"What if?" Dealing with uncertainity in *Repage*'s mental landscape

Mario Paolucci, Jordi Sabater, Rosaria Conte
ISTC-CNR, Via San Martino della Battaglia 44, 00185 Roma, Italy
mario.paolucci@istc.cnr.it, jsabater@iiia.csic.es, rosaria.conte@istc.cnr.it

Abstract

This paper is focused on cognitive decision-making about how to solve inconsistencies and incompleteness in social evaluations (e.g. about potential partners in exchange). We propose a development of Repage*, a computational system presented in [11] for forming and updating social evaluations. The system draws upon a fundamental difference between REPutation and imAGE [3] as a way out from the the trade-off between agents' autonomy and their need to adapt to social environment. A full exploitation of its potentialities includes the activation of a special module, the Analyzer, aimed at solving possible inconsistencies, uncertainties and incompleteness in the output of lower level modules by means of inner simulation. In this work,* Repage*'s Analyzer architecture will be described; some representative examples of problems posed by the Planner to the Analyzer will be discussed and hypotethical simulations will be run within this module to find a solution to uncertainity, avoiding at the same time the exceeding complexities of rule-based reasoning and the costs of reinforcement learning.*

1. Reputation and uncertainity

Reputation co-evolved with human language and social organization [6] as a multi-purpose social and cognitive artifact. Beside allowing for partner selection in exchange, reputation incentives cooperation and norm abiding and discourages defection and free-riding, handing out to cooperators a weapon for punishing transgressors by cooperating at the level of information exchange [4]. The role of reputation as a partner selection mechanism started to be appreciated in the early eighties [8]. However, little understanding of its cognitive underpinnings was achieved at that stage. Evolutionary game theorists ignored the difference between image (i.e. own believed evaluation of a target) and reputation (i.e. nested or meta-evaluation: a belief about how a given target is commonly said to be evaluated). Consequently, the decision to report on reputation to others was ignored.

Hence, the efficacy of preventive social knowledge was not fully appreciated and, what is worse, the role of reputation in updating existing social evaluations was overlooked. The distinction between image and reputation brings about the necessity for a cognitive approach to the subject matter, now perceived to be fundamental in the study of reputation (see for example [7], which however is focused on the mental *effects* of reputation, rather than on the cognitive nature of the phenomenon).

A special field of application, which is becoming more and more important, is the effect of reputation in virtual and agent-mediated markets. Missing the cognitive side of the phenomenon, online reputation systems are found to work poorly and inefficiently [10, 1]. In this paper, we mean to show how *Repage*, a computational system for partner selection [11], can be augmented by a specialised module to repair inconsistencies and overcome uncertainties and incompleteness by internally simulating the effects and the costs of alternative courses of action. Based on a model of REPutation, imAGE, and their interplay, *Repage* provides evaluations on potential partners and is fed with information from others and outcomes from direct experience. This is fundamental to account for (and design) limited autonomous agents as exchange partners[12].

However, the interplay between image and reputation might be a cause of uncertainty and inconsistency. Inconsistencies do not necessarily lead to a state of cognitive dissonance, nor do they always urge the system to find a solution. For example, an inconsistency between own image of a given target and its reputation creates no serious problem to the system. However, a contradiction between own evaluations is sometimes possible: I may get a good impression from a given experience with the target, which may be dismantled later. Or, my direct experience may be confirmed in further interaction, but at the same time it may be challenged by the image I believe others, whom I trust a lot, have formed about the same target. In such a case, I will find myself in a rather awkward condition, especially if that target is one of the few, if not the only, available partners for a necessary transaction. What will I do in such a condition? Will I go ahead and sign a contract, may be a low-cost one, just to acquire a new piece of direct evidence, or will I check the reliability of my informants? Suppose that the latter alternative is chosen on the grounds of a cost to benefit rule. What is meant by checking others' reliability? If their image of the target is better than one's own, and if one should discard direct experience based on the previous rule, what else should one do? As is easily perceived, the picture is rather complex, and the number of possibilities is bound to increase at any step, making the application of rule-based reasoning computationally heavy.

We propose here an innovative approach suggested by simulation modeling and technique. Unlike its usual application, this approach should not be seen as a validation approach to be adopted by the scientist from the outside of the system, but as a decision-making instrument to be incorporated into the agent technology itself as a look-ahead mechanism. Rather than try and see, *Repage*'s Analyzer allows agents to wait and simulate their own mental landscapes, taking into account alternative courses of available action and visualize their effects. Interestingly, this allows a rather sophisticated form of learning to be applied, thanks to which actions that outcompete concurrent alternatives will be reinforced.

