

Human Experiments in Trust Dynamics

Catholijn M. Jonker¹, Joost J.P. Schalken^{1,*},
Jan Theeuwes², and Jan Treur^{1,3}

¹Vrije Universiteit Amsterdam, Department of Artificial Intelligence
De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands
<http://www.cs.vu.nl/~{jonker,joosts,treur}>
{jonker,joosts,treur}@cs.vu.nl

²Vrije Universiteit Amsterdam, Department of Cognitive Psychology
De Boelelaan 1111, 1081 HV Amsterdam, The Netherlands
<http://www.psy.vu.nl/~Theeuwes>, J.Theeuwes@psy.vu.nl

³Universiteit Utrecht, Department of Philosophy
Heidelberglaan 8, 3584 CS Utrecht, The Netherlands.

Abstract

In the literature, the validity of theories or models for trust is usually based on intuition and common sense. Theories and models are not often verified experimentally. The research reported here contributes results of experiments on the dynamics of trust over time depending on positive or negative experiences. In previous research a number of dynamic properties for such trust dynamics were identified, but not verified empirically. As a continuation of this work, now these properties have been verified in an experimental setting. The outcomes of the experiment (involving a substantial number of 238 subjects) are discussed and related to the previously formulated dynamic properties.

1 Introduction

Trust is omnipresent in all our interactions with other people; e.g., [1], [2], [4], [11]. Without trust, the every day social life which we take for granted is simply not possible [11], cited in [7]. Our society, in which each individual plays its own niche role in complex network of social interactions, would grind to a halt due to a lack of cooperation. A difficulty with the concept trust, is that it is impossible to observe it, or to directly relate it to other simple observable facts. This may be one of the reasons why not many sociologists addressed the concept until recently [12].

1.1 Characterising Trust by Entailed Behaviour

Elofson describes trust as “the reliance upon the characteristics of an object, or the occurrence of an event, or the behavior of a person in order to achieve a desired but uncertain objective in a risky situation.” [5]. Luhmann gives a similar definition of trust: “Trust is a reliance in turbulent conditions on some number of certainties and on

* Currently at Vrije Universiteit Amsterdam, Department of Information Management and Software Engineering, De Boelelaan 1081a, 1081 HV Amsterdam, The Netherlands. URL: <http://www.cs.vu.nl/~joosts>. Email: jjp.schalken@few.vu.nl.

other individuals' actions, that affect one's own welfare, that despite conditions largely unknown can be counted on to act in a predictable and presumably benevolent fashion" [11]. Trust can also be defined as "the degree of confidence that you feel when you think about a relationship" [13] as cited in [5]. Or as "an interpersonal or interorganizational state that reflects the extent to which the parties can predict one another's behavior, can depend on one another when it counts; and have faith that the other will continue to act in a responsive manner despite an uncertain future." [16].

The above characterisations of trust look forward in the sense that once a state of trust is there in an agent, they explain how this is used by the agent to make decisions on behaviour. The backward perspective in the sense of the question how a trust state was reached, i.e., how trust is gained or lost over time is left out of consideration in these characterisations.

1.2 Dynamics of Trust Based on Experiences

Trust is not a static mental state, but instead a dynamic one, as trust can change over time. This makes the question of what generates, maintains, substitutes, or collapses trusting relations [6] important. One of the central hypothesis in the research reported here is that trust is based on observed events in the real world. Lewis and Weigert state the same when they state that trust is formed by "observations that indicate that members of a system act according to and are secure in the expected futures constituted by the presence of each other for their symbolic representations." [9]. Eloffson identifies the same origins of trust in "Trust is the outcome of observations leading to the belief that the actions of another may be relied upon" [5].

Events that are observed in the real world can only be interpreted within the context in which they take place. This context defines whether an event helps in achieving one's goals or not and also helps understanding the situation in which the other was placed (e.g., sometimes you cannot really blame someone for not helping you.). Wels and Van Loon point at a similar issue when they say "Every event is created by a different ensemble of interactions; all sense making is relative to this level specificity. Hence, the meaning attached to the events to the concepts vary not only in relation to different actors, but also to different contexts." [15].

