The future of Agent-Based modelling.

Matteo G. Richiardi

Institute for New Economic Thinking and Nuffield College, United Kingdom

Universita' di Torino and Collegio Carlo Alberto, Italy

Contact information: Institute for New Economic Thinking, Eagle House, Walton Well Road, Oxford OX2 6ED, United Kingdom. matteo.richiardi@maths.ox.ac.uk

Abstract. In this paper, I elaborate on the role of agent-based (AB) modelling for macroeconomic research. My main tenet is that the full potential of the AB approach has not been realized yet. This potential lies in the modular nature of the models, which is bought by abandoning the straitjacket of rational expectations and embracing an evolutionary perspective. I envisage the foundation of a Modular Macroeconomic Science, where new models with heterogeneous interacting agents, endowed with partial information and limited computational ability, can be created by recombining and extending existing models in a unified computational framework.

INTRODUCTION. Though there are early antecedents, it's now at least two decades that agent-based (AB) models have been introduced to economics.[1] It is therefore time to ask whether they have marked an impact on the way economics, and especially macroeconomics, is done, and what their future prospects look like. This is all the more relevant given the debate on macroeconomic modelling which was prompted by the Great Recession, and Ricardo Caballero's suggestion that macroeconomics should be in "broad exploration" mode [Caballero 2010].

Forecasting the future is notably a difficult exercise, but a common perception is that AB models have "a bright future past them": they have been charged with high expectations, especially from outside the mainstream literature, which they somewhat failed to live up to.

With respect to mainstream economic models, AB models trade off individual sophistication (in expectation formation and decision making) with complexity in the interaction structure and richness in the institutional details. The rational expectations (RE) hypothesis at the core of mainstream economic models, with its strong consistency requirements –all actions and beliefs must be mutually consistent at all times– requires simplistic models. This is replaced in AB models by the assumption of partial information and limited computational ability, which brings in milder evolutionary requirements: corrections must take place, through learning, selection or reactions in the environment. The resulting increased flexibility in model specification, however, has downsides: (i) assumptions are sometimes deemed arbitrary and disconnected from the literature, suggesting a return to the sort of anarchy that was lamented before the RE revolution [Wickens 2014], (ii) models often exhibit too many degree of freedoms, and are therefore non-falsifiable, (iii) models often lack a sound empirical grounding; when present, this is often limited to some ad hoc calibration [Grazzini and Richiardi 2015]. Other common critiques point to the fact that (iv) models are oftentimes poorly documented and hardly replicable [Leombruni et al. 2006], and (v) writing an AB model requires quite a lot of programming skills; code is often not re-usable and projects are not incremental [Leombruni and Richiardi 2005].

A search on EconLit returns 5,705 articles for Dynamic Stochastic General Equilibrium (DSGE), Real Business Cycle (RBC) and New Keynesian (NK) models since 1980, and only 1,062 hits for AB models (figure 1).[2]
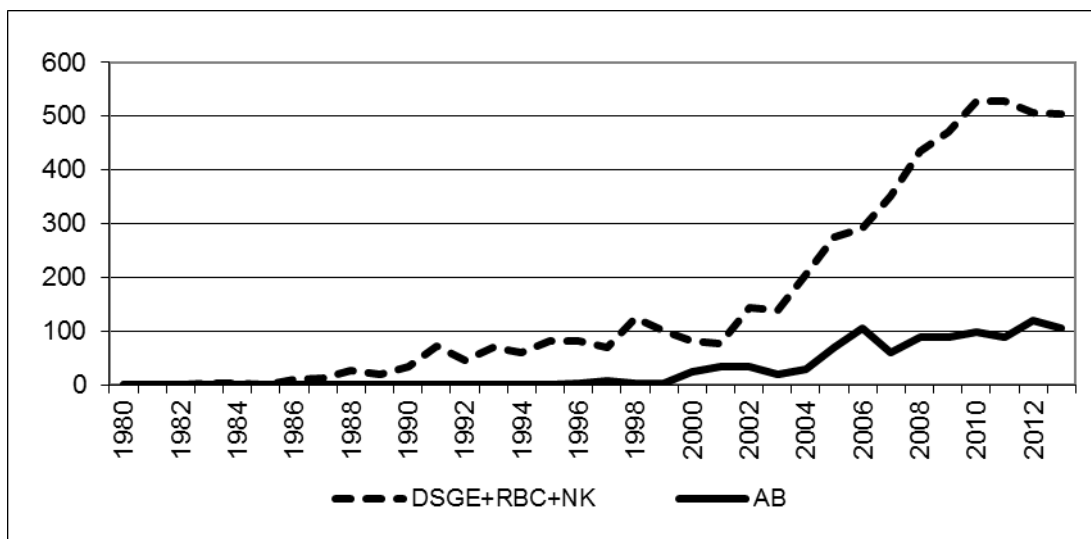
Figure 1: Number of articles retrieved in the EconLit database (all types of publications) searching for "Dynamic Stochastic General Equilibrium", "Real Business Cycle", "New Keynesian", "Agent-based", variants and acronyms, 19/1/ 2015.

Does this graph suggest that AB modelling is ready for take off, as the new consensus in macroeconomics in the early 2000s, or that it has already levelled off? In this paper I argue that in order to reap all the benefits of the AB approach, a change in strategy is required, and this allows to address all the critiques listed above.

Abandoning the consistency requirements of RE equilibrium removes the need for solving the models as one single block: in principle AB models can be fully modularized, with endless possibilities for recombination and extension. However, this feature has not been exploited so far. Existing AB models are mostly one-off exercises which do not travel across research groups and whose "working life" does not usually extend beyond the grant that originated them. Code is not re-used, except possibly by the authors themselves; alternative assumptions are not tested, results are not generalized. In short, knowledge accumulates at a slow pace. The reason for this "modelling individualism", as we might call it, is to be searched in the struggle to win the "modelling race", in line with existing incentives in terms of (short-term) publications and funding opportunities.

Conversely, by fully exploiting the modular nature of AB models, a common computational environment can be built, where different behavioral assumption, markets and institutional frameworks can be implemented as separate simulation modules. Classification of and experimentation with different behavioral assumption, from choice heuristics to learning and expectation formation, will then provide a firm foundation for the modelling assumptions made (critique i above); comparison of different specifications will allow to reduce the number of

parameters to a minimum and focus on lean structures (point ii); adoption of common computational interfaces will allow the development of standard simulation-based estimation procedures (point iii) and will directly address points (iv) and (v).

The remaining of the paper is structured as follows. Section 2 suggests that mainstream macroeconomic models are not "systemic" models, and are thus not suited for analyzing endogenous changes in the economic structure, like those we have experienced during the Great Recession. Section 3 discusses why these models cannot be further extended in order to provide such a vital development. Section 4 explains why AB models, by departing from the strict requirements of RE equilibrium analysis in favor of an evolutionary approach, can provide a valid alternative. Section 5 elaborates on the "double dividend" that can be obtained from moving away from RE, as models can in principle be fully modularized. Section 6 discusses the main challenges that a new Modular Macroeconomic Science faces, as agents have to share common ontologies (an ontology is a part of the agent's knowledge base that describes what kind of things an agent can deal with and how they are related to each other). Section 7 suggests a practical roadmap for testing the new modular approach and build a Modular Macroeconomic Simulator, developed using a readily accessible simulation platform. Finally, section 8 discusses why the modular approach has not yet materialized in AB modelling, and why the AB research community should actively engage in a collaborative effort to help the prediction of a future Modular Macroeconomic Science become true.

