Grounding Reputation Experiments. A Replication Of A Simple Market with Image Exchange

Antonietta di Salvatore[1]
Isaac Pinyol[2]
Mario Paolucci[1]
Jordi Sabater-Mir[2]

Abstract

In this paper we have compared two independent models implemented to analyze the effects of social information transmission. The first of them, Repage, is centered on an articulated theory of social evaluations. The second system, that has been developed in NetLogo, follows the same specifications of the Repage model but with several simplifications. We implemented the second model to confirm the reliability of the most simple results that can be obtained with Repage. We model the spreading of information in a simple market with the presence of liars and the possibility of retaliation. We want to compare the use of the experience alone with the usage of image (agents believed evaluation of a target) to find the good sellers. Both models give the same results in the different simulation sets. Image is shown to be preferable over experience only in particular situations.

Keywords: Reputation, Trust, Social Simulation

1 Introduction

While many authors have simulated the spread of reputation with simple agents (for a review, see (Sabater and Sierra 2005)), more elaborated, cognitive approaches are still needed. The first systems based on cognitive theories like (Conte and Paolucci 2002) have started to appear (see for instance (Sabater *et al.* 2006)), but they are posing new challenges to the way of doing social simulation. Indeed, dealing with Image as distinct from Reputation requires the design of a rather complex social capacity for intelligent autonomous agents. It requires a more complex agent model and architecture than what is usual, endowed with:

- an explicit model of reputation that permits to account for social control (Dunbar 1998);
- explicit modeling of informational autonomous agents: agents cannot be autonomous if they are not designed to weight, partially accept or partially refuse new information arriving from other agents, in a cognitively plausible way.

[1]ISTC-CNR - Institute for Cognitive Science and Technology, Rome, ITALY
[2]IIIA-CSIC - Artificial Intelligence Research Institute, Barcelona, SPAIN

In an unpredictable world, intelligent agents are shown to depend on accurate information (Laffont 1993), (Kluegl and Bazzan 2004) for acting adaptively. More specifically, they depend on accurate social information for interacting into a heterogeneous multiagent world. Memory of past experience is a precious source of information, but is usually acquired at own expenses. As obtaining experience may be fatal in a world of cheaters, agents depend on one another to indirectly acquire information for partner selection, before interacting with, and in order to avoid, the bad guys.

In the last ten years or so (Carter *et al.* 2002), the role of indirectly acquired social information has been appreciated by social scientists to a fairly realistic degree. Indeed, reputation has received a growing attention as network-based social evaluation (Raub and Weesie 1990). However, this way the rate of inaccurate information circulating in a multiagent system increases, and the question is how to put up with such inaccuracy. With news, a lot of junk information spreads around, without societies falling apart. How is this possible?

We have not yet an answer to this question; actually, answering it would require a whole, large scale research program. But we are starting to see the first elements taking their place. Indeed, the attention of the international community is starting to concentrate on these themes, and the first research projects, although on a limited scale, are starting to appear. In particular, we have an interest in simulating a simple market situation in which we plan to study the characteristics of several information levels. We plan to start from direct experience and move up to the communication of direct experience, a field already quite intensely studied (for a review, see (Sabater and Sierra 2005)). In addition, we had a plan to study the effect of more detached forms of communications - gossip, in the sense of communications that do not explicitly indicate the source, and thus diffuses with lower responsibility attribution. To carry on this study, we developed a very detailed module for the management of reputational information, Repage.

However, to understand the effects of complicated cognitive structures working on top of a simulation model, we need to have in advance a firm understanding of the working of that system, in the same sense in which one cannot speculate about perturbations of a system when the behavior of that system is not yet really understood. To this purpose, we have decided to provide an independent implementation of the basic levels of the simulation setting. We have thus obtained two different systems. The first of them, Repage, based on the Jadex agent platform, provides a complete implementation of the simulation experiment in Java, and is centered on an articulated theory of social evaluations. The second system has been developed in NetLogo, following the same specifications of the first experiment but with several simplifications. First, we resort to a simplistic representation of social evaluation instead of the elaborated Repage one. In addition, both because it has been independently implemented, and because of simplifications imposed by the NetLogo simulation platform, the simulation model itself is not really identical to the one implemented in Repage. The simplifications have also the effect of having a faster, more manageable and more scalable platform. This could carry the exploration of different scenarios that for time and memory constraints cannot be currently tested with the Repage system.