In [8], a formal framework to model dynamics of trust is given. In the framework some of the events in the world are considered trust-influencing experiences. Such experiences are either positive or negative, in other words may increase your trust in something or someone, or decrease it. According to this framework the trust that is acquired by an actor depends on two variables, the initial trust and the trust dynamics. The model for trust dynamics specify how an agent adjusts its trust in someone or something based on experiences.

1.3 Trust Dynamics Experiments

In a number of recent articles on intelligent agents and (formal) trust models the validity of certain models is only "proven" with an argumentation that appeals to common sense. The extent to which the models are correct (either in a descriptive or in a normative sense) is not measured. The only way to assess the correctness of a trust model, which claims to describe cognitive processes related to trust, is to perform experiments with human test-subjects. Or as Smet, Wels and Van Loon put

it: 'We need, however, to stop making speculative claims based on grand, but rather unsubstantiated, theorising if we are to make any proper sense out of trust and co-operation.' [14].

In this research the focus is on the question of how to adjust trust based on experiences and to verify that experiences really do influence trust in other persons, organisations or objects. In this paper a description is given of the experiment that has been conducted to determine whether and in what form trust is really influenced by experiences as has been suggested in literature (see for example [9; 5; 8]), and to determine if some regularities occur in the extent and direction to which trust is influenced by experiences.

In this paper, Section 2 describes the design of the experiment. In Section 3 the outcomes of the preliminary validation tests in the experiment are presented. Section 4 presents the outcomes of the final experiments. In Section 5 the outcomes are compared to dynamic properties of trust as have been identified in previous work. Section 6 concludes the paper.

2 Design of the Experiment

Possible angles to examine the dynamics of trust are to focus on the factors that influence how initial trust is established or to focus on how trust is influenced by events. In the research reported in this paper the focus is on how events influence trust.

2.1 Factors Affecting a Trust State

Numerous factors influence the impact an event has on someone's trust. These factors can be divided into two categories: factors that influence how an event is evaluated and factors that influence the relative weight an experience has on someone's trust, in relation to all the other experiences that person had and his or her initial trust value. An example of a factor that could influence how an event is evaluated is kinship: generally people will feel much more left in the lurch if a family member does not cooperate than if a stranger is uncooperative. An example of a factor that could influence the relative weight of an experience is the amount of time passed since the event took place: people might place more trust in a photocopier that failed to work two years ago than they would in a photocopier that malfunctioned just two hours ago.

The difference between the evaluation of an event and the relative weight is that the evaluation of an event is dependent on the context in which the event took place, whereas the relative weight may depend on the total collection of experiences over time of that person. Usually the evaluation of an event remains unchanged after the event has occurred, whereas the relative weight of an experience will usually change over time as the person has acquired more experiences. However, it may be possible for an evaluation of an event to change if a person acquires more information about the context in which the event took place. The focus of this research is on the factors that influence the relative weight an experience has.

Time is one of the important factors that influences the relative weight of an experience. The amount of influence an event has decreases over time, as is

demonstrated by Derrida in the specific context of giving gifts [3], cited in [15]. In [8] it is suggested that the temporal order in which the events occurred, could influence the impact an event has on the trust state. Following this, in this study the temporal order of events has been used as the factor that influences the relative weight of an experience.

2.2 Overview of the Experiment

In the experiment, subjects have been presented with sequences of short stories (written in Dutch) that each describe an event that occurred with an organisation or an object. Within a single sequence all stories deal with the same object or organisation. After each story the subject is asked to state (on a five-points Lickert scale) how much trust he or she has in the object or organisation. The subject is instructed to base her trust in the object or organisation on all stories that have been presented. The difference in the trust values the subject assigns to an object or organisation shows how the trust in that organisation or object has been affected by the experiences. By varying the order of the events, these differences in trust values allow the measurement of how experiences and the order in which the experiences occur, influence a person's trust in an organisation or object.

The questions have been presented to the subjects over the Internet. Distributing the questionnaire over the Internet has two distinctive advantages: first of all the subjects can complete the questionnaire when it suits them and therefore the response rate of potential subjects has been relatively high (the estimated response rate was 35%), even though almost no financial incentive to participate in the experiment has been offered. A potential drawback of distributing the questionnaires over the Internet is that there is little control over the environment in which the subject fills in its questionnaire.

