1

Introduction

Domenico Delli Gatti and Mauro Gallegati

1.1 Hard Times for Dr Pangloss

High and persistent unemployment, over-indebtedness and financial instability, bankruptcies, domino effects and the spreading of systemic risk: these phenomena have taken center stage in light of the Global Financial Crisis.

By construction, the *Neoclassical approach* is much better suited to study the features of the world of Dr Pangloss (Buiter, 1980) than the intricacies of a complex, financially sophisticated economy. This point is well taken in the introduction of a seminal paper by Bernanke, Gertler and Gilchrist published well before the Global Financial Crisis: 'How does one go about incorporating financial distress and similar concepts into macroeconomics? While it seems that there has always been an empirical case for including credit-market factors in the mainstream model, early writers found it difficult to bring such apparently diverse and chaotic phenomena into their formal analyses. As a result, advocacy of a role for these factors in aggregate dynamics fell for the most part to economists outside the US academic mainstream, such as Hyman Minsky, and to some forecasters and financial market practitioners.' (Bernanke et al., 1999, p. 1344).

This candid admission by three of the most distinguished macroeconomics (one of them destined to be Chairman of the Federal Reserve for eight long and turbulent years) – which, incidentally, provides a long overdue implicit tribute to Hyman Minsky – also provides the research question for a model of the *financial accelerator* which has started a non-negligible body of literature in contemporary macroeconomics.

In order to put this development in macroeconomic thinking into context, it is necessary to recall that any mainstream macroeconomic model is based on a Dynamic Stochastic General Equilibrium (DSGE) skeleton, which can support either a Real Business Cycle (RBC) model or a New Keynesian (NK) model.

The latter differs from the former because of the presence of *imperfections*, the most important being imperfect competition and nominal rigidity. The structural form of the standard NK-DSGE framework boils down to a three-equation model consisting of an optimising IS equation, an NK Phillips curve and a monetary policy rule based on changes in the interest rate.

The NK-DSGE framework is, of course, too simple and therefore inadequate to analyse the emergence of a financial crisis and a major recession for the very good reason that neither asset prices nor measures of agents, financial fragility show up anywhere in the model. In order to make the model operational from this viewpoint, *financial frictions* have been incorporated into the basic model in one way or another.

In the last decade we have witnessed an explosion of models with these types of frictions. The story that can be attached to this literature, however, can be told in simple terms. A negative shock triggers a recession and yields a reduction of firms' internally generated funds. Borrowers need more funds, but lenders are less willing to supply loans as the value of firms' collateral is also going down. Hence, firms might be forced to scale activity down. This in turn will lead to lower cash flow, and to a *further deceleration* of activity.

The financial accelerator provides a mechanism of *amplification* of an aggregate shock (i.e., a positive feedback or a self-reinforcing mechanism) based on financial factors. By construction, however, it cannot be a model of the *origin* of a financial crisis and the ensuing recession.

As in all DSGE models, in fact, in models characterised by financial frictions a fluctuation is also determined by an aggregate shock (an impulse) and is channeled to the economy by a propagation mechanism. Moreover, the stability of the steady state makes fluctuations persistent but relatively short lived. Therefore, a great recession may be explained only by a sizable aggregate negative shock and is bound to disappear relatively soon. Recent models incorporating financial frictions trace back the great recession to a major negative shock (possibly of a new type: an 'investment shock', a 'financial shock' instead of the usual Total Factor Productivity shock) which spreads through the economy and becomes persistent because of the financial amplification, but is temporary so that the economy goes back to equilibrium in due time.

This view of the Global Financial Crisis is not convincing. It does not provide a plausible theory of its origin, since the crisis was not the consequence of a global shock, but originated from a small segment of the US financial system (the subprime loans market) and spread to the entire US financial system and to the world economy. Moreover, it does not provide an appropriate characterisation of the actual recovery, which has been unusually long and painful.[1] In fact, during the recession, quantitative forecasts of future GDP growth (also at a very short-time horizon) generated by these models systematically overestimated actual GDP growth.

The financial accelerator story is intriguing, but is not enough to characterise a major crisis. Models with financial frictions yield interesting results, but their scope is necessarily limited because of the built-in features of the DSGE framework. This framework, in fact, abstracts from the complex web of financial and real relationships among heterogeneous agents that characterise modern financially sophisticated economies and are at the root of the spreading of financial fragility economywide. Contemporary macroeconomics, in other words, has assumed away most of the features of the economy which are relevant today.

1.2 The Complexity View

For several years now, a different approach has been developed which conceives the macroeconomy as a *complex system* of heterogeneous agents characterised by bounded rationality, endowed with a limited and incomplete information set, interacting directly and indirectly with other agents and the environment.

