ELSEVIER

Journal of Economic Behavior & Organization
Vol. 64 (2007) 426–447

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

Industrial dynamics, fiscal policy and R&D: Evidence from a computational experiment

Alberto Russo [a], Michele Catalano [b], Edoardo Gaffeo [c], Mauro Gallegati [b,*], Mauro Napoletano [d]

[a] *Scuola Normale Superiore, Pisa, Italy*
[b] *Department of Economics, Università Politecnica delle Marche, Italy*
[c] *Department of Economics and CEEL, University of Trento, Italy*
[d] *ETH-Zentrum, Zürich, Switzerland*

Received 1 October 2005; received in revised form 9 May 2006; accepted 1 June 2006
Available online 16 June 2007

Abstract

We present an agent-based computational model in which bounded rational firms and workers trade on fully decentralized markets for final goods and labor by means of random matching protocols. The model replicates several macroeconomic phenomena regularly observed in the data, with aggregate features emerging from the localized interactions of individual entities. The model is then used as a computational laboratory to run an experiment on the role of fiscal policy in increasing macroeconomic performance.

JEL classification: C63; L16; E62

Keywords: Agent-based computational models; Industrial dynamics; Fiscal policy

1. Introduction

Balancing the roles of individual rationality, institutions and evolutionary forces in understanding long-run economic performance has long been a central concern of the economic profession. While the neoclassical paradigm has traditionally emphasized rationality as the key to economic efficiency, scholars belonging to schools of thought as diverse as the post-Keynesian and the evolutionary ones have popularized the view according to which severe cognitive and procedural

* Corresponding author.
E-mail address: mauro.gallegati@univpm.it (M. Gallegati).

0167-2681/$ – see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.jebo.2006.06.016

bounds on individual rationality prevent agents to adhere to fully optimizing behavior. Institutions and selection mechanisms then naturally enter the scene, the former ones as a means to reduce uncertainty providing a stable – albeit not necessarily efficient – structure to social interactions, the latter ones as the determinants of industrial demography.

For macroeconomists, such a contraposition is closely related to another heated methodological debate, i.e., the one centered on the program of microfoundation of macroeconomics and on the meaning and convenience of the representative agent approach. In a series of inspiring contributions, Kirman (1992, 1997, 1999) has convincingly dubbed the representative agent as unfitted to put individual rationality at the roots of macroeconomics, as well as to solve the problem of aggregation. He advocates instead a vision of the economy in terms of a complex system, in which it is the interaction of a multitude of heterogeneous individuals with limited power of reasoning that generates the emergence of a number of macroeconomic phenomena, which would be otherwise difficult to explain with standard models.

A useful research strategy to model the economy as a complex system consists in exploiting the potentialities of agent-based computational techniques (Epstein and Axtell, 1996; Tesfatsion, 2002). At the simplest level, agent-based models are computer programs that simulate the autonomous behavior of individual entities and the relationships between them. Such virtual environments are particularly powerful and flexible, as they can be employed for advancing theoretical conjectures as well as for testing alternative normative prescriptions in a controlled situation. In fact, we claim that the agent-based approach represents a fruitful methodology to do *realistic* macroeconomics, that is one based on bounded rational, heterogeneous interacting agents adapting to a complex world.

In this paper we present the results of a computational experiment of redistributive policy in an agent-based model in which decentralized transactions occur through bilateral bargains. We model an economy populated by a large number of firms and workers/consumers who execute their trades on the markets for goods and labor, by means of random matching protocols. Individual choices are updated adaptively according to simple routinized procedures, which are sometimes conducive to failure. Hence, this study is rooted in the Nelson and Winter (1982) methodological perspective, according to which substantive and procedural uncertainty force real agents operating in complex markets to adopt heuristic procedures or routines, while natural economic selection forces the exit of agents whose fitness turns out to be insufficient.

Simulations of a baseline model in which firms invest part of their operating profits in R&D exhibit interesting aggregate behaviors emerging from the dispersed interactions of individual entities. In particular, the economy displays: (i) sustained growth characterized by large fluctuations; (ii) key stylized facts of the labor market, like the Beveridge, the Phillips and the Okun curves; (iii) fat-tailed distributions for the growth rates of firms' size, for labor incomes and for firms' equity, as one can typically observe in the data; (iv) a non-monotonic relationship between the amount spent in R&D at an individual level and macroeconomic performance measured by the growth rate of aggregate output.

A key point which deserves to be emphasized is that the facts listed above are largely dependent from the institutional framework – e.g., trading protocols and the information acquisition technology – rather than learning or rationality effects, given that agents are endowed with almost-zero intelligence (Gode and Sunders, 1993). The simple fact of recognizing the importance of the institutional framework in shaping aggregate outcomes leads naturally to the issue of how to *exogenously* design institutions aimed at fostering economic performance or, in other terms, to economic policy. Since the system endogenously generates substantial heterogeneity of firms' size and personal incomes, it seems worthwhile to ask whether one can find some redistributive

scheme which can generate better macroeconomic performances in terms of, e.g., the level and growth rate of aggregate output.

To address this question, we use the baseline model as a computational laboratory to investigate the role of alternative fiscal policies in a macroeconomy characterized by heterogeneous interacting agents. The model economy is thus supplemented with a government who levies taxes on corporate profits and redistribute revenues, under the constraint that the public budget is always balanced. Alternative treatment designs are compared, where we use as treatment variables the tax rate and transfer payments, which can be tailored to benefit unemployed people or firms' R&D investments.

To briefly summarize our results, we find that the aggregate output is non-monotonically related with the level of the flat rate tax on corporate profits – i.e., output increases with the level of the tax rate up to a threshold which can be dubbed as *pseudo-optimal* – as soon as revenues are redistributed on a per-capita basis to subsidize investments in R&D. On the contrary, the net effect in terms of aggregate output is negative for any level of the tax rate if transfer payments are used to provide unemployment insurance. Furthermore, the level of the optimal tax rate in the former case depends critically on the level of private efforts in R&D investment: an economy with a low level of private investment in R&D benefits from a relatively higher tax rate on corporate profits, while the opposite is true if private investments in R&D are sufficiently high.

The remainder of the paper is organized as follows. Section 2 outlines the baseline model and its setup for simulations, while Section 3 illustrates its main aggregate emerging properties. In Section 4 we describe the experimental design aimed at assessing the effectiveness of alternative fiscal policies, and discuss our findings. Section 5 concludes.

2. The model

We consider a sequential economy populated by a large number of firms $i = 1, \ldots, I$, and a large number of workers/consumers, $j = 1, \ldots, J$, who undertake decisions at discrete times $t = 1, \ldots, T$ on the markets for a homogeneous non-storable consumption good and labor services. Firms are run by a single manager each, who shares the following two features. First, they use bounded rational decision rules, in that they choose prices and quantities in an adaptive way. Second, managers employ a common fraction of net profits to make uncertain investments in R&D, with the aim of increasing the productivity of their firm.

Both the goods and the labor markets are characterized by decentralized search and matching processes, which imply out-of-equilibrium dynamics. Thus, due to the absence of market-clearing mechanisms the economy is characterized by the contemporaneous occurrence of persistent involuntary unemployment, unsold production and excess individual demand.