THE GREAT RECESSION POSES A SERIOUS CHALLENGE TO MAINSTREAM MACRO. Dynamic stochastic general equilibrium (DSGE) models, the standard tool for macroeconomic analysis, have been deeply challenged by the Global Financial Crisis of 2008 and the Great Recession that followed. As Jean-Claude Trichet [2010], at the time President of the European Central Bank, stated:

> *"When the crisis came, the serious limitations of existing economic and financial models immediately became apparent. [...] Macro models failed to predict the crisis and seemed incapable of explaining what was happening to the economy in a convincing manner. As a policy-maker during the crisis, I found the available models of limited help. In fact, I would go further: in the face of the crisis, we felt abandoned by conventional tools".*

Brad Delong [2011] reports that when the former U.S. secretary Lawrence Summers was asked to "*name where to turn to understand what was going on in 2008,* [*he*] *cited three dead men, a book written 33 years ago, and another written the century before last*". Summers was pointing to the work of Bagehot [1873], Kindleberger [1978] and Minsky [1982; 1986], well before the New Neoclassical Synthesis, the "gold standard" that emerged in the 1990s from the clash

between two competing business cycle theories, the Real Business Cycle (RBC) and the New Keynesian (NK) paradigm. Somewhat ironically, Olivier Blanchard [2009] declared that "*the state of Macro is good*", just before the Great Moderation almost anagrammatized into an Armageddon for Macroeconomics. The limits of the DSGE approach at the core of the New Neoclassical Synthesis have been analyzed and exposed by a number of prominent voices [e.g. Kirman 2010; Krugman 2011; Stiglitz 2011; for a less technical note, Smith 2014], and not only in the aftermath of the crisis [Colander 2006; Colander et al. 2008]. As Ricardo Caballero [2010] has vividly expressed, the DSGE approach "*has become so mesmerized with its own internal logic that it has begun to confuse the precision it has achieved about its own world with the precision that it has about the real one*". This logic is built upon the implicit view that markets and economies are inherently stable and that they only temporarily get off track, when they are perturbed by external shocks, and fails to recognize that the interaction between heterogeneous agents and the institutional environment can produce endogenous business cycles, possibly leading to systemic crises. In DSGE models the existence of occasionally large booms and busts is explained by the occurrence of large exogenous shocks, while nothing in the model leads to non-normality [De Grauwe 2012]. Indeed, DSGE models are unable to reproduce the observed non-normal distributions of many macro variables, even when fed with shocks drawn from fat tail distributions [Ascari et al. 2013].

Inclusion of financial frictions in a new vintage of DSGE models has improved their ability to reproduce observed features of the big post-crisis recession [Brzoza-Brzezina et al. 2013], like the little sensitiveness of inflation to the output gap [e.g. Christiano et al. 2015; Del Negro et al. 2015]. However, the emergence of the large financial shock that originated the Great Recession in the first place is left unexplained. Similarly, the transition from the Great Inflation to the Great Moderation is mainly explained with "luck" - shocks becoming more benign in the latter period [Fernández-Villaverde et al. 2010]. Hence, the most important macroeconomic events of the past 50 years are "explained" with exogenously given random disturbances. To the extent that "*shocks are a measure of our ignorance*" [Abramovitz 1956], the theory seems to be missing something quite fundamental. In DSGE models there is no causal link between a boom and a bust by construction, except for small self-correcting deviations from the deterministic steady state, so the stable period can be understood as separate from the unstable period. Yet if the boom and the bust are caused by the same process, then understanding one entails understanding the other. DSGE models with financial frictions have indeed yielded predictions that are in line with some of Minsky's insights (e.g. the financial accelerator leading to a "Minsky moment", a sudden major collapse of asset values which shapes the credit and business cycle). However, they fail to incorporate the major intuition of Minsky's Financial Instability Hypothesis: *stability can be destabilizing*, sowing the seeds of its own demise. In tranquil times,

banks become less cautious in extending credit and firms less cautious in borrowing; this results in endogenous financial fragility.

Also, the vast literature on individual heterogeneity within the DSGE framework [Heatcote et al. 2009] is largely confined to investigating the role of idiosyncratic income shocks in affecting labor supply, hence output. A crucial stylized fact that DSGE models fail to reproduce is inequality in the distribution of work – the fact that some individuals are fully employed while others remain unemployed. In downturns and recessions, aggregate demand is what matters, and it can be affected in fundamental ways by the distribution of income: in DSGE models *"there is no discussion of differing marginal propensities to consume, which can serve as the basis of stimulative redistributive policies"* [Stiglitz 2011]. Moreover, there is no consideration of other asymmetries able to produce network and cascade effects, like commercial or credit links between firms.

DSGE MODELS CANNOT BECOME "SYSTEMIC" MODELS. Taking a step back, New Keynesian DSGE models did a reasonable job in describing advanced economies during normal times, though admittedly with a number of ad-hoc (and non-Lucas compliant) patches –with a metaphor, it is as DSGE models were doctors specialized in healthy patients. However, (i) assistance from models –as from doctors– is mostly needed when the conditions are bad and unstable, (ii) even in tranquil times, it is unclear whether DSGE models are of much *practical* relevance: as an example, the "bad luck vs. bad policies" literature on the causes of the Great Inflation showed that monetary authorities were following the Taylor principle even before John Taylor "invented" it, and long before it was rationalized in NK models [Fernández-Villaverde et al. 2010]. Can DSGE be further stretched to analyze economic *systems*, and not only specific economic periods? If capital is added to the most basic models, or the (log)linear approximation –equivalent to the assumption that the economy undergoes only small and gradual changes– is dropped, analytical solutions become out of reach. Introducing more heterogeneity, behavioral articulation, or institutional details, quickly results in intractable optimization, because agents have to solve too complicated forward-looking problems.

This stringent trade-off originates from the RE hypothesis: the assumption that people have probability beliefs that coincide with the probabilities predicted by the model, essentially meaning that agents perfectly understand their environment (the model), and that they are able to compute fixed points in the strategy space. Beyond analytical tractability, this assumption –at the cornerstone of modern macro– allows expectations and hence behavior to react instantaneously to announcements of future policy changes, while in pre-Lucasian models (naïve) adaptive expectations would have been revised only after the policy were actually

implemented, and only gradually [Wickens 2014]. Attempts to devise mechanisms of expectation formation which depart from RE and are sensitive to information [Woodford 2013] are mainly based on the statistical learning approach [Evans and Honkapohja 2009; 2013], where agents are modelled as econometricians in making forecasts. This literature has investigated under what conditions agents are able to learn the true DGP of the process, and consequently converge to the RE equilibria. However, it is still based on a sophisticated view of the agents and it does not take into consideration the broader literature on bounded rationality and the role of heuristics [Gigerenzer and Todd 1999; Gigerenzer et al. 2011], adaptive learning in evolutionary game theory [Durieu and Solal 2012], social learning in networks [Mobius and Rosenblat 2014] and other findings coming from psychology [Della Vigna 2009], which can result in more complex aggregate dynamics.

