Trust-Based Inter-Temporal Decision Making: Emergence of Altruism in a Simulated Society

Tibor Bosse, Martijn C. Schut, Jan Treur, and David Wendt

Vrije Universiteit Amsterdam, Department of Artificial Intelligence,
De Boelelaan 1081a, 1081 HV, The Netherlands
{tbosse, schut, treur, dwendt}@cs.vu.nl
http://www.cs.vu.nl/~{tbosse, schut, treur, dwendt}

Abstract. This paper contributes to the analysis of the question how altruistic behaviour can be in an agent's own interest. The question is addressed by involving a temporal dimension in which altruist behaviour at one point in time can be in the agent's interest at a future point in time, depending on the environment. The claim is that to be able to make reasonable decisions, an agent needs a cognitive system for intertemporal decision making, in relation to a model of the environment. To address this, a society of agents has been modelled, simulated and analyzed. Some of the agents have a cognitive system including a model of the environment based on a dynamic model for trust in other agents. This environment model is combined with a model for intertemporal decision making. Agents with this cognitive system show more altruistic behaviour, they get a larger social network, and become in the end more healthy than agents without such a cognitive system.

1 Introduction

A basic assumption in the evolutionary explanatory framework is that an organism's behaviour serves its own interests, to improve its own well-being and production of offspring. For the explanation of the development of altruistic behaviour from an evolutionary perspective, one easily encounters a paradox; see, e.g., [Sober and Wilson, 1998, pp. 17-23]. As by definition altruistic behaviour is behaviour that is against one's own interests, this paradox can be formulated as: 'an organism serves its interests by behaviour which is against its interests'. One way to solve this paradox is by widening the scope in the temporal dimension. Then the occurrences of the concept 'interest' in the paradox get different time points attached, as follows: altruistic behaviour serves the organism's interests at a future time point by the organism's behaviour which is against its interests at the present time point. So, the organism's present behaviour is seen as an investment to obtain future revenues for the organism itself. As long as the future revenues are at least as important for the organism as the present investment, this may work out fine. It is this approach that is analysed further in this paper; see also [Dennett, 2003, Chapter 7].

In this case a basic assumption is that the environment of an organism has the potentiality or regularity to provide future revenues in return for present investments. This is a nontrivial characteristic of an environment, that often depends on the presence of other organisms in the environment. For example, for species developing agriculture, the activity of sowing in the present, depending on the potential of the seed, leads to growth of food or other products that are in the organism's interest. Another example, which is taken as a case study in this paper, is that other agents are present in the environment that offer the future returns when they are favoured by an agent, depending on their own intertemporal decision making.

Godfrey-Smith [1996, p. 3] relates environmental complexity to the development of cognition, as briefly formulated in his Environmental Complexity Thesis as: 'The function of cognition is to enable the agent to deal with environmental complexity'. For the case considered here, the agent needs a cognitive system that is able to make a decision where a current investment has to be compared to a future revenu. So, it needs cognitive facilities to predict future revenues based on the present world state and the world's regularities, and to compare such predicted future revenues to investments to be made in the present. These processes require nontrivial cognitive capabilities, the more so as the world's regularities usually have a probability aspect in them, that also has to be accounted for in the decision. These cognitive processes are usually called 'intertemporal decision making'; cf. [Loewenstein and Elster, 1992]. To cope with the world's risks that in some cases predicted revenues will not come true, in such decision making the future revenues have to be estimated higher than the present investment, for example, by taking into account a certain interest rate. In the literature on intertemporal decision making, the environmental regularity or probability to indeed provide revenues in return usually is not modelled in a detailed manner, and not adapted on the basis of the agent's experiences. Experiments and models often focus on one subject and its expectations, and do not address how these relate to the real environment. In fact, to estimate the risk of not getting the future revenues in return, the model of intertemporal decision making for the subject should be combined with an environment-dependent model describing how based on its experiences the subject estimates when the environment indeed returns revenues for investments of the subject. In this way the agent can learn and adapt itself to the world's regularities or potentialities.

In this paper, these issues are analyzed and tested by creating an artificial society. As part of the model, for any of the agents also the environment is modelled in a detailed manner as the rest of the society. By formal analysis and simulation it is investigated how agents endowed with a cognitive model for intertemporal decision making can choose for altruistic behaviour by providing services (for free) to other agents in the present and provide revenues in their own interest in the future. The guarantee or probability that revenues indeed are returned by the environment, depends in this case on other agents in the environment receiving the services, that in the future may or may not provide services in return. To estimate the risk of not getting the future revenues in return, the model of intertemporal decision making is combined with an environment model, which in this case is a model for evolution of trust in other agents based on experiences with them (adopted from [Jonker and Treur, 1999, 2005]). If the agent experiences over time that another agent does not provide services, the trust in this agent becomes lower; if it does provide services, trust becomes higher. Having such a dynamic environment model enables the agent to become better adapted to the environment. One of the main properties to verify is whether indeed agents with a cognitive system for trust-based intertemporal decision making do well over time, in contrast to agents that lack such a cognitive system.

