

Chapter 3

The *BAM* Model at Work

The classical theorists resemble Euclidean geometers in a non-Euclidean world who, discovering that in experience straight lines apparently parallel often meet, rebuke the lines for not keeping straight—as the only remedy for the unfortunate collisions which are occurring. Yet, in truth, there is no remedy except to thro over the axiom of parallels and to work out a non-Euclidean geometry.

JOHN M. KEYNES

The General Theory of Employment, Interest and Money, p. 16

3.1 BAM at Work

In this chapter we develop a prototype bottom-up macroeconomic (BAM) model,⁴⁰ which epitomizes the key features at the root of a series of computational investigations of macroeconomic processes conceived as complex adaptive systems (CATS), as recently performed by our research group. Other exemplifications of the CATS approach can be found in Delli Gatti *et al.* (2005), Gaffeo *et al.* (2007) and Russo *et al.* (2007).

In Sect. 3.2 we list the main ingredients of the BAM framework: agents, markets and trading processes. In Sect. 3.3 we carefully describe the *sequence* of actions and interactions which occur in the economy under scrutiny. A pervasive and recurrent feature of this sequence is the *search* process which goes on in each of the market considered: households search for a job on the labor market and for consumption goods on the goods markets, while firms search for a bank loan on the credit market. Search is costly, so that each searching agent can visit only a finite number – i.e., a subset – of potential “providers”: firms which provide job opportunities on the labor market, firms which offer consumption goods on the

⁴⁰ A streamlined version of the present model and a succinct discussion of its features can be found in Delli Gatti *et al.* (2008).

goods market, banks which provide loans on the credit market. Each period the identity of the (finite number of) providers the searcher can visit changes partially at random, so that the network structure is continuously evolving over time, even if the number of “links” (providers) per “node” (searcher) is constant.

All *matching* processes occur in a completely *decentralized* setting. In our framework there is not any centralized auctioneer at work, so that actual transactions can well occur at out-of-equilibrium prices. Moreover, we do not resort to any exogenous “matching function”, a deterministic device which plays the crucial role of coupling agents on the two sides of the labor market in mainstream search-and-matching models of equilibrium unemployment. The main advantage of the BAM model is that one can directly simulate the above-mentioned myriad of dispersed interactions by means of an algorithmic representation, instead of recurring to an aggregate proxy of the behavior of customers trying to buy and of suppliers trying to sell. Sects. 3.4, 3.5 and 3.6 are therefore devoted to an in-depth discussion of the working of the search-and-matching processes in the market for labor services, for bank loans and for consumption goods, respectively.

In Sect. 3.7 we focus on the macroeconomic role of bankruptcy. Financial conditions of firms and banks, in fact, play a crucial role on all the markets considered, either directly or indirectly. When a firm’s or bank’s financial fragility reaches a critical point, i.e. when its net worth turns negative, that economic unit goes bankrupt. Bankruptcy therefore is the most straightforward device to introduce an *exit* mechanism in our virtual economy. An entry process occurs in parallel with exit, so that in our model firms’ demography is fully taken into account.

The (baseline) model described so far is based upon the assumption of a *constant* labor productivity and is capable of reproducing the irregular “short run” fluctuations of aggregate output which is actually characterizing real world economies (as will be shown in Subsect. 3.9.1). In Sect. 3.8 we further introduce an endogenous mechanism for the determination of labor productivity, which links productivity to investment in R&D and the latter to profits. In this case it is easy to show that the model displays both growth and irregular fluctuations. This is the reason why we label this extension the “growth+” model.

Sect. 3.9 is devoted to an analysis of simulations’ results. Since the empirical validity of a model can be assessed comparing theoretical predictions with a selected set of *explananda*, we believe that at least two issues are of key importance in evaluating the empirical success of the BAM model.

First of all, the BAM model should be able to replicate the tendency of the macroeconomy to *self-organize* most of the times, but also to occasionally display severe *coordination failures* so that, say, a great depression can occur because of the transmission of an idiosyncratic shock, i.e. in the absence of a major negative aggregate shock. Macroeconomic models of the usual sort, on the contrary, usually exhibit either regular behavior all the time (whenever a stable equilibrium exists), or permanent degenerate behavior (whenever the previous condition does not hold). In the standard literature the second scenario is discarded a priori so that (short-lived) fluctuations can occur only if an aggregate shock hits the macro-economy and displaces it from its stationary equilibrium.

Second, the BAM model should be able to replicate, at least qualitatively, one or more of the stylized facts of macroeconomic importance that are known to hold for most of the industrialized countries. In particular, we are interested in building a virtual environment able to capture the emergence of aggregate regularities as the result of decentralized interactions of a multitude of heterogeneous agents.

Notice that these criteria for the empirical corroboration of predictions from the BAM model are mainly *qualitative*. A different but complementary strategy consists in adopting *quantitative* methods for *ex-post validation*. We defer to Chap. 4 an exercise in ex-post validation of the model. In the rest of this chapter we focus instead on the qualitative measures just outlined, and in particular we assess the performance of the BAM model in producing:

- a non degenerate dynamics of the aggregate variable of interest (output) punctuated by sudden crises;
- emergent macroeconomic regularities, such as correlated paths of labor productivity and the real wage, Phillips and Beveridge curves and the Okun's law;
- co-movements among aggregate variables and leads-and-lags correlations.

Going to the details, in Subsect. 3.9.1 we discuss results concerning the baseline scenario, while Subsect. 3.9.2 is devoted to the output of simulations of the growth+ model. A check on the robustness of these findings as regards variations in the parameter constellation is postponed to Sect. 3.10. Before it, though, in Subsect. 3.9.3 we perform an assessment exercise by means of actual and simulated data in order to compare the BAM methodological approach to that currently used in modern macroeconomics (DSGE). Finally, Subsect. 3.9.4 describes one of the many possible extensions of the model, with the aim of showing the degree of flexibility of the BAM model.

We hope that the evidence reported in Sect. 3.9 and 3.10 will be sufficient to convincingly convey the belief that identifiable aggregate regularities consistent with the stylized facts may easily appear from the complex interactions of heterogeneous adaptive adjustments on different margins, technological innovation, limited search and out-of-equilibrium decentralized transactions on three interrelated markets.

3.2 The Environment

In order to build an agent-based model, three main ingredients are necessary.

1. The list of the *agents* that populate the model. Generally, pre-determined subsets of the population identify groups or *classes* of agents characterized by specific macroeconomic roles.
2. The *structure* of each agent, which consists of:
 - a list of the *state variables* that describe the agent in every period of the time horizon considered (which translates into a step of the simulation). The “snapshot” of the condition of the agent in a given period, i.e. the vector of

levels of the state variables concerning the specified agent in that period, is the *internal state* of the agent;

- a list of the possible *actions* (the levels of the control variables) that agents can perform. Actions will affect not only their internal state but also the internal state of other agents.

Agents belonging to the same class have the same macroeconomic role and have similar structures. They may be characterized, however, by a specific level of one or more microeconomic (state or control) variables. This allows to preserve individual specificity also within each class.

3. The network of interactions that links agents within the group and among groups. Among group interactions typically occur in virtual or geographically characterized markets.

As to point 1), our model describes a sequential closed economy populated by a finite number ($I + J + B$) of agents grouped into three classes:

- firms, indexed by $i = 1, \dots, I$;
- workers/consumers, indexed by $j = 1, \dots, J$;
- banks, indexed by $k = 1, \dots, K$.

As to point 2), each agent is characterized both by a set of state variables (e.g. productivity, net worth), and by a set of control variables (e.g. notional prices and quantities). Finally, as to point 3), agents undertake decisions at discrete times $t = 1, \dots, T$ on three markets:

- a market for a homogeneous non-storable consumption good;
- a market for labor services;
- a market for credit (bank loans).

Since agents' decision making processes are constrained by imperfect/incomplete information and by limited computational capabilities – a condition which can be labeled with the evocative term of bounded rationality (Simon, 1997; Kahneman and Tversky, 1981) – we assume that actions are not the outcome of an optimization process, but they are chosen *adaptively* according to rules of thumb buffeted by idiosyncratic random disturbances.

Markets are characterized by continuous decentralized search and matching processes (the so-called *procurement process* in the parlance of Tesfatsion [2005]), which imply individual, and *a fortiori* aggregate, out-of-equilibrium dynamics. Even in the absence of a centralized market-clearing mechanism, the economy shows a tendency to *self-organize* towards a *spontaneous order* which is however characterized, depending on the market and the time horizon, by persistent involuntary unemployment, unsold production or excess demands, and credit rationing. While in the standard macroeconomic theory these phenomena are treated as “pathologies” – i.e., departures from a first-best scenario due to imperfections of one sort or another –, in our framework they are emerging properties – i.e., “physiological” outcomes – of the macroeconomy.

The modeling strategy of the BAM framework is built on two pillars. First, the rules of individual behavior and market transactions (that we translate into algorithmic language) are inspired – whenever possible – to the evidence available from survey studies conducted by asking households and business people how they actually behave. Where several competing theories are available, we conform to the dull version of the Occam’s Razor principle known as KISS.⁴¹ Second, as discussed at length above, we do not impose any centralized solving mechanism. Instead, we let the system of adaptive interacting agents evolve autonomously towards self-organizing configurations: in other words, we will not impose the exogenous choice of any equilibrium, but we allow the endogenous *formation* of one of them, if it exists.

3.3 The Sequence of Events

The sequence of events runs as follows:

1. Each operating firm decides on the amount of output to be produced (hence, the amount of labor to be hired) and the price to be charged according to expected demand for consumption goods. Expectations of future demand are updated adaptively, i.e. they are formed on the basis of the firm’s past experience.⁴²
2. A fully decentralized labor market opens. Firms post their vacancies at a certain offered wage, and unemployed workers contact a given number of randomly chosen firms to get a job, starting from the one that offers the highest wage. Firms then have to pay the wage bill in order to start production. Labor contracts expire after a finite number of periods θ . A worker whose contract has just expired applies first to her last employer.

⁴¹ Several interpretations of the acronym KISS circulate, most of them overlapping. The one we prefer is *keep it simple, stupid!*

⁴² Behavioral rules represent by construction the process of *adaptation* of the agent’s actions to changes of the environment. In a behavioral setting therefore expectations formation can be modeled quite straightforwardly as an adaptive scheme: Firms form expectations on future demand only on the basis of the past history of production (which is demand constrained). This adaptive mechanism is inefficient from a rational expectations (RE) viewpoint. In fact, if agents cast adaptive expectations, sooner or later they will incur in systematic errors. In a RE setting, on the contrary, agents are rational, i.e. they are able to elicit all the necessary information – not only the past history of the variable in question – and process it in such a way as to make only random errors which cancel out in the aggregate. Notice, however, that from the statistical point of view REs are conditional expectations of the system’s data generating process (DGP). As such they are inherently liable to errors. In any given situation, according to the RE theory agents endowed with rational expectations should not make mistakes *on average*, but in practice they do. In order for agents to assess whether their specification of the DGP’s conditional mean is right or not, the situation must be repeated over time in such a way as to allow agents to learn and update their expectation formation with the help of an “error correction” procedure. At the same time, however, the DGP is likely to change as well, frustrating agents’ efforts to be “rational”.