For their participation in the experiment the subjects were offered a lollipop and the chance to win a single reward of € 20. In total 238 people participated in the experiments, which brings the estimated response rate to 35% (78 personal invitations have been sent by e-mail to relatives, friends and participants of the pre-tests of the experiment, 250 flyers have been handed out on campus and an open invitation to participate in the experiment has been posted on the Usenet newsgroup of the department of Mathematics and Computer Science, which has an estimated reach of 350 persons).

As the experiment tries to determine the effect of experiences with an object or organisation on the amount of trust a person has in that object or organisation, an operationalisation of the concept amount of trust is necessary. In this experiment we will use a five-point Lickert rating scale to allow subjects to state the amount of trust they have in the object or organisation. The five-point trust rating scale contains the following levels: "veel vertrouwen" (much trust, 5) "redelijk vertrouwen" (a reasonable amount of trust, 4) "neutral" (neutral, 3) "weinig vertrouwen" (little trust, 2) "heel weinig vertrouwen" (very little trust, 1).

In the test the effect of experiences with an organisation or an object on the trust in that organisation or object is measured. The effect of the experiences on trust is measured by exposing the subject to multiple events with a certain organisation or object and measuring the trust in the object or organisation between every event using the trust rating scale described above.

It is time-consuming and expensive to let the subjects have real interactions with objects and organisations, as this would require subjects to undergo the experiment in a lab in which the interactions can be simulated. This is a main reason why in this experiment instead of undergoing real interactions with objects and organisations the subject have been presented with stories that describe certain experiences.

2.3 The Scenarios

For this experiment two scenarios have been written: in one scenario the subjects deal with a photocopier and in the other scenario the subjects interact with a travel agency. Each scenario consists of an introduction and ten distinctive stories, five of which are positive (written to induce trust) and five of which are negative (written to induce distrust).

The topic of the scenarios (a photocopier and a travel agency) have been chosen so that not many people have strong emotional feelings about the subject (as one could have with the Dutch railroad corporation or the tax authority) which could influence the results of the experiment. The stories have been written in a neutral tone, to prevent to explicitly direct the subject to the ‘desired’ response. On top of that stories do not cross-reference each other, as the stories had to be presented in a randomised order.

Photocopier scenario. The first scenario contains an introduction and ten stories, which describe experiences a user can have with a photocopier that uses a debit card. It is assumed that the subjects have had prior experiences with photocopiers that use debit cards (this assumption holds for almost all students). After each event the subject can indicate his or her trust in the photocopier. In the following table the introduction and the some of the stories can be found (both in Dutch and translated into English):

Story no.	Content of story
	Introduction
-	Je moet morgen een voordracht geven over een paper die je geschreven hebt. Hiervoor heb je sheets nodig, en voor elk van de toehoorders een hand-out en een kopie van je paper. Tomorrow you have to give a presentation about a paper you wrote. For this presentation you need transparencies and for each member of the audience a hand-out and a copy of your paper.
	Positive experiences
1	Je kopieert de kaft van het paper apart, omdat dit op speciaal papier gekopieerd moet worden. Om te kijken of de kopieermachine het gekleurde karton, dat je voor de kaft wil gebruiken, aankan maak je eerst één proefkopie. Na een geslaagde proefkopie begin je met kopiëren. You copy the cover of the paper separately, because it needs to be copied on special paper. To see if the photocopier can handle the colored carton, which you want to use as a cover, you first make a single test copy. After a successful test copy you start copying.

	Negative experiences
1	De kopieermachine werkt met een kaartsysteem. Na het maken van een aantal kopieën is je oude kaart op, daarom ga je naar de receptie om een kopieerkaart te kopen. Vervolgens weigert de machine te kopiëren, omdat volgens de machine de nieuwe kaart leeg is.
	The photocopier works with a debit card-system. Having made a number of copies, your old card runs out of debit, so you go to the reception to purchase a new debit card. Thereupon the machine refuses to copy, claiming the new card is empty.

Travel agency scenario. The second scenario contains an introduction and ten stories, which describe experiences a user can have with a travel agency. After each event the subject can indicate his or her trust in the travel agency.

2.4 Balancing the Experiment

In our study the subject received two series of each 10 stories. The first sequence of stories deals with an object, a photocopier, and the second sequence deals with an organisation, a travel agency. In each sequence half the stories were positive (designed to induce trust) and half the stories were negative (designed to induce distrust).