2.1. Sequence of events

The sequence of events occurring in each period runs as follows:

1. At the beginning of the time period t, the generic ith firm checks its financial viability, given by the stock of precautionary reserves accumulated from past retained earnings (net worth) $A_{it} = A_{it-1} + (1 - \sigma)\pi_{it-1}$, where π_{it-1} are nominal profits in period $t - 1$ and σ is the fraction spent on R&D investment. If the net worth is positive, the firm starts a new production cycle. If on the contrary the net worth is lower or equal to zero, the firm becomes technically insolvent

and shuts down due to bankruptcy.[1] In this case, a new firm enters the market. Each entrant is simply a random copy of surviving firms.
2. The productivity of firm i is increased by a random amount whose expected value depends on the investment in R&D, which is in turn a fixed fraction of the last period profits $\sigma\pi_{it-1}$.
3. Starting from the demand it expects to face, each operating firm determines the amount of output to be produced and the amount of labor to be hired. Expectations on future demand are updated adaptively.
4. Workers update the satisfying wage they ask for, taking into account price inflation and their occupational status.
5. A fully decentralized labor market opens. Firms post their vacancies on the basis of their labor demand. Each worker, in turn, sends M applications to as many firms: one to the firm in which he worked during the last period (if employed), and $M-1$ at random. Firms sort workers' applications by the requested wage level (in ascending order), and hire workers to fill all their open vacancies. At the end of the period such a matching protocol determines whether unfilled vacancies or unemployed workers remain after the labor market has closed.
6. Production takes the whole period t, regardless of the scale of output.
7. After the production is completed, a market for the consumption good opens. Firms post their offer price, while consumers are allowed to muddle through searching for a satisfying deal to spend the labor income gained during the period $t-1$. If a firm ends up with excess supply, it gets rid of the unsold goods at zero costs.
8. Firms collect revenues, pay their wage bill, calculate profits, allocate a fraction of their profits to R&D, and update their net worth.

2.2. Production technology, R&D and productivity evolution

Production is carried out by means of a constant returns to scale technology, which employs labor (L_{it}) as the only input:

$$Y_{it} = \alpha_{it} L_{it}, \tag{1}$$

where $\alpha_{it} > 0$ is the firm's labor productivity, that evolves over time depending on firm's R&D investment in innovative technologies. The search for new technological opportunities is assumed to be a risky and uncertain endeavor, with the probability of success depending on the amount of money invested. Let us define the ith firm R&D investment as $\mathrm{RD}_{it} = \sigma\pi_{it}$. The law of motion of productivity follows a first order autoregressive stochastic process:

$$\alpha_{it+1} = \alpha_{it} + z_{it}, \tag{2}$$

where z_{it} is the outcome of a random variable exponentially distributed with mean $\mu_{it} = \mathrm{RD}_{it}/(p_{it}Y_{it})$.[2]

[1] A far from trivial consequence of bankruptcy is the impossibility for the firm to pay the whole wage bill it owes to the workers employed in that period. In this case, workers obtain only a fraction of the revenues of the firm, R_{it}, plus a fraction of the cash flow A_{it} available. In other terms, workers bear the risk of firms' bankruptcy.

[2] The normalization factor $p_{it}Y_{it}$ is aimed at letting the effect of R&D investments on firm's productivity to be independent of the scale of production. In fact, making use of the definition of nominal profits and of Eq. (1) in the main text, it turns out that $\mu_{it} = \sigma/\alpha_{it}(\alpha_{it} - w_{it})$.

2.3. The market for the consumption good

At the beginning of each time period, firms adjust adaptively either their price or their output level. For the sake of simplicity, the "choice" between the two strategies depends on an exogenous random process: each firm extract a number from a uniform distribution with support (0, 1); the firm "decides" to modify the price, keeping unchanged quantities, if the extracted number is smaller or equal to 0.5; the firm changes quantities and leaves the price unchanged in the opposite case. Accordingly, in each period each firm has a probability equal to 1/2 to adjust its price or output level.

We assume that firms operate in a posted offer market. In spite of the good being homogeneous, asymmetric information and search costs imply that consumers may end up to buy from a firm regardless of its price not being the lowest. It follows that the law of one price does not necessarily apply (Stiglitz, 1989). Prices are set considering both the unsold quantities during the last period (S_{it-1}), and the costs incurred in production. More precisely, the ith manager sets his *satisfying* selling price according to the following rule:

$$P_{it}^{s} = \begin{cases} P_{it-1}(1+\eta_{it}) & \text{if } S_{it-1} = 0 \\ P_{it-1}(1-\eta_{it}) & \text{if } S_{it-1} > 0 \end{cases} \tag{3}$$

with η_{it} being an idiosyncratic random variable uniformly distributed on a positive support $(0, H)$. Let us define as P_{it}^{l} the lowest price at which firm i is able to cover its average costs:

$$P_{it}^{l} = \frac{W_{it}}{Y_{it}}, \tag{4}$$

where W_{it} is the wage bill firm i pays at time t. The price determined according to rule (3) is therefore posted on price-tags if and only if $P_{it}^{s} \geq P_{it}^{l}$. In the opposite case, the firm posts the price P_{it}^{l}.

The level of production planned at the beginning of period t (Y_{it}^{d}) depends on expected demand, $Y_{it}^{d} = D_{it}^{e}$. Expectations on total orders to be received are taken adaptively according to:

$$D_{it}^{e} = \begin{cases} Y_{it-1}(1+\rho_{it}) & \text{if } S_{it-1} = 0 \\ Y_{it-1}(1-\rho_{it}) & \text{if } S_{it-1} > 0 \end{cases}, \tag{5}$$

where ρ_{it} is an idiosyncratic shock uniformly distributed on a positive support $(0, H)$, with $H \leq 1$. Thus, demand expectations are revised upward if a manager observes excess demand for its output, and downward when the opposite holds.

Aggregate demand equals total wages paid by firms to workers employed in $t-1$, as we assume that workers express individual demand functions with a unitary marginal propensity to consumption. Given the lack of any market-clearing mechanism and that bargains on the good market are fully decentralized, consumers have to search for satisfying deals. The information acquisition technology is defined in terms of the number Z of firms a consumer can visit without incurring any cost. In other words, search costs are null as the consumer enters the market, continue to be null if he remained confined into his local market of size Z, but they become prohibitively high as soon as a consumer tries to search outside. In what follows, the identity of the Z firms associated to a generic consumer j are picked randomly at any time period t.

Consumers enter the market sequentially, the picking order being determined randomly at any time period t. Each purchaser j is allowed to visit Z firms to detect the price posted by each one of

them. Prices (and the corresponding firms) are then sorted in ascending order, from the lowest to the highest. Consumer j tries to spend all the income gained during the last period in goods of the cheapest firm in his local market. If the cheapest firm has not enough available output to satisfy j's needs, the latter tries to spend his remaining income buying from the firm with the second lowest price, and so on. If j does not succeed in spending his whole income after he visited Z firms, he saves what remain for the following period. For the sake of simplicity, the interest rate is assumed to be equal to 0.

After the market for consumption goods has closed, the ith firm has made sales for Q_{it}, at the price P_{it}. Accordingly, I's revenues are $R_{it} = P_{it}Q_{it}$. Due to the decentralized buying–selling process it is possible that a firm remains with unsold quantities ($S_{it} > 0$). In the following period, the variable S will be considered in adjusting firms' prices or quantities, as explained above.

2.4. The labor market[3]

Firms set their labor demand L_{it}^{d} on the basis of their desired level of production. From Eq. (1), it follows that the number of job openings set by firm i at time t is simply given by

$$L_{it}^{\mathrm{d}} = \frac{Y_{it}^{\mathrm{d}}}{\alpha_{it}}. \tag{6}$$

We assume that workers supply inelastically one unit of labor per period. Each worker sends M applications to as many firms: the first one to the firm in which he worked in the previous period (if employed), and $M-1$ at random (M if unemployed in $t-1$). Workers are therefore characterized by a sort of loyalty to their last employer, on the one hand, and by a desire to insure themselves against the risk of unemployment by diversifying in a portfolio of hiring opportunities, on the other one.

