Evaluating the ReGreT system

Jordi Sabater*
IIIA - Artificial Intelligence Research Institute
CSIC - Spanish Scientific Research Council
08193 - Bellaterra, Barcelona, Spain
{jsabater}@iiia.csic.es

Abstract

The ReGreT system is a trust and reputation mechanism that uses, among other things, social information to improve the calculation of trust and reputation measures. Using a framework called *SuppWorld* designed to test these kind of complex models, we present a set of experiments that evaluate different features of the ReGreT system in several scenarios.

1 Introduction

The scientific research in the area of trust and reputation models for virtual societies is a recent discipline oriented to increase the reliability and performance of electronic communities by introducing in such communities these well known human social control mechanisms.

Up to now, computational trust and reputation models have been considering two different information sources: (i) the direct interactions among agents and (ii) the information provided by members of the society about experiences they had in the past [5, 8, 9, 10]. Those systems, however, forget a third source of information that can be very useful: the social relations. As a direct consequence of the interactions, it is possible (even in not too complex societies) to identify different types of social relations between society members. Sociologists and psychologists have been studying these social networks in human societies for a long time and also how these social networks can be used to analyze trust and reputation. These studies show that it is possible to say a lot about the behaviour of individuals using the information obtained from the analysis of their social network.

The ReGreT system [6, 7] is a trust and reputation mechanism that has been designed to take advantage of these social aspects. ReGreT is intended to be used as a module that extends the capabilities of the agent, adding the

*Jordi Sabater enjoys a Marie Curie Intra-European Fellowship (European 6th framework program) contract No MEIF-CT-2003-500573.

possibility to deal with reputation and trust to improve the agent's behaviour in an e-commerce environment.

In this article we present a initial set of experiments we have performed in order to test some of the capabilities of the ReGreT system. We have designed and implemented a framework we call the *"SuppWorld"* [7] that allows to build scenarios complex enough to exploit the features of systems like ReGreT.

In section 2 we briefly introduce the ReGreT system. In section 3 we present the essential elements of the *SuppWorld* framework. Finally, in section 4 we describe the performed experiments, the results and some conclusions.

2 The ReGreT system in a nutshell

ReGreT [6, 7] is a modular trust and reputation model oriented to complex e-commerce environments where social relations play an important role. You can find a full description of the system in [6, 7]. In this section we will describe only the essential elements to understand the experiments presented in section 4.

The main characteristics of ReGreT are:

- It takes into account direct experiences, information from third party agents and social structures to calculate trust, reputation and credibility values.
- It has a trust model based on direct experiences and reputation.
- It incorporates an advanced reputation model that works with transmitted and social knowledge.
- It has a credibility module to evaluate the truthfulness of information received from third party agents.
- It uses social network analysis to improve the knowledge about the surrounding society (specially when no direct experiences are available).
- It provides a degree of reliability for the trust, reputation and credibility values that helps the agent to decide if it is sensible or not to use them in the agent's decision making process.
- It can adapt to situations of partial information and improve gradually its accuracy when new information becomes available.
- It can manage at the same time different trust and reputation values associated to different behavioural aspects. Also it can combine reputation and trust values linked to simple aspects in order to calculate values associated to more complex attributes.

Some people could argue that current e-commerce scenarios are not complex enough to justify the degree of complexity modelled by the ReGreT system. Although the modular design of the ReGreT system makes possible its use in a

wide range of environments, it is true that in order to fully exploit its capabilities (specially those related with social network analysis) it is necessary a certain degree of complexity in the society. We agree that this complexity is not present in current operative e-commerce environments but we take the stance that in the near future, as the complexity of tasks to be performed by agents will increase, these kind of complex scenarios will become usual.

Figure 1 shows a panoramic view of the ReGreT system. The system maintains three knowledge bases. The outcomes data base (ODB) to store previous contracts and their result; the information data base (IDB), that is used as a container for the information received from other partners and finally the sociograms data base (SDB) to store the sociograms (graphs that represent social relations) that define the agent social view of the world. These data bases feed the different modules of the system.

The *direct trust* module deals with direct experiences and how these experiences can contribute to the trust on third party agents. Together with the reputation model they are the basis for the trust model.

The reputation model is divided in three specialized types of reputation depending on the information source that is used to calculate them. If the reputation is calculated from the information coming from witnesses we talk about the *witness reputation*, if the reputation is calculated using the information extracted from the social relations between partners we are talking about the *neighbourhood reputation*. Finally, reputation based on roles and general properties is modelled by the *system reputation*.