MACROECONOMIC RESEARCH IN 'BROAD-EXPLORATION' MODE. Overall, this discussion brings us back to Caballero [2010], who warns that "*[o]n the methodology front, macroeconomic research has been in 'fine-tuning' mode within the local-maximum of the dynamic stochastic general equilibrium world, when we should be in 'broad-exploration' mode*." AB models are one promising direction for exploration [Freeman 1998; Colander 2005; LeBaron and Tesfatsion 2008; Farmer and Foley 2009; Stiglitz and Gallegati 2011]. From the policy side, Jean-Claude Trichet himself supported this view: "*We need to deal better with heterogeneity across agents and the interaction among those heterogeneous agents. We need to entertain alternative motivations for economic choices. [...] Agent-based modelling dispenses with the optimisation assumption and allows for more complex interactions between agents. Such approaches are worthy of our attention*" [Trichet 2010].

AB models are structural dynamical models characterized by three features [Richiardi 2012]: (i) there are a multitude of objects that interact with each other and with the environment, (ii) these objects are autonomous, that is there are no central, or "top-down" coordination devices (e.g. the Walrasian auctioneer), and (iii) aggregation is performed numerically. Note that neither the presence of micro-foundations nor the computational nature defines the methodology, as this is shared by most state-of-the-art macro models. A crucial role in AB models is played by heterogeneity: this includes partial knowledge of the environment, and limited and differentiated computational ability. Rather than following axioms of logical consistency and computing fixed points, they follow simple heuristics based on psychological plausibility and ecological effectiveness: *ecological rationality* appears when the structure of the boundedly rational decision mechanisms [Altman 2006] matches the structure of information in the environment [Todd and Gigerenzer 2012]. These heuristics are evolved through learning and selection [Hommes 2013], with a crucial distinction being whether learning takes place at an individual or

a social level [Vriend 2000].

In AB models learning need not be only backward-looking [Kirman 2011]. For instance, in an "agenized" version of the Smets and Wouters [2003] model, Sinitskaya and Tesfatsion [2015] compare (i) reinforcement learning, (ii) Q-learning, (iii) a forward-looking rolling-horizon method, and (iv) an adaptive dynamic programming method based on value-function approximation. They find that simpler decision rules can outperform more sophisticated ones – but only if they entail some sort of forward-looking behavior coupled with a relatively long memory. De Grauwe and Markiewicz [2013] compare statistical learning and fitness learning and find that the first cannot replicate the observed disconnection of the market rates from the underlying fundamentals (the latter is not able to replicate volatility clustering). Also, agents can endogenously switch between different learning rules. For instance Anufriev and Hommes [2013], based on the heuristic switching model of Brock and Hommes [1997], distinguish between negative and positive feedback systems and show that in the first adaptive heuristics end up dominating, whereas in the latter trend-following heuristics prevail.

The concept of equilibrium also changes. In rational expectation models, equilibrium is defined as a consistency condition in the behavioral equations: agents must act consistently with their expectations, and the actions of all the agents must be mutually consistent. The system is therefore always in equilibrium, even during a phase of adjustment to a shock. By converse, equilibria in AB models are defined only at the aggregate level and only in statistical terms, when macro-outcomes become stationary [Grazzini and Richiardi 2015].

This methodological stance removes a lot of technical constraints in model building, allowing for more flexibility in model specification. In particular, the use of simple learning mechanisms coupled with evolutionary mechanisms that allow "learning about learning" drastically reduces the difficulty of the choice problem of the agents, and consequently the computational complexity of the model. This permits to introduce more institutional details, as in the modelling of specific markets [Tesfatsion 2011; Li and Tesfatsion 2012] or market clearing mechanisms (e.g. an order book, see Paddrik et al. [2014]), and more complex interactions including, for instance, production or credit networks [Weisbuch and Battiston 2007; Battiston et al. 2007; 2012; Delli Gatti et al. 2010]. AB macro models typically target a higher number of stylized macro facts than DSGE models [Dawid and Neugart 2015; Neugart and Richiardi 2015]. They have arguably been successful in analyzing systemic risk [Geanakoplos et al. 2012] and macroprudential regulations [Teglio et al. 2012; Poledna et al. 2014], the role of innovation policies [Ballot and Taymaz 2001; Dawid et al. 2014; Hommes and Zeppini 2014], endogenous business cycles and stabilization policies in economies with imperfect information and incomplete credit markets [Delli Gatti et al. 2005 and subsequent papers; Dosi et al. 2010, 2013], the stability of general equilibrium [Gintis 2007]. An active role in the development of

AB macro models has been played by the European Commission with the funding of three large-scale projects: EURACE[3], aimed at developing an AB software platform for European economic policy design, POLHIA[4], aimed and analyzing monetary, fiscal and structural policies, and CRISIS[5], aimed at understanding systemic instabilities. However, but for a few exceptions - from Schelling [1969] to Gintis [2007] or Geneakoplos et al. [2010] - research in AB modelling has mainly been confined to the *Journal of Economic Behavior & Organization* and the *Journal of Economic Dynamics & Control*, as if there was a glass ceiling impeding to reach a wider audience.

THE VISION FOR A NEW MODULAR MACROECONOMIC SCIENCE. To rephrase the basic argument, the defining peculiarity of AB models, vis-à-vis mainstream economic models, is a departure from RE equilibrium in favor of an evolutionary approach, where evolution basically means learning (at an individual or social level). Avoiding the need to solve for RE equilibria saves a lot of computing time: this can be spent to complicate the models, introducing more realistic assumptions. Mainstream economics does not like this trade-off, as it reminds of the sort of pre-Lucasian anarchy. However, there is a "double dividend" of moving away from RE, which has not been reaped so far and that may change the nature of the (modelling) game. Because equilibria (defined as stationary states) are not analytically derived or imposed but explored by simulations - which prompts the definition of AB modelling as *generative social science*[6] - parts of the computer program can be easily replaced without compromising the ability to understand the model behavior (that is, without necessarily compromising execution speed). This is, in essence, modularity. Modularity allows the division of labor which, as every economist after Adam Smith knows, is a driving force in boosting productivity. Exploiting modularity means that individual researchers and research groups can develop and refine modules that fit into larger models. Different modules (e.g. labor demand and supply, credit demand and supply, production, consumption, household formation, retirement, etc.) can then be combined within the same modelling environment into new models, as in a Lego™ construction game. Institutional arrangements (e.g. market clearing mechanisms) and behavioral characteristics (e.g. decision rules, expectation formation rules, learning rules) can also be typified and implemented as separate blocks. This makes it possible to combine the labor supply module of author X with labor demand and production modules of author Y, the housing market module of author Z, etc., specifying for each module how agents form their expectation, interact with other agents, and take their decisions. In turns, this allows the systematic comparison and assessment of assumptions: convergence on a common set of modelling choices, possibly as an alternative to those employed in DSGE modelling, is then facilitated because their properties across a wide range of models become known, and consensus

"workhorse" models emerge out of the collaborative effort of a community of researchers.[7]

In other words, the vision for a new Modular Macroeconomic Science entails a meta-modelling strategy aimed at the systematic investigation of the effects of different modelling choices and their systematic empirical validation, possibly drawing from developments across disciplines, including computer science, management science, cognitive science and psychology. Inspiration comes from atmospheric modelling, where additional components (like carbon models, sea-ice and glacial-ice models, atmospheric chemistry models, land-surface models, etc.) are developed and refined, and then added to a core global circulation model.