The ith firm organizes all received applications into two blocks. The first one is composed by all its previous employees, as employers respond to the loyalty of their workforce by assuring them a priority in their hiring policy. The second block of the queue, in turn, is filled in by all other applicants. Inside each block, workers' applications are sorted by wage level in ascending order. Firm i may face two alternative situations:

(a) $L_{it}^{\mathrm{d}} \leq L_{it-1}$, that is the desired labor demand at time t is lower than the number of people employed during the previous period. In this case, the last $L_{it-1} - L_{it}^{\mathrm{d}}$ workers (i.e., the ones asking for higher wages) queuing in the first block are fired, while the remaining are kept. Fired workers have other $M-1$ opportunities to find a job elsewhere.
(b) $L_{it}^{\mathrm{d}} > L_{it-1}$, that is firm i wants to increase its workforce. In this case, i keeps all its past employees and looks for $L_{it}^{\mathrm{d}} - L_{it-1}$ new workers, who are selected from the second block of the queue.

Decentralized labor markets (i.e., one for each firm) are closed sequentially according to an order randomly chosen at each time step. Given that each worker is allowed to sign one labor contract per period, serious coordination failures could arise as the number of workers actually

[3] For recent examples of agent-based computational macroeconomic models where the labor market is analyzed along similar lines, see Fagiolo et al. (2004) and Delli Gatti et al. (2005).

available does not necessarily correspond to the one inscribed in queues, especially for firms which are called to hire their workers late in the sequence.

Workers adjust their satisfying demand wage taking into account the past average price inflation ΔP_{t-1} and their past employment status, according to the rule:

$$w^{a}_{jt} = \begin{cases} w^{a}_{jt-1}(1 + \Delta P_{t-1})(1 + \varsigma_{jt}) & \text{if employed in } t-1 \\ w^{a}_{jt-1}(1 + \Delta P_{t-1})(1 - \varsigma_{jt}) & \text{if unemployed in } t-1 \end{cases}, \tag{7}$$

where ζ_{it} is a random term uniformly distributed on the same support of ρ_{it} $(0, H)$.

The contractual wage rate paid by firm i to employee j is defined on the basis of a bilateral bargaining scheme:

$$w_{jt} = \beta w^{a}_{jt}, \tag{8}$$

where the parameter $\beta \in (0, 1)$ measures the workers' power in the wage bargaining process. Notice that, according to the rules (7) and (8), workers with a history of active employment gets, *ceteris paribus*, a higher wage than workers who experienced several unemployment spells in their recent past. Therefore, the more workers who recently experienced unemployment is hired by firm i, the lower is the per-worker wage firm i has to pay.

3. Properties of the baseline model

In this section, the main properties of the baseline model are explored. Our main goal consists in addressing a couple of questions which are of capital importance whenever one tries to *do* macroeconomic theory without the representative agent:

(a) Is a completely decentralized economy inhabited by far from rational agents viable? At odds with macroeconomic equilibrium solutions based on fixed-point arguments, by viable macroeconomy we mean a system which "[...] *has a 'corridor' of stability and it is only for displacements that take it outside the corridor that it will exhibit serious* [...] *failures*" (Leijonhufvud, 1993, p. 8). The notion of a stability corridor reminds of a situation in which microeconomic agents interact along individual disequilibrium paths, but the aggregate system displays a stable behavior which could be interpreted in terms of a statistical equilibrium, that is "[...] *a state of macroscopic equilibrium maintained by a large number of transitions in opposite directions*" (Feller, 1957, p.356). The ability of a model to predict such a property is of particular importance in frameworks where the Walrasian auctioneer has been completely discarded as an unrealistic assumption. A great advantage of agent-based techniques is that they provide a computational controlled environment for studying when and how a large number of heterogeneous interacting agents coordinate themselves in allocating resources through dispersed trading activities. For obvious reasons of space, in this paper we will deal with normal conditions only, leaving the analysis of out of bounds behavior for future research.

(b) Is the system capable to replicate, at least qualitatively, one or more of the macroeconomic stylized facts which have been shown to hold for the majority of industrialized countries under *normal* economic conditions? In particular, we are interested in building a virtual environment able to capture the emergence of aggregate regularities as the result of the dispersed market interactions of a multitude of heterogeneous agents (Lane, 1993).

In analyzing the properties of our agent-based model, we will first concentrate on one single execution of the model, to subsequently present some robustness checks. Simulations of the baseline model have been run for the following choice of parameter and initial values for variables: $Z=2$; $M=2$; $H=0.1$; $\beta=1$; $\sigma=0.05$; $I=100$; $J=500$; $T=1000$. Notice that in what follows we do not perform any proper calibration exercise, as we are merely interested in assessing the qualitative features of the model.

The plots in Fig. 1 show that the system displays fluctuating growth of aggregate activity (GDP), but also that GDP fluctuations are well proxied by a Gaussian walk, so that the aggregate output moves in the long-run over a stable upward corridor. Fluctuations result from idiosyncratic shocks, which cause localized market failures and the exit of bankrupted firms. Stationary fluctuations characterize also the time series for the growth rates of prices and wages (Fig. 2(a) and (b)). Price inflation is slightly anti-cyclical (the correlation with the output growth rate is $\omega=-0.13$), while the wage inflation is slightly pro-cyclical ($\omega=0.3$).

The evolution of the macroeconomy is matched by the evolution of the industrial market structure. The time path of the Herfindahl–Hirschman (HH) index of industrial concentration registers the persistence of a wide and time-varying ($\omega=-0.23$) heterogeneity of firms' market shares, while the bankruptcy rate oscillates anti-cyclically ($\omega=-0.17$) in a stationary corridor (Fig. 2(b) and (c)).

The longitudinal heterogeneity characterizing agents is also reflected in a significant skewness of wealth distributions. Panels (a) and (b) of Fig. 3 report the Zipf plots (Stanley et al., 1995) – or (log–log) rank–size plot – for the firms' net worth distribution and for wages, respectively, measured at the final time period T. Furthermore, in each graph we report simulated data obtained from a Gaussian distribution with mean and standard deviation equal to the sample mean and standard deviation of the relevant variables. In both case, the upper tail of the empirical distribution is to the right of the Gaussian right tail, signaling that the wealth of richer agents is bigger than would be the case if the wealth distributions were normally distributed.[4] In other terms, wealth distributions for firms and for workers are fat-tailed, as one can normally observe in real data.

Additional experimental evidence on the viability of our model economy comes from an analysis of the variables summarizing how the labor market works. The unemployment rate, the ex ante vacancy rate (i.e., the ratio between the number of job openings and the number of unemployed workers), the average duration of labor contracts (i.e., the number of consecutive periods workers remain employed with the same employee) and the average duration of unemployment all exhibit stationary fluctuations (Fig. 4).

As regards the second evaluative criterion we pose for assessing the suitability of agent-based models for macroeconomic theorizing, notice that our framework displays some interesting aggregate regularities emerging from microeconomic dis-equilibrium interactions. Panel (a) of Fig. 5 shows that average real wage and average labor productivity follow a similar increasing path. Marginal costs fluctuates sensibly over higher frequencies, but on the average the system succeeds in self-organizing as the technological frontier moves up. The functional income distribution shares remain therefore basically constant over the whole simulation.

Panels (b)–(d) of Fig. 5, in turn, illustrate how decentralized adaptive bargaining in the labor market allows the emergence of a negative relationship between the rate of vacancies and the rate of unemployment (the so-called Beveridge curve) (Fig. 5(b)), a negative relationship between the

[4] Both the Bera–Jarque and the K–S tests for normality refute the null hypothesis for each distribution at the 1% significance level.

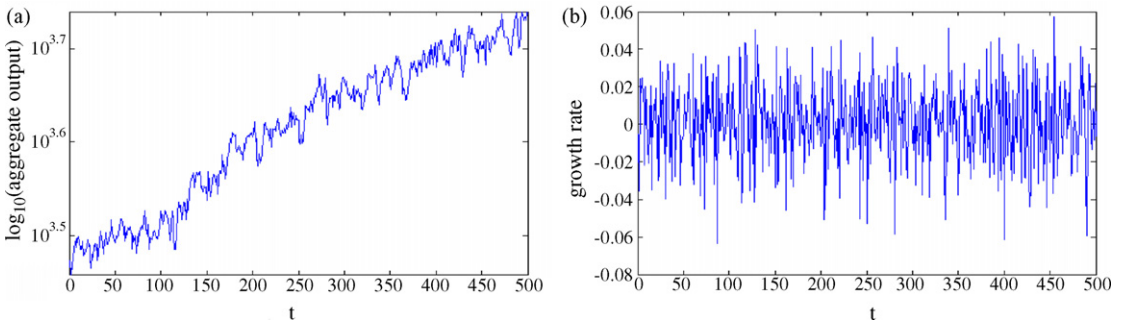

Fig. 1. Simulation results: aggregate activity. All the graphs in this section report simulated data for $t = 501{:}1000$. In other terms, transient dynamics is not presented. (a) Time evolution of aggregate output; (b) growth rates of aggregate output.