In the study half the subjects first received the negative stories for the copier or the travel agency and after that the positive stories for the copier or the travel agency and half the subjects first received the positive stories. This made it possible to study both the increase and decrease in trust. Within a sequence of positive or negative stories the order of the stories has been determined randomly, in order to prevent side effects from the order of the stories.

To prevent carry-over effects of the first part of the test to the second part of the test, half the subjects that received in the first part of the test a sequence with first negative stories and after that positive stories will receive in the second part of the test a sequence with first negative stories and after that positive stories and the other half of the subjects will receive in the second part of the test a sequence with first positive stories and after that negative stories. This way four different paths through the test are created, which is illustrated in the following table.

experiment groups	stories about photocopier	stories about travel agency
1	first five positive stories, followed by five negative stories	first five positive stories, followed by five negative stories
2	first five negative stories, followed by five positive stories	first five negative stories, followed by five positive stories
3	first five positive stories, followed by five negative stories	first five negative stories, followed by five positive stories
4	first five negative stories, followed by five positive stories	first five positive stories, followed by five negative stories

Table 1. Sequences in which the stories are presented to the subject

3 Analysis of the Results: Preliminary Validation Tests

In this section the outcome of the two preliminary validation tests (to test the experience stories) and the final experiment are briefly discussed and the data is analysed to draw some preliminary conclusions. To determine whether the stories that describe interactions with the photocopier and the travel agency have the desired positive or negative effect on the trust in that photocopier and the travel agency, a preliminary test has been executed.

In the preliminary test 80 subjects were presented with a form on which a single description of an experience with the photocopier and a single description of an experience with the travel agency were recorded. After reading each story, the subjects were asked to answer whether they would have more trust, the same amount of trust or less trust in the photocopier or travel agency after the described event.

Answer on photocopier story * Story number on photocopier Crosstabulation												
Count												
		Story category on photocopier										Total
		positive stories					negative stories					
		1	2	3	4	5	1	2	3	4	5	
Answer on photocopier story	more trust	6	2	3	2	3						16
	less trust				1	1	7	6	4	6	5	30
	no change	2	6	5	5	4	1	2	4	2	3	34
Total		8	8	8	8	8	8	8	8	8	8	80

Table 2. Answers on photocopier story: per story

Answer on photocopier story * Story category on photocopier Crosstabulation					
Count					
		Story number on photocopier			Total
		positive stories		negative stories	
		1	2	3	
Answer on photocopier story	more trust	16			16
	less trust	2		28	30
	no change	22		12	34
Total		40		40	80

Table 3. Answers on photocopier story: overall

In the preliminary test each story has been presented eight times. Unwanted carry-over effects from the story about the photocopier onto the responses on the story about the travel agency were prevented, because 50% of the stories about the travel agency were preceded by a positive story about the photocopier and 50% of the

stories about the travel agency were preceded by a negative story about the photocopier.

In Tables 2 and 3 the effectiveness of the stories about the photocopier can be found. In these tables one can see that all negative stories about the photocopier have the desired effect (in all cases the trust either decreases or remains the same). The effect of the positive stories is not as good as that of the negative stories, but when comparing the groups of positive stories and negative stories, the effect is still visible and significant, as Cramer's $V = 0.720$, (cf. [10], pp. 14-15) which means that a significant correlation exists between the category of the stories and its outcome. Therefore all stories depicted in the tables have been used in the further experiments.

In Tables 4 and 5 the effectiveness of the stories about the travel agency can be found.

Answer on travel agency story * Story number on travel agency Crosstabulation													
Count													
		Story number on travel agency										Total	
		positive stories					negative stories						
		1	2	3	4	5	1	2	3	4	5		
Answer on travel agency story	more trust	6	2	5	6	5							24
	less trust	1	1			1	2	5	7	5	8		30
	no change	1	5	3	2	2	6	3	1	3			26
Total		8	8	8	8	8	8	8	8	8	8	8	80

Table 4. Answers on travel agency stories: per story

Answer on travel agency story * Story category on travel agency Crosstabulation					
Count					
		Story category on travel agency			Total
		positive stories		negative stories	
		1	2	3	
Answer on travel agency story	more trust	24			24
	less trust	3		27	30
	no change	13		13	26
Total		40		40	80

Table 5. Answers on travel agency stories: overall

In Table 4 one can see that all stories, except for positive story 2 and negative story 1, have the desired effect on the trust in the travel agency. Comparing the effects of the group of positive and negative stories about the travel agency (see Table 5), it is clear that there is a difference between the effects of the positive and negative stories. This difference is significant, as Cramer's $V = 0,735$ (cf. [10], pp. 14-15).

