Chapter 4
Empirical Validation

> No one who has experienced the intense involvement of computer modeling would deny that the temptation exists to use any data input that will enable one to continue playing what is perhaps the ultimate game of solitaire.
>
> JAMES LOVELOCK
> *Gaia: A New Look at Life on Earth*, pp. 129

4.1 The Empirics of Falsificationism

The aim of this chapter is to understand to what extent the BAM model described and simulated in Chap. 3 is able to reproduce real world phenomena, by applying tools and techniques of *empirical external validation* (see Fagiolo *et al.*, 2007).

As argued at length before, our framework allows us to explore the macroeconomic outcomes generated by the market and non-market interactions of a large population of heterogeneous agents. Starting from simple behavioural rules at the individual level, some aggregate statistical regularities emerge that could not be inferred from the behaviour of single agents in isolation. This *emergent* behaviour often feeds back on individuals, affecting their actions in turn. In other words, micro and macroeconomic behaviours co-evolve in an adaptive way. From this viewpoint, the pattern of the aggregate is not the result of a simple summation and averaging of individual preferences and choices, but it is the product of self-organization "from the bottom-up". A self-organized macroeconomy is furthermore susceptible of abrupt phase transitions when a scenario of criticality occurs.

So far, we have presented simulations carried out using ad hoc parameters' values and the same initial set up for all the agents belonging to the same class. To paraphrase Nicholas Kaldor, we have just started from some hypothesis that could account for some stylized facts (say, the tendency of modern economies to fluctuate around a growing path), without necessarily adhere to historical accuracy (Kaldor, 1965). Using an initial set up of actual Italian data, we currently aim at

D. Delli Gatti, S. Desiderio, E. Gaffeo, P. Cirillo, M. Gallegati:
Macroeconomics from the Bottom-up. © Springer-Verlag Italia 2011

verifying whether the BAM model, simulated over a period for which actual data are fully available, is an acceptable representation of the actual system at micro, meso and macroeconomic scales. In a nutshell, we intend to perform an *ex-post validation* of a microsimulated version of the model. We shall be in a position to conclude that the result of this validation exercise is positive if the model outcome represents a good approximation of the actual data generating process. While the validation process should not be regarded in any case as a formal prerequisite for the development of an ABC model and for a discussion of its predictions (Carley, 1996), a careful assessment of the adequacy and accuracy of the model in matching real world data should help us in resisting the temptation referred to in the epigraph to this chapter.

As we will see, the BAM model is able to reproduce the dynamics of actual data with a good degree of precision along several dimensions. In particular, it nicely replicates some interesting cross-sectional results about size distributions and other remarkable pieces of empirical evidence.

4.2 On the Microfoundation of Agent-based Models

By common wisdom, ABC models possess micro-foundations almost by definition, given that the assumed behaviour of artificial agents is always based on specific micro-rules. Moreover, this kind of models are generally considered satisfactory if and only if they are to some extent able to reproduce aggregate empirical evidence and statistical regularities. However, such a position must be somehow qualified, since ABC models are not naturally and necessarily micro-founded. The right question to be posed is in fact empirical in nature: are we sure that the micro-rules driving the dynamics of the disperse system are really based on empirical evidence and actual data? The answer one can give is essential to decide whether an agent-based model can be validated or not, especially in its microsimulated version when the model itself is initialized and continuously compared with actual data.

According to this approach, a really micro-founded ABC model should be based on the following prerequisites:

- the agents' initial conditions should be initialized with actual data;
- the behavioural rules, structural equations and all the other elements governing the evolution of the system should be based, whenever possible, on statistics-based data analysis. If direct data analysis is not possible, updating rules should rely on surveys' results or on experimental evidence;
- researchers should specify the physical meaning they attach to simulation times. Does a simulation period correspond to one year, to one quarter, or to one day?

Let us see in more detail the practical implications stemming from these three requirements.

If the original endowments and the other starting conditions of artificial agents are initialized with actual data, one of the major criticisms levelled to ABC models (especially for their use as policy tools) – that is, the initial homogeneity (apart from random dislocations) of the agents belonging to the same class – can be easily avoided. While in many cases it is interesting to study the emergence of heterogeneity from homogeneous initial conditions *per se*, it is also true that during a simulation run it is rather difficult to isolate from initial transitional phases the portion in which the system can actually reproduce the desired microeconomic *and* macroeconomic phenomena. By the same token, the initial running-in periods are particularly critical when performing sensitivity analysis, since micro state variables are generally unstable, so that it is not easy to understand their marginal impact in the model.

Similar issues hold as we regards agents' behavioural rules. In many ABC models, microeconomic rules of conduct are so discretional and based on stochastic disturbances that it is practically impossible to discern their impact on the whole system, apart from very trivial cases. Moreover, if behavioural rules are not based on data observations, it remains unclear how could an ABC model be possibly considered as a policy-oriented tool.

