Agent-Based Social Simulation in Markets

Koen Bertels [a] Magnus Boman [b]

[a] *Dept of Business Administration, Univ of Namur, Rempart de la Vierge 8, B-5000 Namur, Belgium*
E-mail: koen.bertels@fundp.ac.be
[b] *Dept of Computer and Systems Sciences, Stockholm Univ and the Royal Inst of Technology, Electrum 230, SE-164 40 Kista, Sweden*
E-mail: mab@dsv.su.se

We show that certain desired behavioural properties of agent-based models can be deterministically induced by an appropriate mathematical structure. We also point out problems related to the handling of parameters, and of the modelling of time, in agent-based models. Our purpose is to illustrate some problems of agent-based social simulations in markets, as a first step towards the more ambitious goal of providing a methodology for such simulations.

1. Introduction

Agent-based computer simulation of social phenomena is rapidly becoming popular in the area of electronic commerce, because of its potential for predicting individual and flock behaviour in markets. It is certainly true that the use of agent-based modelling (ABM) opens up new scientific perspectives, but we must keep in mind the fundamental limitations and pitfalls of computer-based modelling. We provide some examples of problems and issues, as has been reported previously for game theory within the agent community [1]. In particular, agent-based social simulation (ABSS) as a scientific method can be problematic [2].

One of the widely proclaimed advantages of ABM is its ability to exhibit so-called emergent properties: novel, coherent, global, dynamical, and ostensive properties [3]. This concept refers to the observed behaviour of the model, which was not explicitly programmed and is therefore unanticipated in the sense that the theory did not predict it. Not only properties and behaviour can be emergent. The process by which organisations emerge, autopoieses, can be described as striving towards organisations not built on stable traditional structures, the prime example being the emergence of a global friction-free marketplace on which artificial agents can act as electronic intermediaries [4],[5].

We will explore some models which have generated interesting results. We do not intend to question the obtained results but only to clarify how certain results are more the deterministic consequence of the underlying structure than an emergent property. The following stances will be taken and discussed.

- certain desired behavioural properties can be deterministically induced by an appropriate mathematical structure (section 3)
- exploring and completely understanding the parameter regime of a model can be problematic and very labour intensive (section 4)
- methodological aspects of alignment efforts and the modelling of time are in need of further study (section 5)

2. A Naïve Categorisation

There is a difference between ABM in a simulation environment, on the one hand, and ABSS on the other. Programmers in the first category typically spend thousands of lines of code on each agent, and usually consider a small (upper-bounded) number of agents at a time. The most popular languages are C and Java, and complexity is usually a practical, rather than a theoretical concern. For instance, the capacity of the network affects the chances of efficient agent communication more than does the fact that traversing the message generation loop of an agent is NP-hard. Agents may be heterogeneous, communication may be unreliable, and the sensors might have to cope with noise. Moreover, while ABM in a simulation setting merits the study of diffusion, self-organisation, and emergence, this is typically not the final aim. Instead, one resorts to simulations because the target environment is inaccessible, like the surface of Mars [6]. Simulations are often planned to be replaced by physical agents, e.g., robots when the environment is at a desired level of control.

A typical platform for ABM simulations is the RoboCup soccer server [7]. A team there consists of eleven players (plus an optional coach) that all could be written in a different programming language. Agent behaviours are diversified, usually through roles, e.g., left-wing defender. Advanced behaviours, such as dynamic role switching and traits adapting over time are programmed (if only with great difficulty). At RoboCup'99, most teams were written in object-oriented code, in some cases with an eye on organisational attitudes [8],[9]. In this and similar domains, social interaction can easily be observed by using the provided visualisation software. The measure for team performance is complex in that it depends on the performance of the opposing team, and in that all deliberation is made under uncertainty, but straightforward in that the overall goal of all teams is the same: score more often than the opposing team does. In RoboCup, a number of additional measures are easily obtained after each game, such as ball possession, number of successful passes, and number of successful interceptions. Emergent behaviour is usually a side-effect, caused by the behaviour of the opposing team. For example, unexpectedly defensive play in a team might result from the opposing team repeatedly making offside traps, which the team detects. Realisations in terms of physical agents in RoboCup are chiefly made in terms of competitions with physical robots, which always display a lower level of sophistication in multi-agent interaction and deliberation than does the agents in the simulated league (cf. [10]).