CHALLENGES. A necessary requirement for AB modelling to evolve into a new Modular Macroeconomic Science is the development of appropriate application programming interfaces (APIs) to allow communication between different modules of the same agent, or between different agents. These are a set of routines, protocols, and tools which define functionalities internally used by the simulated agents (e.g. learning algorithms) or used by the agents to interact with other agents (exchange of information, goods and services) that are independent of their respective implementations. Note that I am not envisaging here the adoption of one common simulation platform, which besides being unattainable, is also undesirable. Competition between different simulation platforms, from general purposes programming languages (e.g. C++, Java, Python, Scala), possibly integrated with simulation specific libraries (e.g. Mason[8], RePast[9], JAS-mine[10]), to general-purpose mathematical software (e.g. Mathematica, Mathlab or Matcad, or agent-based specific languages (Netlogo[11]), is a good thing because it foster improvements and provides cross-validation of the different implementations. What is needed is rather an abstract protocol to which the different implementations should adhere. This protocol should also be neutral about the hardware characteristics of the implementation (synchronous vs. asynchronous). Here I just provide some preliminary thoughts about how such a protocol might look like.

As a preliminary remark, note that a protocol for modelling economic agents is different from a standard for agents' communication in multi-agent systems. For instance, the Foundation for Intelligent Physical Agents (FIPA), an IEEE Computer Society standards organization, provides a collection of standards for the interoperability of heterogeneous agents and the services they represent. The FIPA standards defines how messages should look like, but fall short of providing requirements about their content. Available standards deal with agents' interaction, but they have nothing to say about the structure and behavior of *economic agents*. Modularization requires, possibly on top of existing communication standards, the definition of appropriate APIs for such economic agents (*economic APIs*).

Also, the economic APIs are different from protocols related to the internal implementation of the agents. These protocols solve important problems and should be adopted to ensure an efficient and robust software architecture, but they do not allow by themselves the modularization of the models.[12]

From a general perspective, economic agents are characterized by *states*, *parameters* and *actions*. More specifically, the states include the following: (i) *individual characteristics* –in the case of a person: age, sex, location, abilities, family status, work status, health status, etc.; in the case of a firm: size, location, industry, labor force composition, etc. – possibly including the networks agents are linked to; (ii) *endowments* (money, time, assets); (iii) *expectations* about relevant variables; (iv) some measure of *fitness* (or *utility*).

Actions can modify an agent's internal state, and possibly other agents' states as well. States can also change due to factors which are independent of an agent's actions (think of ageing, or depreciation of human and physical capital). Irrespective to whether they change according to own actions, actions of others, or independently of any actions, I refer to states as governed by state-specific *law of motions*, on the basis of some individual-specific *parameter*s.[13] Among the actions an agent can take, of particular importance for an economic agent is the ability to make *transactions*, that is, engaging in market exchanges. Transactions involve deliberately swapping, or commit to swap at future dates –possibly conditional on specific conditions being met– some endowments for others, and can also be thought of as *contracts[14]*. The main characteristics that makes endowments the object of transactions is that they can be cumulated (though they might depreciate). Nothing precludes an agent from having negative endowments (that is, borrowing). Stocks and flows of endowments are registered in a *balance sheet.*

Finally, choices about actions (including transactions) are determined by *decision rules*, individuals-specific functions that take as an input the state of an agent, possibly the states of other agents, and the parameters.

Therefore, what is needed to model an economic agent is a set of parameters; a list of individual characteristics; a list of endowments; for each of the endowments, a list of (conditional) commitments to future transactions[15]; a list of actions; a set of decision rules; for each forward-looking variable entering a decision rule, a vector of expectations about its future values; for each state, a law of motion; a list of fitness measures.[16]
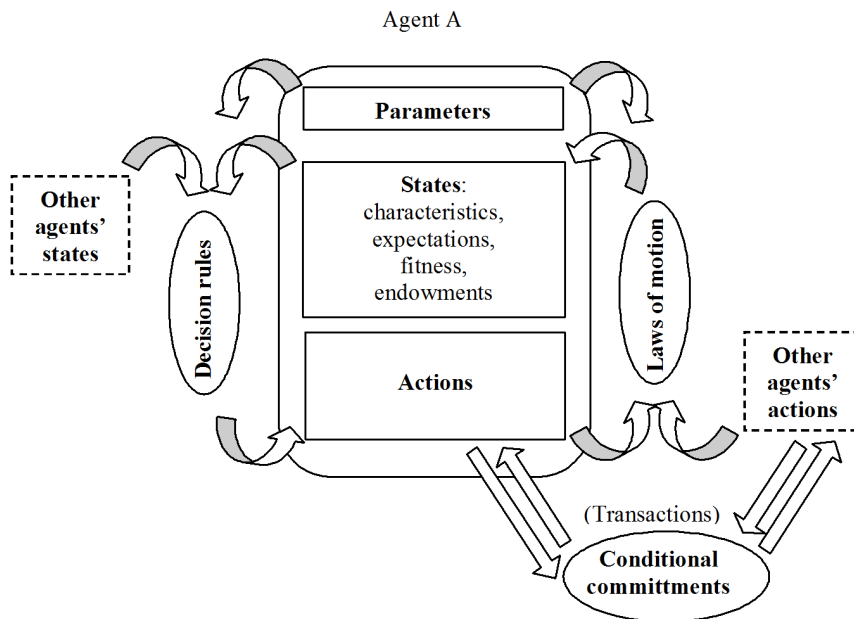
Figure 2 provides a schematic representation.

Figure 2: Scheme of an economic agent.

Modularizations, in this context, means :

i)   as many components of an agent as possible are classified and implemented as separate objects, so that an agent can be equipped with alternative, ready to use, specifications; when none is deemed appropriate, modifications can easily be introduced, and saved as new modules;

ii)   interaction between different agents, possibly playing different roles, follows simple protocols that allow to change the internal functioning of an agent without having to modify those of other agents.

Point (i) essentially involves the typification of laws of motion and decision rules. Among the first, a prominent role is given to expectation formation, where different mechanisms, from rational to adaptive or naïve expectation, can be implemented. Decision rules on the other hand relate to different theories of agents' behavior, from zero-intelligence [Gode and Sunders 1993] to the Belief-Desire-Intention model often employed in the literature on Multi-Agent Systems [Woolridge 2009], to sophisticated forward-looking behavioral engines, possibly including RE optimization. The definition of a common API for these functions is relatively straightforward: the agent passes the function a set of inputs (states and parameters), and receives back an output (a new state, or an action). As long as different functions accept the same types of inputs and produce the same type of output, they can share a common interface, though in practice there

will be a different interface for each (type) of law of motion or decision rule considered (e.g. whether to look for a job as a dependent employee or as a self-employed, whether to create a new business, whether to stay or not in a market, what to consume, whether to buy or sell stocks, etc.).

Point (ii) –the definition of abstract models of interactions– requires modelling *markets*. Here, I restrict my attention to search models of market interaction, where a match between supply and demand has to be found. Although markets are characterized by different institutional arrangements (auctions, limit order markets, bargaining, fixed price posting, etc.), one can think of defining just two interfaces depending on whether search is unilateral or bilateral.

The common starting point is that one side of the market posts an offer, in terms of a *bundle of characteristics*. These might involve the quality, quantity and price of a consumption good; the wage (or a maximum wage), work hours, duration of the contract, and other job characteristics (e.g. location, skills required, agreeableness) for a vacancy; the ask price plus the quantity and maturity structure / expiration date of a financial instrument. I then distinguish between two types of offer processing.