Finally, any modelling choice on the physical meaning one must assign to simulation time periods is strictly related to the preceding points. If a model is initialized with actual data and micro-founded rules, comparisons of simulation results and actual data are much easier, and the physical meaning to be attached to simulation times can be defined naturally. For example, if we consider yearly actual data and we discover that a distance metric between them and simulations results is minimized on the average every two simulation periods, then we can assert that each simulation period corresponds to a semester. If, on the contrary, such a distance is minimized in every single period then one simulation time is equal to one actual year. It goes without saying that a correct understanding of the physical meaning of simulation time is also essential for policy purposes.

These practical issues are at the heart of a methodological distinction between two general categories of ABC models: *i*) *pure* ABC models, which must not necessarily be taken to the data to preserve their theoretical relevance as soon as microspecifications are plausible, and that are mainly useful to explain through generative techniques the mechanics of emergent aggregate phenomena under different rules of behaviour; *ii*) *applied* ABC models, which are conceived from the start in a form which admits microsimulation. While pure agent-based models, apart from some heuristical observations regarding emerging statistical regularities and evidences, are not externally validable almost by definition, applied ABC models need and have to be descriptively validated (i.e., compared with actual data) and calibrated (i.e., all the parameters' values must be chosen with the aim of better approximating reality). Only microsimulated ABC models can be effectively validated in a complete statistical sense, either *ex-ante* (in terms of the microfoundation on data and empirical evidence of behavioural rules), *ex-post* (in terms of a comparison over time between simulations' outcomes and actual data), and *simultaneously* (by the means of parameter calibration).

In order to accomplish this task, both the availability and the quality of microeconomic data is essential. As we shall discuss in the next lines, the BAM model of Chap. 3 possesses some components for which actual data are not (or not completely) readily available, in particular those referring to the entry and exit processes of firms and to the characteristics of the transactional networks linking agents in the markets for consumption goods and credit.[49] In what follows those parts will be considered in terms of degrees of freedom of the model, and thus not taken into account in our validation experiment. The bulk of the analysis will focus instead on firm-level balance sheet data, industrial dynamics and labour incomes.

4.3 Data Description

The greater part of the validation experiments and of the empirical analysis we present in this chapter are based on firm-level observations from the CEBI database. CEBI, which is probably the most comprehensive Italian dataset collecting balance sheet information on business units, was originally set up by the Bank of Italy, and it is actually maintained by *Centrale dei Bilanci Srl*.[50] In particular, we shall consider annual data, over the period 1988–2002, on a balanced panel of about 25,000 Italian non-financial firms, all satisfying the following requirements:

- reliable data on the level of real productive assets (capital);
- a net worth of at least 20,000 euros;
- at least 5 employees and labour costs of at least 20,000 euros per year.

Reliability of the data at the firm level essentially means that we exclude outliers characterized by unrealistic levels and changes of total assets, defined as the sum of net worth and total liabilities. The thresholds imposed at the level of 20,000 euros of net worth and of 5 employees aims at eliminating from the pool all really small firms, whose behaviour is usually erratic. For each firm and year, we can retrieve from CEBI the following variables from the dataset: firm ID, number of employees, total labour costs, equity, total assets, current liabilities, interest payments, sales, operating turnover, profits and losses, ROI, ROE, and the gearing ratio.

The use in our validation experiment of a balanced panel comes with both strengths and weaknesses. On the one hand, it assures continuity in the data, thus simplifying the analysis and allowing for clear-cut conclusions. On the other hand, it rules out by construction entry-exit processes. This is a significant limitation, since bankruptcy plays a crucial role in the BAM model. A possible way out of this problem could be the use of a complementary dataset composed of a subset of firms, from which to pick a substitute every time a firm exits from the panel. Even

[49] This does not mean that these data could not be collected in principle. Think for instance to the mass of proprietary data which could be easily retrieved (… if priced at their marginal costs) from the electronic records registering the transactions executed in large chain-stores by possessors of fidelity cards.

[50] See the *Centrale dei Bilanci* website at http://www.centraledeibilanci.it.

abstracting from the obvious critique that these substituting firms would not represent new entries in a proper sense, this solution would create even bigger problems concerning the "representativeness" of the firms' size distribution and its stability. For all these reasons, we decide to stick to the simpler "balanced" approach.

To complete the information about the labour market, we shall also use the so-called WHIP (Work Histories Italian Panel) database, developed by the Laboratorio Revelli of the University of Turin.[51] WHIP is a database of individual work histories, based on administrative archives at INPS (the Italian Social Security Institute), in which the main episodes of the working career of each individual in the sample are observed and recorded. The reference population is made up by all the people (both Italian and foreign citizens) who have worked in Italy even only for part of their career. The complete list of available observations includes: private employee working contracts; atypical contracts; self-employment activities as artisan, trader and some activities as freelancer; retirement spells and unemployment spells during which the individual received social benefits, like unemployment subsidies or mobility benefits. The sample is really large and representative of the total population of employees in the private sector. Workers in the public sector and professionals (lawyers, engineers, etc.) – who are covered by a social security provider different from INPS – are not recorded in WHIP, however.

