochastically search for new machines and more efficient production routines.

Finally, total consumption, investment, change in inventories, and total product are computed by aggregating individual time- t quantities.

3.2 Production and investment: the consumption-good sector

Each consumption-good firm $j = 1, 2, \dots, F_1$ produces a homogenous good using machines and labor under constant returns to scale. Planned output depends on myopic demand expectations of the form:

$$D_j^e(t) = D_j(t-1), \quad (1)$$

where $D_j(t-1)$ is the demand of firm j at time $t-1$.⁴

According to the expected demand, the desired level of inventories (N_j^d) and the inventories (N_j) inherited from the previous period, firms fix their desired level of production (Q_j^d):

$$Q_j^d(t) = D_j^e(t) + N_j^d(t) - N_j(t-1), \quad (2)$$

with $N_j^d(t) = \theta D_j^e(t)$, $0 \leq \theta \leq 1$.

The current stock of capital determines the maximum level of production achievable by each firm. Given the desired level of production, firms compute the desired stock of capital as:

$$K_j^d(t) = \frac{Q_j^d(t)}{u^d}, \quad (3)$$

where u^d is the desired level of capacity utilization.

Consumption-good firms decide whether to expand⁵ their stock of capital following an (S,s) model. They compute their trigger (K_j^{trig}) level of capital as follows:

$$K_j^{trig} = K_j(t)(1 + \alpha), \quad (4)$$

with $0 < \alpha < 1$. Firms then plan to increase their capital stock only if the desired capital stock is higher than the trigger one:

$$EI_j(t) = \begin{cases} 0 & \text{if } K_j^d(t) < K_j^{trig}(t) \\ K_j^{trig}(t) - K_j(t) & \text{if } K_j^d(t) \geq K_j^{trig}(t) \end{cases}, \quad (5)$$

where $EI_j(t)$ is the expansion investment.

The stock of capital of each consumption-good firm is heterogeneous, since it is composed of various vintages of machines which differ in terms of (labor) productivity $A_{i,\tau}$, where i denotes their producer and τ their generation (see

⁴Different extrapolative expectation-formation rules based on both firm-specific past demand and aggregate market signal are explored in Dosi et al. (2006). Interestingly, one finds that increasing the computational sophistication of agents does not improve either the performance of the economy, as measured by average growth-rates, or the stability of growth patterns over time.

⁵We assume that there are no secondary markets for capital goods. Hence, firms have no incentives to reduce their capital stock.

Section 3.5 below for details). Let $\Xi_j(t)$ be the set of all types of machines belonging to firm j at time t . Firm j 's capital stock is defined as:

$$K_j(t) = \sum_{A_{i,\tau} \in \Xi_j(t)} g_j(A_{i,\tau}, t),$$

where $g_j(A_{i,\tau}, t)$ is the absolute frequency of machine $A_{i,\tau}$. Given the nominal wage $w(t)$, the unit labor cost of each machine is computed as:

$$c(A_{i,\tau}, t) = \frac{w(t)}{A_{i,\tau}}.$$

Scrapping policies follow a payback-period routine. The replacement of an incumbent machine depends on its degree of “technological” obsolescence and on the price of new capital goods. More formally, firm j will scrap machines $A_{i,\tau} \in \Xi_j(t)$ if they satisfy:

$$RS_j(t) = \left\{ A_{i,\tau} \in \Xi_j(t) : \frac{p^*(t)}{c(A_{i,\tau}, t) - c^*(t)} \leq b \right\}, \tag{6}$$

where p^* and c^* are, respectively, the price and unit labor cost of new machines, and b is a strictly positive payback-period parameter. Moreover, firms scrap machines that are older than η periods (η positive integer). Firms compute their replacement investment by pooling the machines satisfying Eq. 6. The level of firm investment (I_j) is the sum of expansion and replacement investment. Summing up the actual investment of all consumption-good firms, we get aggregate investment (I).

Consumption-good firms choose their capital-good supplier according to the price and productivity of the currently produced machines. Since the capital-good market is characterized by imperfect information, consumption-good firms receive “brochures” depicting the characteristics of machines only from a subset of capital-good firms (cf. Section 3.3). Each consumption-good firm compares the available machines, chooses the one with the highest productivity/price ratio and sends its investment orders to the correspondingly capital-good firm.