The first type (unilateral search) is first-come, first-served, and defines a one-to-one (1-to-1) relationship between the two sides of the market: whoever accepts the offer first gets the deal, and the offer is immediately withdrawn from the market. This is typically the case of goods and services (when there are more than one item on sale, the offer is repeated subject to availability). The *API for 1-to-1 market transactions (unilateral search)* only requires one side of the market to post offers, and the other side of the market to select. How offers are advertised (e.g. by shelf display, through intermediaries, by word-of-mouth, etc.), how agents come across offers (e.g. by random walking in a supermarket, by systematic search, etc.) and how they select one specific offer among the possibly many they have known about is related to the internal functioning of the agent and does not require to be standardized though an interface.[17]

The second type of offer processing involves bilateral search: collecting expressions of interests, and then selecting a transaction partner among those who applied. It defines a many-to-many (m-to-m) relationship between the two sides of the market, which collapses to a one-to-many relationship when the interested applicants are allowed to make only one application each. Typical examples are job vacancies, financial instruments, and the marriage market. In this case, those who are looking for the good, service, security, job or mate need also specifying a bundle of characteristics (including a reservation price), which determine what offers are considered, and a maximum number of applications they send. Then, markets are distinguished in centralized or decentralized. In *centralized markets* (e.g. clearing houses, auctions) the market is a specific agent that matches demand and supply. No further action is required on the part of the

agents. In *decentralized markets*, agents on one side of the market must select a trading partner themselves (generally in competition with other agents), in an asynchronous manner. Again, the specific way selection is done is a problem of the agents and does not affect the interface. Hence, the *API for m-to-m market transactions (bilateral search)* requires both sides of the market to specify bundles of characteristics. Matching is then performed by a third agent (the market maker, in centralized markets) or by one side of the market (in decentralized markets).

It is important to stress that the APIs sketched above are limited to search-and-match markets, which are a fairly general class of markets but by no means exhaust all the possible market structures. For instance, the Walrasian auctioneer, where the two sides of the market respectively specify a supply and demand schedule, remains outside their range of application.[18] One could argue that a Walrasian auctioneer is not compliant with the AB methodology (see above), and stop worrying; alternatively, one could write a specific API for Walrasian markets, or "tweak" the APIs in order to include the possibility of specifying demand and supply functions. The latter is actually quite simple to implement, as it is sufficient that both sides of the markets specify sets of bundles, rather than individual bundles –combinations of price and quantities that define a demand or a supply schedule. As it is often the case, there is a trade-off between generality and simplicity.

PRACTICAL STEPS. Developing the economic APIs cannot be done only at a theoretical level, and requires testing on a large enough number of (simple) models. This in turns implies the development of a Modular Macroeconomic Simulator, a computational environment where the test models are completely modularized, in terms of markets considered and individual behavior (e.g. expectation formation and learning abilities). The Modular Macroeconomic Simulator has to be platform-specific, though it would be good to have different implementations using different simulation platforms. I will now discuss the potential of *JAS-mine*, a simulation platform that I have contributed to develop, to evolve into a Modular Macroeconomic Simulator.

JAS-mine ("Java Agent-based Simulation library. Modelling in a Networked environment"[19]) is a toolkit for AB and dynamic microsimulation modelling, in line with a trend of convergence of the two methodologies [Richiardi 2013]. The platform is *generic* and allows for all types of discrete-event simulations, not necessarily macroeconomic models. As such, it is a *simulator*, but not yet fully *modular* nor specifically *macroeconomic*. The platform provides standard tools which are frequently used both in AB modelling and dynamic microsimulations; these include scheduling of events, design of experiments, run-time monitoring and visualization, I/O communication (it features an embedded relational database, with flexible methods for

persisting the output of the simulations, in addition to automated CSV table creation), statistical analysis, GUI with plots and graphs etc. The main value added of the platform, however, is to give the researcher a guide / template on how to structure a generic simulation model, and in this respect makes it a good candidate for evolving into a Modular Macroecomic Simulator. In particular, JAS-mine favors the separation of data representation and management, which is automatically taken care of by the simulation engine, from the implementation of processes and behavioral algorithms, which should be the primary concern of the modeler. This results in quicker, more robust and more transparent model building. To be more precise, JAS-mine extends the Model-Observer paradigm introduced by the Swarm experience [Askenazi et al. 1996] and introduces a new layer in simulation modelling, the Collector:

- The *Model* deals mainly with specification issues, creating objects, relations between objects, and defining the order of events that take place in the simulation.

- The *Collector* collects the data and compute the statistics both for use by the simulation objects and for post-mortem analysis of the model outcome, after the simulation has completed.

- The *Observer* allows the user to inspect the simulation in real time and monitor some pre-defined outcome variables as the simulation unfolds.

This three-layer methodological protocol allows for extensive re-use of code and facilitates model building, debugging and communication. Flexibility in model design and implementation, scalability of the code (that is, the code complexity must increase approximately linearly with the complexity of the underlying model, and must remain highly readable for debugging, cooperative development, and documentation), efficiency in data exchange and input-output communication, and compatibility with external software solutions and tools were the main specifications for the project.[20] To meet those targets, JAS-mine is based on the following architectural choices: (a) object-oriented programming language (OOP), which provides a natural and intuitive way of modelling populations of agents; (b) strict adherence to the open-source paradigm; (c) use, whenever possible, of standard, open-source tools already available in the software development community (rather than development of an ad-hoc grammar and syntax); (d) close model-data integration (introduction of an ORM layer to connect the object-oriented simulation with an underlying SQL relational database management system).

In particular, JAS-mine already includes alternative implementations of standard functionalities used in AB and microsimulation modelling, which can be used by the agents in the simulation following a modular approach. These include a library implementing a number of different matching methods, to match different lists of agents; a library implementing a number of

different alignment methods, to force the microsimulation outcomes meeting some exogenous aggregate targets; a library implementing a number of common econometric models, from continuous response linear regression models to binomial and multinomial logit and probit models, with tools to bootstrap the estimated coefficients to allow for uncertainty analysis;

Turning JAS-mine into a Modular Macroeconomic Simulator is however a long-term project, which has just been spelled out.[21] Again, it should be stressed that any platform / language could work as a testbed for developing the economic APIs, the only requirement possibly being, in order to foster collaborations, an open-source nature and a small upfront learning cost.[22]

CONCLUSIONS. It can be argued that the popularity of DSGE macro-models is due only in part to their theoretical and empirical appeal. Other important determinants are:

1. DSGE models are quite homogenous: they all follow the same methodological assumptions and modelling logic.

2. Despite their being complicated models, writing a DSGE model is relatively easy, thanks to the availability of ready-to-use software tools (e.g. DYNARE, a Matlab plugin. An archive with more than 60 macroeconomic models implemented in DYNARE, with common diagnostic for systematic model comparison, is maintained at www.macromodelbase.com).

3. Because of 1) and 2), and because it is relatively easy to publish DSGE-based papers, many Ph.D. students embark on DSGE modelling. This further reinforces their general acceptance.

In contrast, AB models are heterogeneous and often criticized as ad-hoc. There is little awareness of the different modelling choices. Writing an AB model requires quite a lot of programming skills; code is often not re-usable and projects are not incremental.