The system also incorporates a credibility module that allows the agent to measure the reliability of witnesses and their information. This module is extensively used in the calculation of *witness reputation*.

All these modules work together to offer a complete trust model based on direct knowledge and reputation. However, the modular approach in the design of the system allows the agent to decide which parts it wants to use. For instance, the agent can decide not to use *neighbourhood reputation* to calculate a reputation value or rely only on *direct trust* to calculate the trust on an agent without using the reputation module.

Another advantage of this modular approach is the adaptability that the system has to different degrees of knowledge. The system is operative even when the agent is a newcomer and it has an important lack of information. As long as the agent increases its knowledge about the other members of the community and the social relations between them, the system starts using other modules to improve the accuracy of the trust and reputation values. This allows the system to be used in a wide range of scenarios, from the most simple to the most complex. If the information is available, the system will use it.

In the ReGreT system, each trust and reputation value has an associated reliability measure. This measure tells the agent how confident the system is on that value according to how it has been calculated. Thanks to this measure, the agent can decide, for example, if it is sensible or not to use the trust and reputation values as part of the decision making mechanism.

The last element in the ReGreT system is the *ontological structure*. We

consider that trust and reputation are not single and abstract concepts but rather multi-facet concepts. The *ontological structure* provides the necessary information to combine reputation and trust values linked to simple aspects in order to calculate values associated to more complex attributes. For example, the reputation of being a good flying company summarizes the reputation of having good planes, the reputation of never losing luggage and the reputation of serving good food. In turn, the reputation of having good planes is a summary of the reputation of having a good maintenance service and the reputation of frequently renewing the fleet. Note that each individual can have a different *ontological structure* to combine trust and reputation values and a different way to weigh the importance of these values when they are combined.

3 The SuppWorld framework

There are several frameworks that are used to test computational trust and reputation models [3, 8, 4, 2, 1]. However, none of these frameworks provides a scenario rich enough to test all the dimensions of the ReGreT system. The main drawback in these frameworks is that they are not prepared to naturally support social relations among members. This is because they are designed to test trust and reputation models that do not consider this aspect. Social relations play an essential role in the ReGreT system. Therefore, we decided to design and implement a new framework adapted to the special characteristics of systems like ReGreT. In this section we briefly present *SuppWorld* [7], the framework we have used to test the ReGreT system and perform the experiments detailed in section 4. This framework allows us to build simple scenarios oriented to test the basic aspects of trust and reputation models and also complex scenarios where social relations acquire a great relevance.

3.1 The framework

The *SuppWorld* framework is built around the idea of a supply chain. In a typical supply chain, agents trade by buying one level below in the chain, adding value to the purchased goods, and selling the manufactured goods up to the next level in the chain. The *SuppWorld* allows the design of scenarios based on this structure and oriented to the study of negotiation, trust and reputation models.

The activity in a *SuppWorld* scenario is organized in several scenes. The heartbeat of a *SuppWorld* scene is measured in 'ticks'. All the activities performed during a tick, from the point of view of the simulation, are supposed to run in parallel. Each scene is active during a fixed amount of ticks. The scenes can be combined to recreate different supply chain configurations. A 'round' is defined as a cycle in the execution of the scenes that compound a given scenario.

The possible scenes in a *SuppWorld* scenario are:

- Markets. As the name suggests, in these scenes is where agents trade. A market is represented as a toroidal grid where each cell is owned by a

single seller. While sellers stay always in its cell during a market session, buyers can move freely from one cell to another. Agreements between sellers and buyers materialize after a negotiation process. It is possible to have more than one buyer in a cell. In that case, only one buyer can negotiate with the seller and the others have to wait in a queue. While they are waiting, they can exchange information about the reputation of other sellers. The exchange of information is performed in turns and each agent has a limited number of opportunities to query the others.

- Conventions. Here an agent can exchange information with other agents of the same type. As in the waiting queue in a cell, the exchange of information is performed in turns and each agent has a limited number of opportunities (equal to the number of ticks that is running the convention scene) to query the others.

- Entrance of rough material. A supply chain is not a closed environment. It is necessary to allow the entrance of rough material that, properly transformed, will follow its path up to the final consumers.

- Production process. In each link of the supply chain there is a process of transformation that adds new value to the goods. This process is simulated in this scene.

- Entrance of money. Similarly to the entrance of rough material, it is necessary at least one point to introduce money to the supply chain. With this scene you can simulate the action of final consumers or a "pay day" for these final consumers if they are explicitly represented.

3.2 The agents behaviour

There are two elements in a *SuppWorld* scenario that determine the behaviour of an agent.