We can conclude based on the above crosstables that although the effects of the stories are not perfect (in which case positive stories would only increase trust and

negative stories would only decrease trust), the overall effect of the stories is significant and therefore the stories can be used for the experiment.

4 Analysis of the Results: Final Experiment

In total 294 subjects started with the final experiment. This means that 294 people opened the first web-page of the questionnaire of the final experiment. From these 294 subjects 238 subjects (81%) completed the full questionnaire. The other 19% of the subjects were either not able to complete the questionnaire because of technical problems, decided to stop during the experiment or did not respond to a question within a given time limit of 15 minutes between each question. Only the data obtained from subjects that fully completed the questionnaire has been used for analysis.

After the subject has been presented with both pre-tests, the main part of the test begins. In the main test the effect of experiences with an organisation or an object on the trust in that organisation or object is measured. The effect of the experiences on trust is measured by describing various experiences in small stories and instructing the subject to state his or her trust in the object or organisation after having went through such an experience. Trust in a certain object or organisation is stated using the five-points trust rating scale described in Section 2. Each scenario consisted of an introduction and ten distinctive stories, five of which were positive (written to induce trust) and five of which were negative (written to induce distrust). For more details, see Section 2. In this following section the results of the main part of experiment are presented.

In the Figures 1 to 4 below the dynamics of trust, for both the photocopier and the travel agency are plotted. In the plots the median of the trust values are displayed. We can immediately see that trust increases when the subject had a positive experience and that trust decreases as negative experiences are received. Moreover, as we can see, there is a clear difference between the between the plots that start with negative experiences and the plots that start with positive experiences.

To determine the significance of the difference, a 2-way between subjects ANOVA test was performed on the means of the positive and negative experiences within a single scenario. The ANOVA test takes both the experience (positive or negative) and the order in which experiences are presented (positive-first or negative-first) into account. From the results we can see that both factors have an effect on trust in the object or organization at a significance level beyond 0,001. Furthermore we can conclude that there is no significant interaction between experience (positive or negative) and the order in which experiences are presented.




Fig. 1. Dynamics of trust in Photocopier: positive experiences first




Fig. 2. Dynamics of trust in Photocopier: negative experiences first

Having examined and established the order effect of experiences for mean trust scores, now the duration of the effect is established.

For the stories about the photocopier, the negative effect, of first receiving negative stories, on the trust score after receiving positive experiences remains significant for the whole trail. The positive effect, of first receiving positive stories, on the trust score after positive experiences remains significant only for the first three turns, after which the difference becomes insignificant (at the 0.05 significance level).

For the stories about the travel agency, the negative effect, of first receiving negative stories, on the trust score after receiving positive experiences remains significant a single round. The positive effect, of first receiving positive stories, on the trust score after positive experiences remains significant only for the first turn, after which the difference becomes insignificant (at the 0.05 significance level).




Fig. 3. Dynamics of trust in Travel Agency: positive experiences first




Fig. 4. Dynamics of trust in Travel Agency: negative experiences first

5 Relating the Outcomes to Previous Work

In [8] a number of (possible) dynamic properties of trust have been identified. In this section the outcomes of the experiments are compared to the most relevant of these dynamic properties.

Positive and negative trust extension

After a positive experience an agent will trust more or the same amount, but never less. After a negative experience an agent will trust less or the same amount, but never more.

From the figures in Section 3, Tables 3 and 5 it can be seen that in a large majority of cases these two properties hold. In particular, negative trust extension always holds, both for the photocopier and the traveling agency. Positive trust extension holds in

95% of the cases for the photocopier and in 92.5 % of the cases for the travelling agency. Other evidence that by and large these properties hold can be found in the graphs in Section 4. In none of the graphs depicting the median values a transition can be found that violates one of the two properties. Of course these results heavily depend on the chosen stories as being positive or negative. For example, if a story was classified as positive, whereas it only is felt as a slightly positive experience, it would be reasonable to assume that a very high level of trust can decrease to a slightly less high (but still positive) trust level by this experience.