To resume, our purpose in this work is twofold:

- draw a confirmation of our hypotheses on the simulation model, confirming the reliability of the simplest results obtainable with Repage;
- keep exploring the simulation model outside of the parameter area reachable by Repage.

After the presentation of the background theory that motivates our efforts, we will present the simulation model that we have implemented and we will discuss shortly how the independent implementations produced algorithmic differences. Then we will show results from experiments run on both platforms.

2 The Repage model

2.1 Theory of Reference

Repage (Sabater *et al.* 2006) is a computational system based on a cognitive theory of reputation (Conte and Paolucci 2002) that proposes a fundamental difference between image and reputation. This theory claims that althougth both image and reputation are social evaluations, image is a simple evaluative belief and reputation a metabelief (a belief about others' beliefs). Image refers to an evaluation that the agent holding it acknowledge as true. It represents the agent's opinion (goodness or badness) of certain target agent with respect to an specific norm, standard or skill. In contraposition, reputation is a belief about the existence of an evaluation that circulates in the society. For this reason, to acknowledge the existence of a reputation does not imply to accept the evaluation itself. For instance, agent A might have a very good image of agent B as a seller, and at the same time accept that it is said that agent B is a bad seller.

In this paper, we do not deal with reputation. However, the difference between image and reputation carries important consequences that we take into account in the experiments shown in this paper. Communicating an image about somebody implies that the source agent is giving its personal opinion, and therefore, there implicitly is a commitment about the truth of the communication. In this sense, if the recipient agent realizes that the information received as image is false, it may retaliate the source by not giving him/her accurate information or avoiding future hypothetical cooperation. This commitment is not present in a communicated reputation, since as we stated, to accept a reputation does not imply to accept the evaluation itself.

In the next subsections we describe the Repage architecture and its integration with the other parts of a computational agent.

2.2 Repage Architecture

In the Repage architecture we find three main elements, a memory, a set of *detectors* and the *analyzer* (see figure 1). The memory is composed by a set of references to the predicates hold in the main memory of the agent. Only those predicates that are relevant for the calculus of reputation and image are considered.

In the memory, predicates are conceptually organized in levels and inter-connected. Each predicate that belongs to one of the main types (image, reputation, shared voice, shared evaluation, valued communication and outcome) contains an evaluation that refers to a certain agent in a specific role. For instance, an agent may have an image of agent A (target) as a seller (role), and an image of the same agent A as informant.

We maintain the value associated to a predicate as a tuple of five numbers (summing to one) plus a strength value. Each number has an associated label in the rating scale: very bad (vb), bad (b), neutral (n), good (g) and very good (vg). We call this representation a *weighted labeled tuple* and it represents a probability distribution.

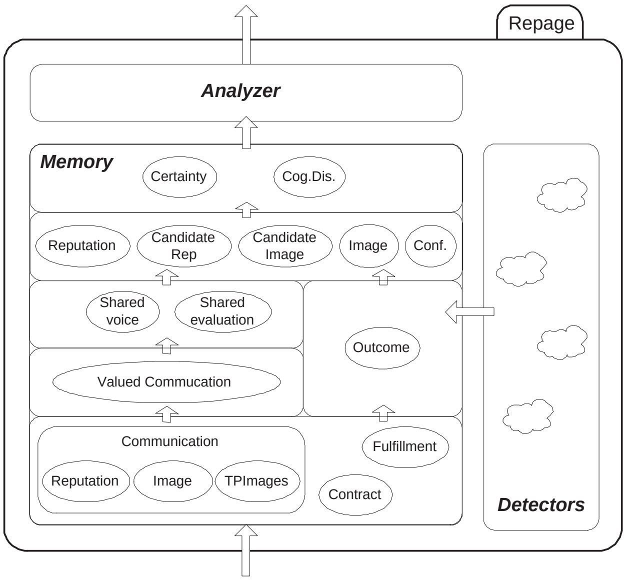

Figure 1: The Repage Architecture.

