Agent based computational model of trust

Alexander Gorobets and Bart Nooteboom

ERIM REPORT SERIES *RESEARCH IN MANAGEMENT*	
ERIM Report Series reference number	ERS-2004-108-ORG
Publication	December 2004
Number of pages	12
Email address corresponding author	alex-gorobets@mail.ru
Address	Erasmus Research Institute of Management (ERIM) Rotterdam School of Management / Rotterdam School of Economics Erasmus Universiteit Rotterdam P.O.Box 1738 3000 DR Rotterdam, The Netherlands Phone: +31 10 408 1182 Fax: +31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

BIBLIOGRAPHIC DATA AND CLASSIFICATIONS		
Abstract	This paper employs the methodology of Agent-Based Computational Economics (ACE) to investigate under what conditions trust can be viable in markets. The emergence and breakdown of trust is modeled in a context of multiple buyers and suppliers. Agents adapt their trust in a partner, the weight they attach to trust relative to profitability, and their own trustworthiness, modeled as a threshold of defection. Adaptation occurs on the basis of realized profit. Trust turns out to be viable under fairly general conditions.	
Library of Congress Classification (LCC) LCC Webpage	Mission: HF 5001-6182	
	Programme: HF 5546-5548.6, HF 5549-5549.5	
	Paper:	HB846.3 Transaction costs
Journal of Economic Literature (JEL) JEL Webpage	Mission: M	
	Programme : M 10, L 2, M 12	
	Paper:	C 53 Forecasting and Other Model Applications L 14 Transactional Relationships
Gemeenschappelijke Onderwerpsontsluiting (GOO)		
Classification GOO	Mission: 85.00	
	Programme: 85.05, 85.08, 85.62	
	Paper:	31.80 Toepassingen van de wiskunde
Keywords GOO	Mission: Bedrijfskunde / Bedrijfseconomie	
	Programme: Organisatieleer, prestatiebeoordeling, personeelsbeleid	
	Paper: transactiekosten, vertrouwen, wiskundige modellen	
Free keywords	Agent-based computational economics, transaction costs, trust	

Agent based computational model of trust

Alexander Gorobets and Bart Nooteboom

Erasmus University Rotterdam, the Netherlands
alex-gorobets@mail.ru, b.nooteboom@uvt.nl

This paper employs the methodology of Agent-Based Computational Economics (ACE) to investigate under what conditions trust can be viable in markets. The emergence and breakdown of trust is modeled in a context of multiple buyers and suppliers. Agents adapt their trust in a partner, the weight they attach to trust relative to profitability, and their own trustworthiness, modeled as a threshold of defection. Adaptation occurs on the basis of realized profit. Trust turns out to be viable under fairly general conditions.

Keywords: agent-based computational economics, transaction costs, trust.

1. Introduction

The viability of trust between firms in markets is a much-debated issue (for a survey, see Nooteboom 2002). Economics, in particular transaction cost economics, doubts the viability of trust. Thus, refuting skepticism from TCE would make the strongest case for trust, and that is the project of this article. In this article we employ TCE logic, but we also deviate from TCE in two fundamental respects.

First, while TCE assumes that optimal forms of organization will arise, yielding maximum efficiency, we consider that problematic. The making and breaking of relations between multiple agents with adaptive knowledge and preferences may yield complexities and path-dependencies that preclude the achievement of maximum efficiency.

Second, while TCE assumes that reliable knowledge about loyalty or trustworthiness is impossible (Williamson 1975, 1993), so that opportunism must be assumed, we expect that under some conditions trust is feasible, by inference from observed behaviour, and that trustworthiness is viable, in yielding profit. To investigate this, the methodology of ACE enables us to take a process approach to trust (Gulati 1995, Zucker 1986, Zand 1972), by modeling the adaptation of trust and trustworthiness in the light of experience in interaction.

The analysis is conducted in the context of transaction relations between multiple buyers and suppliers, where buyers have the option to make rather than buy, which is the classical setting for the analysis of transaction costs.

We employ a model developed from an earlier model from Klos and Nooteboom (2001). In this model, agents make and break transactions relations on the basis of preferences, based on trust and potential profit.

The article proceeds as follows. First, further specification is given of technical details of the model. Next, we specify the experiments. The article closes with conclusions.

2. The model

2.1 Preference and matching

In the literature on trust distinctions are made between different kinds of trust, particularly between competence trust and intentional trust (see Nooteboom 2002). Intentional trust refers, in particular, to (presumed) absence of opportunism. That is the focus of TCE and also of the present article. We focus on the risk that a partner will defect and thereby cause switching costs. In our model trust is interpreted as a subjective probability that expectations will be fulfilled (Gambetta 1988), which here entails realization of potential profit. Thus, expected profit (E) would be: E = profitability.trust. In the model, agents are assumed to have differential preferences for different potential partners (cf. Weisbuch et al., 2000), on the basis of a generalized preference score:

$$\text{score}_{ij} = \text{profitability}_{ij}^{\alpha_i} \cdot \text{trust}_{ij}^{1-\alpha_i}\,, \qquad (1)$$

where: score_{ij} is the score i assigns to j, $\text{profitability}_{ij}$ is the profit i can potentially make 'through' j, trust_{ij} is i's trust in j and $\alpha_i \in [0, 1]$ is the weight i attaches to profitability relative to trust, i.e. the 'profit-elasticity' of the scores that i assigns; i may adapt the value of α_i from each timestep to the next.