2. Value added of reputation

The social cognitive perspective on reputation presented in this paper aims to model both the cognitive properties and the social aspects of reputation, that is, its transmission. In order to model both, it is necessary to understand the interrelationships between two different types of social evaluation, i.e. image and reputation.

Image and reputation are distinct objects. Both are social evaluations. They concern other agents' (targets) attitudes toward socially desirable behavior, and may be shared by a multitude of agents. Image is an evaluative belief; it tells that the target is "good" when it displays behavior *b*, and that it is "bad" in the opposite case. (Social) evaluations may concern physical, mental and social properties of targets. In particular, agents may evaluate a target as for its capacity as well as its willingness to achieve a shared goal. The interest/goal *with* regard to which a target is evaluated may also be a distributed or collective advantage.

At the meta-level, a given agent may have a belief about others' evaluations of a target. When these evaluations converge, we say that the image of the target is shared among a given set of agents. In the latter case, we speak about a shared image. Notice, that a shared image of a target is not yet a reputation, since the latter results from a process of communication among agents sharing a common environment. Although it does not coincide with one's own image of a target, a shared image is likely to be accepted by the individual agent, especially if those that share it enjoy a good evaluation in the latter's eyes. *Reputation* is instead a shared voice, i.e. a belief about others saying that a given target enjoys or suffers from a shared image. Whereas image consists of a set of evaluative beliefs [9] about the characteristics of a target, reputation is the process and the effect of transmission of a target image. More specifically, the image relevant for social reputation concerns a subset of the target's characteristics, i.e. its willingness to comply with social norms. To understand the difference between image and reputation, the mental decisions based upon them must be analyzed. They consist of three decisions, epistemic, to accept the beliefs that form either a given image or acknowledge a given reputation, pragmatic-strategic, to use these beliefs in order to decide whether and how to interact with the target, and memetic, to transmit these beliefs to others.

This difference is not inconsequential: to spread news about someone's reputation does not bind the speaker to commit himself to the truth value of the evaluation conveyed but only to the existence of rumours about it. Therefore, unlike ordinary sincere communication, only the acceptance of a meta-belief is required in communication about reputation. And unlike ordinary deception, communication about reputation implies *(i)* no personal commitment of the speaker with regard to the main content of the information delivered; if the speaker reports on t's bad reputation, he is by no means stating that t deserved it; and *(ii)* no responsibility with regard to the credibility of (the source of) information ("I was told that t is a bad guy"). Two points ought to be considered here. First, the source of the meta-belief is implicit ("I was told"). Secondly, the set of agents to whom the belief p is attributed is non defined ("t is ill/well reputed").

Of course, this does not mean that communication about reputation is always sincere. Quite on the contrary, one can and does often deceive about others' reputation. But to be effective the liar neither commits to the truth of the information transmitted nor takes responsibility with regard to its consequences. If one wants to deceive another about somebody's reputation, one should report it as a rumour independent of or even despite one's own beliefs! As a consequence of this analysis, we can see how, unlike other (social) beliefs, reputation may spread in a population even if the majority does not believe it to be deserved. Meta-beliefs spread without first-level beliefs spreading.

3. Current systems

Applications of reputation abound in two sub-fields of information technologies, i.e. computerized interaction (with a special reference to electronic marketplaces) and agent-mediated interaction. Large systems like eBay show a characteristic bias to under-provided positive evaluations [10], suggesting that factual cooperation among users at the information level may lead to a "courtesy" equilibrium [4]. Interpretative hypotheses have been suggested [3], pointing to eBay-like systems as centralized image systems rather than reputation-based ones. Analogous considerations are made with regard to other systems, although systems based upon networks of agents have been shown to provide partial exceptions [15].

Models of reputation for multi agent systems applications [2, 13, 5] clearly present interesting new ideas and advances over conventional online reputation systems, and more generally over the notion of global reputation, or centrally controlled image. The "agentized environment" is likely to produce interesting solutions that may apply also to online communities. This is so for two main reasons. First, in this environment two problems of order arise: to meet the users' expectations (external efficiency), and to control agents' performance (internal efficiency). Internal efficiency is instrumental to the external one, but it re-proposes the problem of social control at the level of the agent system. In order to promote the former, agents must control, evaluate, and act upon one another. Reliability of agents implements reliability of users. Secondly, and consequently, the agent system plays a double role, it is both a tool and a simulator. In it one can perceive the consequences of given premises, which may be transferred to the level of users' interactions. In a sense, implemented agent systems for agent-mediated interaction represent both parallel or nested sub-communities.