Consumption-good firms must bear production costs before selling their output. Hence, they must finance production as well as investment. In tune with the spirit of the evolutionary perspective – and of many New Keynesian models – we assume imperfect capital markets with credit rationing. Hence, firms will initially employ their stock of liquid assets (NW_j) in order to finance production. If liquid assets are not sufficient, they will borrow the necessary amount at the interest rate r . The borrowed amount cannot let the debt/sales ratio exceed the value of Λ . If production is not rationed, firms finance investment using their residual stock of liquid assets and, if necessary, their residual debt availability. If firms are not able fully to finance investment, they privilege expansion investment over replacement investment.

Given their current stock of machines, consumption-good firms compute their average productivity (π_j) and their unit cost of production (c_j). Average productivity reads:

$$\pi_j(t) = \sum_{A_{i,\tau} \in \Xi_j(t)} A_{i,\tau} \frac{g_j(A_{i,\tau}, t)}{K_j(t)},$$

while unit cost of production will be given by:

$$c_j(t) = \frac{w(t)}{\pi_j(t)}.$$

Firms fix the price as a mark-up (μ_j) on their unit cost of production:

$$p_j(t) = (1 + \mu_j(t))c_j(t). \quad (7)$$

The mark-up is flexible: it changes across time according to the past variation of firm market share (f_j):

$$\mu_j(t) = \mu_j(t-1) \left(1 + \frac{f_j(t-1) - f_j(t-2)}{f_j(t-2)} \right).$$

Given their average productivity and their production, consumption-good firms determine their labor demand (L_j^D):

$$L_j^D(t) = \frac{Q_j(t)}{\pi_j(t)}.$$

Denoting by S_j total sales of firm j , profits (Π_j) read:

$$\Pi_j(t) = p_j(t)S_j(t) - c_j(t)Q_j(t) - rDeb_j(t),$$

where Deb_j is the stock of debts. The variation of the stock of liquid assets of consumption-good firms depends on their profits as well as on their investment choices:

$$NW_j(t) = NW_j(t-1) + \Pi_j(t) - cI_j,$$

where cI_j is the amount of internal funds employed by firm j to finance investment.

3.3 Machine production

In the previous section, we described how the demand of capital goods is generated. Let us now turn to the machine producing sector.

Each machine-tool firm $i = 1, 2, \dots, F_2$ sells its latest generation of products characterized by labor productivity coefficient $A_{i,\tau}$, with $\tau = 1, 2, \dots$. At the beginning of the period, capital-good firms update their historical client (HC_i) sending a “brochure” containing information about the price and productivity of their currently produced machine. Moreover, they also send it to a random sample of consumption-good firms (NC_i):

$$NC_i(t) = (1 + \kappa)HC_i(t), \quad (8)$$

with $0 < \kappa < 1$. Of course, $NC_i(t) \leq F_1 - HC_i(t)$. The production process employs labor only, under constant returns to scale. The unit cost of production depends on firm labor productivity (B_i):

$$c_i(t) = \frac{w(t)}{B_i(t)}.$$

The price (p_i) is equal to the unit cost of production. The demand of capital-good firms depends on the investment choices of consumption-good firms. According to the orders they receive, capital-good firms fix the level of production (Q_i) and hire workers accordingly:

$$L_i^D(t) = \frac{Q_i(t)}{B_i(t)},$$

where L_i^D is the labor demand of firm i . Machine production requires time: machines are delivered to consumption-good firms at the end of the period.

3.4 The consumption-good market

Since consumption-good firms make their production decisions according to their demand expectations, they can obviously make mistakes which are revealed by variations in inventories. If in the previous period, they produced too much ($Q_j(t-1) > D_j(t-1)$), they accumulate stocks. On the contrary, if they were not able fully to satisfy their past demand ($Q_j(t-1) < D_j(t-1)$), their “competitiveness” (E_j) at time t is reduced:

$$E_j(t) = -\omega_1 p_j(t) - \omega_2 l_j(t), \tag{9}$$

where l_j is the level of unfilled demand inherited from the previous period and $\omega_{1,2}$ are non-negative parameters. The average sectorial competitiveness (\bar{E}) is obtained by weighting the competitiveness of each firm with its past market share ($f_j(t-1)$):

$$\bar{E}(t) = \sum_{j=1}^{F_1} E_j(t) f_j(t-1).$$

Under imperfect information, consumers take time to adjust imperfectly to relative consumption-good prices. Thus, market shares evolve according to a replicator dynamics. More specifically, the market share of each firm will grow (shrink) if its competitiveness is above (below) the industry-average competitiveness:

$$f_j(t) = f_j(t-1) \left(1 + \chi \frac{E_j(t) - \bar{E}(t)}{\bar{E}(t)} \right), \tag{10}$$

with $\chi \geq 0$.