Other properties in [8] address the flexibility of trust: can negative trust be made positive (again) by offering the appropriate types of experiences, and vice versa? From our analysis in [8] the following more or less opposite properties would be possible

Degree of trust dropping or trust gaining N

After N negative events the trust will be negative.

After N positive events the trust will be positive.

Negative or positive trust fixation of degree N

After N negative events the agent will never trust anymore and its trust will remain the least possible. After N positive events the agent will forever trust (even when faced with negative events) and its trust will remain maximal.

These trust fixation properties are more or less the opposite of the previous ones. In [8] we could not indicate which of them would be more realistic. On the basis of the experiments, now it is suggested that trust fixation does not occur, at least in contexts as investigated. In the graphs depicted in Section 4, negative or positive trust fixation of degree N does not occur for $N < 6$. For higher N it was not tested in the experiments.

From the graphs in Section 3 it can be seen that for the photocopier, 3 negative experiences in a row are sufficient to get a negative trust (no matter how positive trust was), and for the travelling agency 2 negative experiences are sufficient. So, $N = 3$, resp. $N = 2$ apply in these cases, i.e., for the photocopier the property 'degree of trust dropping 3' holds, and for the traveling agency 'degree of trust dropping 2' holds. For the positive side, in both cases 2 positive experiences are sufficient to get trust positive again, so the property 'degree of trust gaining 2' holds for both cases.

An effect that does occur, however, in the photocopier context, is that after a series of negative experiences (see Figure 2), the level of trust does not become as high as in the case of no negative experiences (see Figure 1). More refined properties than the ones above can be formulated to account for this relative form of trust fixation. Notice that in the traveling agency context this effect does not occur.

6 Discussion

In papers on trust models the validity of models is usually based on intuition and common sense only. The extent to which models are correct is rarely verified experimentally, which is considered a lack in the literature on trust; cf. [14]. The research reported here contributes results of experiments on the dynamics of trust

over time depending on positive or negative experiences. In [8] a number of dynamic properties for trust dynamics were identified, indeed mainly on the basis of intuition and common sense. As a continuation of this work, now these properties have been verified in an experimental setting. Even if sometimes these properties may seem clear or self-evident at first sight, without any empirical verification they remain speculative.

By the above experimental results it can be shown that positive experiences can be identified that (usually) have an increasing or at least nondecreasing effect on trust, and negative experiences that have a decreasing or at least non-increasing effect. Here it appears easier to destroy trust than to build trust: the designed negative experiences show a stronger negative effect on trust than the positive effect shown by the designed positive experiences (see Tables 2 to 5).

Moreover, it is shown that trust can be flexible in the following sense: trust that has become positive can be made negative if a number of subsequent negative experiences occur, and trust that has become negative can become positive if a number of subsequent positive experiences occur. This may give an indication for handling trust in open system applications. If it is noticed that an agent has encountered a number of negative experiences, then it can be arranged that this agent should have a number of positive experiences first, for example, by paying extra attention to this agent and offer special services.

A number of issues can be investigated in more depth. First of all, it may be investigated in how far the same patterns can be found in other contexts. There may be contexts where, in contrast to the contexts used in our experiments, trust fixation does occur. For example, after a number of serious negative experiences with your partner in a relationship, trust may have gone forever, and not be (re)gained by positive experiences.

A related issue is to investigate further the notions of positive and negative experiences. What types of experiences qualify as such? One may be tempted to consider a positive (resp. negative) experience by definition as an experience that usually increases (resp. decreases) trust. However, a more independent definition would be more valuable. For example, the positive experiences are experiences leading to satisfaction of a certain type and level, the negative ones are experiences leading to frustration of a certain type and level? This also implies that a more fine-grained scale between positive and negative may be relevant. A more fine-grained scale would also enable to classify the experience stories based on the results in Tables 2 to 5 (which show that the stories differ in their impact on trust).

