Evolutionary Cooperation in a Multi-Agent Society

Marjolein de Vries, Pieter Spronck
Tilburg center for Cognition and Communication
Tilburg University
Tilburg, The Netherlands
Email: mdv@marjoleindevries.com, p.spronck@uvt.nl

Abstract. According to the evolutionary ethics theory, humans have developed a moral sense by means of natural selection. A "moral sense" supposes some form of cooperation, but it is not immediately clear how "natural selection" could lead to cooperation. In this research, we present a multi-agent society model which includes a cooperative act in order to investigate whether evolution theory can explain to what extent cooperative behavior supports the emergence of a society. Our research shows that cooperation supports the emergence of a society in particular when taking into account individual differences.

Keywords: evolutionary ethics ▪ evolutionary algorithms ▪ multi-agent societies ▪ simulation ▪ cooperation ▪ individual differences

1 Introduction

For many years, philosophers, biologists and psychologists have been studying how the human race got instilled with a moral sense [16]. A moral sense can be defined as "a motive feeling which fueled intentions to perform altruistic acts" [16]. It drives humans to exhibit moral behavior such as performing acts of cooperation. According to the evolutionary ethics theory, humans have developed a moral sense by means of evolution and natural selection [16]. Natural selection implies competition and thus can be argued not to lead to cooperation [13], which has made the existence of cooperation in human behavior a challenge to explain [1]. Moreover, according to the Social Exchange theory, humans require some cost-benefit mechanism in order to successfully engage in a cooperative act [3].

While the ideas behind evolutionary ethics are appealing, it is difficult to determine their value experimentally. Performing experiments with organisms, such as humans, on evolutionary ethics is hard, as evolution takes place over hundreds of thousands of years and experiments on ethics with humans are unethical themselves. However, simulations may provide a means to perform such experiments. In particular, the artificial life approach using a multi-agent society may be suitable to perform such experiments, as agent-based models are an effective tool for modeling the emergence of social phenomena [21].

Previous research on evolutionary ethics by Spronck and Berendsen [18], based on the work of Epstein and Axtell [6] and Mascaro [11], shows that artificial multi-agent societies may provide new insights in evolutionary ethics theory. Nonetheless, no specific act of cooperation was present in their computational model, even though social and ethical behavior implies some form of cooperation.

In our research we present a multi-agent society model based on the model of Spronck and Berendsen [18], in which agents are able to perform a specific cooperative act. With our model, we aim to investigate to what extent evolution theory can explain the emergence of cooperation in the formation of a society. Because research has shown that individual differences strongly affect evolutionary outcomes [2], we will differentiate between regular cooperation and cooperation in which individual differences are taken into account.

The outline of this paper is as follows: Section 2 provides the theoretical background needed for the set-up of the computational model, which is described in Section 3. Section 4 presents our results. In Section 5, our results are discussed. Section 6 formulates the conclusion of this paper.

2 Theoretical Framework

Section 2.1 describes the research area of Evolutionary Ethics. Section 2.2 outlines the concept of cooperation.

2.1 Evolutionary Ethics

In "The Origin of Species," Darwin [4] describes a new way of looking at our contemporary behavior from an evolutionary approach. According to Darwin, one of the most important features of species is self-reproduction, a system of positive and negative feedback which we refer to as natural selection [19]. Individuals having any advantage over others, and thus having a better fitness, have the best chance of surviving and procreating. Darwin referred to this process as "survival of the fittest" [14]. From a sociobiological approach, natural selection can explain patterns of social behavior of humans and other species [19]. According to the evolutionary ethics theory, natural selection has caused humans to have a moral sense, i.e. "a motive feeling which fueled intentions to perform altruistic acts" [16]. This moral sense drives humans to behave morally, such as performing acts of cooperation.

Natural selection implies competition, as every organism is designed to enhance its own fitness at the expense of the fitness of others [13]. Contrariwise, humans are amongst the most communal and socially dependent organisms alive on our planet [16]. The aspect which distinguishes us from most other species is that we cooperate even with individuals which are not genetically closely-related in order to achieve common goals [9]. Moreover, evolutionary ethicists state that if members of the same species execute a process of cooperation instead of struggling with each other, more personal benefit can be achieved [17]. Evolutionary biologists have been fascinated by this phenomenon for several decades [13].

