Arianna Dal Forno and Ugo Merlone (2002)

A multi-agent simulation platform for modeling perfectly rational and bounded-rational agents in organizations

Journal of Artificial Societies and Social Simulation vol. 5, no. 2
<http://jasss.soc.surrey.ac.uk/5/2/3.html>

To cite articles published in the *Journal of Artificial Societies and Social Simulation,* please reference the above information and include paragraph numbers if necessary

Received: 23-Dec-2001 Accepted: 20-Feb-2002 Published: 31-Mar-2002

Abstract

This paper presents an agent-based simulation framework for the analysis of the equilibria that emerge in a complex structure such as an organization; we can think of some of these equilibria as *corporate culture*. We concentrate on modeling the effort exerted by heterogeneous agents in an organization, and how the interaction between them may lead to a common level of effort (corporate culture). The simple model we propose is a system in which agents interact in a dynamic, adaptive and evolving way. Such a model encompasses many of the peculiarities which make organization modeling a hard task, because it would involve a difficult mathematical problem with solution highly sensitive to parameters. The computational approach, by contrast, allows us to overcome some of these difficulties and to consider easily both perfectly rational and bounded-rational agents; in this way we are able to study the interactions between different types of agents and interpret them in the relevant economic frame. Consequently we can observe how different compositions of the population may lead the system to different common behaviours; the implications of our findings are both descriptive and normative, and shed light on some core problems of the economics of organization design.

Keywords:
Bounded Rationality; Corporate Culture; Multi-agent System

Introduction

1.1

The problem of incentives and compensation is cardinal in modern economic literature; many papers address these problems considering moral hazard in agency relationships. Other papers address the problems of optimal form of hierarchy in firms and organizations and give interesting results (for a survey see e.g. Macho-Stadler and Pérez-Castrillo 1997 and Radner 1992 and for different approaches in modeling bounded rationality see Rubinstein 1998). The approach we use is different; we provide a very simple model of organization and do not consider explicitly hierarchy. The model of organization we introduce may be interpreted as an Interaction Game (e.g. Morris 1997) with agents playing bounded rationality strategies. This approach is useful to shed light on aspects that are usually neglected such as the interpretation of equilibria as corporate cultures.

1.2

In the economic studies about firms and organizations, corporate culture has been almost ignored. Nevertheless there are a few important exceptions such as Kreps (1990), where a game theoretic model is suggested to analyze corporate culture. For a survey and an analysis of the economic literature about corporate culture the reader may refer to Hermalin (2001). Different is the adaptive behaviour approach Young (1993, 2001) proposes. In his model, players are not fully rational and adapt their play to their past experiences. In his results the convention which eventually emerges depends also on the evolution path. The emergence of corporate culture when a population consists of heterogeneous agents is studied in Dal Forno and Merlone (2001). A single shot game played by different types of agents is considered and some results about equilibria arising when not all agents are fully rational are provided. By contrast there is a rich literature on modeling organizational culture using a computational approach, for a survey the reader may refer to Carley and Prietula (1994).

1.3

Steers and Black (1994) define corporate culture as the basic assumptions and beliefs that are shared by the members of a group or organization and that are used as a norm. The organization's problem is to identify a rule that allows relatively efficient transactions to take place and devise some way to communicate that rule to all current and potential trading partners. The organization has an interest in preserving and promoting a sufficiently high level of effort provided by employees. This may imply establishing a sort of good reputation to allow for future beneficial transactions; nevertheless it is quite common to observe situations where there is no interest in reputation and members' main interest is free riding.

1.4

According to Kreps (1990) corporate culture may play the role to establish general principles that should be applied to select a particular equilibrium when problems of multiplicity occur. Culture is a complex phenomenon and consists of many different facets. The literature has focused on different aspects; Harrison and Carrol (1991) consider a model of cultural transmission. They studied cultural socialization and found that cultural systems are robust even in the presence of system turnover. Carley (1991) develops a dynamic simulation model of the interaction shared knowledge cycle in which individuals interact, communicate and adapt to new information. Her model allowed the study of group stability and endurance in different societies where no change in membership and no new ideas are considered.

1.5

In our model we do not consider any particular task. In contrast we focus on the effort exerted by organization members since any activity is (usually) affected by the effort exerted by the employees. In this sense our approach transcends other models even if, for the same reason, it may be more abstract. The organizational structure we consider may be compared to the organized anarchy studied in Cohen, March and Olsen (1972).

Nevertheless some important differences arise: in the garbage can model a decision-making task is considered, effort is exogenously determined, and homogeneous agents are considered. In our model we do not focus on any particular task and, most important, effort is determined endogenously by mutual agent interaction. Finally, Cohen, March and Olsen used their model to simulate and study different organizational structures. In our model, in contrast, we study the impact of heterogeneous agents in a fixed organizational structure. In our modeling we are interested more in the process of choice rather than the results of rational choices by agents (Simon 1982). In this sense we consider different types of bounded rationality agents. When considering different types of agents it may become very difficult to predict the evolution of the whole system using the mathematical approach. In fact, even if it is relatively simple to provide a mathematical model of single types of agents, the whole system depends on too many parameters.

1.6

In this paper we describe the multi-agent simulation platform we developed to simulate complex interactions between different types of agents in a dynamic setting. The simulation approach allows us to overcome some of the difficulties mentioned above and easily consider both perfectly rational and bounded-rational agents. This approach may help in studying how the micro behaviours of different agents may evolve into a general pattern of behaviour of the whole system. In particular we are interested in studying the dynamical evolution of the effort exerted by members of an organization. The paper is organized as follows. First we present the theoretical model we consider. Second we describe the simulation platform and discuss the type of agents we implement. Finally we illustrate and discuss some of the results we obtained with our simulations providing their economic interpretation.