In a complex economy, an idiosyncratic shock – i.e., a shock to a specific agent – can well be the source of an epidemic diffusion of financial distress. In other words, idiosyncratic shocks do not cancel out in the aggregate, especially if the macroeconomy is characterised by an underlying network structure and the idiosyncratic shocks hit crucial nodes of the network. Therefore a recession may not be caused only by an aggregate shock.

To be specific, in a credit network, the financial accelerator can lead to an avalanche of bankruptcies due to the positive feedback of the bankruptcy of a single agent on the net worth of the 'neighbours' linked to the bankrupt agent by credit relationship. This is, of course, ruled out by construction in a framework with a representative firm and a representative bank.

In order to deal with these issues, one has to start with a population of heterogeneous agents. *Heterogeneity*, therefore, is key in modelling the phenomena which we want to investigate.

[1] The idea of a 'secular stagnation' pioneered by L. Summers, which is gaining ground in the profession, is based exactly on the unusual length and painfulness of the recovery from the Great Recession.

1.3 Heterogeneity in a Neoclassical World

The way in which financial frictions have been incorporated in current macroeconomic models provides an example of a recurrent pattern in the development of contemporary macroeconomics. Research issues brought to the fore by new macroeconomic facts are incorporated into an analytical edifice based on neoclassical foundations by means of appropriate twists of some assumptions or by additional assumptions, as epicycles in Ptolemaic astronomy. This is the way in which 'normal science' (in Kuhn's wording) adjusts to real macroeconomic developments. In this way, there is nothing truly new under the sun.[2]

Heterogeneity has been incorporated in Neoclassical models since the early 1990s. This is an impressive achievement, as the Neoclassical approach is utterly unsuitable for the study of this issue. In a Neoclassical Representative Agent-Rational Expectations world, equilibrium prices depend on a relatively small number of state variables and shocks. Forming Rational Expectations of future prices in such an environment is a daunting but not impossible task, as the Representative Agent in the model has the same information of the modeller herself, the 'true' model of the economy included (at least in reduced form).

Things are much more complicated in a multiagent setting. In this case, equilibrium prices are in principle a function of the entire distribution of agents (e.g., households' wealth). Hence, to form expectations, agents need to know the entire distribution at each point in time, i.e., the law of motion of this distribution. An impossible task. This is the well known 'curse of dimensionality'.[3] Neoclassical Heterogeneous Agents Models have been developed in order to study the causes and consequences of income and wealth inequality in a DSGE framework.[4]

These papers by Imrohoroglu, Hugget, Aiyagari essentially relax only the Representative Agent assumption (and only as far as households are concerned), but generally retain all the other conceptual building blocks of

2 As Max Planck put it: 'Normal science does not aim at novelties of fact or theory and, when successful, finds none.'

3 One possible way is to keep the number of agents low, i.e., to reduce the dimensionality of the problem (two types of agents). An example, among many, is the NK-DSGE framework with Limited Asset Market Participation, where households can be either free to access financial markets or financially constrained. In the latter case, households cannot smooth consumption by borrowing and lending.

4 These models are also known as Bewley models (according to the terminology proposed by Ljungqvist, 2004) or Standard Incomplete Markets models (Heathcote, 2009). Notice that Heterogeneous Agents and Incomplete Market go hand-in-hand, for reasons which will become clear momentarily.

DSGE models (intertemporal optimization, continuous market clearing and general equilibrium).[5]

In Aiyagari (1994) households are heterogeneous because of idiosyncratic shocks to earnings. If the markets were complete, these shocks would be insurable and therefore they would not impinge on average or aggregate consumption (they would wash out in the aggregate).[6] Under these circumstances, the long-run or ‘equilibrium’ distribution of wealth would be both indeterminate and irrelevant (because any distribution of wealth would yield the same average behaviour).

If markets were incomplete, on the contrary, the possibility to insure against idiosyncratic risk would be limited and therefore: (1) idiosyncratic risk would impact on consumption (and macroeconomic performance), (2) the equilibrium distribution of wealth would be determinate. In Aiyagari’s model particularly, inequality yields precautionary savings which impact positively – through investment – on growth.

Research in this field has been extended in at least three directions: (1) the analysis of other sources of heterogeneity (besides idiosyncratic shocks to earnings), e.g., innate characteristics; (2) the analysis of additional ways to insure idiosyncratic shocks; (3) the impact on aggregate fluctuations (see Heathcote, 2009). Focusing on the third line, we will single out the pioneering paper by Krusell and Smith (Krusell and Smith, 1998) as typical of the approach.