To create an artificial society of multiple agents performing trust-based inter-temporal decision making, several modelling approaches have been used. First, the LEADSTO simulation environment [Bosse et al., 2005] was used for rapid prototyping, i.e., to create a high-level declarative specification of the simulation model, involving a small number of agents (six in this case). When this specification turned out to show coherent behaviour, it was used as a blueprint to create a large-scale version of the model in the NetLogo environment [Wilensky, 1999]. Finally, to analyse the simulation model and its outcomes, the predicate logical Temporal Trace Language [TTL, see Bosse et al., 2006] was used to specify a number of global dynamic properties that are relevant for the domain in question. An example of such a property is "agents that anticipate on the future eventually become fitter". Using the TTL Checker Tool, these properties have been checked automatically against the traces resulting from the LEADSTO simulations and NetLogo simulations.

In Section 2 the model for trust-based inter-temporal decision making is presented in detail. Section 3 describes the simulation model, which was designed and formally specified at a conceptual, temporal logic level in the form of local dynamic properties (in LEADSTO) for the basic mechanisms, and implemented in NetLogo. Section 4 discusses some of the simulation results. In Section 5 global dynamic properties are formulated and formalised in TTL, and by further formal analysis logical interlevel relations between such a global property and the local properties specifying the basic mechanisms are established, thus providing insight in how the global property depends on these local properties. Section 6 is a discussion.

2 Trust-Based Intertemporal Decision Making

The adaptation mechanism introduced in this paper is a decision theoretic function that includes a dynamic trust factor. The function is used for inter-temporal decisions, i.e., those decisions that compare investments or revenues at one point in time to those at another point in time; e.g., deciding whether to buy a car with a not-so-nice colour which you can take home immediately or waiting half a year for a car in your desired colour. Often uncertainty is involved about events that are more distant in time; they may depend on the environment's dynamics which can be unpredictable. For example, the event of making an investment in the present is certain, while an expected revenue in the future may be uncertain, and depend on the environment. The particular decisions concerned here involve the cooperation between agents, where the inter-temporal aspect is the expected reciprocity of cooperation: if I help you now, you will help me later. For such patterns to occur, as part of the agent's cognitive system an adequate decision model is essential. Such a model should include an environment model to predict future revenues upon present investments. In our model, the environment model has the form that agents maintain trust in other agents: they adapt their trust in other agents' willingness to help based on experienced (non)cooperations over time. Thus our agent model consists of two main parts: one part concerns the intertemporal decision making (to cooperate or not), and a second part concerns the updating of trust in other agents based on experiences. Both are described below.

We are concerned with the following decision situation. Consider a set of agents, where each agent is working towards its own benefit. At every point in time, agents are able to ask each other whether they want to cooperate or not. Let us assume that agent *x* requests agent *y* to cooperate. Such a cooperation has cost *c* for agent *y* and provides reward *f* to agent *x*. In response to this request, agent *y* evaluates the benefit of cooperation (i.e., calculates if the reward outweighs the cost). Based on this evaluation, agent *y* either accepts the request (agent *x* receives *f* and agent *y* pays *c*) or declines the request (agents *x* and *y* neither receive nor pay anything). Note that the evaluation function contains a reciprocity factor: if agent *y* cooperates with agent *x* then agent *y* can reasonably assume that agent *x* will later return the favour.

Inter-temporal choice is a decision in which the realisation of outcomes may lie in the imminent or remote future. Recently, inter-temporal choice has caught the attention in the literature on behavioural decision-making [Loewenstein and Elster, 1992]. Before this, results on the subject were mainly due to the research contributions in related fields, like economics and animal psychology. The standard agent model for decision-making over time is a framework called time discounting [Loewenstein and Elster, 1992], which works according to the same principles as interest that one receives on a bank account: I calculate a delayed reward back to its current value based on the interest that I would receive for it.

We use a similar agent model for inter-temporal decision making here, extended to our particular decision situation (involving reciprocity for cooperation) in two main ways. Firstly, the decisions involve an explicit model the agent has of (regularities in) the environment, in this case incorporating the other agents. This results in parameters for *trust* of the agent in other agents. As explained below, the value of this parameter evolves over time as a consequence of monitoring (regularities in) the environment over time, i.e., the experienced (non)cooperations. Secondly, the individual decisions are concerned with choosing between (1) a possible reward in the remote future and (2) having no immediate cost, rather than choosing between an immediate and delayed reward (as investigated traditionally in time discounting). In the model, the discounted value $f_{discounted}$ of a future reward is calculated by: $f_{discounted} = f * 2^{(1 - \beta)(t / n * (1 - (tr + 1) / 2))}$, where

f : REAL	= future reward,
$\beta \in [0,1]$	= discount factor,
t : INTEGER	= duration after which the future reward is received,
n : INTEGER	= duration of cooperation, and
$tr \in [-1,1]$	= trust in the agent who asks you to cooperate.