The second category consists mainly of researchers in the social sciences. Among these researchers, there is an on-going debate on whether one should run simulations with a large number of simple agents, or run simulations with a small number of complex agents. While the model might be extremely complicated in the first case, each agent is relatively simple, consisting of less than a hundred lines of code in a functional language, or in forms suitable for rapid development, e.g., Java applets or Tcl/Tk. In the second case, agent modelling is typically inspired by cognitive models of reasoning, as in the SOAR language [11]. While in the second category, a cellular automaton typically represents an agent (or in the case of Sugarscape, the environment), in the first, it is only the part of the code used for the modelling of states and transitions [12]. Other important differences include the size of each agent, and the required effort that goes into design and optimisation. Moreover, neither time, nor graceful degradation are constraints in the second category. Note that the modelling of time might be a constraint, however, but this poses problems of its own, as we will show below. Noise is usually introduced as deviation from the programmed agent behaviour, i.e. it is used as a qualitative measure. Uncertainty is correspondingly modelled as the extent of such deviations. While naïve, this modelling has its advantages, the most important being that noiseless simulations can be run for control group generation.

An interesting distinction has been made by Axtell [13], who restricts ABSS to computational models with explicitly soluble equations. Axtell then separates his treatment of what amounts to our first category above into those agent-based models where the equations can be written down, and those where they cannot. While the mathematical conditions can indeed merit such separations (cf. [14]), one should not pretend that there is a strong separation between our two categories with respect to choice of environment or language—the individual variation between implementations is too large—but most programming platforms do not appeal to both. There are different programming environments and platforms, each of which is appropriate for a certain class of problems. SDML (see `www.cpm.mmu.ac.uk/sdml/`) allows the construction of relatively deep models involving a small number of agents, whereas Swarm (see `www.santafe.edu/projects/swarm/`) and Ascape (see `www.brook.edu/ES/dynamics/models/ascape/`) typically allow larger-scale simulations with more shallow representations of agent behaviour.

Unfortunately, performance measures are not always available, and often badly described. Because of its bottom-up characteristics, this type of simulation also predicts macro-level behaviour with great difficulty. In the worst case, emergent behaviour reduces to an excuse for the designer not having a semantics—in the sense of a description of the space of possibly exhibited behaviours—for the program prior to execution. If there is some way to describe this space, the situation is different (cf. [14]). We can then model the different possible outcomes, their plausibility, and develop metrics for deviations: "An appeal to emergence is thus a way to describe the need to go to the macro level and its unique dynamics, laws, and properties in order to explain more adequately what is going on. The construct of emergence is therefore only a foundation on which to build an explanation, not its terminus" [3], p.58. Roles are usually simplified to a single rule governing behaviour, as in when to co-operate and when to defect in an iterated game.

3. Knowing the Equations

In this section, we address the issue of a model's behaviour which is deterministically induced by its mathematical structure. It is normal to construct a model in function of the phenomena one wants to study. However, it is important to be fully aware of the structural properties and its impact on the overall be-

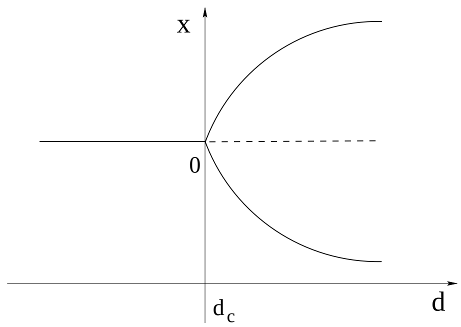

Figure 1. Pitchfork Bifurcation.

