

Trust-Based Selection of Partners

Joana Urbano, Ana Paula Rocha, and Eugénio Oliveira

LIACC / Departamento de Engenharia Informática
Faculdade de Engenharia, Universidade do Porto
Rua Dr. Roberto Frias, 4200-465 Porto, Portugal,
{joana.urbano, arocha, eco}@fe.up.pt

Abstract. The community of multi-agent systems has been studying ways to improve the selection of partner agents for joint action. One of such approaches consists in estimating the trustworthiness of potential partners in order to decrease the risk inherent to interacting with unknown agents. In this paper, we study the effect of using trust in the process of selecting partners in electronic business. We empirically evaluate and compare different trust-based selection methods, which either use trust in a preselection phase previous to the negotiation, in the negotiation process, or in both of these stages. We here briefly introduce a computational model of trust that uses a simple machine learning mechanism to dynamically derive the expected tendencies of behavior of potential candidate partner agents. The results obtained in our comparison study allow us to point to the best trust-based selecting methods to use in specific situations.

Keywords: Computational trust, selection of partners, multi-agent systems.

1 Introduction

Trust is an important area of research in several disciplines, including sociology, psychology, philosophy, economics, distributed systems and distributed artificial intelligence. In fact, some authors considers it a public good that enables production and exchange and that is vital for the survival of the society [1].

The research area of multi-agent systems has been proposing computational trust models that help recognizing the different social behaviors of the communities in artificial societies. These models are considered essential for making more informed decisions in these societies of agents, reducing the risk associated to the information asymmetry problem in open and dynamic environments.

Some of these models provide probabilistic or heuristic-based aggregation engines that compute the trustworthiness of the agents in evaluation based on the available evidences on these agents [2,3,4]. Other models are proposed with the aim of being resistant to attacks, such as fraud, badmouthing, collusion, and other forms of deceptive behavior [5]. Still other computational models propose to incorporate in their trust reasoning important concepts imported from the

social sciences area, such as forgiveness, prejudice, asymmetry, regret, erosion and coherence (e.g. [6][7][8][9]).

Independently of the original purpose, the majority of these proposals is based in the strong assumption that the use of computational trust mechanisms to select partners enhances the decision process and gives higher values of utility to the selecting agent. Moreover, these models are empirically evaluated in scenarios where service customers seek the best provider of services, using, in their selection process, no other differentiating factor than the estimated trustworthiness of the candidate agents.

However, there may exist specific real-world situations where the most trustworthy agent is not the one that offers the best payoff to the selecting agent. Let us consider two hypothetical examples. In the first example, firm A is a manufacturer of t-shirts and firms B and C are providers of fabric. Firm A knows, from experience, that B rarely fails a contract. In the same way, it also knows that C is less reliable, and sometimes it delays a delivery; however, firm C offers better utility (possibly derived from better quality of the product or better shipment and payment conditions) than firm B when it does not breach the contracts. In this case, the fact that B is more trustworthy than C can mean that B is more useful to A than firm C?

The second example depicts a recruitment scenario and is related to the use of trust in the selection decision as a *prefiltering* activity. In the example, firm D has one position open for Java programmers for which it has received more than three hundred applications. The firm has the possibility of preselecting the best candidates according to their trustworthiness, before pursuing to a deeper and more expensive analysis of the candidates. In this case, how many candidates shall be returned by the filtering process?

In this paper, we address the questions raised in the examples introduced above. In particular, we study the effect of using different methods based on trust for selecting partners. This study is enhanced by considering two distinct situations: in the first one, the proposals received by a buyer in a negotiation process are relatively similar and yield comparable utility to the buyer. In the second situation, the proposals are more disparate.

The remainder of this paper is structured as follows. Section 2 introduces related work. Section 3 presents the scenario and notation used in this paper and Section 4 revisits the computational trust model that serves as basis to our study. Section 5 presents the experiments and the main results of our study. Finally, Section 6 concludes the paper and presents future work.