If the discounted future reward evaluates higher than (or equal to) the current (immediate) cost, the agent decides to cooperate. In other words:

If ($f_{discounted} \geq c$), then cooperate, else do not cooperate

where c : REAL is the immediate cost.

The next sections show how it was tested how agents that use this decision function develop in a multi-agent society. The prediction is that these agents will show altruistic behaviour, will establish a larger social network than agent without such a decision function (i.e., agents that are not able to estimate the future reward, and thus never cooperate), and will eventually get a higher fitness.

Agents adjust their trust values in other agents according to the following principle: if I ask you to cooperate and you accept, then I increase my trust in you; if you decline, then I decrease my trust in you. For modeling such adaptation of trust over time, we use a trust function that was presented in [Jonker and Treur, 1999, 2005]. This function, as applied here, takes the response of an asked agent (accept/decline) to determine how to revise the trust value. Such a response $e \in [-1,1]$ evaluates to 1 if the agent accepts or -1 if the agent declines. A scaling factor $\delta \in [0,1]$ (which is constant throughout the experiments) determines how strongly an agent is committed to its trust values: a higher δ means that an agent puts much weight on its current trust value and lets an (non)cooperation experience not weigh so heavily; and vice versa. In the model, when the outcome of a request to cooperate is known, we calculate the trust value tr_{new} as follows: $tr_{new} = \delta * tr + (1 - \delta) * e$, where

$tr \in [-1,1]$	= current trust value,
$\delta \in [0,1]$	= scaling factor (constant),
$e \in [-1,1]$	= the response of the agent who you asked to cooperate.

Thus, each agent maintains a list of trust values for all other agents in the environment. The model also includes a *cooperation threshold* $ct \in [-1,1]$ such that agent *x* only requests cooperation with agent *y* if trust of agent *x* in agent *y* is above this threshold.

3 Simulation Model

This section describes the simulation model that was used to investigate the impact of trust-based inter-temporal decision making on fitness. In accordance with the explanation in Section 2, a society of agents is analyzed, where each agent can ask the other agents for a favour. An example of such a favour is that an agent asks another agent for help during a house removal event. Figure 1 briefly sketches the different possible scenarios in one round of interaction between two agents x and y.

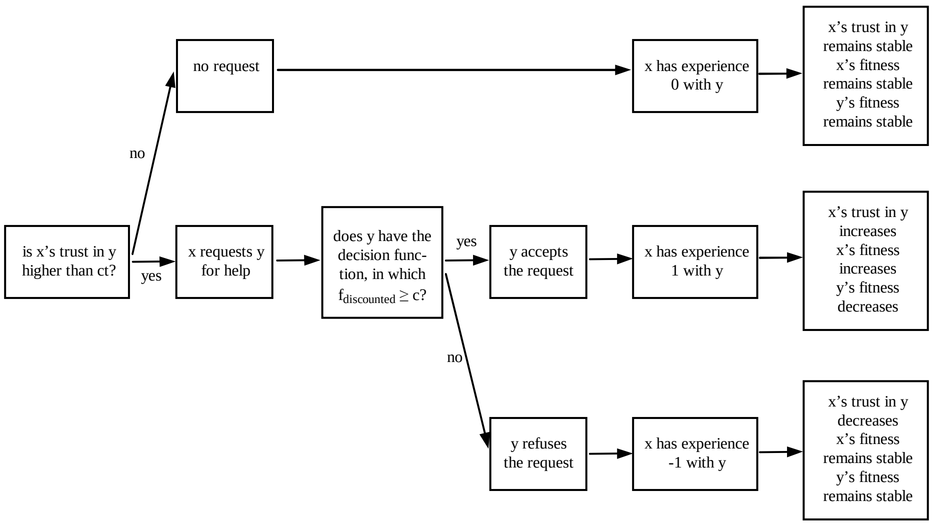

Fig. 1. Scenario sketch of one round of interaction.

First, a conceptual model has been created of this domain in LEADSTO, modelling a society of six agents. In this model, three agents used the decision function introduced in Section 2, whereas the other three agents were not equipped with this function. These three agents simply always selected the current reward, thus they never cooperated. Next, a large-scale simulation of the domain has been implemented in NetLogo, involving 200 agents. In this model, again half of the agents used the inter-temporal decision function, whilst the other half did not. Other parameter settings (both in the LEADSTO as in the NetLogo simulation) were as follows. In the decision function, c=10, f=14, β=0.20, t=10, n=10. In the trust function, δ=0.9. The request threshold ct (see Figure 1) is -0.3, and the daily decrease in fitness d is 3.

