Greed, Envy, Jealousy. How we can make Resource Management more efficient.

Mike Farjam

Radboud Universiteit Nijmegen

M. D. Farjam

Studentennummer: 0733857

Cursus: Bachelorscriptie Kunstmatige Intelligentie

Docent: Haselager, W.F.G., Sprinkhuizen-Kuyper, I.G.

Universitair Onderwijsinstituut voor Psychologie en Kunstmatige Intelligentie

Radboud Universiteit Nijmegen

12 Augustus 2011

Abstract

Highly social animals like humans developed features such as greed, envy, and jealousy through evolution. Assuming that the concept of jealousy has already been learned, experiments are performed in an artificial life environment. They show what the benefits of jealousy are for a multiagent system and how the underlying principles can be used making AIs more effective with respect to resource management. Furthermore they show under which circumstances (such as the population size or the possibility to punish greed) jealousy turns into a useful feature in a multiagent system. Abstract concepts like population size are translated back into real world concepts to show possible applications of artificial jealousy. Simulations show that the benefits in resource-management outweigh the costs of having a jealous system.

"You shall not covet your neighbor's house; you shall not covet your neighbor's wife, or male or female slave, or ox, or donkey, or anything that belongs to your neighbor." (Exodus 20:17)

We can find greed, envy and jealousy in highly social animals and across human cultures (Morris, Doe & Godsell, 2008). The question is why jealousy evolved as a universal feature in social interaction while it is viewed as something negative. This paper proposes that jealousy evolved because it brings some advantage to the species. This point was made explicit for the first time by Wilson (1978) and modifications of this point are popular in modern sociobiology and psychology (Lelord, 2003).

It is the aim of this paper to show through artificial life how greed, envy, and jealousy influence resource management in societies. Observations made in an artificial life environment that was programmed can be transferred to programs that autonomously share resources such as calculating capacities and make them more efficient. In the following paragraphs we will first show what we mean by jealousy and other related terms, then how we think that we can simulate jealousy. Finally we will describe how we expect jealousy to influence simulations we want to perform in an artificial life environment.

To ascribe a mental state, such as being jealous, to agents we use two criteria. Not just the behavior must be the same as in humans that we call jealous, but also the reason(s) for the behavior (Fraley, 2001). For humans the reason to be jealous is that they think that others have something they deserve more (Russell, 1930). Humans adjust the behavior (for instance, some kind of punishment) through which they express jealousy to different situations. Behavior depends, for example, on the costs of punishing or on the general disparity in the environment. Factors such as these will be incorporated into the simulations to understand how jealousy interacts with such factors. Later on, it will be shown how to translate these factors back, not only to problems in artificial intelligence, but also to generic sociological and psychological questions. By doing so, we hope to find mechanisms through which resource management (distribution of things valuable to an agent) become more effective. Throughout this paper we will reference two scenario's to which we

can apply the concepts, problems, and actions occurring in the artificial life environment. Through these scenarios the reader on the one hand will see what certain concepts in this experiment stand for, and on the other hand see possible applications of the findings. These scenarios are not intended to prove anything. They surely have their shortcomings and may be interpreted differently. However they show the external validity of this paper.

Scenario 1 is taken from everyday life. Humans live in societies where some earn more money than others do, have more free time than others, etc. Jealousy and dissatisfaction occurs when disparity is too high. People can punish greedy and wealthy people in many ways such as excluding them socially or trying to take them to court. In both cases punishment will cost both parties resources (e.g. money, time, friends). People in such societies often try to avoid the negative consequences of jealousy and instead prefer to decrease disparity by sharing some of their resources (through such institutions as foundations, social welfare, amicable agreements). Humans may also deviate strongly in the amount of money they need to survive. For instance, while earning amount X as a single may be sufficient, a father of a large family may need much more to make ends meet.

Scenario 2 is taken from the field of mobile communications. Mobile telephones find the broadcasting tower that is closest to them and try to connect to it. At some point some broadcasting towers may be doing nothing whereas others at another place (e.g. the subway station) are chronically overtaxed. Wouldn't it be nice if towers found a way to balance the work in a way that would provide a good connection to every mobile phone without overtaxing any one connection point? Figure 1 shows a scenario where a tower is doing nothing although it could easily "help"(share its free capacities) another tower by taking over some of the other tower's mobile phones, even though it is not closest to these phones. Making towers jealous at other towers, which have more free capacity than they have, could help.
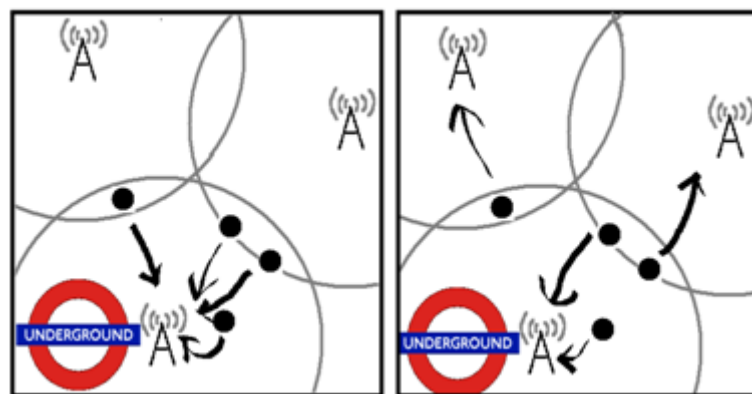
Figure 1: Scenario 2 (dots are phones, arrows show the tower used)

De Jong (2009) showed in his PhD.-thesis that fairness is beneficial for the total reward of a group when resources have to be shared. Humans automatically and often subconsciously punish unfair behavior to force individuals to act fair. Punishment can be social (e.g. excluding individuals from the group) or material (e.g. a penalty or fine).

De Jong showed in his experiments that individuals will show more egoistic behavior if there are no such possibilities. On an inter-agent level we need possibilities to punish selfish behavior to rein in egoistic, destructive behavior within the group.

On an intra-agent level there seems to be an intuitive link between greed and envy. Why envy someone when one does not want more? The next step towards envy is to make our agents "feel" less content with something they have, while observing others who have more. Fehr and Schmidt (1999) developed a utility-function with exactly these properties. Through their utility function they can explain why we often feel better off getting no reward, than getting a small reward while observing others getting a much higher one.