Krusell and Smith circumvent the curse of dimensionality in a very smart way. They summarise the shape of the agents’ distribution by means of a finite number of its moments. In this way they can abstract from the actual distribution and be as precise as they wish in describing its shape: the larger the number of moments considered, the more granular the description of the distribution. For simplicity, they use only two moments: The first moment (mean) captures the central tendency of the distribution; the second moment (variance) captures the dispersion, one aspect of the degree of heterogeneity characterising the distribution. When dispersion is low, the mean of the distribution is almost sufficient to describe the distribution itself. Therefore

[5] See Rios-Rull (1995) for a review of early work in this area.

[6] Completeness and homothetic preferences imply a linear relationship between consumption and wealth at the individual level, i.e., Engel curves are linear. In this case, perfect approximation applies: average consumption will be a linear function of average wealth. Only the first moment of the distribution of wealth is necessary to determine average (and aggregate) consumption. Heterogeneity, as captured by the variance and higher moments of the distribution, is irrelevant. Of course this is no longer true if the relationship between consumption and wealth at the individual level is nonlinear. If the relationship were concave, for instance, an increase in the dispersion of wealth would lead to lower consumption on average – thanks to Jensen inequality – even if the mean were preserved.

higher moments of the distribution can be safely ignored and one can think of the economy as if it were a Representative Agent world, identified with the Average Agent.

Agents in a Krusell-Smith economy are ‘near rational’ as they optimise using only the moments of the distribution. Forming near rational expectations of future prices in such an environment is a daunting but not impossible task, as equilibrium prices are functions only of the moments of the distribution instead of the entire distribution.

In this model there is *approximate aggregation*: ‘in equilibrium all aggregate variables … can be almost perfectly described as a function of two simple statistics: the mean of the wealth distribution and the aggregate productivity shock’ (Krusell and Smith, 1998, p. 869). Using only these measures, near-rational agents are able to minimise the forecast errors (therefore higher moments of the distribution do not affect the decision of the agents).

Moreover, Krusell and Smith show through simulations that macroeconomic time series generated by the model are almost indistinguishable from those generated by a Representative Agent model. Hence macroeconomic fluctuation can be sufficiently described by fluctuation of the mean; higher moments of the distribution do not add much to the picture. In other words, taking heterogeneity on board does not add much to the accuracy of the model. Only the first moment of the distribution has macroeconomic consequences. In a sense, this is a very smart way of resurrecting the moribund Representative Agent and the macroeconomic literature based on it.

However, as shown by Heathcote (2009), with reference to fiscal shocks, there are indeed real-world circumstances in which heterogeneity has important macroeconomic consequences, even in Neoclassical multiagent models.

1.4 Agent-Based Models (ABMs)

The research agenda of the Neoclassical multiagent literature is very specific, dictated by the self-imposed guidelines on the way in which economic theorising should take form in the Neoclassical approach. Heterogeneity, therefore, is key in these models, but is dealt with in a restricted, disciplined environment. This may be considered a virtue of the approach, but can also be a limitation as issues and problems which are indeed important ‘out there’ in the real world are left out of the admissible set of issues and problems to be dealt with. *Agent Based Models* (ABMs) have a completely different origin and a much more flexible agenda. ABMs are the analytical and computational tools developed by an interdisciplinary network of scientists – physicists, economists, computer scientists – to explore the properties of *Complex Adaptive Systems* (CAS), i.e., 'systems comprising large numbers of coupled elements the properties of which are modifiable as a result of environmental interactions . . . In general CAS are highly nonlinear and are organised on many spatial and temporal scales' (1st workshop on CAS, Santa Fe, 1986).

In ABMs a multitude of objects, which are heterogeneous under different profiles, interact with each other and the environment. The objects are autonomous, i.e., there is no centralised (top-down) coordinating or controlling mechanism. Therefore, ABMs cannot be solved analytically. The output of the model must be computed and consists of simulated time series.

Agent-based Computational Economics (ACE) is the application of agent-based (AB) modelling to economics or: 'The computational study of economic processes modelled as dynamic systems of interacting agents' (Tesfatsion and Judd, 2006, p. 832). The economy, in fact, can be conceived of as a complex adaptive system.

Behavioural equations may or may not be derived from optimization. AB modellers generally prefer to assume that agents are characterised by bounded rationality; they are 'not global optimisers, they use simple rules (rules of thumb) based on local information' (Epstein, 2006a, p. 1588).[7]

No equilibrium condition is required (out-of-equilibrium dynamics). This is, in a sense, a consequence of the assumption according to which there is no top-down coordinating mechanism in ABMs. The Walrasian auctioneer, who is gently nudging the agents towards an equilibrium position, is indeed a metaphor of such a top-down coordinating mechanism. AB modellers, in fact, generally prefer to assume that markets are systematically in disequilibrium. In principle, however, at least some markets may converge to a statistical equilibrium.