2 Related Work

The majority of the papers in the area of computational trust assumes that trust is the only dimension to take into attention when selecting partners. Falcone and Castelfranchi [10], Kerschbaum et al. [5] and Maximilien and Singh [11] refer that trust must be used additionally to other relevant dimensions, but do not provide a practical study on the complementary use of such dimensions.

Gujral et al. [12] and Griffiths [13] propose models of partner selection based on multi-dimensional trust but do not refer the preselection phase. The work by Padovan et al. [14] develops a scenario that depicts a small value chain. The selection of partners is performed by ranking the received offers by the assessed offer price, which includes the expected value of loss based on a reputation coefficient. This work does not consider preselection.

The work by Kerschbaum et al. [5] addresses the problem of member selection in virtual organizations and considers the possibility of selection of candidate partners based on the reputation of agents, prior to the negotiation phase. The authors also consider the use of trust in the negotiation phase, both as another negotiation dimension, such as price and delivery time, or as a factor in deciding between equally well-suited candidates. However, the empirical evaluation of their trust model is focused on testing its resistance to attacks, and they do not model negotiation in their experiments.

Our work goes further than the related work in the sense that it provides an empirical study on the effect of using different trust-based selection methods – including preselection and the use of trust in the negotiation phase – on the utility of the selecting agents.

3 Scenario and Notation

The scenario used in this paper simulates an Electronic Institution (EI) through which buyer agents select the best suppliers of textile fabric using a simple one round, multi-attribute negotiation protocol. In this section, we describe this scenario and formalize its key concepts.

Every buyer registered in the EI has a business need, which is assigned randomly at setup. This need is represented by a fabric and associated values of quantity, price and delivery time.

The set of possible fabrics is given by $F = \{cotton, chiffon, voile\}$. The values of quantity, price and delivery time are assigned randomly from sets $Q = \{q \in \mathbb{N} : q \in [v_{quant,min}, v_{quant,max}]\}$, $P = \{p \in \mathbb{N} : p \in [v_{price,min}, v_{price,max}]\}$ and $D = \{d \in \mathbb{N} : d \in [v_{dtime,min}, v_{dtime,max}]\}$, respectively. The values $v_{i,min}$ and $v_{i,max}$ define the minimum and maximum values allowed for attribute i , respectively.

This way, buyers announce their needs in the form of a call for proposals (*cfp*), as defined next.

Definition 1. Call for proposals $cfp \in F \times Q \times P \times D$

A call for proposals *cfp* is an ordered tuple from the 4-ary Cartesian product $F \times Q \times P \times D$.

All suppliers registered in the EI are able to provide any type of fabric. When a buyer sends a *cfp* to a defined set of suppliers, each one of these suppliers generates a proposal with its own values for the quantity, price and delivery time

attributes. These values are generated randomly following a uniform distribution in the range $[v_{i,p,min}, v_{i,p,max}]$, where $v_{i,p,min}$ and $v_{i,p,max}$ are defined in equations 1 and 2, respectively.

$$v_{i,p,min} = \max((1 - \delta) \times v_{i,cfp}, v_{i,min}) . \quad (1)$$

$$v_{i,p,max} = \min((1 + \delta) \times v_{i,cfp}, v_{i,max}) . \quad (2)$$

In Equation 1, $v_{i,cfp}$ is the value defined in the *cfp* for attribute i (quantity, price or delivery time), and $\delta \in [0, 1]$ is a *dispersion* parameter that allows to define how distant the generated proposal is from the preferences of the buyer, as stated in the *cfp*.