Aggregate consumption (cf. Section 3.6) shapes the demand-side of the market and it is allocated to consumption-good firms according to their market share:

$$D_j(t) = C(t) f_j(t). \tag{11}$$

3.5 Entry, exit, and technical change

At the end of every period, firms with zero market shares and/or negative net assets die and are replaced by new firms. Hence, the number of firms in both sectors remains constant across time. We also assume that each entrant is a random copy of a survived firm.

The economy is fuelled by a never-ending process of technical change. At the end of each period, capital-good firms try both to develop the next generation of their machines (product innovation) and to discover more efficient production routines (process innovation). The result of their efforts is strongly uncertain.

As far as product innovation is concerned, firms develop a prototype the labor productivity ($A_{i,new}$) of which may be higher or lower than the one of the currently manufactured machine. More formally, we let:

$$A_{i,new} = A_{i,\tau}(1 + \epsilon_1), \quad (12)$$

where $\epsilon_1 \sim U[\iota_1^-, \iota_1^+]$, with $-1 < \iota_1^- < 0 < \iota_1^+$. We also posit that firm i will release the next generation machine only if it is more productive (i.e. $A_{i,new} > A_{i,\tau}$). If the firm decides to produce the new machine, the index τ is accordingly incremented by one unit.

Similarly, firms stochastically search for new production routines. Firms compare the incumbent and the new production routines affecting their own labor productivity:

$$B_{i,new} = B_i(t)(1 + \epsilon_2), \quad (13)$$

where $\epsilon_2 \sim U[\iota_2^-, \iota_2^+]$, with $-1 < \iota_2^- < 0 < \iota_2^+$. If $B_{i,new} > B_i(t)$, the firm adopts the new routine, otherwise it keeps on producing with the old one.

3.6 Macro dynamics and consumption scenarios

The dynamics generated at the micro-level by individual decisions and interaction mechanisms induces, at the macroeconomic level, a stochastic dynamics for all aggregate variables of interest (e.g. output, investment, consumption, unemployment, etc.).

The labor market is not cleared by real wage movements. As a consequence, involuntary unemployment may arise. The aggregate supply of labor (L) is exogenous and inelastic. The aggregate demand of labor is the sum of machine- and consumption-good firms' labor demands:

$$L^D(t) = \sum_{j=1}^{F_1} L_j^D(t) + \sum_{i=1}^{F_2} L_i^D(t).$$

Hence, aggregate employment (Emp) reads:

$$Emp(t) = \min(L^D(t), L). \quad (14)$$

Table 1 Initial conditions

Description	Symbol	Value
Market wage	$w(0)$	1
Consumer price index	$cpi(0)$	1.3
Average labor productivity	$\bar{A}(0)$	1
Mark-up	$\mu(0)$	0.3
Liquid assets	$NW_{i,j}(0)$	10000
Capital stock	$K_j(0)$	800
Labor supply	$L(0)$	1000000

The wage rate is determined by both institutional and market factors. More precisely, it depends on indexation mechanisms upon consumption prices, average productivity, and adjustments to unemployment rates as follows:

$$w(t) = w(t - 1) + \left(1 + \psi_1 \frac{\Delta cpi(t)}{cpi(t - 1)} + \psi_2 \frac{\Delta \bar{A}(t)}{\bar{A}(t - 1)} + \psi_3 \frac{\Delta U(t)}{U(t - 1)} \right), \quad (15)$$

where cpi is the consumer price index, \bar{A} is average labor productivity and U is the unemployment rate. The system parameters $\psi_{1,2,3}$ allow one to characterize various institutional regimes for the labor market.