haviour. We will discuss a model which aims to simulate the buying behaviour of end consumers in the Marseille Fishmarket [15]. The authors particularly wanted to be able to model the loyalty of consumers as they tended to systematically go back to the fish store that served them well. To this purpose, they used the probabilistic physical mean field theory. The preference of the consumer is given by the following differential equation:

$$\frac{dJ_j}{dt} = -\gamma J_j + < \pi > \tag{1}$$

where γ is a learning parameter of the buyer and π is the mean profit of the buyer when going to reseller j. In the simplified case where there are only two sellers on the market, equation 1 becomes:

$$\frac{d\Delta}{dt} = f(\Delta, \beta) = -\gamma\Delta + \pi tanh(\beta\frac{\Delta}{2}) \tag{2}$$

where Δ is the difference between the preferences for either two sellers and β is a parameter.

The authors then ran this particular model which revealed the existence of two market regimes in function of the parameter β. When β has a value beyond a certain threshold value, they found that the market organises itself into a stable configuration where buyers systematically go to the same merchant. When β does not exceed this value, the buyers never become loyal to one seller and their preferences are distributed randomly. The question we want to address is whether or not this result can be considered an emergent property of the simulation. However, the particular behaviour of the model is typical for a pitchfork bifurcation [16] where a system at some point can have two different solutions or states.

As shown in the figure, the system can then be in only one of these states and it is not possible to predict in advance what branch of the bifurcation will lead to the end state. Such a bifurcation can arise when certain mathematical conditions of the model are satisfied. As the model tries to describe the evolution of the buyers' preferences, we look at equation 2. It is known that a pitchfork bifurcation will occur at point $(0, \beta_c)$, when the following conditions are met [17]:

$$\frac{\delta f(0, \beta_c)}{\delta \Delta} = 0 \tag{3}$$

$$\frac{\delta^3 f(0, \beta_c)}{\delta \Delta^3} \neq 0 \tag{4}$$

$$\frac{\delta^2 f(o, \beta_c)}{\delta \Delta \delta \beta} \neq 0 \tag{5}$$

where β_c is the critical value of β given by $\frac{2\gamma}{\pi}$.

This implies that whenever a value for the β-parameter larger than the critical threshold is chosen, the system will evolve into, and remain in, either of the two branches of the bifurcation. In that case, the buyers are said to have chosen a particular reseller. In the case where the threshold has not been exceeded, the buyer will be indifferent with respect to the two resellers.

4. Knowing the Parameter Regime

Parameters are constants in equations which normally remain fixed but can be changed in function of a desired behaviour. We will show by means of a simple difference equation how the value of one parameter can drastically change its behaviour. The equation is the Verhulst equation which has the following quadratic structure:

$$x_{t+1} = \alpha x_t (1 - x_t) \tag{6}$$

The equation was introduced to explain the population growth in terms of its birth and death rate, x_t represents the normalised population at time t, and α is a parameter. It is easy to see that this equation has two solutions, the trivial solution $x^* = 0$ and $x^* = 1 - 1/\alpha$. We can therefore expect to always end up in either state, irrespective of the starting value of x_0. However, the situation is far more complex and it turns out that this simple equation can have different kinds