3. If internal financial resources (net worth) are in short supply with respect to the wage bill – i.e. if there is a financing gap – the firm can access a fully decentralized credit market. Borrowing firms contact a given number of randomly chosen banks to get a loan, starting from the one which charges the lowest interest rate. Each bank sorts the borrowers’ applications for loans in descending order according to the financial soundness of firms, and satisfy them until all credit supply has been exhausted. The contractual interest rate is calculated applying a mark-up (which is itself a function of financial viability) on an exogenously determined baseline interest rate. After the credit market is closed, if financial resources – both internal and external – are not enough to pay for the wage bill of the population of workers, some workers remain unemployed or are fired.
4. Production takes one time period, regardless of the scale of production/firm’s size.
5. After production is completed, the market for goods opens. Firms post their offer price, and consumers contact a given number of randomly chosen firms to purchase goods, starting from the one which posts the lowest price. If a firm ends up with excess supply, it gets rid of the unsold goods at zero costs. The good in fact is perishable and cannot be stored in a warehouse to be sold in the future.
6. Firms collect revenues and calculate gross profits. If gross profits are high enough, they “validate” debt commitments, i.e. firms pay back both the principal and the interest to the bank. If net profits are positive, firms pay dividends to the owners. In a “growth+” variant of the present model (to be discussed in Sect. 3.8), firms invest a fraction of net profits in R&D in order to increase their productivity before distributing dividends.
7. Earnings after interest payments and dividends are retained profits, which are employed to increase net worth. Net worth at the end of a period, in fact, is the sum of all retained profits accumulated in the past. Firms and banks are financially viable – and therefore survive – if their net worth is positive. If, on the contrary, net worth is negative, they go bankrupt, shut down and exit the market. Lenders, therefore, have to register a bad debt (non-performing loan).
8. A string of new firms/banks equal in number to the bankrupt ones enters the market. Their size at entry is smaller than the average size of exiting agents.

3.4 The Labor Market

The i -th firm carries on production by means of a constant return to scale technology, with labor L_{it} as the only input:

$$Y_{it} = \alpha_{it} L_{it}, \alpha_{it} > 0 \quad (3.1)$$

where α_{it} is labor productivity. While in this section productivity is considered as a parameter, in general it can change according to a simple rule of technological

updating, which in turn depends on profitability and the availability of financial resources to carry on R&D expenditure. Heterogeneous financial conditions, therefore, imply heterogeneous productivity levels. The case of an endogenous, financially driven, productivity will be dealt with in the next section.

From equation (3.1), it follows that the desired workforce – i.e. the demand for labor, L_{it}^d , expressed as the number of workers the firm is allowed to hire – is simply given by:

$$L_{it}^d = \frac{Y_{it}^d}{\alpha_{it}}, \quad (3.2)$$

where Y_{it}^d is the desired level of production. In other words, the desired workforce represents the labor requirement that must be fulfilled to reach the desired scale of production. We will show in Sect. 3.6 how the latter is determined.

At the beginning of period t , each firm advertises the opening of vacant positions, and the associated offered wage. In order to determine the effective number of vacancies, note that at the beginning of period t the i -th firm is endowed with an *actual* workforce equal to $L_{it}^o = L_{it-1} - \hat{L}_{it-1}$, where L_{it-1} represents workers employed at the firm in $(t-1)$, while \hat{L}_{it-1} is the number of workers whose labor contract has just expired. If the desired labor force is larger than the actual one, the firm creates a number of vacancies equal to $V_{it} = L_{it}^d - L_{it}^o$. Hence, the amount of open vacancies is:

$$V_{it} = \max(L_{it}^d - L_{it}^o, 0). \quad (3.3)$$

Workers with an active contract can be fired only if the firm's funds (both internal and external) are not enough to pay for the desired wage bill.

We assume that workers supply inelastically one unit of labor per period, and that only unemployed workers can search for a new job. In other words, we rule out on-the-job search. Each unemployed worker sends M applications to as many firms. If her contract has just expired, she applies first to the firm in which she worked in the previous period and, after that, she will send the remaining $M-1$ applications to as many firms chosen at random. New unemployed workers are therefore characterized on one hand by a sort of loyalty to their last employer, and on the other hand by a desire to insure themselves against the risk of unemployment by diversifying the portfolio of hiring opportunities. Of course, loyalty to the past employer does not make any sense if the worker has just been sacked, or if she has lost her job because of a bankruptcy. In all these cases, as well as when the worker is actually living a long spell of unemployment, she simply sends M applications to as many randomly chosen potential employers.

Once the offered contractual terms of vacant positions have been publicized to all applicant workers, each worker chooses to enter a settlement stage only with the firm offering the highest wage, out of the M firms she visited. Contracts are closed sequentially according to an order randomly chosen at each time step. Since each worker is allowed to sign one labor contract per period and the labor market microstructure is completely decentralized, serious “coordination failures” could

arise due to two different reasons. First, the number of unemployed workers actually searching for a job in the aggregate does not necessarily correspond to the number of vacancies, so that aggregate excess supply or demand for labor is a frequent market outcome. Second, some firms – typically those that offer relatively high wages – may experience an excess of requests for employment with respect to actual vacancies, while some other firms – mainly those that post relatively low wages and hire workers late in the sequence – may end up in the opposite situations and some vacancies may remain unfilled.

When hired, a worker is asked to sign a contract that determines her nominal wage for a fixed number of periods. The contractual wage offered by firm i in period t is determined according to the following rule:

$$\begin{aligned} w_{it}^b &= \max(\hat{w}_t, w_{it-1}) \text{ if } V_{it} = 0 \\ w_{it}^b &= \max(\hat{w}_t, w_{it-1}(1 + \xi_{it})) \text{ if } V_{it} > 0 \end{aligned} \quad (3.4)$$

where \hat{w}_t is the minimum wage (set by a mandatory law), while w_{it-1} is the wage offered to the cohort of workers employed the last time the firm hired. ξ_{it} is an idiosyncratic shock uniformly distributed on the interval $(0, h_\xi)$. The minimum wage is periodically revised upward, in order to catch up with inflation. In other words, wages are fully indexed. Wages set in the past that happen to fall below the current minimum wage are automatically aligned to the latter.⁴³ Workers paid the minimum wage therefore are fully insured against eroding purchasing power due to inflation. The indexation of the minimum wage may hamper the capability of firms to seek and preserve profitability, in a sort of wage-price spiral. For instance, in periods of *tight* labor market, firms that are expanding their workforce hiring new workers increase their price to preserve profit margins. Higher prices, in turn, drive the minimum wage up, offsetting the efforts of the firms. The process works in the opposite direction when the labor market is *loose*.

The design of the labor market we choose is somehow consistent with the findings reported by numerous surveys of firms' wage-setting policies. First, there is clear evidence of nominal wage downward rigidity. Firms are particularly reluctant to cut nominal wages even during recessions because they are afraid that lower wage rates would increase turnover and decrease labor effort (Campbell and Kamlani, 1997; Bewley, 1999). Second, downward rigidity is observed also for the salary of the newly hired workers, probably for reasons of perceived equity (Bewley, 1999). Akerlof and Shiller (2009) interpret this downward rigidity of the nominal wage as one instance of money illusion.

⁴³ In simulations we set the duration of contracts θ to 8 periods, while the minimum wage is revised every 4 periods. If we assume that one simulation period corresponds to a quarter, this means that labour contracts last two years, while the minimum wage is revised annually.

3.5 The Credit Market

At the beginning of period t , the generic firm i is endowed with an amount of retained past profits or net worth equal to A_{it} (see equation 3.12 below). If its desired wage bill W_{it} is larger than its net worth, the firm looks for a bank loan, $B_{it} = W_{it} - A_{it}$. The demand for credit therefore is simply given by:

$$B_{it} = \max(W_{it} - A_{it}, 0). \quad (3.5)$$

Due to transaction costs, the search for loans on the part of the firm is restricted: each firm can in fact apply for a loan only to a fixed number $H < K$ of banks. In a sense, if we extend to the credit market the conceptual apparatus originally introduced for the analysis of search and matching on the labor market, these are “credit applications” coming from agents in need of external finance.

Each time period t , the k -th bank will extend a total amount of credit C_k equal to a multiple of its equity base: $C_{kt} = E_{kt}/\nu$, where $0 < \nu < 1$ can be interpreted as a *capital requirement* coefficient. The reciprocal of ν therefore represents the *maximum allowable leverage* for the bank. For simplicity, we assume for the moment that the capital requirement coefficient is determined by a regulatory authority, and is uniform across banks. If we apply to the credit market the conceptual apparatus we used for the analysis of search and matching processes in the labor market, C_k represents the amount of “credit vacancies” posted by the k -th bank.

Banks advertise credit opportunities consisting of credit vacancies and the associated “price”, i.e. the nominal interest rate. We assume that a generic bank k offers to firm i a standard single-period debt contract, which defines an interest rate r_{it}^k and the corresponding repayment schedule:

$$\begin{aligned} B_{it} (1 + r_{it}^k) & \text{ if } A_{it+1} > 0 \\ R_{it+1} & \text{ if } A_{it+1} \leq 0 \end{aligned} \quad (3.6)$$

where R_{it+1} is the amount the bank succeeds in retrieving in case the borrower’s net worth becomes insufficient, i.e. if the firm goes bankrupt. To be more precise, the contractual interest rate offered by bank k to firm i is determined as a mark-up over a policy rate set by the central monetary authority \bar{r} :

$$r_{it}^k = \bar{r} (1 + \phi_{kt} \mu(\ell_{it})). \quad (3.7)$$

The mark-up is a function:

- of the specificity of the k -th bank, modeled as random variations in its operating costs and captured by the random variable ϕ_{kt} , an idiosyncratic shock uniformly distributed on the interval $(0, h_\phi)$;
- of the financial fragility of the borrower, captured by the term $\mu(\ell_{it})$, $\mu' > 0$, where $\ell_{it} = \frac{B_{it}}{A_{it}}$ is the borrower’s leverage.

The last term implies that the mark-up the bank charges over the policy rate reflects a risk premium increasing with the financial fragility of the borrower.

Equation (3.7) can be interpreted in the light of the theory of the “external finance premium” pioneered by Bernanke and Gertler (1989, 1990). In the presence of ex post asymmetric information and costly state verification, the higher the borrower’s financial fragility, the more frequent the auditing activity of the bank should be, and the higher the interest rate charged to the borrower. Alternatively one can think of (3.7) as the reduced form of a model in which a commercial bank can insure against potential losses due to lending by borrowing, at least to a certain extent, from a central bank acting as a lender of last resort. The policy rate, in this case, is the rate at which the central bank refinances the commercial bank. A by-product of this interpretation is that in principle firms can always find external funds and can be credit rationed only when total credit supply is small (i.e. v is large), since banks can obtain additional funds from the central monetary authority and price-discriminate among borrowers, via interest rates, according to their quality.

A firm which needs external finance can explore a segment of the market for bank loans by randomly picking H banks out of the population of K banks. Once the terms of the credit opportunities at the H banks have been revealed, the firm chooses the bank offering the lowest interest rate. We assume that the demand for credit is divisible, so that if the most preferred bank is in short supply of credit the firm can resort to the remaining $H-1$ banks. If total resources are still not sufficient to pay for the wage bill, the firm will be allowed to fire redundant workers at zero costs.

Contract settlements are closed sequentially, according to an order randomly chosen at each time step. Since the credit market microstructure is completely decentralized, once again serious “coordination failures” could arise. First of all, the amount of credit demanded in the aggregate does not necessarily correspond to the credit supply. Second, some banks may experience an excess of demand for loans with respect to “credit vacancies” – generally those banks that post relatively low interest rates – while some other banks may end up in the opposite situation and some vacancies may remain unfilled, especially in the case of banks which post relatively high interest rates. Some firms will therefore be rationed.

3.6 The Market for Consumption Goods

At the beginning of each period, the i -th firm adjusts its control variables, i.e. the price or the quantity supplied, to adapt to changing business conditions. In spite of the good being homogeneous, asymmetric information and search costs imply that consumers may end up buying from a firm even if its price is not the lowest. It follows that the conditions for perfect competition are not satisfied, and the law of one price does not apply (Stiglitz, 1989). Each firm has a certain degree of market power on its own local market.

For simplicity, we assume that a firm can change either the price or the quantity, but not both of them at the same time. In other words, the strategies consisting in “changing the price” and in “changing the quantity” are mutually incom-

patible. This assumption is based on the evidence of survey data on price and quantity adjustment of firms over the business cycle (Kawasaki *et al.*, 1982; Bhaskar *et al.*, 1993).