The Theoretical Model

2.1

The model of organization we propose is very simple. The organization consists merely of a set of agents. They may perform their task only through cooperation with some other agent. When two agents meet and form a team, they perform their tasks. Both agents in a team observe each others' effort and adapt their future behaviour to the outcome of their cooperation. For an economic analysis of teams see Marschak and Radner (1972).

2.2

The interpretation is simple. Consider an employee in an organization performing different tasks with different people. Assume that in performing his tasks he may observe the effort exerted by his different partners and decide on his effort as a function of his past experiences. Obviously, if he meets only shirking partners he will have a strong incentive to shirk too. By contrast, should he interact only with people exerting high effort he may have both a (moral) incentive to match their effort and an incentive to free ride. We are interested in situations where agents adapt their behaviour by comparing their private profit and effort to their partners'.

2.3

We consider the different equilibria of this simple model. According to Hermalin (2001) a possible interpretation of Kreps (1990) is that corporate culture is a way of ensuring coordination in some kind of games (for details the reader may refer to Hermalin (2001:229), or Kreps (1990)). Therefore we are interested in equilibria where effort exerted by agents is somehow coordinated. According to our interpretation such equilibria correspond to the emergence of corporate culture.

2.4

In Kreps (1990) the organization is considered as a single decision-making entity. The entity's problem is to identify a rule on which a viable reputation can be based and to communicate that rule to current and potential future trading partners. While Kreps (1990) considers corporate culture as a rule relative to external relations, we consider it as a rule shared by the members of the same organization and considered important. For a definition of corporate culture see Steers and Black (1994).

2.5

The approach we follow departs from the game-theoretical one since we consider (especially) agents that are not rational in the sense required by game theory. For a definition of rationality see Meyerson (1991).

2.6

As is common, we assume that output is increasing with effort and the marginal output of effort is decreasing. The utility of output is shared among the members of the team[1]. On the other hand greater effort means greater disutility to the agent; we also assume that the marginal disutility of effort is increasing.

2.7

To perform our simulations we used the following profit function[2]:

$$\pi_i = 5\sqrt{e_i + e_j} - e_i^2$$

where e_i is agent i's effort. This formulation satisfies the common assumptions we stated.

2.8

This may be also interpreted as a game. For details about the one-shot game when no local interactions are considered see Dal Forno and Merlone (2001).

2.9

In the certainty case, if the only verifiable variable is final profit, each agent may infer its partner's effort and profit. Vice versa, when the verifiable variable is partner's effort, each agent may use it to compute its profit. Depending on its type the agent may take into account this information to adapt its future effort. In the uncertainty case, by contrast, we assume that agents may both overestimate and underestimate their partner's efforts and profit.

The Simulation Platform

3.1

We developed a multi-agent simulation platform based on the programming environment C++ Builder 5. Our platform is intended to provide a simple way to implement different agent behaviours and observe their dynamical interactions in the model we presented above[3]. This platform consists of the following main components:

1. a window with main controls of the simulation
2. a window displaying agents' location
3. a window displaying agents' efforts

3.2

The user may edit the parameters of the model modifying their values in the main window. Parameters may be summarized as:

a. World parameters: the dimensions of the world
b. Agent parameters:
    - Number of implemented types. Up to six different types of agents may coexist at the same time. The different types are detailed in 5.4 - 5.11. A coloured label indicates the different kinds of agent appearing in the simulation; agents of the same type are displayed with the same colour
    - Number of agents for each desired type
    - The class of behaviour related to each type of agent. This is referenced as a number and when the cursor is positioned on the coloured label a short description of the type assigned is displayed
c. View parameters:
    - Resolution used to display output windows: the user may choose to display each agent either as a pixel or a small square
    - Colour display mode: effort may be displayed using either a graded red or a chromatic scale. The user may switch scale at any time during the simulation
    - Average effort: a small indicator moving below the graded scale displays population average effort in real time.
d. Simulation parameters:
    - Lag time between each iteration
    - Noise: agents effort may be affected by noise
    - The percentage of noise may vary from ±1% to ±50%