To improve on the current state of affairs and establish AB modelling as a standard practice in macroeconomics, two strategies seem possible. The first one can be described as a search for a *silver-bullet*: writing an AB model which is so good that everybody has to pay attention. This appears to be the strategy followed by the EURACE, POLHIA and CRISIS project. The problem with this strategy is that it easily leads to gigantic "multi-purpose" models that try to hit too many different targets (research questions) at the same time, with the consequence that it is very difficult to understand what matters, in explaining the model results.

The second strategy refer to a *meta-modelling* approach: comparing many (relatively simple) AB models, which depart from each other –and possibly from the mainstream literature– only in

limited respects, so as to understand the effects of each deviation from the assumptions usually made in the literature.

The first strategy tends to bypass the mainstream, while the latter makes continuous reference to the mainstream and often builds upon it. In particular, the meta-modelling strategy allows to understand systematically the impact of deviations from the standard assumptions of macroeconomic models: heterogeneity, imperfect information about the economic environment, replacement of RE with specific learning routines. Whilst the appeal of the silver-bullet strategy is not to be understated, I believe that the future of AB (macro-) modelling rests with the meta-modelling strategy: creating a common computational environment that facilitates model comparison and model development, by adopting a fully modular approach. This requires that models follow a Lego™ approach, with new models created by combining modules and features of existing models, effectively creating a new Modular Macroeconomic Science.

This vision, if successful, has the potential for radically transforming how research in AB macroeconomic modelling is done. I will therefore conclude by discussing why it has not been advanced before, and why it would be an extraordinary breakthrough for advancing the entire macroeconomic field.

To start with, *modularity is precluded in a DSGE setting* because of the assumptions that agents (i) know a lot about the macroeconomic environment in which they are embedded, and (ii) are able to solve complicated optimization problems where the direct and indirect effects of individual choices have to be taken into account before actually implementing them, implying that all model components are fully connected and the model has to be solved as one single block.

Second, *modularity has not been exploited so far in AB modelling* because of the high fixed cost involved in developing a flexible simulator, where modules can be easily combined, replaced or extended. The relative novelty of the methodology, and the incentives for immediate returns in terms of publications and funding, have resulted in highly specific computational architecture and "disposable" models. Also, modularization involves the definition of appropriate economic APIs, to allow the interoperability of different modules, and the implementation of many small-scale models, possibly taken from the literature, in order to test the APIs and produce an initial library of modules. Both tasks have a public good nature, and require a lot of infrastructure building.

As for what concerns the possible *benefits of a modular approach* for AB modelling, these include (i) more efficient model building, stemming from the possibility to integrate components developed by previous researchers –e.g. the housing market of Geanakoplos et al. [2012] with the financial market of Poledna et al. [2014] and the labor market of Richiardi

[2006], (ii) more general results, stemming from systematically testing and comparing alternative specifications –e.g. the financial market of Poledna et al. [2014] with that of Teglio et al. [2012], or different learning algorithms as in Sinitskaya and Tesfatsion [2015], (iii) better documentation of the model structure, (iv) increased cooperation of different research groups.

Exploiting AB modelling to its full potential will clarify to what extent the methodology is an alternative or a complement to mainstream modelling approaches, beyond methodological disputes.

REFERENCES

Abramovitz, Moses. 1956. Resource and Output Trends in the U.S. since 1870. *American Economic Review* 46(2): 5-23.

Altman, Morris. 2006. *Handbook of Contemporary Behavior Economics: Foundations and Developments*. M E Sharpe, Armonk, NY

Anufriev, Mikhail, and Cars H. Hommes. 2013. Evolutionary selection of expectations in positive and negative feedback markets. *Journal of Evolutionary Economics* 23(3): 663-688.

Ascari, Guido, Giorgio Fagiolo, and Andrea Roventini. 2013. Fat-Tail Distributions and Business-Cycle Models. *Macroeconomic Dynamics*. Available on CJO2013. doi:10.1017/S1365100513000473.

Askenazi, Manor, Roger Burkhart, Chris Langton, and Nelson Minar. 1996. The swarm simulation system: A toolkit for building multi-agent simulations. Santa Fe Institute Working Paper no. 96-06-042.

Bagehot, Walter. 1873. *Lombard Street: A Description of the Money Market*. Henry S. King and Co., London.

Ballot, Gérard, and Erol Taymaz. 2001. Training policies and economic growth in an evolutionary world. *Structural Change and Economic Dynamics* 12(3): 311-329.

Battiston, Stefano, Domenico Delli Gatti, Mauro Gallegati, Bruce Greenwald, Joseph E. Stiglitz. 2007. Credit chains and bankruptcy propagation in production networks. *Journal of Economic Dynamics and Control* 31(6): 2061-2084.

------------. 2012. Default cascades: When does risk diversification increase stability? *Journal of Financial Stability* 8(3): 138-149.

Blanchard, Olivier, 2009. The State of Macro. *Annual Review of Economics* 1(1): 209-228.

Brock, William A, and Cars H. Hommes. 1997. A Rational Route to Randomness. *Econometrica* 65(5): 1059-1096.

Brzoza-Brzezina, Michal, Marcin Kolasa, and Krzysztof Makarski. 2013. The anatomy of standard DSGE models with financial frictions. *Journal of Economic Dynamics and Control* 37 (1): 32-51.

Caballero, Ricardo J. 2010. Macroeconomics after the Crisis: Time to Deal with the Pretense of Knowledge Syndrome. *Journal of Economic Perspectives* 24(4): 85-102.

Christiano, Lawrence J, Martin S. Eichenbaum, and Mathias Trabandt. 2015. Understanding the Great Recession. *American Economic Journal: Macroeconomics* 7(1): 110-167

Colander, David. 2005. The Future of Economics: the Appropriately Educated in Pursuit of the Knowable. *Cambridge Journal of Economics* 29: 927-941.

Colander, David, ed. 2006. *Post Walrasian Macroeconomics. Beyond the Dynamic Stochastic General Equilibrium Model*. Cambridge University Press, Cambridge.

Colander, David, Peter Howitt, Alan Kirman, Axel Leijonhufvud, and Perry Mehrling. 2008. Beyond DSGE Models: Toward an empirically based macroeconomics. *The American Economic Review* 98(2): 236–240.

Dawid, Herbert, Philipp Harting, and Michael Neugart. 2014. Economic convergence: Policy implications from a heterogeneous agent model. *Journal of Economic Dynamics & Control* 44(C): 54-80.

Dawid, Herbert, and Michael Neugart. 2015. Agent-Based Macroeconomic Modelling and Policy Analysis: The Eurace@Unibi Model. In: Chen, Shu-Heng, and Mak Kaboudan, eds. *Handbook on Computational Economics and Finance*, Oxford University Press, Oxford.

De Grauwe, Paul. 2012. Booms and Busts in Economic Activity: A Behavioral Explanation. *Journal of Economic Behavior & Organization* 83: 484-501.

De Grauwe, Paul, and Agnieszka Markiewicz. 2013. Learning to Forecast the Exchange Rate: Two Competing Approaches. *Journal of International Money and Finance* 32: 42–76.

Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide. 2015. Inflation in the Great Recession and New Keynesian Models. *American Economic Journal: Macroeconomics* 7(1): 168-96

DellaVigna, Stefano. 2009. Psychology and Economics: Evidence from the Field. Journal of Economic Literature 47(2): 315-372.