For expositional simplicity we assume that each strategy is ex-ante equally likely. In principle, however, we could attach a probability to each strategy which could be calibrated on real data. For instance, the available evidence suggests that liquidity constrained firms – i.e. firms with a limited cash-flow – quantity adjustments are more likely during recessions than during booms, whereas the reverse is true for price adjustments; i.e. constrained firms are less likely to cut prices in recessions.

In our model, the adaptation of each strategy depends on signals coming from the internal condition of the firm and/or from the market environment. The information set relevant for price or quantity adjustment of the i -th firm at time t consists of two components:

- The level of excess demand/supply in the previous period. Excess supply is signaled by the accumulation of an inventory of unsold goods ($S_{it-1} > 0$). Since the good is perishable, this inventory cannot be carried over to t and therefore it is *temporary*. Moreover, we assume that the firm can get rid of the inventory at no cost. If demand happens to be equal to supply or if there is excess demand, there will be no inventory ($S_{it} - 1 = 0$). In the former case, in principle, the firm has an incentive to reduce the price or reduce the quantity – we will be more precise momentarily – while in the latter case there is room for a price increase or an increase in quantity. There is a lower bound to a reduction of the price which is represented by the minimum price the firm has to charge to cover average costs.
- The deviation of the individual price from the average price $P_{it-1} - P_{t-1}$ during the last transaction round. If this deviation is positive (negative), the firm recognizes that it is charging a price higher (lower) than its competitors and therefore may be induced to reduce (increase) the price or the quantity to avoid (facilitate) a massive migration of consumers in favour of (from) its rivals. Also in this case a reduction of the price is bounded from below: the price cannot be lower than the minimum price the firm has to charge to cover average costs.

Internal conditions (i.e. the level on the temporary inventory or the individual price) are private knowledge, while the aggregate price is common knowledge.

In principle we have four cases. As we said above, we assume that price changes and quantity changes cannot occur simultaneously. Therefore, we associate *either* a price change *or* a quantity change to each case.

- a) In case inventories are positive (excess supply) and the individual price is high with respect to the average, the firm will reduce the price (until the lower bound is reached) keeping the quantity unchanged.
- b) In case inventories are zero (excess demand) and the individual price is low with respect to the average, the firm will increase the price keeping the quantity unchanged.

- c) In case inventories are positive (excess supply) and the individual price is low with respect to the average, the firm form an expectation of lower demand today (in t) than yesterday (in $t-1$) and therefore will reduce the quantity supplied keeping the price unchanged.
- d) In case inventories are zero (excess demand) and the individual price is high with respect to the average, the firm forms an expectation of higher demand today than yesterday and will increase the quantity keeping the price unchanged.

In cases *a*) and *b*) the firm has an unambiguous incentive to change the price in the suggested direction. In case *c*) the firm could in principle cut the price to allure consumers instead of cutting production, but this move would reduce profitability. In case *d*) the firm could in principle increase the price to reduce demand instead of increasing production, but this move would induce a loss of customers. The strategy of changing prices in cases *c*) and *d*) moreover is based on the implicit assumption that the firm is able and willing to manipulate demand through price changes, a situation that we can rule out on the ground of bounded rationality.

Cases *a*) and *b*) are incorporated in the following price rule:

$$P_{it}^s = \begin{cases} \max[P_{it}^l, P_{it-1}(1+\eta_{it})] & \text{if } S_{it-1} = 0 \text{ and } P_{it-1} < P \\ \max[P_{it}^l, P_{it-1}(1-\eta_{it})] & \text{if } S_{it-1} > 0 \text{ and } P_{it-1} \geq P \end{cases} \quad (3.8)$$

where η_{it} is an idiosyncratic random variable uniformly distributed on the support $(0, h_\eta)$, and P_{it}^l is the lowest price at which firm is able to cover average costs:

$$P_{it}^l = \frac{W_{it} + \sum_k r_{kit} B_{kit}}{Y_{it}}. \quad (3.9)$$

Cases *c*) and *d*) trigger quantity adjustments. In this case, the level of production planned or “desired” at the beginning of period t (Y_{it}^d) is equal to expected demand, $Y_{it}^d = D_{it}^e$. Expectations on future total orders – and therefore the scale of production – are revised adaptively according to the following rule:

$$D_{it}^e = \begin{cases} Y_{it-1}(1+\rho_{it}) & \text{if } S_{it-1} = 0 \text{ and } P_{it-1} \geq P_{t-1} \\ Y_{it-1}(1-\rho_{it}) & \text{if } S_{it-1} > 0 \text{ and } P_{it-1} < P_{t-1} \end{cases} \quad (3.10)$$

where ρ_{it} is an idiosyncratic shock uniformly distributed on the support $(0, h_\rho)$. Thus, expectations are revised upward if a manager observes excess demand for its output and its price is already above the average price on the market, and downward when the opposite holds true.

The four cases and the associated adjustments are represented in Fig. 3.1. Point A is the “equilibrium” of the firm/market in this particular setting. It is characterized, on the one hand, by $P_{it} = P_t$. This means that all the agents charge




Fig. 3.1 Price and quantity adjustments for a generic firm i

the same price so that there is no incentive to change individual prices.⁴⁴ Moreover, $D_{it} = Y_{it}$, i.e. demand and supply are equal, so that involuntary inventories are equal to zero.

In the region characterized by a), $P_{it} < P_t$ and $D_{it} > Y_{it}$ (i.e. $S_{it} = 0$): the firm has an incentive to increase the price (in order to catch up with its competitors) and, in principle, also an incentive to increase the quantity produced. In fact, since expectations are formed adaptively, the firm simply adds a stochastic increment to its current output level to determine the future expected level of demand: $D_{it+1}^e = Y_{it}(1 + \rho_{it+1})$. There is room therefore for quantity adjustment.

We have assumed, however, a separation between the domains of quantity and price adjustments so that, *in this case*, we inhibit quantity adjustment. This is the reason why the horizontal arrow is dotted. By increasing the individual price today, in fact, the firm will lower demand in the future so that the absorption of the increased volume of output is not granted. The other three scenarios and the implied adjustments of prices and quantities can be inferred straightforwardly from the figure.

It is clear from the arrows that in a sense there is an implicit tendency for the firm to move towards an “equilibrium”. Having inhibited some of the possible price or quantity adjustments, this tendency would be characterized by a spiraling pattern on the price-quantity space. We have implicitly ruled out therefore monotonic convergence, which would be a likely occurrence in case the dotted arrows were solid ones. Notice, however, that the “equilibrium” itself is changing over time.

Total households’ income is the sum of the wage bill paid to workers employed in t and of dividends distributed to shareholders. Since profits are realized at the end of period $t-1$, accounting consistency implies that dividends also are distributed in that same period.

⁴⁴ In a monopolistic competition setting characterized by Bertrand competition, this would correspond to a Symmetric Bertrand-Nash Equilibrium.

The marginal propensity to consume out of labor income c is a decreasing function of worker's total wealth, defined as the sum of labor income plus all accumulated past savings, and is defined by the following:

$$c_{jt} = \frac{1}{1 + \left[\tanh \left(\frac{SA_{jt}}{SA_t} \right) \right]^\beta}, \quad (3.11)$$

where SA_t and SA_{jt} are average and consumer j 's actual savings, respectively. These savings, in turn, are due to a typical precautionary motive in the face of income uncertainty: households hold assets to smooth their consumption in case of unpredictable declines in income associated with spells of unemployment.

In line with the empirical evidence from the Consumer Expenditure Survey (Souleles, 1999), as well as with predictions from the theory of consumption under uncertainty (Carroll and Kimball, 1996), the marginal propensity c of our artificial consumers is assumed to decline with personal wealth.

Given the absence of any aggregate market-clearing mechanism, consumers have to search for satisfying deals on a fully decentralized goods market. The information acquisition technology affects the number Z of firms a consumer can visit without incurring transaction costs. In other words, transaction costs are equal to zero if the consumer does not cross the border of her local market of size Z , but they become prohibitively high as soon as a consumer tries to search outside it. In what follows, the identity of the Z firms associated to a generic consumer j at any time period t is determined by a combination of chance and deterministic persistence. The search mechanism in fact works as follows:

- Consumers enter the market sequentially, the picking order being determined randomly at any time period t .
- Each consumer j is allowed to visit Z firms to assess the price posted by each one of them. In order to minimize the probability to be rationed, she visits for sure the largest (in terms of production) firm visited during the previous round, while the remaining $Z-1$ firms are chosen at random. Thus, consumers adopt a sort of preferential attachment scheme, whereby preference is given to the biggest firms.
- Posted prices (and the corresponding firms) are then sorted in ascending order, from the lowest to the highest. Consumer j tries to spend a fraction c out of the labor income earned in period $t-1$ and of accumulated past savings in goods of the firm charging the lowest price in his local market.
- If the cheapest firm has not enough output to satisfy j 's needs, the latter tries to spend her remaining income buying from the firm with the second lowest price, and so on.
- If j does not succeed in spending her whole income after she visited Z firms, she saves (involuntarily) what remains for the following periods. For the sake of simplicity, the interest rate on savings is assumed to be equal to 0.

The search and matching process described above is based upon an evolving network structure. The links connecting firms and consumers are in fact continu-

ously changing over time. In particular, the mechanism that governs the choice of a seller on the part of the buyers yields a sort of *preferential attachment*. The firm which posts the lowest price in fact attracts a large fraction of consumers and crowds out competitors, gaining the ability to stay on the market in a predominant position also in the future. After the market for consumption goods has closed, the i th firm has sold Y_{it} , at the price P_{it} . Accordingly, i 's revenues are $R_{it} = P_{it}Y_{it}$. Due to the decentralized buying-selling process among firms and consumers, it is possible that a firm remains with unsold quantities ($S_{it} > 0$). In the following period, the variable S will be used as a signal in adjusting firms' prices or quantities, as explained above.

3.7 Bankruptcy, Exit and Entry

At the end of period t , each firm computes profits $\pi_{it} - 1$. Should they be positive, firm's shareholders receive dividends $Div_{it} - 1$, which are calculated as a fixed fraction δ .

The residual, i.e. retained profits, are added to net worth inherited from the last period, A_{it-1} . Therefore, the law of motion of net worth of a profitable firm is:

$$A_{it} = A_{it-1} + \pi_{it-1} - Div_{it-1} \equiv A_{it-1} + (1-\delta) \pi_{it-1}. \quad (3.12)$$

As we have seen above, net worth is used to finance the wage bill. If internal funds are insufficient, firms can borrow external funds from banks.⁴⁵ The higher the amount of debt relative to net worth – i.e., the leverage ratio – a firm records, the higher is the probability of bankruptcy, *ceteris paribus*. If net worth turns out to be negative, i.e. if the firm records a loss (negative profit) and this loss is such as to wipe out all net worth accumulated in the past, the firm becomes technically insolvent and is declared bankrupt. In the case of the bankrupt firm – say firm f – therefore $\pi_{ft-1} < 0$ and

$$A_{ft-1} < -\pi_{ft-1}, \text{ so that } A_{ft} = A_{ft-1} + \pi_{ft-1} < 0. \quad (3.13)$$

As a consequence, the bankrupt firm exits the market. In line with a large literature on capital market imperfections, then, net worth is the key variable to assess the firm's viability. When a firm is not viable any more, i.e. when it goes bankrupt, it exits the market. For this reason, bankruptcy is the most straightforward mechanism to model exit. From the viewpoint of complexity, the dynamics of operating cash flows drives the selection mechanism.

⁴⁵ For simplicity, we assume that firms do not attempt to raise funds by issuing new equities. This admittedly extreme assumption can be grounded in asymmetric information on the stock market. The manager of the equity issuing firm, in fact, can assess the fundamental value of the firm much better than the potential shareholder and this asymmetric information scenario is common knowledge. In this setting *equity rationing* can occur, i.e. the firm may eventually rule out the issuing of new equities because the shareholders would purchase the new shares only at too low a price.

