Individual behavior and macro social properties.
An agent-based model

**Riccardo Boero · Marco Castellani ·
Flaminio Squazzoni**

Abstract The paper aims at presenting an agent-based modeling exercise to illustrate how small differences in the cognitive properties of agents can generate very different macro social properties. We argue that it is not necessary to assume highly complicated cognitive architectures to introduce cognitive properties that matter for computational social science purposes. Our model is based on different simulation settings characterized by a gradual sophistication of behavior of agents, from simple heuristics to macro-micro feedback and other second-order properties. Agents are localized in a spatial interaction context. They have an individual task but are influenced by a collective coordination problem. The simulation results show that agents can generate efficiency at a macro level particularly when socio-cognitive sophistication of their behavior increases.

Keywords Agent-based model · Behavioral heuristics · Socio-cognitive properties · Social patterns

1 Introduction

Can a small difference in the cognitive features of agents generate a significant difference in the macro social properties? This paper aims at presenting a simple agent-based modeling exercise to illustrate this issue. The model is based on a stylized

R. Boero
Department of Economic and Financial Studies "G. Prato", University of Torino, Turin, Italy
e-mail: boero@econ.unito.it

M. Castellani · F. Squazzoni (✉)
Department of Social Sciences, University of Brescia, Via San Faustino 74/B, 25122 Brescia, Italy
e-mail: squazzon@eco.unibs.it

M. Castellani
e-mail: castella@eco.unibs.it

coordination problem: agents are endowed with individual goals, they move across space to try to reach their goal, but are localized in an ever changing interaction context in which neighbors affect each other at a local level, so that interactions strongly influence micro and macro properties. This means that macro patterns of the system depend on the self-organized capacity of agents to solve coordination problems at a local level.

At the same time, the model allows emphasizing two relevant issues in the computational social science literature. First, we argue that, for taking into account cognitive properties in agent-based models that matter in the analysis of social phenomena, the introduction of highly complicated cognitive architectures behind agents behavior is not necessary. We show that even small modifications towards the cognitive sophistication of agents behavior generate significant and interesting macro consequences to be studied. This evidence can contribute to the vivid dispute on simplification or complexification of models between cognitive and social scientists within computational social sciences, and in social sciences in general (Conte 1999; Conte et al. 2001; Sun 2001; Sawyer 2005; Gilbert 2005; Hedström 2005; Epstein 2006; Buchanan 2007).

We show that in our model agents find a stable solution of the coordination problem at a system level just in particular parameters conditions at the micro level, and, especially, when agents are endowed with socio-cognitive behavioral capacities. Again, this evidence can be emphasized because of our particular modeling exercise: starting with simple behaviors and increasing the cognitive sophistication of agents step-by-step by means of small modifications.

To explore these issues, we have used an agent-based model (hereafter, ABM). In the last ten years, there are relevant examples of ABMs of social phenomena, both in empirical and theoretical sense, with effective and interesting results at the level of explanatory power (Axelrod 1997; Conte et al. 1997; Prietula et al. 1998; Hedström 2005; Epstein 2006; Gilbert 2008). This holds also for the particular issues investigated in this paper. For instance, we have followed some relevant examples of ABMs, like Macy (1990), Flache and Hedgselmann (1999), and Jager et al. (2000), to name a few, where a strong emphasis has been put on the relevance of bounded and adaptive rationality of agents for the emergence of observed macro patterns and indicators. ABMs can be of help in overcoming some limits of mathematical and statistical standard tools and narrative languages mostly used in social science research. They stand beside mathematical and statistical methods in term of rigor, but are flexible and intuitive like narrative or verbal language (Gilbert and Terna 2000; Gilbert 2008). They help formalizing models in which heterogeneous, autonomous, adaptive, and spatial localized agents, interacting in a decentralized social context, can generate macro structures or patterns that show some intriguing emergent social properties. This paper aims also at giving a modest example of the potentials of these tools for illuminating social science issues in theoretical terms.

The paper is organized as follows. In Sect. 1 we introduce the structure of the model and the rules of behavior followed by agents. In Sect. 2 we show the different simulation settings in which the model is executed. Section 3 focuses on simulation results. Section 4 elaborates on some discussions and concluding remarks, with a particular attention to the issues outlined in this introduction.

2 The model and the simulation settings

Our model consists of 101 agents located in a toroidal two-dimensions finite space of cells. They stay or move step-by-step across the cells trying to fit their own ideal temperature value, which is assigned heterogeneously and randomly at the start of the simulation.[1] They can see just eight neighboring cells, and they can move just towards empty cells. Moving across space, they produce heat and they leave it in the cell where they are. Heat diffuses synchronously across neighboring cells, equally in all directions, at a rate proportional to the heat differential between cells, and evaporates at a rate proportional to the amount of heat. An evaporation rate, constant over the environment, determines how heat fades. Heat spreads in space, according to a fixed value, from 0 to 1, and it is constantly updated according to the following rule: $H(t + 1) = e * (H(t) * k(\text{Avg}H(t) - H(t)))$, where e is the evaporation rate $(0..1)$, k is the diffusion rate $(0..1)$, and $\text{Avg}H$ is the average temperature in the eight neighboring cells.

Agents interact so that each decision at the single agent level brings about a change in the space of possibilities of others. Moving across cells, trying to reach their individual ideal temperature, they produce continuously a growth of the system's temperature. Agents have two fundamental information: they know the temperature of the 8 neighboring cells surrounding them, and also that the system's temperature depends on the action, position and interaction of agents. The temperature of the cells is the outcome of the interaction amongst agents and is influenced by past interactions. The fit between the ideal temperature of the agent and the current environmental temperature of the cell is calculated by an *attainment discrepancy value* (Mezias 1988; Murphy et al. 2001): the more this value tends to 0, the more the agent has found a perfect environmental fit.

Agents are endowed with different behavioral heuristics. Heuristics work in the agents decision to move or not and to select one cell or another. In a broad sense, they can be considered as "rules of thumb", or "fast and frugal" heuristics that drive agents actions over various computational steps (Gigerenzer and Selten 2001).

2.1 Agents behavioral heuristics

We introduce four different heuristics that are assigned to the population of agents at the beginning of each simulation setting. The first one is labeled *H1*: if the cell temperature is lower than agents ideal temperature, and the *attainment discrepancy* thus is different from zero, agents move to the neighboring cell with the highest temperature, and, if the foregoing condition is negative, they switch to the neighboring cell with the lowest temperature. The second one is *H2*: if the *attainment discrepancy* is different from zero, among the neighboring cells, agents select the one that has a temperature closer to their ideal one. The third one is *H3*: if the *attainment discrepancy* is different from zero, agents can randomly move around the cells, until the *attainment discrepancy* at time t is $\leq$ to their previous one at time t_{-1}. As soon as they succeed,

---

[1] In our simulations random values are always generated according to a uniform distribution in the range set by parameters.

they stop in a cell. The last one is *H4*: if the temperature of cells in which agents are located is higher than their ideal temperature level, and if agents have other neighboring agents (this can cause a further increase of the temperature of the cells), then agents move away from neighboring agents, selecting the first empty cells opposite to them; *vice versa*, if the temperature of cells in which agents are located is lower than their ideal temperature level, and if agents have other neighboring agents (this can cause a further increase of the temperature of the cells), then agents remain in the same cells. If agents do not have any neighbors, they randomly move. Finally, in order to introduce an error effect, we assume that agents have a 20% individual probability of random moving towards one of their empty neighboring cells, whatever the heuristics.