At each time step, all buyers and suppliers establish a strict preference ranking over all their alternatives. Random draws are used to settle the ranking of alternatives with equal scores. The matching of partners is modeled as follows. On the basis of preferences buyers are assigned to suppliers or to themselves, respectively. When a buyer is assigned to himself this means that he makes rather than buys. In addition to a preference ranking, each agent has a 'minimum tolerance level' that determines which partners are acceptable. Each agent also has a quota for a maximum number of matches it can be involved in at any one time. A buyer's minimum acceptance level of suppliers is the score that the buyer would attach to himself. Since it is reasonable that he completely trusts himself, trust is set at its maximum of 1, and the role of trust in the score is ignored: $\alpha = 1$. The algorithm used for matching is a modification of Tesfatsion's (1997) deferred choice and refusal (DCR) algorithm and it proceeds in a finite number of steps, as follows:

1. Each buyer sends a maximum of o_i requests to its most preferred, acceptable suppliers. Because the buyers typically have different preference rankings, the various suppliers will receive different numbers of requests.
2. Each supplier 'provisionally accepts' a maximum of a_j requests from its most preferred buyers and rejects the rest (if any).
3. Each buyer that was rejected in any step fills its quota o_i in the next step by sending requests to next most preferred, acceptable suppliers that it has not yet sent a request to.
4. Each supplier again provisionally accepts the requests from up to a maximum of a_j most preferred buyers from among newly received and previously provisionally accepted requests and rejects the rest. As long as one or more buyers have been rejected, the algorithm goes back to step 3.

The algorithm stops if no buyer sends a request that is rejected. All provisionally accepted requests are then definitely accepted.

2.2 Trust and trustworthiness

Trust, taken as inferred absence of opportunism, is modelled on the basis of observed absence of defection. Following Gulati (1995), we assume that trust increases with the duration of a relation. As a relation lasts longer, one starts to take the partner's behavior for granted, and to assume the same behavior (i.e. commitment, rather than breaking the relation) for the future. An agent *i*'s trust in another agent *j* depends on what that trust was at the start of their current relation and on the past duration of their current relation:

$$t_i^j = t_{\text{init},i}^j + (1 - t_{\text{init},i}^j)\left(1 - \frac{1}{fx + 1 - f}\right), \qquad (2)$$

where t_i^j = agent *i*'s trust in agent *j*,

$t_{\text{init},i}^j$ = agent *i*'s initial trust in agent *j*,

x = the past duration of the current relation between agents *i* and *j, and*

f = trustFactor.

This function is taken simply because it yields a curve that increases with decreasing returns, as a function of duration *x*, with 100% trust as the limit, and the speed of increase determined by the parameter *f*.

In addition, there is a base level of trust, which reflects an institutional feature of a society. It may be associated with the expected proportion of non-opportunistic people, or as some standard of elementary loyalty that is assumed to prevail. If an agent *j*, involved in a relation with an agent *i*, breaks their relation, then this is interpreted as opportunistic behavior and *i*'s trust in *j* decreases; in effect, *i*'s trust drops by a percentage of the distance between the current level and the base level of trust; it stays there as *i*'s new initial trust in *j*, $t_{\text{init},i}^j$ until the next time *i* and *j* are matched, after which is starts to increase again for as long as the relation lasts without interruption.

The other side of the coin is, of course, one's own trustworthiness. This is modelled as a threshold τ for defection. One defects only if the advantage over one's current partner exceeds that threshold. It reflects that trustworthiness has its limits, and that trust should recognize this and not become blind (Pettit 1995, Nooteboom 2002). The threshold is adaptive, as a function of realized profit.

2.3 Costs and profits

In sum, the way profits are made is that buyers may increase returns by selling more differentiated products, and suppliers may reduce costs by generating production efficiencies.

There are two sources of production efficiency: economy of scale from a supplier producing for multiple buyers, and learning by cooperation in ongoing production relations. Economy of scale can be reaped only in production with general-purpose assets, and learning by cooperation only in production that is specific for a given buyer, with buyer-specific assets. This yields a link with the fundamental concept, in TCE, of 'transaction specific investments'. We assume a connection between the differentiation of a buyer's product and the specificity of the assets required to produce it. In fact, we assume that the percentage of specific products is equal to the percentage of dedicated assets. This is expressed in a variable $d_i \in [0, 1]$. It determines both the profit the buyer will make when selling his products and the degree to which assets are specific, which determines opportunities for economy of scale and learning by cooperation. Economy of scale is achieved when a supplier produces for multiple buyers. To the extent that assets are specific, for differentiated products, they cannot be used for production for other buyers. To the extent that products are general purpose, i.e. production is not differentiated, assets can be switched to produce for other buyers. In sum, economy of scale, in production for multiple buyers, can only be achieved for the non-differentiated, non-specific part of production, and economy by learning by cooperation can only be achieved for the other, specific part.

Both the scale and learning effects are modelled as follows:

$$y = \max\left(0, 1 - \frac{1}{fx + 1 - f}\right), \tag{3}$$

where

for the scale effect, f=scaleFactor, x is general-purpose assets of supplier j summed over all his buyers and scale efficiency $y = e_{s,j}$;

for the learning effect, f=learnFactor; x is the number of consecutive matches between supplier j and buyer i and learning efficiency $y = e^i_{l,j}$.