Of course, new firms are also entering the market. We assume that each bankrupt firm is replaced by a new entrant whose initial condition (size at entry) is set below the average size of incumbent firms.⁴⁶ This one-to-one replacement of bankrupt firms with entrant firms is essentially a working hypothesis, which allows us to keep the total firms' population constant. We can offer a rationale for the assumption, however, based on two widely accepted stylized facts (Sutton, 1997). First, in each established (mature) industry, there is a tendency for the number of firms to settle down around a roughly constant level, below the maximum recorded in that sector's history. Second, the inflow and outflow of firms are highly correlated: Geroski (1991), for example, reports a correlation coefficient of 0.796 for a sample of 95 industries in United Kingdom in 1987. Implicitly we are assuming a correlation equal to 1.

Due to firms' bankruptcies, banks will record non-performing loans (bad debt). Bad debt on the bank's book is equal to a certain share of the bankrupt firm's equity. For example, if the bank is financing 50% of firm's debt and the firm goes bankrupt, the bank will write down its assets' value for an amount equal to 50% of firm's equity. Consequently, a law of motion for banks' equity can be defined as well:

$$E_{kt} = E_{kt-1} + \sum_{i \in \Theta} r_{kit-1} B_{kit-1} - BD_{kt-1} \quad (3.14)$$

where Θ is bank k 's loan portfolio, r_{kit-1} is the interest rate charged to firm i at time $t-1$ and $BD_{kt-1} \leq \sum_{i \in \Theta} B_{kit-1}$ represents bank's bad debt. As for firms, it may happen that bank's equity becomes negative. In this case the Government bails the bank out, replacing it with a random copy of surviving banks.

3.8 The "Growth+" Model: R&D and Productivity

A key insight of modern growth theory is that technological progress is an incentive-responsive activity pursued directly at the firm level. In this section we discuss a simple variation of the baseline framework to allow for the endogenous evolution of productivity, and we label this case the "growth+" scenario.

In order to implement this variant of the basic *BAM* model, we assume that productivity evolves over time according to a first-order autoregressive stochastic process:

$$\alpha_{it+1} = \alpha_{it} + z_{it} \quad (3.15)$$

where z_{it} is the realization of a random variable, exponentially distributed with parameter $\frac{1}{\mu_{it}} = \frac{p_{it} Y_{it}}{\sigma_{it} \pi_{it}}$. The parameter σ_{it} is the fraction of gross nominal positive

⁴⁶ To compute the average size of the incumbent firms we use the truncated mean at 10%. This means the lower and upper 5% of the firms' population are ruled out.

profits (π_{it}) which is used to fund investments in R&D. Hence, μ_{it} is R&D expenditure per unit of output, or R&D expenditure intensity. It follows that in our setting the higher R&D intensity is, the higher the expected increase in productivity results.

In simulations, σ_{it} will be modeled as an exponential function decreasing with the firm's financial fragility, defined as the ratio between the current wage bill and internal financial resources A_{it} , and normalized such that $\sigma_{it}(0) = 10\%$. As a consequence, fluctuations in R&D expenditure can be traced back either to changes in profits or to endogenous changes in the behavioural parameter σ_{it} . Equation (3.15) and the operational underlying assumptions can be thought of as a reduced form reflecting theoretical and empirical considerations suggested by a profusion of studies on the determinants of corporate R&D investment (Reynard, 1979; Fazzari and Athey, 1989, Greenwald *et al.*, 1990), according to which investment in research activity for the sake of technical progress is inversely related to financial fragility.

In the “growth+” model the law of motion of net worth (3.12) must be amended to take into account not only the payment of dividends, but also R&D expenditures:

$$A_{it} = A_{it-1} + (1 - \sigma_{it-1}) * (1 - \delta) \pi_{it-1}. \quad (3.16)$$

In terms of the computational model, the growth mechanism can be switched off by simply posing the parameter $\sigma = 0$ for all firms.

3.9 Simulation Results

We are now ready to explore the key properties of the BAM model. We run several sets of simulations using the constellation of parameters presented in Table 3.1. The choice of parameter values has been constrained merely by the need to rule out patently unrealistic dynamic behavior, i.e. degenerating paths identifiable by visual inspection and conventional empirical standard.⁴⁷ In particular, no attempt has been made at this stage to calibrate the model – for instance, by means of genetic algorithms – in order to force the output of simulation to replicate some pre-selected empirical regularities. As we will see momentarily, in spite of this limitation the model works pretty well along several margins. An analysis of robustness to changes in parameters through Montecarlo methods will be carried out in Sect. 3.10.

⁴⁷ Examples of degenerate dynamics we want to avoid are extremely volatile aggregate GDP dynamics, average rates of bankruptcy and unemployment over 50%, and average rates of annualized inflation outside the $\pm 10000\%$ range.

Table 3.1 Parameter values used in simulations

	<i>Parameter</i>	<i>Value</i>
<i>I</i>	Number of consumers	500
<i>J</i>	Number of firms	100
<i>K</i>	Number of banks	10
<i>T</i>	Number of time periods	1000
c_P	Propensity to consume of <i>poorest</i> people	1
c_R	Propensity to consume of <i>richest</i> people	0.5
σ_P	R&D investment of <i>poorest</i> firms	0
σ_R	R&D investment of <i>richest</i> firms	0.1
h_ξ	Maximum growth rate of wages	0.05
H_η	Maximum growth rate of prices	0.1
H_ρ	Maximum growth rate of quantities	0.1
H_ϕ	Maximum amount of banks' costs	0.1
<i>Z</i>	Number of trials in the goods market	2
<i>M</i>	Number of trials in the labor market	4
<i>H</i>	Number of trials in the credit market	2

3.9.1 The Baseline Scenario

We first simulate a baseline version of the model obtained by switching off R&D expenditure – i.e., $\sigma_P = \sigma_R = 0$ – so that productivity is constant. In the four panels of Fig. 3.2 we present the output of a representative simulation concerning: (a) the (log) real GDP; (b) the rate of unemployment; (c) the annual inflation rate and (d) the ratio of labour productivity to the real wage. In order to get rid of transients, only the last 500 simulated periods have been considered.

The time path of aggregate activity is characterized by irregular fluctuations around a roughly constant mean. The model is able to generate an alternation of booms and recessions as a non-linear combination of *idiosyncratic* shocks affecting individual decision-making processes. The account of business cycles offered by the present model is at odds with that provided by DSGE models, according to which fluctuations in aggregate activity are explained by random changes in *aggregate* variables such as TFP growth (as in RBC-DSGE models) or monetary, investment or mark-up shocks (NK-DSGE approach).

Sudden, deep and rather short recessions are due essentially to the bankruptcy of big firms, which spread through subsequent shockwaves to the economy as a whole. In fact, the bankruptcy of a firm, say α , yields:

- A negative *demand spillover*. The loss of employment generated by the failure of firm α , in fact, brings about a reduction of demand – financed out of the wages previously paid to α 's workforce – for the products of other firms, say β and γ . These firms will experience a reduction of sales and, other things being equal, of profits. The accumulation of net worth of forms β and γ , there-

fore, will slow down and their fragility (and vulnerability to idiosyncratic shocks) will in principle increase.

- *A non-performing loan.* The bank which has extended loans to α will record a bad debt on its balance sheet. The accumulation of net worth at the bank, therefore, will slow down and the supply of loans will change in the same direction due to the target capital requirement ratio. This means that also β and γ may eventually face a constraint on the amount of credit they can get from the bank.

Even though we have not made any serious attempt at calibration, the BAM framework displays neither pathological phenomena, nor degenerate dynamics. The unemployment rate ranges between 2% and 12%, while the yearly rate of inflation is on average equal to 5%, and turns occasionally into moderate deflationary episodes. The average real wage and labour productivity follow a similar pattern so that – as shown in panel (d) – their ratio settles around a long run constant value of approximately $2/3$. Since we did not impose any aggregate equilibrium relationship between the two variables, the (average) constancy over time of income shares turns out to be an emerging feature of our self-organizing system of heterogeneous interacting agents.




Fig. 3.2 Emergent macroeconomic dynamics from a representative simulation of the baseline model. (a) Real GDP; (b) rate of unemployment; (c) annualized rate of inflation; (d) productivity/real wage ratio

Other interesting aggregate stylized facts emerging from simulated decentralized interactions are shown in the four panels of Fig. 3.3. Panel (a) illustrates the presence of a negative relationship between the rate of wage inflation and the rate of unemployment, i.e. a standard (albeit quite flat) Phillips curve. The negative correlation between the two variables is weak (-0.10) but statistically significant. Panel (b) shows a negative relationship between the output growth rate and the unemployment growth rate – i.e. a typical Okun curve. A third emerging regularity regarding the labour market is the Beveridge curve reported in Panel (c), in which it is shown that a negative relationship appears as we plot the rate of vacancies (here approximated by the ratio between the number of job openings and the labour force at the beginning of a period) against the rate of unemployment. Also in this case the goodness of fit is not particularly satisfactory, but the negative correlation between the two variables, albeit weak (-0.27), is once again statistically significant. Finally, Panel (d) shows the firms' size distribution, with size measured by total production. As in the real world, the distribution is highly skewed to the right: small and medium sized firms dominate the economy; large firms are relatively rare, but their production represents a large part of total supply.

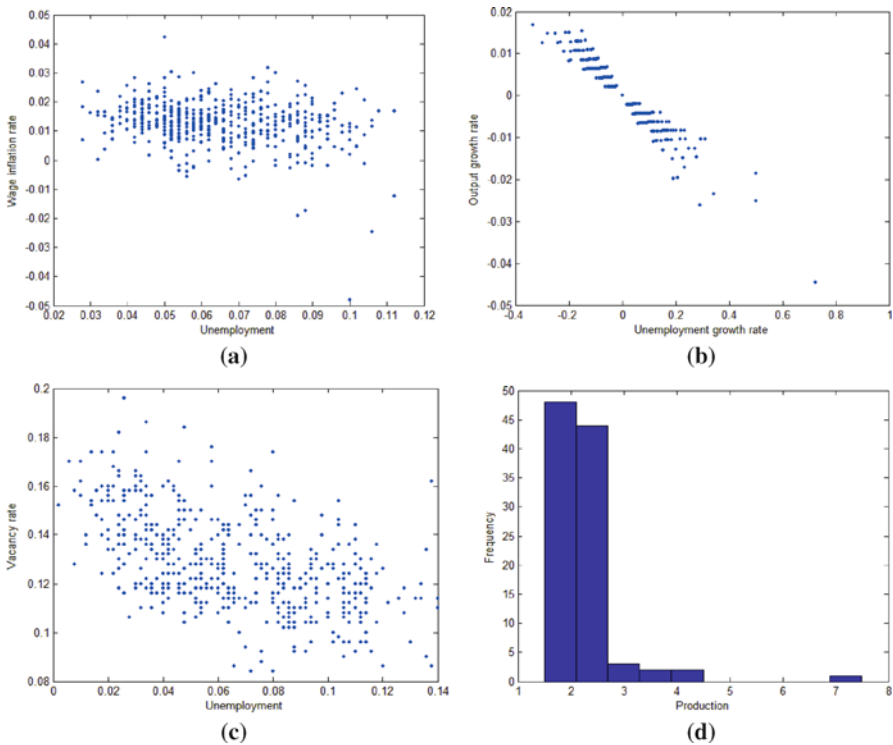


Fig. 3.3 Emergent macroeconomic dynamics from a representative simulation of the baseline model. Phillips (a), Okun (b) and Beveridge (c) curves, and the firms' size distribution (d) generated by simulations

3.9.2 Profits, R&D and Productivity

In this subsection we present results for the “growth+” version of the model, in which firms invest in R&D ($\sigma_R > 0$), so that productivity evolves over time as described in Sect. 3.7. In Fig. 3.4 we present simulation results on the dynamics of GDP, the rate of unemployment, the rate of inflation, the productivity of labour and the real wage.