Humans attribute mental states to themselves and others and are thereby able to predict the behavior of others (Premack & Woodruff, 1978). By knowing what envy is and the assumption that others "feel" envy for the same reasons, they are able to predict when envy will occur as a result of their actions. If it is general knowledge that envy will lead to punishment against the one that is envied, agents will take the punishment in consideration for their actions and expected utility.

An artificial-life environment is programmed in which agents try to gather (greedy) as much resources as possible. Furthermore they have the possibility to punish agents that have more resources than themselves or share resources when others have less. Agents will die after a certain time, but can prolong their life when they get resources.

In the experiment conducted we assumed that jealousy already has been learned. Agents in the artificial-life environment will automatically punish other agents in their neighborhood when they are much richer than themselves. Furthermore agents will try to prevent punishment by sharing with agents that are much worse off than themselves.

The question remains why one should make the effort to implement greed and jealousy and not simply let agents share resources with each other. There are at least two big advantages to taking a detour before implementing sharing. First of all, the decision to share resources very often consumes resources. When a computer decides to take over a computation, making this decision and initializing variables will consume computation time. The same holds for humans or animals. Sometimes it is more effective to do something ourselves even though we do not have much time, than explaining the task to someone else who has more time. Simply sharing resources for the sake of equilibrium may be ineffective. Adding an agent's egoistic perspective while deciding whether it will benefit from sharing or not may ultimately spare resources. The second advantage of jealousy

above simple sharing is that jealousy will lead to punishment and punishment provides a possibility to learn effective sharing. In a training phase, agents would learn through punishing each other how they can avoid being punished through sharing. Even after training, punishment will enable agents to adapt their sharing-behavior to a changing environment. In a changing environment this flexibility in sharing may exceed the amount of resources that get abstracted from the system through punishment (just as sharing costs resources, punishment does as well).

We propose that jealousy, although it has a bad reputation, will make resource management more effective than simple greed alone without jealousy. How does resource management change in a multiagent system when there is jealousy and punishment? We will take a close look at the behavior changes per agent, and the changes for different kind of environments or agent populations. We hope to find certain kind of environments where the effect of jealousy is the highest and want to propose parameters for interactions (such as costs of punishment) that make jealousy most effective.

Method

Software

The open-source software "breve 2.7.2." (Klein, 2002) is used to implement an environment to simulate jealousy in agents. The program is written in breve's own language. This language is object oriented, has a huge library with functions that are normally needed in artificial life programming, and it fits the demands of this experiment very well. A master object initializes the environment and will start a new simulation with different parameter settings whenever an old simulation is completed. For later statistical analysis, a record is kept of the age that each agent in a simulation reached, and the parameters with which the simulation world was initialized. The parameters of the simulations will be explained in the following paragraphs.

The environment

All simulations where conducted in an identical environment, a world with a square area surrounded by walls. Within this area fruits (representing resources) are popping up randomly. Figure 2 is a snapshot from one of the simulations. Fruits start with a certain amount of energy and lose energy with every iteration of the simulation. When the energy of a fruit is zero the fruit disappears. The amount of fruits that pop up per iteration depends on the value of the variable `fruitsAvailable` and varies between low (if no fruits are eaten 0.15% of the area are covered with fruits), midlow, midhigh and high (1.05% of the area is covered with fruits if no fruits are eaten). Agents will be set in this environment at random positions. The amount of agents varies between

simulations. As with the fruits, 4 values are possible for agent variable (between 0.5% and 1.5% of the area will be initially covered by agents).
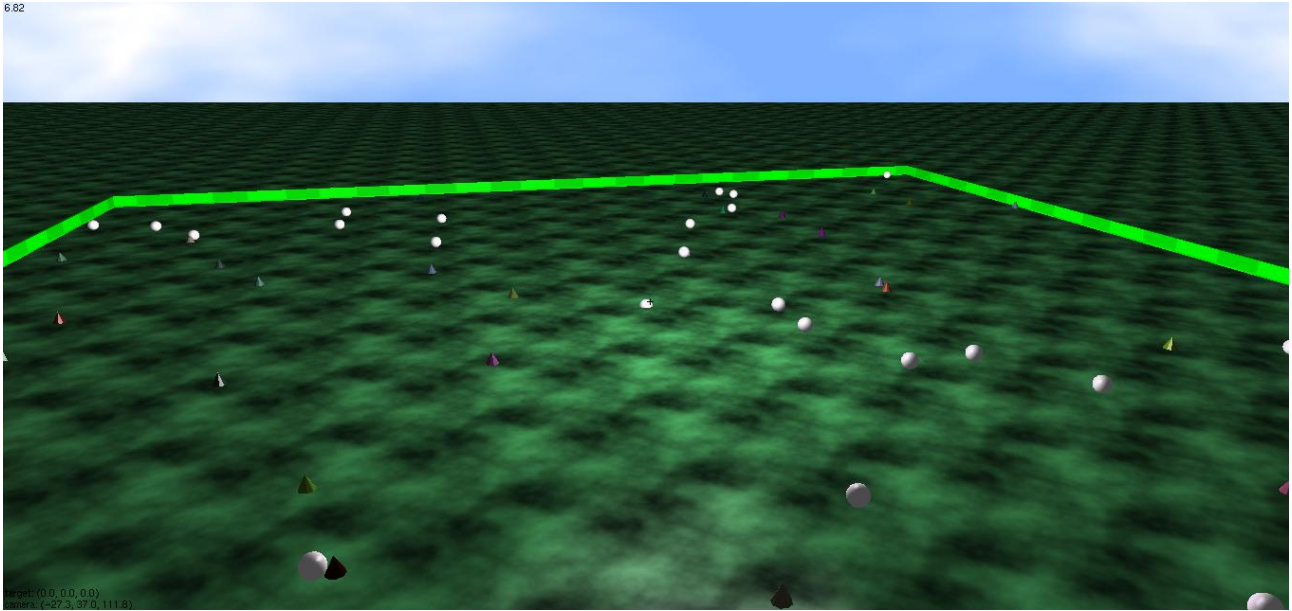
Figure 2: Balls = fruits, colored cones = agents, area is surrounded by green walls

Agents

Initially, agents move in this environment randomly and have the same amount of energy. Within a fixed range, agents know how "rich" other agents are and have the possibility to punish/share with them. This area is called the agent's neighborhood. Furthermore agents can see fruits within this range and try to approach any fruit that is visible and close to them. When they reach the fruit they will eat it, meaning that they take the energy that is left in the fruit and add it to their own.