A research method used for evolutionary experiments is the artificial life method. Artificial Life as a research field attempts to resemble biological phenomena by means of computer simulations in order to create theoretical grounding for biological theories

[15]. Artificial societies are inhabited by autonomous agents, which are capable of exhibiting emergent behavior. Emergent behavior is behavior which is not programmed into the agents by the designer, but instead is created using a process of self-adaptation. This resembles the assumed automatic emergence of moral behavior within humans [7]. De Weerd, Verbrugge and Verheij, for example, show that agent-based models are suitable for investigating emerging social phenomena [21].

In order to research ethical behavior using artificial agent societies, a specific definition of which actions are called "ethical" is needed. According to Jaffe [8], a distinction between the benefit to society 'S' and the compound benefit to the individual 'I' needs to be made. Jaffe [8] classifies actions as either:

"True altruism", if $S \geq 0$ and $I < 0$
"Social investment", if $S \geq 0$ and $I \geq 0$
"Destructive egoism", if $S < 0$ and $I \geq 0$
"Destructive behavior", if $S < 0$ and $I < 0$

The common moral sense of a society can be defined as the average of executed actions in a particular period in a society [18]. A 'moral' society would predominantly perform actions of true altruism and social investment.

Earlier research done on ethics in agent societies have been performed by Epstein and Axtell [6], who experimented with an artificial society built by defining simple local rules. Their research shows that these simple rules can lead to complex emergent behavior in an agent society. Mascaro [11] proposed using genetic algorithms to let the behavioral rules evolve over time, resembling the evolution process described by Darwin. Mascaro also presented basic environments in which agents are able to perform a finite number of actions.

Spronck and Berendsen [18] based their model for researching the emergence of moral behavior on the work of Epstein and Axtell [6] and Mascaro [11]. In their society model, agents are able to perform four actions: wander, forage, steal and share. They found some specific configurations in which agents would steal very little, and would even share for a small percentage of their actions. However, a specific act of cooperation was missing from this model, even though social and moral behavior implies cooperative acts. Therefore, in our research, a cooperative action will be added to the model of Spronck and Berendsen [18].

2.2 Cooperation

To design a cooperative action, we looked at earlier research on cooperation from an evolutionary perspective. Such research has been done on the (iterated) Prisoner's Dilemma [12]. However, it is hard to draw conclusions from the Prisoner's Dilemma, as it is only a model of a simple problem with a finite number of options [18]. Moreover, it only models the behavior of two agents, and not the behavior of a society. Therefore, a different proposal for a design of a cooperative act should be made.

In the interest of designing an act of cooperation for agents in a society, a representation of the concept of cooperation is needed. A cooperative social act includes some individual costs, but also supplies an amount of collective group benefits which is divided equally over all the cooperating members [7]. A cooperative act enables not only an agent's own objectives, but also the objectives of other agents [5]. This kind of cooperation is closely related to Social Exchange, which is defined as "cooperation between two or more individuals for mutual benefit" [3]. During an exchange, an individual is generally obligated to pay a cost in order to be able to receive a benefit. This requires humans to have specialized cognitive mechanisms for deciding whether the benefits of a potential exchange outweigh its costs.

To be able to fully exploit the benefits of cooperation, it might help if there exist individual differences in a society. Research has shown that individual differences strongly affect evolutionary outcomes such as survival and mating success [2]. Diversity creates value and benefit for team outcomes [10]. Introducing individual differences will have an effect on the outcome of a group process such as cooperation.

3 Society model

In this section, we describe the design of our society model. The society model is based on the model of Spronck and Berendsen [18]. In Section 3.1, the general set-up of the environment is explained. Section 3.2 describes the design of the agents and their attributes. Section 3.3 construes the set of rules each agent possesses and Section 3.4 illustrates the mating process. Section 3.5 states by what means we have measured the effect of a cooperative act in a multi-agent society. Extensive details of the model are given by de Vries [20].