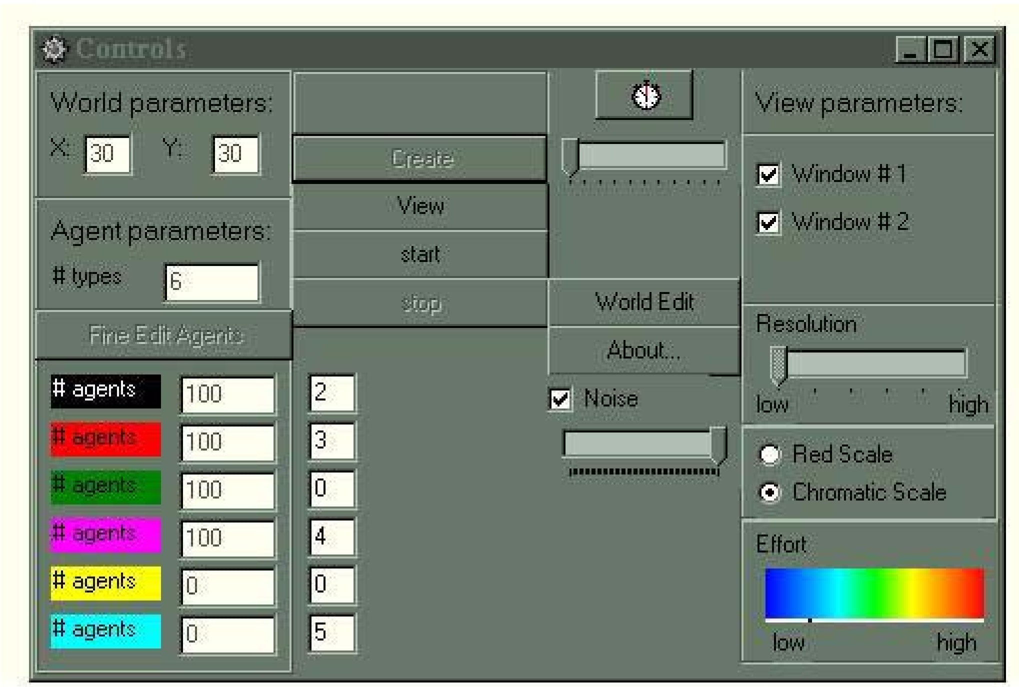

Figure 1. The main window with controls for the simulation parameters

3.3

Output windows display the simulation. Currently two windows are implemented.

1. An agents window, providing a spatial visualization of the agents in the simulation. Agents' colours match the agent type colour on the main window;

this way it is immediately possible to understand the agent type location. In this window colours do not change since the agents' type remains the same over the simulation. This window is particularly useful when agents are allowed to move.

2. An efforts window, providing the corresponding effort exerted by every agent at each period. Agent effort is displayed according a graded scale, where red represents high effort while blue represents low effort. This window depicts the evolution of the system and colours change according to the effort agents exert.

In both windows empty grid locations are displayed using white.

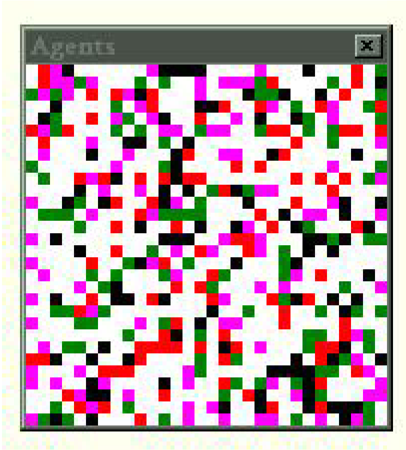

Figure 2. The toroidal grid with agents (different colours representing different types and behaviours of agents)

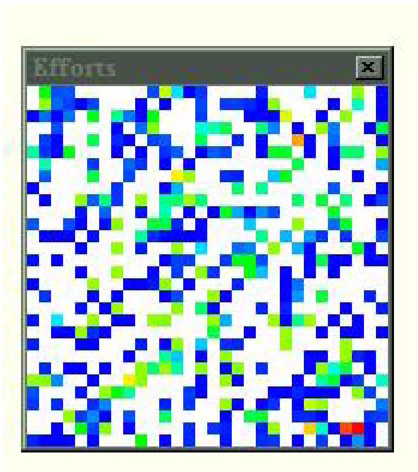

Figure 3. Efforts window with agents' effort displayed with a chromatic scale

General Structure of the Simulation

4.1

Agents move and operate on a toroidal grid. At each turn each agent performs two phases:

1. *movement*: determines a random direction (North, South, East and West) and if the corresponding adjacent position is not occupied moves to it.

2. *interaction*: if two adjacent agents are facing each other they form a team, play the game and according to their type adapt their future effort.

4.2

The general structure of the simulation can be sketched as follows.

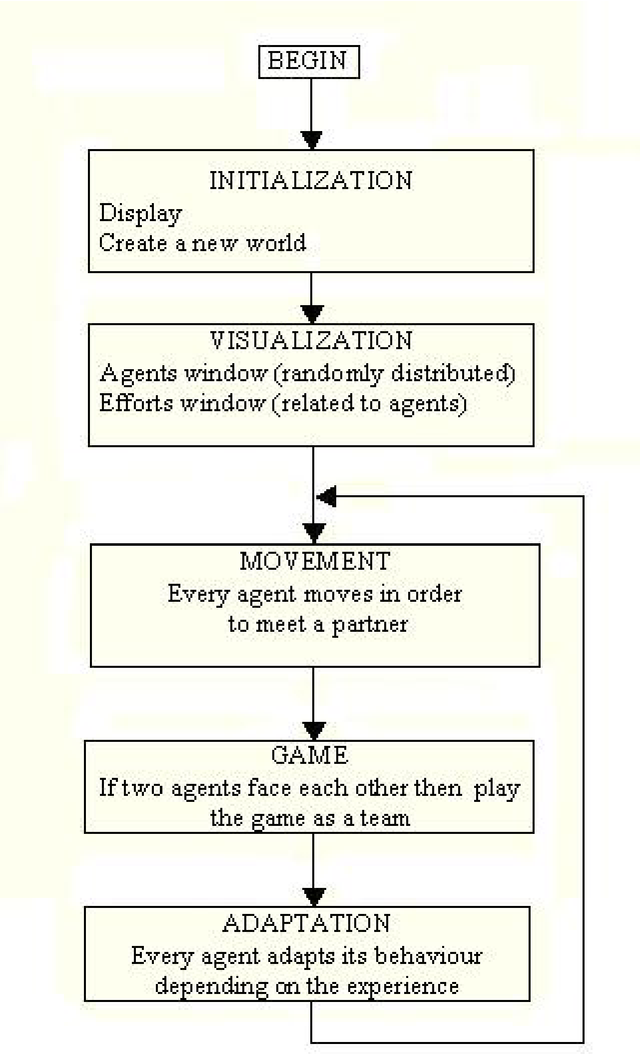

Figure 4. General structure of the simulation

Initially, according to world parameters (chosen by user) agents are randomly distributed on the grid. Each agent's initial effort is determined accordingly to their type; for most types of agents initial effort is determined randomly. Agents' position and effort may be observed on the relative display window. At every round each agent moves randomly in order to meet a partner and form a team. The only constraint on possible movement is that no more than one agent may stay in a cell of the grid.