The first is what we call the *alignment* of the agent. The *alignment* defines the basic behaviour of the agent in aspects like how contracts will be fulfilled or the truthfulness of the information given to other members of the society.

The second are the social relations between the agent and the other members of the society. In a *SuppWorld* scenario we consider three types of social relations among their members:

- *Competition*. This is the type of relation found between two agents that pursue the same goals and need the same (usually scarce) resources. In this situation, agents tend to use all the available mechanisms to take some advantage over their competitors, for instance hiding information or lying.

- *Cooperation*. This relation implies significant exchange of sincere information between the agents and some kind of predisposition to help each other if possible. Notice that we talk about "sincere" information instead

of "true" information because the agent who gives the information *believes* it is true. We consider that two agents cannot have at the same time a competitive and a cooperative relation. This is also the relation type that identifies groups of agents.

- *Trade.* This type of relation reflects the existence of commercial transactions between two agents and is compatible either with cooperative or competitive relations. For the moment this is the only social relation that agents can identify by themselves in a *SuppWorld* scenario.

The social relations can change the basic behaviour defined by the *alignment*. For instance, an agent that normally provides false information to the others, will tell the truth to an agent that has a cooperative relation with it.

4 Experiments

In this section we will present a set of experiments to show how the ReGreT system behaves in different situations. We will present several scenarios to experiment with different parts of the system. The experiments are designed using the *SuppWorld* framework described in the previous section.

The set of experiments can be divided in two main blocks that correspond to two different scenarios. In the first scenario we consider a society with no other social relations among agents that trade relations. This scenario will allow us to test and show how the direct trust and the witness reputation work. In the second scenario we add cooperative and competitive relations among the members of the society and we analyze the ReGreT social credibility module.

4.1 The common framework

Although as we have said there are two different scenarios, both share a common structure. In this section we describe this structure.

The *Suppworld* base scenario used in all the experiments is composed by two grid markets. The first grid is the home of "producers" and the second grid the home of "manufacturers". The manufacturers buy products in the first market and then sell them in the second market. The buyers in the second market are simulated by a single entity that always buy all the product available from manufacturers. Figure 2 shows the sequence of events in the base scenario.

The first step is the entrance of rough material in the producers storage facilities. Then, during the production process scene the producers generate the product that will be sold to manufacturers in the market scene. Before starting the market scene, there is a convention of manufacturers that allow them to exchange information about producers. The market scene is the central event. Here the manufacturers negotiate with the producers to buy products. The round finishes with a production scene for manufacturers and the sale of the manufacturer product to the entity that simulates final consumers.

We have made some simplifications to the scenario in order to make the analysis of the results simpler and clearer.

Because we want to focus our analysis on the manufacturers, producers are designed to never lose money and to have always product in stock. Therefore, it never happens that a manufacturer cannot buy because the producer does not have products to sell. To ensure that, we give to the producers enough storage capacity to satisfy the demand, renewing this capacity each round without taking into account their cash.

Producers only sell one type of product. The presence of more than one product implies that manufacturers need some kind of strategy to adjust the amount of product of each type if they do not want to lose money. This strategy is very relevant for the final success of a buyer in a *Suppworld* scenario and, at this stage, complicates the analysis of trust and reputation mechanisms unnecessarily.

The last simplification is in the negotiation process. In a *SuppWorld* scenario, a contract has four issues: price, quantity, quality and transport type. Each issue has a minimum and maximum value that is fixed by the engineer for each experiment. This minimum and maximum values are used to ensure that agents are always negotiating within a controlled range for each issue. Since presently we do not want negotiation to interfere with the analysis of trust and reputation, all producers and manufacturers use the same negotiation engine with the same parameters. These parameters are adjusted to obtain the agreement exactly when the value for each issue is in the middle of its range. In our example this happens when $Price = 30$, $Quantity = 10$, $Quality = 3$ and $Transport\_type = 3$. This specific point in the negotiation has been chosen to confer symmetry on the definition of good and bad behaviours when the agents fulfill a contract and it is not relevant in the analysis of trust and reputation.

Once the producer and the manufacturer arrive to an agreement, producer has to send the product to the manufacturer. At this point, producers can decide not to fulfill what was agreed. To simplify, the quantity and the transport type are fixed. Therefore, the producer can only influence the price and the quality. This fact fixes another parameter that is common for all the experiments: the ontological structure of the buyers. This ontological structure is shown in figure 3 and defines a good seller as a seller that offers good prices and good quality, giving the same relevance to both aspects.