3.1 The Environment

The society is situated in an environment which consists of a two-dimensional 30×30 grid with correspondingly 900 cells. The environment is bounded by the edges, which means agents close to the boundary have a limited environment. At the start of the experiment, the environment contains 200 agents and 9 cells contain nutrition. During a round, each of the 900 cells can contain nothing, an agent, nutrition or both an agent and nutrition.

A 'round' is the unit of time which is used in the society model. Each run in the model lasts 20,000 rounds. In every round, an agent executes one action. At the beginning of each round, new nutrition is be added to the environment. Moreover, every agent is allowed to start the mating process (if the agent meets some specific requirements), which creates a new agent. The mating process counts as one action, which means that after mating an agent is not able to execute another action anymore. The order in which agents are allowed to perform actions is randomly decided. At the end of the round, the health and age of each agent changes, which could result in agents dying when their health or age is respectively too low or too high.

In order to create a dependency and a possible urge to cooperate between agents two nutrition types are available in the environment: food and water. Half of the cells in the environment are assigned to contain water and the other half to contain food. On a turn, agents are only able to forage one type of nutrition. Assigning the cells to one of the nutrition types is done randomly.

3.2 Agent attributes

Each agent has an **age**. This is the number of rounds the agent has been alive. The age of an agent at birth is 0, and an agent dies when it reaches the age of 100. Moreover, agents have a **health** attribute represented as an integer. An agent dies when its health reaches zero or less. An agent created at the beginning of a run has a health of 100. An agent created by the mating process inherits 25% of both its parents' health. The health attribute is used as fitness measure. Agents lose 5 units of health at the end of each round. Furthermore, agents have a field of **vision**, which determine how many cells an agent can take into account when making decisions. All agents can look in four main directories (north, east, south and west) and have a vision of three cells.

Furthermore, agents have a **food basket** and a **water basket**, of which its value represents the amount of units of the corresponding nutrition type. When agents receive nutrition, the acquired units are added up to the corresponding baskets. At the end of a round, agents are able to digest nutrition from their baskets and increase their health. An agent can only consume water and food in equal amounts, thereby implying a need for both nutrition types. Agents created at the beginning of a run have a value of zero for each basket, and agents created by the mating process inherit 25% of each of the baskets of both parents.

3.3 Agent Rules

Agents have a rule set of seven rules. Rules consist of zero, one or two tests with a specified action, thus having the following structure: $(Test \wedge Test\ ) \rightarrow Action$. Agents created at the beginning of a run get a randomly made rule set. Agents created by the mating process inherit rules from their parents (Section 3.4).

Each test begins with specifying an agent on which the test is performed. A choice is made among (1) the agent itself, (2) the closest agent, (3) the furthest agent, (4) the healthiest agent, (5) the weakest agent, (6) the agent with the most food, (7) the agent with the least food, (8) the agent with the most water and (9) the agent with the least water. Then, an attribute is chosen from (1) health, (2) age, (3) food availability, (4) water availability, (5) food basket and (6) water basket. Moreover, a mathematical operator is selected from smaller than, larger than, equals and wildcard. If a test contains a wildcard, it always yields "True". Next, a whole number is selected.

Finally, an action is executed by an agent if both tests are applicable. The following actions are available, where cooperation is not available in the baseline experiment.

Wander. The agent moves to a random cell within its vision. In the case that the new cell contains nutrition, the agent puts it in the corresponding basket with a maximum of 30 units.

Forage Food. The goal for the agent with this action is to search for food within its vision. The agent then moves to the cell with food and puts it in the food basket with a maximum of 30 units.

Forage Water. The goal for the agent with this action is the same as Forage Food, only this action is aimed at water instead of food.

Steal Food. The agent steals an amount of 15 units from the food basket of the specified agent. The amount of units is based on the model of Spronck and Berendsen [18]. If the specified agent does not have 15 or more units in its food basket, the maximum quantity possible is stolen. The agent generally aims the steal action at the agent specified in the second test. If that agent is the agent itself, the aim of the action is the first specified agent or otherwise a random agent within the active agent's vision.