Delli Gatti, Domenico, Corrado Di Guilmi, Edoardo Gaffeo, Mauro Gallegati, Gianfranco Giulioni, and Antonio Palestrini. 2005. A new approach to business fluctuations: heterogeneous interacting agents, scaling laws and financial fragility. *Journal of Economic Behavior and Organization* 56: 489-512.

Delli Gatti, Domenico, Mauro Gallegati, Bruce Greenwald, Alberto Russo, and Joseph E. Stiglitz. 2010. The financial accelerator in an evolving credit network. *Journal of Economic Dynamics and Control* 34(9): 1627-1650.

Delong, J. Bradford. 2011. Economics in Crisis. *The Economists' Voice*, May.

Dosi, Giovanni, Giorgio Fagiolo, and Andrea Roventini. 2010. Schumpeter meeting Keynes: a policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics & Control* 34: 1748-1767.

Dosi, Giovanni, Giorgio Fagiolo, Mauro Napoletano, and Andrea Roventini. 2013. Income distribution, credit and fiscal policies in an agent-based Keynesian model. *Journal of Economic Dynamics & Control* 37: 1598-1625.

Durieu, Jacques, and Philippe Solal. 2012. Models of Adaptive Learning in Game Theory. In Arena, Richard, Agnès Festré, and Nathalie Lazaric. *Handbook of Knowledge and Economics*, Edward Elgar, Cheltenham.

Epstein, Joshua M. 2007. *Generative Social Science: Studies in Agent-Based Computational Modelling*. Princeton University Press, Princeton, NJ.

Evans, George W., and Seppo Honkapohja. 2013. Learning as a Rational Foundation for Macroeconomics and Finance. In Frydman, Roman, and Edmund S. Phelps, eds. *Rethinking Expectations: The Way Forward for Macroeconomics.* Princeton University Press, Princeton, NJ.

------------. 2009. Learning and Macroeconomics. *Annual Review of Economics* 1(1): 421-451.

Farmer, J. Doyne, and Duncan Foley. 2009. The economy needs agent-based modelling. *Nature*, 460: 685-686.

Farmer, Roger, Vladimir Khramov, and Giovanni Nicolò. 2013. Solving and Estimating Indeterminate DSGE Models. NBER Working Papers 19457.

Fernández-Villaverde, Jesus, Pablo A. Guerrón-Quintana, and Juan F. Rubio-Ramírez. 2010. Fortune or Virtue: Time-Variant Volatilities Versus Parameter Drifting in U.S. Data. NBER Working Paper 15928.

Freeman, Richard B. 1998. War of the models: Which labor market institutions for the 21[st] century? *Labor Economics* 5: 1-24.

Geanakoplos, John, Robert Axtell, J. Doyne Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M. Palmer, and Chun-Yi Yang. 2012. Getting at Systemic Risk via an Agent-Based Model of the Housing Market. *American Economic Review* 102(3): 53-58.

Gigerenzer, Gerd, Ralph Hertwig, and Thorsten Pachur, eds. 2011. *Heuristics: The foundations of adaptive behavior.* Oxford University Press, New York.

Gigerenzer, Gerd, and Peter M. Todd 1999. *Simple Heuristics that Make us Smart*. Oxford University Press, New York.

Gintis, Herbert. 2007. The Dynamics of General Equilibrium. *The Economic Journal*, 117: 1280-1309.

Gode, Dhananjay K, and Shyam Sunders. 1993. Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy* 101(1): 119-137.

Grazzini, Jakob, and Matteo G. Richiardi. 2015. Consistent Estimation of Agent-Based Models by Simulated Minimum Distance. *Journal of Economic Dynamics and Control* 51: 148-165.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante. 2009. Quantitative Macroeconomics with Heterogeneous Households. Annual Review of Economics, 1(1): 319-354.

Hommes, Cars H. 2013. *Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems*. Cambridge University Press, Cambridge.

Hommes, Cars H., and Paolo Zeppini. 2014. Innovate or Imitate? Behavioral technological change. *Journal of Economic Dynamics & Control* 48(C): 308-324.

Kindleberger, Charles (1978). *Manias, Panics, and Crashes: A History of Financial Crises*. New York: Basic Books.

Kirman, Alan P. 2010. The economic crisis is a crisis for economic theory. *CESifo Economic Studies* 56(4): 498–535.

Kirman, Alan P. 2011. Learning in Agent-based Models. *Eastern Economic Journal* 37(1): 20-27.

Krugman, Paul. 2011. The profession and the crisis. *Eastern Economic Journal* 37: 303–312.

LeBaron, Blake, and Leigh Tesfatsion. 2008. Modelling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents. *The American Economic Review* 98(2): 246-250

Leombruni, Roberto, and Matteo G. Richiardi. 2005. Why Are Economists Sceptical About Agent-Based Simulations? *Physica A*, 355(1): 103-109 .

Leombruni, Roberto, Matteo G. Richiardi, Nicole J. Saam, and Michele Sonnessa. 2006. A Common Protocol for Agent-Based Social Simulation. *Journal of Artificial Societies and Social Simulations* 9(1): art 15.

Li, Hongyan, and Leigh Tesfatsion. 2012. Co-learning patterns as emergent market phenomena: An electricity market illustration. *Journal of Economic Behavior & Organization* 82: 395–419.

Minsky, Hyman P. 1982. *Can 'it' Happen Again? Essays on Instability and Finance*. M E Sharpe, New York.

------------. 1986. *Stabilizing an Unstable Economy*. Yale University Press, New Haven.

Mobius M, Rosenblat T. (2014). Social Learning in Economics. *Annual Review of Economics* 6:827-47.

Neugart, Michael, and Matteo G. Richiardi. 2015. Agent-based models of the labor market. LABORatorio Revelli WP 125/2012. In Chen, Shu-Heng, and Mak Kaboudan, eds. *Handbook on Computational Economics and Finance*, Oxford University Press, Oxford.

Paddrik, Mark E., Roy Hayes Jr., William T. Scherer, and Peter Beling. 2014. Effects of Limit Order Book Information Level on Market Stability Metrics. Office of Financial Reasearch, *OFR Working Paper* 14-09.

Poledna, Sebastian, Stefan Thurner, J. Doyne Farmer, and John Geanakoplos. 2014. Leverage-induced systemic risk under Basle II and other credit risk policies. *Journal of Banking & Finance* 42: 199-2012.

Richiardi, Matteo G. 2006. Toward a Non-Equilibrium Unemployment Theory. *Computational Economics* 27(1): 135-160.

------------. 2012. Agent-based Computational Economics. A Short Introduction. *The Knowledge Engineering Review* 27(S2): 137-149.

------------. 2013. The missing link: AB models and dynamic microsimulation. In: Leitner S, Wall, F (eds). *Artificial Economics and Self Organization. Agent-Based Approaches to Economics and Social Systems*. Springer, Lecture Notes in Economics and Mathematical Systems, vol. 669, Berlin.

Richiardi, Matteo G., and Ross Richardson. 2016. Agent-based Computational Demography and Microsimulation using JAS-mine. In Grow, André, and Jan van Bavel. Agent-Based Modelling in Population Studies. Springer Series on Demographic Methods and Population Analysis, Berlin.