We differentiate five types of producers according to their behaviour in the fulfillment of contracts (what we have referred in section 3.2 as the "alignment" of the agent):

- SAINT: 30% of the times favours the partner by decreasing the price to the minimum and increasing the quality to the maximum.
- GOOD: 30% of the times favours the partner by decreasing the price and increasing the quality a quantity equal to 1/4 of the issue's range.
- NEUTRAL: Always fulfill the contracts as agreed.

- BAD: 75% of the times cheats the partner by increasing the price and decreasing the quality a quantity equal to 1/4 of the issue's range.
- EVIL: 75% of the times cheats the partner by increasing the price to the maximum and decreasing the quality to the minimum.

When they are not cheating/favouring the partner, all the producers fulfill the contracts as agreed.

Given that the agreement contract is fixed, it is possible to characterize the alignment of producers using the contract they send when cheat or favour the partner. Table 4.4 shows the relation between the alignment of a producer and the fulfillment of a contract (remember that we do not allow producers to modify the values for *Quantity* and *Transport_type*).

The position of manufacturers in the producers' grid at the beginning of each round is decided randomly. The movement of manufacturers over the grid is also random.

To summarize, the base scenario for our experiments is a market grid populated by a fixed amount of sellers (producers) that offer all of them the same kind of product. At the beginning of each market session, a set of buyers (manufacturers) are distributed randomly over that grid. During the market session they move around randomly and buy things. However, sellers not always fulfill contracts as agreed. Manufacturers have only trust and reputation mechanisms (the ReGreT system) to fight against cheaters. At the end of each round, manufacturers sell all the product to final consumers (simulated by a single entity). Depending on how each contract was fulfilled they will win or lose money in that transaction.

With the restrictions we have imposed, you can see the interaction among producers and manufacturers as a Prisoner's dilemma game where only producers decide if they want to cooperate (fulfill the contract as was agreed) or defect (cheat the partner by applying worse conditions in the fulfillment). Here, however, the producers can choose among different degrees of cooperation and defection.

4.2 Scenario I: Direct trust and witness reputation

The objective of these experiments is to see how an agent can improve its performance by using direct trust and witness reputation (without the credibility module). The performance of an agent is measured by the amount of cash that has won (or lost) after a fixed number of rounds.

We have compared four types of manufacturers:

- Manufacturers that always negotiate with producers. (Negotiate always)
- Manufacturers that use direct trust to decide if it is worth it or not to negotiate. (DT)
- Manufacturers that use direct trust and witness reputation. (DT + WR)

- Manufacturers that use direct trust and witness reputation but always provide wrong information to other manufacturers. (DT + WR (LIERS))

The first three types are tested in an environment where manufacturers always say the truth and do not try to lie the others. The fourth type, however, reproduces a situation where manufacturers always give wrong information.

We performed an initial set of experiments where the full population of manufacturers had, in turn, each one of the profiles commented before. We were expecting an increase of performance for the first three types of manufacturers but when we arrived at the manufacturers that were using witness reputation, we found that their performance was worse than the performance of manufacturers that were using only direct trust.

The reason for that is the nature of witness reputation. Suppose the situation where all the manufacturers start from scratch and witnesses only give information based on their own direct experiences. Even assuming the witnesses always tell the truth, a manufacturer only can provide reliable information after several rounds. For that time, however, all the other manufacturers have had a similar number of direct experiences and witness information is useless. In this scenario, witness reputation is redundant. Moreover, if producers do not have a fixed behaviour (that is, bad agents not always behave badly) you can have witnesses that are distributing wrong information. In that case, the use of witness information is self-defeating. This is what was happening in our experiments.

As you can imagine, the situation is even worse if witnesses also spread opinions based on reputation.

The conclusion is that witness reputation is only useful if you can guarantee that there is a pool of individuals that can provide well founded information. It is not worth it (and can be even self-defeating) the use of witness reputation in those situations where there is a general lack of knowledge.

Knowing that, we repeated the experiments but this time using the following procedure. All manufacturers use only direct trust for a number of rounds until we are sure that the general knowledge about the market among manufacturers is good enough. Then we selected a set of manufacturers and initialize their internal status (as if they were just arriving at the market), changing the parameters of the ReGreT system according to the characteristics we want to study. This is the moment when the experiment really starts. We monitor the performance of these selected set of manufacturers during a fixed number of rounds.

The main parameters of the experiments are the following:

- 16 producers (4x4 grid).
- 64 manufacturers (8x8 grid).
- 50 rounds. Each round has 30 ticks.
- 64 manufacturers, 10 are randomly selected to be monitored as explained before.