Steal Water. The goal for the agent with this action is the same as Steal Food, only this action is aimed at water instead of food. The target agent is specified in the same way as for the Steal Food action.

Share Food. With this action, the agent gives away 25% of its basket with food to the specified agent. The target agent is specified in the same way as for the Steal Food action.

Share Water. With this action, the agent gives away 25% of its basket with water to the specified agent. The target agent is specified in the same way as for the Steal Food action.

Cooperate. Agents have to pay a cost of x units of their fullest basket. Agents are only allowed to engage in the cooperative act if they have enough units in their basket to pay the cost. These costs are removed from the environment, which means they are lost to the population of agents. After paying the costs, agents are randomly assigned to either forage food or forage water. Then, agents have to hand in the units of food or water they just have foraged. At the end of the round, the total nutrition collected is divided equally amongst all the participating agents. The cost that an agent pays to perform the cooperative action is thus the entry fee of x units of their fullest basket. The benefit of the cooperative act is that the agent is guaranteed to receive both food and water, instead of at most one type of nutrition. This act satisfies the description of cooperation made by Gintis [7] and by Doran, Franklin, Jennings and Norman [5].

According to Jaffe's dimensions [8], we can classify wandering, foraging and cooperating as *social investment*, stealing as *destructive egoism* and sharing as *true altruism.*

In a second experiment, individual differences are introduced into the society. To reflect the concept of individual differences, each agent is instilled with a preference for foraging water or food. When the agent has a preference for foraging nutrition y, the agent receives $\frac{3}{2}$ times its foraged units of the type y and $\frac{1}{2}$ times its foraged units of the other nutrition type. Agents at the start of a run get a randomly assigned preference. Agents created by the mating process have a 50% chance to receive each of their parents' preferences, with a 2% chance of being mutated. While engaging in the cooperative act, agents are now assigned to forage the nutrition type which has their preference.

An example rule is *(closest agent age > 28) and (agent itself health < 40) then forage*. A child inherits its rule sets by means of a uniform crossover operators which operates on the different tests in the rule, where each rule has a 2% chance of having a test being mutated.

Rules are ordered on complexity, and agents start with the most complex rule when executing their rules. Agents go through their rules according to their complexity until one of the rules apply, then the according actions is executed. The specified agent for the steal and share actions is generally the agent specified in the second test of the rule. If the specified agent is the agent itself, the specified agent of the first rule is selected and otherwise a random agent within an agent's vision is selected. If all rules do not apply, the agent executes the action *foraging* as a way of using an instinct.

3.4 Mating

An agent tries to mate with another agent before a round starts. If an agent already has mated with another agent, it is not able to mate again in the same round. Mating is only possible if an agent has an age of at least 18, a minimum health of 50 and at least 30 units of nutrition in each of its two baskets. The agent chooses randomly from suitable agents in his vision.

A child agent is created by the mating process and placed on a cell nearby its parents if possible. A child inherits 25% of each of its parents' health and 25% of each of its parents' baskets. A child inherits a combination of the rules of its parents by means of a uniform crossover operator. Moreover, each rule has a 2% chance of having a test part being mutated.

3.5 Measurements

Statistics are gathered during the runs in order to investigate the common moral sense of a society. Demographics of the society, statistics of the executed actions by the agents, and information about the executed rules are taken into account when answering the research questions.

We define the common moral sense as the averages of the moral evaluation of the actions in the rule set. This means that the common moral sense exists of certain percentages of true altruism, social investment, destructive egoism and destructive behavior. An ethical moral sense would consist of mostly true altruism and social investment. Because it is only interesting to look at agents who are strong enough to survive and to mate, we only take the statistics of the agents who survived at least to the mating age. Moreover, only the executed actions by the agents in the last 5,000 rounds (of 20,000) are used for determining the common moral sense.

4 Results

Section 4.1 describes the baseline behavior. Section 4.2 describes the experiment with a cooperative act, and the experiment in which individual differences are also introduced. The results of the forage, steal and share actions for both nutrition types are grouped together for the purpose of clarity (e.g., "forage" is the sum of "forage food" and "forage water").