4.3

It must be noted that, while in Carley (1991) interaction is determined by relative similarity between individuals (motivation), in our model interaction is mostly random[4]. This is an important point from the organizational point of view. In an organization the principal would prefer to incentivise productive interactions even if this does not necessarily mean those between the same type of agents.

4.4

After the movement step, each agent faces a particular direction: if two agents in adjacent cells face each other, this couple forms a team. Both agents exert their (non negative) effort following their particular behavioural rule. Agents can observe the profit as a result of their joint effort and estimate the effort played by the partner. Then both may modify their behaviour according to their type and then move again in order to meet another partner.

The Agent Class

5.1

Agents interact in the simple environment we have described. They are modeled by defining a class which considers all the relevant variables. Obviously, it is possible to introduce new types of agents with more complex behaviours. Furthermore it is possible to add new variables to this class.

```
class Agent {
  public:
    Agent(int _type = 0 );
    ~Agent();
    void  RandAgent(int* grid, int modx, int mody, int
num);
    void  RandMove(int* grid, int modx, int mody);
    void  WorkAgent(float parteff);
    void  RandDir();

  protected:
  private:
   int type;                                       // type of
agent
   int posx;                                       // x position
on the grid
   int posy;                                       // y
position on  the grid
   int dir;                                        // direction
N=1, E=2, S=3, W=4, No Interact=0
   int step;                                       // allowed
motion
   int col;                                        // effort
colour
   int colt;                                       // type
colour
   int evol;                                       // the type
of evolution the agent has reached
   int numinc;                                     // number of
interactions;
   double effort;                                  // current
effort
   double profit;                                  // current
profit
   double cumprof;                                 // cumulated
profit
```

```
   double leffort;                              // last
effort provided by agent
   double cumeffort;                            // cumulated
partner effort
   double lprofit;                              // last
profit
   double aeffort;                              // last
antagonist effort
   double aprofit;                              // last
antagonist profit
   double neffort;                              //
neighbourhood average effort
   double nprofit;                              //
neighbourhood average profit
};
```

5.2

Some of the variables listed above are not currently used in the implemented types. In particular it is possible to implement agents capable of switching strategies using the variable `evol`.

5.3

The public member functions implemented in this class are:

- RandAgent: performs the random positioning of the agents on the grid.
- RandDir: randomly assigns the direction the agent faces; this function is important in determining the teams' composition and the agents' movement.
- WorkAgent: determines the effort and the reaction of agent whenever it meets another agent; this function implements the behaviour of all the types of agents.

5.4

The class is very flexible and, in particular, it is open to modeling new types of agents. At the moment the following types of agents have been implemented:

1. *null effort*: this agent always exerts the same almost null effort
2. *shrinking effort*: this agent halves the effort provided by its last partner
3. *replicator*: this agent exerts the same effort its last partner exerted in the previous interaction
4. *rational*: this agent exerts the best reply for its last partner effort
5. *profit comparator*: this agent compares its profit to its last partner's one; it increases its effort if it gave a higher profit
6. *high effort*: this agent always exerts the same high effort
7. *average rational*: this agent exerts the best reply to the average effort of its partners
8. *winner imitator*: this agent starts with high effort but copies its partner's effort when this one proves to yield a higher profit
9. *effort comparator*: this agent compares its effort to its last partner's one; it increases its effort if it is inferior to its partner's one and vice versa
10. *averager*: it averages its effort with its last partner's effort

5.5

Null effort [5] and *high effort* are fixed effort agents; they do not adapt their behaviour to their past experience. They do not consider profit, and stick to a norm. For example, we can think of high-effort agents as if they were driven by a moral norm. Obviously, a world consisting of only high-effort agents would be optimal from the principal point of view. Furthermore we found that introducing high effort agents in otherwise homogeneous populations may improve overall productivity[6].

5.6

Rational and *average rational* agents follow the game theoretic paradigm of best reply function. Obviously, given the profit function we consider, a population consisting only of such agents converges to the Nash equilibrium of the game (for mathematical details see Dal Forno and Merlone 2001). *Replicator* agents play a tit for tat strategy; a homogeneous population of replicator agents does not converge to any stable equilibrium. In contrast, if a fixed effort agent is introduced into such a population the whole population will end up playing the same effort. *Shrinking effort* agents lower their effort and tend to free ride their partner. A homogeneous population consisting of these agents would converge to play the null effort equilibrium.

5.7

In the literature different varieties of learning behaviour have been discussed. In the natural selection approach, a reproductive advantage is given to agents with high-payoff. Even if our approach is not evolutionary, we expect that some agents may imitate high-payoff agents. We introduced *winner imitator* agents to consider this kind of behaviour. We consider different types of agents that adapt their efforts comparing with their neighbour's. *Profit comparator* agents increase their effort by 10% if their profit is higher than their partner's otherwise they decrease it by 10%. *Effort comparator* agents are similar but compare efforts instead of profits. In both cases a homogeneous population reaches the null effort equilibrium. Finally *averager* agents adapt to their neighbours more softly. They consider their last partner's effort and average it with their own.