Schelling, Thomas C. 1969. Models of Segregation. American Economic Review 59(2): 488-493.

Sinitskaya, Ekaterina, and Leigh Tesfatsion. 2015. Macroeconomies as Constructively Rational Games. *Journal of Economic Dynamics & Control,* 61 (December):152-182.

Smets, Frank, and Raf Wouters. 2003. An Estimated Dynamic Sochastic General Equilibrium Model of the Euro Area. *Journal of the European Economic Association* 1(5): 1123-1175.

Smith, Noah. 2014. The most damning critique of DSGE. *Noahpinion blog*, 14 January 2014.

Stiglitz, Joseph. E. 2011. Rethinking macroeconomics: What failed, and how to repair it. *Journal of the European Economic Association* 9:591–645.

Stiglitz, Joseph E., and Mauro Gallegati. 2011. Heterogeneous Interacting Agent Models for Understanding Monetary Economies. *Eastern Economic Journal* 37(1): 6-12.

Teglio, Andrea, Marco Raberto, and Silvano Cincotti. 2012. The Impact Of Banks' Capital Adequacy Regulation On The Economic System: An Agent-Based Approach. *Advances in Complex Systems* 15(supp02): 1250040.

Tesfatsion, Leigh. 2011. Agent-based Modelling and Institutional Design. *Eastern Economic Journal* 37(1): 13-19

Todd, Peter M., and Gerd Gigerenzer. 2012. *Ecological Rationality: Intelligence in the World*. Oxford University Press, Oxford.

Trichet, Jean-Claude. 2010. Reflections on the nature of monetary policy non-standard measures and finance theory. Speech at the Opening address at the ECB Central Banking Conference, Frankfurt, 18 November 2010.

Vriend, Nicolaas J. 2000. An illustration of the essential difference between individual and social learning, and its consequences for computational analyzes. *Journal of Economic Dynamics & Control* 24(1): 1-19.

Weisbuch, Gérard, and Stefano Battiston. 2007. From production networks to geographical economies. *Journal of Economic Behavior & Organization* 64(3-4): 448-469.

Wickens, Michael R. 2014. How did we get to where we are now? Reflections on 50 years of macroeconomic and financial econometrics. *University of York Discussion Papers in Economics* No. 14/17.

Wolfson, Michael, Steve Gribble, and Reed Beall. 2016. Exploring Contingent Inequalities – Building the Theoretical Health Inequality Model. In Grow, Andre', and Jan van Bavel. Agent-Based Modelling in Population Studies. Springer Series on Demographic Methods and Population Analysis, Berlin.

Woodford, Michael (2013). Macroeconomic Analysis Without the Rational Expectations Hypothesis. *Annual Reviews of Economics* 5: 303-346.

Woolridge, Michael. 2009. *An Introduction to Multiagent Systems*. John Wiley & Sons, Chichester.

NOTES.

[1] For a discussion of the roots of AB modelling, see Richiardi [2013].

[2] While basically all DSGE, RBC and NK models are macro models, most AB models fall outside the macro literature (many pertain to ecological and consumers' choice modelling or finance). When adding the keywords "macroeconomy" or "macroeconomics" to the search, the fraction of AB models to DSGE/RBC/NK models falls to 3.6%.

[3] FP6-STREP grant 035086, http://cordis.europa.eu/project/rcn/79429_en.html. See also: http://www.wiwi.uni-bielefeld.de/lehrbereiche/vwl/etace/Eurace_Unibi/.

[4] FP7-SSH grant 225408, http://cordis.europa.eu/project/rcn/89951_en.html.

[5] FP7-ICT grant 288501, http://cordis.europa.eu/project/rcn/101350_en.html.See also: http://www.crisis-economics.eu/.

[6] "*If you didn't grow it, you didn't explain it*" [Epstein, 2007].

[7] As pointed out by one referee, there is always a risk in having a restricted number of workhorse models dominating the scene. However, the computational nature of AB models implies that it will always be relatively easy to explore departures from the standard models.

[8] https://cs.gmu.edu/~eclab/projects/mason/.

[9] http://repast.sourceforge.net/.

[10] www.jas-mine.net.

[11] https://ccl.northwestern.edu/netlogo/.

[12] To this regard, the literature on AB modelling seems to be completely unaware of both the problems and the available solutions. In particular, implementation of agents as UML states machines, an evolution of finite state machines (FSMs), is worth consideration. FSMs are important in event-driven programming because they make event handling explicitly dependent on both the event-type (e.g. the contracts an agent is offered) and on the state of the system (e.g. the agent's internal state). FSMs can drastically cut down the number of execution paths through the code, simplify the conditions tested at each branching point, and simplify the switching between different modes of execution. Conversely, using event-driven programming without an underlying FSM model can lead programmers to produce error prone, difficult to extend and excessively complex application code. Though traditional FSMs are an excellent tool for tackling smaller problems, they tend to become unmanageable for more involved systems. Due to the phenomenon known as state and transition explosion, the complexity of a traditional FSM tends to grow much faster than the complexity of the system it describes. This happens because the traditional state machine formalism inflicts repetitions. UML state machines address exactly this shortcoming of the conventional FSMs. They provide a number of features for eliminating the repetitions so that the complexity of a UML state machine no longer explodes but tends to faithfully represent the complexity of the reactive system it describes.

[13] The distinction between an individual characteristic and a parameter is sometimes arbitrary. Generally speaking, characteristics can evolve over time, while parameters remain constant. However, it is possible to think of immutable characteristics (like gender, or "human capital potential"), and time-dependent parameters (like technological parameters). Distinguishing between parameters and states has mainly a conceptual value, and has no impact on the functioning of an agent.

[14] Breaking a contract should result in a revision of the list of commitments. Note that transactions do not need to be consensual (for instance, breaking a contract is typically a one-sided decision).

[15] Conditional commitments can be represented as *maps*, where the conditions (including time) define the keys, and value determines the change in the specific endowment considered.

[16] Meta-decision rules can be devised that change a decision rule on the basis of the states.

[17] It is thus in principle possible that agents look for a newly released smart phone in shopping malls, while the producer only sells it over the internet: the consumers do not get their craved for little object, while the launch of the new phone ends up in a flop. This will translate in poor measures of fitness (a state), both for the consumers and the producer, possibly prompting a change in the respective search strategy (an action).

[18] I thank Doyne Farmer for having pointed this out.

[19] See Richiardi and Richardson (2015), available for download at www.jas-mine.net.

[20] An exercise aimed at testing the performance of the simulation platform with respect to scaling involved the implementation in JAS-mine of a complex mixed AB-microsimulation model of the two-way relationship between health and economic inequality [Wolfson et al., 2016], calibrated on both US and Canadian cities. The JAS-mine implementation can run 5 million agents with a time-step equivalent to 1 day for 500 years (182,500 time-steps) in 50 minutes on a standard laptop (using less than 4GB of RAM).

[21] As a referee noted, a first step would be to consolidate the achievements in AB modelling obtained so far. I cannot stress enough the desirability of some collaborative effort in order to formulate standards or best practices regarding model documentation, testing, accessibility of code and tools, and replicability of experiments for the plenty of models that already exist - see also [Leombruni et al., 2006].

[22] This effectively restricts the choice to platforms based on commonly used programming languages, as Java, C++, and possibly R, Matlab and Python (the latter however have some limitations as far as speed of execution is concerned).