We consider two scenarios according to the composition of aggregate consumption. In the **work-or-die scenario**, only employed workers earn an income that they fully consume:

$$C(t) = w(t)Emp(t). \quad (16)$$

In the **social-security scenario**, unemployed workers do not starve, but receive a fraction of the market wage from an unmodeled “public” sector:

$$C(t) = w(t)[Emp(t) + \varphi(L - Emp(t))], \quad (17)$$

with $0 < \varphi < 1$.

As mentioned above, our model straightforwardly belongs to the evolutionary family. Since in general, analytical, closed-form, solutions can hardly be obtained, one must resort to computer simulations to analyze the properties of the (stochastic) processes governing the coevolution of micro and macro variables.⁶

To do so, one should in principle address an extensive Monte Carlo analysis in order to understand how the statistics of interest change together with initial conditions and system parameters. However, sensitivity exercises show that, in our model, across-simulation variability is quite low and no chaotic pattern is detected. Hence, we confidently present below results concerning averages over a limited number of replications (typically $M = 50$) as a robust proxy for the behavior of any statistics we compute. Tables 1 and 2 report the values employed for initial conditions and parameters. Our choice of initial conditions is done in such a way that the economy evolves over a steady state in absence

⁶On the methodology of analysis of evolutionary / agent-based computational economics models, see e.g. Lane (1993) and Pyka and Fagiolo (2007).

Table 2 Benchmark parametrization

Description	Symbol	Value
Size of consumption-good industry	F_1	200
Size of capital-good industry	F_2	50
Econometric sample size	T	600
Replicator dynamics coefficient	χ	-0.5
Competitiveness weights	$\omega_{1,2}$	1
Innov. uniform distrib. supp. lower bound	$\iota_{1,2}^-$	-0.5
Innov. uniform distrib. supp. upper bound	$\iota_{1,2}^+$	0.5
Wage setting Δcpi weight	ψ_1	0
Wage setting $\Delta \bar{A}$ weight	ψ_2	1
Wage setting ΔU weight	ψ_3	0
Desired inventories parameter	θ	0.1
Desired level of capacity utilization	u^d	0.75
Trigger rule	α	0.1
Payback period parameter	b	8
Maximum machine age	η	19
Maximum debt/sale ratio	Λ	2
Consumption-firm sample coefficient	κ	0.01
Interest rate	r	0
Wage share	φ	0.33

of technical change. All results presented below are robust to variations of the parameters within a reasonably large neighborhood of the benchmark parametrization reported in Table 2.⁷

4 Simulation results

In this section we explore the extent to which the foregoing model is able to account for the empirical regularities presented in Section 2. To do so, we shall compare simulation results under the *work-or-die* and the *social-security* scenarios described above.

To begin with, let us look at the outcomes of the model when technical change is turned off. In this case, the model behaves like the Solow growth model: the economy is always in steady state and, since population is fixed, the output growth rate is zero. At the microeconomic level, the initial configuration with homogeneous firms remains unaltered as there is neither entry nor exit.

As soon as one turns on technical change, self-sustaining patterns of growth emerge (cf. Figs. 1 and 2). The economy evolves in a permanent disequilibrium state characterized by entry and exit of heterogeneous firms interacting both within and among industries.

⁷In Dosi et al. (2005, 2006) we perform extensive Monte Carlo sensitivity analyses on the most relevant parameters of the closest antecedent of the present model. The picture emerging from Monte Carlo studies confirms all results presented below.

Fig. 1 Work-or-die scenario. Level of output, investment and consumption




Fig. 2 Social-security scenario. Level of output, investment and consumption




Fig. 3 Work-or-die scenario. Bandpass-filtered output, investment and consumption





Fig. 4 Social-security scenario. Bandpass-filtered output, investment and consumption



Simulated aggregate time-series possess in this case statistical properties well in line with empirically-observed ones. More precisely, if we separate the business cycle frequencies of the series by applying a bandpass filter,⁸ we observe the typical “roller coaster” shape that characterizes real data (see Figs. 3 and 4 and Section 2 above). In the *social-security* scenario, simulated series of aggregate investment appear to be more volatile than output (SF1), whereas the opposite seems to happen in the *work-or-die* scenario. Finally, aggregate investment and consumption display a procyclical behavior in both scenarios.