After receiving the proposals from the suppliers, the buyer calculates the utility of each one of them. The utility of a proposal, μ_p , is given by the complement of the *deviation* between the client preferences specified in the *cfp*, for all the negotiable items price, quantity and delivery time, and what is offered in the received proposal (cf. Equation 3).

$$\mu_p = 1 - \frac{1}{k} \times \left(\sum_i^k \frac{|v_{i,cfp} - v_{i,p}|}{v_{i,max} - v_{i,min}} \right) . \quad (3)$$

In Equation 3, which is adapted from [15], $v_{i,p}$ is the value of the negotiation attribute i of the current proposal in evaluation and k is the number of negotiation attributes considered.

After calculating the utilities of all received proposals, the buyer makes a decision concerning the selection of the *best* proposal. In this paper, we analyze three different approaches for the selection of the best proposal: i) proposals are sorted by their utility (as calculated in Equation 3), and the best proposal is the one that has the highest utility; ii) proposals are sorted by the trustworthiness of the proponent suppliers, and the best proposal is the one which corresponds to the highest value of trustworthiness; and iii) proposals are sorted by the weighted sum of their utility and the trustworthiness of the corresponding proponents, and the best proposal is the one that presents the highest value for this weighted sum. Methods ii) and iii) assume that, previous to the evaluation phase, the buyer estimates the trustworthiness τ of all suppliers that presented a proposal, using the computational trust algorithm presented in Section 4. In addition, method iii) defines the weighting parameter $\omega_\tau \in [0, 1]$, which allows to configure the importance assigned to the trustworthiness component in this selection method (cf. Equation 4).

$$weighting\ sum : \omega_\tau \times \tau + (1 - \omega_\tau) \times \mu_p . \quad (4)$$

In addition to the process described above, we must refer that the buyers have the possibility to preselect the supplier agents that will receive the *cfp*'s, by filtering them by their trustworthiness. After this filtering is done, the selection of the best proposal proceeds as described before.

Finally, after the selection of the best proposal, the buyer establishes a contract with the selected supplier, stipulating that the latter must provide the

Table 1. The set of all handicaps considered in our scenario

Handicap	Description
HFab	handicap in specific fabric
HQt	handicap in high quantities
HDt	handicap in low delivery times
HFabQt	handicap in specific fabric <i>and</i> high quantities
HFabDt	handicap in specific fabric <i>and</i> low delivery times
HQtDt	handicap in high quantities <i>and</i> low delivery times

fabric at the conditions of quantity, price and delivery time described in its proposal.¹

Suppliers can either fulfill or violate the contracts associated to their business interactions, according to their model of behavior. The sample space of outcomes is thus given by $O = \{f, v\}$, where outcome $o = f$ corresponds to a fulfillment of the contract and outcome $o = v$ corresponds to a contractual breach.

In this paper, we model the behavior of suppliers using probabilities. Every supplier that is registered in the EI has an intrinsic degree of performance reflecting the fact that it has some *handicap* in providing specific components in certain *circumstances*. Therefore, at setup, each supplier is randomly assigned a handicap following a uniform distribution over all possible handicaps considered in this paper, which are informally described in Table 1.

The outcome of the interaction between the buyer and the selected supplier is further used to update the value of the trustworthiness of this supplier.

4 The Trust Model

In this section, we briefly describe the computational trust model that serves as basis to our study. It consists of two main components. The first component is *Sinalpha* ([7]), a general aggregator that we have developed that computes the trustworthiness scores of the agents in evaluation based on the trust evidences available on these agents.

The second component of the model is *Contextual Fitness (CF)*, a situation-aware, machine learning-based component that we have developed [16] in order to refine the trustworthiness scores computed by *Sinalpha*, taking into account the current *situation* in assessment. Equation 5 shows the formula used to compute the trustworthiness of agent ag in the specific situation s .

$$\tau(ag, s) = \tau_{sinalpha}(ag) * \tau_{CF}(ag, s) \quad (5)$$

In the equation above, $\tau_{sinalpha}(ag) \in [0, 1]$ gives the trustworthiness score as computed by *Sinalpha*, and $\tau_{CF}(ag, s) \in \{0, 1\}$ gives the value returned by the situation-aware tuner.