5.8

With exception of fixed effort and average rational agents, the agents implemented so far are Markovian: they update their behaviour considering only the last interaction outcome. Furthermore most of them are bounded rationality agents since they use different simplifications in making their choices (Simon 1955). Consequently some of the agent rules may be interpreted as rational in Simon's sense of procedural rationality (Simon 1982).

5.9

The different types provide sufficient variability to have a complex system arising from simple individual behaviours.

Noise

5.10

We considered the certainty case: each agent can perfectly observe its partner's effort and profit. In the uncertainty case, by contrast, we assume that agents may either overestimate or underestimate their partner's effort and profit. Each agent observes his partner's perturbed effort while perceiving no perturbation of its own effort. We assume that the partner's effort may be underestimated and overestimated with the same probability. We model the uncertainty case introducing noise. Noise acts as a multiplicative factor of agents' effort. When set at the minimum this factor may vary in the range 0.99-1.01, while when set at the maximum it may vary in the range 0.50-1.50.

Simulation Experiments

6.1

Some of types studied in Dal Forno and Merlone (2001), and many other types have been implemented in this platform. The simulation allowed us to study the dynamic evolution of population consisting of different types of agent. Many different simulations are possible. In the following we mention some of the most interesting according to their economic interpretation. In particular we are interested in examples where a corporate culture emerges.

Corporate culture

6.2

Some types of agents allow us to observe the emergence of equilibria where all agents exert the same effort. We interpret these equilibria as corporate culture. Some simple examples are:

- a population consisting only of *rational* agents converges to the Nash equilibrium of the game
- a population consisting of only *shrinking effort* agents converges to a corporate culture where all the agents play the null effort
- a population consisting in all *replicators* except one fixed effort agent converges to a corporate culture where all the agents play the effort exerted by the fixed effort agent
- a population consisting in all *replicators* except one *rational* agent converges to the Nash equilibrium of the game

6.3

The last two results are important because they show that, in our model, some populations may be highly sensitive to turnover. Consider, for instance, a replicator population where a fixed effort agent is subject to turnover. The whole population would be led by the effort provided by the new agent. This is different from Harrison and Carrol (1991) results.

6.4

When considering heterogeneous populations the emergence of corporate culture may be difficult to observe. In some cases the population does not converge to any equilibrium with the same effort level. In other cases the effort the agents exert depends on their neighbourhood. An example of this situation is the following. Consider a population consisting of effort comparator (in black) and fixed high effort agents (in red). It may be observed that the effort of the effort comparator agents is higher where there is a high concentration of fixed effort agents. In this case, obviously, there is no emergence of corporate culture, as we can see in the following figure.

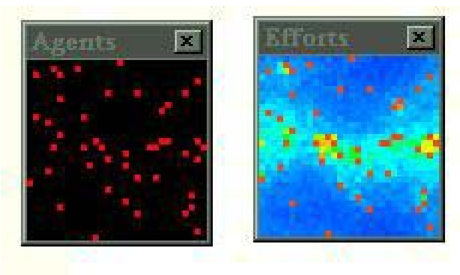

Figure 5. Local effects prevent the emergence of a corporate culture

6.5

More generally we found that when some fixed high effort agents are scattered in a homogeneous population they may boost the productivity of other agents.

6.6

Let us consider the average effort in a population consisting only of homogeneous agents. If we introduce high effort agents, the average effort of the resulting heterogeneous population depends on the proportion of high effort agents introduced.

6.7

If the homogeneous population consists of null effort agents the result is rather predictable: the average effort is simply a linear combination of the null effort and the high effort, weighted by the relative proportions of agents, as shown in the following table:

Table 1: Average effort introducing high effort agents in a population of null effort agents

% *High Effort*	0.0%	0.6%	5.6%	33.3%	66.7%	100.0%
Expected Effort	0.00010	0.01122	0.11126	0.66707	1.33403	2.00100
Observed Effort	0.00010	0.01092	0.12283	0.67948	1.32605	2.00100

6.8

Results are different when the initial population consists of the other different type of agents we implemented. Some results are presented in the following table:

Table 2: Average effort introducing high effort agents in a population of shrinking effort agents

% *High Effort*	0.0%	0.6%	5.6%	33.3%	66.7%	100.0%
Expected Effort	0.00010	0.01122	0.11126	0.66707	1.33403	2.00100
Observed Effort	0.00010	0.01235	0.20056	0.98380	1.58420	2.00100

6.9

We do not report all the our simulation data, because they are quite similar; nevertheless Figure 6 summarizes results relative to all bounded-rationality types. It may be observed that the presence of high effort agents boosts the average effort.