The main difference between this scenario and the baseline one (Fig. 3.1) is the time path of aggregate activity, which is now characterized by an alternation of aggregate booms and recessions along a *long-run* growth path. The reason for this dynamic pattern is obvious. Output growth is now driven by productivity growth stochastically depending on R&D investments. The latter, in turn, depend on the firms’ financial conditions: the higher profits, the greater expenditure in R&D and the quicker the pace of productivity. As regards fluctuations, inflation, unemployment, productivity and the real wage, what we said about the baseline scenario applies here as well. Sudden stops of growth and short recessions are due essentially to the bankruptcy of large firms, which spread through the macroeconomy as

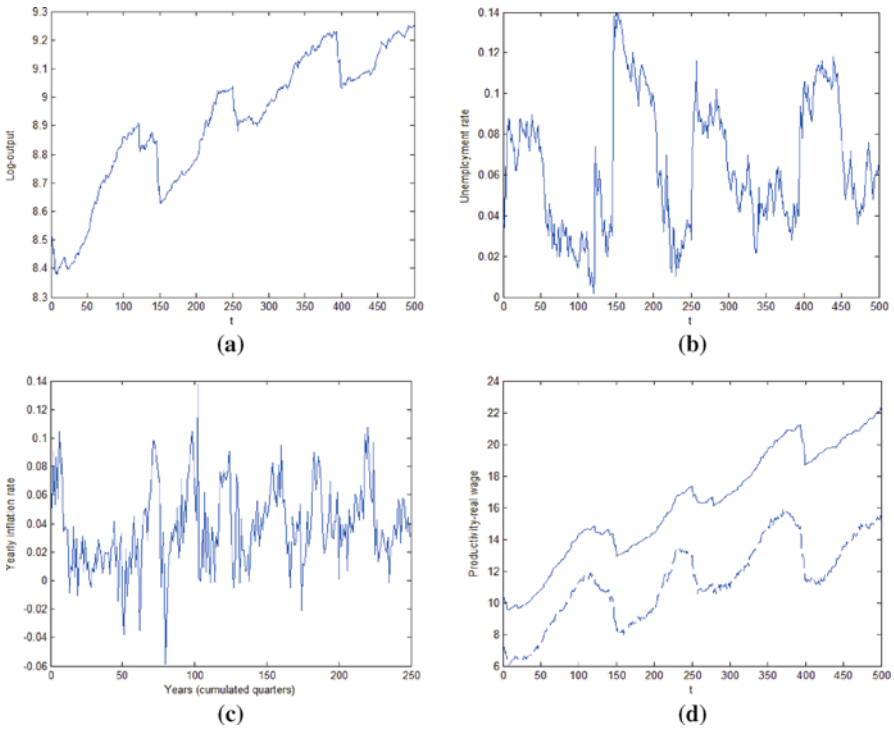


Fig. 3.4 Emergent macroeconomic dynamics from a representative simulation of the “growth+” model. (a) Real GDP; (b) rate of unemployment; (c) annualized rate of inflation; (d) productivity/real wage ratio

explained in the previous subsection. If we let each simulated time period correspond to one quarter, in our simulations the per-year probability to experience an economic disaster (i.e., a drop in real GDP of 15% or higher) is ranging between 0.8% and 1.7%. These figures are essentially in line with estimates reported by Barro (2006), according to whom the per-year probability of a big depression in OECD countries in the 100 years immediately before the global recession of 2008/09 is in the range 1.5–2%. Notice, however, that Barro includes wars into his calculation of major disruptions. Furthermore, in line with the long-run experience of industrialized countries, simulated data suggest that great depressions represent transitory disturbances, in that the long-run real GDP growth path is not significantly affected by major displacements.

Simulations illustrate that the likelihood and severity of economic disasters are increasing with the relevance assigned to the preferential-attachment scheme followed by consumers when searching for the best bargain in the goods market (see above, Sect. 3.6). This makes sense: if customers spread more equally over the market, the probability of finding a really big firm – and *a fortiori* the probability of finding a really big firm on the verge of bankruptcy – is lower. In fact, a preferential attachment scheme generates *auto-catalyticity*, a property a simple unit possesses whenever the time variations of the quantities characterizing it are proportional (via stochastic factors) to their current values. The performance of the macro system is then dominated by the micro units which happen to experience the highest auto-catalytic stochastic positive and/or negative growth rate, rather than by the behavior of a typical or *representative* element. The system is endowed with a kind of multiplier, which accelerates both positive and negative growth.

In Fig. 3.5 we present the Phillips, Okun and Beveridge curves emerging from the simulation of the “growth+” variant. Panel (a) shows the emergence of a Phillips curve. The negative correlation between the rate of wage inflation and the unemployment rate is small (-0.19), but statistically significant. Panel (b) shows the Okun curve. The Beveridge curve is reported in Panel (c): also in this case the goodness of fit is not that high, but the negative correlation between the two variables is statistically significant. Finally, Panel (d) shows the firms’ size distribution. The shape of the latter is highly skewed to the right, as in the corresponding panel of Fig. 3.3.

In addition to the features characterizing the size distribution, a significant body of empirical literature (see e.g. Amaral *et al.*, 1997; Bottazzi and Secchi, 2005) has revealed that the observed distribution of firms’ growth rates is tent-shaped and can be well represented by an asymmetric Laplace (i.e. double exponential) distribution. Though in general the theoretical functional form is excessively regular to capture empirical extreme values, which are generally distributed around much fatter tails than predicted by a Laplace, nonetheless the latter returns an extremely good fit in central portions of the data support.

Fig. 3.6 allows us to visually assess the ability of simulated data in replicating this empirical regularity. If we focus on the cross-sectional outcome in the last simulation period, the (log)rank-output growth rate (Panel a) is clearly tent-shaped for the bulk of the distribution, while both tails happen to be sensibly

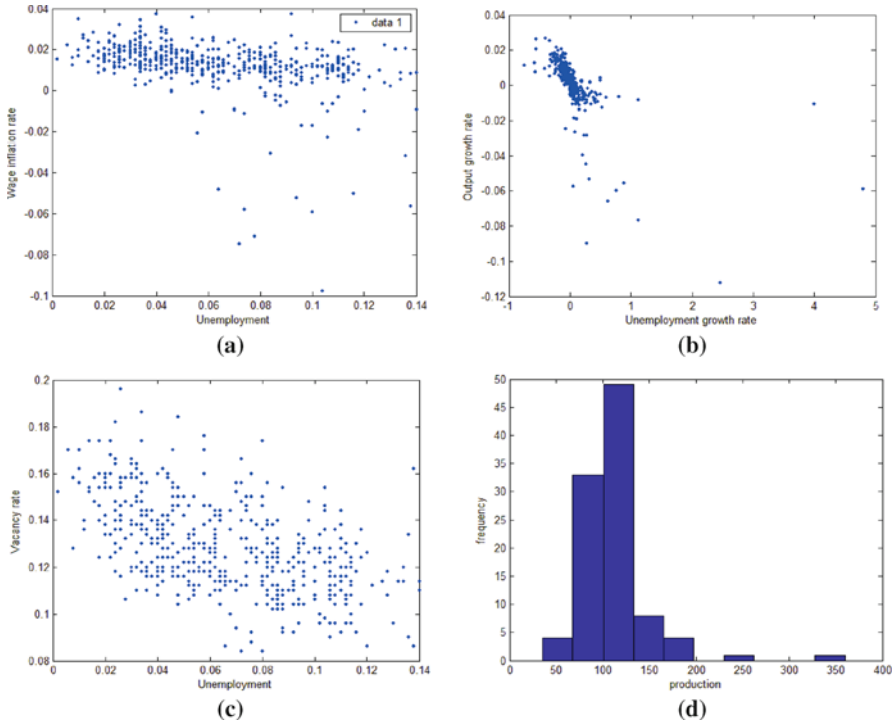


Fig. 3.5 Emergent macroeconomic dynamics from a representative simulation of the “growth+” model. Phillips (a), Okun (b) and Beveridge (c) curves, and the firms’ size distribution (d) generated by simulations

fatter than predicted by the Laplace model. Furthermore, this regularity is robust to a change in the variable used to measure firms’ size: a similar pattern emerges for the (log)rank-size diagram of the growth rate net worth (Panel *b*). The last two panels of the figure report simulated evidence for two additional aggregate variables: the average real interest rate (Panel *c*) and the number of firms which go bankrupt each period (Panel *d*). Their low-frequency fluctuations are clearly synchronized, as both of them peak in correspondence of aggregate slumps, a point which deserves to be further explored.

We start by observing that in this framework a recession is first and foremost the outcome of a wave of bankruptcies. The dynamics of aggregate economic activity is due to the combination of exogenous small idiosyncratic shocks, on the one hand, and of the endogenous systemic evolution stemming from the complex interaction of the financial stance of individual firms and the market structure, on the other one. All decisions regarding production plans are influenced by changes in financial positions: in a deep sense, we might say that business cycles are endogenous and financially-driven. Because of the stochastic nature of firms’ productivity and the time-varying composition of the corporate sector, the frequency and amplitude of business fluctuations change over time; accordingly, the rela-

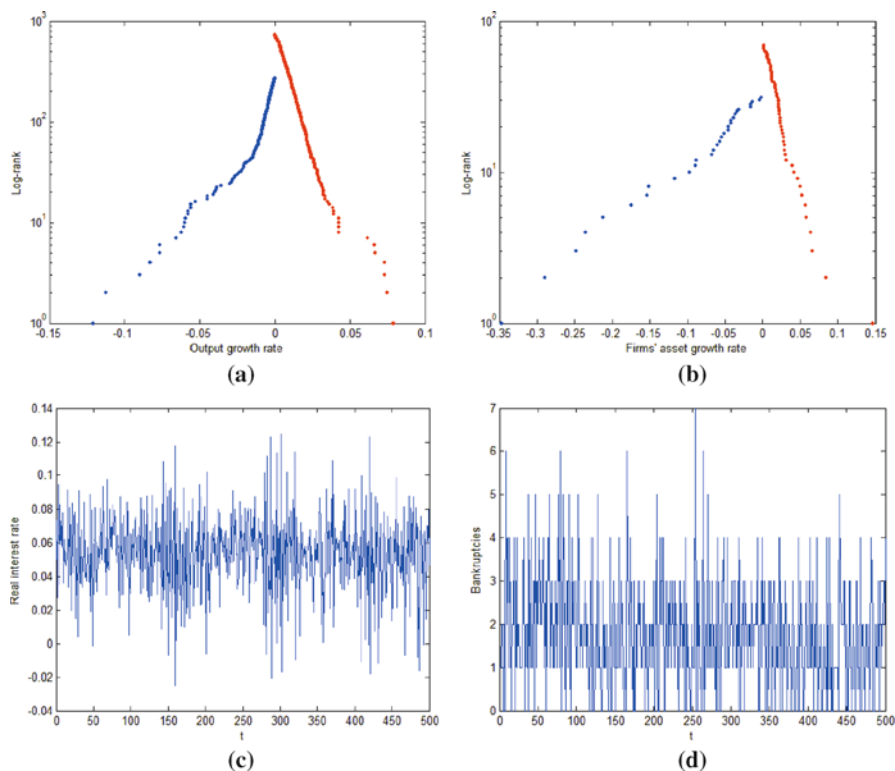


Fig. 3.6 Emergent macroeconomic dynamics from a representative simulation of the “growth+” model. (a) Distribution of output growth rates; (b) distribution of firms’ net worth growth rate; (c) average real interest rate; (d) number of firms’ defaults

relationship between the aggregate output and some measure of financial fragility (here, the cross-sectional average of wage-bill/total-equity ratios), though preserving over time the same qualitative pattern, changes from cycle to cycle. In particular, the endogenous nature of fluctuations can be described in terms of Hyman Minsky’s “financial instability” hypothesis (Minsky, 1982), according to which a crisis is the result of two contextual tendencies. First, during expansions economic units tend to increase the risk embedded in their balance sheets, as they shift their liability structures from a hedge (units which can fulfill all of their contractual payment obligations by means of cash flows) to a speculative (units that can fulfill their payment obligations on “interest account”, but cannot repay the principle out of cash flows) or even to a Ponzi (units whose cash flow is not enough to fulfill either the repayment of principle or the interest due on outstanding debts) position. Second, as the weight of speculative and Ponzi financing increases, the system as a whole becomes more and more sensible to falls in profits and to rises in interest rates.