2.2 Types of agents and simulation settings

We introduce different simulation settings with different types of agents. The difference among these settings rests on the growing sophistication of cognitive features of agents behavior. We start from agents that simply apply the single heuristics described before (called Agents1), then we get on to agents able to change the heuristics over time according to more sophisticated cognitive properties (called Agents2), to finally end up with agents able to take advantage of even more sophisticated second-order properties (called Agents3). For sophistication, we mean the introduction of small differences in the modeling of agents towards the gradual increasing of their cognitive capacity to detect some particular macro features of the interaction context in which they are embedded, as well as their capacity of using this information to modify their behavior (Conte and Castelfranchi 1996; Gilbert 1996, 2002). By introducing such second-order properties, we explicitly introduce in the model a macro-micro feedback that influences agents behavior. It is reasonable to suppose that such second-order properties are typical of systems in which intelligent social agents interact (Conte 1999).

Agents1 behave according to one of the four heuristics (*H1*, *H2*, *H3*, *H4*) and can not change them over time. Agents2 act as follows: if their *attainment discrepancy* is 10% ≥ of the average level of the agents population *attainment discrepancy*, they are pushed to shift heuristics, following a random selection. Here, we introduce an explicit macro-micro feedback that influences agents behaviors, without assuming other complicated cognitive properties, like memory or learning of agents. Agents2 use heuristics as a toolbox (Gigerenzer and Selten 2001). Agents3 have the same functioning rules of Agents2, but are divided in two groups: "the greens" and "the blues". They are able to recognize if they are a "green" or a "blue" and to detect agents in their own neighborhood. If they find neighbors of the same group, their *attainment discrepancy* decreases over time, *vice versa* it increases. This is to introduce a higher level cognitive property, which is similar to the tag mechanism investigated in cooperation studies by Hales (2000), Riolo et al. (2001), among many others. To conclude, the differences introduced in these simulation settings are small modifications of behavioral sophistications of agents at the micro level, while there is no change in the other simulation parameters.

We have elaborated on seven simulation settings. In the first four, Agents1 follow respectively *H1*, *H2*, *H3*, *H4* (see from Fig. 1 to Fig. 4 and Figs. 8, 9, and 12). In the

Table 1 Simulation parameters

Parameter	Value
Number of agents	101
World size (in cells)	80×80
Diffusion constant	1
Evaporation rate	0.99
Probability of random moves	0.2
Maximum level of ideal temperature	31,000
Minimum level of ideal temperature	17,000
Maximum level of heat generated by agents	10,000
Minimum level of heat generated by agents	3,000

fifth, Agents1 use all the heuristics at the same time (that is to say that 25% of the population behaves according to *H1*, 25% to *H2*, and so forth—see Figs. 5, 10, 11 and 12). This is to observe the combined effect of heuristics and of heterogeneity at the population level. The sixth setting is based on Agents2 using the heuristic toolbox, while the seventh setting shows the case of Agents3 (Figs. 6, 7, 13, 14 and 15).

The model parameters are summarized in Table 1. As we have already mentioned, there is no change in these parameters across simulation settings. The only thing that changes across simulation settings is therefore agents behavior. Simulations have been run until stable patterns were generated within the system.

3 Simulation results

We use two types of resulting data. The first ones are visualizations of local/global spatial patterns generated by agents over time. They are represented in Figs. 1 to. 7, where we have included some representative snapshots of macro distribution patterns over simulation runs. In the figures, the white zone means a space without agents, the halo represents a zone of temperature concentration around agents distribution patterns, while bright points represent agents localized into cells. This first type of data is used for observing and comparing spatial macro patterns across simulation settings. The second type of data represents the set of quantitative data that we have used to dissect some explanatory mechanisms on the relation between individual behavior and system efficiency. They are: average values of attainment discrepancy over time; maximum values of attainment discrepancy over time; and percentage of agents not moving over time. They are represented in Figs. 8 to 15.

To begin with the analysis of macro visualizations, by comparing the spatial patterns represented in Figs. 1 to 7, we can outline first that Agents1 using *H2* are able to produce more clear-cut spatial patterns than other Agents1. As represented in Fig. 2, around run 300, the system is characterized by a regionalization of spatial patterns, with some greater patterns that begin to emerge. After a growing concentration (see run 700), around run 1500, agents are clusterized into a collective pattern that reaches a stable form. The best performance of *H2* was expected on this. As a matter of fact,
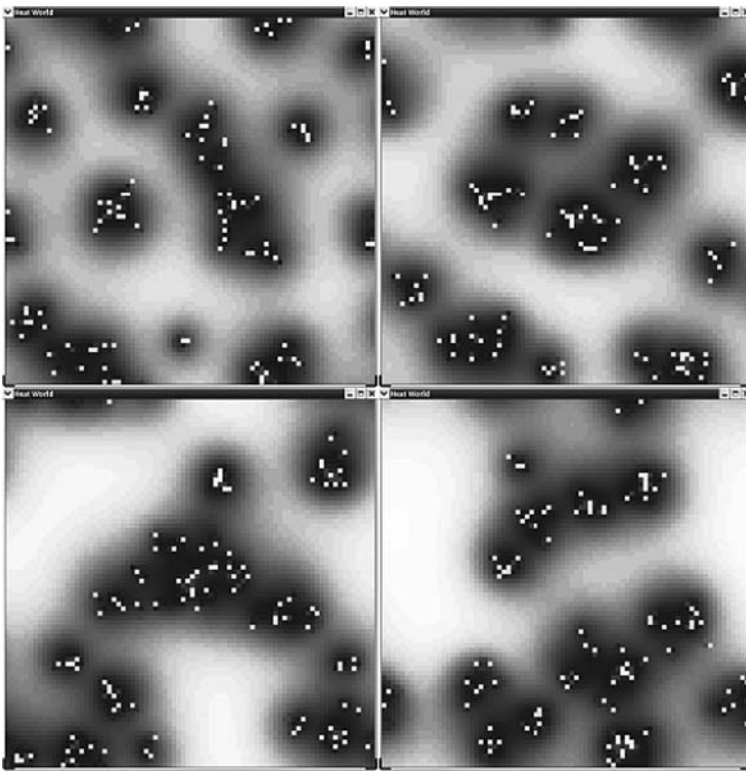
Fig. 1 Simulation setting 1 (Agents1 using *H1*). From *top-left* to *top-right*, and from *down-left* to *down-right*, patterns at 300–700, 1500–5000 simulation run

H2 houses a rule that allows agents to fill directly the cell that is closest to their ideal temperature. On the contrary, for example, *H1* is not of help in generating a big macro stable clusterization instead of many small ones (see Fig. 1). While in the first simulation setting, with *H1*, the halo tends to disappear earlier (around run 700) than in the second simulation setting, with *H2* (around run 1500), the micro clusterizations are in a steady state. A possible explanation of this is that *H1* implies for agents a continuous adjustment around the neighboring cell, since it is based on the attempt of moving towards the cell that satisfies the "basic need" of agents (if it feels hot it looks for the coldest cell and vice versa), without paying attention to the level of the gap between ideal temperature of agents and actual temperature of the cells. The same could be said for *H4*: agents are just pushed away from an unsatisfactory situation, whereas *H2* is focused on allowing agents to directly reach an optimal situation.