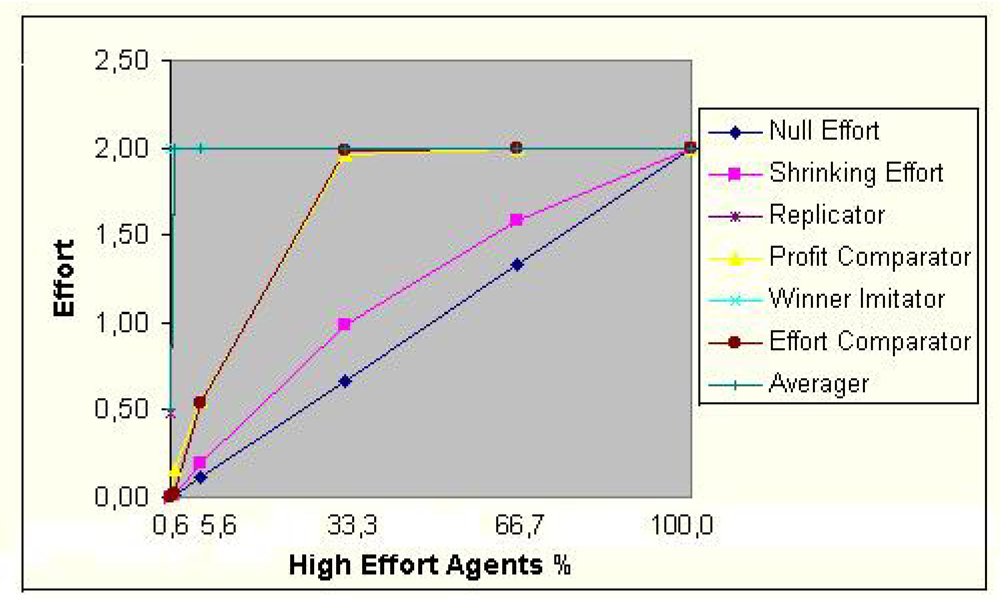

Figure 6. The effect of high effort agents on bounded rationality agents

6.10

The situation is different when considering rational or average rational agents. The average observed effort is boosted only when a small proportion of fixed agents is introduced, while for greater proportions rational agents reduce their effort. This result is reported in the following table.

Table 3: Average effort introducing high effort agents into a population of rational agents

% High Effort	0.0%	0.6%	5.6%	33.3%	66.7%	100.0%
Expected Effort	0.92101	0.92701	0.98101	1.28100	1.64100	2.00100
Observed Effort	0.92101	0.92823	0.97348	1.25872	1.59225	2.00100

Figure 7 summarizes results relative to 'rational' agents. Furthermore it must be noted that our simulations are consistent with the theoretical results we found in Dal Forno and Merlone (2001).

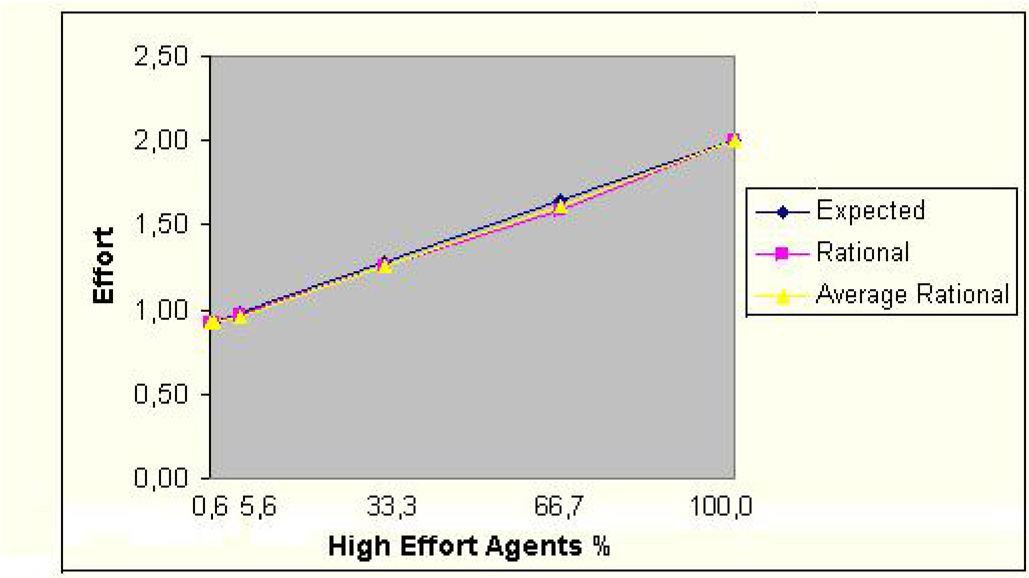

Figure 7. The effect of high effort agents on rational agents

6.11

A first consequence is that interaction between different types of agent may be productive from the principal's point of view. The insight is that it may be rational for the principal to monitor just few agents giving them a strong incentive to work hard. This may help solving the problem mentioned in note 4.

Noise effects in the simulation

6.12

As we said above, agents' effort may be affected by noise. As can be expected the presence of noise does not help to improve the global effort of the organization. A particular interesting result is the following.

6.13

Consider a population consisting only of winner imitators. In absence of noise all the agents keep exerting the same high effort they started with. With noise, the whole population degenerates to an equilibrium where all agents exert low effort.

6.14

Observing perturbed efforts may lead agents to believe their partner obtained a higher profit by shirking. Since they tend to imitate people with high profit they lower their efforts. In the long run misunderstandings are detrimental to the system and overall effort degenerates to the null effort. The following figure illustrates how noise may lower the organization global effort.