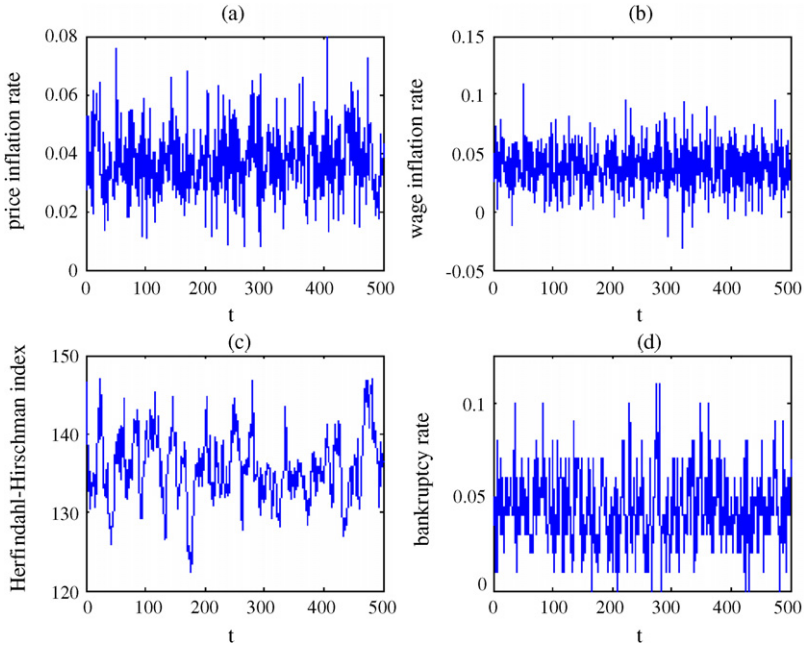

Fig. 2. Simulation results. (a) Price inflation; (b) wage inflation; (c) Herfindhal-Hirschman index; (d) bankruptcy rate.

rate of wage inflation and the rate of unemployment (Phillips curve) (Fig. 5(c)), and a negative relationship between the rate of unemployment and the rate of growth of aggregate output (Okun curve) (Fig. 5(d)). As discussed at length in Fagiolo et al. (2004), standard macroeconomic theories rooted in the representative agent approach fail to jointly explain these three aggregate regularities, whose existence in real data is confirmed by plenty of empirical evidence.

A recent strand of empirical work (see, e.g., Amaral et al., 1997; Bottazzi et al., 2001) has shown that the empirical probability distribution of the (log) growth rates of firms is tent-shaped, and it could be well fitted by means of a Laplace (symmetric exponential) functional form. Both tails decay much slower if compared to the Gaussian case, which in turn would result if the rate of growth of firms were time independent. This striking regularity turns out to be particularly robust, as it has been largely confirmed for different data sets, different countries, different sectors, and different proxies to measure firms' size. As shown in Fig. 6, the empirical density for simulated firms growth rates (in terms of final sales) has a clear triangular shape, thus adding another item to the list of stylized facts our model is able to replicate. A Laplace fit returns estimates for the location and the scale parameters equal to $\mu = 0.0338$ and $a = 0.194$, respectively.

The robustness of these qualitative results has been checked by recurring to Monte Carlo techniques. First, we run 100 independent simulations for different values of the initial seed generating the pseudo-random numbers. We then study the moments of the distributions of the statistics of interest. Results reported in Table 1 confirm that the findings discussed so far are indeed quite robust. A second Monte Carlo exercise has been run to analyze the role of the R&D in shaping macroeconomic performance. We focus on the relationship between the propensity of firms to invest in R&D, the growth of the economic system and the average values of major economic variables as the rate of unemployment, the rate of vacancies, the price inflation, and the

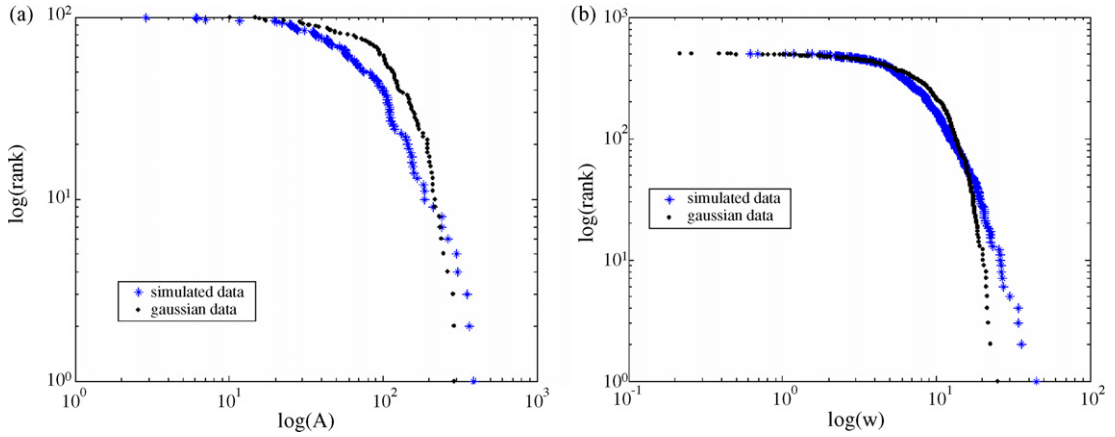

Fig. 3. The wealth distribution. For workers, total wealth is given by sum of the labor income and of involuntary savings from previous periods. (a) Firms' net worth distribution; (b) worker's wealth distribution.

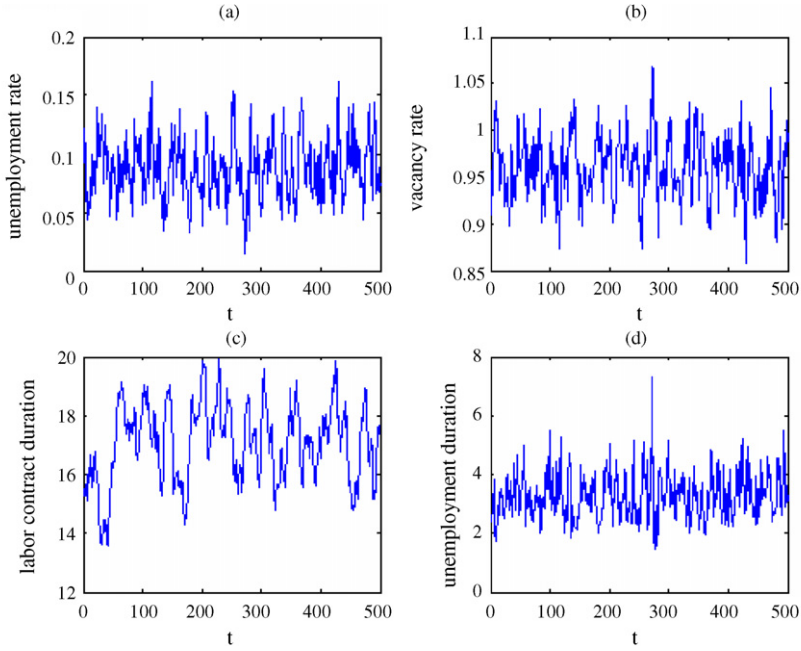

Fig. 4. Simulation results. (a) Unemployment rate; (b) vacancy rate; (c) labor contract duration; (d) unemployment duration.