4.4 Validation Procedure

The validation procedure we propose in this context is a mix of established statistical techniques and other innovative methods only recently applied in this area (see Bianchi *et al.*, 2007, 2008), and it is largely based on well-known results developed in the fields of computational models' evaluation (Kleijnen, 1995) and extreme value theory (Embrechts *et al.*, 1997). In particular, the iterative validation procedure followed below is made up of three steps:

1. a microsimulation of the model using actual data to calibrate it;
2. a descriptive validation of the simulated output;
3. if the output validation is satisfactory, then stop; otherwise re-calibrate the model and re-start from point 1.

Microsimulation is a technique operating at the level of individual units, like households or firms, that are always required to be initialized with actual data. Within the model, each unit is represented by a record containing a unique identifier and a set of associated attributes. A collection of rules are then applied to these units, leading to simulated changes in states and behaviours. Rules may change according to deterministic drivers, such as changes in tax liabilities resulting from changes in tax regulations, or because of stochastic variations. In both cases, as a result of applying these empirically based rules of action over many

[51] The dataste is available at http://www.laboratoriorevelli.it under the WHIP section.

time steps one gets numerical values for the state variables characterizing each unit, as well as for the distributional features of any change.

Bianchi *et al.* (2007) have proposed the use of microsimulation techniques as a means for validating agent-based models. The idea is rather simple: if an agent-based model is meant to reproduce empirical evidence to some extent, once it is initialized with actual data it should be able to reproduce real world phenomena with a certain level of precision. In other words, the model is first set up with the available actual data. Then it is simulated without any further intervention. Finally, simulated and actual data are compared at the end of the simulation and for each intermediate time period. From this point of view, any ABC model is well validated if its microsimulated version can sufficiently approximate reality. The higher the ability of the model in fitting real-world data, the more the model is judged as reliable. The microsimulated agent-based model can then be the starting point of a calibration exercise, meant to reduce the distance between actual and simulated data. Different techniques are available for these purposes. A good example is represented by indirect inference, as shown in Bianchi *et al.* (2007) and Gilli and Winker (2003).

There are also several graphical and statistical tools available either to perform a descriptive analysis of results, as well as to compare actual and simulated data in order to determine to what extent the microsimulated model reproduces observed data. First of all, the main characteristics of the data can be assessed by means of simple distributional fitting exercises (histograms, smoothed density plots, Zipf's plots, box-plots). Second, a comparison between simulated and actual data can be based both on graphical tools (quantile-quantile [Q-Q] plots, probability-probability [P-P] plots, mean excess function versus threshold [MEPLOT] plots), and on related tests aimed at weighing against the basic descriptive statistics and position means of the two data sets (Kruskal-Wallis tests and empty box tests). Finally, the corresponding distributions can be further analyzed and compared using formal goodness-of-fit tests, like the Kolmogorov-Smirnov (based on the supremum metric), the Cramér-Von Mises (which makes use of a quadratic measure of distance), and the Anderson-Darling (very powerful but not distribution-free, since critical values depend on the distribution under scrutiny). In the next sections we shall show how these tools and techniques can be used to assess the validity of the BAM model.

4.5 Firms' Size and Growth

A profusion of empirical investigations in industrial dynamics have detected two cross-sectional empirical regularities, which are amazingly widespread across countries and persistent over time:

1. The distribution of firms' size (measured in terms of total assets, sales and employment) is in general right skewed, and it can often be well described by a power law probability density function (Axtell, 2000; 2001; Gaffeo *et al.*, 2003; Okuyama *et al.*, 1999; Quandt, 1966; Ramsden and Kiss-Haypal, 2000; Simon, 1955).

2. Firms' growth rates are Laplace distributed, a model which can be nested in the larger and more flexible Subbotin's family of distributions (Stanley *et al.*, 1996; Bottazzi and Secchi, 2005).

Delli Gatti *et al.* (2008) have already shown analytically and computationally that facts 1 and 2 (in addition to many other distributional regularities) are at the heart of several financial and business cycle facts. To be considered a model capable of generating a macrostructure of interest, also the BAM model should be able to replicate such an empirical evidence. The first part of the validation exercise is therefore focused on this.

Let us start by considering total assets. Accepting a maximum deviation of ± 15% between observed and simulated data, in 2002 (the last year for which we have data) we succeed in reproducing 20888 firms, that is almost 84% of the sample. This is a remarkably good percentage, and similar values can be found in all previous years. Also the distribution obtained by pooling observations over time is correctly fitted for 20391 productive units over a total of 24867 (82%).[52] Fig. 4.1 reports the Zipf's plot of actual and simulated data for the year 2002, showing that the two distributions are almost overlapping. Both observed and simulated capital distributions are particularly skewed, with really fat right tails. Most of the actual firms which are not correctly approximated by their artificial clones are concentrated in the very upper tail, indicating that the model is not fully adequate in reproducing the dynamics of very large companies. This difficulty is likely due to the peculiar structure and dynamics of large firms, whose behaviour is known to be sensibly different from those of other firms (Cirillo and Hüsler, 2009).