With every iteration of the simulation, agents lose energy. Agents will die as soon as their energy level is 0. The average age that agents reached in a simulation before dying was taken as the performance measure of the simulation. Possible values for the simulation's variables were chosen such that the amount of fruits which can be consumed will never exceed the amount of energy that is lost per iteration in the long run. Eating is just postponing certain death but cannot stave it off indefinitely.

On average, agents in every simulation consume the same amount of energy per iteration, but per agent there is a deviation from the average possible, depending on the variable `heterogeneityRescNeed`. Four values are possible ranging from 0 (every agent within the simulation consumes the same amount of energy per round) to 0.2 (the energy consumed ranges between 80 – 120 % of the average).

Besides eating and moving in the environment agents can interact with other agents in their neighborhood. The neighborhood was defined within a radius of 1/50 of the edge length of the

simulation area. Within their neighborhood agents know the energy level of other agents and can compare it to their own level. Agents share resources with the poorest neighbor when the difference is too high and positive and if they will have more than 50 energy-points left after sharing. When the difference is too high, negative and agents would have a reserve of < 50 energy-points left after sharing, they punish the agent that is the richest in their neighborhood (jealousy). Whether a difference is seen as too high or not depends on the variable `toleranceHeterogeneity`. The variable can have four values ranging from 0-20. A value of 0 means that even the smallest difference will lead to action, while a value of 20 means that action will only be taken when the difference exceeds 20. Punishing (as proposed in De Jong (2009)) and sharing costs energy. How expensive an action is depends on the variable `actionEffectRatio` with 1, 7.6, 13.6 and 20 as possible values. A value of 1 means that punishment will cost the punisher as much as the punished, and a value of 20 means that punishing will only cost 1/20 of the effect it has on the one that gets punished.

Table 1 gives an overview of the variables and concepts used in the simulations and how they would correspond in scenario 1 and 2. All independent variables (in italics) were made discrete for easier analysis later on.

Concept in simulation	Values	Scenario 1	Scenario 2
agent	Object	human	tower
age (performance)	positive integer	money	call quality
neigboorhood	part of area	other humans visible	towers within certain range
jealous	boolean	jealousy	jealousy algorithm
populationDensity	0.5, 0.83, 1.16, 1.5	humans in area	number of towers in area
fruitsAvailable	0.15, 0.45, 0.75, 1.05	GDP, Jobmarket	1/ (phones in area)
actionEffectRatio	1.0, 7.6, 13.6, 20	legal costs	computations needed for an action
heterogeneityRescNeed	0, 0.66, 1.33, 2	how many people do I have to feed?	crowded areas – calmer areas
toleranceHeterogenity	1.0, 7.6, 13.6, 20	feeling about acceptable disparity	acceptable differences in workload

Table 1: Overview of concepts (*independent variables italic*)

Simulations

248 simulations were performed. Every simulation received a randomly chosen value for the 6 independent variables as input. The 248 simulations together imitated about 10 000 agents and each simulation had a similar number of agents. For every agent the age reached + the parameter values of the environment it lived in (in terms of the independent variables) were stored for later analysis. In each simulation agents got a random starting position and throughout the simulations fruits were positioned randomly.

Results

All data was tested for normal distribution. The jealous and the not-jealous group had a high skewness meaning that some agents died early while some became very old. This seems sensible as

there would be more resources left for other agents as one agent dies, leading to plenty of food for the last agents. It is remarkable that the jealous group scored higher in skewness (2.8) and kurtosis (9.0) than the not-jealous group (1.5 and 5.2 respectively). This means that age is more homogeneously distributed in the not-jealous simulations than in the jealous ones. Taking a closer look at the histograms (figure 3) it appears that this is the result of a larger (although still relatively small) group of agents in the jealous group who live much longer than the average agent. It seems that the agents who live longer (perhaps those with low energy-costs/iteration) benefit from their remaining society-members more in the jealous societies and become more effective in their resource-management through jealousy in the later stages of the simulation. If we leave out the highest quartile of agents in both groups the distribution and average is almost the same.