- 30 rounds preparing the market, 20 rounds monitoring the selected group of manufacturers.
- For each round there is a 1 tick convention of manufacturers.

Maintaining this configuration we made 5 sets of experiments, each set corresponding to a different alignment configuration for producers. Concretely we tried the following alignment configurations:

- 100% of producers are EVIL.
- 50% are BAD, 50% are EVIL.
- 50% are NEUTRAL, 50% are EVIL.
- 100% are BAD.
- 50% NEUTRAL, 30% BAD and 20% EVIL.

We tested each alignment configuration with the different type of manufacturers presented before, that is, manufacturers that always negotiate with producers, manufacturers that use direct trust to decide if it is worth it or not to negotiate, manufacturers that use direct trust and witness reputation and manufacturers that use direct trust and witness reputation but in an environment where witness information is always wrong.

The results of these experiments are shown in figure 4 and figure 5. These plots show, for each alignment configuration, the average performance of the selected manufacturers from the moment they are initialized to the end of the experiment.

This time, for all the alignment configurations, we get the expected result. There is a great improvement in the performance of manufacturers that use direct trust to decide if it is worth it or not to negotiate with a specific producer compared with the manufacturers that always negotiate. As it would be expected, this difference is bigger in hostile environments like the first three but also quite significant in more normal environments like the latest shown in figure 5. Also remarkable is the improvement of performance for those manufacturers that were using witness reputation in a collaborative environment compared with manufacturers that only were using direct trust.

Finally, you can observe that the performance of agents that use witness reputation in an environment where witness information is always wrong is always worse than the performance of manufacturers that use only direct trust. This supports our hypothesis that if there are no guarantees that witness information has a minimum quality it is better to ignore it. These results also validate the results obtained in the experiments with manufacturers that use witness reputation in a collaborative environment.

In order to check if these results can be extended to bigger societies we repeated the last experiment in a bigger scenario. The relevant parameters for this experiment are:

- 64 producers (8x8 grid).
- 256 manufacturers (16x16 grid).
- 90 rounds. Each round has 60 ticks.
- 256 manufacturers, 10 are randomly selected to be monitored as explained before.
- 50 rounds preparing the market, 40 rounds monitoring the selected group of manufacturers.
- For each round there is a 20 ticks convention of manufacturers.

The results are shown in figure 6. As you can observe, the results are very similar to those obtained for the medium size environment.

4.3 Scenario II: Social credibility

In this scenario we will add competitive and cooperative relations among the members of the society. Competitive and cooperative relations in a *SuppWorld* scenario (as in real life) have a direct impact on the behaviour of the agents. Because we want to analyze the social credibility module, we will focus on the impact that cooperative and competitive relations have in the truthfulness of the information given by witnesses.

We want to observe how the use of the social credibility module of the ReGreT system can improve the performance of the manufacturers in such kind of scenarios.

In a *SuppWorld* scenario, competitive and cooperative relations define the probability to deliberately provide wrong information. To compute this probability, an agent takes into account the competitive and cooperative relations among the agent that is making the query (the source), the subject of the query (the target) and the witness (the agent itself). In those cases where there is no cooperative and competitive relation among the source, the target and the witness, the witness computes the probability to cheat according to its alignment.

For this experiment we have compared three types of manufacturers:

- Manufacturers that use direct trust to decide if it is worth it or not to negotiate. (DT)
- Manufacturers that use direct trust and witness reputation. (DT + WR)
- Manufacturers that use direct trust, witness reputation and the social credibility module. (DT + WR + SCr)

We have to say that the manufacturers that use the social credibility module have a perfect knowledge of the society. In a real world, the sociograms usually have some mistakes that would decrease the performance of the module.

The main parameters of the experiments are the following:

- 16 producers (4x4 grid).
- 64 manufacturers (8x8 grid).
- 50 rounds. Each round has 30 ticks.
- 64 manufacturers, 10 are randomly selected to be monitored.
- 20 rounds preparing the market, 30 rounds monitoring the selected group of manufacturers.
- For each round there is a 10 tick convention of manufacturers.

We maintain a fixed configuration of 25% SAINTs and 75% EVILs in the case of producers and a 100% of SAINTs in the case of manufacturers (thinking that the behaviour of the manufacturers is also conditioned by the social relations).

In the first set of experiments shown in figure 7 we can see the performance of the manufacturers given different environments, each one more competitive than the previous. The degree of competition in the environment is measured by the density of the competitive sociogram shared by all the manufacturers. For instance, a density of 0.5 means that a manufacturer has a competitive relation with 50% of the agents (manufacturers and producers) in that society.