As a consequence, solutions applied to the problems encountered in this environment are validated more severely, against both external and internal criteria. Second, their effects are observable at the level of the virtual community, with a procedure essentially equivalent to agent based simulation and with the related advantages. Third, and moreover, solutions may be not (only) implemented between the agents, but (also) within the agents, what greatly expands the space for modeling.

So far, however, these potentialities have not been fully exploited. Rather than research-based systems for reputation, models have been aimed to ameliorate existing tools, implemented for computerized markets. Agent systems can do much more than this: they can be applied to answer the question as to (a) what type of agent, (b) what type of beliefs, and (c) what type of processes among agents are required to achieve useful social control. More specifically, what type of agent and processes are needed for which result: better efficiency, encourage equity (and hence users' trust), discourage either positive or negative discrimination (or both), foster collaboration at the information level or at the object level (or at both), etc. The solutions proposed are interesting but insufficient attempts to meet the problems left open by online systems. We believe that what is still strongly needed is a cognitive, non-atomized, interraction- and gossip- oriented theory of reputation.

4. Social evaluation and its modification

To overcome the limitations inherent within current systems, we start our proposal with a classification of social evluation, in acccord with the general theory presented in 2. What we need here is a representation of a social evaluation [9] that will allow to represent both communications and personal evaluation.

In most real social situations, there is no way to make social evaluation precise; the main exception being economic science, blessed by the invention of money. To the contrary, most of human social skills are based on imprecise and incomplete data, made even more vague by the tendency to misrepresent frequencies. However, it is hard to deny that the performance of human society is impressive, and we are still fare from understanding it in detail. The purpose of *Repage*, an implemented system currently under experimentation, is that of providing advances both in designing and in understanding how a cognitive approach could be more advantageous than a purely rational one.

To decide how to handle social evaluation, on the basis of the theory presented above, we have to consider them under three different aspects. The first aspect is the *type* of the evaluation - as discussed above, personal experience, image, third party image, shared voice and all the other cognitive constructs have different functional properties, that call for a clear and sharp distinction. In our model. different types of information go through different paths in the cognitive network that represent the memory. This is a sharp distinction; interaction between different types of information is regulated externally. The assumption here is that there will be no intrinsic noise or lack of precision in distinguishing between the types - for example, the agents will not confuse the results of direct experience with related information, nor they will confuse reputation information - what other agents say that "is generally told" - with information having a well defined source - image or third party image.

The second aspect is the *subject* (or *role*) the evaluation concerns. Are we considering our target as a seller or as an informant? In our system, surely we want to keep different aspect separated, at least to some measure. In the example that will consider, we separated evaluation of actual performance (the seller) from evaluation in the field of information (the informer, separately for image and for reputation). This distinction also shows interesting functional properties: changes in the evaluation of somebody as an informer will reverberate on all evaluations the agent has from that source, adjusting their strength accordingly. Again, we treat this distinction sharply - a piece of information can regard an agent as an informer or as a seller, but not both.

The third aspect is the content of the evaluation: is John good or very good? To store the content, we considered the use of a simple number, as in e-Bay style evaluation and as in most reputation systems. This sharp representation, however, is quite unplausible in inter-agent communications, that are one of the central focuses of *Repage*; one is not told that people is saying around that Jane is "0.234 good". While we always identify precisely the type of information communicated and the role discussed, we want to leave some space in the evaluation itself to capture the lack of precision coming (a) from vague utterances, i.e. "I believe that agent t is good, I mean, very good - good, that is", and (b) from noise in the communication or in the recollection from memory. For these reasons, we decided to model the actual value of an evaluation with a fuzzy number, represented by a tuple of positive real values summing to one. These values express the membership of the evaluation to a rating scale. For this version of the model, we tried with five different levels, ranging from *very bad* to *very good*. Moreover, we add to the number a value indicating the strength of belief in the evaluation.