In addition, the model is also able to generate a microeconomic landscape consistent with the micro “stylized facts” mentioned in Section 2. So, for example, the skewed size distributions which emerge in the simulations are not statistically different from the empirically observed ones in either scenarios (cf. the rank-size plot in Fig. 5).⁹


Furthermore, well in tune with the empirical evidence, pooled firm growth rates exhibit the typical “tent-shaped” pattern, characterized by tails fatter than the Gaussian benchmark (cf. Fig. 6). More precisely, we have fitted our simulated firm growth-rate distributions with the Subbotin family of densities.¹⁰ We find that simulated growth rates are well proxied by Subbotin densities with estimates for the shape-parameter that robustly suggest a departure from normality in both the *work-or-die* and the *social-security* scenarios (with $\hat{\beta} = 0.32$ and $\hat{\beta} = 0.24$, respectively). Notice that our estimates actually

⁸See Baxter and King (1999). Cf. also Dosi et al. (2005) for a discussion of the properties of alternative filtering techniques.

⁹We employ consumption-good firm sales as a proxy of firm size. Before pooling our data, we normalize each observation by the year-average of firm size in order to remove any time trends in our data. This allows us to get stationary size and growth distributions across years. Due to space constraints, we show the rank-size plot and the firm growth rate distribution plot for the *work-or-die* scenario only.

¹⁰Subbotin densities include as special cases the Normal (shape parameter $\beta = 2$) and the Laplace ($\beta = 1$) distributions. More on the application of the Subbotin family to the fitting of firm growth rates is in Bottazzi and Secchi (2006).

Fig. 5 Work-or-die scenario. Pooled (year-standardized) sales distributions. Log rank vs. log size plots. *M-G*: model-generated distribution



entail growth-rates distributions with tails even fatter than those empirically observed. We argue that this result is due to the different statistical features of real-world and simulated firm growth data samples. On the one hand, in empirically observed growth-rate distributions, small firms are typically not included in the sample and any entry-exit turbulence is washed away by considering surviving firms only. Conversely, in our simulated data: (i) we do not set any lower bound to the size of firms; and (ii) we consider also the entry and exit of firms. Both features of simulated data tend to increase the proportional “lumpiness” of growth shocks. In fact, simulation results show that, if one suitably builds balanced samples of simulated firm growth rates, the estimated shape-parameter turns out to increase and replicate its empirical counterpart (e.g., $0.5 \leq \hat{\beta} \leq 1$).

Let us now turn to a more detailed study of simulated aggregate time series. More specifically, we shall investigate the issue whether aggregate output, investment, consumption, etc. display statistical properties similar to the empirically observed ones, as summarized in SF1-4.

Fig. 6 Work-or-die scenario. Pooled (year-standardized) firm growth rates. Binned densities of simulated growth rates vs. Laplace fit. *M-G*: model-generated growth rates




Table 3 Output, investment and consumption statistics

Statistic	GDP	Consumption	Investment
(a) Work-or-die scenario			
Avg. growth rate (%)	1.8 (0.007)	1.8 (0.006)	1.8 (0.005)
DF test (logs)	-0.0988	0.9914	0.3692
DF test (bpf 6,32,12)	-5.6450*	-4.8685*	-6.2572*
Std. dev. (bpf 6,32,12)	1.1720	0.6198	0.3306
Rel. std. dev. (GDP)	1	0.5288	0.2821
(b) Social-security scenario			
Avg. growth rate (%)	1.8 (0.0006)	1.8 (0.0005)	1.8 (0.0017)
DF test (logs)	2.6816	5.8739	-0.3739
DF test (bpf 6,32,12)	-6.3837*	-6.0359*	-6.8881*
Std. dev. (bpf 6,32,12)	0.1358	0.0946	0.4357
Rel. std. dev. (GDP)	1.00	0.70	3.21

Standard deviations in parentheses. DF test specification: no intercept term, no linear trend, not augmented

*Significant at 1%

We begin by focusing on the average growth rate (AGR) of the economy:

$$AGR_T = \frac{\log Y(T) - \log Y(0)}{T + 1}, \quad (18)$$

where Y denotes aggregate output and T is the econometric sample size.¹¹

We then compute Dickey-Fuller (DF) tests on output, consumption and investment in order to detect the presence of unit roots in the series (all results refer to averages computed across $M = 50$ independent simulations). In both scenarios, the AGRs of output, consumption and investment are strictly positive ($\approx 1.8\%$, see Table 3).¹² DF tests strongly suggest that output, consumption, and investment are non-stationary. This result is robust to alternative specifications of DF tests (e.g., considering an intercept term, adding a linear trend, etc.).