The network of dependences specifies which predicates contribute to the values of others. Each predicate (except those at the bottom level) has a set of antecedents and at the same time contributes to the calculation of other predicates. The *detectors*, inference units specialized in each particular kind of predicate, receive notifications from predicates that have changed or that appear in the system, like new communications or new fulfillments, and use the dependences to recalculate the new values and to populate the memory with new predicates. The aggregation of evaluations is done with a simple weighted mean. Mapping the sorted discrete set with integers from 0 to 4, let w_i be the probability (or weight) of element i of the evaluation w, let n be the number of evaluations to aggregate, and let w^j be the jth evaluation, the resulted evaluation w satisfies the following equation:

$$\forall i : 0 \leq i \leq 4 : w_i = \frac{\prod_{j=1}^{n} w_i^j}{\sum_{i=0}^{4} \prod_{j=1}^{n} w_i^j} \tag{1}$$

Furthermore, each predicate has associated a strength that is function of its antecedents and of the intrinsic properties of each kind of predicate. As a general rule, predicates that resume or aggregate a bigger number of predicates will hold a higher strength. However, strength is closely related to bias factors, rules that for instance, give more importance to direct experiences that indirect experiences, and that may come from sociology or psychology theories, or from simple common sense. In other words, bias rules tune the intrinsic performance of Repage in the meta-situation defined by being in the mind of an autonomous social agent that participates in a society and needs the interaction with others in order to achieve her/his goals.

At the first level of the Repage memory we find a set of predicates not evaluated yet by the system.

- Contracts: agreements of the future interaction between two agents. For instance, in an e-Commerce environment, an agent may expect that the maximum quality of a product that for sure the seller is saying will offer.

- Fulfillment: the result of the interaction. In the same e-Commerce example, the fulfillment would be the real quality of the product the agent got.

- Communications: Information that other agents may communicate about others evaluations. These communications may be related to three different aspects: the image that the informer has about a target, the image that according to the informer a third party agent has, and the reputation that the informer has about the target.

In level two we have two kind of predicates:

- Valued communication: The subjective evaluation of the communication received that takes into account, for instance the image the agent may have of the informer as informant. Communications from agents whose credibility in terms of image or may be reputation are low, will not be considered as strongly as the ones coming from well reputed informers.

- Outcome: The agent's subjective evaluation of the direct interaction. From a fulfillment and a contract a *detector* builds up an outcome predicate that evaluates the particular transaction.

In the third level we find two predicates that are only fed by valued communications. On one hand, a shared voice will hold the information received about the same target and same role coming from communicated reputations. On the other hand, shared evaluation is the equivalent for communicated images and third party images.

Shared voice predicates will generate candidate reputation, and share evaluation together with outcomes, candidate image. In this fourth level candidate reputation and candidate images aren't strong enough to become a full reputation and image respectively. New communications and new direct interactions will contribute at this level to enrich these predicates and therefore "jump" to images and reputations.

The last level implements cognitive dissonances and certainties. From the point of view of the agent, different pieces of relevant information may conclude in contradictory information (cognitive dissonance) or the opposite, certain information. In the case of dissonance, the *analyzer* will propose actions to the agent in order to solve the contradiction. We refer to ((Paolucci *et al.* 2005)) for a more detailed explanation about how the *analyzer* works.