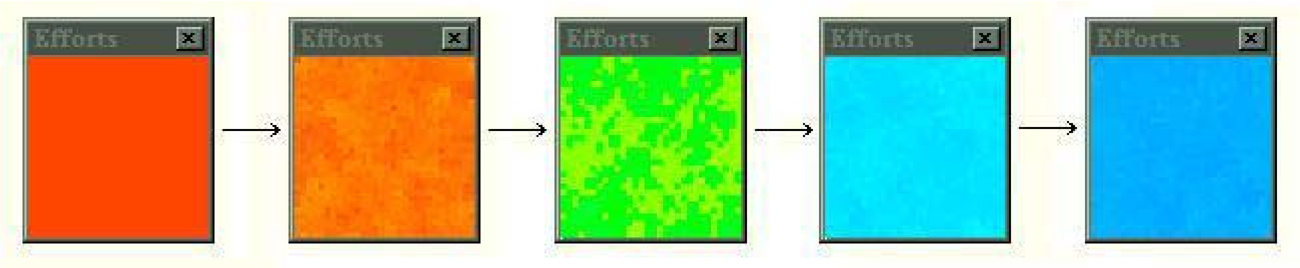

Figure 8. Effects of noise on a simple population

6.15

In our simulations we observed that noise dramatically affected the average effort of the whole system. Consider a homogeneous population of winner imitators. In presence of noise agents tend to lower their effort over time. The rate of lowering depends on the noise level: the higher the level the higher the decreasing rate. Our results are summarized in the following figure:

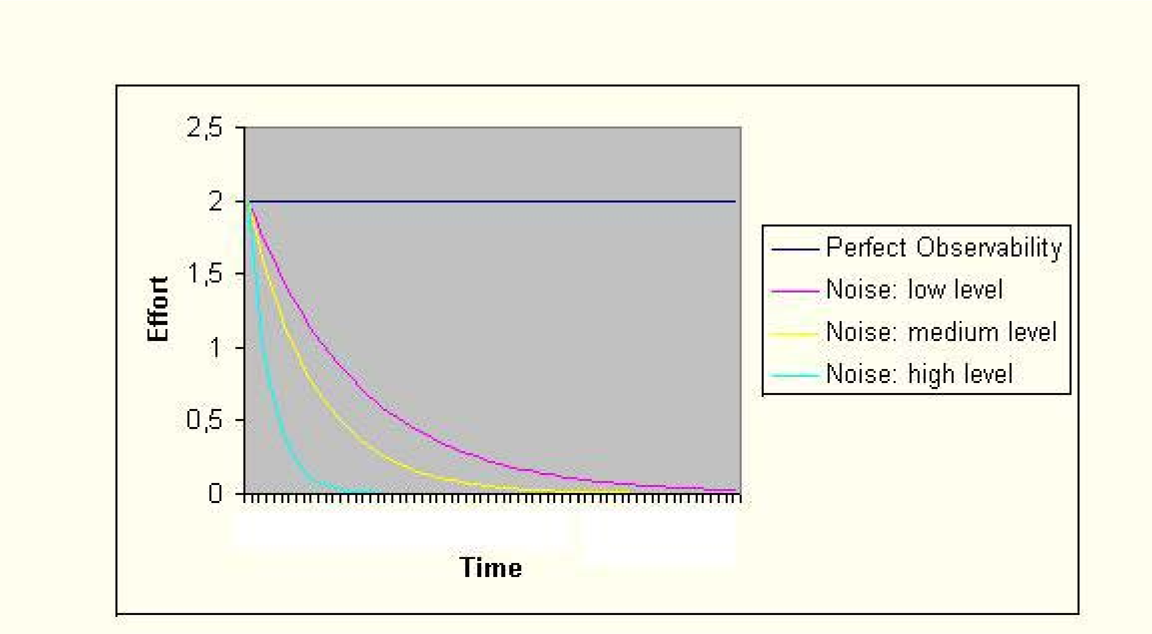

Figure 9. Noise effects on a winner imitators population

Theoretical Results

6.16

Finally our simulation platform was useful in providing a way to observe some theoretical results and to suggest further directions of research. For example in Dal Forno and Merlone (2001), it was proved that, when a rational agent and a fixed effort agent are paired, the higher is the provided fixed effort, the lower is the optimal effort of the rational agent. It is possible to replicate the same result by simulation. Consider a low and a high fixed effort agent together with a homogeneous population consisting of rational agents. The population effort converges to the Nash equilibrium, while it is possible to observe that rational agents paired with the low effort agent work harder and those paired with the high effort agent shirk.

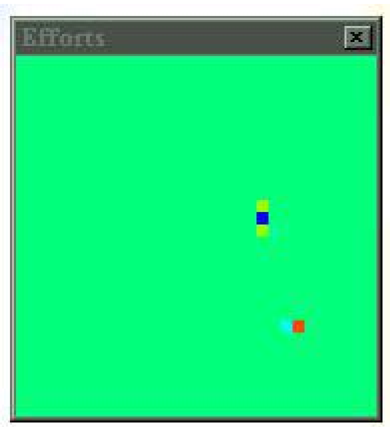

Figure 10. Rational agents shirk when interacting with high effort agents

Mechanistic vs. Organic Structures

6.17

When no grid location is empty, agents cannot move and they always interact with the same neighbours. This situation may be interpreted as a Mechanistic System of Organization where tasks tend to be rigidly defined. By contrast, when some locations are empty, agents are free to move and this may be interpreted as a more Organic System of Organization where tasks are continually adjusted and redefined through interaction of organizational members (for a comparison of Mechanistic and Organic Systems of Organization see Steers and Black 1994:367-369). In this case agents follow a zero mean random walk. They tend, in the long run, to remain in the same place and interact with same mates. As a consequence relations tend to remain almost fixed even if they may change slightly over time. Obviously, when agents are free to move local effects are moderated and zones where high effort is exerted are less stable.