Although the data is not completely normally distributed, figure 3 shows that it comes close. The size of the sample used for analysis is quite large and t- and F-tests are robust to violations of normal distribution when the size of the sample gets large. Therefore both kind of tests were performed, while keeping in mind that p-values have to be interpreted with care.
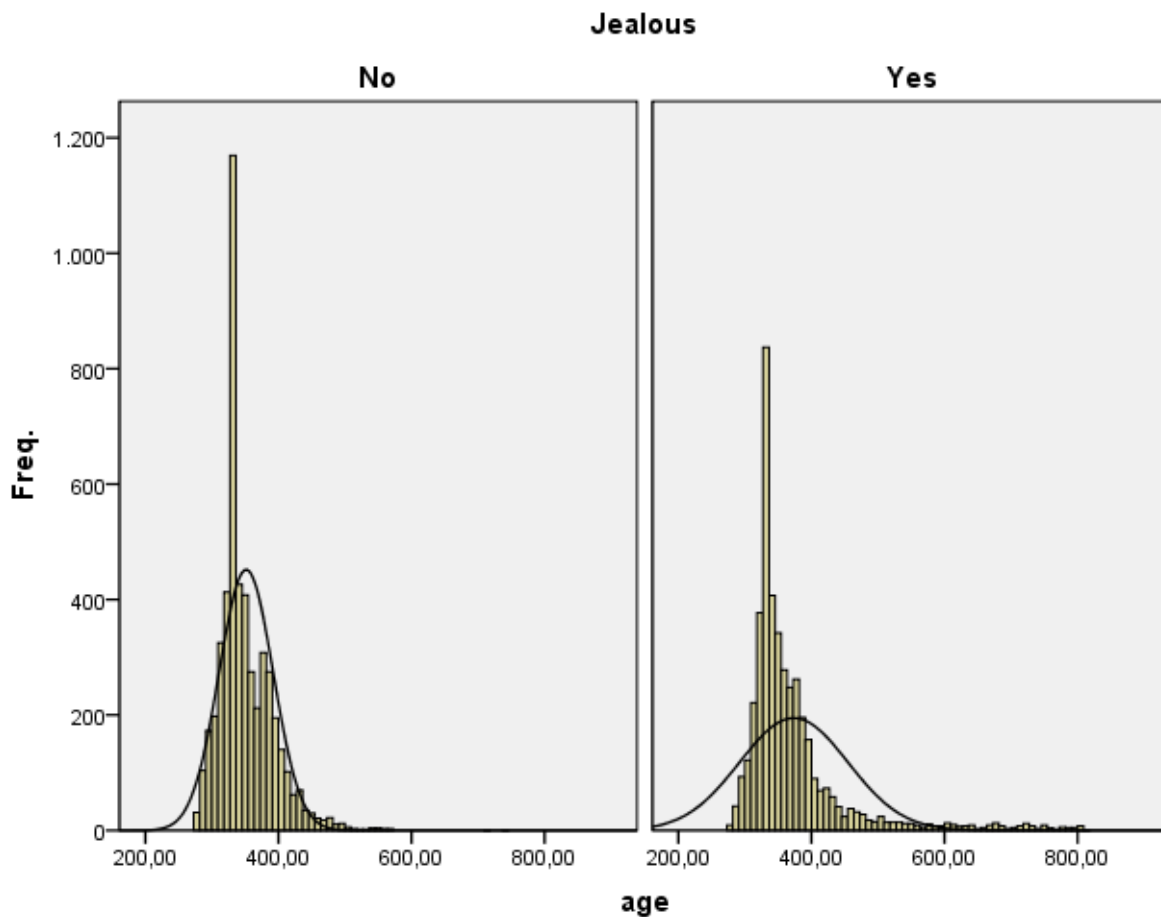

Figure 3: Distribution of age (normal distribution curve as solid line)

Performing a t-test brings significant results (p < 0.001) showing that the group of jealous agents (M = 373.37) performs better in terms of average life length than the not-jealous group does

(M = 351.39). A Mann-Whitney U test (not assuming normal distribution) leads to the same results. The jealous group had an average rank of 4172.54 the other group 6002.46 (p < 0.01).

To analyze the role of each independent variable an ANOVA was performed with all independent variables included. The overall tests showed that all variables together could explain variances of age in the sample (p < 0.001). Every variable had a significant influence on the average age of agents. Furthermore all differences between values of variables were significant. All p-values where below 0.001, which is not surprising when taking into account the large sample-size used. Looking at the effect-sizes it appears that all independent variables and their interactions could only explain ¼ of variance found in age of agents. This means that an agent's life was influenced greatly by the random factors in the simulations (their initial position, the position of other agents and location of fruits). Hence effect-sizes found for the effects of each variable and the interaction of the variables are likely small. The largest effect-sizes of any individual variable was found for jealousy (0.02) and tolerance-Heterogeneity (0.021). Effect sizes for ActionEffectRatio (0.014) and availabilityOfResources (0.018) were also quite high, whereas populationDensity (0.003) and heterogeneityRescNeed (0.004) had a fairly small effect-size. This paper only looked at unique effects of variables on age and all two- way interactions with jealousy. These effects together explained 17% of the variance of age.

General findings

For all variable values of the independent variables (except for the lowest value of ActionEffectRatio), agents had a higher life expectancy in the jealous group compared to the not-jealous group. No matter what the values of the independent variables were, jealousy had a positive influence for the age of the average agent.

Another general effect of jealousy appears to be that it flattens the distribution of age (also visible in Figure 3). For most independent variables and their values there is a peak of agents dying at an age of ~300 iterations. This peak is lower for most independent variables in the jealous group compared to the not-jealous group. In the following paragraphs the interactions between jealousy and the other independent variables will be explored in more detail.

Population density

In the group of not-jealous agents, average life expectation gradually decreases with an increase of population density. This seems plausible as there is less to eat for every agent when there are more other agents around. Agents in the jealous group live longer for all variable values of population density than the agents in the not-jealous group. Interestingly, the average age of agents in the jealous group increases with a higher population density. However there seems to be no linear

relationship between population-density and age.

Figure 4 shows a fluctuation in average age from 0.83 to 1.67 in the jealous group. The difference between the two values within the jealous group is significant. We can just speculate on the reasons for this fluctuation and its implication. We expect that the interaction jealousy X populationDensity acts together with variables such as availabilityOfResources on these values. However, as described above, we do not look at higher order interactions in this paper and leave this to future research.
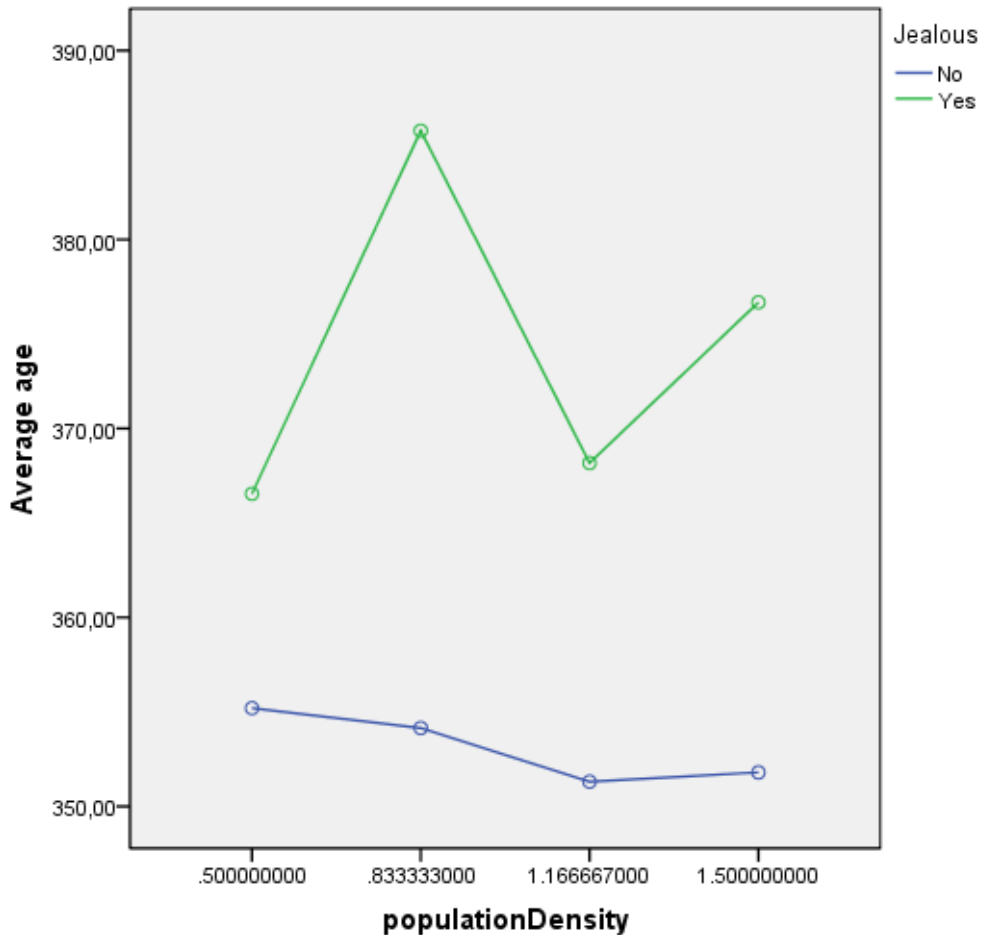