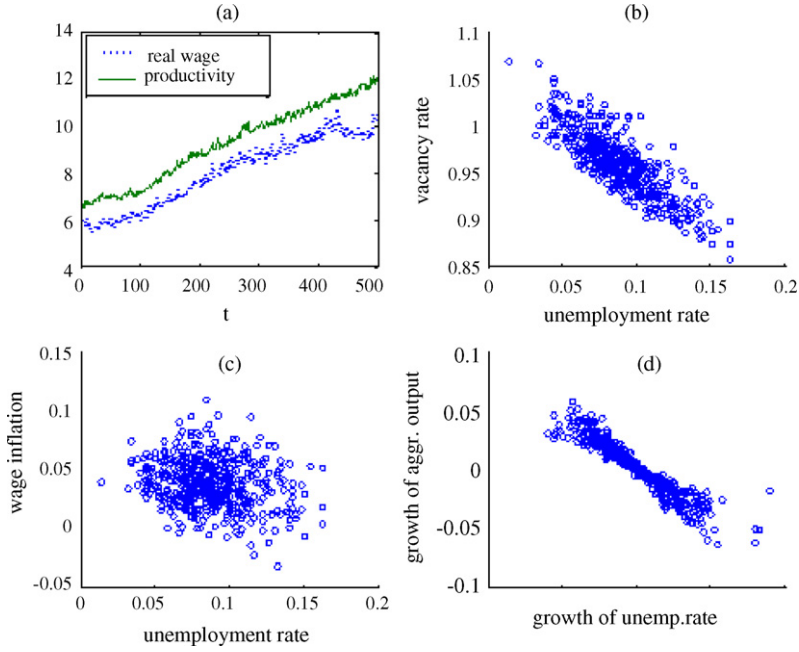

Fig. 5. Simulation results: emergent macroeconomic properties. (a) Real wage-productivity relationship; (b) Beveridge curve; (c) Phillips curve; (d) Okun curve.

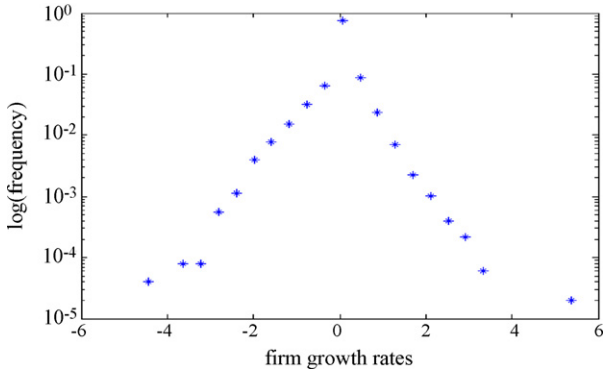

Fig. 6. Simulation results: Laplace distribution of firms' growth rates.

Table 1
Mean, standard deviation and coefficient of variation across 100 independent Monte Carlo simulations of the baseline model

	Mean	S.D.	CV
Growth rate	0.0013	0.0001	0.0978
Growth volatility	0.0223	0.0008	0.0372
Price inflation	0.0365	0.0003	0.0085
Wage inflation	0.0378	0.0002	0.0065
Unemployment rate	0.0870	0.0013	0.0149
Vacancy rate	0.9593	0.0020	0.0021
Labor contract duration	16.843	0.2000	0.0119
Unemployment duration	3.1909	0.0487	0.0153
HH index	135.06	0.8071	0.0060
Bankruptcy ratio	0.0420	0.0752	0.0177
Laplace location parameter	0.0352	0.0008	0.0232
Laplace scale parameter	0.1936	0.0065	0.0334

bankruptcy rate. To accomplish this task we perform 100 simulations of the model with increasing levels of the R&D parameter σ, starting from 0% through 99% with 1% steps.[5]

Simulations show that the relationship between efforts in R&D and the growth rate of aggregate activity is monotonically increasing, but also that R&D investment display decreasing marginal returns (Fig. 7). In our model, a higher fraction of profits invested in searching for new technological opportunities corresponds to a lower rate of net worth accumulation, which in turn leads to a higher financial fragility and a higher probability of bankruptcy (cf. panel (d) of Fig. 8). For sufficiently high level of R&D investments, the two forces tend to counteract.

An increase in R&D investment has monotonic but non-linear effects also on the unemployment rate, the inflation rate and the rate of vacancies. Recall that in this model technological progress is purely labor saving. It does not come as a surprise, therefore, that higher investments in R&D are associated to a higher rates of unemployment (Fig. 8(a)). By the same token, a higher productivity forces firms to open less job positions (Fig. 8(c)), and drives a higher price competition which

[5] The setting for all parameters other than the fraction of profits invested in R&D is the same used for the simulation of the baseline model.

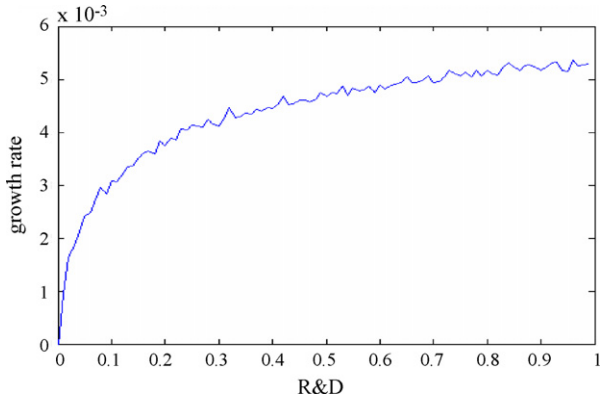

Fig. 7. Monte Carlo evidence on the relationship between R&D and the GDP growth rate.

determines a lower rate of inflation (Fig. 8(b)). The emergence of non-linear effects of increases in R&D spending is an interesting result by itself. It must be stressed, however, that in our artificial economy the aggregate labor-saving effect of R&D is somehow overemphasized, as the assumption of a single industry rules out any employment reallocation across sectors.

As an additional robustness check, we explore the parameter space for the searching costs both in the labor (M) and the goods markets (Z). In particular, we run 49 simulations for combinations

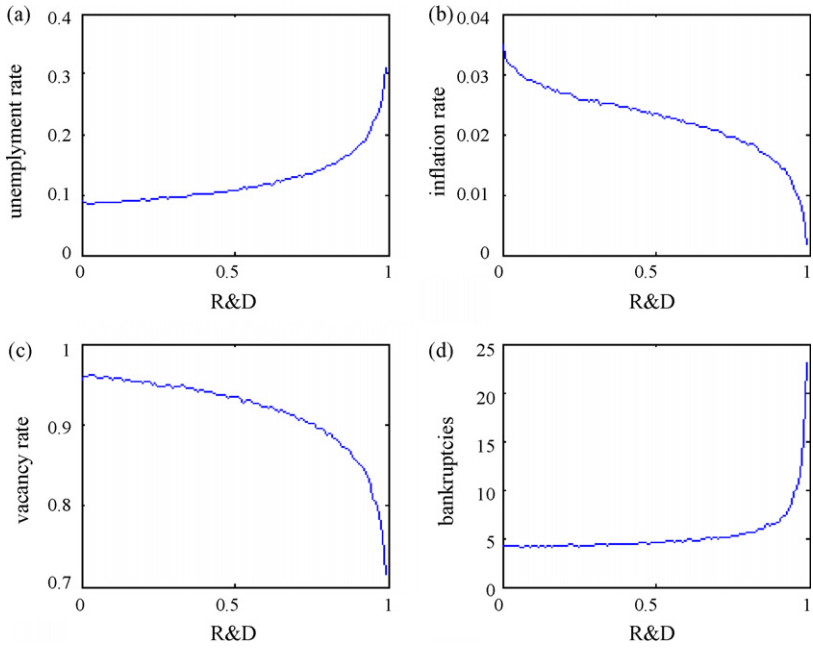

Fig. 8. Montecarlo evidence on the relationship between R&D and other relevant variables: (a) unemployment rate; (b) inflation rate; (c) vacancy rate; (d) bankruptcies.