Conclusion

7.1

The simulation platform we have designed provides a simple way to study the complex interactions between different types of agents. This way we could perform some simulations and replicate some results previously proved only theoretically. We were interested to study the equilibria of such a system. Our attention was focused on those where the effort exerted by agents converged. According to our interpretation such equilibria may be considered as corporate culture.

7.2

We could observe some simple situations where corporate culture could emerge. In contrast in some other examples it was not possible to observe such an emergence. It would be extremely interesting to find conditions under which corporate culture may or may not occur. Nevertheless agents heterogeneity makes it a difficult task. Further research should deal with this problem both by simulation and theoretically.

7.3

Many different types of agents have been introduced. Our first simulations considered the evolution of homogeneous populations. We perturbed these situations in different ways; for example we introduced a few fixed effort agents to study their effect on the population. We found that usually some fixed high effort agents scattered in a homogeneous population may boost productivity of the other agents. We could observe a sort of evolution in our agents even if it was not explicitly considered.

7.4

Furthermore we studied how noise could affect the behaviour of this system and its equilibria. The results were interesting since we observed that under particular circumstances noise could make population effort totally collapse.

7.5

Both these results are interesting. The first suggests a possible way to improve overall performance: the principal may decide to monitor just few agents giving them a strong incentive to work hard. The second, by contrast, suggests that in the presence of noise, the situation may degenerate. The implication is that, as in real world, great effort should be made to keep noise effects under control.

7.6

In further research it would be interesting to introduce new types of agents such as those that explicitly switched strategies and to implement the possibility to move agents on the grid. It would be also interesting to consider a "garbage can model" where efforts change dynamically. Furthermore our platform is open to model different types of interactions; it would be interesting to move beyond "random interaction" and consider interaction depending on relative similarity (Carley 1991). Finally, it could be interesting to extend our approach to models of organization with a hierarchical structure.

Acknowledgements

The authors thank two anonymous referees for helpful comments and suggestions.

Notes

[1] Team members may share only a part of the output. It must be noted that the output each agent receives depends on the aggregate effort.

[2] Obviously, it is possible to implement a different profit function.

[3] The .exe program and the agent class code is available upon request from the authors.

[4] The reader will note that in our model interaction depends on the location of agents. The principal may mitigate the random interaction effects by an optimal location of agents. But this cannot be done without knowing the agents' types: a vicious circle arises.

[5] For numerical reasons null effort agents actually exert an almost null effort: 0.0001.

[6] The only exception is a rational agents population. See Dal Forno and Merlone (2001) for theoretical details and 6.16.

References

CARLEY K (1991), A Theory of Group Stability. *American Sociological Review*, 56(3). pp. 331-354.

CARLEY K M and Prietula M J Eds. (1994), *Computational Organization Theory*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, Publishers.

COHEN M D, March J G and Olsen J P (1972), A Garbage Can Model of Organizational Decision Making. *Administrative Science Quarterly*, 17. pp. 1-25.

DAL FORNO A and Merlone U (2001), Incentive Policy and Optimal Effort: Equilibria in Heterogeneous Agents Populations, *Quaderni del Dipartimento di Statistica e Matematica Applicata*, n.10.

HARRISON J R and Carrol G R (1991), Keeping the Faith: A Model of Cultural Transmission in Formal Organizations. *Administrative Science Quarterly*, 36(3). pp. 552-582.

HARRISON J R and Carrol G R (2001), Modeling Organizational Culture: Demography and Influence Networks. In Cooper C L, Cartwright S and Earley P C (Eds.), *The International Handbook of Organizational Culture and Climate,* Chichester, UK: John Wiley & Sons.

HERMALIN B E (2001), Economics and Corporate Culture. In Cooper C L, Cartwright S and Earley P C (Eds.), *The International Handbook of Organizational Culture and Climate,* Chichester, UK: John Wiley & Sons.

KREPS D M (1990), Corporate culture and economic theory. In Alt J E and Shepsle K A (Eds.), *Perspectives on Positive Political Economy*, Cambridge: Cambridge University Press, pp. 91-143.

MACHO-STADLER I and Pérez-Castrillo D (1997), *An Introduction to the Economics of Information (Incentive and Contracts)*. Oxford: Oxford University Press.

MARSCHAK J and Radner R (1972), *Economic Theory of Teams*. New Haven, Conn.: Yale University Press.

MEYERSON R (1991), *Game Theory*. Cambridge Massachusetts: Harvard University Press.

MORRIS S (1997), Interaction Games: a Unified Analysis of Incomplete Information, Local Interaction and Random Matching. Working Paper. University of Pennsylvania.

RADNER R (1992), Hierarchy: The Economics of Managing. *Journal of Economic Literature*, Vol. XXX. pp.1382-1415.

RUBINSTEIN A (1998), *Modeling Bounded Rationality*. Cambridge Massachusetts: the MIT Press.

SIMON H A (1955), A Behavioral Model of Rational Choice, *The Quarterly Journal of Economics*, 69(1). pp. 99-118.

SIMON H A (1982), *Models of Bounded Rationality, vol.2 Behavioral Economics and Business Organization*. Cambridge Massachusetts: MIT Press.

STEERS R M and Black J S (1994), *Organizational Behavior*. Fifth Edition, New York, NY: HarperCollins College Publishers.

YOUNG H P (1993) The Evolution of Conventions, *Econometrica*, Vol. 61, No. 1., pp. 57-84.

YOUNG H P (2001), *Individual Strategy and Social Structure*. Princeton, NJ: Princeton University Press.

Return to Contents of this issue