4.1 Baseline Behavior

The simulation was executed 20 times, which results in 20 runs where each run consists of 20.000 rounds. Table 1 shows the percentage of executed actions, number of mature agents and health of mature agents averaged over 20 runs and the standard deviation.

Table 1. Percentage of executed actions and descriptive statistics in the baseline behavior,

	Average	St. dev.
Wander	0.13	0.07
Forage	53.55	11.82
Steal	46.15	11.76
Share	0.27	0.23
Nr of mature agents	151.73	12.69
Health of mature agents	95.80	4.99

Table 1 shows that the baseline behavior consists of mostly foraging and stealing. Therefore, we can conclude that the common moral sense of the baseline society is mainly egoistic, because the agents only use actions which are helping themselves. The results of the baseline behavior in this research are quite similar to the results of the baseline experiments of Spronck and Berendsen [18].

4.2 Cooperation experiments

The simulation of the society with the introduction of a cooperative act is executed with different costs for engaging in the cooperative act, where the cost range from 0 to 4 units. For each cost value, 20 runs were executed. Table 2 shows the percentage of executed actions, number of mature agents, health of mature agents and the compound benefit (the sum of food and water units) received from the cooperative act, all averaged over 20 runs. With a cost of 4 units, the cooperative was hardly executed anymore. Note that a cooperative act with a cost of 0 units cannot be called a true cooperative act, because it does not require any cost for engaging in the cooperation, while by definition a cooperative act must have a cost associated with it.

Table 2 shows that the cooperative act is executed less and the execution of the foraging act increases when the cost increases. Clearly the benefits of participating in the cooperative act do not always outweigh its costs. No statistically significant linear dependence of the mean of compound benefit on the cost was detected. Table 2 also shows that the average number of agents decreases with the introduction of the cooperative act when comparing to the baseline behavior. We assume that this is a consequence of the fact that the environment loses nutrition (which is paid as a cost for the cooperative act), thereby having less nutrition to nourish the agents in the society. Even when the cooperative act is no longer executed according to Table 2, there might still be agents which execute it, but which are not counted in Table 2 as they do not reach the age of maturity.

The common moral sense is quite similar to the baseline experiment. Nevertheless, agents do choose using the cooperative act over using the individual action 'foraging' when the costs are low, which results in somewhat more social behavior than in the baseline experiment.

Table 2. Percentage of executed actions and descriptive statistics in the cooperative experiment, standard deviation between brackets

	Cost 0	Cost 1	Cost 2	Cost 3	Cost 4
Wander	0.16	0.14	0.14	0.16	0.13
	(0.05)	(0.10)	(0.10)	(0.11)	(0.10)
Forage	17.20	30.33	40.82	45.50	58.08
	(27.22)	(27.44)	(24.32)	(14.12)	(12.95)
Steal	41.39	47.20	41.90	48.66	41.53
	(9.52)	(10.22)	(9.54)	(14.12)	(12.88)
Share	0.58	0.29	0.34	0.23	0.17
	(0.50)	(0.20)	(0.34)	(0.15)	(0.13)
Cooperate	40.85	22.19	16.95	5.59	0.19
	(20.85)	(22.90)	(20.20)	(10.32)	(0.18)
Nr of mature	143.09	134.00	129.27	133.33	131.72
agents	(4.62)	(3.89)	(4.03)	(4.40)	(4.37)
Health of	106.57	99.95	97.02	93.84	90.69
mature agents	(8.25)	(10.52)	(9.36)	(4.92)	(3.53)
Compound	7.63	8.53	8.42	8.91	n/a
benefit	(0.73)	(0.63)	(1.01)	(0.84)	(n/a)

Table 3 shows the results, averaged over 20 runs, when providing agents with a foraging preference as a reflection of the concept of individual differences. Table 3 shows that the cooperative act is now chosen considerably more than any of the other acts, in all tests. The benefits received from the cooperative act now outweigh the costs more convincingly, as also can been seen in the increase in benefit comparing with the previous cooperation experiment. With the use of linear regression, we can see that the cost significantly predicts the compound benefit, $\beta = 0.83$, $t(3) = 6.84$, $p < .01$. The cost also explains a significant proportion of variance in the health scores, $R^2 = 0.94$, $F(1, 3) = 46.74$, $p < .01$. Table 3 also shows that the introduction of the foraging preference attribute results in a larger mature population and an increase in the average health.