The basic building blocks of a LEADSTO simulation model are so-called Local Properties (LP's, in contrast to the Global Properties that are shown in Section 5). The format of these local properties is defined as follows. Let α and β be state properties of the form 'conjunction of atoms or negations of atoms', and e, f, g, h non-negative real numbers. Then, the notation $\alpha \twoheadrightarrow_{e, f, g, h} \beta$, means:

If state property α holds for a certain time interval with duration g, then after some delay (between e and f) state property β will hold for a certain time interval of length h.

Some local properties that were used for our conceptual model are shown below:

LP1 Trust Adaptation

Trust is adapted on the basis of experiences. Here, ‘delta’ is a constant, e.g. 0.9.

∀x,y:agent ∀tr:real ∀e:real

has_trust_in(x, y, tr) ∧ has_experience_with(x, y, e) $\twoheadrightarrow_{0,0,1,1}$
has_trust_in(x, y, delta × tr + (1 - delta) × e)

LP3a Decision Function

If you have the ability to evaluate the future reward, and the discounted value of the future reward is evaluated higher than the current cost, then accept a request for cooperation.

∀x,y:agent ∀tr:real ∀a:real

has_trust_in(x, y, tr) ∧ request_from_to(y, x) ∧ has_decision_function(x) ∧
has_alpha(x, a) ∧ current <= future × $2^{(-(1-a)\times(1-(tr+1)/2))}$ $\twoheadrightarrow_{0,0,1,1}$
accept_cooperation_with(x, y)

LP3c Primitive Decision Function

If you don’t have the ability to evaluate the future reward, then refuse a request for cooperation.

∀x,y:agent ∀tr:real

has_trust_in(x, y, tr) ∧ request_from_to(y, x) ∧ not has_decision_function(x) $\twoheadrightarrow_{0,0,1,1}$
refuse_cooperation_with(x, y)

The complete set of local properties in LEADSTO can be found at http://www.cs.vu.nl/~tbosse/egal/egal-AppendixA.pdf. When this declarative specification turned out to show coherent behaviour, it was used as a blueprint to implement the model in NetLogo.

4 Simulation Results

This section discusses the results of the LEADSTO simulation and the Netlogo simulations for 6, 25 and 200 agents. The results of the LEADSTO simulation and Netlogo-25 and -200 simulations are available at http://www.cs.vu.nl/~tbosse/egal/egal-AppendixB.pdf. In these cases (which are not shown here due to a lack of space), initially, all agents request each other for help. However, only the agents with the inter-temporal decision function are willing to cooperate; they accept all requests, thereby showing some kind of altruistic behaviour. The other agents show egoistic behaviour: they refuse all requests. As a result, the trust in the cooperating agents increases, whilst the trust in the non-cooperating agents decreases. This development continues for a while, until the trust in the non-cooperating agents is so low that even the agents with the inter-temporal decision function (the cooperating agents) are not willing to help them anymore. However, they still continue helping the other cooperating agents. Thus, a group emerges of agents that are helping each other, whilst the other agents get isolated: they do not interact with any agent anymore. As a consequence, the fitness of the agents with the decision function (which was first rather low, since these agents were initially exploited by the other agents) recovers, and the fitness of the agents without decision function gets lower and lower. These results confirm the hypothesis that agents that have the ability for inter-temporal decision making will show altruistic behaviour, which leads to a bigger social network, and eventually to a higher fitness.

Figure 2 shows the results for the Netlogo simulation with 6 agents. In Figure 2 the following measurements are shown: fitness – the average fitness of the agents for the short-termers[1] and long-termers, respectively; trust – the average trust of each type of agent (short-term or long-term) in the other type of agent; requests – total number of requests that agents have done

[1] This Section uses the term ‘short-termers’ to indicate those agents that do not operate based on a decision function, but simply never accept cooperation requests. The ‘long-termers’ use the decision and trust-update function presented earlier to decide on accepting cooperation requests.

(cumulative over time); cooperations – total number of cooperations (cumulative over time). A cooperation ($type_1$ -> $type_2$, e.g., long -> long) is an accept from agent *y* on a request from agent *x*, where $type_1$ is the type of agent *x* and $type_2$ is the type of agent *y*.