In order to assess whether the goodness of fit reported above is due to the adequacy of the model rather than to the characteristics of actual data, we further perform an empirical analysis aimed at verifying how the real firms comprised in our sample changed in size during the time span considered. Indeed, it could be the case that the model does a good job in fitting the real-world system in a given year merely by chance. If the variability of actual data over time is small, therefore, one could erroneously attribute a dynamic similarity between the simulation outcomes and the data to the model's generative capability. On the contrary, if the actual data vary a lot over time and − in spite of this − the BAM model does a god job in fitting them, the confidence we can pose on its adequacy is necessarily higher.

In Fig. 4.2 one can perform a comparison, by means of Zipf's plots, among the cross-sectional distributions of actual balance sheet data on total assets in 1988, 1995 and 2002. The inference from simple visual inspection is quite clear: there is a substantial difference between the three sets of data. If we accept the usual ± 15% confidence interval, less than 17% of the firms (essentially the ones that have changed a little over time) can be considered as contemporaneously fitting the data in all samples. In fact, even if all distributions seem to belong to

[52] As in Ijiri *et al.* (1977), the use of pooled distributions is possible since the yearly distributions show similar slopes.

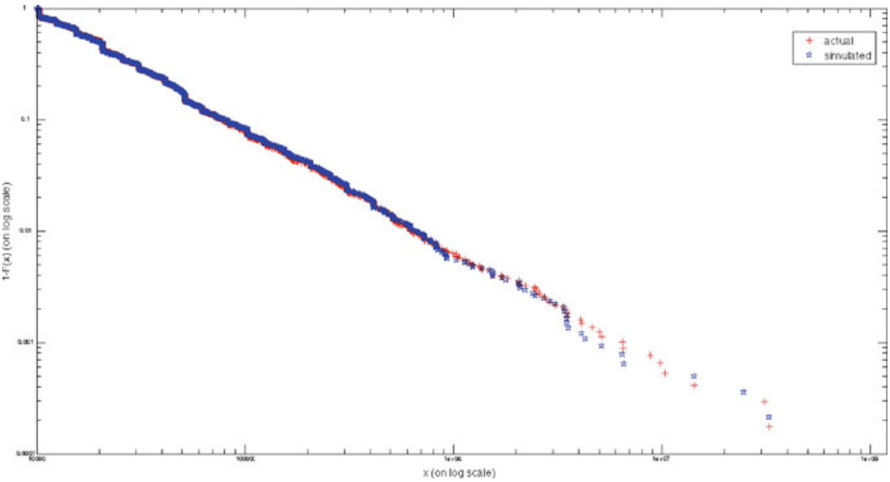

Fig. 4.1 Zipf's plot of total assets distributions: actual and simulated in year 2002

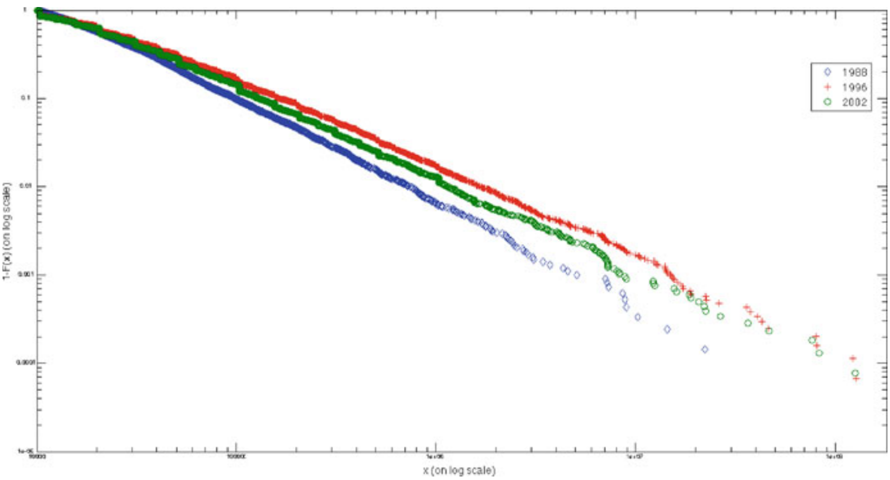

Fig. 4.2 Evolution of actual total assets in years 1998, 1995 and 2002

the same family, their shape and scale parameters are sensibly different. Several additional analytical tools, such as the Kolmogorov-Smirnov's statistics and the Anderson-Darling's test for power laws, confirm this result.

Another informative test can be easily built by calculating the average growth rate of actual firms between 1988 and 2002, to subsequently let the artificial firms – calibrated with actual balance sheet data for the starting year 1988 – grow at this factor in each subsequent simulation period. If the actual growth process is trivially linear, this experiment should allow artificial firms to reproduce actual data in 2002 with an acceptable level of precision. However, the result obtained in this way is rather poor: only 44% of the sample in 2002 is correctly fitted, against a percentage of 84% of the BAM model.