Having decided what kind of fuzzy number we want to use, we need to define carefully how to operate on them by weighting, aggregating, and comparing. For all these issues we propose solutions (mostly standard ones) from the literature [14], with special care to find out solutions that - coherently with the interpretation of uncertainty in this representation - leave unchanged number when aggregated with complete uncertainity, that is, with a completely flat fuzzy number. The reader is referred to [12] for details.

5. *Repage* architecture and implementation

In *Repage* implementation, the agent's memory is defined as a graph of predicates connected by their relations. To reflect their dependencies, the predicates in the *Repage* memory are conceptually organized in different levels and inter-connected. Predicates contain a fuzzy evaluation, belonging to one of the main types (image, reputation, shared voice, shared evaluation), or to one of the types used for their calculation; these include valued info, evaluation related from informers, and outcomes. Special-purpose predicates, dependent on the application domain (for example, a contract not yet fulfilled), exist in the lower layer; they do not necessarily contain an evaluation. Each predicate (except the special purpose ones) has a role and a target; for example, an image of an agent (target) as informer (role). It also has a *strength* value associated to it, in most cases function of the strength of its antecedents.

The network of dependencies specifies which predicates contribute to the values of other predicates. Each predicate in the *Repage* memory has a set of antecedents and a set of consequents. If an antecedent changes its value or it is

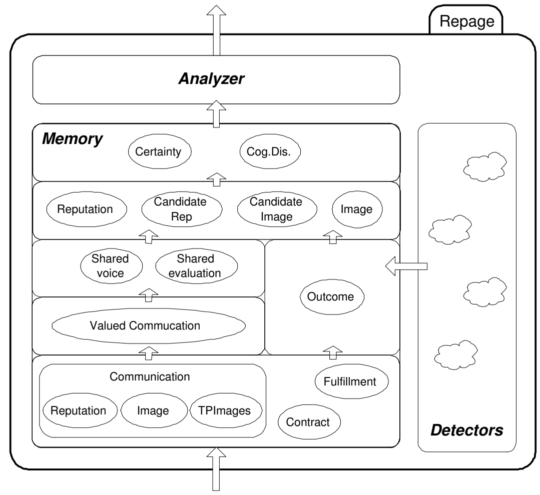

Figure 1. *Repage* **architecture**

removed, the predicate is notified. Then the predicate recalculates its value and notifies the change to its consequents.

The *Repage* architecture, as shown in figure 1, is composed of three main elements: besides the memory, there is a set of components called *detectors* and the *analyzer*.

The *detectors* are inference units specialized in certain predicates. They populate the *Repage* memory (and consequently the main memory of the agent) with new predicates inferred from those already in the memory.

A few examples from current implementation will be shown here. The reference domain is virtual market. Of course, other examples, such as barter, the search for mates, and the labor market are by far more inspiring. But reputation is currently used in electronic marketplaces, and it is this application that needs to be ameliorated. At first level three types of information are needed: *contracts*, not necessarily formal - can be just an agreement between two agents; *fulfillments*, the results of contracts with a quality evaluation; *informers' communication*, coming from third party agents (informers). This information can be related to three different aspects: the image that the informer has of the target, the image that according to the informer other agents have of the target (third party image) and finally the reputation of the target, that will contribute to the building of a shared voice.

Two *detectors* work at this level. One can infer a new *outcome* from a contract and its fulfillment (also considering *how* the contract was fulfilled); the other, given a certain communication, generates a *valued information*, weighted according to the reliability of the informant, that is, the *image* of the informer.

This lead us to the next conceptual level. In this level we find two predicates: *shared voice* and *shared evaluation*.

A shared voice is the main element to build a reputation. It is build from communicated reputation from third party agents. On the other hand we have the shared evaluation that is build from communicated images. The shared evaluation together with the outcomes are themain elements to build images.

In the fourth level there are five types of predicates: *Candidate Image*, *Candidate Reputation*, *Image*, *Reputation* and *Confirmation*. As the names indicate, *candidate images* and *candidate reputations* still do not have enough support to become real images and reputations (either because elements contributing to them are insufficient, or because information is inconsistent). There is a specialized *detector* for each type. Once a *candidate image* or a *candidate reputation* reaches a certain level of strength it becomes a full *image/reputation*. The idea behind the last predicate, *Confirmation*, is that it mirrors how good a previous information was. A communication, the truth value of which is known to the recipient, feedbacks on the *image/candidate image* of the information sender as an informer. The *Confirmation* is similar to an outcome where the contract is the communication provided by the informer and the fulfillment is the *image* the agent has about the target.