It is interesting to see that even with a density of only 0.2, the performance of manufacturers that use witness reputation and direct trust is worse than the performance of those that use only direct trust. As expected, the social credibility module improves considerably the performance of the manufacturers that use witness reputation.

With a density of 0.5 the situation is similar. However the performance of the manufacturers using witness reputation (with or without the social credibility module) is worse than in the previous case (density = 0.2). This happens because by increasing the number of agents with a competitive relation we are decreasing the number of reliable witnesses and, as a consequence, we are decreasing the amount of useful information that can be used by the witness reputation module. In the last experiment we force this situation by fixing a density of 0.8. Here, even using the credibility module, the use of witness reputation is not worth it. Again, as we observed in the *Scenario I*, not always the use of witness reputation implies an improvement in the performance.

We repeated the same set of experiments with the same parameters but this time considering only cooperative relations. Figure 8 shows the obtained results.

It seems that cooperative relations (as they are defined in the *SuppWorld* framework) do not have the same relevance that competitive relations. The performance of the agents that use the social credibility module is very similar in the three experiments. The agents that use witness reputation without the social credibility module experiment a decrease in performance when the number of cooperative relations increases but it is not comparable with the decrease of performance due to competitive relations.

4.4 Conclusions

The aim of the experiments we have presented is to show how the *SuppWorld* framework can be used to test the ReGreT system and, in general, any complex trust and reputation model. These experiments are a small sample of the full set of experiments that are necessary to validate a system with the complexity of ReGreT. However, even knowing they are not complete, it is possible to extract some interesting conclusions.

As we have seen in the experiments of the *Scenario I* not always the use of witness reputation contributes to improve the performance of the agent. There are situations where the use of witness reputation is self-defeating. Something similar happens with the use of social information.

One of the most important factors to be taken into account to decide if it is worth it or not the use of alternative sources of information is the cost to obtain direct experiences. In a scenario where it is easy and cheap to get direct experiences it is difficult that the use of other solutions to compute trust and reputation values compensates the problems associated to them. However, as long as the cost of direct experiences increases, the use of other sources of information like witness information or social knowledge become more and more important. This means that we cannot rely on a static mechanism to combine the different sources of information if we want a generic trust and reputation system we have to go for a dynamic method to combine the different sources of information that can adapt to the characteristics of the environment.

References

[1] J. Carbo, J.M. Molina, and J. Davila. Trust management through fuzzy reputation. *Int. Journal in Cooperative Information Systems*, pages in–press, 2002.

[2] C. Castelfranchi, R. Conte, and M. Paolucci. Normative reputation and the cost of compliance. *Journal of Artificial Societies and Social Simulation (JASSS)*, 1(3), 1998.

[3] S. Marsh. *Formalising Trust as a Computational Concept*. PhD thesis, Department of Mathematics and Computer Science, University of Stirling, 1994.

[4] L. Mui, A. Halberstadt, and M. Mohtashemi. Notions of reputation in multi-agent systems: A review. In *Proceedings of the first international joint conference on autonomous agents and multiagent systems (AAMAS-02), Bologna, Italy*, pages 280—287, 2002.

[5] J. Sabater and C. Sierra. Regret: A reputation model for gregarious societies. In *Proceedings of the Fourth Workshop on Deception, Fraud and Trust in Agent Societies, Montreal, Canada*, pages 61—69, 2001.

[6] J. Sabater and C. Sierra. Reputation and social network analysis in multiagent systems. In *Proceedings of the first international joint conference on autonomous agents and multiagent systems (AAMAS-02), Bologna, Italy*, pages 475—482, 2002.

[7] Jordi Sabater-Mir. *Trust and Reputation for Agent Societies*. PhD thesis, IIIA-CSIC, 2003.

[8] M. Schillo, P. Funk, and M. Rovatsos. Using trust for detecting deceitful agents in artificial societites. *Applied Artificial Intelligence*, (Special Issue on Trust, Deception and Fraud in Agent Societies), 2000.

[9] Bin Yu and Munindar P. Singh. A social mechanism of reputation management in electronic communities. In *Cooperative Information Agents (CIA), Boston, USA*, pages 154—165, 2000.

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

Alignment	Price	Quantity	Quality	Trans.Type
Agreement	30	10	3	3
Saint	10	10	5	3
Good	20	10	4	3
Neutral	30	10	3	3
Bad	40	10	2	3
Evil	50	10	1	3

Table 1: Relation between the alignment and the fulfillment of a contract.