Once the issue of whether the BAM model is able to generate macroscopic regularities of interest from empirically-based microspecifications has been positively addressed, we are in a good position to further assess the degree of matching between the model and the real data. To do this, we make use of several graphical and analytical tests aimed at checking if the observed and simulated samples can be considered belonging to the same distribution. For instance, as shown in Fig. 4.3 a Q-Q plot computed for the year 2002 clearly supports the hypothesis that both samples are drawn from the same statistical distribution. Similar findings – not shown here to save space – can be obtained in all the other years. These results can be further corroborated by means of alternative graphical methods (for instance, box-plots), and formal statistical testing procedures (like generalized Kolmogorov-Smirnov [Prabhakar *et al.*, 2003] tests at the 95% confidence level).

As we exclude from the samples – for the reasons reported for instance in Cirillo and Hüsler (2009) – the right upper tails with a cut-off value at the top 12%, both the actual and the simulated data follow a Singh-Maddala distribution. Table 4.1 reports the maximum likelihood point estimates of the relevant parameters for four representative equally spaced years. Similar estimates can be obtained for all the other years not reported here.

The MEPLOT reported in Fig. 4.4 suggests that the right tails of the two distributions display a Paretian behaviour, as signalled by the upward slope of their mean excess function. In order to see whether the two samples possess a similar behaviour, we therefore estimate the shape parameter characterizing their right tails by means of the semiparametric Hill's method.

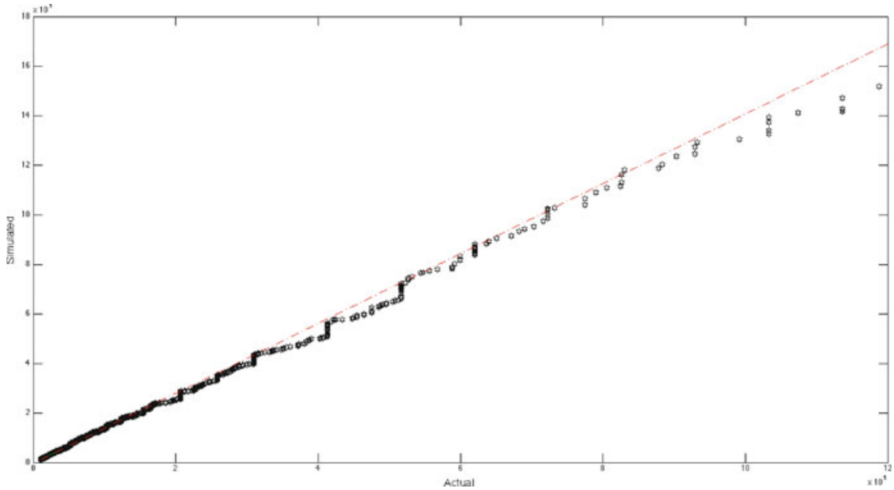

Fig. 4.3 Q-Q plot for actual and simulated total assets in year 2002

Table 4.1 Estimates of the parameters of the Singh-Maddala distribution for actual and simulated data in four different years

Year	*Data*	α	β	γ	ρ
1988	Actual	1.61	93.114	0.96	7.22
	Simulated	1.63	92.737	0.96	7.26
1992	Actual	1.64	97.333	0.98	6.89
	Simulated	1.66	98.122	1.01	6.93
1996	Actual	1.81	99.282	0.93	7.02
	Simulated	1.85	100.071	0.91	7.04
2002	Actual	1.98	101.668	0.99	6.97
	Simulated	1.96	102.002	0.97	6.96

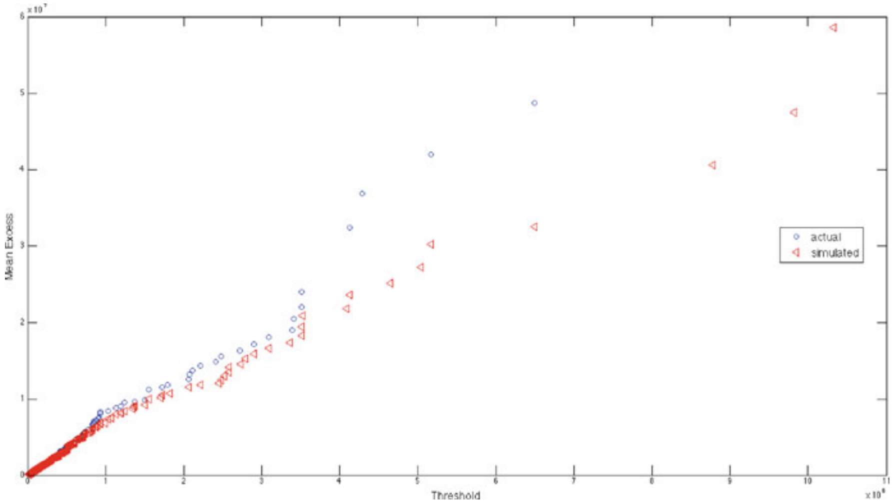

Fig. 4.4 Mean Excess Function plot for actual and simulated total assets

The point estimate we obtain for simulated data is $\alpha = 1.81$, while for the actual sample the corresponding value is $\alpha = 1.79$. The two parameters are quite similar and belong to the Paretian support ($0.8 < \alpha < 2$). Notice, however, that the distribution for simulated total assets shows a slightly thinner tail, a finding which further confirms the difficulty of the present version of the BAM model in replicating the behaviour of really large firms, and a pointer for further research.