We then detrend the time series obtained from simulations with a bandpass filter (6,32,12) and compute standard deviations and cross-correlations between output and the other series. Our simulated figures for relative standard deviations show that the model is able to match SF2 (i.e. consumption is less volatile than output) in both scenarios. However, in the *work-or-die* scenario, output appears to be more volatile than investment. This result stems from the fact that our simulated economy does not contain any mechanism that contributes to stabilizing effective demand, e.g. service industries and, especially, the government sector. When, as happens in the *social-security*

¹¹ All results refer to $T = 600$ time-periods, cf. Table 2. This econometric sample size is sufficient to allow for convergence of recursive moments of all statistics of interest.

¹² Note also that the ex post identity between savings and investments is always satisfied.

Table 4 Correlation structure

Time series (bpf 6,32,12)	GDP (bpf 6,32,12)								
	t-4	t-3	t-2	t-1	t	t+1	t+2	t+3	t+4
(a) Work-or-die scenario									
GDP	-0.19	0.14	0.54	0.87	1.00	0.87	0.54	0.14	-0.19
Consumption	-0.09	0.24	0.58	0.85	0.95	0.86	0.59	0.23	-0.12
Investment	-0.18	-0.17	-0.08	0.11	0.38	0.64	0.78	0.73	0.50
Change in stocks	0.03	0.17	0.31	0.40	0.39	0.28	0.10	-0.10	-0.24
Net investment	0.06	0.29	0.49	0.63	0.69	0.63	0.46	0.21	-0.07
Employment	-0.13	0.17	0.52	0.81	0.96	0.89	0.64	0.28	-0.08
Unemployment rate	-0.01	-0.25	-0.49	-0.67	-0.75	-0.71	-0.53	-0.26	0.04
(b) Social-security scenario									
GDP	-0.38	-0.03	0.44	0.84	1.00	0.84	0.44	-0.03	-0.38
Consumption	-0.35	-0.04	0.38	0.78	0.98	0.89	0.54	0.08	-0.32
Investment	-0.19	-0.15	-0.03	0.18	0.42	0.58	0.58	0.40	0.11
Change in stocks	-0.18	0.01	0.24	0.41	0.44	0.31	0.09	-0.11	-0.21
Net investment	-0.26	0.06	0.45	0.76	0.87	0.73	0.41	0.02	-0.28
Employment	-0.40	-0.10	0.33	0.73	0.96	0.91	0.61	0.17	-0.23
Unemployment rate	0.40	0.09	-0.33	-0.73	-0.96	-0.91	-0.61	-0.18	0.22

scenario, we include a proxy for the foregoing stabilizing factors, investment turns to be more volatile than GDP, thus satisfying SF1 as well.

As far as cross-correlations are concerned, consumption appears to be procyclical and coincident in both scenarios (cf. Table 4). This matches SF3. Change in inventories appears to be procyclical and coincident in the *social-security* scenario (SF3), whereas it is slightly leading in the *work-or-die* scenario. Investment is instead procyclical and leading in both scenarios. However, this result is entirely due to the dynamics of replacement investment. Indeed, net investment is always procyclical and coincident (SF3).

Fig. 7 Work-or-die scenario. Model generated (*M-G*) vs. empirical data (*S-W*: Stock and Watson 1999) cross-correlations