Furthermore, the common moral sense of the society with a cost of 0, 1 or 2 units contains of more *social investment* than the common moral sense in the baseline experiment. Consequently, we can conclude that the introduction of a cooperative act with a cost of 1 unit and the foraging preference attribute results in a superior society with the largest reduction of *destructive egoism*. In the light of the dimensions provided by Jaffe (2004), we can define its common moral sense as 61.05% *social investment*, 38.57% *destructive egoism* and 0.49% *true altruism*. Therefore, this experiment shows that with the introduction of a new action and attribute, the common moral sense can become significantly less anti-social.

Table 3. Percentage of executed actions and descriptive statistics in the individual differences experiment, standard deviation between brackets

	Cost 0	Cost 1	Cost 2	Cost 3	Cost 4
Wander	0.16	0.16	0.18	0.16	0.16
	(0.09)	(0.06)	(0.07)	(0.07)	(0.06)
Forage	4.98	3.07	2.05	6.58	12.33
	(10.17)	(7.95)	(3.46)	(14.52)	(16.40)
Steal	36.36	38.57	41.51	44.84	48.95
	(18.57)	(11.66)	(7.37)	(8.17)	(7.50)
Share	0.85	0.49	0.79	0.86	0.52
	(0.85)	(0.27)	(1.50)	(1.87)	(0.53)
Cooperate	58.00	57.98	55.69	47.73	38.22
	(25.44)	(15.15)	(8.93)	(12.76)	(16.99)
Nr of mature agents	170.10	158.67	148.25	134.18	126.06
	(23.90)	(11.59)	(9.33)	(10.19)	(9.00)
Health of mature agents	111.93	112.42	113.42	111.58	109.91
	(6.64)	(4.43)	(2.05)	(4.80)	(6.66)
Compound benefit	8.25	9.71	10.71	11.14	11.70
	(2.35)	(1.43)	(1.16)	(1.24)	(1.52)

5 General Discussion

Our research, which uses computer science to investigate questions in social sciences, supports the Social Exchange theory, as the cost-benefit behavior of the agents is evident in all our experiments. Because agents have to pay a cost for the cooperative act, agents need to receive a benefit which outweighs the cost, i.e., a benefit which is enough to maintain health. Therefore, the compound benefit should increase as the cost increases, as more compound benefit is needed to make up for the paid cost. We can see that the benefit is significantly predicted by the cost. Future research could implement actions with more complex cost-benefit calculations.

Moreover, the results of our research support the positive effect of individual differences on team outcomes [2], as the individual differences experiments yielded an increase in benefit and a morally superior society.

The research also introduces a new way of representing a cooperative act. Our society model provides an alternative to modelling cooperation by means of the (iterated) Prisoner's dilemma [12]. Moreover, many agent-based models on cooperation consist of agents having a shared goal on which they have to work together [5]. However, in our research, agents do not have a shared goal but rather individual goals on which they can achieve more effectively with the help of others.

In this paper, the results for the actions concerning both nutrition types are grouped together. When looking at the distinctive results for food and water [20], it can be seen that in some cases for example more food was foraged than water but more water was stolen than food. One explanation could be that agents need a minimal amount of food and water in order to reach their maximum age. For this purpose, it does not matter if an agent has more of one nutrition type than the other, as long as the lowest one is sufficient. Therefore, it is possible that the evolutionary process results in a society in which the two nutrition types are not foraged equally. This does not prevent a healthy society from emerging.

We succeeded in evolving a common moral sense that leans towards mainly social behavior, even with our simple agent model. Naturally, our society model is in no way comparable with human society. Many subtle influences of a human society are left out in the computational society, but could be included in future work. Future research could investigate the effect of other actions, attributes and settings.