Figure 4: Interaction-plot jealous X popoulationDensity

Availability of resources

Generally, both groups profit from a higher availability of resources. More resources means more energy per agent and thereby a longer life. Again all values of the variable show a higher average age for the jealous group. This difference is stable for almost all values of the variable availabilityOfResources (around 20 iterations). Jealousy not just increases the performance for all values of availabilityOfResources but also flattens the early peak at 300 iterations.

The most noticeable difference between the two groups can be seen for the highest value of

availabilityOfResources. The early peak of agents dying vanishes in the jealous group showing that the number of agents dying stays quite constant throughout the simulations.
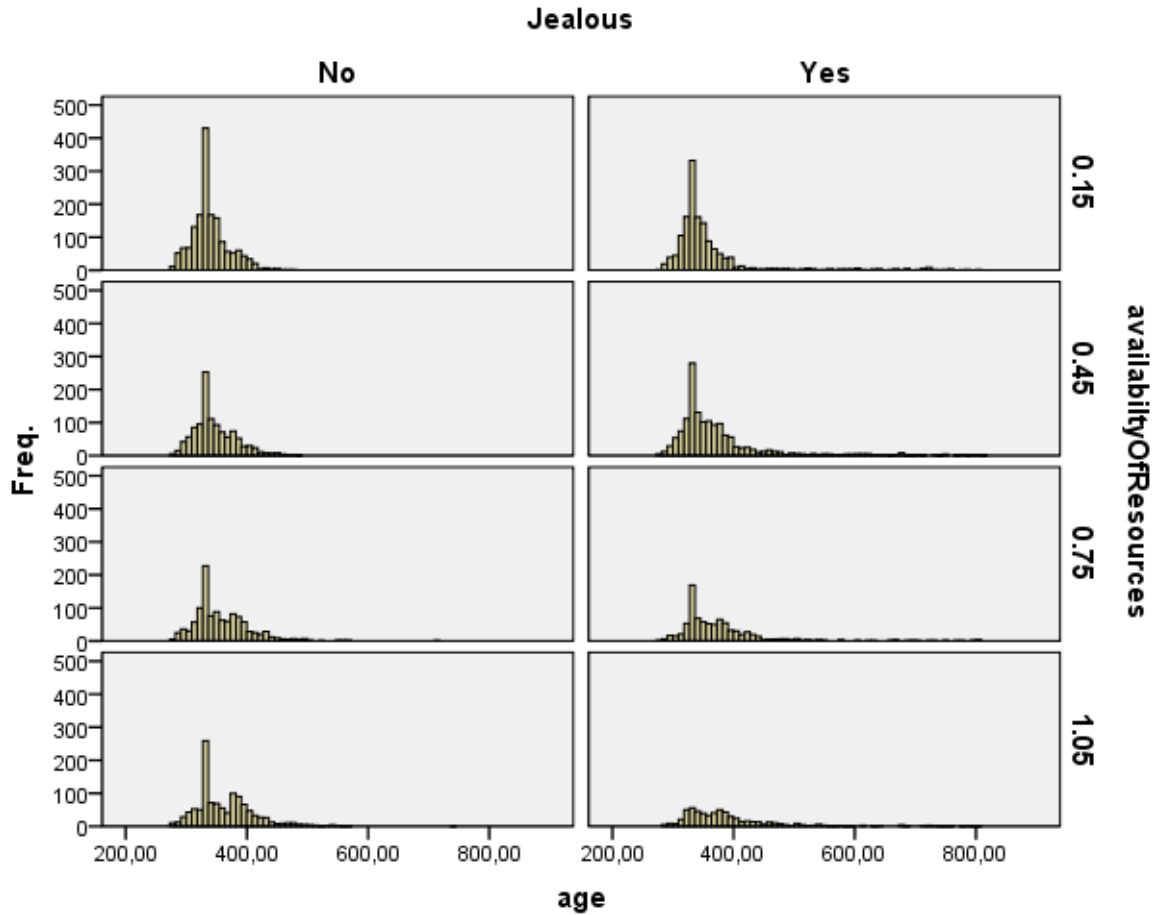

*Figure 5: Age in all subgroups of "availabilityOfRescources" and jealousy*

Heterogeneity resource need

As with all other independent variables the jealous group outperforms the not-jealous group at all values of the variable. The data does not show any bigger interaction between jealousy and "HeterogeneityRescNeed". Accordingly, this interaction had the smallest effect-size of all two-way interactions (0.002). The distribution of the jealous and not-jealous groups show no major differences.