In the last level, we find the last two predicates: *cognitive dissonance* and *certainty*. A *cognitive dissonance* is a contradiction between two pieces of information that are relevant for the individual and refer to the same target; it generates an instability in the mind of the individual. Depending on how strong and relevant the *cognitive dissonance* is, the individual is pushed to take special actions for solving it. Although these actions are context dependent, they are always oriented to confirm the grounds of the elements that are causing the *cognitive dissonance*. On the other side, a *certainty* predicate implies full reliance.

6. The analyzer

The main task of the analyzer is to propose actions that (i) can improve the accuracy of the predicates in the *Repage* memory and (ii) can solve *cognitive dissonances* trying to produce a situation of *certainty*. To support decision making, *Repage*'s analyzer offers to the agent employing it a functionality modeled over a cognitive model of hypotetical short-term reasoning. The system will accept a request for advice about the clarification of the evaluation of an agent *target* in a specified *role* (i.e. seller, buyer, informer). On such a request, the system will analyse different hypoteses on the result of specific possible actions, taking into account all the actions with a potential to provide a modification in the informational state about the target in the cosidered role. In the currrent implementation, this amount to deciding whether it is better to try and get another *outcome* with the target (direct experience) or asking around to other agents; in the second case, all potential informants' informational power will be compared, including new, currently unevaluted, informants.

The agent cannot know in advance what kind of information will result from its actions - that is, if an outcome will be positive or negative, if an informant will support the target or judge it untrustworthy. As a consequence, evaluations should be made in the two possible directions (we are presuming here that neutral information or a neutral outcome will not affect the current state in a substantial way). Thus, in all cases mentioned above, two different hypoteses (good result, bad result) will be simulated and the knowledge added will be considered as the result of the two contributions.

Technically, what will happen is that for each possible action (again, interacting so to create an outcome or asking around to specific informers), two working (deep) copies of the current memory will be istantiated, to which the hypotetical new information (good/bad) will be added. This will cause a simulation of the effects that this addition would cause to the current state. These effects will then be evaluated about change in the image of the target in the specified role, be measuring (in fuzzy terms) the distance between the image before and after the simulation. Results for good/bad situations will be aggregated to form an evaluation of the informtative potential of the action under exam. Due to the modular nature of the project, the Analyzer will not try to combine informational value with other costs and benefits of the action under evalutation; the decision making, possibly including comparison with the results of other specialised modules, is left to the wrapper agent architecture. At this point, the system will simply order the possible action by their informative potential; this ordering, obtained by simulation of memory modification, will be a *situated* ordering, depending in a deep way from the current memory state - this is a definite advantage with respect to unsituated orderings or abstract biases. Moreover, internal simulation will automatically take into account any variation in the state of the agent, for example due to learning, imitation, or goal modification. To obtain the same result with a rule-based system, one should take into account explicitly all such potential modifications, a task that will rapidly get unmanageable for any non-toy system; but, even more important, one will lose any modularity - modification in the system will need a rewrite of the rules, while the Analyzer as a simulator would implicitly take modifications into account.

6.1. The analyzer in execution

In this section, we will present two example situations that illustrate how the Analyzer works. The authors are carrying on the development work as a Sourceforge

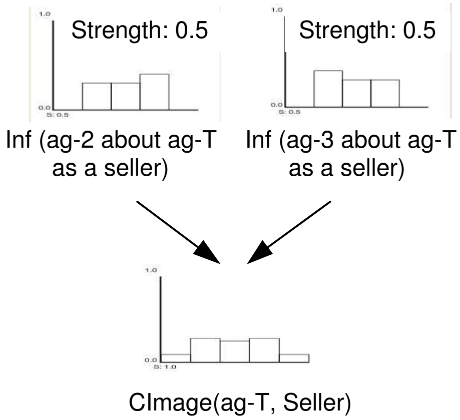

Figure 2. Initial status; "flat" fuzzy values are not enough to decide.

project; the technically interested reader can download the version used to produce the examples by public CVS at cvs.sourceforge.net:/cvsroot/repage/papers/IAT2005.

The base scenario is an eMarket populated by several agents: the agent current evaluator (*ag-0*), the agent that is being evaluated as a potential partner (target: *ag-T*) and several informants.