Finally, we remark that when experimenting with agent models, one has to be careful not to "build in" the conclusions that one seeks. In our experiments, the evolved behavior is the result of only an agent's ability to survive and procreate, which is only dependent on age and health. Therefore, no preference for a specific act is built in our model.

6 Conclusion

The goal of this paper was to investigate to what extent evolution theory can explain the emergence of cooperation in the formation of a society. Our research shows that cooperation supports the emergence of a society, in particular when taking into account individual differences. With the introduction of individual differences, benefits from engaging in cooperation increased. The positive effect of individual differences on the society can be predicted by the Social Exchange theory [3], as individual differences lead to higher benefits from cooperation and thus more agents engaging in cooperative acts. Our research thereby supports the importance of individual differences for group outcomes [2].

References

1. Axelrod, R., Hamilton, W.D.: The Evolution of Cooperation. Science. 211(4489), 1390-1396 (1981).
2. Buss, D. M.: How Can Evolutionary Psychology Successfully Explain Personality and Individual Differences? Perspectives on Psychological Science. 4(4), 359-366 (2009).
3. Cosmides, L.: The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. Cognition. 31(3), 187-276 (1989)
4. Darwin, C.: On the origin of Species. John Murray, London (1859)
5. Doran, J.E., Franklin, S., Jennings, N.R., Norman, T.J.: On cooperation in multi-agent systems. The Knowledge Engineering Review. 12(3), 309-314 (1997)
6. Epstein, J. M., Axtell, R.: Growing artificial societies. MIT Press, Cambridge (1996)
7. Gintis, H.: The Hitchhiker's Guide to Altruism: Gene-culture Coevolution, and the Internalization of Norms. Journal of Theoretical Biology. 220(4), 407-418 (2003)
8. Jaffe, K.: Altruism, Altruistic Punishment and Social Investment. Acta Biotheoretica. 52(3), 155-172 (2004)

9. Kurzban, R., Neuberg, S.: Managing Ingroup and Outgroup Relationships. In: Buss, D. M. (ed).: The Handbook of Evolutionary Psychology, pp. 653-675. John Wiley & Sons, Hoboken (2005)
10. Mannix, E., Neale, M. A.: What Differences Make a Difference? Psychological Science in the Public Interest. 6(2), 31-55 (2005)
11. Mascaro, S.: Evolutionary ethics. Master's Thesis, School of Computer Science and Software Engineering, Monash University, Victoria, Australia (2001)
12. Mascaro, S., Korb, K.B., Nicholson, A.E., Woodberry, O.: Evolving Ethics: The New Science of Good and Evil (2010) Retrieved from http://www.csse.monash.edu.au/~korb/chap1.pdf
13. Nowak, M. A.: Five Rules for the Evolution of Cooperation. Science. 314(5805), 1560-1563 (2006)
14. Paul, D. B.: The selection of the "Survival of the Fittest". Journal of the History of Biology. 21(3), 411-424 (1988)
15. Pfeifer, R., Scheier, C.: Understanding Intelligence. The MIT Press, London (2001)
16. Richards, R. J.: A Defense of Evolutionary Ethics. Biology and Philosophy. 1, 265-293 (1986)
17. Ruse, M.: Evolutionary Ethics: A Phoenix Arisen. Zygon. 21(1), 95 – 112 (1986)
18. Spronck, P., Berendsen, B.: Evolutionary Ethics in Agent Societies. International Journal of Social Robotics. 1(3), 223 – 232 (2009)
19. Tooby, J., Cosmides. L.: Conceptual Foundations of Evolutionary Psychology. In: Buss, D. M. (ed).: The Handbook of Evolutionary Psychology, pp. 5-67. John Wiley & Sons, Hoboken (2005)
20. De Vries, M. J.: The Introduction of a Cooperative Act into an Evolutionary Ethics Agent Society, Bachelor's thesis, Tilburg University, The Netherlands (2014). Retrieved from: http://www.marjoleindevries.com/bachelors-thesis
21. Weerd, H. de, Verburgge, R., Verheij, B.: Agent-Based Models for Higher-Order Theory of Mind. Advances in Intelligent Systems and Computing. 229, 213-224 (2014)