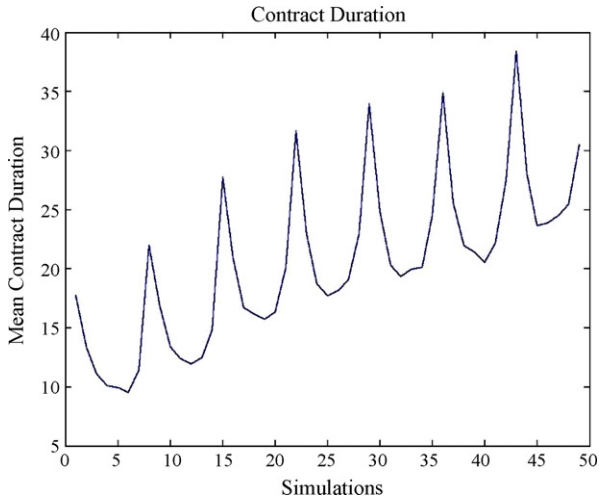

Fig. 9. Average duration of labor contracts for different levels of the searching costs.

of M and Z evolving in the space (1, 2, 4, 6, 8, 10, 20),[6] and we collect the time average of key macroeconomic variables across each simulation. Table 2 reports our results. As far as the GDP growth rate is concerned, it appears that lower search costs on the labor market have limited positive effects, while surprisingly an increase of Z does not seem to have any appreciable consequence on the aggregate growth performance. In turn, different combinations of M and Z determine large shifts of the Phillips curve: as search costs on both markets are lowered, inflation increases and unemployment decreases, but the elasticity of unemployment to variations of M and Z is much higher than that of inflation. This last result is echoed in the sensible increase on the ex ante vacancy rate as search costs shrink, signalling that the labor demand becomes higher than the labor supply as both consumers and workers are allowed to search more intensively. Such a soaring efficiency of the labor market in allocating idle resources for higher values of M and Z is somehow countervailed, however, by simultaneous higher rates of bankruptcy due to a stronger competition in the goods market.

A final interesting result is presented in Fig. 9, where we plot the mean of labor contracts' duration for each of the 49 simulations. Data are organized such that the first moment of the duration is collected for seven different values of M, given Z, and Z increases as we move from the left to the right along the x-axis. Two key facts emerge. First, for any given value of Z the average duration of labor contracts follows a U-inverted pattern, with peaks invariably given at $M = 1$ and troughs at $M = 8$ or 10. Second, the average length of contracts raise for higher values of Z, given M.

The reason for such a cyclical behaviour along an increasing trend rests on the combined effects exerted by intertemporal underbidding and market competition as searching costs on both markets decrease. As search costs in the labor market increase, the number of applications from outsider workers at each firm rises, as well as the probability to hire a new employee (in addition to the old ones) if that firm aims at increasing its labor demand. Recall that a new worker with a recent experience of unemployment, once hired, is likely to ask for a wage lower than that requested by

[6] The reason we limit M and Z at 20 is that for very low search costs (i.e., $M, Z > 20$) our artificial economy displays full employment and, as a result, to divergent dynamics between average wage and labor productivity.

Table 2
Simulated time average of key variables for different levels of search costs

M	Z						
	1	2	4	6	8	10	20
GDP growth rate							
1	0.0025	0.0026	0.0022	0.0024	0.0023	0.0023	0.0022
2	0.0024	0.0024	0.0023	0.0023	0.0023	0.0024	0.0023
4	0.0024	0.0024	0.0023	0.0024	0.0024	0.0024	0.0024
6	0.0025	0.0024	0.0024	0.0024	0.0026	0.0026	0.0024
8	0.0026	0.0025	0.0025	0.0025	0.0027	0.0027	0.0025
10	0.0026	0.0026	0.0027	0.0027	0.0026	0.0026	0.0027
20	0.0028	0.0028	0.0028	0.0029	0.0028	0.0028	0.0028
Inflation rate							
1	0.0239	0.0274	0.0304	0.0309	0.0316	0.0321	0.0327
2	0.0277	0.0301	0.0320	0.0324	0.0332	0.0334	0.0339
4	0.0310	0.0326	0.0342	0.0346	0.0349	0.0356	0.0359
6	0.0328	0.0343	0.0356	0.0364	0.0365	0.0367	0.0375
8	0.0340	0.0356	0.0368	0.0371	0.0373	0.0376	0.0381
10	0.0354	0.0362	0.0374	0.0381	0.0379	0.0386	0.0385
20	0.0386	0.0395	0.0395	0.0382	0.0393	0.0386	0.0357
Rate of unemployment							
1	0.1482	0.1156	0.0896	0.0796	0.0747	0.0720	0.0643
2	0.1127	0.0866	0.0671	0.0615	0.0565	0.0536	0.0491
4	0.0769	0.0596	0.0465	0.0407	0.0353	0.0325	0.0281
6	0.0578	0.0437	0.0318	0.0252	0.0226	0.0192	0.0159
8	0.0434	0.0322	0.0213	0.0170	0.0132	0.0124	0.0086
10	0.0330	0.0229	0.0136	0.0104	0.0093	0.0073	0.0049
20	0.0071	0.0039	0.0017	0.0013	0.0009	0.0005	0.0004
Vacancy rate							
1	0.9007	0.9378	0.9721	0.9902	0.9978	1.0023	1.0214
2	0.9266	0.9575	0.9896	1.0012	1.0102	1.0202	1.0397
4	0.9551	0.9778	1.0050	1.0201	1.0353	1.0462	1.0669
6	0.9716	0.9921	1.0225	1.0390	1.0517	1.0647	1.0881
8	0.9849	1.0041	1.0356	1.0500	1.0654	1.0717	1.1037
10	0.9946	1.0150	1.0474	1.0611	1.0713	1.0841	1.1064
20	1.0337	1.0501	1.0803	1.0938	1.1070	1.1174	1.1288
Bankruptcy rate							
1	0.0419	0.0396	0.0415	0.0456	0.0469	0.0483	0.0582
2	0.0451	0.0409	0.0463	0.0504	0.0534	0.0599	0.0724
4	0.0520	0.0490	0.0561	0.0639	0.0713	0.0773	0.0914
6	0.0594	0.0564	0.0670	0.0752	0.0806	0.0880	0.1039
8	0.0643	0.0625	0.0754	0.0816	0.0861	0.0907	0.1121
10	0.0688	0.0677	0.0798	0.0864	0.0922	0.0975	0.1107
20	0.0776	0.0819	0.0975	0.1038	0.1129	0.1189	0.1253

insiders. This allows him to jump immediately at the top of the queue, while old employees shift downward. As a consequence, old insiders will be the first to be fired as the firm will eventually cut its labor demand when facing adverse market conditions. Therefore, longer queues due to lower search costs in the labor market entail shorter labor contracts, on average. Notice that this effect is reversed for very high M, as the economy approaches full employment.

By the same token, more competition on the goods market (i.e., higher Z) implies that firms have a lower probability to increase their labor demand on a period by period basis, so that the average lengths of contracts is positively related to Z, for any given M.

4. A computational experiment of fiscal policy

The results for the agent-based computational model detailed above seem to give some credit to the idea that the complexity approach is a workable way to think about macroeconomic issues. The interplay of locally interacting agents and of simple institutional mechanisms allows the system to generate regular aggregate dynamics and leads to the emergence of several macroeconomic stylized facts. Notice that jointly explaining all the facts we showed by means of a standard general equilibrium model would imply the introduction of a set of highly unrealistic assumptions, first of all that of requiring a supra-natural cognitive capacity to the representative agent. On the contrary, the agent-based methodology we employ builds precisely on – and exploits the – existence of heterogeneity among rule-of-thumb agents. This point leads naturally to a closely related issue. As discussed in Kirman (1992), the representative agent device is particularly badly equipped for analyzing redistributive policies, while the possibility to keep track of the dynamical path of individual entities is one of the major strengths of agent-based techniques.