Function (3) expresses decreasing returns for both scale and experience effects. For the scale effect, it shows positive values along the vertical axis only for more than 1 general-purpose asset. This specifies that a supplier can be more scale-efficient than a buyer producing for himself only if the scale at which he produces is larger than the maximum scale at which a buyer might produce for himself. For the learning effect, a supplier's buyer-specific efficiency is 0 in their first transaction, and only starts to increase if the number of transactions is larger than 1. If a relation breaks, the supplier's efficiency due to his experience with the buyer drops to zero.

All this results in the following specification of profit. The number of general-purpose assets that a supplier j needs in order to produce for a buyer i, is equal to $(1-d_i)(1-e_{s,j})$. The number of buyer-specific assets that a supplier j needs, to produce for a buyer i, is equal to $d_i(1-e^i_{l,j})$. Thus, the profit that can potentially be made in a transaction between a buyer i and a supplier j is:

$$p_i^j + p_j^i = (1+d_i) - (d_i(1-e^i_{l,j}) + (1-d_i)(1-e_{s,j})). \tag{4}$$

The first part of (4) specifies returns and the second part specifies costs. It is assumed that the agents involved share the profit equally. In other words, we allow for defection but not for the threat of defection with the purpose of increasing one's share in jointly produced added value.

2.4 Adaptation

An agent is adaptive if 'the actions of the agent in its environment can be assigned a value (performance, utility, or the like); and the agent behaves in such a way as to improve this value over time' (Holland and Miller 1991: 365). In this model, agents adapt the values for $\alpha \in [0, 1]$ (weight attached to profit relative to trust) and τ [0, 0.5] (threshold of defection) from one time step to the next, which may lead to changes in the scores they assign to different agents. Here, adaptation takes place on the basis of past, realized profit. While τ

could conceivably rise up to 1, a maximum of 0.5 was set because initial simulations showed that otherwise relations would get locked into initial situations, with little switching. Note that this biases the model in favour of opportunism.

At each time step, each agent assigns a 'strength' to each possible value of α and τ. This expresses the agent's confidence in the success of using that particular value. The various strengths always add up to constants C_α and C_τ, respectively. At the start of each timestep, the selection of values for α and τ is stochastic, with selection probabilities equal to relative strengths, i.e. strengths divided by C_α and C_τ, respectively. The strengths of the values that were chosen for α and τ at the start of a particular timestep are updated at the end of that timestep, on the basis of the agent's performance during that timestep, in terms of realized profit: the agent adds the profit obtained during the timestep to the strengths of the values that were used for α or τ. After this, all strengths are renormalized to sum to C_α and C_τ again (Arthur 1993). The idea is that the strength of values that have led to high performance (profit) increases, yielding a higher probability that those values will be selected again. This is a simple model of 'reinforcement learning' (Arthur 1991, Arthur 1993, Kirman and Vriend 2000, Lane 1993).

2.5 The algorithm

The algorithm of the simulation is presented by the flowchart in Fig.1. This figure shows how the main loop is executed in a sequence of discrete time steps, called a 'run'. Each simulation may be repeated several times as multiple runs, to even out the influence of random draws in the adaptation process. At the beginning of a simulation, starting values are set for certain model parameters. The user is prompted to supply the number of buyers and suppliers, as well as the number of runs, and the number of timesteps in each run. At the start of each run, all agents are initialized, e.g. with starting values for trust, and selection probabilities for α and τ. In each timestep, before the matching, each agent chooses values for α and τ, calculates scores and sets preferences. Then the matching algorithm is applied. In the matching, agents may start a relation, continue a relation and break a relation. A relation is broken if, during the matching, a buyer does not send any more requests to the supplier, or he does, but the supplier rejects them.

After the matching, suppliers that are matched to buyers produce and deliver for their buyers, while suppliers that are not matched do nothing. Buyers that are not matched to suppliers produce for themselves ('self-matched', in 'make' rather than 'buy'). Afterward, all buyers sell their products on the final-goods market. Profit is shared equally with their supplier, if they have one. Finally, all agents use that profit to update their preference rankings (via α and τ), used as input for the matching algorithm in the next timestep. Across timesteps realized profits are accumulated for all buyers and suppliers, and all the relevant parameters are tracked.

Fig. 1. Flowchart of the simulation.