In this market, ag-0 wants to buy a certain product that is sold by ag-T. The only thing *ag-0* knows about *ag-T* as a seller is a couple of image informations coming from two unknown agents. However, as shown in figure2, these informations are not very relevant and the resulting candidate image is not enough to take a decision about *ag-T*.

At this point *ag-0* has two options to improve the knowledge about *ag-T*: (i) to try to interact directly with it or (ii) to ask new opinions about *ag-T* as a seller to other agents.

6.1.1 Situation 1

In this first situation, ag-0 has knows no trusted agents. To help with this decision, the analyzer simulates two options. First, what if *ag-0* interacts with *ag-T*? The analyzer considers the possibility that the interaction goes well (that is, a good outcome would be generated) and that the interaction fails (that is, a bad outcome would be generated). For each possibility, the distance between the candidate image of *ag-T* value before and after the hypothetical interaction is calculated and aggregated. This distance value is a measure of how much the hypothetical information is moving the candidate image value either toward the good or the bad extremes. Therefore, because what we want is to solve a situation of uncertainty, as bigger this value the better.

Second, what if *ag-0* asks for information to another agent? Here, the first thing the Analyzer has to do is to identify the possible informants. For each informant, the analyzer simulates an information coming from it again considering two alternatives, positive and negative. A distance value is calculated but this time for each possible informant. In this situation, *ag-0* do not have images or reputations of agents as informants so only the case of an unknown informant is considered.

All this process is summarized in section a) of figure3. The first column (Outcome Sim.) corresponds to the simulation of a direct interaction with *ag-T*. The figures of fuzzy numbers show the hypothetical candidate image of ag-T after a successful (top figure) or a failed (bottom figure) interaction. Similarly the second column shows the hypothetical candidate image of *ag-T* after a positive (top figure) or negative (bottom figure) information about *ag-T* as a seller coming from an unknown informant. In this case, the action contributing more to clarify the image about *ag-T* as a seller is the direct interaction with this agent.

6.1.2 Situation 2

Let's suppose now that *ag-0* has received also four images from different agents saying *ag-1* is a very good informant. Thanks to these informations, *ag-0* builds a positive image of *ag-1* as an informant, that increases the consideration of its communications.

If now we reproduce the steps the Analyzer has followed before to see which is the best action, the picture is quite different. As shown in section b) of figure3 now the option of asking to another agent (specifically the option of asking to *ag-1*) is better than the direct interaction. This is due to the extra value of asking to an agent (*ag-1*) that has a good image as an informant. Notice that the Analizer in this situation also evaluates the possibility of asking to an unknown agent like before, and would do the same for each possible informant if there were more.

7. Conclusions and future work

In this paper, the *Repage* system was presented as a tool for integrating image and reputation information in partner selection among autonomous but socially adaptable agents. The social cognitive theory on which the system is based was briefly described, and the system's architecture illustrated. In particular, we presented the Analyzer module that tries to solve uncertain situation by simulating the potential informative gain that would be caused by the agent's actions.

In the future, we plan several developments of the basic architecture. A direction that looks promising is that of learning in all the memory manipulation modules. Furthermore, artificial simulative experiments comparing *Repage* with a mere image-based systems ought to be carried out in