A second observation is that Agents1 using the entire heuristics toolbox are able to generate two clustered patterns (see Fig. 5). No particular features are caused by their heterogeneity at a population level. The result is just a simple composition of single dynamics. A third evidence is that Agents2, which are endowed with sophisticated cognitive features (the capacity to elaborate on macro information), clusterize in a single spatial pattern around run 500 (see Fig. 6), without halo. In this case, the sum of
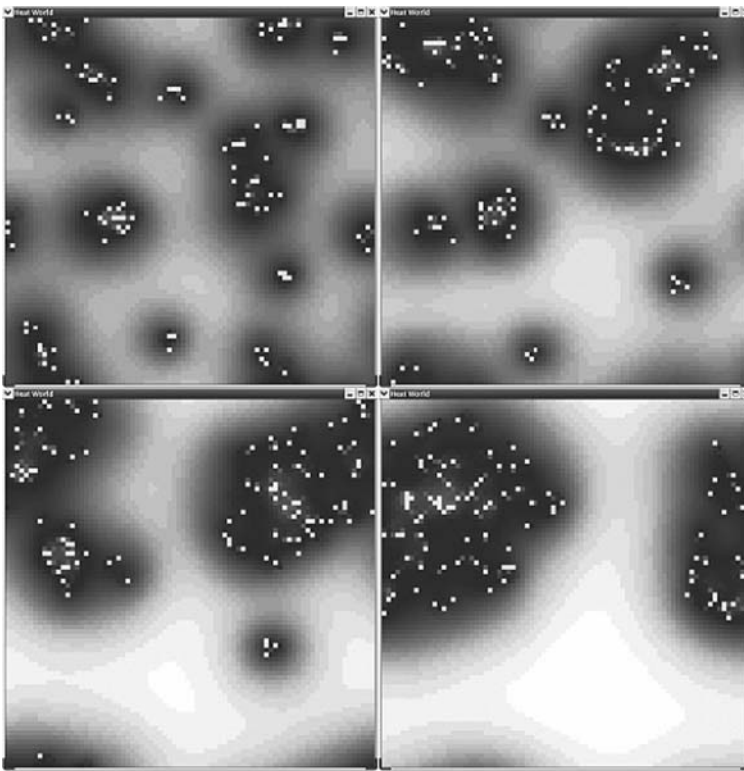
Fig. 2 Simulation setting 2 (Agents1 using *H2*). From *top-left* to *top-right*, and from *down-left* to *down-right*, patterns at 300–700, 1500–5000 simulation run

two factors could explain this situation: the possibility of changing the initial heuristic (as in the fifth simulation setting, at the beginning of the simulation, each heuristic is assigned to the 25% of population), and the presence of *H2*.

Agents3, which have the capability of belonging to and recognizing groups, generate, around run 700, a spatial distribution pattern characterized by two regions: the first one is wider and located in the top-side of the system, while the second one is located more horizontally in the down-side (see Fig. 7). Each region is composed by agents belonging to both groups. Agents of two groups give rise to a chain figuration. Agents warm up each other without melting themselves. Around run 1500, spatial patterns are well formed and more stable, while around run 5000 they finally stabilize: just one agent is again on move. Around 17,200, the system crystallize itself, with each agent stopped within cells. Compared with the spatial patterns of the other simulation settings, this one shows a specific emerging property: each agent finds locally an appropriate cell that fits its temperature, and the system globally finds a stable equilibrium. As we will see with quantitative data below, these last agents have found a stable solution of the coordination problem that is embodied in the structural properties of this model.
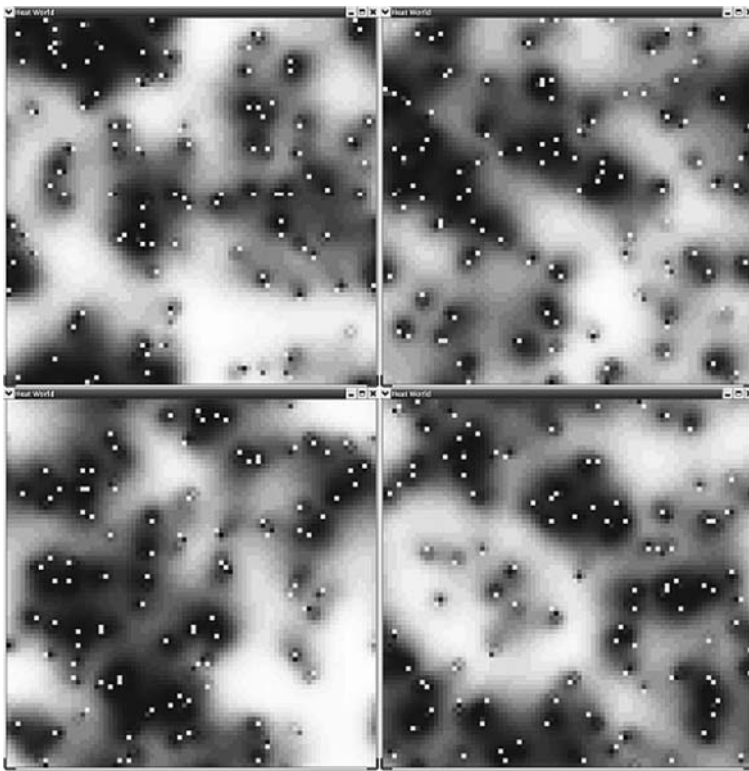
Fig. 3 Simulation setting 3 (Agents1 using *H3*). From *top-left* to *top-right*, and from *down-left* to *down-right*, patterns at 300, 700, 1500 and 5000 simulation run

Passing to the quantitative data, we have first created two comparison settings to get more control on the simulation settings 1 to 5. Such comparison settings are useful for understanding the system's behavior in extreme cases. The first comparison setting is called "Non-moving": in this case, agents are blocked in the position randomly assigned at the beginning of the simulation. The second setting is called "Randomly moving": in this case, agents always move randomly across the cells at each run.

From a descriptive viewpoint, it is worth reporting here that agents in the "Non-moving" setting create regions of temperature around the randomly assigned localization of agents. After a few run, the system gets a stable condition, due to the heating of agents and the evaporation rate. On the contrary, agents in the "Randomly moving" setting generate a completely unstable system, where agents are always moving, with no homogeneous spatial patterns.