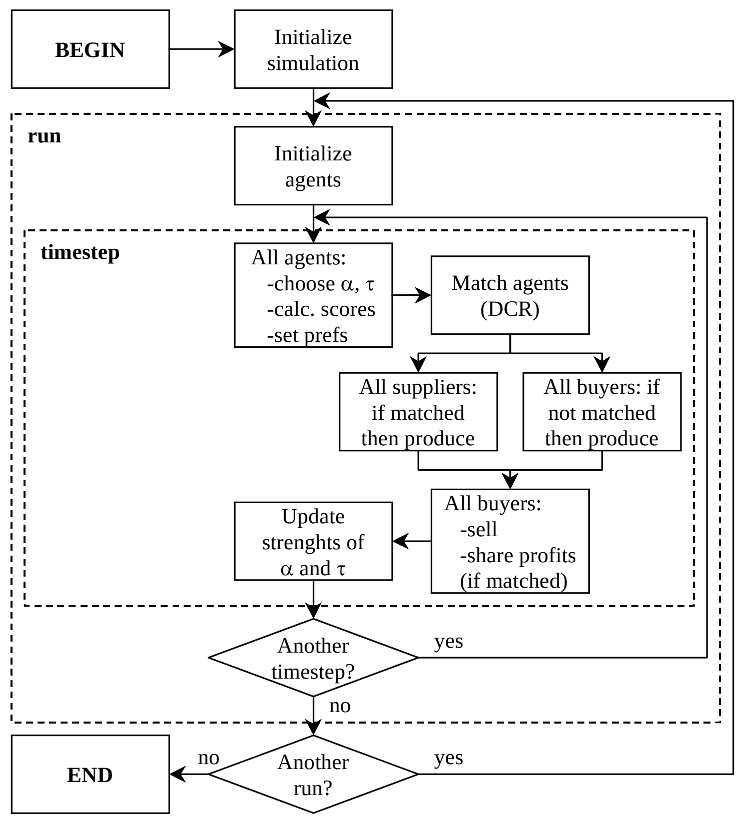

3. Experiments

3.1 Hypotheses

The goal of the experiments is to test the following hypotheses.

While according to TCE maximum efficiency can be achieved, we expect:

Hypothesis 1: due to complexities of interaction maximum efficiency can rarely be attained.

According to TCE, high asset specificity leads to more make rather than buy. We expect the same result in our extended framework, according to the same argumentation.

Hypothesis 2: When trust is low, higher asset specificity/differentiated products yields less outsourcing.

In agreement with TCE we expect:

Hypothesis 3: The more trust, the more collaboration in ‘buy’, rather than ‘make’.

More specifically:

Hypothesis 3a: The lower the weight attached to profit relative to trust (α), the more collaboration (buy rather than make), and the more loyalty (less switching).

Hypothesis 3b: The higher the threshold of defection (τ), the more collaboration (buy rather than make), and the more loyalty (less switching).

Counter to TCE we expect:

Hypothesis 4: Even in markets, where profit guides adaptation, high trust (low α; high τ) may be sustainable.

Recall that if during the matching between buyers and suppliers a buyer decides to 'buy' rather than 'make', he can follow two different strategies. One is an opportunistic scale strategy, where the buyer seeks a profit increase on the basis of economy of scale, by trying to find a supplier who serves more than one buyer. This entails much switching and less emphasis on loyalty and trust. The other strategy is the learning by cooperation strategy, seeking an increase of profit in ongoing relations. This entails less switching and more emphasis on loyalty and trust. Thus, in manipulating the strength of the scale effect relative to the effect of learning by cooperation, we can bias the model towards opportunism or loyalty. This interacts with the degree of asset specificity/specialization, since economy of scale applies only to general purpose assets, and learning by cooperation only to specific assets. Note that there is an overall bias towards the opportunistic scale strategy, in that economy of scale is immediate, thus yielding a more immediate return in profits, while learning by cooperation takes time to build up. Thus, we are stacking the odds in favour of the TCE theory that we criticize. However, this does seem to be a realistic feature, supporting the intuition that trust is more viable in a long-term perspective.

3.2 Model parameters

Each simulation run involves 12 buyers and 12 suppliers and continues for 100 timesteps. In order to reduce the influence of random draws, each run is repeated 25 times and results are averaged across all runs. Initially, results are also averaged for the two classes of agents: buyers and suppliers, in order to explore systematic effects. Each buyer's offer quota was fixed at 1, and each supplier's acceptance quota was set to 3. In previous experiments with each supplier *j*'s acceptance quota set to the total number of buyers, the system quickly settled in a state where all buyers buy from a single supplier. For this reason, suppliers were only allowed to have a maximum of three buyers. This limits the extent of the scale economies that suppliers can reach. A maximum number of buyers may be associated with competition policy setting a maximum to any supplier's market share.

To test Hypothesis 1, we analyse outcomes in terms of cumulative profit, to see to what extent maximum attainable profits are in fact realized, and how this varies across runs of the model. To test Hypothesis 2, we consider different values for the percentage of specific assets/differentiated products: d = 25, 45, and 65%. In addition, to test Hypotheses 3 and 4, we vary initial trust in the range 10, 50 and 90%, initial threshold for defection (τ) from 0 to 0.5, initial weight attached to profit relative to trust (α) from 0.0 to 1.0.

We present and discuss averages, across runs as well as agents (all buyers, all suppliers), as an indication of overall results. We present the results in the order of different starting values of trust. This reflects different institutional settings, from high to low trust 'societies'. Here, we can see to what extent those are stable or shift. In particular, the question is whether high initial trust can be sustained (Hypothesis 3), and whether perhaps distrust can evolve into trust.

3.3 High initial trust

First, we consider an initial situation of high, 90% trust across all agents. This reflects a society with a general assumption of high trustworthiness. First, we take intermediate initial expected values for α (0.5) and τ (0.25). Next to the variation of degree of specificity (d = 0.25, 0.45, 0.65), we vary the strength of economy of scale (scale factor sf) and learning by cooperation (lf), as follows:

- both medium strength (lf = sf = 0.5)
- high learning (lf = 0.9), medium scale (sf = 0.5). This is expected to favour a learning by cooperation strategy, with high loyalty
- medium learning by cooperation (lf = 0.5), high scale (sf = 0.9). This is expected to favour a scale strategy, with less loyalty.

The results are given in Table 2.