The integration of Repage with the other parts of our deliberative agents is made through the Repage memory that is always synchronized with the agent main memory. Each time Repage adds a new predicate this appears in the main memory of the agent and the other way around. Therefore the predicates that access Repage and the other modules like the decision making module or the communication module are the same.

3 Description of the experiments

We have designed the simulation experiments as the simplest possible setting where accurate information is a *commodity*, meaning that information is both valuable and scarce. All unnecessary detail has been removed, generating a generic economic metaphor of a very simple agent-based market setting with instability. The experiments include only two kind of agents, the buyers and the sellers.

All agents perform actions in discrete time units (turns from now on). In a turn, a buyer will perform a communication request and one purchase operation. In addition, the buyer will answer all the information requests that has received previously. We fix the number of buyers to 25.

Goods are represented by an utility factor that we interpret as quality (but, given the level of abstraction used, could as well represent other utility factors as quantity, discount, timeliness) with values between 1 and 100.

Sellers are characterized by a constant quality and a fixed stock, that is decreased at every purchase; They are essentially reactive, their functional role in the simulation being limited to providing an abstract good of variable quality to the buyers. Sellers exit the simulation when their stock is exhausted and are substituted by a seller with the same characteristics. In addition, there is a maximum window of inactivity; Sellers disappear if they are not able to sell anything during a time longer than this window. At the beginning, they are generated in three groups. Disappeared sellers are substituted with individuals drawn from the same pool. We start with a fixed number of good sellers (sellers that provide quality 100) and bad sellers (sellers that provide quality 0), and the rest is drawn from a group where the provided quality of those agents is uniformly distributed between 0 and 100.

The sellers' disappearance is the reason for the necessity of information; Reliable communication allows for faster discovery of the better sellers. This motivates the agents to participate in the information exchange. In a setting with permanent sellers (infinite stock), once all buyers have found a good seller, there is no reason to change and the experiment freezes. With finite stock, even after having found a good seller, buyers should be prepared to start a new search once the stock of the good seller ends. On the other hand, limited stock makes good sellers a scarce resource, and this constitutes a motivation for the agents of not distributing information. One of the interests of the model is in the balance between these two factors.

In the following sections we explain how the agents carry on their decisions - where to buy, what kind of information to ask and to whom, and so on. We will see also that in the agent implementation there are several differences between the two platforms.

4 Repage decision model

In this section we are going to describe the decision making process that follows a JADEX buyer implementing the Repage model in the experiment described above. We differentiate between the situation where there is no exchange of information (what we call Level-0, L0) and where there is an exchange of images (Level-1, L1).

4.1 Information representation

Repage stores the information as a set of predicates connected through a network of dependency. An image in Repage differs from the Netlogo implementation in how its value is represented. Instead of being a single number between -1 and 1, it is a probability distribution over the discrete sorted set: Very Bad, Bad, Normal, Good, Very Good. Each image has also associated a strength that reflects how reliable is that image. In order to take decisions the agent uses the center of mass(CM) of the probability distribution to apply different thresholds. This measure gives a real number included in the interval $(0,4) \in I\!R$ that indicates to which value of the discrete set tends the evaluation represented as probabilistic distribution, taking into account the mapping defined in the section 2 where $0 \leftrightarrow vb, 1 \leftrightarrow b, 2 \leftrightarrow n, 3 \leftrightarrow g, 4 \leftrightarrow vg$. Let w be an evaluation represented as probabilistic distribution, the function CM is defined as:

$$CM(w) = \sum_{i=0}^{4} i w_i \tag{2}$$

4.2 Level-0, no exchange of information

The only decision an agent has to take in this situation is to choose which is the seller that can offer a better quality. The agent bases its decision in the following algorithm:

1. From all possible sellers, a subset of "good enough" sellers is selected. These list of "good sellers" is built using the images of the sellers. A seller is considered to be "good enough" if the image that the agent has on it as a seller has a center of mass greater than 3.7 (that is, *vg* if we look at the corresponding label); in this experiment, agents are looking for the few sellers with very high quality.