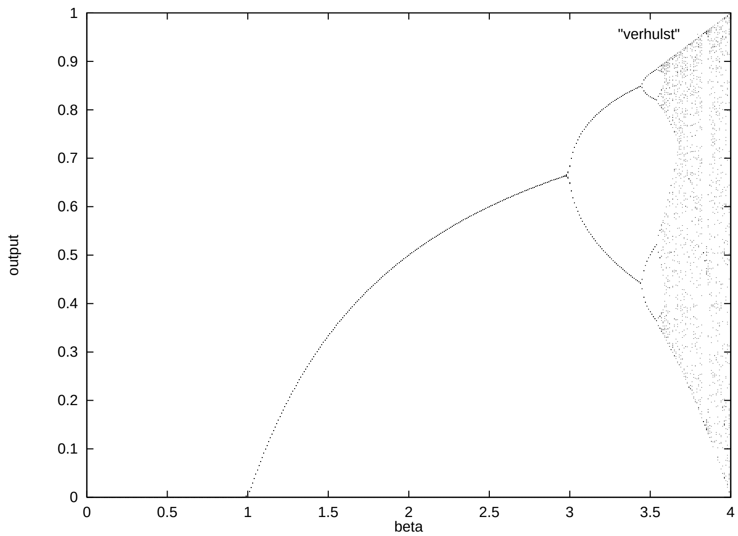

Figure 2. Bifurcation diagram for the Verhulst equation.

of solutions, depending on the value of the α-parameter. The following solutions are possible. In function of the values of the parameter α, we can distinguish between four different kinds of behaviour of our system:

1. $\alpha < 1$: the output of the system will always be 0.

2. $1 < \alpha < 3$: the output is unique: $x^* = 1 - 1/\alpha$.

3. $3 < \alpha < 3.569999$: the system enters a limit cycle of period 2. This means that the system will eventually produce the sequence of values $x_1^*, x_2^*, x_1^*, x_2^*, ...$ and as α is increased, the periodicity will double and become 4, 8, 16, ...

4. $\alpha > 3.569999$: the system enters a chaotic state corresponding to an infinite eriodicity, meaning that the system will have an endless number of states.

We can graphically sketch these different solutions by means of a bifurcation diagram, shown in Figure 2. On the x-axis, the different values of α are plotted and the y-axis gives the value of the system's solution after a number of iterations. Similar phenomena occur in backpropagation neural networks [18]. These networks are often used to model the learning behaviour of agents.

5. Modelling Time

A wide-spread report on alignment by Huberman and Glance shows the importance of modelling the passing of time [19]. The authors replicate a simulation purporting to produce spatial patterns of co-operate and defect moves in a Prisoner's Dilemma, using synchronous updates. Since Huberman and Glance reason that synchronous updates are rare in the real world, they move to asynchronous updates. Their simulation then quickly reaches a fixed state where all agents defect. Since so many simulations, and simulation environments, depend on synchronous updating, one must ask whether (a) this is a realistic proviso, and (b) what the effects would be of switching to asynchronous updating.

While Sugarscape is an excellent tool for studying nonequilibrium dynamics, it is typically used with synchronous updating, just as most other tools. The conditions governing asynchronous updates have to be distributed over the agent space in any tool claiming to model either real life or a synthetic environment without a global clock. This stands in contrast to the existence proofs, search methods, and metrics related to Walrasian equilibria in economics, where the auctioneer represents an opportunity for global updating and control, as in the market-oriented programming approach [20],[21]. Here, the general equilibrium price is the main emergent property. In Sugarscape simulations, what the authors call collective structures instead emerge: "tribes of agents, stationary wealth distributions, and collective patterns of movement, for example" [12], p.17. Even in these relatively simple simulations, such emergent phenomena result from nonlinear interactivity among agents. If we consider collective patterns, for instance, there is a growing amount of examples of social simulation studies exhibiting them and rooting them in sociological and psychological terminology [22]. As the Huberman and Glance study shows, however, one should always investigate the replication problem also with respect to time. A methodology supporting this kind of investigations is today widely acknowledged as useful for the entire field of social simulations, and we have addressed this question in a separate paper [23].