ABMs are built from the bottom-up. At the micro-level, the behaviour of heterogeneous agents is captured by simple, often empirically based heuristics which allow for adaptation, i.e., gradual change over time in response to changing circumstances. Aggregate variables are determined by means of summation or averaging across the population of heterogeneous agents. Instead of untying the Gordian knot of aggregation, in ABMs this is cut by allowing the computational tool to do the job. Due to interaction and nonlinearities, statistical regularities emerge at the macroscopic level that cannot be inferred from the primitives of individuals. These *emergent properties* are at the core of macroeconomics in a complex setting. Generally, aggregate variables in

[7] In principle, however, behavioural rules can be either grounded in bounded rationality (rules of thumb) or can be derived from specific optimization problems (optimal rules).

macroeconomic ABMs (e.g., GDP) show a tendency to self-organise towards a stable aggregate configuration occasionally punctuated by bursts of rapid change. The self-organisation of the macroeconomy can be represented by a *statistical equilibrium* in which the aggregate spontaneous order is compatible with individual disequilibrium. The equilibrium of a system no longer requires that every single element be in equilibrium by itself, but rather that the statistical distributions describing aggregate phenomena be stable, i.e., in ‘a state of macroscopic equilibrium maintained by a large number of transitions in opposite directions’ (Feller, 1957, p. 356). This is not general equilibrium in the standard meaning, i.e., a state in which demand equals supply in each and every market.

In a macroeconomic ABM – i.e., an ABM applied to the macroeconomy – a ‘crisis’, i.e., a deep downturn followed by a long and painful recovery, is a macroscopic phenomenon which spontaneously emerges from the web of microscopic interactions. In a macro ABM, in other words, big shocks are not necessary to explain big recessions, an appealing property indeed in light of the Global Financial Crisis.

The real-world phenomena that are conceived of as rare ‘pathologies’ in the Neoclassical view – high and persistent unemployment, over-indebtedness and financial instability, bankruptcies, domino effects and the spreading of systemic risk – turn out to be the spontaneous emerging macroscopic consequence of complex interactions in a multiagent framework with heterogeneous agents.

1.5 Plan of the Book

The main aim of this book is to provide an introduction to Agent-Based Modelling methodology with an emphasis on its application to macroeconomics.

The book is organised as follows. In Chapter 2 we answer the most basic questions: What is an ABM? When is it necessary and/or appropiate to build such a model? The chapter ends with a succinct overview of a very early example of ABM: Schelling’s model of racial segregation.

In Chapter 3 we provide a formal characterisation of an ABM as a recursive model. We put ABMs in the wider context of simulation models and introduce notation and key concepts to describe the agents’ state and behavioural rules in ABMs in rigorous terms.

Chapter 4 is devoted to a general overview of rationality, the determination of behavioural rules and expectation formation in contemporary macroeconomics, from Keynesian aggregative models to macroeconomic ABMs passing through monetarist, New Classical and New Keynesian models. This survey allows us to put the AB methodology into context and paves the way to the more detailed analysis of behavioural rules and learning processes in Chapter 5.

Chapter 5, in fact, digs deeper into the definition and description of the agent's rationality and learning processes. Where do behavioural rules come from? Agents in real economies are intentional subjects. In order to decide on a line of action (a behavioural rule), in fact, they must form mental representations of their environment and of the behaviour of other agents. Behavioural rules in the real world, therefore, must be related to the cognitive processes that guide actions. Learning is a key ingredient of these cognitive processes.

Chapter 6 deals with the issue of interaction, which is key in ABMs. In a sense, the chapter is an introductory overview of network theory, a rapidly expanding field both in mainstream and complexity theory. ABMs, in fact, are often based on an underlying network structure, e.g., of trading relationships, credit relationships, supply chain, etc.

Chapter 7 explores the research outcome of an ABM, i.e., the model behaviour. The AB researcher, in fact, sets up the 'rules of the game' – i.e., she builds the model – but does not know in advance the implications of those rules, e.g., the statistical structure of the output of simulations. The chapter presents techniques to gain understanding about the model behaviour – the Data Generating Process implicit in the ABM – which are quite underexplored in the AB literature. In a model which requires simulations, only inductive knowledge about its behaviour can be gained, by repeatedly running the model under different samples from the parameter space.

Chapter 8 is devoted to the empirical validation of ABMs. Empirically validating an ABM means, broadly speaking, 'taking the model to the data', essentially in the form of empirical and/or experimental data. Validation may concern the model inputs and/or outputs. Input validation is essentially the assessment of the 'realism' of the assumptions on which the model rests while output validation is the assessment of the capability of the model to replicate in artificial or simulated data the stylised facts of economic reality under consideration. Output validation is a joint test on the structure of the model and the values of the parameters. This means that input and output validation are connected.

Chapters 9 deals with the issue of estimation of ABM parameters, an intriguing new field which aims to align the empirical validation techniques of ABMs to that of standard macroeconomic models where estimation tools are readily available and relatively easy to implement.