Table 2. Buyers' maximum normalized profits for different learn and scale factors

d	#S.per buyer	Buyers max. normal. Profit			
		l.f.=0.5;s.f.=0.5	l.f.=0.9;s.f.=0.5	l.f.=0.5;s.f.=0.9	l.f.=0.9;s.f.=0.9
0.65	1	0.98	0.994	0.978	0.99
0.45	1	0.91	0.92	0.89	0.90
0.25	1	0.80	0.81	0.77	0.78

High initial trust is sustained, and in fact increases from 0.9 to the maximum of 1.0

Table 2 supplies maximum normalized profit actually achieved in the course of time. It is obtained by dividing the buyers' profits by the maximum attainable (theoretical) profit they can potentially make in each experiment, which depends on differentiation, and on learn and scale factors. Maximum attainable profit is the profit a buyer makes when he has an infinite relation with a supplier who produces to an infinite number of buyers. The latter factor is limited because suppliers' acceptance quota is set to 3. Usually maximum actual profit is achieved at the last steps of simulation because of adaptation processes in relations between buyers and suppliers. At the start point the normalized profit is about 52% for high d and 61% for low d.

Table 2 shows that buyers perform better for high d then for low d. This is partly built-in: differentiated products are assumed to give higher profit margins. The outcome also results from the effect of economy of scale for general purpose (non-specific assets) and learning by cooperation for specific assets. Maximum scale effect is achieved when d is low. Here, the maximum arises in a situation where which 12 buyers together buy from only 4 suppliers (each, i.e. one third of all suppliers producing for the maximum of three buyers). Because the optimal network configuration, where suppliers produce for 3 buyers, emerges rarely, buyers organize closer to the optimum when d is higher. Maximum profit is approached more closely when products are more differentiated, because then buyers are less sensitive to the optimal configuration of network between agents. Then, a buyer has less scope for increased efficiency by getting into an arrangement of one supplier producing for him as well as two other buyers.

Now we turn to the hypotheses. Maximum actual profit never achieves maximum attainable profit, which confirms Hypothesis 1. The high levels of initial trust are sustained, and in fact increase, on average, from 0.9 to the maximum of 1.0, which confirms Hypothesis 3. Other results, not specified in the table, show that here there is maximum outsourcing: each buyer has a supplier, even at high levels of asset specificity (d=0.65). This is in agreement with Hypothesis 2. At high levels of trust, outsourcing takes place even at high levels of asset specificity. For all levels of asset specificity (d), in each run at least one supplier produced for the maximum of 3 buyers, on average across runs 10% of suppliers did this, 15% of suppliers produced for 2 buyers, 40% for 1 buyer, and 35% for 0 buyers. The results indicate that in this high-trust society buyers follow the strategy of learning by cooperation and loyalty for all d, without switching between suppliers, even for the low value d=0.25, where only 25% of assets are subject to learning by cooperation.

So far, we assumed intermediate levels for the initial weight attached to profit (α) and for the threshold of defection (τ). Now we analyze the effects of varying those values: α = 0.0 and 1.0; τ = 0.0 and 0.5. Learn and scale factors are fixed at the average level, i.e. 0.5. The purpose of this experiment is to explore the effects on outsourcing and profit of initial values of focus on profit and threshold of defection. According to hypotheses 3a and 3b, a low value for α and a high value for τ would favour a trust strategy of durable collaboration and loyalty, but the net effects are difficult to judge a priori. A high threshold of defection would tend to favour outsourcing and stable relations, particularly when initial trust is also high, provided such relations are profitable. However, a high weight attached to profit may prevent outsourcing, whereby effects of loyalty do not get a chance to arise. Here, we use the power of the ACE methodology to explore outcomes of processes that are too complex to compute analytically. The results are given in Table 3. Here, we also supply the average number of suppliers per buyer, as an indicator of the extent of outsourcing.

Table 3. Buyers' maximum normalized profits for different α and τ

d	Buyers max. norm. Profit & #Suppl. per buyer							
	α=0.0; τ=0.0		α=0.0; τ=0.5		α=1.0; τ=0.0		α=1.0; τ=0.5	
0.65	0.96	1	0.99	1	0.96	0	0.99	0.3
0.45	0.91	1	0.91	1	0.85	0	0.92	0.4
0.25	0.80	1	0.80	1	0.82	0.5	0.84	0.6

When α=0, agents put their emphasis on trust and follow the strategy of learning by cooperation for all d. The distribution of suppliers between buyers in this case is the same as before (Table 2). Each buyer has ongoing transactions with the same supplier but when loyalty is equal to zero (τ=0) buyers sometimes break relations with suppliers for high d because then profit doesn't exceed the level of when they make. These buyers try to switch to other suppliers but they don't succeed because all agents are concentrated on trust built up in the past of their current relation. Opportunistic buyers then return to their initial partners and as a result they lose in profit slightly, for high d, because of switching costs. If loyalty is high (τ=0.5) there is no switching for any level of d, and agents try to generate as much profit as possible in stable relations by using the advantage of loyalty and trust, in learning by cooperation.