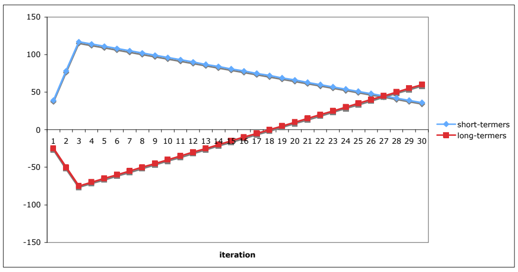

Development of fitness over time

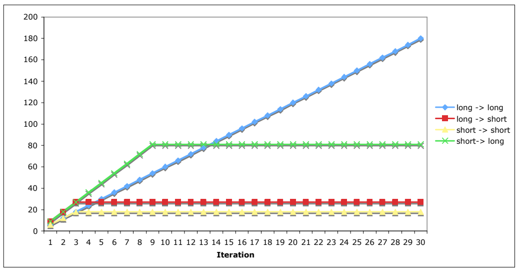

Number of requests over time

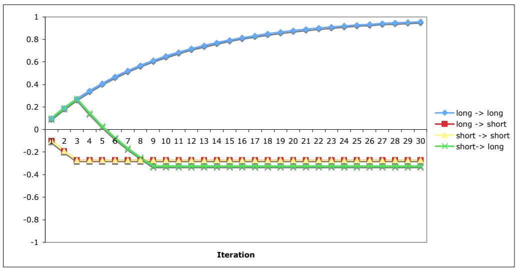

Development of trust over time

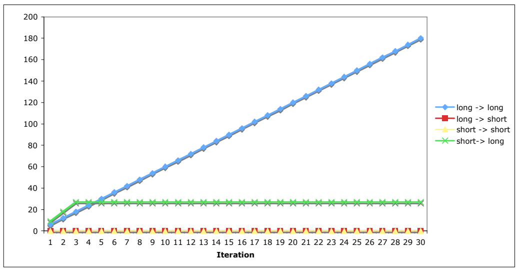

Number of cooperations over time

Fig. 2. NetLogo results for simulation with 6 agents.

We can do a number of observations looking at the 6-agents and 200-agents Netlogo traces. Firstly, we see that the results of the simulation with 6 agents are consistent with the results of the LEADSTO simulation. Most importantly, we observe 1) the same trends and 2) that there is a turning point at the third iteration, after which the trust of the long-termers in the short-termers is sufficiently low that they do not cooperate anymore. Secondly, the results are scaleable (from 6 to 200 agents) with some noteable differences for the various measures. For *fitness*: we see that the fitness of the short-termers increases rapidly at first, and then decreases – but slower for 200 agents than for 6 agents. This is a result of the fact that short-termers receive reward from many more agents in the first three iterations and (because the metabolism is equal for both traces) can live off this much longer. The point at which the fitness curves cross each other is also earlier for 200 agents than with 6 agents. This is because the long-termers also benefit from the fact that after iteration 3, they receive cooperation from many more other long-termers, leading to faster fitness increase. For *trust*: we do not observe differences between the two traces. For *number of requests*: most noteably, we see that the (short -> long) curve crosses the (long -> long) curve much later for the 6 agents trace (at iteration 14 instead of 9). Initially, the (short -> long) curve grows quicker than the (long -> long) curve because the growth-factor of the (short -> long) is the number of long-termers and the growth-factor of the (long -> long) is the number of long-termers – 1 (since a long-termer does not do a request to itself). Eventually they cross, because the (short -> long) curve does not grow further after iteration 9 because of the fixed cooperation threshold. For *number of coopera-*

tions: in both traces we see that the long-termers never cooperate with the short-termers, and short-termers never cooperate with each other. The long-termers always cooperate with each other. The short-termers cooperate with the long-termers up till iteration 3.

To better illustrate the emergence of groups in the society, the cooperation between the agents can be visualised, using the organisation visualisation tool by [Hoogendoorn at el., 2005]. A screenshot of this tool is depicted in Figure 3, for the results of the NetLogo simulation of 25 agents. Here, the nodes denote the agents, the edges denote cooperation. The size of the nodes and the thickness of the edges indicate the number of cooperations the agents were involved in during the whole simulation. Figure 3 clearly illustrates that the long-termers have established a network, whilst the short-termers have become isolated.

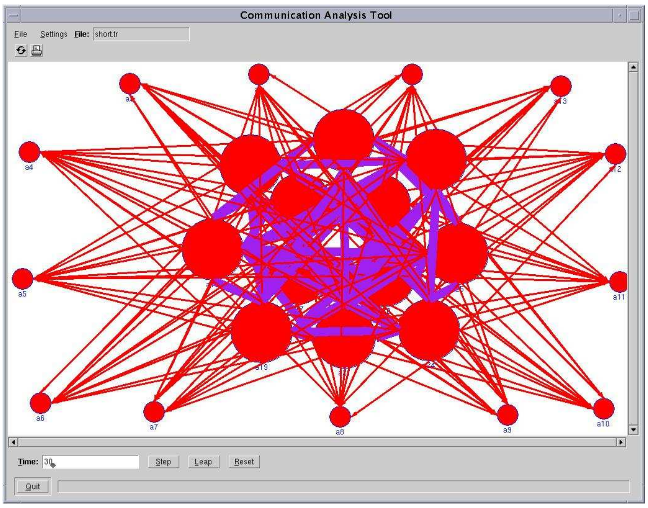

Fig. 3. Emergence of groups in the society.

5 Analysis

This section addresses further formal analysis of the simulation model and its results. In particular, a number of global dynamic properties have been formalised in TTL, and it was verified whether the local temporal properties defining the simulation model entail these global properties. This type of analysis can be performed in two ways: (1) using a checker tool, it can be verified automatically whether the global properties hold for the generated simulation traces, and (2) interlevel relations can be established between the local properties (see Section 3) and the global properties, which can be verified by mathematical proof. These two types of analysis are addressed by Section 5.1 and 5.2, respectively.