2. If the list of "good enough" sellers is not empty and the agent has decided not to risk[1], one randomly selected seller from the "good enough" list is chosen as a partner.

3. If the list of "good enough" sellers is empty (no good sellers have been detected) or the agent has decided to risk, then one seller from the whole list of agents is selected. Each seller has a probability of being selected that is proportional to the center of mass of its image value or to a default probability for those that are completely unknown.

4.3 Level-1, exchange of images

As described before, the decision making process at this level can be divided in two different parts:

- The decisions related to choose a seller or an informer.
- The decisions about how the agent should answer the possible questions.

4.3.1 Choosing a seller and asking questions

To decide which is the best seller and which is the best informant[2] the agent uses the same process described for level-0. Of course, when the agent needs to choose a good informant, the process is the same but considering the list of possible informants instead of the list of possible sellers.

[1]The notion of risk in this context is associated to the exploration willingness. An agent with a high risk propension will tend to interact with unknown sellers instead of remaining with known possibilities

[2]in our experiments the informants are always the buyers.

Once the agent has selected an informant the next step is to decide the topic of the query. In our experiments there are two options each one with a 50% of probabilities to be selected by the agent: (i) ask about how honest is a concrete informant or (ii) ask for some image (either good or bad) of a seller. In the first option, the agent specifies who is the object of the query (in our case, it is always the most unknown informant, that is, the one with the image closest to the absolute center of mass) while in the second the agent is asking information about sellers in general and is the informant who decides about whom it will give information.

4.3.2 Giving answer to questions

As we have seen, from time to time, buyers receive questions coming from other buyers. At that point they need to decide what to answer. The first step is to decide if they will lie or not. In these experiments, lying means answering a value that is the oposite of the one the agent thinks it is the truth. An agent lies if it is a cheater or if the querier is recognized as being a cheater (retaliation).

If the question is about the image of a good/bad seller, the agent selects from the known sellers those that are clearly good or clearly bad (for building these lists the agent uses the same procedure presented for level-0). The answer will be the image of a random seller from the list of good sellers, a image of a random seller from the list of bad sellers if the good sellers list is empty or an "I don't know" if both lists are empty. If the agent decided to lie, the image is inverted before send it to the querier.

If the question is about the image of a certain buyer as informant the agent will send the available image (inverted if the agent decided to lie) or "I don't know" if there is no image formed.

4.3.3 Information update

Each time the agent receives a new information from an informer or interacts with a seller, the new knowledge is introduced in the Repage memory. This new knowledge will activate the Repage inference engine (the *detectors*, see section 2) that will recalculate each evaluation affected by the new income information.

5 NetLogo decision model

In this section we will describe the decision making process that NetLogo agents follow to participate in the experiments. For the description, we differentiate between what we call Level-0, where there is no exchange of information, and Level-1, where there is exchange of images.

5.1 Information representation

In general, we represent information in NetLogo as simple lists of scalar numbers. Each buyer owns three different information kinds:

1. a list $S1$ of sellers with whom it had direct experience (i.e. at least one purchase), and of whom it knows the actual quality; this list provides the observed quality;

2. a list $S2$ of sellers about whom only (positive or negative) indirect information is available; this list contains only boolean (good/bad) information;

3. a list I containing informers'image, represented by a numerical value. An image larger than zero corresponds to a good informer while a negative image corresponds to a liar.

The list $S1$ have no more than one seller with good image both Level-1 and Level-0. Indeed when a buyer finds a seller with satisfying quality it will buy from this seller until the last one depletes its stock and disappears from the market. At this point the information about the old seller is removed from list $S1$ and the buyer has to find a new good seller.

5.2 Level-0, no exchange of information

At this level each buyer has to choose a seller to buy from and after it has remember its quality. The agent bases its decision on the following algorithm, in order:

- if available in $S1$, it buys from the seller with satisfying quality;
- then, if in $S1$ there are only sellers with not satisfying quality, it chooses radomly from unknown sellers (any one except those in $S1$) and, finally, it updates $S1$ with the reference about the new seller.