Action-effect ratio

The not-jealous group stays almost constantly at an average age of 352, which is sensible as this variable does not affect the behavior of not-jealous agents. Here we find the only case where not-jealous agents outperform jealous agents. At a very low ratio the cost of jealousy carries more weight than the benefits of sharing. Furthermore the data suggests that there is a value for the ration at which it maximizes the advantages of jealousy. Whether this is a local or global maximum cannot

be seen from the data. At the ratio of 13.66 we find the second largest difference (43 iterations) between the jealous and not-jealous group, not just for the variable "ActionEffectRatio", but also for all other sub groups that can be made within independent variables. Although small, the ActionEffectRatio variable had the largest two way interaction with jealousy (effect size = 0.01).
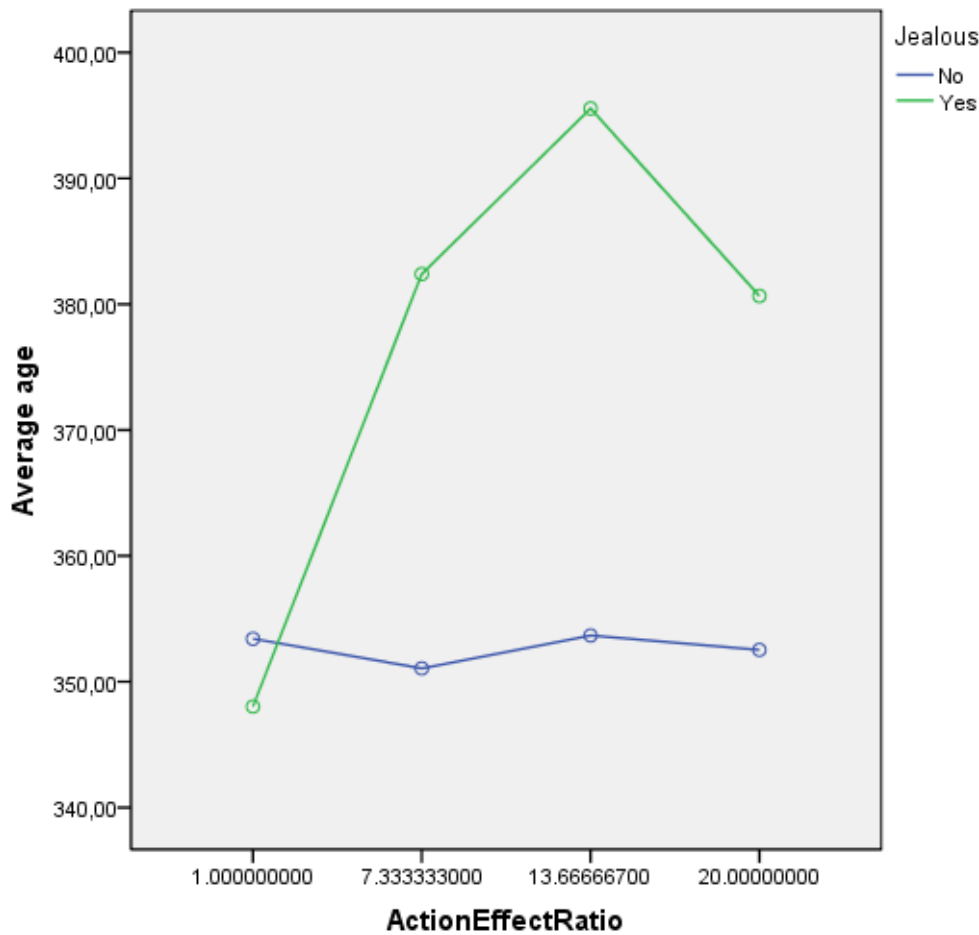
Figure 6: Interaction-plot jealous X ActionEffectRatio

Tolerance for heterogeneity

As with "ActionEffectRatio" the variable ToleranceHeterogeneity only influences the behavior of agents in the jealous group, therefore the average age stays constant at ~352 iterations.

A tolerance of 0 leads to the highest difference between the jealous and not-jealous group. This difference is the highest (55 iterations) of all differences found in the comparison between jealousy and the other independent variables. It seems that no tolerance for disparity makes jealousy the most effective. The two-way interaction scored second in terms of effect-size (0.009). When agents become too tolerant for disparity in their neighborhood there is almost no difference between the two groups. The relationship between tolerance and age in the jealous group appears to be linear and negative, suggesting that jealousy becomes less effective as agents tolerate disparity.

Although figure 7 suggests that jealousy should occur immediately, figure 6 tells us that

jealousy will lead to suboptimal results when punishing becomes too cheap. In other words `be jealous, but don't always let your actions be determined by it`.
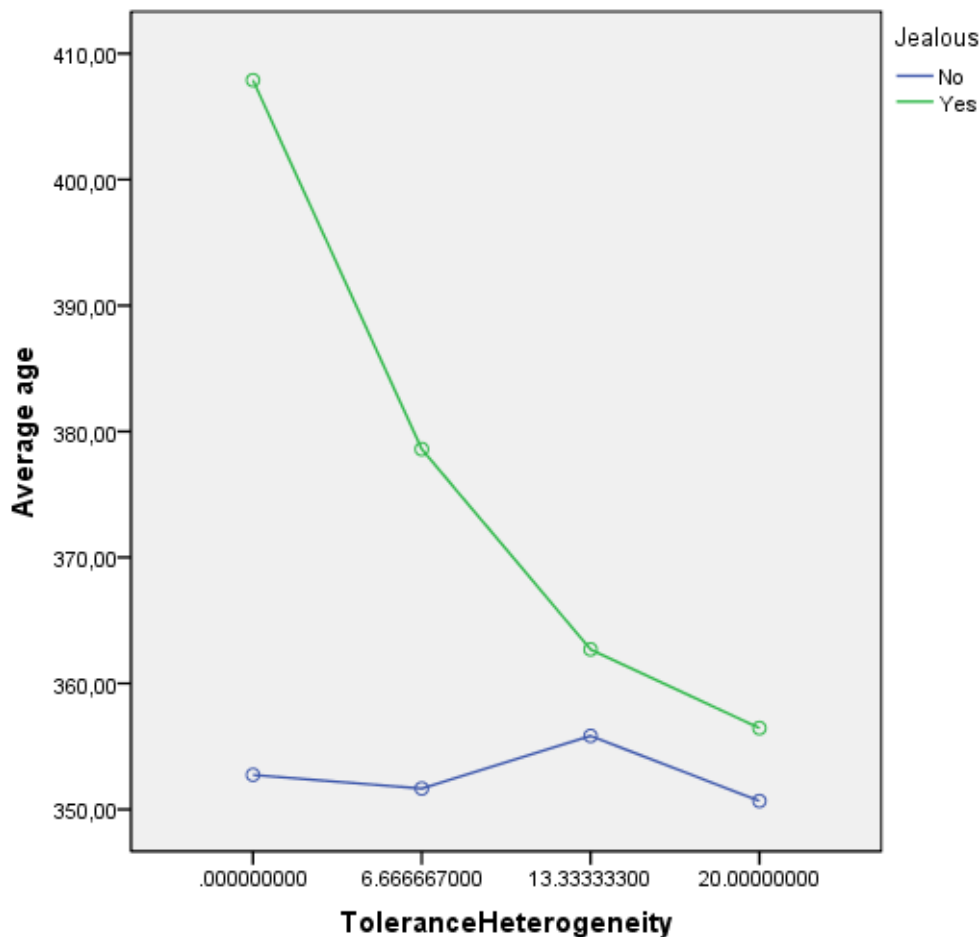