When α=1, agents focus on profitability rather than on trust, and buyers follow two strategies simultaneously: some of them buy from suppliers and others make themselves.
If τ=0.0 approximately half of buyers have suppliers for d=0.25 and these buyers follow the scale strategy, seeking a supplier who already serves two buyers, and trying to match with him. As a result, in this case 17% of all suppliers produce for three buyers. For d=0.45 and d=0.65 buyers prefer to make themselves, mostly because outsourcing is only preferred as relations with suppliers last longer and generate economies of learning, but this is unlikely to

happen at zero loyalty. However, because of high initial trust buyers try to reach suppliers sometimes and then lose profit a little because of switching costs. If τ=0.5 the proportion of buyers who have suppliers increases for all d: 60% of buyers have suppliers for d=0.25, 40% for d=0.45 and 30% for d=0.65. However, the distribution of suppliers over buyers is different for all d. When d=0.25 approximately 20% of suppliers produce for three buyers and therefore profit is higher than in the case with τ=0.0. When d=0.45 about 12% of suppliers produce for three buyers and 5% of suppliers produce for one buyer and when d=0.65 suppliers produce only for one buyer and it is about 30% of them. Therefore for low and average d more buyers follow the scale strategy because high loyalty allows them to keep stable relations with matched suppliers and generate higher profit than in the case with zero loyalty. For high d one part of buyers (70%) produce themselves and other part (30%) follow the strategy learning by cooperation because economies of learning are more important than scale effect.

In sum, overall the results confirm hypotheses 3a and 3b, but there are complicated interaction effects that can be calculated only by simulation. Counter to hypothesis 3a, a high weight attached to profitability relative to trust (α) does not always favour opportunism. Once a buyer accumulates efficiency by learning by cooperation, an emphasis on profit also favours loyalty, not to lose the benefit of learning by cooperation.

3.4 Average and low initial trust

Now we turn to 'societies' with a lower, average level of initial trust: X = 0.5. Learn and scale factors are again fixed at the average level, i.e. 0.5. The main outcome here is that buyers make for high and average levels of specific assets (d), and buy only for low levels. This confirms Hypothesis 2. The results are specified in Table 4.

Table 4. Buyers' maximum normalized profits for average initial trust

d	#Suppl. Per buyer	Buyers max. normal. Profit
0.65	0	0.99 trust remains at 0.5
0.45	0	0.87 trust remains at 0.5
0.25	1	0.80 trust increases to 1.0

At first sight, it may seem counter-intuitive that trust increases from an average to the highest level under low specific assets (d = 0.25), since then the effect of learning by cooperation is lowest, so that the rewards of a trust strategy seem lowest. The explanation is as follows. Under average trust, suppliers are more attractive than buyers consider themselves only for low d, because potential losses in a case of switching are smaller for low d than for high d. For high levels of specificity, buyers never enter into relations with outside suppliers, and thus never profit from collaboration and forego opportunities for the build-up of trust.

Compared with the corresponding case in the high trust world (first column, Table 2), normalized profits are the same for high and low values of d, but lower for intermediate values. The network configuration of suppliers and buyers for low d is the same as in the case of high initial trust: 10% of suppliers produce for 3 buyers, 15% of suppliers produce for 2 buyers, 40 % for 1 buyer, 35% for 0 buyers. Buyers follow the learning by cooperation strategy in ongoing relations without switching.

In the case of low initial trust, i.e. X=0.1, buyers produce themselves (have no suppliers) even for a low level of specific assets. This again confirms hypothesis 2. The results are specified in Table 5. The result is a drop of normalized profits for low d, compared to the medium and high trust cases. All opportunities for learning by cooperation in collaboration are foregone.

Table 5. Buyers' maximum normalized profits for low initial trust

d	#Suppl. Per buyer	Buyers max. normal. Profit
0.65	0	0.99 trust remains at 0.1
0.45	0	0.87 trust remains at 0.1
0.25	0	0.63 trust remains at 0.1

Overall, the results show that under high trust (Table 2, column 1) maximum realized normalized profit is higher than under low trust (Tables 4 and 5) for intermediate and low levels of asset specificity (0.91 and 0.80 vs. 0.87 and 0.63), and marginally lower for high asset specificity (0.98 vs. 0.99). Overall, this confirms the central hypothesis 4 that trust can well be viable in markets.

4. Conclusions

The overall outcome is that in interactions between agents maximum efficiency is seldom achieved and that both trust and opportunism can be profitable, but they go for different strategies. This suggests that there may be different societies, going for different strategies, of switching or of loyalty, which settle down in their own self-sustaining systems.

High initial trust dictates buy relative to make for all levels of specific investments. For high specific investments, buyers' maximum profit is almost the same as in the cases of average or low initial trust. Low initial trust imposes make relative to buy, but buyers' maximum profits for low specific investments are smaller than in the case of high initial trust. Overall, across all parameter settings, profit tends to be higher under high than under low trust.

In addition to the expected results, incorporated in the hypotheses, the model yields a few unanticipated results. One is that buyers organize closer to maximum possible efficiency for high levels of specific investments. The reason is that for low levels of specific investments there is more scope for scale effects, but this is difficult to attain by having suppliers supply to the maximum number of buyers. A strong effect of learning by cooperation, a high weight attached to trust, and high loyalty favour the learning by cooperation strategy for high levels of specific investments, while a high weight attached to profit and high loyalty favour the scale strategy for low and average levels of specific investments. Finally, it is not always the case that a high weight attached to profitability relative to trust does not always favour opportunism. Once a buyer begins to profit from learning by cooperation, an emphasis on profit may also lead to loyalty, in an ongoing relationship.