5.3 Level-1, exchange of images

At this level, agents can exchange information and this implies they have to take more decisions:

- the decisions to choose a seller and an informer;
- the decision about how the agent should answer incoming questions.

5.3.1 Choosing an informer and asking questions

To ask for informations, the buyer will choose as informer the one with the higher non negative image in list I. If there is none, it will choose at random among all unknown informers. Finally, if all informers are given a negative image, it will abstain from asking. Note that this quite improbable case can be reached for recognizer cheaters, since agents lie to recognized cheater per retaliation. Like Repage, once the agent has selected an informant, it can ask for a seller's image or for an informer's image both with the 50% of probability.

5.3.2 Giving answer to questions

Informers give priority to verified informations. On a seller's request, when possible, they will send a reference to a seller they had direct experience with. They will send reported image only when they have not such reference. The honest informer will check the requester's image. If the requester is unkwon or has not a negative image, it will choose, in order one of the following alternatives:

- if available, it will send a positive reference to the seller with satisfying quality;

- if not, it will send a negative reference to the lower image seller in $S1$;
- if $S1$ is empty, it will send a reported good information from $S2$, giving precedence to positive images;
- If $S2$ is also empty, no information will be sent.

Buyers that are informational cheaters or that are answering an agent that they repute as a cheater (image less than zero) follow exactly the same routine as for good informers. Once the information to send is selected, they reverse its value: a satisfactory seller will be reported as bad, and so on.

5.3.3 Information update

Once an information is obtained , the requester updates its knowledge. A new piece of information may regard both an already known agent or an agent on which the requester has not yet any information.

For what regards sellers, if already known, the new information will be checked for compatibility. If the check is positive, the image of the informer in I will be increased of 0.1; in the opposite case, it is decreased by the same amount. If unknown, after a repetition check (if the receiver has already received the same evaluation from the same informer, it will disregard the new one), the new information will be registered.

For informers, if the target is not known, it will be created with an image of 0.1/-0.1, depending on the good/bad content of information. If it is already known, its image will be updated but only for half of the previous amount.

Remember that the source of information cannot have negative image since agents will never request informations to informers considered as liars.

Note that in the case a buyer is acting on the basis of received information, they will also test its validity. If the information results as accurate, it will upgrade the reliability level of the informer: it will introduce that informer in I with value 0.1, or increase its image of 0.1 if already present. In the opposite case, the image will be reduced of the same amount, and at the same time they will remove all informations received from that source. Note also that, in both cases, after the interaction the actual quality of the seller is known (registered in $S1$).

5.3.4 Choosing a seller to buy from

Choosing a seller follows a sligtly different procedure from that explained for Level-0. To choose its seller each buyer opts for the best information available, i.e. in order one of the following:

- if available, the seller with satisfying quality;
- then, if available it will choose one of the sellers on wich it has good reference (from $S2$)
- finally, it will choose randomly from any unknown seller (except those in $S1$ and in $S2$).

If the buyer chooses its seller from $S2$ it will verify the received information and after it will update its knowledges (about the informer and the seller).

6 Research Questions

In the present paper, we will focus on the different influence of information on large and small offer in markets; in other words, we will study the performance of buyers with growing number of sellers, keeping constant the number of good sellers. We will explore the two settings $L0$, in which agents can use only personal experience, and $L1$, where we allow for communications (including a fixed number of informational cheaters). In addition, we will show both results obtained with the Repage platform (R) an with the NetLogo platform (N).

To describe an experiment, we need the number of buyers NB, the number of sellers NS, and the stock for each seller (S).

The hypotheses at this point are quite simple:

H1 Advantage of communication: L1 shows an improvement over L0

H2 Data reproduction: Data obtained from the N and the R platform are comparable

H3 Seller effect: while in a small world it is very easy to find out the good sellers, with large numbers of sellers the search is more difficult.