Figure 7: Interaction-plot jealous X ToleranceHeterogeneity

Decision-tree

To understand which variables influence the effectiveness of jealousy most, a C4.5 algorithm built a decision-tree from all simulations where agents were jealous. The C4.5 algorithm was trained with 200 randomly chosen simulations to decide whether a certain set of values for the independent variables would lead to an above average lifetime in the simulation (M = 363). The performance of the decision-tree was tested with the remaining 48 simulations that were not in the training-set. Performance was quite high (95% with a depth of 3) for the group of jealous agents.

The advantage of having this decision-tree is that we can say which features of the environment are most important for making jealousy effective. Figure 8 shows the classifications made by the decision tree, where a value of 1 stands for the classifications where the simulation will lead to an above average age for agents.

The decision-tree shows that the most important variable to predict whether a jealous population will become old is the availability of resources (R). This makes sense as there is an

intuitive link between food (resources) and survival (age). More interesting though is diving down one further level on the tree. On the left side we can see that jealousy can help a population to live longer even when availability of resources is low. By keeping the costs for punishing/sharing low (thus the "action/effect ratio"(AER) high) 6 out of 11 populations can overcome the negative effects of a low availability of resources through jealousy. Numbers get even better if we consider not just the action-effect ratio but also the tolerance for heterogeneity (TH) within the population. When TH is low and AER is high, 6 out of 7 populations will have an above average life expectation. In the not-jealous group only 1 out of 16 had a life expectation above average with the same availability of resources. The right part of the tree does not give that much information, except for the fact that heterogeneity of resource need (HRN) needs to be carefully chosen in environments with higher availability of resources.
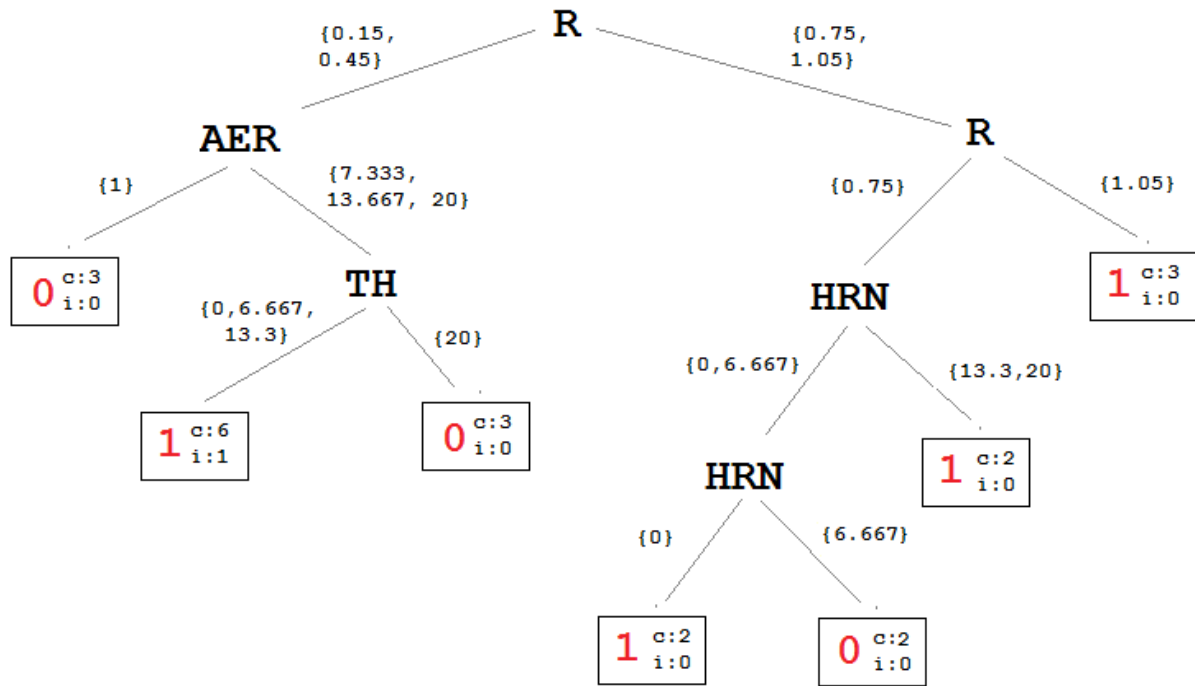

Figure 8: Decision-tree for jealous agent simulations (variable values in brackets)

Discussion

Simulation findings support the hypothesis that jealousy can help making resource management effective. Furthermore the simulations show how jealousy interacts with certain other variables of the environment. Through the help of the two scenario's (mentioned in the introduction) we will discuss what the external validity of these findings might be.