Another reason for being careful with the modelling of time is the effect it has on the so-called micro/macro link fundamental to sociology and to social simulations [24]. A core question in MAS research is how global or macro goals (or rationality, or control) can be brought about by local or micro goals [25]. The active area of criticalities [26] affect sociological (see, e.g., [27]) as well as computer science [28] considerations of this question. Examples of bold statements about how well a particular tool may cope with the micro/macro link abound in the MAS literature. In most domains, however, the micro/macro problem cannot be addressed in a solely bottom-up fashion. If the agents adapt, plan, constrain, or take any other form of deliberative action, micro design can be shown to override macro design of agent properties that affect individual agent choices and that are programmed bottom-up. Such overrides can be demonstrated mathematically, e.g., in terms of classical decision theory [29].

Just as sociology is not readily applicable to just any ABSS, one might encounter problems with applying microeconomics. For instance, several studies have reported results by referring to theorems of social welfare theory as being relevant to the simulations ran. A pitfall here is that the theory has over the years been adapted to fit the real world, in order to be able to account for real world phenomena, and that these adaptations are unnecessary or even fallacious when artificial agents replace humans. For instance, the interpersonal non-comparability of individual utilities has no straightforward counterpart in inter-agent communication among utilitarian agents [30]. A more non-encompassing model in fact does fine, the stability of which might be studied, e.g., in terms of robustness under varying auction protocols [31].

6. Conclusion

In this paper, we have pointed out a number of potential pitfalls in agent-based social simulations of which one should be aware when creating an agent-based model. An appropriate mathematical structure can and will in a deterministic way generate a certain desired or undesired behaviour. So when observing certain phenomena during the simulation, we should always first try to look for a "deterministic" answer in the sense that we must attempt to determine what part of the model induced it. As we showed in our discussion on the parametrisation problem, a small structural part of the model can have a great impact on the overall behaviour. Knowing and hence extensively studying and testing the model is of the utmost importance. Given the fact that even relatively simple models can have multiple, multiple-valued parameters, parametric studies may be a practical impossibility. If agent-based modelling is to become a standard and a basis for useful tools, then these issues must be given proper attention.

Acknowledgments

The authors would like to thank Lars Rasmusson and Harko Verhagen for comments on an earlier draft. Luc Neuberg contributed significantly to the understanding of the Marseille Fishmarket model. Magnus Boman carried out this research within the PROMODIS programme (see `www.dsv.su.se/~mab/PROMODIS/Index.html`), sponsored by NUTEK (The Swedish National Board for Industrial and Technical Development).

References

[1] Bjørn Lomborg. Game theory vs. multiple agents: The iterated prisoner's dilemma. In Cristiano Castelfranchi and Eric Werner, editors, *Artificial social systems*, number 830 in Lecture notes in AI, pages 69–93. Springer-Verlag, 1994.

[2] Robert Axelrod. Advancing the art of simulation in the social sciences. In R. Conte, R. Hegselmann, and P. Terna, editors, *Simulating Social Phenomena*, pages 21–40. Springer-Verlag, 1997.

[3] Jeffrey Goldstein. Emergence as a construct: History and issues. *Emergence*, 1(1):49–72, 1999.

[4] Duane P. Truex, Richard Baskerville, and Heinz Klein. Growing systems in emergent organizations. *Communications of the ACM*, 42(8):117–123, 1999.

[5] Yannis Bakos. The emerging role of electronic marketplaces on the Internet. *Communications of the ACM*, 41(8):35–42, 1998.

[6] Luc Steels. Cooperation between distributed agents through self-organisation. In Yves Demazeau and Jean-Pierre Müller, editors, *Decentralized AI*, pages 175–196. Elsevier Science, 1990.

[7] Itsuki Noda, Shoji Suzuki, Hitoshi Matsubara, Minoru Asada, and Hiroaki Kitano. Overview of RoboCup-97. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, number 1395 in Lecture Notes in AI, pages 20–41. Springer Verlag, 1997.

[8] Silvia Coradeschi, Tucker Balch, Gerhard Kraetzschmar, and Peter Stone, editors. *RoboCup99 Team Descriptions: Simulation League*. Linköping Univ Electronic Press, 1999.