Quantitative results presented in Figs. 8 to 15 are average values over 100 simulations made with different values for the random numbers generators. In those figures, the *X* axis presents simulation time on a logarithmic scale. Figure 8 shows the system average values of *attainment discrepancy* in the cases of Agents1 behaving according to *H2* and *H3* and of the two comparison settings just mentioned above. The resulting series about *H1* and *H4* are not depicted in Figs. 8 and 9 for the sake of clarity, but it
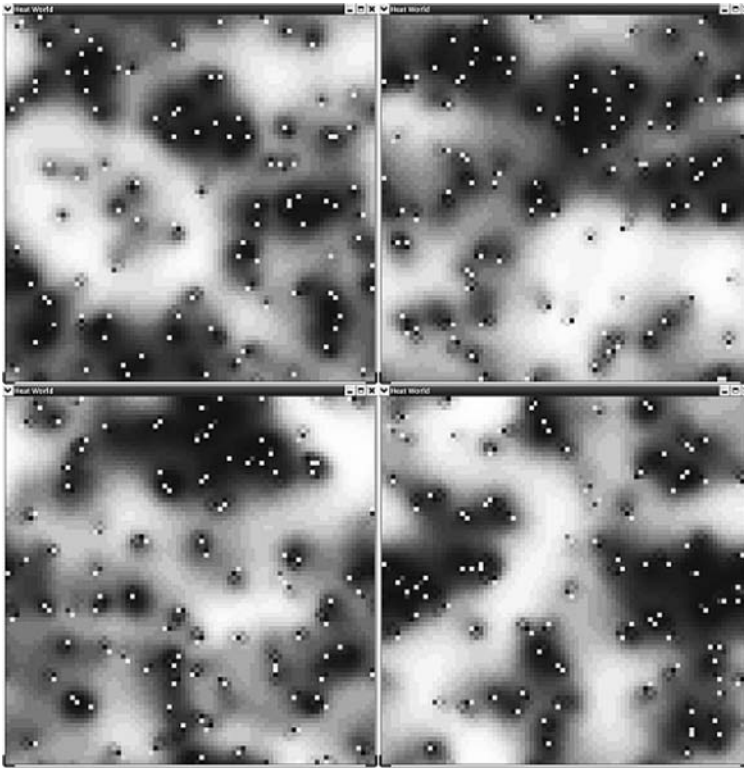
Fig. 4 Simulation setting 4 (Agents1 using *H4*). From *top-left* to *top-right*, and from *down-left* to *down-right*, patterns at 300–700, 1500–5000 simulation run

is worth noting that *H1* gives raise to results almost identical to *H2*, just slightly less efficient (i.e., with higher values of discrepancy), while *H4* is not statistically different from *H3*. The main result is that, at a system level, the most efficient heuristics are *H1* and *H2*. Agents behaving according to *H3* and *H4* obtain values which are only slightly better than random agents, and far worse than non-moving agents.

Figure 9 shows, for the same simulation settings considered in Fig. 8, the maximum value of *attainment discrepancy* at the population level. This helps in verifying that the efficiency of the different heuristics does not imply larger inequality at the population level. As a matter of fact, better average values of discrepancy are positively correlated with better maximum values of discrepancy. Figures 10 and 11 present similar data for the simulation setting 5, the spatial patterns of which have been summarized in Fig. 5. In these figures, simulation setting 5 has been put in relation to the best heuristic found so far (*H2*) and to the two comparison settings. The main result is that setting 5 obtains average values of discrepancy that are in between the ones obtained with the single heuristics. Moreover, the results present a slight worsening of the distribution of discrepancy at the population level. As Fig. 11 clearly points out, the maximum level of discrepancy is higher for setting 5 than for randomly moving agents.
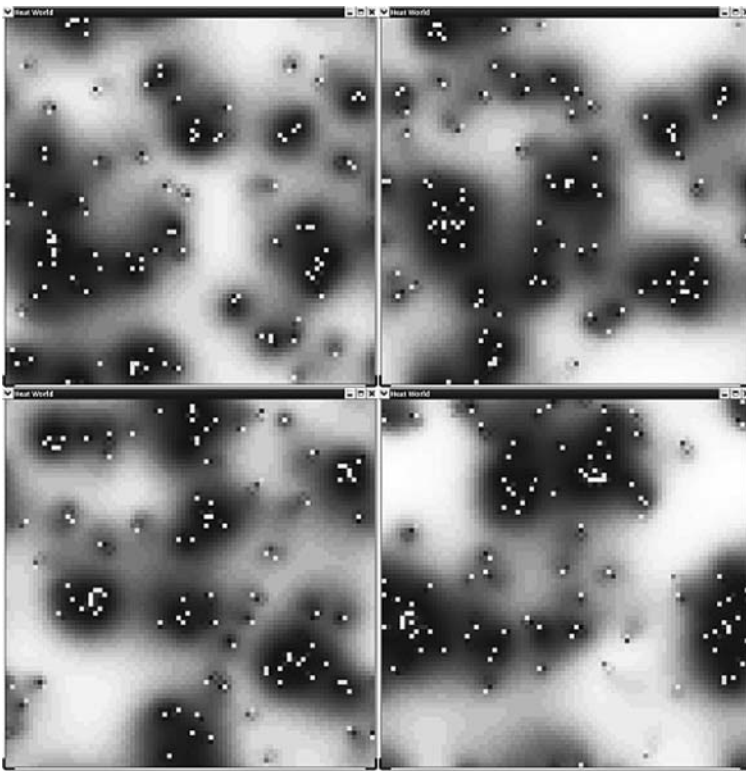
Fig. 5 Simulation setting 5 (Agents1 using all the heuristics). From *top-left* to *top-right*, and from *down-left* to *down-right*, patterns at 300–700, 1500–5000 simulation run

Figure 12 shows data about the percentage of agents in the population which are not moving over time. For now it is just worth noting that the fact of being stopped in a localization does not mean an improvement in systemic outcomes (as the comparison settings could suggest) and, on the contrary, heuristics generating lower performances are the ones "stopping" the highest number of agents.

The sum of these evidences should suggest that simple but efficient rules of thumbs, like *H2*, at a micro level are capable of giving rise to an efficient pattern at a macro level. As compared to other heuristics, *H2* is smart and simple because it houses an optimal path between the present situation and the desired outcome. In theoretical terms, our results on the efficiency generated by *H2* at a macro level could support the well-known argument that the evolution of human societies has most likely to do with the discovery and the social diffusion of simple and smart rules of thumbs (Barkow et al. 1995; Gigerenzer and Selten 2001).