The following step consists in the empirical validation of the net worth of firms. If we agree to the usual ± 15% limit band for deviations between actual and simulated data in 2002, we are able to reproduce the individual histories of 74% of the population. This figure, although largely satisfactory in itself, is somewhat lower than that for total assets. Such a result is likely due to two basic aspects of the BAM model which have been introduced at this stage for the sake of computational tractability: 1) the equity market is not present in the model, and this influences the overall dynamics of firms' balance sheet items, especially of the largest ones; 2) the BAM model lacks enough structure in the formulation of how R&D investments are funded out of retained profits, which are currently modelled in terms of linear relationships.

Apart from these drawbacks which must be fixed in future versions, from a distributional-fitting point of view the BAM model works pretty well also in the case of net worth. As shown by the Zipf's plots reported in Fig. 4.5, the two distributions for the actual and simulated data − calculated in the last period of simulation (2002) − present an unambiguous Paretian behaviour. Hill's estimates of the shape parameters are $\alpha = 1.73$ for actual data, and $\alpha = 1.77$ for simulated data.

The results we get about loans are very similar to those of total assets, since we succeed in fitting about 80% of the firms. Similarly to total assets and net worth, both graphical and analytical tests support the idea of a common density function for both the actual and the simulated debt's data generating processes. It is interesting to notice that, as in Fujiwara (2004), the distribution of loans is in fact a power law. The Hill's estimates of the shape parameters of the Paretian right tails are $\alpha = 1.67$ and $\alpha = 1.64$ for the actual and for the simulated data, respectively, signalling a slight overestimate of the model as regards the largest firms.

The next issue to be explored is the growth dynamics of firms, measured once again in terms of total assets. Several empirical studies have pointed out that the empirical distribution of firms' growth rates is invariably tent-shaped (Hall, 1987; Axtell, 2000). In particular, the best-fitting candidates are the truncated-Lévy model (Gabaix *et al.*, 2003) and the Laplace (or double-exponential) model (Amaral *et al.*, 1997), which has been recently treated as a specialization of the

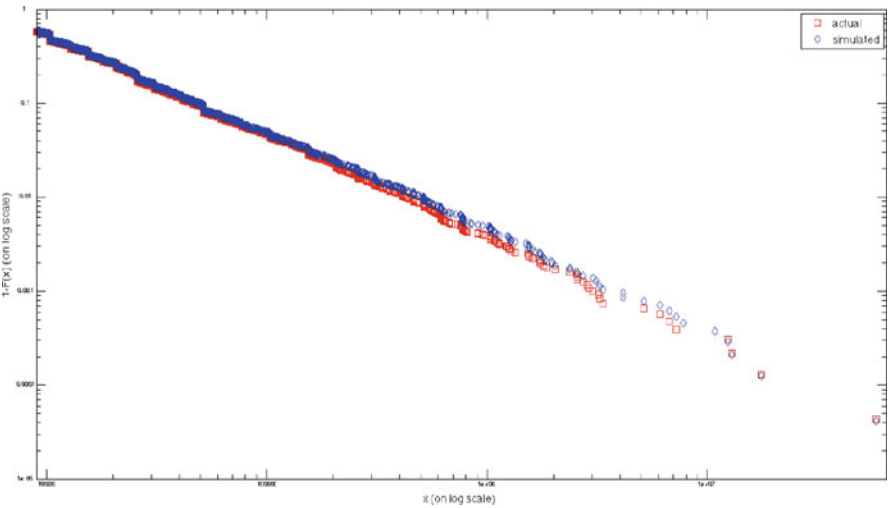

Fig. 4.5 Zipf's plot of net worth distributions: actual and simulated in year 2002

more general Subbotin (Subbotin, 1923) model (Bottazzi and Secchi, 2005). The functional form of Subbotin family of distributions is:

$$f(x,a,b) = \frac{1}{2ab^{\frac{1}{b}}\Gamma\left(1+\frac{1}{b}\right)} e^{-\frac{1}{b}\left|\frac{x-\mu}{a}\right|^{b}} \tag{4.1}$$

where μ is the mean, b and a are two different shape parameters, and Γ is the standard Gamma function. If $b \to 1$, the Subbotin distribution reduces to a Laplace, whereas for $b \to 2$ it approaches a Gaussian. The point estimates for the three Subbotin's parameters − obtained by means of maximum likelihood methods − as we fit this model to our actual and simulated data are reported in Table 4.2.