References

Arthur, W. Brian (1991), Designing economic agents that act like human agents: A behavioral approach to bounded rationality, *American Economic Review*, 81/2, 353-359.

Arthur, W. Brian (1993), On designing economic agents that behave like human agents, *Journal of Evolutionary Economics*, 3/1, 1-22.

Gambetta, D. (ed.) (1988), *Trust: The Making and Breaking of Cooperative Relations*, Basil Blackwell: Oxford.

Gulati, R. (1995), Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances, *Academy of Management Journal*, 38/1, 85-112.

Holland, J.H. and J.H. Miller (1991), Artificial adaptive agents in economic theory, *American Economic Review*, 81/2, 365-370.

Kirman, A.P. and N.J. Vriend (2000), Evolving market structure: An ACE model of price dispersion and loyalty, *Journal of Economic Dynamics and Control*, ...

Klos, T.B. and B. Nooteboom, (2001), Agent based computational transaction cost economics, *Journal of Economic Dynamics and Control*, 25, 503-526.

Lane, D.A. (1993), Artificial worlds and economics, part II, *Journal of Evolutionary Economics*, 3/3, 177-197.

Lewicki, R.J. and B.B. Bunker (1996), Developing and maintaining trust in work relationships, in R.M. Kramer and T.R. Tyler (eds.), *Trust in organizations: Frontiers of theory and research*, Sage Publications: Thousand Oaks, 114-139.

Nooteboom, B. (2002), *Trust: forms, foundations, functions, failures and figures*, Edward Elgar: Cheltenham UK.

Pettit, Ph. (1995), The virtual reality of homo economicus, *The Monist*, 78/3, 308-329.

Tesfatsion, L.S. (1997), A trade network game with endogenous partner selection, in: H.M. Amman, B. Rustem and A.B. Whinston (eds.), *Computational Approaches to Economic Problems*, *Advances in Computational Economics*, Vol. 6, Kluwer: Dordrecht, 249-269.

Weisbuch, G., A.P. Kirman and D.K. Herreiner (2000), Market organisation and trading relationships, *Economic Journal*.

Williamson, O.E. (1975), *Markets and Hierarchies: Analysis and Antitrust Implications*, The Free Press: New York.

Williamson, O.E. (1993), Calculativeness, trust, and economic organization, *Journal of Law and Economics*, 36/1, 453-486.

Zand, D.E. (1972), Trust and managerial problem solving, *Administrative Science Quarterly*, 17/2, 227-239.

Zucker, L.G. (1986), Production of trust: Institutional sources of economic structure, 1840-1920, in: B.A. Staw and L.L. Cummings (eds.), *Research in Organizational Behavior*, Vol. 8, JAI Press: Greenwich, Conn, 53-111.

Publications in the ERIM Report Series Research* in Management

ERIM Research Program: “Organizing for Performance”

2004

Learning And Governance In Inter-Firm Relations
Bart Nooteboom
ERS-2004-003-ORG
http://hdl.handle.net/1765/1122

Organisational Learning And Multinational Strategy
Bart Nooteboom
ERS-2004-004-ORG
http://hdl.handle.net/1765/1123

Density And Strength Of Ties In Innovation Networks: A Competence And Governance View
Bart Nooteboom and Victor A. Gilsing
ERS-2004-005-ORG
http://hdl.handle.net/1765/1124

Innovation, learning and cluster dynamics
Bart Nooteboom
ERS-2004-006-ORG
http://hdl.handle.net/1765/1125

Empirical Tests Of Optimal Cognitive Distance
Stefan Wuyts, Massimo G. Colombo, Shantanu Dutta, and Bart Nooteboom
ERS-2004-007-ORG
http://hdl.handle.net/1765/1126

Entrepreneurship in Transition: Searching for governance in China’s new private sector
Barbara Krug and Hans Hendrischke
ERS-2004-008-ORG
http://hdl.handle.net/1765/1128

Exploring Emotional Competence: Its effects on coping, social capital, and performance of salespeople
Willem Verbeke, Frank Belschak and Richard P. Bagozzi
ERS-2004-014-ORG
http://hdl.handle.net/1765/1174

The Impact of Business Ownership Change on Employee Relations: Buy-outs in the UK and the Netherlands
Hans Bruining, Paul Boselie, Mike Wright and Nicolas Bacon
ERS-2004-021-ORG
http://hdl.handle.net/1765/1263

∗ A complete overview of the ERIM Report Series Research in Management:
https://ep.eur.nl/handle/1765/1

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship

Towards a Dynamic (Schumpeterian) Welfare Economics
Wilfred Dolfsma
ERS-2004-026-ORG
http://hdl.handle.net/1765/1264

The Three-Step Test-Interview (TSTI): An observational instrument for pretesting self-completion questionnaires
Tony Hak, Kees van der Veer and Harrie Jansen
ERS-2004-029-ORG
http://hdl.handle.net/1765/1265

Measuring the Knowledge Base of an Economy in terms of Triple-Helix Relations among 'Technology, Organization, and Territory'
Loet Leydesdorff, Wilfred Dolfsma & Gerben van der Panne
ERS-2004-034-ORG
http://hdl.handle.net/1765/1300

Paradoxes of Modernist Consumption – Reading Fashions
Wilfred Dolfsma
ERS-2004-035-ORG
http://hdl.handle.net/1765/1330

Some Economics of Digital Content
Wilfred Dolfsma
ERS-2004-036-ORG
http://hdl.handle.net/1765/1331