Figure 13 shows average values of *attainment discrepancy* in the case of Agents2 and Agents3, which are endowed with more sophisticated cognitive features. Results of these agents are compared to results of Agents1 exploiting all the heuristics or the best one (*H2*). It could be expected that Agents2, given their cognitive capability of switching between heuristics, obtain better results than Agents1. The evidence is that
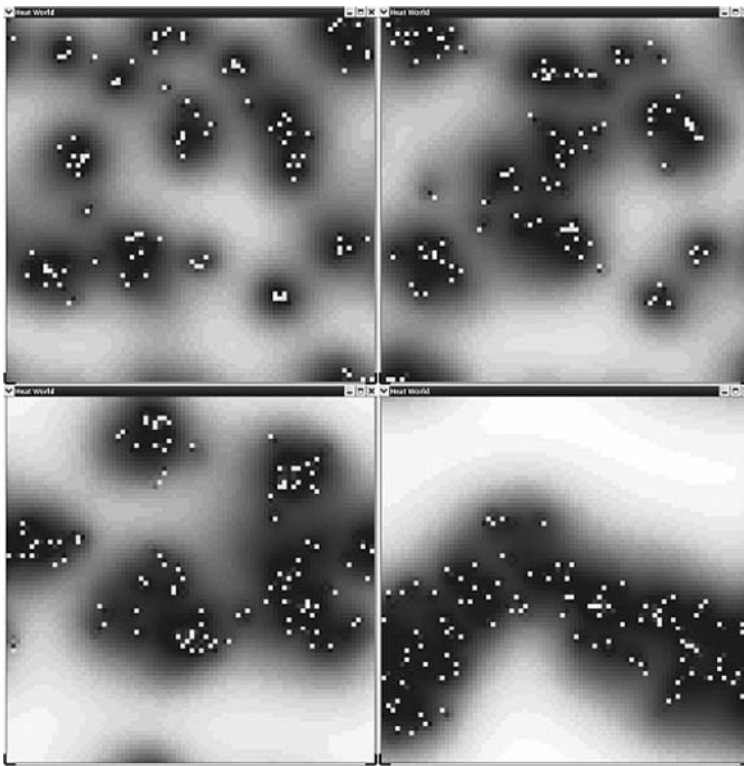
Fig. 6 Simulation setting 6 (Agents2 using heuristic toolbox). From *top-left* to *top-right*, and from *down-left* to *down-right*, agents patterns at 300–700, 1500–5000 simulation run

Agents2 obtain very similar results to Agents1 exploiting the most efficient heuristic (*H2*). Compared to Agents1 with all the heuristics, the improvement is particularly significant whenever the distribution of discrepancy in the population is taken into account. As Fig. 14 points out, the maximum level of discrepancy of Agents2 is strongly lower than the one generated by a system populated by Agents1 behaving according to all the four basic heuristics. This compared bad performance can be explained if one takes into account that Agents2 behavior is more sophisticated in cognitive and systemic terms, and that they tend to adopt more efficient heuristics. As a matter of fact, Agents2 are examples of boundedly rational social agents. They make use of a macro information (the average level of *attainment discrepancy* at a population level) as an inference to change their behavior. Essentially, they use the evidence that agents at a population level enjoy a better situation than their own one, they infer that their worst situation depends on behavior they have adopted, and change their heuristics trying to improve their individual situation. Such a behavior depends on cognitive features more sophisticated as regards to Agents1 and their simple heuristics mentioned above, since this conveys a typical feature of human agents that behave in social settings characterized by uncertainty, that is, the use of information on others (information that is often partial, imperfect, and biased) as a reference for decision
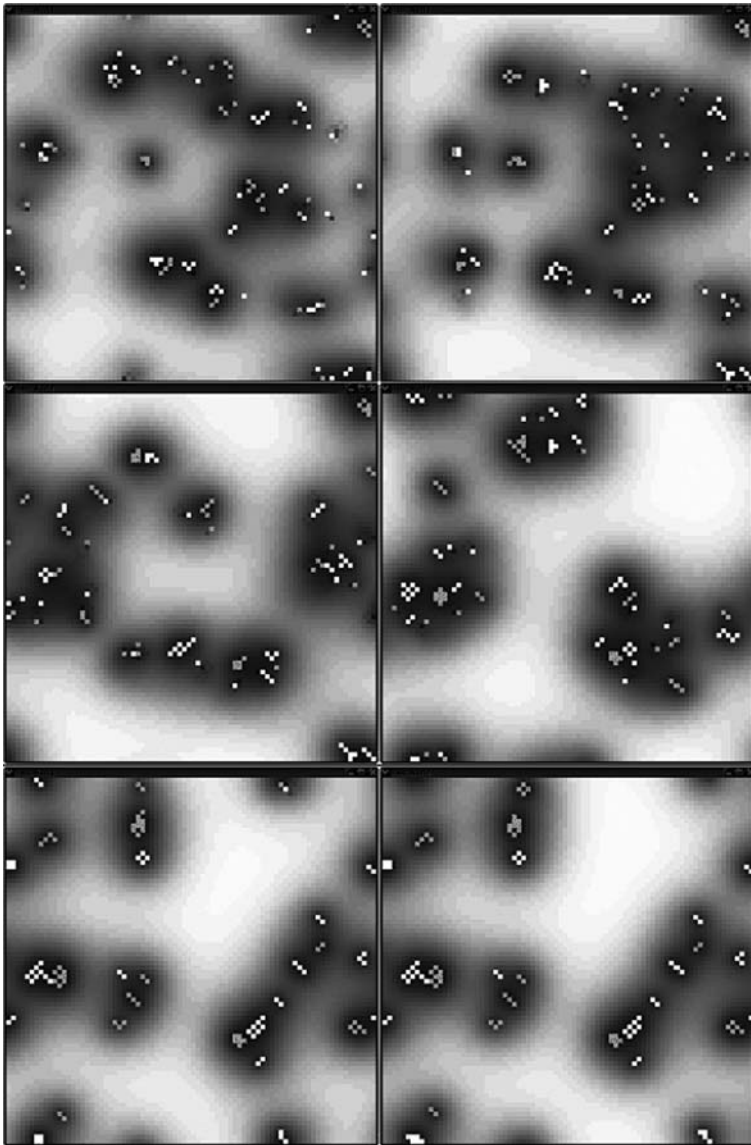
Fig. 7 Simulation setting 7 (Agents3). From *top-left* to *top-right*, from *middle-left* to *middle-right*, and from *down-left* to *down-right*, agents patterns at 300–700, 1500–5000, and 15,000–17,200 simulation run

making. As we know from the neo-institutional sociological literature, this same often holds true also among organizations embedded within organizational fields (Powell and DiMaggio 1991).