¹ The negotiation mechanism we present in this paper is deliberately simple, as it does not constitute the focus of this work. We assume that the conclusions derived from our study using this mechanism are still valid in the presence of other, more complex negotiation protocols.

The mode of operation of the *situation-aware* component is based on the dynamic extraction of *tendencies of failure* from the past behavior of the agent in evaluation. In order to extract these tendencies, we developed an algorithm that uses the *information gain* metric [17]. This metric is used in the ID3 algorithm [17] as a machine learning classification algorithm; however, we use it in an incremental way, by generating a new tree every time a selecting agent needs to assess the trustworthiness of agent *ag* in evaluation.

Before we proceed to the description of the *CF* algorithm, we first give the formal notion of trust evidence.

Definition 2. Trust evidence $evd \in Evd$

A trust evidence $evd \in Evd$ is an ordered tuple from the 6-ary Cartesian product $Evd = AG \times AG \times F \times Q \times D \times O$, where AG is the set of all agents registered in the *EL*.

Using this definition, we can define $Evd^{ag} \subset Evd$ as the subset of all trust evidences that are available to the selecting agent about agent *ag* in evaluation, such that $Evd^{ag} = AG \times AG^* \times F \times Q \times D \times O$, where $AG^* = \{ag\}$.

The *CF* algorithm is illustrated in Algorithm 1.

Algorithm 1. The algorithm of the situation-aware component

```

1: function CF ( $s, Evd^{ag}$ ) returns a value in  $\{0, 1\}$ 
2:    $s$ : context of current situation
3:    $Evd^{ag}$ : set of trust evidences on agent  $ag$ 
4:    $tree^{ag} \leftarrow \text{generateTree}(Evd^{ag})$ 
5:   for each negative rule  $nr_i$  in  $tree^{ag}$  do
6:      $t_{neg} \leftarrow \text{extract negative tendency from rule } nr_i$ 
7:     if there is a match between  $t_{neg}$  and  $s$  then
8:       return 0
9:   return 1

```

Observing Algorithm 1, we verify that it first generates a classification tree from the set of evidences Evd^{ag} , using the evidence outcome as class attribute (line 4). This tree classifies the elements of Evd^{ag} in different classes, corresponding to the elements of the set O of all possible evidence outcomes. Then, for each branch in the tree corresponding to *negative* classes (line 5), a tendency of failure t_{neg} is extracted (line 6). If this tendency matches situation s in assessment (line 7), this means that the agent has a *tendency to fail* in situations similar to the current one, and the algorithm returns the value 0 (line 8). Otherwise, it returns the value 1 (line 9).

Being an incremental process, the algorithm allows for the extracted tendencies of behavior of the evaluated target to change dynamically with the size of the historical data on the agent, being, this way, very responsive to the changes of behavior of the agents in assessment. Another good property of this algorithm is that it is able to extract negative tendencies of behavior since the first evidences available, achieving good performances with very small datasets.

5 Experiments

We ran a set of experiments in order to analyze the effect of using trust on the selection phase of automatic negotiation processes. Most of the papers on computational trust show the benefits of using trust in the selection of partners, but these are described exclusively in terms of the number of successful transactions. In these experiments, we compared different selection methods, including those that do not use trust, those that use trust in a preselection phase previous to the negotiation, those that use trust in the negotiation process and, finally, those that use trust in a preselection phase *and* in the negotiation process.

5.1 Testbed and Methodology

All experiments described in this paper were performed using the Repast simulation tool [18] and the scenario described in Section 3.

We ran six different experiments, according to the selection methods in evaluation. Table 2 presents these experiments.

As can be observed in Table 2, we tested two different filtering approaches (experiments 2, 3, 5 and 6): the first one preselected 10% of the most trustworthy suppliers registered in the EI, and the second one preselected 50% of this population. In experiments 1 and 4, no trust-based preselection was performed and all suppliers were allowed to proceed to the negotiation phase.