Learning Opportunities And Learning Behaviours Of Small Business Starters: Relations With Goal Achievement, Skill Development, And Satisfaction
Marco van Gelderen, Lidewey van der Sluis & Paul Jansen
ERS-2004-037-ORG
http://hdl.handle.net/1765/1429

The Process Of New Service Development – Issues Of Formalization And Appropriability
Wilfred Dolfsma
ERS-2004-051-ORG
http://hdl.handle.net/1765/1445

On And Off The Beaten Path: How Individuals Broker Knowledge Through Formal And Informal Networks
Rick Aalbers, Wilfred Dolfsma & Otto Koppius
ERS-2004-066-LIS/ORG
http://hdl.handle.net/1765/1549

Governance Modes For Systemic Innovation. Service Development In Mobile Telecommunications
J. van den Ende and F. Jaspers
ERS-2004-067-ORG
http://hdl.handle.net/1765/1539

Performance Management: A model and research agenda
Deanne N. den Hartog, Paul Boselie & Jaap Paauwe
ERS-2004-068-ORG
http://hdl.handle.net/1765/1779

Human Resource Function Competencies In European Companies
Paul Boselie and Jaap Paauwe
ERS-2004-069-ORG
http://hdl.handle.net/1765/1451

"Best Practices…in spite of Performance" Just a matter of Imitation?
Jaap Paauwe and Paul Boselie
ERS-2004-070-ORG

Web-Based Organizing In Traditional Brick-And-Mortar Companies: The Impact On HR
Jaap Paauwe, Elaine Farndale and Roger Williams
ERS-2004-071-ORG
http://hdl.handle.net/1765/1609

Longevity in services: the case of the Dutch warehousing companies 1600-2000
Hugo van Driel, Henk Volberda and Sjoerd Eikelboom
ERS-2004-072-STR/ORG
http://hdl.handle.net/1765/1571

Honing and Framing Ourselves (Extreme Subjectivity and Organizing)
Sławomir Magala
ERS-2004-076-ORG
http://hdl.handle.net/1765/1583

Cross-cultural compromises, multiculturalism and the actuality of unzipped Hofstede
Sławomir Magala
ERS-2004-078-ORG
http://hdl.handle.net/1765/1584

Perceptions about the ISO 9000 (2000) quality system standard revision and its value: The Dutch experience
T. van der Wiele, J. Iwaarden, R. Williams and B. Dale
ERS-2004-081-ORG
http://hdl.handle.net/1765/1736

Mystery shopping: A tool to develop insight into customer service provision
M. Hesselink, J. van Iwaarden and T. van der Wiele
ERS-2004-082-ORG
http://hdl.handle.net/1765/1737

A transparent role of information systems within business processes: A case study
Menno Verboom, Jos van Iwaarden and Ton van der Wiele
ERS-2004-083-ORG
http://hdl.handle.net/1765/1738

Understanding the development of temporary agency work in Europe
Bas Koene, Jaap Paauwe and John Groenewegen
ERS-2004-086-ORG
http://hdl.handle.net/1765/1803

Central Unification versus Local Diversity: China's Tax Regime, 1980s-2000s
Ze Zhu and Barbara Krug
ERS-2004-089-ORG
https://ep.eur.nl/handle/1765/1787

The evolution of high-technology in China after 1978: Towards technological entrepreneurship
M.J. Greeven
ERS-2004-092-ORG
http://hdl.handle.net/1765/1785

Effects of multiple network ties Knowledge transfer and sharing in a network: The effects of multiple ties
Irma Bogenrieder
ERS-2004-093-ORG
http://hdl.handle.net/1765/1781

Multiple Inclusion and Community Networks
Irma Bogenrieder and Peter van Baalen
ERS-2004-094-ORG
http://hdl.handle.net/1765/1782

The Performance Of Team Start-Ups In The First Phases Of The Life Course
Erik Stam And Veronique Schutjens
ERS-2004-097-ORG
http://hdl.handle.net/1765/1784

How To Be Better Prepared For A Paradigm Shift In Economic Theory, And Write Better Articles In The Meantime
Pat Welch and Wilfred Dolfsma
ERS-2004-101-ORG
http://hdl.handle.net/1765/1811

Is China a Leviathan?
Ze Zhu And Barbara Krug
ERS-2004-103-ORG
http://hdl.handle.net/1765/1821

Market and Society: How do they relate, and contribute to welfare?
Wilfred Dolfsma, John Finch and Robert McMaster
ERS-2004-105-ORG
http://hdl.handle.net/1765/1824

Determinants Of Entrepreneurship In Europe
Isabel Grilo And Roy Thurik
ERS-2004-106-ORG
http://hdl.handle.net/1765/1823

Agent based computational model of trust
Alexander Gorobets and Bart Nooteboom
ERS-2004-108-ORG

A Matter of Life and Death: Innovation and Firm Survival
Elena Cefis and Orietta Marsili
ERS-2004-109-ORG

Using Networks For Changing Innovation Strategy: The Case Of IBM
Koen Dittrich, Geert Duysters and Ard-Pieter de Man
ERS-2004-111-ORG
http://hdl.handle.net/1765/1836

China's emerging tax regime: Devolution, fiscal federalism, or tax farming?
Barbara Krug, Ze Zhu and Hans Hendrischke
ERS-2004-113-ORG