Figures 13 and 14 stress the performance of Agents3, that is to say those that are capable of recognizing tags of neighbors. The different spatial patterns and behavior give raise to lower average and maximum values of discrepancy. Figure 15 shows that
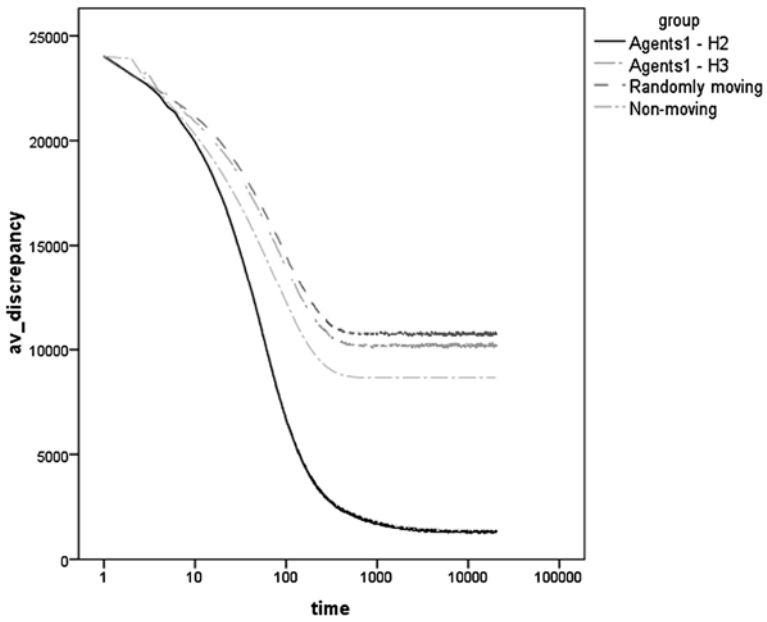
Fig. 8 Average values of *attainment discrepancy* over time: Agents1 using *H2*, *H3* and comparison settings
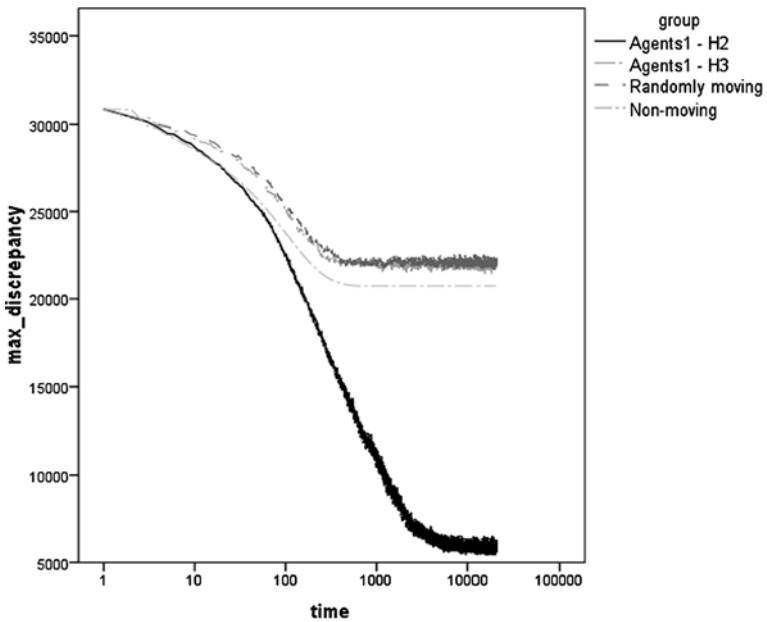

Fig. 9 Maximum values of *attainment discrepancy* over time: Agents1 using *H2*, *H3* and comparison settings
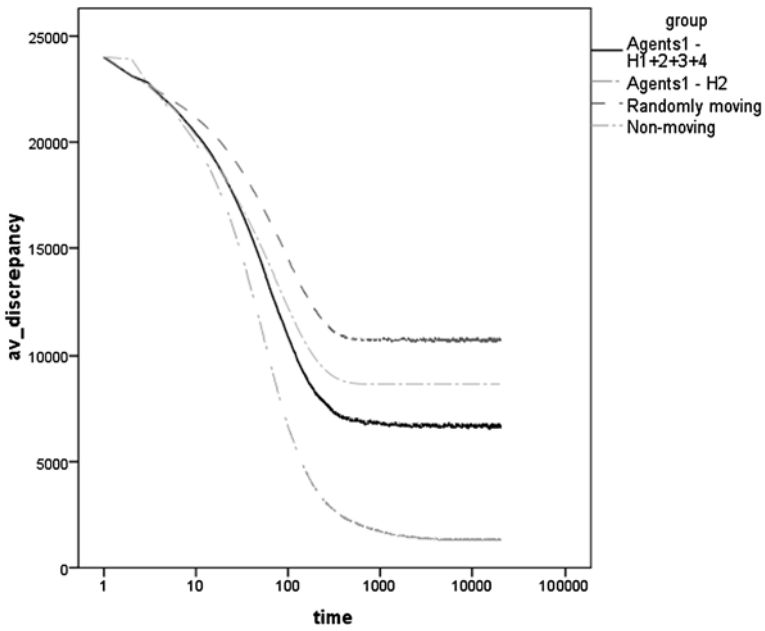
Fig. 10 Average values of *attainment discrepancy* over time: Agents1 using *H2* only, all 4 heuristics and comparison settings
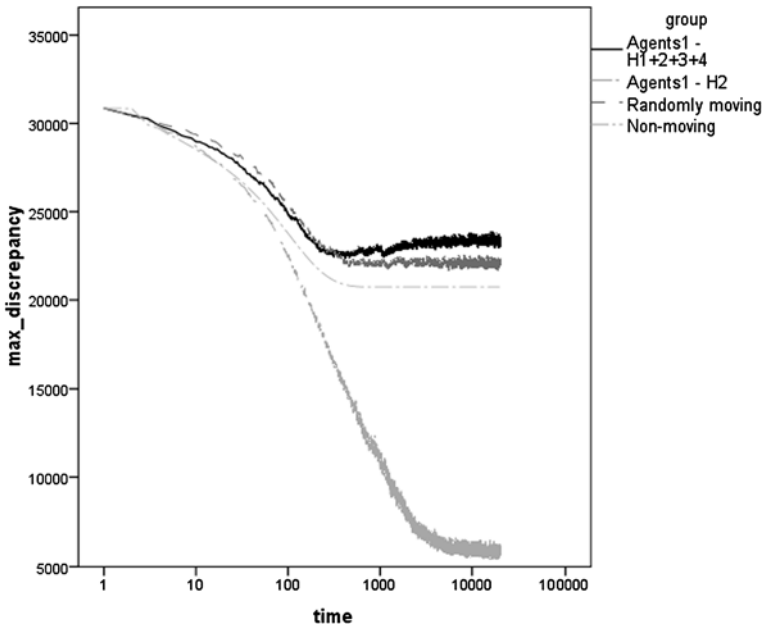