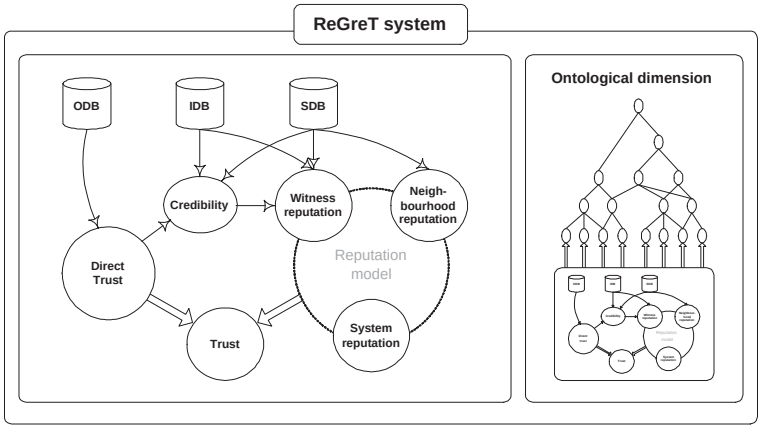

Figure 1: The ReGreT system.

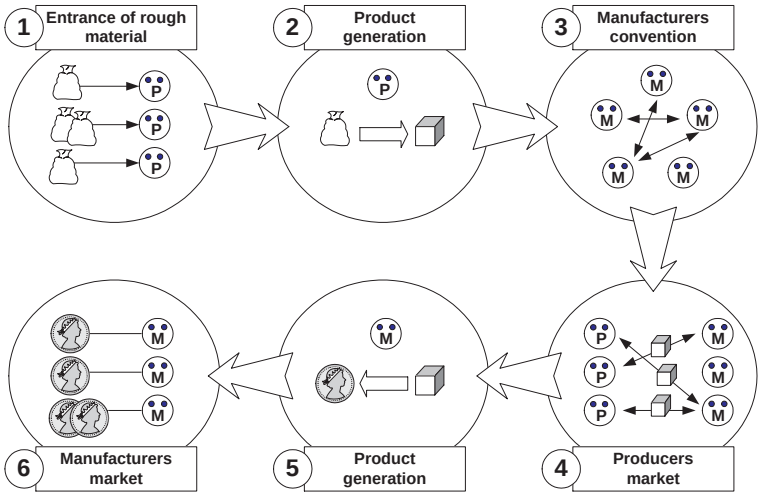

Figure 2: *SuppWorld* base scenario.

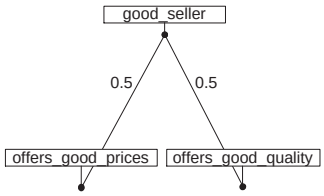

Figure 3: Ontological structure for a buyer.

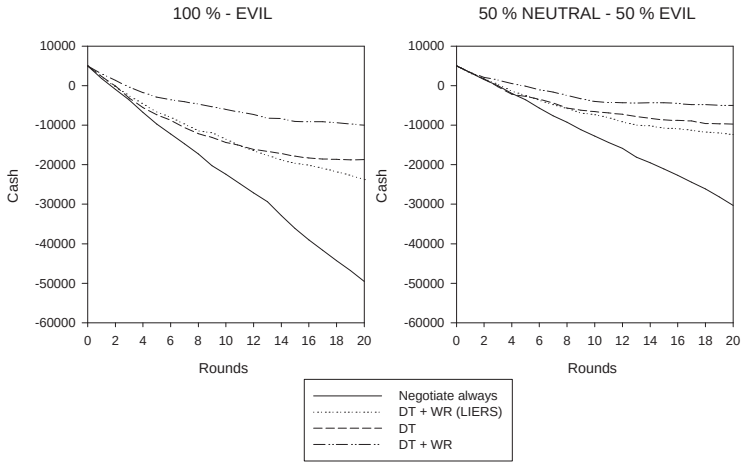

Figure 4: DT and WR experiments (I).

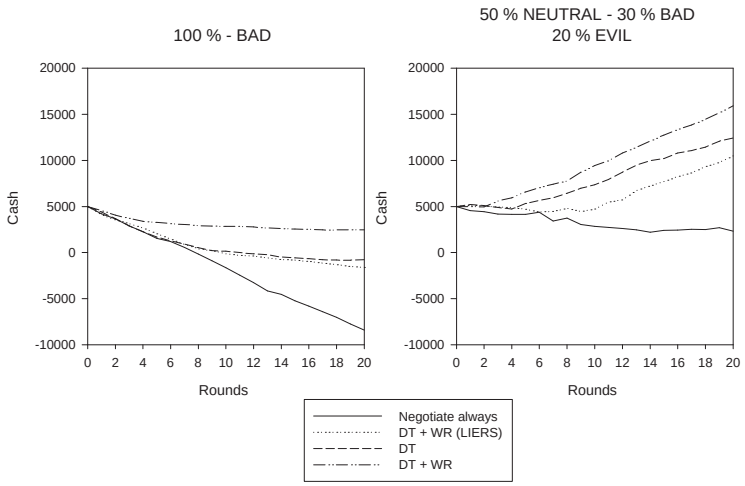

Figure 5: DT and WR experiments (II).