5.1 Checking Global Properties

A number of global properties have been identified that are relevant for the domain of trust-based inter-temporal decision making. These properties have been formalised in the TTL language. Three of them are shown below:

FM Fitness Monotonicity

If x has the cognitive system for decision making, then there exists a time t such that for all t1 and t2 after t, with t1 < t2, the fitness of x at t2 is higher than the fitness of x at t1.

∃t ∀x:AGENT ∀t1,t2≥t ∀f1,f2:REAL
[[∀t state(γ, t) |= has_decision_function(x)] & t1<t2 & state(γ, t1) |= has_fitness(x, f1)
⇒ ∃f2≥f1 state(γ, t2) |= has_fitness(x, f2)]

Here, for example, state(γ, t1) |= has_fitness(x, f1) denotes that in trace γ in the state at time point t the agent x has fitness f1.

DMAF Decision Making Agents get Fitter

Eventually, all agents with the cognitive system for decision making, will be more healthy than the agents without this system.

∀t ∀x,y:AGENT ∀f1,f2:REAL
[state(γ, t) |= has_decision_function(x) ∧ not has_decision_function(y) &
state(γ, last_time) |= has_fitness(x, f1) ∧ has_fitness(x, f2)
⇒ f1>f2]

NDMA Network of Decision Making Agents

All agents with the cognitive system for decision making will always cooperate with each other.

∀t ∀x,y:AGENT
[state(γ, t) |= has_decision_function(x) ∧ has_decision_function(y) ∧ requests_from_to(x, y)
⇒ state(γ, t+1) |= accept_cooperation_with(y, x)]

A specific software environment has been built, which takes as input a set of traces and a formalised property in TTL, and verifies whether the property holds for the traces [Bosse et al., 2006]. Using these kinds of checks, the above properties have been checked against the traces mentioned in Section 4 involving 6 and 25 agents. They all turned out to hold, which validates the above statements, such as “decision making agents get fitter”, for the simulation traces.

5.2 Interlevel Relations

This section aims at getting insight in why on the basis of the mechanisms as modelled in the local properties (see Section 3) the global properties are obtained. This is done by an analysis, for a given global property, based on a hierarchical AND-tree of dynamic properties at different levels, in which the branching specifies interlevel relations. This tree represents an argumentation why the global property holds, given as premises that the local properties hold. The tree considered here focuses on the highest level property FM shown in the previous paragraph, and is shown in Figure 4.

Roughly spoken the argumentation runs as follows. There are two ways to affect fitness. One way is (1) to increase it by earning profit by investing in cooperations with other cooperative agents. Another way to affect fitness is (2) to decrease it by investing in cooperations with non-cooperative agents so that the investment gives no return. Initially, it can be unclear which contribution to fitness dominates. However, by maintaining trust in other agents based on the experiences in cooperation, the agent learns to discriminate the agents in the categories (1) and (2), and thus decides not to invest in the noncooperative agents anymore. Therefore, after some time point contribution (1) to the fitness dominates, and thus fitness becomes monotonically increasing. This rough outline of the argumentation was detailed as follows, making use of properties at one level lower:

CTF Fitness Change by Cooperation Results

Cooperation profit contributes to fitness increase, and cooperation loss contributes to fitness decrease.

∀t1,t2 ∀x:AGENT ∀f1,d1,d2
[state(γ, t1) |= has_fitness(x, f1) &
cooperation_profit_over(γ, x, d1, t1, t2) & cooperation_loss_over(γ, x, d2, t1, t2)]
⇒ state(γ, t2) |= has_fitness(x, f1+d1-d2-d*(t2-t1))]

CP Cooperation Profit
If x has the cognitive system for decision making, then there exists a time t such that for all t1 and t2 after t, with t1 < t2 and length L1 ≤ t2-t1 ≤ L2, between t1 and t2 the amount of revenues contributed by other cooperative agents y to x in cooperations is at least M higher than the amount of expenses invested by x in cooperation with these agents y.

$\forall x$ [has_decision_function(x) $\Rightarrow$
$\exists t\ \forall t1,t2 \geq t$ [L1 ≤ t2-t1 ≤ L2 $\Rightarrow \exists d$ cooperation_profit_over(γ, x, d, t1, t2) & d ≥ M]

Here L1, L2 and M are constants that can be given specific values.

CL Cooperation Loss
If x has the cognitive system for decision making, then there exists a time t such that for all t1 and t2 after t, with t1 < t2 and length L1 ≤ t2-t1 ≤ L2, between t1 and t2 the amount of loss due to other noncooperative agents z to x in cooperations becomes lower than M.