The whole story can be appreciated by looking at Panel (a) of Fig. 3.7. The inception of big recessions – here signaled by shaded vertical areas – is in general




Fig. 3.7 Recessions (grey bands) and market structure from a representative simulation of the “growth+” model. (a) Financial fragility measured by the ‘wage-bill to equity’ ratio; (b) ratio between the market price and the market-clearing price; (c) firms’ heterogeneity measured by the coefficients of variation of posted prices; (d) dispersion of the equity and sales distributions

heralded by a substantial increase of the cross-section mean leverage ratio, which decreases as the downturn ensues. During expansions, in turn, the financial fragility goes through two phases: it goes down steadily at first, to subsequently increase at an accelerating pace.

In an attempt to provide a chronological description of the intertwined dynamics of financial fragility and aggregate output during a financially-driven business cycle, we identify four different phases for any cyclical movement from trough to trough. The system goes through two distinct stages as the economic activity moves from a cyclical trough up to a peak along an expansion – a *period of tranquillity* (or *financially-hedge* phase) and a *financially fragile boom* period – and two distinct stages as the economy moves along a recession from a cyclical peak down to a trough, namely a *speculative recession* period and a *safe recession* (or *hedge depression*) period.

At the bottom of the cycle – i.e., at the lower turning point – the average debt-to-equity ratio is on a descending gradient, as the cascade of bankruptcies characterizing the now-ending recession has already “cleared up” the corporate sector, forcing all financially unsound (Ponzi) firms to exit the market. As the balance

sheets of survivors become more and more robust due to such a natural selection mechanism, output and profits increase, while debt commitments become lighter. This scenario describes a virtuous circle – thus, a *period of tranquillity* – in which the growth of output and profits is paralleled by a decline of debt.

Positive profit opportunities tend to reduce risk-awareness, inducing firms to expand their production and to increase their workforce, generating positive demand spillovers and making their demand for external finance stronger. As a result, debts in the aggregate start increasing, and their escalating amount eventually determines a transition towards a *financially fragile boom* period, characterized by high leverage ratios and a growing sensitivity of firms' balance sheets to accidental falls in profits or increases in interest rates. The aggregate economic activity reaches its cyclical peak when the deterioration of individual balance sheet positions is such that a *normal* flow of idiosyncratic shocks starts transforming the rising number of speculative firms into Ponzi units, so that a higher-than-usual number of Ponzi units fail. This leads to an endogenous downturn, triggered by a new cascade of bankruptcies. A new recession begins.

Right after the upper turning point – during the *speculative recession* period – the sharp decline in profits starts to depress output and productivity growth. Firms' financial conditions are still unsound, and their debt-to-equity ratio goes further up. Only when the average financial soundness improves due to exists (bankruptcies) and deleveraging – i.e., when the debt-to-profit ratio starts declining – the recession becomes *safe* or financially robust. At the end of the robust depression, profits becomes greater than debt commitments, a turning point in the business cycle occurs and a new recovery sets in.

If we employ in our artificial world the Minsky's taxonomy of firms as regards their financial conditions, we find that in each simulated period approximately two thirds of the firms are hedge, while Ponzi firms represent less than one tenth of the whole population. Of course, the remaining units are speculative. While the ratio of the number of financially fragile (the sum of speculative and Ponzi) firms to that of hedge ones is rather stable over time, the cross-section mean of debt-to-equity ratios (that is, *systemic* financial fragility) is significantly pro-cyclical. This apparent inconsistency can be solved as one thinks about the role heterogeneity plays in our artificial world: during periods of positive growth some really big firms emerge, as suggested by a proxy of the market concentration index (Panel *(d)*). Even though the number of financially fragile firms is moderately stable over time, average fragility can go up as the economy expands simply because the financial position of a small number of very large firms eventually starts to become more and more unsound. Thus, our model corroborates the prediction of a substantial increase of overall financial fragility during “prosperous times” (the ascending phase of the business cycle), which is generally seen as the cornerstone of the financial instability hypothesis. Furthermore, unforeseen disturbances will trickle down across the whole distribution of agents because of aggregate demand spillovers, modifying this way the macroeconomic behavior. If composition effects are large enough, the response of the system to an identical shock changes

over the business cycle, as it depends on the actual distribution of firms in terms of the balance between their internal and external finance.

Given that the degree of competition among firms and the distribution of profit opportunities interact with the dynamics of systemic financial fragility, an additional issue worth exploring is the evolution of market power over the cycle. As an indicator of market structure we employ the ratio of the actual (average) price index to the (homogeneous) equilibrium price, defined as the price that an imaginary Walrasian auctioneer would cry in order to equate the quantities demanded by households and supplied by firms (Panel (b)). An increase of the ability to price profitably above the competitive level – an increase of the value of the ratio depicted in Panel (b) – translates into an increase of market power. Over a typical cycle, the latter crosses three different phases. During a *robust expansion*, competition is becoming more and more fierce and actual price(s) tend to converge towards the market-clearing level. Such a convergence reaches its lower limit – with the actual-clearing price ratio remaining well above 1 – as the system enters a *fragile expansion*. It is only after a new recession sets in that individual prices start again to wander away from the fictitious Walrasian equilibrium level, as a stream of new bankruptcies shakes the market and the competitive pressure decreases accordingly. This in turn lowers significantly the standard deviation of prices (Panel (c)).

A somewhat opposite dynamics can be detected as regards the degree of heterogeneity of active firms, measured both in terms of their level of equity (net worth) and of sales. As depicted in Panel (d), during upswings dispersion increases steadily because the system dynamics is dominated by the micro units which happen to experience the highest stochastic autocatalytic growth rate, and can grow very rapidly. On the contrary, during recessions dispersion reduces, as a certain number of large financially fragile firms are forced to exit due to bankruptcy, just to be replaced by new entrants characterized by a relatively homogeneous initial size.

3.9.3 *Measuring the Performance of the BAM Model by Means of DSGE Methodology*

Standard macroeconomic theory faces enormous difficulties in jointly explaining the rich list of phenomena we have just overviewed. For instance, basically all mainstream theories attempting to explain the Great Depression which hit the world economy during the 1929–39 period treat this episode basically as an outlier, and rely on a rather ad-hoc combination of severe frictions, technological and policy shocks to explain it (Chary *et al.*, 2002). BAM models, on the contrary, can naturally accommodate the alternation of phases of smooth growth and deep crises as instances of the same underlying dynamical process. For instance, in Panel (d) of Fig. 3.6 one can appreciate that the time series of firms' bankruptcies

remains roughly constant during the whole simulation, even when the system experiences severe breakdowns. This feature of the model reveals the importance of heterogeneity, since a recession does not depend on the mere number of bankrupt firms but on their size: the same economic process can thus produce small or large recessions according to the size of bankrupt firms.

It must be noticed, however, that an appropriate comparison between the *BAM* family of models and more traditional DSGE models can be made only if a common testing methodology is employed. According to DSGE scholars, the explanatory performance of business cycle models has to be measured in terms of their ability to replicate aggregate phenomena at cyclical frequencies along three dimensions: persistence, volatility and co-movements of key variables with aggregate output. In this section we explore the ability of the *BAM* virtual economy to challenge DSGE models by mainly focusing on the latter dimension.

In particular, to make a more direct comparison we stick to qualitative measures of success. This could sound odd, since DSGE models are usually taken to the data by comparing quantitative theoretical predictions with figures summarizing key features of cyclical fluctuations in real economies. This impression is largely false, however. Since none formal metric is in general offered to measure the closeness of the model data to the real data, the assessment presented in almost every DSGE paper is ultimately qualitative. Hence, instead of reproducing the familiar table of figures based on actual and simulated data, we prefer to illustrate the performance of our *BAM* model in replicating first-order features of real economies with graphical methods.

For easiness of comparison with a conspicuous literature, the empirical benchmark used against simulation outcomes is the postwar U.S. economy. In particular, filtered-detrended quarterly data for real GDP, employment, labour productivity, real wages, inflation and bank loan interest rates obtained from the Federal Reserve web-based FRED[®] database have been used to calculate correlations at different leads and lags. Results are reported in the first five panels of Fig. 3.8, where we plot the cross-correlations with output at four leads and lags of: (a) employment, (b) labour productivity, (c) the price index, (d) the interest rate on loans and (e) the real wage. Each panel is completed by the corresponding function calculated from real data and a $\pm 20\%$ band, which is conventionally assumed as signaling a lack of correlation.

Our model does a remarkably good job in four cases out of five. From simulations we find that employment and productivity are highly correlated with contemporaneous output; prices are slightly negatively correlated and anticipate output; while the interest rate is a-cyclical. All these patterns mimic the evidence for the U.S. economy remarkably well. The simulated real wage turns out to be procyclical, as in real data, but fails to anticipate cyclical movements of aggregate activity by two to three quarters. Finally, Panel (f) presents the transitory impulse-response functions, calculated by means of an AR(2) estimate, for the actual (solid line) and the model-generated (dashed line) output, respectively. The simulated model can mimic the hump-shaped response of cyclical output to transitory

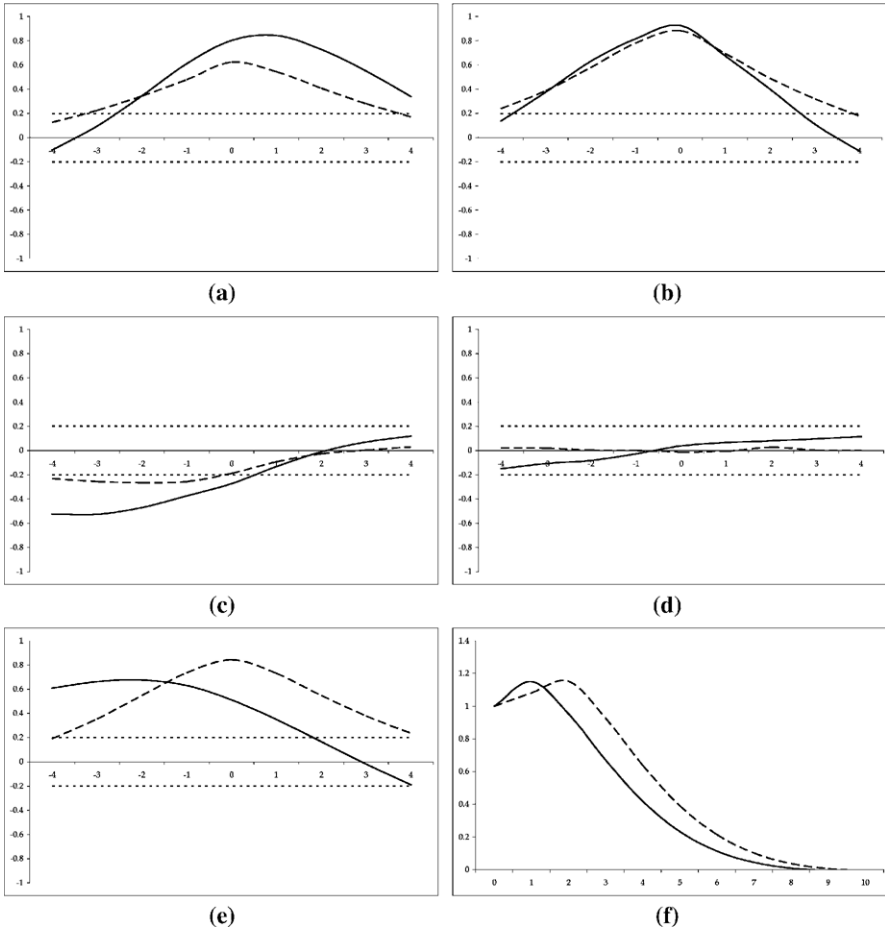


Fig. 3.8 Cyclical features of model-generated and real data. Solid lines show sample moments, while dashed lines show moments generated by simulations. (a) Employment; (b) productivity; (c) price index; (d) interest rate; (e) real wage; (f) GDP cyclical component impulse-response function

shocks – a feature that first-generation RBC models failed to capture (Cogley and Nason, 1995) – though the peak in real data anticipates the simulated one by one quarter. The trend-reverting dynamics is nevertheless really similar.