In all experiments, we used 20 buyer agents and 50 supplier agents. Every supplier had a 95% chance to succeed in case it did not present a handicap in the situation embedded in the *cfp*. This probability dropped to 5% when its handicap matched the *cfp*'s situation.

Each experiment was composed of 30 *episodes*, and at every episode each buyer started a new negotiation cycle by issuing a new *cfp*. At the first episode of each experiment, the repository of trust evidences was cleaned, which means that the trustworthiness of all suppliers was set to zero. Finally, we ran every experiment 20 times. At every new run, the buyer agents changed their preferential values regarding their business needs, by randomly picking up new values from sets F , Q , D and P .

In order to enhance our study on the effect of using trust in selection processes, we considered two different values for the dispersion parameter δ : 0.2 and 1.0 (cf. equations 1 and 2). As mentioned in Section 3, parameter δ is used to

Table 2. Different types of experiments, based on the places where trust was used

#	Selection Method	Preselection	Negotiation
1	No Trust	—	—
2	Trust in preselection (10%)	✓	—
3	Trust in preselection (50%)	✓	—
4	Trust in negotiation	—	✓
5	Trust in preselection (10%) and in negotiation	✓	✓
6	Trust in preselection (50%) and in negotiation	✓	✓

configure how distant the proposals generated by the suppliers are from the conditions specified in the received *cfp*. In these experiments, the value 0.2 was used to configure small deviations, which means that all the proposals received by the buyer agent were close to its preferential values for current interaction; in opposition, the value of 1.0 allowed for a greater dispersion in the utility of the proposals received by the buyer agent.

5.2 Evaluation Metrics

In order to evaluate and compare each one of the selection methods considered in the experiments, we used six different performance metrics. The first metric was the *utility of the interaction* (μ_t), given in Equation 6. We averaged this utility over all buyers and all episodes.

$$\mu_t = \begin{cases} \mu_p, & \text{if } o = f, \\ 0, & \text{if } o = v. \end{cases} \quad (6)$$

The second metric was the *number of positive outcomes* (o^+) obtained by all buyer agents in an episode, averaged over all episodes. The third metric was the *number of different suppliers* (Δ_{sup}) selected by all buyers in one episode, averaged over all episodes. The fourth and the fifth metrics measured the *trustworthiness of the supplier* and the *utility of the proposal* selected by a buyer in one episode (τ_s and μ_s , respectively), averaged over all buyers and all episodes. Finally, the sixth metric was the *number of unfitted choices* (ζ) performed by a buyer, averaged over all buyers and all episodes. This latter metric is related to the *CF* component of our computational trust model. It concerns the choice of a supplier that the buyer knows has an handicap in the current business conditions.

5.3 Results

In this section, we start by presenting the results obtained for a dispersion value (δ) of 0.2, and then we present the values obtained for $\delta = 1.0$.

Experiments with $\delta = 0.2$. The first part of the experiments was performed using $\delta = 0.2$. We first measured the average utility of the proposals *received* by a buyer in one episode and averaged it over all buyers and all episodes. The value we obtained for this average was 0.93, with a standard deviation of 0.03. These values were obtained consistently for all the selection methods tested. Their meaning is that the suppliers offered proposals with approximated utility and close to the buyers' preferences.

Table 3 presents the results obtained in this first set of experiments for the metrics described in Section 5.2.

In experiments 4.x, 5.x and 6.x, the utility of the interaction (μ_t) is a weighted sum of the trustworthiness of the supplier and the utility of its proposal. In the experiments, we used two different values for the weight of the trust component, $\omega_\tau = 0.1$ and $\omega_\tau = 0.5$ (cf. Equation 4).