7 Simulation Runs and Result Analysis

We have run many simulations to compare the different experimental conditions. We have consider low stocks varying from 10 to 20, and low and high number of sellers, from 150 to 1100. Simulations are run for 100 time steps; at the end, we register the total quality earned from all agents.

We will show, for several combination of the parameters, the temporal curves for the four experimental situations - N_L0, N_L1, R_L0, R_L1. In the plots, we show the total earned quality averaged over the agents and over the number of steps. We will examine four situations, combining low/high stocks and low/high number of sellers. In these experiments there are no cheaters.

7.1 Low number of sellers

From Figure 2, we can start to see that the two implementations (N lines and R lines) are actually quite convergent (less than 3 points in a scale of 60-70). In this case, in both implementations we can see an improvement of L1 over L0 that is more accentuated when the stock is lower. NetLogo implementation shows a bigger difference between both situations than the Repage one, however we can observe that both implementations behave in a similar way.

7.2 High number of sellers

From Figure 3, the relative position of Repage and Netlogo is inverted. Notice that the general achieved quality decreases with respect to the previous subsection, since higher number of sellers implies more difficulty finding good sellers. For this reason, it takes more time steps to observe a difference between L0 and L1, but still in both implementations the impact of exchanging image is evident, improving the gained quality.

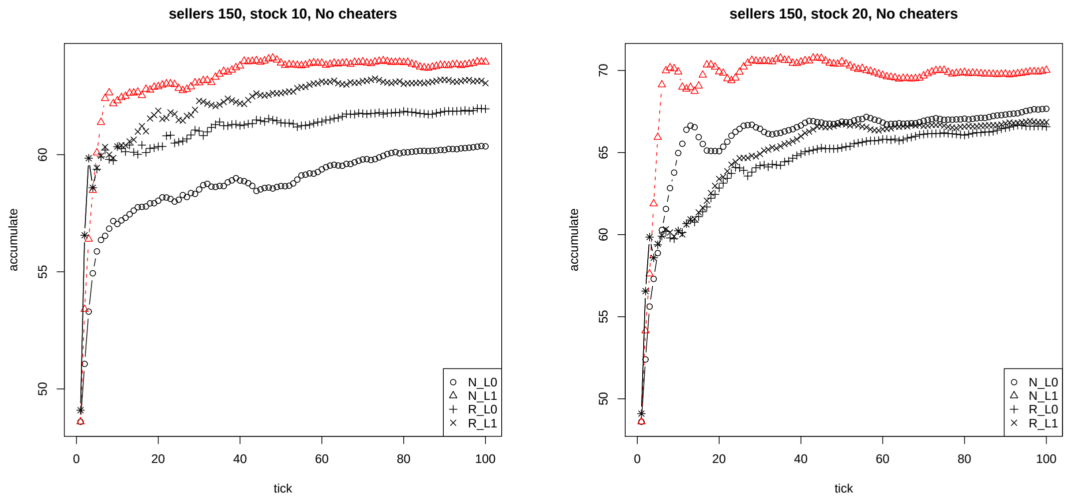

Figure 2: Accumulated average quality per turn with 150 sellers

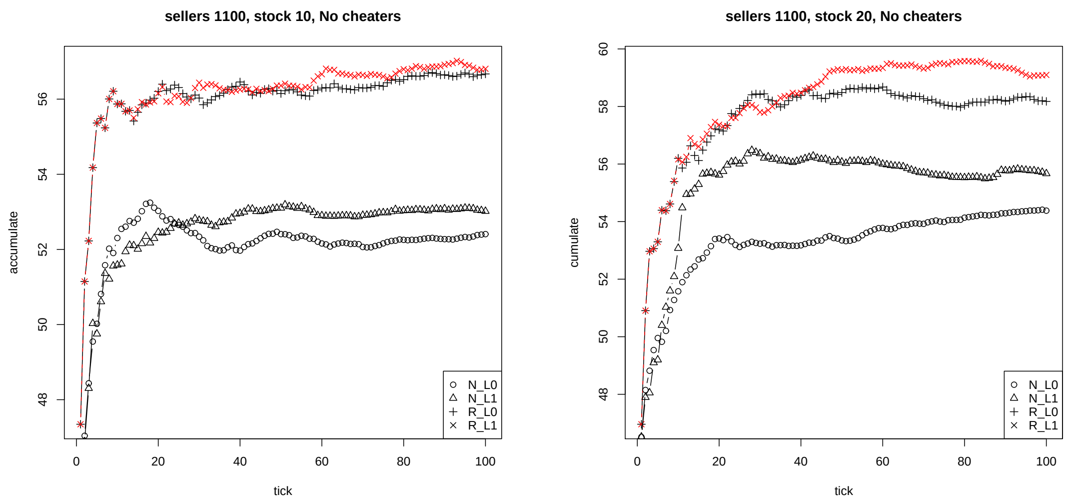

Figure 3: Accumulated average quality per turn with 1100 sellers

8 Conclusions and Future Work

In this paper, we have shown a successful re-implementation of a model designed for simulations with complex agents. The simplified (N) model would not be able to produce the necessary elaboration to deal with an actual cognitive theory of reputation, what the complex (R) model was designed to do. On the other side, with the simplified model we have been able to explore parameter areas that would have been inaccessible to Repage.

The real added value of this work, however, lies in the improved understanding and increased trust in our basic model, that would not have obtained if we would limit to a single implementation. When starting to make simulations to study a complex effect - and in most social cases, complex effect are the only interesting ones - one must be assured to have a deep grasp of the reference model *before* the introduction of complex social artifacts. We can now consider this basic "image" level confirmed, paving the way to future experiment in Repage for the study of reputation.