Recalling that all these results have been obtained without any serious effort to properly calibrate the model, we argue that the BAM basic setup proves to display rich and interesting aggregate and disaggregated dynamics under rather general conditions. Furthermore, as we have just showed it can also successfully challenge the explanatory power of DSGE models when confined to their same ground.

3.9.4 Consumption and Buffer Stock

Models built as agent-based computational laboratories offer distinctive opportunities as an experimental tool. In this subsection we provide an illustration of the flexibility of the *BAM* model by exploring its features as we employ an alternative assumption on households' behavior. Namely, we introduce a variation characterized by individual consumption functions based on simple buffer-stock saving rules (Deaton, 1991; Carroll, 1997), in order to examine in particular their effects on the personal wealth distribution. We will see that in this case the ability of the model to reproduce stylized facts is even improved if compared to the baseline version.

The individual marginal propensity to consume (MPC) c is now derived from an adaptive rule, without any mean-field interaction. In practice, each consumer is supposed to possess a personal desired 'total savings-income' ratio, that she strives to keep constant along her lifetime:

$$\frac{S_{t+1}}{W_t} = h \quad (3.17)$$

where S and W represent total savings and income, respectively. If income at time t increases (decreases), consumers will try to increase (decrease) their savings at time $t+1$ as well. Thus, the actual MPC can change from time to time since it depends on the current income growth rate.

Two alternative spending rules have been tested in simulations. In the first one, consumption depends upon current income only, that is $C_t = c_t W_t$. In the alternative version, consumption is financed drawing on both income and savings, $C_t = c_t (W_t + S_t)$. Interestingly enough, we found that the two rules yield identical long-run results and, consequently, we decided to present here only results obtained by means of the simplest $C_t = c_t W_t$ consumption rule.

We define the desired stock of future savings at time $t+1$ as past savings plus retained income at time t :

$$S_{t+1} = S_t + (1-c)W_t \quad (3.18)$$

where c is the individual MPC. Plugging (3.18) into (3.17), and defining $W_t = W_{t-1}(1+g_t)$ where g_t is the income growth rate at time t , we get:

$$\frac{S_t}{W_{t-1}(1+g_t)} + (1-c) = h. \quad (3.19)$$

If we define $S_t/W_{t-1} = h + d_t$, where d_t is the time t divergence between desired and actual savings-income ratios, we finally obtain the expression for the time t MPC for a generic household:

$$c_t = 1 + \frac{d_t - hg_t}{1+g_t}. \quad (3.20)$$

Consumption is then simply defined as $C_t = c_t W_t$.




Fig. 3.9 Fitting of the complementary CDF of personal incomes in correspondence of the last simulation period

Once the buffer-stock based consumption rule is employed, the BAM model keeps its ability to return all the basic emergent macroeconomic features shown in Figs. 3.4, 3.5 and 3.6. Furthermore, we discover that its degree of realism is even improved once the personal wealth distribution is considered. In fact, a huge body of recent theoretical and empirical work (Kleiber and Kotz, 2003) has persuasively shown that three statistical functional forms can be considered as the best-fitting candidates to model real data on personal incomes and wealth: *i*) the four-parameter Generalized Beta II (GB2) distribution; *ii*) the Dagum (D) distribution; and *iii*) the Singh-Maddala (SM) distribution. Thus, a natural way to further assess the ability of the modified BAM model to replicate reality consists in applying these three statistical models to cross-section simulated data for personal income. In addition, we test also the κ -Generalized distribution recently introduced by Kaniadakis (2001), and successfully employed in the analysis of income distribution estimation by Clementi *et al.* (2007, 2008).⁴⁸ Results from such a distribution fitting exercise can be observed in Fig. 3.9. All the statistical models appear to match remarkably well the simulated data, especially the D and the SM distributions. Even though at this stage we are not even trying to confront punctual esti-

⁴⁸ GB2, D and SM can be derived as the solution of a differential-equation Pearson system, with D and SM being three-parameter specializations of the GB2 distribution. The κ -Generalized distribution, on its part, is obtained from the entropy constrained maximization of a deformed exponential function.

mates obtained from real data with estimates for simulated data, from a qualitative point of view this last result confirms once again the amazing ability of the BAM model to generate macroeconomic stylized facts.

3.10 Robustness

In this last section we present some computational tests aimed at checking the robustness of simulation results to changes in the random seeds and in the values of some key parameters (Subsect. 3.10.1). Finally, we explore how our findings are affected by variations of two crucial aspects: the consumers' preference attachment mechanism and the entry mechanism (Subsect. 3.10.2).

3.10.1 *Exploration of the Parameter Space*

In a typical agent-based model an exhaustive robustness check – a procedure also known as *model verification*, aimed at: *i*) confirming the central results of the simulated model and/or revealing possible output variations when the input parameters are changed; *ii*) guiding future work by drawing attention to the most promising directions for further research – should be performed along the whole grid of parameters and random number seeds through extensive Montecarlo simulations (Fagiolo *et al.*, 2007). According to an increasing consensus among practitioners, for each vector in the parameter space a high number of independent simulations should be run, each one for a different seed of the random number generator. Then, after calculating all the relevant statistics of the simulated data, one should compute their mean and variance across simulations. If the latter is sufficiently small, one can state that the model is stable, and each simulation can be interpreted as representative of the underlying data generating process (DGP). Clearly, such a procedure is extremely demanding. For instance, suppose that in a model there are just 10 relevant parameters, and that each parameter can assume 10 different values (a rather simplifying assumption). As a result, one obtains that the constellation of the parameter space is given by 10^{10} vectors. If we perform 20 different runs for each one of them to take into account the possible effects of changing the random seeds, the total number of simulations would amount to $2 \cdot 10^{11}$!

Our strategy for robustness checking is far more modest, as we employ the two different techniques involved in a proper model verification procedure, namely *internal validity* and *sensitivity analysis*, in two separate steps. In a first exercise we run a certain number of independent simulations, each one with a different random seed, using the particular parameter vector shown in Table 3.1. If the

random seeds employed for the random number generator do not cause large variability of the outcome sample points, the model can be deemed as sufficiently accurate. Second, we choose a selected subset of parameters, and we run several simulations to quantify how changes in the values of the input parameters alter the output. The model is then believed to be *good* if the output values of interest do not vary significantly despite significant changes in the input values.

The aggregate behavior emerging from an averaging of outcomes over 20 alternative random-seed simulations show that the results we have discussed so far are significantly robust. The key qualitative time-series features of growth and cyclical fluctuations remain unaffected, and the cross-simulation variance calculated for typical macroeconomic variables (GDP, productivity, inflation, real wage, unemployment, interest rates, bankruptcy rates) is remarkably small. The distribution of the firms' size (both in terms of sales and net worth) calculated in correspondence of the last simulation period is definitely invariant in its significant departure from normality and its strong positive skewness. Finally, a Phillips curve, an Okun law and a Beveridge curve continue to emerge from each simulation and on average.

Fig. 3.10 reports the structure of co-movements at four leads and lags, plus the contemporaneous one, between the de-trended values of the GDP and of the other five variables already considered in Fig. 3.8. It largely corroborates our previous findings regarding the procyclicality of unemployment, productivity and the real wage, as well as the substantial a-cyclicality of the aggregate price index and of the real interest rate. Furthermore, the signs of the configuration of non-contemporaneous correlation coefficients already found for the baseline simulation is largely confirmed as we control for the stochastic dimension of the model. A final remark is in order to highlight the simulation outcome that proves to be most challenging, namely the auto-regressive structure of the de-trended output and its relative hump-shaped impulse-response pattern. At odds with the result shown in Panel (f) of Fig. 3.8, when we consider an average over cross-section simulations, the movement in the log of detrended GDP can be best approximated by an AR(1) structure (with an autoregressive parameter around 0.8). Of course, this calls for further investigations to assess when and how endogenous aggregate positive feedback loops operates in this world.

As regards the second step, we choose to perform a univariate sensitivity analysis, according to which the model outcomes are analyzed with respect to the variation of one parameter at a time, whereas all the other parameters of the system remain constant. For each parameter we run at least four alternative scenarios, with values chosen on rather coarse grids. To somehow summarize our main findings, the parameters that prove to be crucial – in that alternative parameter values change simulation results significantly – are the ones related to the duration of labour contracts, to the number of opportunities any unit is allowed to locally explore as it searches for market transactions (*local* markets), and to the total size of the economy. Let us see them in more detail.

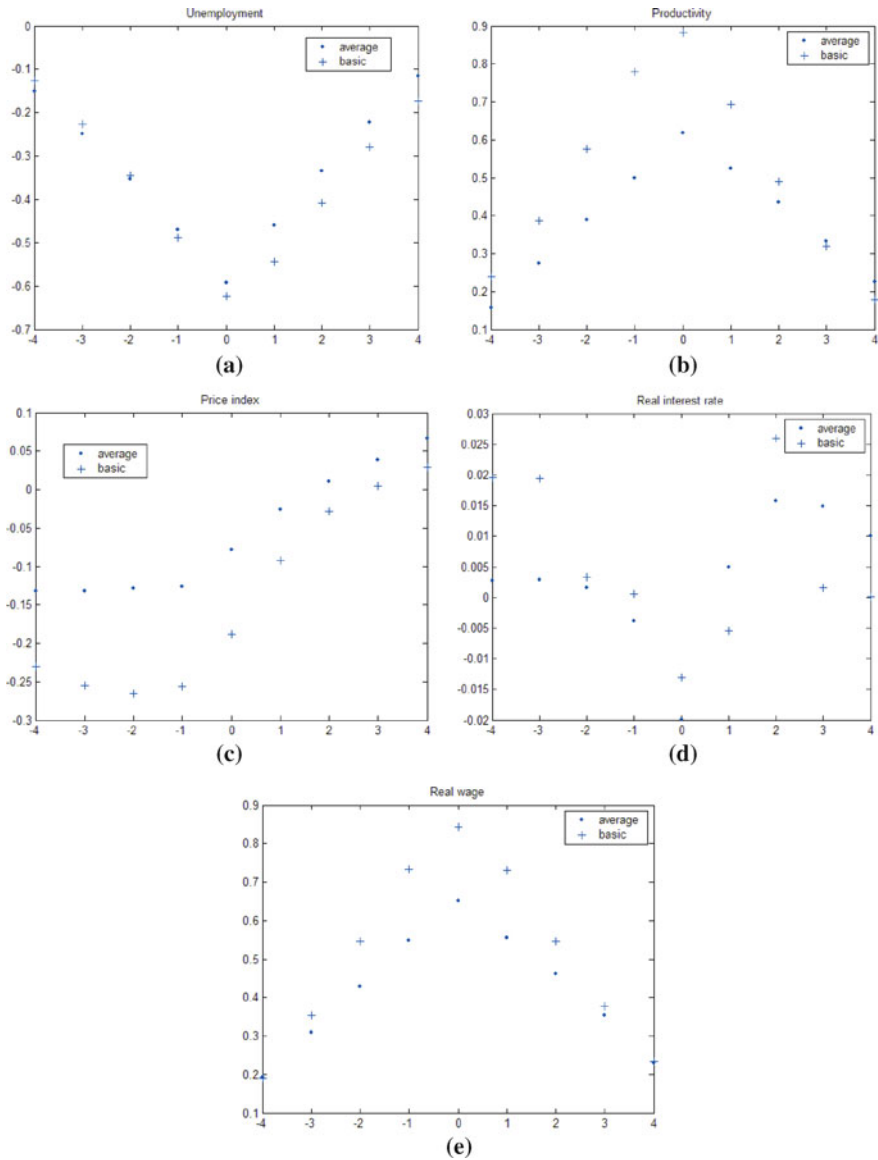


Fig. 3.10 Baseline (+) and cross-simulation mean (°) co-movements at four leads and lags. (a) Unemployment; (b) productivity; (c) price index; (d) interest rate; (e) real wage