In **Scenario 1** we draw parallels between the simulations and real life in human societies. The results suggest that jealousy has a worse reputation than it deserves, as it can help sharing

14

resources among larger groups. To make jealousy effective the costs of expressing this jealousy (*ActionEffectRatio* in the simulation) should be rather low. Jealousy towards people that do not share their wealth is mostly expressed by social exclusion or in extreme cases through legal actions. Lowering the costs of legal actions for poor people is, obviously for a good reason, part of many western societies. For social exclusion, the excluded should lose more friends than the excluder to make jealousy effective. Another important finding is that the positive effects of jealousy unfold best when the tolerance for disparity of wealth is low (*ToleranceHeterogenity* in the simulation). This could be an important argument in public debates about social redistribution (a society without any tolerance for disparity is the self-proclaimed communist ideal). The findings hold for societies with a high diversity in resource needs as well as for societies with quite homogeneous resource needs. From a politically point of view, the fact that jealousy keeps the average wealth stable with increasing population density might also be very interesting. Lastly, it appears that without jealousy average wealth decreases when *population density* increases.

In **Scenario 2** we show how multiple broadcasting towers decide which phones in their area to take over and which to leave for other towers. The amount of calls a tower has to handle is comparable with the resource need of agents in the simulations. Resources are free capacity of a tower (negatively correlated to phones in the area of the tower). The simple heuristic "every phone should be assigned to the tower to which the phone is closest" is not very effective. If towers could be jealous towards other towers that have more free capacity than they do, towers could get a better arrangement (thus higher connection quality per phone). Similar to the findings in scenario 1 and in the simulations, towers should not be too tolerant to differences in free capacity for jealousy to have maximum effectiveness (*ToleranceHeterogenity*). Furthermore, the algorithms that implement the punishment should not require too many resources (high *Action-EffectRatio*), otherwise jealousy may have negative consequences for overall connection quality. The findings for population density suggest that effectiveness of jealousy does not depend on the density of towers in the area as (e.g. jealousy can be effective with any *populationDensity*).

The question remains whether we can say that jealousy is effective just because the jealous group had a higher average life length. If we do not take the average as our criterion but e.g. the median we find that the two groups almost do not differ from each other. There may be cases were we are just interested in the average performance of a multiagent system or a society but there a clearly also scenarios where it is the median that is most important. Scenario 2 may be a good example, where extreme good quality for some phones, but bad quality for many may be worse than a mediocre quality for most phones even if the average quality is lower.

One could argue that the effect-sizes found are very small (especially for experiments in such a controlled environment) and that the results therefore are not relevant. However, despite the

small effect-sizes, jealousy boosted performance of agents by 20 iterations (7%) in all simulations and under optimal circumstances by even more than 50 iterations (~20 %). An additional reason for the small effect-sizes might be that, while programming the environment, the simulations included as many random variables (location of fruits, location of agents, etc.) as possible to make them more realistic as a model, while still being applicable for dynamic environments. In the two scenarios described there may be many comparable random factors.

In our analysis we only looked at unique effects of the independent variables and two-way interactions with jealousy. Future research will help us to get a more realistic idea of the effects of jealousy by looking at higher-order interactions with jealousy and other independent variables.

A general remark we have to make on the data used and the conclusions drawn is that the dataset did not contain all possible values for the independent variables. Statements like "if x increases, average age increases as well" may not be generally true. The simulations used just a limited range of values for each variable. Relations that seem linear may just be a part of the function that describes the relation between the independent variable and age. Furthermore, no conclusive statements about optimal values for these variables can be made, as maxima that were found may just be local. Limiting the range of variables (like availability of resources) was needed though to make sure that agents die eventually. Future experiments, with a different environment and other performance measures than age, may be needed to validate findings and overcome artifacts introduced in the simulations.

Lastly, the simulations assumed that being jealous was already learned. The programmed agents could not adapt their behavior yet. This was not needed as the environment did not change its behavior within one simulation. But in the two scenarios described the environment will change permanently: i.e. population density changes, sometimes there are more free resources than at other times, etc. The question how agents can handle jealousy flexibly might make for further interesting study.

Humans have more civilized ways to express jealousy than the agents had. Punishments create costs of resources not just for the one who punishes and is being punished, but also extracts resources from the community in general. Humans therefore try to find other ways to express their concerns than physical punishment. Any negative effects punishment had on the performance of the society could be lowered by finding "more civilized" ways to express oneself than just harming the other directly. We expect that this will increase the performance drastically.

Many other independent variables and higher-order interactions have to be analyzed to get a deeper understanding of jealousy. The amount of variables and interactions analyzed here was quite small as this paper was meant to make a first attempt towards realizing resource management through jealousy. Critics can also argue that the operationalization and implementation of jealousy

has its shortcomings (which it surely has), but nevertheless this paper showed that a rough concept of jealousy can increase performance. Not just that, jealousy allows agents to weigh their own interests against the interests of others. We have to keep in mind that all computations were done locally and autonomously by the agents, computations have to be fairly easy for punishment to provide a learning mechanism through which an agent can learn to be jealous in a changing environment. This all makes jealousy a concept that can effectively increase performance in multi-agent systems and we encourage others to study the effect of jealousy on resource management further... we promise not to be jealous!

Reference

De Jong, S. (2009). *Fairness in Multiagentsystems.* Unpublished doctoral dissertation, University Maastricht, the Netherlands.

Fehr, E., & Schimdt, K. (1999). A Theory of Fairness, Competition and Cooperation. *Quart J. of Economics, 114,* 817–868.

Fraley, L. F. (2001). Strategic interdisciplinary relations between a natural science community and a psychology community. *The Behavior Analyst Today, 2,* 209–324.

Klein, J. (2002). Breve: a 3D simulation environment for the simulation of decentralized systems and artificial life. *Proceedings of Artificial Life VIII, the 8th International Conference on the Simulation and Synthesis of Living Systems.* The MIT Press.

Lelord, F., & André, C. (2003). *La Force des émotions.* Paris: Poche.

Morris, P. H., Doe, C.. & Godsell E. (2008). Secondary emotions in non-primate species? Behavioral reports and subjective claims by animal owners. *Cognition and Emotion, 22,* 3-20.

Premack, D. G. & Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences, 1,* 515-526.

Russell, B. (1930). *The Conquest of Happiness.* New York: H. Liverwright.

Wilson , E. O. (1975). *Sociobiology: The new Synthesis.* Oxford, England: Belknap Press of Harvard U Press.