Table 4.2 Estimates of the Subbotin parameters

	Observed data	*Simulated data*
μ	0.001	0.001
a	0.067	0.061
b	1.006	1.013

At a first glance, the results obtained for observed and simulated growth rates display several similarities. The two means μ are very close to zero, as expected. Since the estimated values of b approach 1 in both cases, the two distributions are in the field of attraction of the Laplace model, excluding normality and, indirectly, ruling away the possibility that a pure Gibrat's law holds true either in reality and in simulations. The values for the Laplacian shape parameter a, finally, are quite similar, even if simulated data show slightly thinner tails. Overall, we can say that the BAM model is able to mimic firms' growth dynamics with a remarkably degree of accuracy. As already noted, the bulk of the problems in the empirical validation of industrial dynamics are concentrated in mimicking the behaviour of firms over the far right tail.

4.6 Interest Rates

Some other interesting results can be obtained by analyzing the functioning of the credit market, focusing in particular on the temporal path followed by interest rates on loans. The graph in Fig. 4.6 allows a visual comparison between the average nominal interest rate obtained in the micro-simulated version of the BAM model and the average interest rate paid by actual firms, as extrapolated from the CEBI database, over the whole time horizon 1988–2002. In particular, we notice that in both cases from 1988 to 1992 the average interest rate paid by firms on their debts has increased, to subsequently decrease steadily. The nearness of the

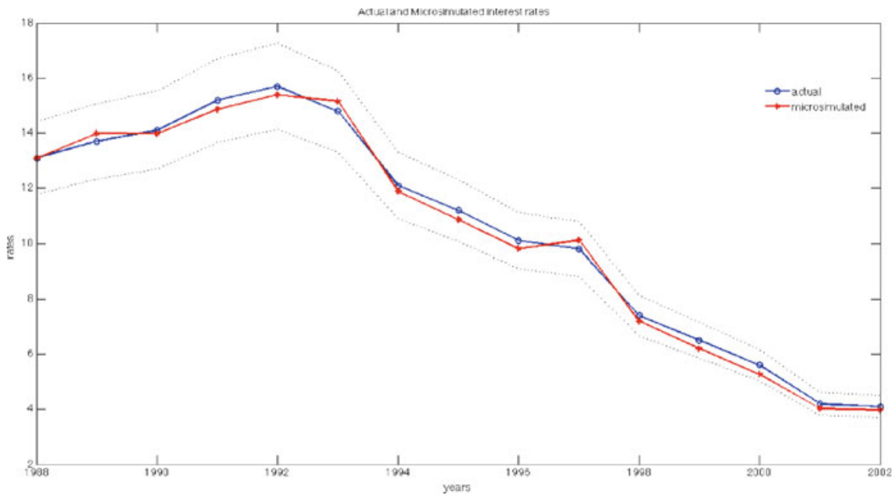

Fig. 4.6 Average interest rate paid on debt: actual and simulated data, 1988–2002

two time series is really remarkable, as the distance between them is always comprised inside a ± 6% deviation range from the mean calculated for the actual data.

In addition to the average rate, in simulations we are also able to track at any time step the interest rate paid to the banking sector by each single firm. If we accept the usual ± 15% deviation from the actual path as a measure of success, the model is able to reproduce the single-firm interest rate dynamics for 57% of the companies comprised in the CEBI dataset, a result we believe to be largely satisfactory given the simplicity of the market microstructure we use in the version of the BAM model discussed so far. Obviously, given the incompleteness of the data at our disposal, we cannot empirically validate other aspects of the model, such as the demand and the supply of loanable funds.

4.7 The Labour Market

The assessment of the labour market is for sure the most difficult part of the whole validation experiment. The ultimate reason is that the data available for the labour market are in general not as accurate and complete as the one for firms' balance sheets we have employed in the previous sections. In particular, the two key problems are that it is rather difficult to obtain accurate microeconomic data for effective wages and other labour-related variables (benefits, hours worked, and so on) on the one hand, and that of it is far from immediate to match these kind of data with the ones on production and credit referred to a particular firm, on the other one.

To partly overcome this difficulty, in this last section we perform an empirical validation of the BAM model's labour market by combining information from the

CEBI and the WHIP databases. As a matter of facts, the WHIP database contains panel-type information about the distributions of Italian wages by firms' dimensions, defined in terms of the number of employees. For all the five dimensional classes available ([0–9], [10–19], [20–199], [20–999] and [1000+]) for every year, a preliminary analysis of the data suggest that the wage distributions are three-modal and can be well approximated by a mixture of three Gaussian distributions, even if better fits – for example by means of Bernstein's polynomials – are obviously available (see Cirillo, 2009). Thanks to the conditions we have previously imposed on the CEBI dataset, we proceed to exclude all firms belonging to the two smallest classes. In other words, in order to guarantee consistency with our previous analysis, we eliminate from the WHIP dataset all firms with less than 20 employees. Fig. 4.7 shows the wage distribution in 1988 for the three dimensional classes remained.

The labour market is then initialized as follows. First, for every firm we sample the wages received by its total work-force from the wage distribution of the dimensional class to which it belongs. Second, for every firm the sampling process is stopped when the total cost of wages is as close as possible to the actual one, as obtained from the CEBI dataset. In particular we accept the standard ± 15% difference. This means that for each single firm we may have a small difference between actual and initialized data but, on the average, actual total labour costs are well replicated by the micro-simulated model. If during the simulation a firm hires new workers, their wages are sampled from the corresponding wage distribution as usual.