Local credit markets. As we increase the number of banks each firm can borrow from – in particular, as we raise the parameter H from its baseline value from 2 to 3, 4, and 6 – the general properties of the model (in terms of output, productivity, unemployment, inflation, real wages, bankruptcy rates, and so on) do not manifest any significant variation. It must be noted, however, that an increase in H forces

the cyclical component of the price index to be coincident with the aggregate output, while the right tail of the size distribution of firms' net worth becomes more and more similar to a Pareto distribution. As the number of potential partners on the credit market is reduced to 1, on the contrary, the size distribution looks more similar to an exponential. A plausible explanation for this feature is as follows. When search costs in the credit market are lower, and accordingly the number of different banks a firm can visit is higher, the probability that firm has to be rationed is relatively smaller, all other things being equal. In terms of the whole population, therefore, firms can fully exploit their proportional growth potential (autocatalicity), and the right tail of the firms' size distribution assumes a Pareto-like behavior.

Local consumption goods markets. The second experiment consists in increasing the number of firms which consumers can visit before purchasing (Z). As we increase Z from 2 to 3, 4, 5 and 6, competition among firms increases, and the function exerted on firms' growth by the preferential attachment mechanism becomes less and less effective. In particular, the real wages become lagging, their co-movement with output similar to those of the price index and, as it is logical, the kurtosis of the firm's size distribution decreases dramatically. Moreover, production displays smoother patterns, without sudden booms or crashes. This happens because in a more competitive environment truly big firms cannot emerge, and consequently systemic risk is more evenly spread across producers.

Local labour markets. The functioning of the labour market is regulated by two crucial parameters: the number of workers' applications (M) and the labour contract length. As far as the former is concerned, we start our sensitivity experiment by decreasing the number of allowable applications from 4 to 3 and 2, discovering that prices switch from being anti-cyclical and leading to pro-cyclical and lagging. Aggregate output shows an higher degree of instability, since firms have a lower probability to fill in their vacancies – and thus to produce planned output – while the upper tail of firms' size distribution appears to become more Pareto-like. Strong path-dependency in the labour market allows the formation of “advantaged” (with a higher probability to fill in their vacancies), and thus more performing, firms. This interpretation is indirectly confirmed as we *increase* the number of applications (to 5 and 6): tougher competition on the labour market and a higher probability to find workers make firms all alike, and their size distribution scales much more as an exponential, or even a uniform. In addition, as one can expect, competition between firms in hiring workers tends to push the real wage up, sometime even above average productivity.

Employment contracts duration. Another relevant parameter tuning the functioning of the labour market is the duration of the employment contracts signed by firms and workers, which in the baseline simulation we set to 8 periods. In order to control for both a very flexible and a quite rigid labour market, we have first decreased it to 6, 4 and 1, to subsequently increase it to 10, 12 and 14. Since we

interpret each simulation period as a quarter, the sensitivity experiment thus covers contract durations stemming from one quarter to three years and half. While for intermediate values of the parameter the main statistical properties of the model do not change significantly, the opposite is true for the extreme values, which produce degenerate dynamics. More precisely, decreasing the labour contract length produces a continuous process of creation and dissolution of the network linking workers and employers. This ever-changing network reduces path-dependence, causing co-movements to become less and less pronounced, except for the unemployment rate and real wages, that basically keep on showing the same properties of the baseline simulation. With a contract length of 6 and 4 periods, output becomes smoother and its cyclical component definitely loses the AR(2) structure. Moreover, because of the lessening of path-dependence the bulk of operating firms tends to distribute more uniformly. It is worth noting that, in spite of a more flexible labour market, on the average unemployment increases and output decreases, revealing the presence of coordination failures on a grand scale due to aggregate demand spillovers. In fact, during downturns firms can easily fire workers; consequently, the economy experiences a sensible reduction of the aggregate demand that causes firms to further revise downwards their production plans and labour demand for the subsequent simulation iterations. On the contrary, if firms are forced by longer contracts to hoard labour and to pay wages also during recessions, aggregate demand reduces less, thus preventing the triggering of a vicious circle. The actual functioning of this mechanism is further confirmed by pushing it to the extreme: when labour contracts last only one period, that is when firms are given full freedom of firing, the number of bankruptcies and the unemployment rate reach very high values, and in most of the simulations the whole economy collapses, signaling the presence of fatal market failures.

A different reasoning applies when the labour market is rigid (in our case when the contract duration is equal to 12 periods, or higher). In this case, simulated co-movements contrast sharply with the ones calculated for real data, and time series dynamics are often degenerated. The supply side of the model is now the weakest ring of the chain: because of long contractual commitments, firms cannot resort to firing when they are financially fragile and go bankruptcy more easily. This leads to an overall macroeconomic breakdown.

The size and the structure of the economy. A last sensitivity experiment concerns the role played on simulation outcomes by the absolute size of the economy and its composition. In our context, this amounts to vary the total number of agents populating the economy on the one hand, and the relative frequency of classes (firms, households, banks) over the whole on the other one. In order to shed light upon these issues, we first run small groups of simulations multiplying sequentially the number of all agents by 2, 5 and 10 *without* changing the proportions among the three classes of economic units. As the size of the economy is scaled up, the average growth rate and the statistical properties expressed in terms of co-movements are very similar to their counterparts calculated for the baseline simulation, whereas the time series of macroeconomic variables display rather

smoother cyclical fluctuations. The negative relationship between aggregate volatility and the economy's mass we find in simulations can be rationalized intuitively – since macroeconomic volatility tends to reduce as microeconomic specific volatility is averaged out over an increasing number agents – and it is also consistent with a large number of empirical studies based on cross-section international data finding a significant negative relationship between the GDP's variance and the country's size (Barro, 1991; Head, 1995, Canning *et al.*, 1998).

As a second step, we proceed to vary the structural composition of the economy by muting the relative frequencies of the classes of agents operating in this world. In particular, we run three groups of simulations doubling the size of just one class per time, while the size of the other two classes are kept fixed to their baseline values. Interestingly enough, the three experiments lead to different outcomes. Doubling the number of banks does not exert any significant variation to the model's outcomes. When the number of households is increased, in turn, the leads-and-lags co-movement analysis shows a scenario quite similar to that of the baseline simulation, but time series appear to grow much faster – and with a higher volatility – thanks to the enlarged availability of workforce. Conversely, an increase of the proportion of firms has the effect of slowing down the average rate of growth of the economy. This happens because of an increased competition on the credit market (with more rationings occurring), on the labour market (with more unfilled vacancies) and especially on the supply side of the consumption goods market (with lower prices, revenues and profits). Since R&D investments conducing to productivity enhancements are financed out of retained profits, in this world a fiercer competition eventually tends to reduce growth opportunities.

3.10.2 Preferential Attachment in Consumption and the Entry Mechanism

The last part of this section is devoted to an inspection of the influence exerted on the model's output by two mechanisms, one regulating the choice of the preferred supplier exerted by consumers, and the other one regulating the entry process of new firms as bankrupt firms leave the market.

Recall that in the baseline BAM model discussed above, consumers are allowed to search for a satisficing deal inside a local market composed of Z firms. At each time period, $Z-1$ of them are chosen randomly, while the last one is in any case the largest (in terms of its scale of production) firm visited during the previous round. This mechanism corresponds to a localized preferential attachment (PA) scheme, and in our context it plays a double role. From the point of view of consumers, maintaining the largest firm they have knowledge of inside their search space allow them to minimize the risk of being rationed, *ceteris paribus*. Since it is not directly influenced by pricing concerns, the common preference for larger firms creates a type of non-market – or *social* – interaction among

consumers: the higher the number of people who have previously chosen a certain firm, the higher the probability that I choose that firm as well. The localized PA scheme, in turn, provides a structure to the topology of the market interactions' network linking firms and consumers. In particular, it endows the economy with a high degree of granularity, with the largest firms becoming even larger as they take advantage of the loyalty of customers and grow to a size not attainable under a pure random network.

To control for the influence exerted by the localized PA scheme on the structure of business fluctuations, we run 20 independent simulations of a pure-random-network version of the BAM model, holding all else constant. The experiment tells us that in the absence of the localized PA mechanism there is a sensible gain in stability, in that the volatility of all relevant macroeconomic variables decrease steadily. As a matter of example, the time series for the GDP obtained from a representative simulation is shown in Fig. 3.11: growth is still fluctuating, but deep crisis disappear completely. The reason why this happens is intuitive: the PA mechanism increases the path-dependence of choices and, at the same time, it makes the economy's volatility greater since it allows the formation of very large firms, whose behavior deeply affects the entire system for the reasons we have explored before. Thus, we can argue that the topology of the networks structuring an economic system plays an important role in its functioning and performance: social interaction *matters* and cannot be ignored without consequences. Furthermore, the localized PA scheme can be suitably tuned to calibrate our agent-based model by means of real data on macroeconomic volatility.

The new co-movement structure resembles its counterparts of the baseline simulation, with the exception of the price index and of real wages that now become lagging, even if for any practical purpose they can still be deemed as a-cyclical. The autoregressive structure of output's cyclical component turns firmly




Fig. 3.11 Log of GDP for a representative simulation without the preferential attachment in consumption mechanism

into an AR(1) process with an auto-regressive coefficient around 0.4, while the firms' size distribution becomes significantly less skew. Hence, moving from a PA scheme to a fully random network linking consumers and sellers produces results that are similar to those obtained previously when we lowered transaction costs on local consumption goods markets. As we perform small clusters of simulations for different points in the parameter space, we discover that the model's outcome preserves all its key features. In particular, this holds true when the sizes of local markets agents can visit each time period are either increased and decreased (that is, when searching costs are varied).

Finally, we turn to evaluate the consequences produced by the firms' entry-exit mechanism on the model's outcomes. In Subsect. 3.9.2 we provided an explanation of emergent output fluctuations based on the endogenous dynamics of financial fragility. However, one could wonder whether business cycle dynamics (i.e., the recurrence of booms, busts and recoveries) actually depends just on endogenous mechanisms, or if it relies also upon the exogenous and automatic introduction in the system of new well capitalized firms whenever bankrupt firms exit. Consequently, in order to explore this issue we run a modified version of the model where firms' profits are heavily taxed by an unmodeled internal revenues office, but revenues are not redistributed into the system. This trick is basically intended to increase the firms' financial fragility, thus producing a higher probability of insolvency. If the automatic entry process is really distortional, the model should display a somehow better performance as the number of bankrupt firms is increased. In spite of the higher systemic financial fragility, in fact, the massive entrance of new financially-sound firms should counterbalance the negative effects caused by the transfer of firms from a speculative to a Ponzi position. Actually, in correspondence of a higher average number of bankruptcies the overall economic behavior shows a substantially worse performance, both in terms of lower growth and of higher volatility. Hence, we argue that the automatic entry mechanism is likely to be neutral, confirming the endogenous explanation of economic dynamics.