Table 3. Results obtained with $\delta = 0.2$

#	Selection Method	μ_t	o^+	Δ_{sup}	τ_s	μ_s	ζ
1	No Trust	0.69	0.70	0.84	0.17	0.98	0.21
2	Trust in preselection (10%)	0.82	0.85	0.35	0.80	0.96	0.00
3	Trust in preselection (50%)	0.79	0.81	0.75	0.41	0.98	0.01
4.1	Trust in negotiation ($\omega_\tau = 0.1$)	0.82	0.87	0.23	0.83	0.95	0.00
4.2	Trust in negotiation ($\omega_\tau = 0.5$)	0.79	0.85	0.11	0.90	0.93	0.00
5.1	Trust in presel. (10%) & in neg. ($\omega_\tau = 0.1$)	0.83	0.88	0.18	0.88	0.95	0.00
5.2	Trust in presel. (10%) & in neg. ($\omega_\tau = 0.5$)	0.82	0.88	0.11	0.90	0.93	0.00
6.1	Trust in presel. (50%) & in neg. ($\omega_\tau = 0.1$)	0.83	0.87	0.22	0.85	0.95	0.00
6.2	Trust in presel. (50%) & in neg. ($\omega_\tau = 0.5$)	0.83	0.89	0.11	0.91	0.93	0.00

From the results presented in Table 3, we verify that the selection method that did not rely on trust got worse results for the metric utility of the interaction ($\mu_t = 0.69$), as it would be expected. This method selected the suppliers by the utility of their proposals, which allowed for the selection of proposals with very high values of utility ($\mu_s = 0.98$) and for a high degree of exploration of new partners ($\Delta_{sup} = 0.84$). However, the trustworthiness of the selected suppliers was in average very low ($\tau_s = 0.17$), and a relevant number of unfitted choices was done ($\zeta = 0.21$). In consequence, the number of positive outcomes was relatively low ($o^+ = 0.70$).

The results presented in Table 3 also show that the mixed use of trust, both in preselection and in the negotiation phase (experiments 5.x and 6.x), got the best results in terms of the utility of interaction ($\mu_t \approx 0.83$), for all combinations of the degree of filtering (10% and 50%) and ω_τ . In this case, we verified that although reinforcing the trust component in the negotiation phase ($\omega_\tau = 0.5$) allowed for higher values of the trustworthiness of the selected suppliers (τ_s), relaxing this value ($\omega_\tau = 0.1$) allowed for higher values of the utility of the selected proposals (μ_s). Also, the difference between filtering the 10% or the 50% more trustworthy agents was not relevant for $\delta = 0.2$.

Finally, we observed that both the use of standalone, stricter preselection (10%) and the use of trust in negotiation with $\omega = 0.1$ allowed for similar good results of μ_t (0.82), and approximated values of o^+ , τ_s and μ_s . The use of standalone, more relaxed preselection (50%) and the use of trust in negotiation with $\omega = 0.5$ got lower values of μ_t (0.79), with the first method exploring more the utility of the proposals ($\mu_s = 0.98$) in detriment to the trustworthiness of suppliers ($\tau_s = 0.41$), and the latter having an opposite behavior ($\mu_s = 0.93$ and $\tau_s = 0.90$).

Experiments with $\delta = 1.0$. In the second part of the experiments, we wanted to evaluate the effect of each one of the selection methods when the dispersion in the utilities provided by different suppliers was bigger. For that, we configured δ to have value 1.0. In this case, the measured value for the average utility of the received proposals was 0.73, with a standard deviation of 0.11, showing a higher variance in the proposals made by the suppliers.