Fig. 11 Maximum values of *attainment discrepancy* over time: Agents1 using *H2* only, all 4 heuristics and comparison settings
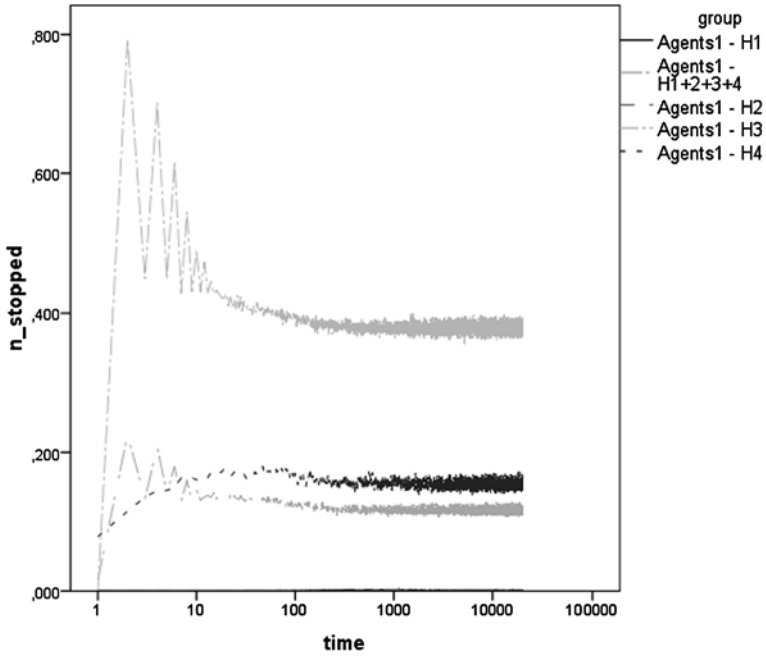
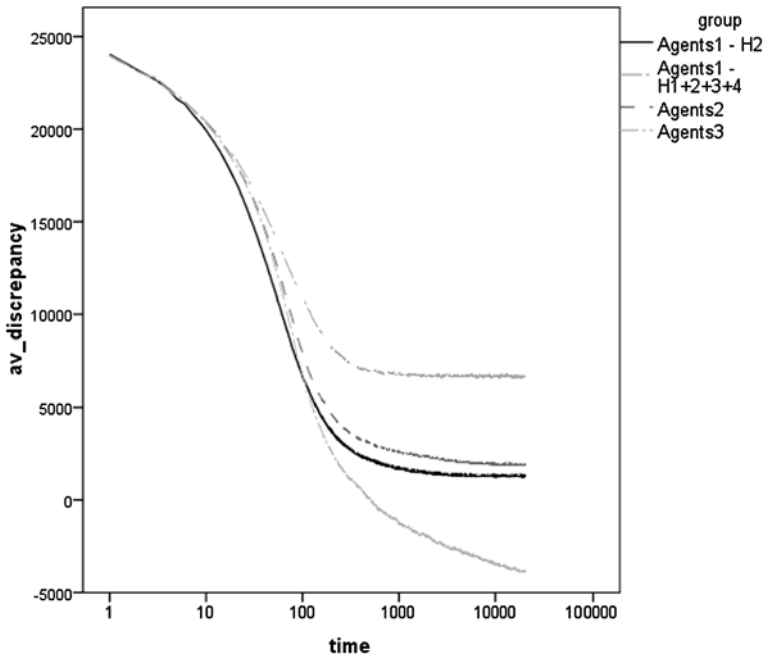**Fig. 12** Percentage of agents non moving over time: Agents1 using single heuristics and all 4 heuristics

Fig. 13 Average values of *attainment discrepancy* over time: Agents1 (using *H2* only and all 4 heuristics), Agents2 and Agents3
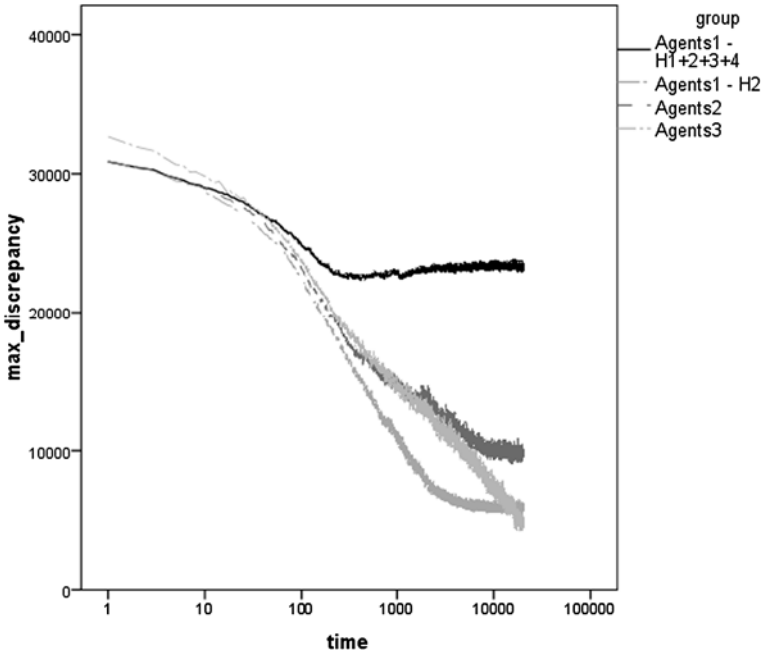
Fig. 14 Maximum values of *attainment discrepancy* over time: Agents1 (using *H2* only and all 4 heuristics), Agents2 and Agents3
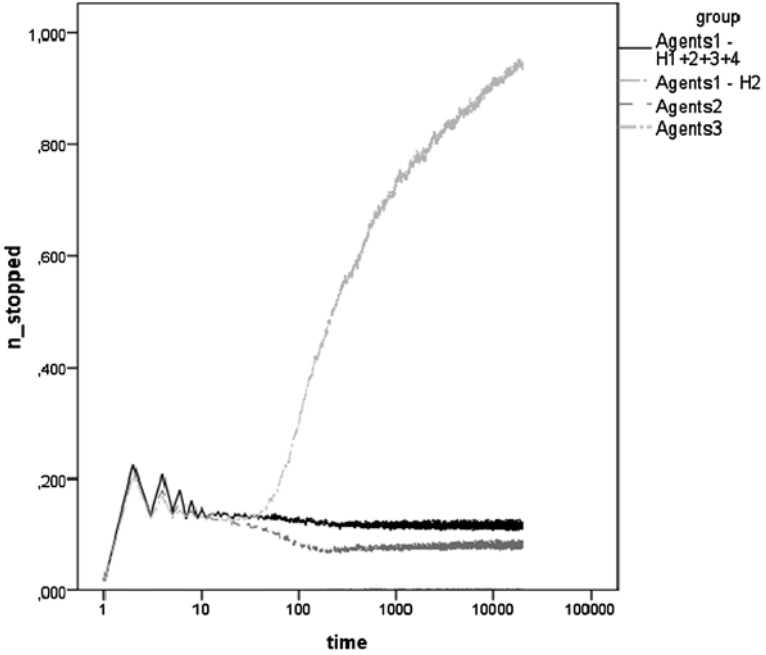

Fig. 15 Percentage of agents non moving over time: Agents1 (using *H2* only and all 4 heuristics), Agents2 and Agents3

Agents3 are the only ones for which the achievement of stable spatial distribution patterns across space goes positively with efficient data on the *attainment discrepancy*. By introducing the possibility of detecting others' features and following a group-like sense of belonging, we have further increased the sophistication of agents behavior towards second order social properties. The difference between Agents2 and 3 is that Agents3 are subjected to a macro-micro feedback which is less global and more local (the macro influence of two groups in fact works at a neighbor level), but based on a type of information that is more precise (the possibility to detect the tag of the entire neighborhood).

It is worth emphasizing that only in this simulation setting the system has been able to achieve a self-organized efficient solution to the coordination problem that characterizes the structure of the model, by matching individual efficiency and system effectiveness. Essentially, here, agents have been able to achieve their individual goals at best, giving rise, at the same time, to an equally distributed well-being at a systemic level. Our results can reinforce the evidence of the strength of tag mechanisms fully illustrated in the literature on agent-based models of cooperation, that is, the fact that things change at a macro level if agents are able to detect their social environment, by identifying other agents tags and referring action to this (Hales 2000; Riolo et al. 2001).