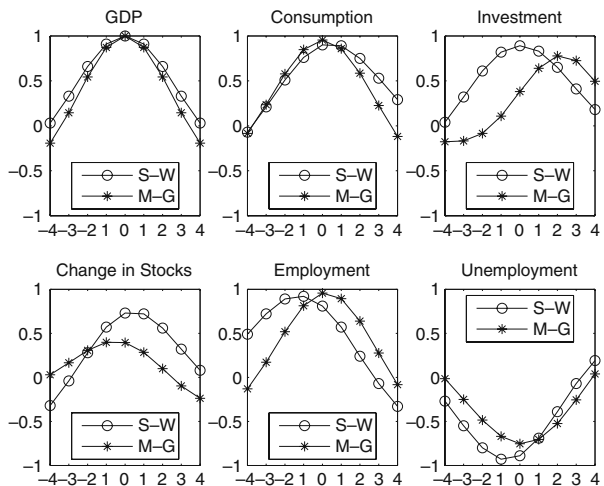
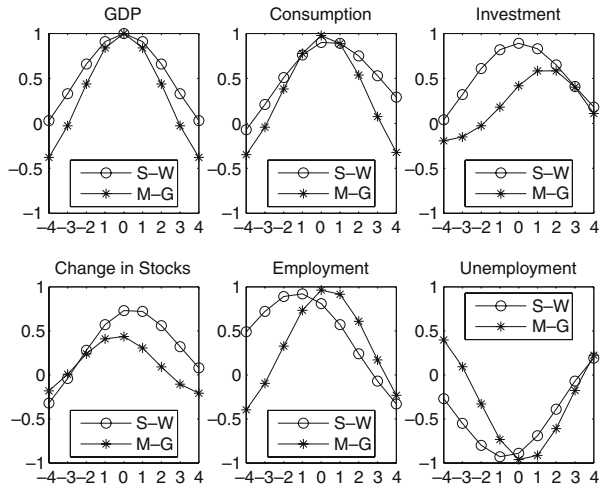


Fig. 8 Social-security scenario. Model generated (*M-G*) vs. empirical data (*S-W*: Stock and Watson 1999) cross-correlations



Finally, our simulated cross-correlation patterns are also *quantitatively* in line with those obtained by Stock and Watson (1999) on U.S. data (see Figs. 7 and 8).

Notwithstanding the fact that we did not model the labor market in detail, empirically-plausible employment series do arise. Indeed, employment turns out to be procyclical, whereas unemployment is countercyclical (SF4). Notice, however, that the two variables appear to be coincident. This result may stem from the complete lack of frictions that characterizes the labor market in our model. Indeed, since in every time period firms can hire and fire workers without limitations, production fluctuations pour out in the labor market with no lags.

Furthermore, we checked whether our model is able to match micro-economic stylized facts on productivity dynamics (SF10-11). To do so, we computed – at each t – the standard deviation of labor productivities across consumption-good firms in both scenarios. Our results (not shown) indicate that significant asymmetries persist throughout the history of our simulated economy (in tune with SF10). Moreover, firm-productivity auto-correlations remain significantly larger than zero for many lags, thus suggesting persistence in micro productivity differentials (cf. SF11).¹³

Finally, we explored the distributional properties of pooled, aggregate-output growth rates. In both the *work-or-die* scenario and – for a wide range of φ parameter values – in the *social-security* scenario, the estimation of the Subbotin shape parameter (β) robustly reveals departures from normality. Fat tails emerging in aggregate output growth rates are thus in line with the empirical evidence discussed in Fagiolo et al. (2008).

¹³Firm-productivity auto-correlations (up to lag 6) are computed by considering normalized productivity of firms that survived for at least 40 periods in the last 100 periods of any simulation run.

5 Conclusions

In this work, we have explored the properties of an evolutionary, agent-based model of output and investment dynamics. In the model, aggregate dynamics is driven by a population of heterogeneous boundedly-rational firms, which operate in two vertically linked sectors producing “machines” and a consumption good. Technical progress occurs in the machine-tool industry and diffuses in the consumption-good industry via time-consuming investment by firms.

The results, despite the simplicity of the model,¹⁴ appear to be surprisingly in tune with a rather long list of empirical “stylized facts” – concerning both the properties of aggregate variables and the underlying microeconomics.

First, self-sustained growth emerges together with fluctuations in macroeconomic variables characterized by statistical properties similar to the empirically observed one. Indeed, simulations show the emergence of self-sustaining, fluctuating patterns of output growth out of the interactions among firms operating in market regimes that strongly depart from perfect competition. Furthermore, the properties of simulated series of investment, consumption, employment, etc. all closely match their empirical counterparts.