Another issue is the notion of trust used by the subjects. In our experiments a kind of folk-psychological trust notion was assumed in each of the subjects. A more fundamental approach would not use the word 'trust' in the experiments, but would define trust by the decisions (e.g., in relation to specific goals or tasks of the agent) that are made based on certain trust levels, and ask the subjects about these decisions, instead of their level of trust. However, in such an approach other factors (other than trust) affecting such decisions may have to be taken into account as well.

An important further question is in how far subjects are equal in the way in which they show trust dynamics. Our hypothesis is that substantial individual differences between subjects may exist, for example in initial trust attitudes (e.g., positive or negative initial bias), in trust-steadiness (how sensitive trust is w.r.t. new

experiences), or in (non-initial) positive or negative biases in trust dynamics. The number of subjects in our study was not low (238), but to obtain statistics for individual differences, a reasonable number per type of subject is needed. Further experimental work is planned to address this issue.

Acknowledgements

The authors are grateful to the anonymous reviewers putting forward a number of further suggestions that have been used to improve the text.

References

1. Barber, B., *The Logic and Limits of Trust*. New-Jersey: Rutgers University Press, 1983.
2. Castelfranchi, C. and Falcone, R., Principles of trust for MAS: cognitive anatomy, social importance, and quantification. In: *Proceedings of the International Conference on Multi Agent Systems 1998 (ICMAS'98)*. IEEE Press, 1998.
3. Derrida, J., *Given time. Counterfeit money*. Chicago: Chicago University Press, 1992.
4. Deutsch, M., Trust, trustworthiness, and the F scale. In: *Journal of Abnormal and Social Psychology*, vol.: 61, no. 1, 1960, pp. 138-140.
5. Elofson, G., Developing trust with intelligent agents: An exploratory study. In: *Proceedings of the First International Workshop on Trust*, 1998, pp. 9-19.
6. Gambetta, D., Foreword. In: (Gambetta, D.) *Trust: Making and breaking of cooperative relations*. Oxford: Basil Blackwell, 1988.
7. Good, D., Individuals, Interpersonal relations, Trust. In: (Gambetta, D.) *Trust: Making and breaking of cooperative relations*. Oxford: Basil Blackwell, 1988.
8. Jonker, C. M., and Treur, J., Formal analysis of models for the dynamics of trust based on experiences. In: F. J. Garijo, M. Boman (eds.), *Multi-Agent System Engineering, Proceedings of the 9th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, MAAMAW'99*. Lecture Notes in AI, vol. 1647, Springer Verlag, Berlin, 1999, pp. 221-232.
9. Lewis, D. and Weigert, A., Social atomism, holism and trust. In: *Sociological Quarterly*, 1985, pp. 445-471.
10. Liebetrau, A.M. (1983), *Measures of Association*, In (John L. Sullivan and Richard G. Niemi, eds.): *Sage University Paper series on Quantitative Applications in the Social Sciences*, vol 32. Beverly Hills and London: Sage Publications, Inc.
11. Luhmann, N., *Trust and Power*. New York: John-Wiley, 1979.
12. Luhmann, N., Familiarity, confidence, trust: problems and alternatives. In: (Gambetta, D.) *Trust: Making and breaking of cooperative relations*. Oxford: Basil Blackwell, 1988.
13. Rempel, J. K., Holmes, J. G. and Zanna, M. P., Trust in close relationships. In: *Journal of Personality and Social Psychology*, vol.: 49, 1985, pp. 95-112.
14. Smets, P., Wels, H. and Van Loon, J., Trust and co-operation in theory and everyday life. In: Smets, P., Wels, H. and Van Loon, J. (eds.), *Trust & Co-operation, Symbolic exchange and moral economies in an ages of cultural differentiation*. Amsterdam, The Netherlands: Het Spinhuis, 1999.
15. Wels, H. and Van Loon, J., Making sense of trust and co-operation: theoretical perspectives. In: Smets, P., Wels, H. and Van Loon, J. (eds.), *Trust & Co-operation, Symbolic exchange and moral economies in an ages of cultural differentiation*. Amsterdam, The Netherlands: Het Spinhuis, 1999.

16. Zalman, G. and Moorman, C., The importance of personal trust in the use of research.
In: *Journal of Advertising Research*, October-November, 1998, pp 16-24.