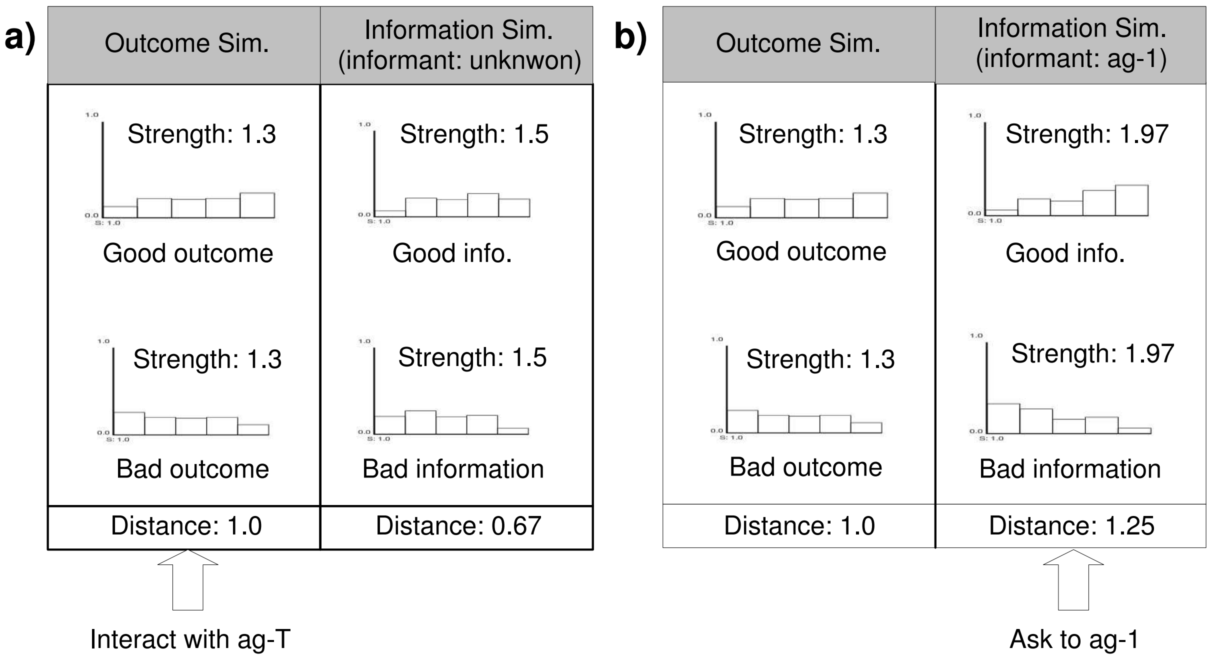

Figure 3. The simulations of the Analyzer

different domains of applications, including both cooperative settings like organisations and electronic marketplaces. Finally, future developments of *Repage* ought to concern its integration with other components of a social agent, with special reference to learning and social adaptation, on one hand, and personalised inclinations and biases on the other.

References

[1] G. Bolton, E.Katok, and A. Ockenfels. How effective are online reputation mechanisms? an experimental investigation. Technical report, Max-Planck-Institut, 2002. http://www.mpiew-jena.mpg.de/esi/discussionpapers/.

[2] J. Carbo, J. Molina, and J. Davila. Comparing predictions of sporas vs. a fuzzy reputation agent system. In *3rd International Conference on Fuzzy Sets and Fuzzy Systems, Interlaken*, pages 147—153, 2002.

[3] R. Conte and M. Paolucci. *Reputation in artificial societies: Social beliefs for social order*. Kluwer Academic Publishers, 2002.

[4] R. Conte and M. Paolucci. Social cognitive factors of unfair ratings in reputation reporting systems. In *Proceedings of the IEEE/WIC International Conference on Web Intelligence - WI 2003*, pages 316—322, 2003.

[5] D.Huynh, N. Jennings, and N. Shadbolt. Developing an integrated trust and reputation model for open multi-agent systems. In *Proceedings of the Workshop on Trust in Agent Societies (AAMAS-04), New York, USA*, pages 65—74, 2004.

[6] R. Dunbar. *Grooming, Gossip, and the Evolution of Language*. Harvard Univ Press, October 1998.

[7] J. Grunig and C. Hung. The effect of relationships on reputation and reputation on relationships: A cognitive, behavioral study, 2002. Ppresented at the PRSA Educator's Academy 5th Annual International, Interdisciplinary Public Relations Research Conference.

[8] D. Kreps and R. Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1982.

[9] M. Miceli and C. Castelfranchi. *Human cognition and agent technology*, chapter The Role of Evaluation in Cognition and Social Interaction. Amsterdam:Benjamins, 2000.

[10] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system, 2001. Working Paper for the NBER Workshop on Empirical Studies of Electronic Commerce.

[11] J. Sabater, M. Paolucci, and R. Conte. Repage. In *Proceedings of the workshop RAS05 - Reputation in Agent Societies as part of 2004 IEEE/WIC/ACM International Joint Conference on Intelligent Agent Technology (IAT'04) and Web Intelligence (WI'04)*, Beijing, 2005.

[12] J. Sabater, M. Paolucci, and R. Conte. Repage: Reputation and image among limited autonomous partners. Submitted to JASSS, 2005.

[13] J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings AAMAS-02, Bologna, Italy*, pages 475—482, 2002.

[14] R. R. Yager. On the determination of strength of belief for decision support under uncertainity–part ii:. *Fuzzy Sets and Systems*, 142(1):129–142, 2004.

[15] G. Zacharia. Collaborative reputation mechanisms for online communities. Master's thesis, Massachusetts Institute of Technology, September 1999.