[9] Alexis Drogoul and Anne Collinot. Applying an agent-oriented methodology to the design of artificial organizations: A case study in robotic soccer. *Autonomous agents and multi-agent systems*, 1(1):113–129, 1998.

[10] Magnus Boman. Agent programming in RoboCup'99. *AgentLink Newsletter*, 1(4):20–22, 1999.

[11] Allen Newell. *Unified theories of cognition*. Harvard Univ Press, 1990.

[12] Joshua M. Epstein and Robert Axtell. *Growing Artificial Societies*. Brookings Institution Press and the MIT Press, 1996.

[13] Robert Axtell. Varieties of agent-based computational models in the social sciences. In *Workshop on Agent Simulation*. Univ of Chicago, October 15-16 1999.

[14] Steen Rasmussen and Christopher Barrett. Elements of a theory of simulation. Technical Report 95-04-040, Santa Fe Institute, 1995.

[15] J.-P. Nadal, G. Weisbuch, O. Chenevez, and A. Kirman. A formal approach to market organization: Choice functions, mean field approximation and maximum entropy principle. In Jacques Lesourne and André Orlean, editors, *Advances in Self-Organization and Evolutionary Economics*, pages 149–159. Economica, 1998.

[16] E.C. Zeeman. Bifurcation and catastrophe theory. *Contemporary Mathematics*, 9:207–272, 1982. Notes by Catriona Glenton.

[17] H. Lorenz. *Nonlinear Dynamical Economics and Chaotic Motion*. Springer-Verlag, 1989.

[18] K. Bertels, L. Neuberg, S. Vassiliadis, and G. Pechanek. Chaos and neural networks: The backpropagation paradigm. *Artificial Intelligence Review*, to appear.

[19] Bernardo A. Huberman and Natalie S. Glance. Evolutionary games and computer simulations. *Proc National Academy of Sciences USA*, pages 7716–7718, 1993.

[20] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford Univ Press, 1995.

[21] Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

[22] Amedeo Cesta, Maria Miceli, and Paola Rizzo. Help under risky conditions: Robustness of the social attitude and system performance. In Mario Tokoro, editor, *Proc ICMAS'96*, pages 18–25, 1996.

[23] Koen Bertels and Magnus Boman. Why we need a COMA in agent based social simulation. Presented at the AgentLink SIG meeting on Agent-Mediated Electronic Commerce, paper available upon request, September 1999.

[24] Thomas C. Schelling. Social mechanisms and social dynamics. In Peter Hedström and Richard Swedberg, editors, *Social Mechanisms*, pages 32–44. Cambridge Univ Press, 1998.

[25] Rosaria Conte and Cristiano Castelfranchi. *Cognitive and Social Action*. UCL Press, 1995.

[26] Per Bak and Kim Sneppen. Punctuated equilibrium and criticality in a simple model of evolution. *Physical Review Letters*, 71(24):4083–4086, 1993.

[27] Malcom Gladwell. The tipping point. *New Yorker*, June 3 1996.

[28] Koen Bertels, Jean-Marie Jacques, and Magnus Boman. Implications of self-organised criticality and resilience measures to risk and crises management. Presented at the FRN Workshop on Systems Shocks—Systems Resilience, Abisko, Sweden, May 22-26, 2000.

[29] Love Ekenberg, Mats Danielson, and Magnus Boman. From local assessments to global rationality. *Intelligent Cooperative Information Systems*, 5(2-3):315–331, 1996.

[30] Amartya Sen. Social choice theory. In K.J. Arrow and M.D. Intriligator, editors, *Handbook of Mathematical Economics*, volume III, chapter 22. Elsevier, 1986.

[31] Ken Steiglitz, Michael L. Honig, and Leonard M. Cohen. A computational market model based on individual action. In Scott H. Clearwater, editor, *Market-Based Control*, chapter 1. World Scientific, 1996.