More simulations are currently on process in order to study the effect of cheater agents that send not accurate information. This study is closely related to the study of reputation we mentioned in the previous paragraph. The fact of including cheaters in an environment where the only allowed communications are personal opinions implies that agents use less this capability, and therefore random choices are preferred over asking questions. We claim that the use of reputation may help to solve the lack of communication in these situations. A NetLogo re-implementation will help us to play with all these concepts, but a previous validation is needed, and for this reason more simulations are done including in scenarios different percentages of cheaters.

9 Acknowledgments

This work was partially supported by the European Community under the FP6 programme (eRep project CIT5-028575, OpenKnowledge project FP6-027253 and SocialRep project MERG-CT-2005/029161). Jordi Sabater-Mir enjoys a RAMON Y CAJAL contract from the Spanish Government.

Bibliography

References

CARTER J, Bitting E, and Ghorbani A A (2002) Reputation formalization for an information sharing multi agent system. *Computational Intelligence*, 18(4):515—514, 2002.

CONTE R and Paolucci M (2002) *Reputation in artificial societies: Social beliefs for social order.* Kluwer Academic Publishers, 2002.

DUNBAR R (1998) *Grooming, Gossip, and the Evolution of Language.* Harvard Univ Press, October 1998.

KLUEGL F and Bazzan A L C (2004) A case study on the role of information for implicit coordination. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS-04), New York, USA*, 2004.

LAFFONT J J (1993) *The Economics of Uncertainty and Information.* Cambridge, MA: MIT Press, 4th edition, 1993.

PAOLUCCI M, Sabater-Mir J and Conte R (2005) 'what if?' dealing with uncertainty in repage's mental landscape. In *The 2005 IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT-2005)*, pages 372–378, September 2005.

RAUB W and Weesie J (1990) Reputation and efficiency in social interactions: An example of network effects. *American Journal of Sociology*, 96, 1990.

SABATER J and Sierra C (2005) Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, 2005.

SABATER J, Paolucci M and Conte R (2006) Repage: Reputation and image among limited autonomous partners. *Journal of Artificial Societies and Social Simulation*, 9(2), 2006.