Thus, in this section we present the results of a computational experiment aimed at exploring the qualitative consequences of alternative fiscal policies in our model economy. The baseline model of Section 2 is augmented with a public agency, whose role consists in collecting tax revenues from firms and to subsequently redistribute them during the same period. We will consider two alternative scenarios. In Treatment 1, the agency applies a flat tax rate (τ) to firms' profits; tax revenues are then redistributed as unemployment benefits. In Treatment 2, on the contrary, tax revenues are redistributed to firms on a per-capita basis as cash incentives for investing in R&D. All the other characteristics of the model are unchanged. From a theoretical viewpoint, Treatment 1 is aimed at capturing the influence on macroeconomic performance of policies aimed at fostering aggregate demand, while Treatment 2 concentrates on the effects of supply-side policy interventions.

Fig. 10 clearly shows that in Treatment 1 a redistributive policy is harmful for aggregate growth. The average growth rate of aggregate output decreases almost monotonically as the profit tax rate is increased from 0% to 99%. For sufficiently high tax rates, the GDP growth becomes negative.

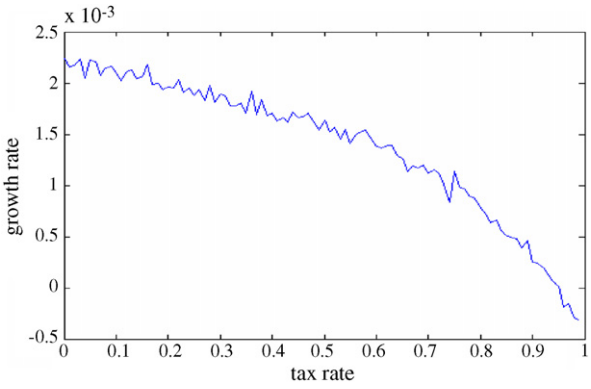

Fig. 10. Treatment 1: relationship between tax rate and GDP growth rate.

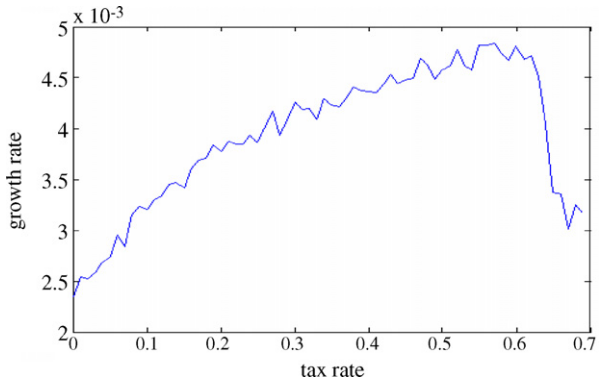

Fig. 11. Treatment 2: relationship between tax rate and GDP growth rate.

The reason for the negative performance of redistributive schemes aimed at fostering aggregate spending is that the growth-enhancing effect of higher demand implied by Eq. (5) is more than offset by lower investment in new technological opportunities, as the fraction of firms' after-tax profits destined to R&D investments shrinks proportionally to increases in the tax rate.

A different picture emerges as we move to Treatment 2. The graph in Fig. 11 shows that a redistributive scheme aimed at promoting investment in more productive technologies affects positively the growth rate of aggregate output. The effect is reversed as the tax rate is raised to more than 60%, however.

The turning point in the graph of Fig. 11 takes place as the growth-enhancing effect associated to the redistribution of resources aimed at R&D investment is more than counteracted by the increase in systemic financial fragility associated to lower individual net worth. This fact is reflected in a boost of bankruptcies as the tax rate rises slightly above 60% (Fig. 12(d)) on the one hand, and in the negative *modal* firms growth rate (i.e., the location parameter of the Laplace firms growth rate distribution) as the tax rate approaches that same level (Fig. 13), on the other one. Notice that the tax rate is almost neutral on the modal firms growth rate as it is increased from 0% to 40%.

The 60% tax rate threshold causes qualitative changes also in the dynamics of the rate of unemployment, whose increase becomes steeper (Fig. 12(a)). When the growth rate of aggregate output declines, the increase in unemployment caused by higher investment in labor-saving technological improvements is coupled to that caused by a lower labor demand as the growth rate of aggregate activity slows down. The same mechanism is at the root of the discontinuity in the relations for the rate of vacancies and the rate of inflation (Fig. 12(b) and (c)).

To fully appreciate how redistributive policies aimed at promoting R&D affect macroeconomic performance, the relation between the level of taxation on profits and of private effort in R&D is further explored. Recall that all the results from the fiscal experiment shown so far have been obtained for $\sigma = 0.05$. In additional simulation runs, the same setting as in Treatment 2 is employed to perform a comparison between two alternative scenarios. First, we allow firms to increase the fraction of profits they allocate to R&D (scenario A). Second, we set the parameter σ to zero (no private R&D investment), and merely allow each firm to invest in R&D the cash it receives as the fiscal agency redistribute fiscal revenues on a per-capita basis. In this way it is possible to appreciate the role of heterogeneity in levels of R&D investment on aggregate growth. Fig. 14 shows that a public intervention in terms of horizontal equality redistribution yields a lower performance than the private-R&D solution, for any rate of private investment and any rate of taxation. After the tax rate exceeds the level $\tau = 64\%$, the growth gap increases sensibly.

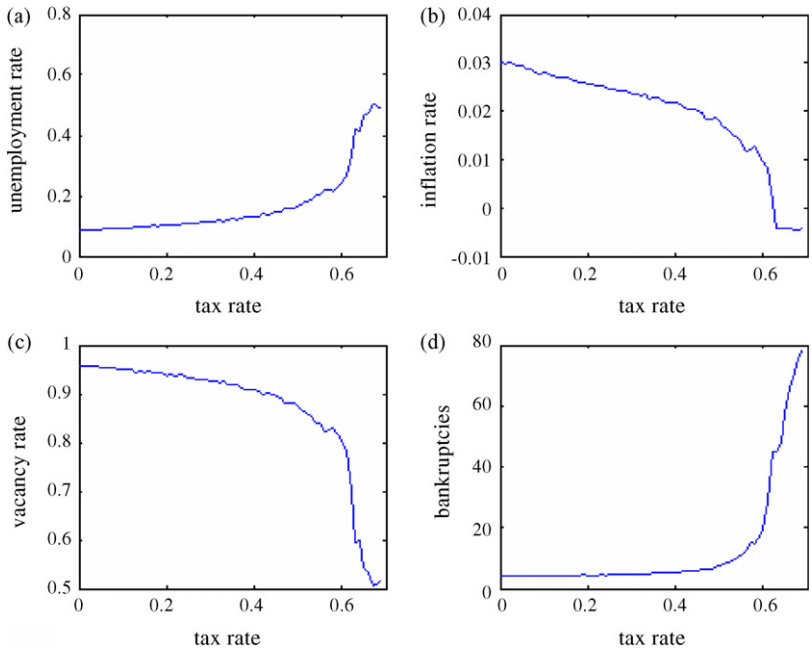

Fig. 12. (a–d) Treatment 2: relationship between tax rate and other relevant variables.

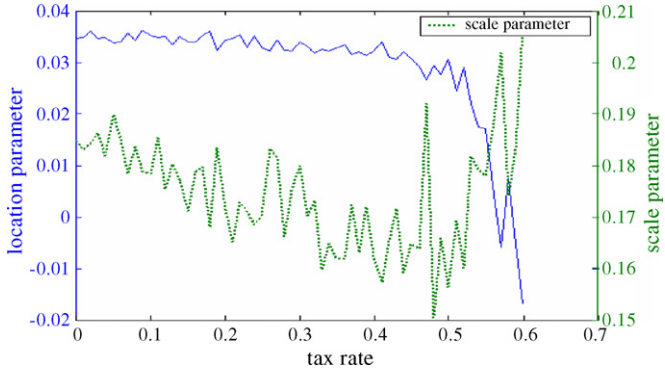

Fig. 13. Treatment 2: shape of the firms growth rate distribution for different level of the tax rate.