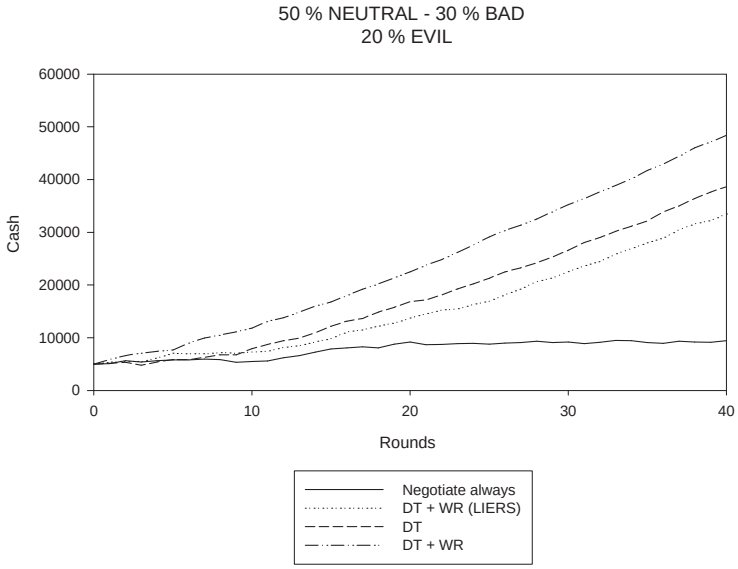

Figure 6: Large scenario.

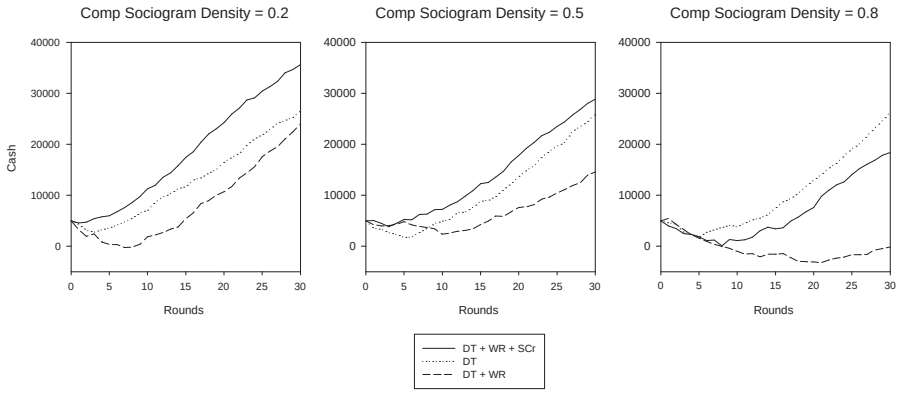

Figure 7: Social credibility and competitive relations.

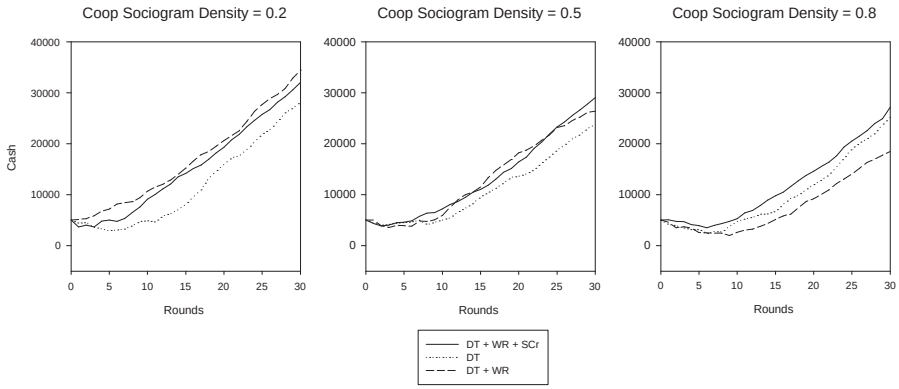

Figure 8: Social credibility and cooperative relations.