$\forall x$ [has_decision_function(x) $\Rightarrow$
$\exists t\ \forall t1,t2 \geq t\ \forall d$ [L1 ≤ t2-t1 ≤ L2 & cooperation_loss_over(γ, x, d, t1, t2) $\Rightarrow$ d≤M]

The property CP relates to the following two lower level properties (of which the formalisation has been omitted):

CTMT Investment in Cooperation Leads To More Trust
For an agent y with the cognitive system for decision making, if x invests in cooperation with y, then trust of y in x will increase

MTTMC More Trust Leads to More Cooperation
For an agent x with the cognitive system for decision making, more trust of x in y in leads to more investment of x in cooperation with y.

The property CL relates to the following two properties (of which the formalisation has been omitted):

NCTLT Non-Cooperation with a Cooperative agent Leads To Less Trust
For an agent y with the cognitive system for decision making, if x does not invest in cooperation with y, then trust of y in x will decrease.

LTTLC Less Trust Leads to Less Cooperation
For an agent x with the cognitive system for decision making, less trust of x in y in leads to less investment of x in cooperation with y.

Based on the properties defined above, the logical relationships are as follows. Suppose x and y both have the cognitive system for decision making, and an initial amount of trust of x in y and of y in x is available, and y is cooperative. Moreover, assume that the properties CTMT and MTTMC hold. Then:

1. The initial trust of x in y leads to some investment by x in cooperation with y by property MTTMC.
2. This investment of x in cooperation with y leads to more trust of y in x, by CTMT.
3. The increased trust of y in x leads to more investment of y in cooperation with x, by MTTMC.
4. This positive feedback process continues for some time, and thus will increase the amount of revenues of cooperation of agent x (and of y) after some time point to a level that profit is obtained from the cooperation, which shows that property CP is implied by CTMT and MTTMC.

A similar argument can be made for x with respect to a noncooperative agent z. Assume that the properties NCTLT and LTTLC hold. Then:

5. The initial trust of x in z leads to some investment by x in cooperation with y by property MTTMC.
6. As z is noncooperative, this does not lead to more investment of z in cooperation with x.
7. Because of noncooperation of y with x, x will decrease its trust in z by NCTLT.
8. The decreased trust of x in z leads to less investment of x in cooperation with z, by LTTLC.
9. This negative feedback process continues for some time, and thus will decrease the amount of loss of cooperation of agent x with z after some time point, showing that property CL is implied by NCTLT and LTTLC.
10. If after some time point over any interval there is profit from cooperation, and the loss becomes lower, and profits and losses affect fitness in positive, respectively negative manners, then after a point in time the fitness at t2>t1 will be higher than the fitness at t1. This expresses that property FM is implied by CP, CL, CPFI, and CLFD.

By this argumentation the logical interlevel relationships as shown in Figure 4 are obtained. The semantics of this tree is as follows: if a certain property is connected to a number of lower level properties, then the conjunction of the lower level properties (logically) entails the higher level property. As the picture shows, eventually the intermediate properties can be related to the local properties as shown in Section 3.

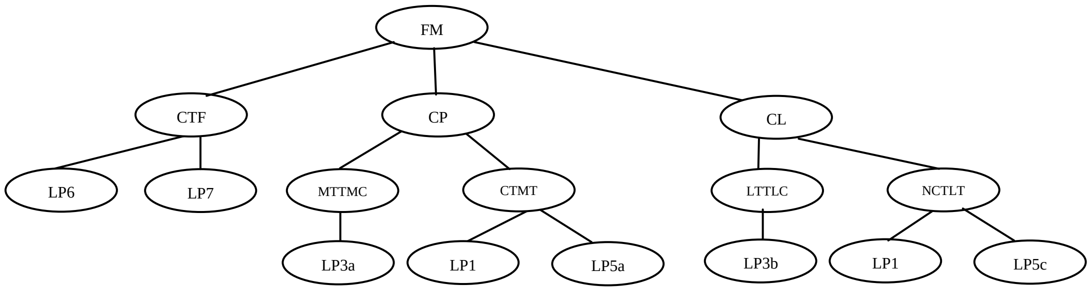

Fig. 4. AND-tree of interlevel relations for Global Property FM.

6 Conclusion

The work reported in this paper contributes to the analysis of a paradoxical question from an evolutionary perspective: how can altruistic behaviour be in an agent's own interest? The question is addressed by involving a temporal dimension in which altruistic behaviour at one point in time can be in the agent's interest at a future point in time, depending on the environment. The claim is that to be able to make reasonable decisions, an agent needs a cognitive system for intertemporal decision making, in relation to a model of the environment to predict when indeed it may expect to provide revenues for the agent's investments by its altruist behavior.

To address this, a society of agents has been modelled, simulated and analyzed. Some of the agents have a certain cognitive system for decision making that enables them to choose for altruistic behaviour. Part of this cognitive system is a model of the environment to predict whether future revenues may be expected in return, which is based on a dynamic model for trust in the other agents based on experiences with them over time, adopted from [Jonker and Treur, 1999, 2005]. This environment model is combined with a model for intertemporal decision making taken from the literature; e.g., [Loewenstein and Elster, 1992].