4 Concluding remarks

The simulation results described in the preceding paragraph show that small modifications of agents behavior towards more sophisticated cognitive features have allowed us generating very different self-organized patterns at a macro level. This evidence is consistent with other computational social science ABMs (Flache and Hedgselmann 1999; Macy 1990; Jager et al. 2000), where the macro observed outcomes are understood as the result of more or less complicated cognitive features at a micro level. It is worth noting that these examples allow appreciating a theoretical approach to micro-foundations of social agents behavior that is alternative to mainstream economics and rational choice theory, and grounded on the recent influential arguments suggested by experimental and evolutionary psychology (Barkow et al. 1995; Gigerenzer and Selten 2001). In particular, our simulation results have confirmed that smart and simple rules of thumbs are important for a coordinated and stable evolution of social systems. Moreover, results show that increasing the sophistication of agents behavior towards socio-cognitive properties means the system capability of computing macro global stability and of generating micro-macro efficiency and coordination through self-organizing decentralized patterns.

In this regard, it is worth outlining that these results have been achieved through simple and gradual modifications introduced in the behavioral architectures of our agents. As we have stressed, there is no need for assuming highly complicated cognitive architectures to bring in relevant cognitive properties that matter for computational social science purposes. The modifications that we have introduced were just on specific parts of agents behavior: from particular differences in simple heuristics, as in simulation settings 1 to 5, to the cognitive capability of figuring out a property of

the context and to use this information for changing behavior, as in simulation setting 6, up to the recognition of a group tag (the color), as in the last simulation setting.

From a modeling viewpoint, our examples should be viewed as an attempt to demonstrate how it is possible to understand the relevance of cognitive and social properties of agents behavior in simulation models without assuming and introducing right from the start very complicated cognitive architectures, difficult to evaluate, use, and understand, at least for social science purposes (Gilbert 2005; Epstein 2006). This is particularly relevant when using agent-based models, since the high degree of flexibility of these tools could tempt the social scientist in building quite complicated models without sound analytical reasons to do it (Hedström 2005).

Finally, our results show that the efficacy of a system clashes with the socio-cognitive intelligence of its agents. This can be viewed by comparing all the simulation settings with the last one. In this last one, we have introduced a second-order social property in the agents behavior, making agents more intelligent towards their social environment. This evidence suggests that the exercise of exploring and complicating agent-based foundations of simulation models can be useful in understanding some properties that are typical of social agents. The consequence of this for the debate on the peculiarity of the social science domain as regards to physical sciences is far from our humble purpose. But, the overall impression is that this is a point which computational social scientists sooner or later will be called to reply on.

References

Axelrod R (1997) The complexity of cooperation. Agent-based models of competition and collaboration. Princeton University Press, Princeton

Barkow JH, Cosmides L, Tooby J (1995) The adapted mind: evolutionary psychology and the generation of culture. Oxford University Press, London

Buchanan M (2007) The social atom. Why the rich get richer, cheaters get caught, and your neighbor usually looks like you. Bloomsbury, London

Conte R (1999) Social intelligence among autonomous agents. Comput Math Organ Theory 5(3):203–228

Conte R, Castelfranchi C (1996) Simulating multi-agent interdependencies: a two-way approach to the micro-macro link. In: Troitzsch KG, Mueller U, Gilbert N, Doran J (eds) Social science microsimulation. Springer, Berlin, pp 394–415

Conte R, Edmonds B, Moss S, Sawyer RK (2001) Sociology and social theory in agent-based social simulation: a symposium. Comput Math Organ Theory 7:183–205

Conte R, Hegselmann R, Terna P (eds) (1997) Simulating social phenomena. Springer, Berlin

Epstein JM (2006) Generative social science. Studies in agent-based computational modeling. Princeton University Press, Princeton

Flache A, Hedgselmann R (1999) Rationality vs learning in the evolution of solidarity networks: a theoretical comparison. Comput Math Organ Theory 5(2):97–127

Gigerenzer G, Selten R (2001) Rethinking rationality. In: Bounded rationality-the adaptive toolbox. MIT Press, Cambridge

Gilbert N (1996) Holism, individualism and emergent properties. An approach from the perspective of simulation. In: Hegselmann R, Mueller U, Troitzsch KG (eds) Modelling and simulation in the social sciences from the philosophy of sciences point of view. Kluwer Academic, Dordrecht, pp 1–27

Gilbert N (2002) Varieties of emergence. In: Sallach D (ed), Social agents: ecology, exchange, and evolution. Agent 2002 Conference. University of Chicago and Argonne National Laboratory, pp. 41–56

Gilbert N (2005) When does social simulation need cognitive models? In: Sun R (ed) Cognition and multi-agent interaction: from cognitive modeling to social simulation. Cambridge University Press, Cambridge, pp 428–432

Gilbert N (2008) Agent-based models. Sage Publications, London

Gilbert N, Terna P (2000) How to build and use agent-based models in social science. Mind Soc I(1):57–72

Hales D (2000) Cooperation without space or memory: tags, groups and the Prisoner's dilemma. In: Moss S, Davidsson P (eds) Multi-agent based simulation. Springer, Berlin

Hedström P (2005) Dissecting the social. On the principle of analytical sociology. Cambridge University Press, Cambridge

Jager W, Janssen MA, de Vries BJM, de Greef J, Vlek CA (2000) Behavior in commons dilemmas: homo-economicus and homo-psychologicus in an ecological-economic model. Ecol Econ 35:357–379

Macy MW (1990) Learning theory and the logic of critical mass. Am Soc Rev 55:809–826

Mezias SJ (1988) Aspiration level effects: an empirical investigation. J Econ Behav Organ 10:389–400

Murphy PR, Mezias S, Chen YR (2001) Adapting aspirations to feedback: the role of success and failure. In: Lant TK, Shapira Z (eds) Organizational cognition. Computation and interpretation. Lawrence Erlbaum Associates, London, pp 125–146

Powell WW, DiMaggio P (eds) (1991) The new institutionalism in organizational analysis. University of Chicago Press, Chicago

Prietula MJ, Carley KM, Gasser L (eds) (1998) Simulating organizations. Computational models of institutions and organizations. MIT Press, Cambridge

Riolo RL, Cohen MD, Axelrod R (2001) Evolution of cooperation without reciprocity. Nature 414:441–443

Sawyer RK (2005) Social emergence. Societies as complex systems. Cambridge University Press, Cambridge

Sun R (2001) Cognitive science meets multi-agent systems: a prolegomenon. Philos Psychol 14(1):5–28

Riccardo Boero received a First-Class Honours Degree in Economics (University of Torino, Italy) in 2000, a PhD in Economics in 2003 (University of Pavia, Italy) and in Sociology in 2007 (University of Surrey, UK). His field of research is computational and experimental economics.

Marco Castellani received a Degree in Business Administration in 1998 and a PhD in Economic Sociology in 2003 (University of Brescia, Italy). He is research fellow in Economic Sociology at University of Brescia. His field of research is the understanding of cognitive processes in social phenomena.

Flaminio Squazzoni is Assistant Professor of Economic Sociology (University of Brescia). His main field of research is social simulation. He is Review Editor of JASSS.