Second, the microeconomic picture stemming from simulations is quite in tune with the one observed in empirical works. Persistent heterogeneity in firm size and growth patterns, as well as persistence in productivity differentials, are all robustly detected. Moreover, distributional properties of firm size and growth patterns mimic real-world evidence.

Finally, evolutionary microfoundations are shown to exhibit a macrodynamics with strong Keynesian features. Indeed, investment and production decisions induce in the model demand propagation effects much alike Keynesian “multiplier” effects. Conversely, adaptive expectations on demand drive investments in ways closely resembling the Keynesian “accelerator”. The resulting aggregate demand fluctuations endogenously give rise to business cycles.

References

- Bartelsman E, Doms M (2000) Understanding productivity: lessons from longitudinal microdata. *J Econ Lit* 38:569–594
- Baxter M, King R (1999) Measuring business cycle: approximate band-pass filter for economic time series. *Rev Econ Stat* 81:575–593
- Bottazzi G, Secchi A (2006) Explaining the distribution of firm growth rates. *RAND J Econ* 37:235–256
- Brenner T, Werker C (2007) A taxonomy of inference in simulation models. *Comput Econ* 30:227–244

¹⁴For example, the demand side of the economy, as well as the labor market, are not explicitly modeled. A more detailed microfoundation of such admittedly neglected ingredients is one of the main points in our agenda.

- Caballero R (1999) Aggregate investment. In: Taylor J, Woodford M (eds) *Handbook of macroeconomics*. Elsevier Science, Amsterdam
- Castaldi C, Dosi G (2004) Income levels and income growth. Some new cross-country evidence and some interpretative puzzles. Working paper 2004/18, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy
- Chiaromonte F, Dosi G (1993) Heterogeneity, competition, and macroeconomic dynamics. *Struct Chang Econ Dyn* 4:39–63
- Dosi G (2005) Statistical regularities in the evolution of industries. A guide through some evidence and challenges for the theory. Working paper 2005/17, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy
- Dosi G, Fabiani S, Aversi R, Meacci M (1994) The dynamics of international differentiation: a multi-country evolutionary model. *Ind Corp Change* 3:225–242
- Dosi G, Fagiolo G, Roventini A (2005) Animal spirits, lumpy investment and endogenous business cycles. Working paper 2005/04, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy
- Dosi G, Fagiolo G, Roventini A (2006) An evolutionary model of endogenous business cycles. *Comput Econ* 27:3–34
- Dosi G, Freeman C, Fabiani S (1994) The process of economic development: introducing some stylized facts and theories on technologies, firms and institutions. *Ind Corp Change* 3:1–45
- Fagiolo G, Moneta A, Windrum P (2007) A critical guide to empirical validation of agent-based models in economics: methodologies, procedures, and open problems. *Comput Econ* 30: 195–226
- Fagiolo G, Napoletano M, Roventini A (2008) Are output growth-rate distributions fat-tailed? Some evidence from OECD countries. *J Appl Econ* (in press)
- Hubbard R (1998) Capital-market imperfections and investment. *J Econ Lit* 36:193–225
- King R, Rebelo S (1999) Resuscitating real business cycles. In: Taylor J, Woodford M (eds) *Handbook of macroeconomics*. Elsevier Science, Amsterdam
- Kirman, A (1989) The intrinsic limits of modern economic theory: the emperor has no clothes. *Econ J* 99:126–139
- Lane DA (1993) Artificial worlds and economics, part I and II. *J Evol Econ* 3:89–107 and 177–197
- Mankiw GN, Romer D (eds) (1991) *New Keynesian economics*. MIT, Cambridge
- Napoletano M, Roventini A, Sapio S (2006) Are business cycles all alike? A bandpass filter analysis of the Italian and US cycles. *Riv Ital Econ* 1:87–118
- Nelson R, Winter S (1982) *An evolutionary theory of economic change*. The Belknap Press of Harvard University Press, Cambridge
- Pyka A, Fagiolo G (2007) Agent-based modelling: a methodology for neo-schumpeterian economics. In: Hanusch H, Pyka A (eds) *The Elgar companion to neo-Schumpeterian economics*. Edward Elgar, Cheltenham
- Stock J, Watson M (1999) Business cycle fluctuations in U.S. macroeconomic time series. In: Taylor J, Woodford M (eds) *Handbook of macroeconomics*. Elsevier Science, Amsterdam