It turned out that the agents with this cognitive system enabling them to anticipate on the future show more altruistic behaviour. As a result, these agents get a bigger social network, and become in the end more healthy than the agents without such a cognitive system. This is in accordance with the theory of how altruism emerged in Nature as a result of more elaborated capabilities of mankind for inter-temporal decision making; e.g., [Dennett, 2003]. Among the agents with the ability to anticipate on the future, different variants can be identified. In future work, it will be investigated how a society with such different variants develops.

We are aware of the fact that there exists a substantial body of work on the evolution of co-operation, with most notably the seminal works by Trivers [1971] and Axelrod [1981, 1984]. Trivers elaborates on the mathematics of *reciprocal altruism* – the form of altruism in which an organism provides a benefit to another in the expectation of future reciprocation – and includes human reciprocal altruism as a case study to illustrate the model. Axelrod's work investigates the question under what conditions cooperation will emerge in a world of egoists without central authority. He demonstrates that "altruistic" strategies such as Tit-for-Tat perform well in iterated prisoner's dilemma (IPD) tournaments. Although these works address the same fundamental question on reciprocal altruism as in this paper (how can altruistic behaviour be in an agent's own interest?), our perspectives are significantly different. First, the context of our simulations differs from the IPD in a number of ways. For example, it is more realistic (the domain of providing services is less artificial than that of prisoners) and it is not symmetric (it is possible that agent A does a request to agent B, while B does not do a request to A). Second, we addressed a question that is narrower than the one posed by Axelrod. Whereas Axelrod investigated which types of strategies performed best in a heterogeneous population with many different strategies, we took the theory of Dennett [2003] as inspiration in order to explicitly compare two types of agents: those with a decision function, and those without. For the former type, we consider the decision function of the involved individuals to be based on the decision-theoretic notion of utility discounting. Ainslie [2001] reports on the existence of such a function (exponential or hyperbolic) based on collected field data. Dennett [2003] then uses such function as the basis of an evolutionary explanation of free will. To our best knowledge, to include such a discount utility function in the context of reciprocal altruism has not been addressed explicitly before. Also, we explicitly include an adaptive trust parameter to the decision function, which can be considered a novel contribution to the research on this topic. Nevertheless, we predict that strategies such as Tit-for-Tat can be modelled as special cases of our trust-based intertemporal decision function (e.g., using specific parameter settings such as $\delta=0$). In the future, we will perform experiments with heterogeneous populations, including Tit-for-Tat agents.

References

1. Ainslie, G. (2001). *Breakdown of Will*. Cambridge University Press, 2001.
2. Axelrod, R. (1981). The Evolution of Cooperation. *Science.* 211(4489):1390-6, 1981.
3. Axelrod, R. (1984). *The Evolution of Cooperation*. Basic Books, New York, 1984.
4. Bosse, T., Jonker, C.M., Meij, L. van der, and Treur, J. (2005). LEADSTO: a Language and Environment for Analysis of Dynamics by SimulaTiOn. In: Eymann, T., et al. (eds.), *Proc. of the Third German Conf. on Multi-Agent System Technologies, MATES'05*. Springer LNAI, vol. 3550, 2005, pp. 165-178.
5. Bosse, T., Jonker, C.M., Meij, L. van der, Sharpanskykh, A., and Treur, J. (2006). *Specification and Verification of Dynamics in Cognitive Agent Models*. In: Proceedings of the Sixth International Conference on Intelligent Agent Technology, IAT'06. IEEE Computer Society Press, 2006, pp. 247-254.
6. Dennett, D.C. (2003). *Freedom Evolves*, New York: Viking Penguin, 2003.
7. Godfrey-Smith, P. (1996). *Complexity and the Function of Mind in Nature*. Cambridge University Press, 1996.
8. Hoogendoorn, M., Treur, J., and Yolum, P. (2005). A Labelled Graph Approach to Support Analysis of Organizational Performance. In: *Proceedings of the International Workshop on Agent-based Technologies and Applications for Enterprise Interoperability, ATOP'05*, 2005, pp. 49-60.
9. Jonker, C.M., and Treur, J. (1999). 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.
10. Jonker, C.M., and Treur, J. (2005). *A Framework for Specification and Analysis of Models for the Dynamics of Trust based on Experiences.* Technical Report, Vrije Universiteit Amsterdam, 2005.
11. Loewenstein, G.F., and Elster, J. (1992). *Choice over time*. Russel Sage Foundation, New York, 1992.
12. Sober, E., and Wilson, D.S. (1998). *Unto Others: The Evolution and Psychology of Unselfish Behaviour*. Harvard University Press, Cambridge, MA, 1998.
13. Trivers, R.L. (1971). The Evolution of Reciprocal Altruism. *Quarterly Review of Biology.* 46:35-57, 1971.
14. Wilensky, U. (1999). NetLogo. http://ccl.northwestern.edu/netlogo. Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston, IL, 1999.