Because of the lack of suitable data, in this block of micro-simulations the labour market mechanisms of the original BAM model is somehow simplified. In

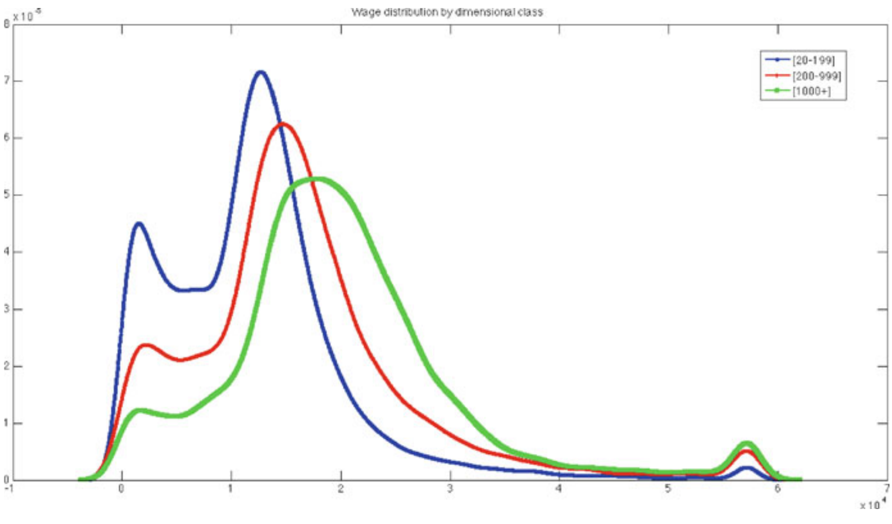

Fig. 4.7 Wage distributions for the classes of dimension [20–199], [20–999] and [1000+] in year 1988

particular, we abstract from the duration of labour contracts, while artificial firms have no particular constraints but the consistency of their total labour costs with those of their actual counterparts. As far as wages are concerned, their growth rate (the quantity ξ_{it} of the model) is set equal to the average actual growth in the WHIP dataset, for the three dimensional classes considered.

As we have seen in the previous sections, our agent-based model is able to reproduce the empirical evidence obtained from actual data with a good degree of precision. Also in the case of the labour market the model seems to be quite accurate, and its weaknesses should be interpreted as useful pointers for further research. In particular, while it is neither possible nor reasonable to look for validating aggregate empirical laws (Beveridge and Okun curves, for example) because of the actual incompleteness of the data, it is interesting to analyze the behaviour of the micro-simulated wage distribution. Specifically, it seems worthy to assess what happens at the end of the simulation, in order to gauge whether the model is able at least to replicate the three-modal feature of actual data. Fig. 4.8 reports the comparisons between the actual and the micro-simulated wage distributions in 2002 for the upper three dimensional classes.

Actual and micro-simulated distributions are quite similar, the scale is consistent, three-modality is present and the distances among the empirical and simulated cumulative density functions are not sufficiently small to pass a Kolmogorov-Smirnov test at the standard 5% significance level, but its statistic is almost borderline at the 10% statistical level. The availability of improved data would have allowed us to better calibrate the model's parameter using simultaneous techniques such as indirect inference or simulated annealing (Gourieroux and Monfort, 1996; Das and Chakrabart, 2005).

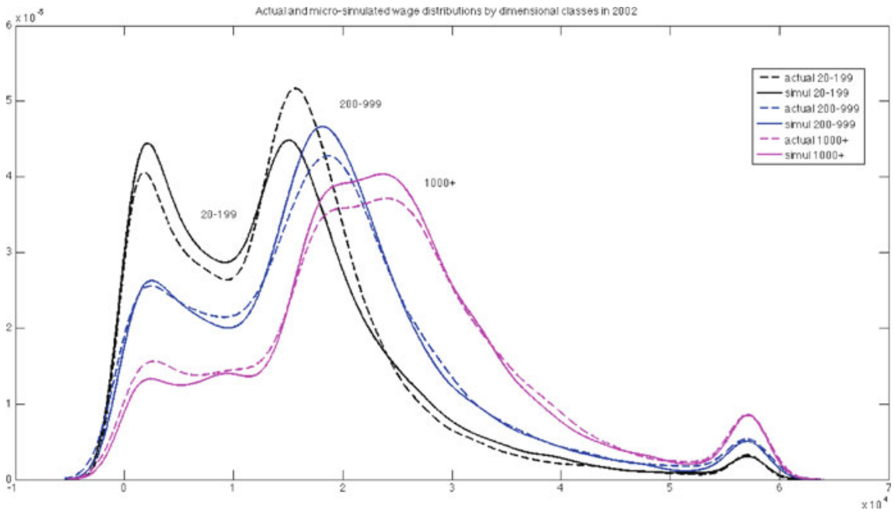

Fig. 4.8 Comparison of wage distributions between actual and simulated data for the classes of dimension [20–199], [20–999] and [1000+] in year 2002