As a final exercise, we investigate interactions between private and public contributions to R&D investment. We find that the gain of redistributing tax revenues as transfers to firms is a function of the level of private effort in searching for more productive technologies. Some results of a simulation exercise are reported in Table 3. As the fraction of profits allocated to R&D by firms increases, the net benefit of collecting revenues as a proportion of profits and to redistribute them on a per-capita basis decreases. The turning point tax rate – that is, the threshold tax rate over which macroeconomic performance deteriorates – decreases from roughly 60% when private firms invest in R&D a fraction of their profits comprised between 5% and 15%, to 55% as firms increase their R&D investment to 25% of their profits.

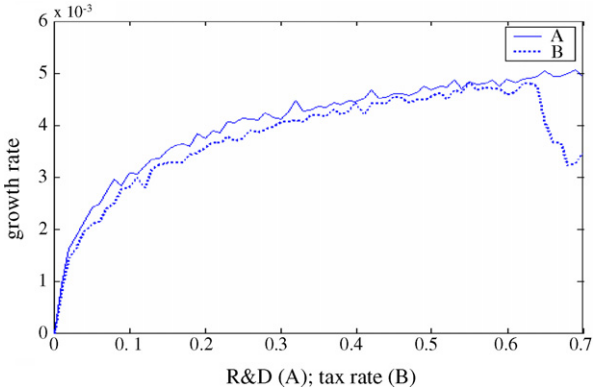

Fig. 14. Treatment 2: relationship between R&D and GDP growth rate in two alternative scenarios: (A) private R&D; (B) public incentives to R&D.

Table 3
Behavior of major economic variables for different levels of tax and private R&D rates

	Tax rate				
	5%	15%	30%	45%	60%
$\sigma = 0.05$					
Growth rate	0.0027	0.0034	0.0043	0.0045	0.0048
Inflation rate	0.0292	0.0269	0.0235	0.0199	0.0097
Unemployment rate	0.0888	0.0981	0.1143	0.1445	0.2421
Vacancy rate	0.9550	0.9430	0.9267	0.8949	0.8054
Bankruptcy rate	0.0407	0.0421	0.0477	0.0597	0.19319
$\sigma = 0.15$					
Growth rate	0.0037	0.0040	0.0043	0.0047	0.0044
Inflation rate	0.0268	0.0254	0.0229	0.0181	−0.004
Unemployment rate	0.0924	0.1016	0.1212	0.1599	0.4262
Vacancy rate	0.9511	0.9404	0.9181	0.8778	0.5861
Bankruptcy rate	0.0419	0.0433	0.0489	0.0640	0.4821
$\sigma = 0.25$					
Growth rate	0.0041	0.0042	0.0045	0.0049	0.0032
Inflation rate	0.0255	0.0241	0.0220	0.0154	−0.004
Unemployment rate	0.0961	0.1066	0.1270	0.1823	0.5136
Vacancy rate	0.9467	0.9344	0.9122	0.853	0.4908
Bankruptcy rate	0.0423	0.0452	0.0516	0.0789	0.6496

5. Conclusion

In this paper we present an agent-based computational model in which the macroeconomy is modeled as a "[. . .] complex system where aggregate behaviour is determined by the complicated interaction among individuals operating in a simple way at the micro level" (Kirman, 1999, p. 5). A large number of heterogeneous firms and workers interact through randomly determined dispersed trades in the markets for labor and for the final consumption good. Agents update their

choice variables by means of routinized strategies, and a selection mechanism is at work which forces the exit of unprofitable firms.

Simulations reveal that a fully decentralized economy can easily display regular behavior on a macroeconomic scale, even if the Walrasian auctioneer does not operate. Our economy replicates some interesting stylized facts both at the industrial dynamics level and at a macroeconomic scale. We subsequently employ our baseline model as a computational laboratory to perform an experiment of fiscal policy. A government agency is allowed to collect resources form firms through a flat tax rate, and to redistribute them on a per-capita basis to firms or to unemployed workers. In our virtual economy, using tax revenues to subsidize aggregate demand returns a bad result in terms of aggregate activity's growth rate. The opposite holds if tax revenues are employed to subsidize R&D investment.

Obviously, our results are merely qualitative and strictly related to the assumptions we made at the start regarding the degree of individual rationality, institutions and selection mechanisms. Much more important is the methodological message implied by our exercise: we claim that agent-based techniques represent a practical and feasible approach to do positive complexity-inspired macroeconomic research, while normative proposals can be tested in a controlled environment by using agent-based virtual economies as computational laboratories. From this point of view, future research on the framework introduced in this paper will be devoted to exploit several alternative assumptions on the tax code (e.g., flat rate versus proportional taxation; labor income taxation) and on the redistributive choices of the government (e.g., picking-the-winner industrial policy; *selective* government expenditure).

Acknowledgments

We thank two anonymous referees for a number of valuable suggestions. The usual disclaimer applies.

References

Amaral, L., Buldyrev, S., Havlin, S., Leschhorn, H., Maas, P., Salinger, M., Stanley, E., Stanley, M., 1997. Scaling behavior in economics. I. Empirical results for company growth. Journal de Physique 7, 621–633.

Bottazzi, G., Dosi, G., Lippi, M., Pammolli, F., Riccaboni, M., 2001. Innovation and corporate growth in the evolution of the drug industry. International Journal of Industrial Organization 19, 1161–1187.

Delli Gatti, D., Gaffeo, E., Gallegati, M., Palestrini, A., 2005. The apprentice wizard: monetary policy, complexity and learning. New Mathematics and Natural Computation 1, 109–128.

Epstein, J., Axtell, R., 1996. Growing Artificial Societies: Social Science from the Bottom Up. Brookings/MIT Press.

Fagiolo, G., Dosi, G., Gabriele, R., 2004. Matching, bargaining, and wage setting in an evolutionary model of labor market and output dynamics. Advances in Complex Systems 7, 157–186.

Feller, W., 1957. Introduction to Probability Theory and its Application. Wiley, New York.

Gode, D., Sunders, S., 1993. Allocative efficiency of markets with zero-intelligence traders: markets as a partial substitute for individual rationality. Journal of Political Economy 101, 119–137.

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

Kirman, A., 1997. The economy as an interactive system. In: Arthur, B., Durlauf, S., Lane, D. (Eds.), The Economy as an Evolving Complex System II. Addison-Wesley, Reading, pp. 491–532.

Kirman, A., 1999. Interaction and markets. In: Gallegati, M., Kirman, A. (Eds.), Beyond the Representative Agent. Edward Elgar, Cheltenham, pp. 1–44.

Lane, D., 1993. Artificial world and economics, Part I & II. Journal of Evolutionary Economics 3 (89-108), 177–197.

Leijonhufvud, A., 1993. Towards a not-too-rational macroeconomics. Southern Economic Journal 60, 1–13.

Nelson, R., Winter, S., 1982. An Evolutionary Theory of Economic Change. Harvard University Press, Cambridge.
Stanley, M., Buldyrev, S., Havlin, S., Mantegna, R., Salinger, M., Stanley, E., 1995. Zipf plots and the size distribution of firms. Economics Letters 49, 453–457.
Stiglitz, J., 1989. Imperfect information in the product market. In: Schmalensee, R., Willig, R. (Eds.), Handbook of Industrial Organization. North-Holland, Amsterdam, pp. 769–847.
Tesfatsion, L., 2002. Agent-based computational economics: growing economies from the bottom up. Artificial Life 8, 55–82.