Table 4. Results obtained with $\delta = 1.0$

#	Selection Method	μ_t	o^+	Δ_{sup}	τ_s	μ_s	ζ
1	No Trust	0.66	0.71	0.83	0.17	0.93	0.21
2	Trust in preselection (10%)	0.73	0.87	0.36	0.80	0.84	0.00
3	Trust in preselection (50%)	0.73	0.80	0.76	0.41	0.92	0.02
4.1	Trust in negotiation ($\omega_\tau = 0.1$)	0.75	0.83	0.63	0.58	0.91	0.00
4.2	Trust in negotiation ($\omega_\tau = 0.5$)	0.67	0.88	0.14	0.88	0.77	0.00
5.1	Trust in presel. (10%) & in neg ($\omega_\tau = 0.1$)	0.73	0.87	0.32	0.83	0.85	0.00
5.2	Trust in presel. (10%) & in neg ($\omega_\tau = 0.5$)	0.66	0.86	0.13	0.89	0.77	0.00
6.1	Trust in presel. (50%) & in neg ($\omega_\tau = 0.1$)	0.77	0.85	0.59	0.64	0.90	0.00
6.2	Trust in presel. (50%) & in neg ($\omega_\tau = 0.5$)	0.66	0.86	0.14	0.89	0.77	0.00

Table 4 presents the results obtained in this second set of experiments for the metrics described before.

The results obtained and presented in Table 4 show relevant differences from the results obtained with $\delta = 0.2$. In fact, the combined use of trust in preselection and in negotiation did not achieve the same good performance as observed with $\delta = 0.2$, for $\omega_\tau = 0.5$. As illustrated in Table 4, in experiments 5.2 and 6.2, the buyers kept selecting the same trustworthy agents again and again ($\Delta_{sup} \approx 0.14$), showing a rather parochial behavior. This had the cost of decreasing the utility of the selected proposals ($\mu_s = 0.77$) in a significant manner, with just a slight improvement in the trustworthiness of the selected suppliers ($\tau_s = 0.89$). In a general case, we can observe in Table 4 that all trust methods that used trust in negotiation with a strong weight for the trust component ($\omega_\tau = 0.5$) got as little value for μ_t as the selection approach that did not use trust at all. In the same way, approaches using more restricted preselection (10%) exhibited significantly lower values of μ_t than their counterparts using $\delta = 0.2$.

The results obtained also show that the combined use of a more relaxed filtering of suppliers (50%) and a lower weight of the trust component ($\omega_\tau = 0.1$) had again achieved the best result for the average utility of interaction ($\mu_t = 0.77$). This approach allowed for a better equilibrium between the trustworthiness of the selected suppliers and the utility of the selected proposals.

5.4 Interpretation of Results

The results obtained and presented in the sections above allow us to conclude that parochialism in partner selection is acceptable when the proposals in evaluation are not too disparate ($\delta = 0.2$). In this case, selection methods strongly supported on trust reveal to be good choices, as they are able to select more reliable partners without the expense of losing utility.

However, we have shown that when the standard deviation of the utility of the received proposals is about 11% of the mean, the excessive use of trust is not acceptable, as parochialism prevents buyers from exploring partners that offer deals with higher utilities. In both the situations that we have studied, a

method that preselects half of the population of candidate suppliers and then moderately uses trust in negotiation revealed to be a better choice (experiments 6.1 in tables 3 and 4).

6 Conclusions

Recently, different agent-based trust models have been proposed as support mechanisms to the selection of partners. These proposals are based on the hard assumption that trust enhances the selection process, but no studies were presented on the role of trust in the presence of other selection differentiation factors.

In this paper, we empirically evaluated and compared different selection methods based on trust. We concluded that methods that strongly rely on trust are not adequate when the proposals in evaluation are disparate. The best solution seems to be the trust-based preselection of about half of the candidate partners, followed by a selection process where the weight of the trust component must be adjusted to the estimated dispersion of the proposals' utilities.

As future work, we intend to explore other prefiltering options and different other combinations of the parameters configured in the experiments. We also intend to explore other computational trust models in our study.

Acknowledgments. This research is funded by FCT (Fundação para a Ciência e a Tecnologia) project PTDC/EIA-EIA/104420/2008. The first author enjoys a PhD grant with reference SFRH/BD/39070/2007 from FCT.

References

1. Dasgupta, P.: Trust as a Commodity. In: Gambetta, D. (ed.) *Trust: Making and Breaking Cooperative Relations*, Department of Sociology, University of Oxford, pp. 49–72 (2000)
2. Jøsang, A., Ismail, R.: The Beta Reputation System. In: *Proceedings of the 15th Bled Electronic Commerce Conference* (2002)
3. Huynh, T.D., Jennings, N.R., Shadbolt, N.R.: An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems* 13, 119–154 (2006)
4. Yu, B., Singh, M.P.: An Evidential Model of Distributed Reputation Management. In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: part 1., AAMAS 2002*, pp. 294–301(2002)
5. Kerschbaum, F., Haller, J., Karabulut, Y., Robinson, P.: Pathtrust: A trust-based reputation service for virtual organization formation. In: Stølen, K., Winsborough, W.H., Martinelli, F., Massacci, F. (eds.) *iTrust 2006*. LNCS, vol. 3986, pp. 193–205. Springer, Heidelberg (2006)
6. Marsh, S., Briggs, P.: Examining Trust, Forgiveness and Regret as Computational Concepts. In: Golbeck, J. (ed.) *Computing with Social Trust*. Human-Computer Interaction Series, pp. 9–43. Springer, London (2009)
7. Urbano, J., Rocha, A.P., Oliveira, E.: Computing confidence values: Does trust dynamics matter? In: Lopes, L.S., Lau, N., Mariano, P., Rocha, L.M. (eds.) *EPIA 2009*. LNCS, vol. 5816, pp. 520–531. Springer, Heidelberg (2009)

8. Melaye, D., Demazeau, Y.: Bayesian dynamic trust model. In: *Multi-Agent Systems and Applications Iv, Proceedings*, vol. 3690, pp. 480–489 (2005)
9. Joseph, S., Sierra, C., Schorlemmer, M., Dellunde, P.: Deductive coherence and norm adoption. *Logic Journal of the IGPL* 18, 118–156 (2010)
10. Castelfranchi, C., Falcone, R., Pezzulo, G.: Trust in information sources as a source for trust: a fuzzy approach. In: *Procs. of the Second International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2003*, pp. 89–96 (2003)
11. Maximilien, E.M., Singh, M.P.: Agent-based trust model involving multiple qualities. In: *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems (July 2005)*
12. Gujral, N., DeAngelis, D., Fullam, K.K., Barber, K.S.: Modeling multi-dimensional trust. In: *Procs. of The Workshop on Trust in Agent Societies at AAMAS 2006*, pp. 35–41 (2006)
13. Griffiths, N.: Task delegation using experience-based multi-dimensional trust. In: *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS 2005, New York, NY, USA*, pp. 489–496 (2005)
14. Padovan, B., Sackmann, S., Eymann, T., Pippow, I.: A prototype for an agent-based secure electronic marketplace including reputation-tracking mechanisms. *Int. J. Electron. Commerce* 6, 93–113 (2002)
15. Rocha, A.P., Oliveira, E.: An electronic market architecture for the formation of virtual enterprises. In: *Proceedings of the IFIP TC5 WG5.3 / PRODNET Working Conference on Infrastructures for Virtual Enterprises: Networking Industrial Enterprises*, pp. 421–432. Kluwer, B.V., Deventer, The Netherlands (1999)
16. Urbano, J., Rocha, A.P., Oliveira, E.: Trustworthiness tendency incremental extraction using information gain. In: *Proceedings of the 2010 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, WI-IAT 2010*, vol. 2, pp. 411–414. IEEE Computer Society, Los Alamitos (2010)
17. Quinlan, J.R.: Induction of Decision Trees. *Mach. Learn.* 1, 81–106 (1986)
18. North, M., Howe, T., Collier, N., Vos, J.: A declarative model assembly infrastructure for verification and validation. In: Takahashi, S., Sallach, D.L., Rouchier, J. (eds.) *Advancing Social Simulation: The First World Congress*. Springer, Heidelberg (2007)