

VU Research Portal

Evolution of a Risk Coefficient in Artificial Societies

van der Goes, V.A.N.M.; Eiben, A.E.

published in

2010 IEEE Congress on Evolutionary Computation (CEC)
2010

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

van der Goes, V. A. N. M., & Eiben, A. E. (2010). Evolution of a Risk Coefficient in Artificial Societies. In *2010 IEEE Congress on Evolutionary Computation (CEC)* (pp. 1-7). IEEE.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Evolution of a Risk Coefficient in Artificial Societies

Vincent van der Goes and A.E. Eiben

Abstract—In this paper we investigate how life expectation influences the development of risk attitude within an artificial species. Our hypothesis is that agents with a very long life span are likely to become more risk averse because they have more to lose. To assess this hypothesis we set up a simple system, based on Sugarscape, where risk attitude is an inheritable (hence, evolvable) property. Performing numerous simulations with different versions of this system we found that long-lived agents consistently and clearly evolve a more risk averse behavior than short-lived agents. Perceiving evolution as a general force towards optimal behavior, these results indicate that increased risk avoidance is a generally good strategy for agents with a higher life expectation. This finding can be used to explain various real-world phenomena. For instance, it can clarify the fact that people tend to adopt risky strategies when their life is in danger.

I. INTRODUCTION

Behavior under uncertainty and risk is an important topic in the social sciences. Uncertainty is a pervasive element of everyday life and present in even the most simple situations. Consumers daily need to make decisions involving chance. Within both the fields of economics and psychology, models have been developed to capture human behavior under such conditions.

In this paper we investigate how factors in biological evolution can influence the development of risk attitude within a species. In particular, we are interested in the role of life expectation. Is a short-lived species likely to exhibit a different type of behavior under risk than a long-lived species? Our hypothesis is that since agents with a very long life span have more to lose, they are likely to become more risk averse.

Multi-agent systems provide a suitable tool to research this kind of question [4] [6]. To answer this question we set up a simple system of agents with basic properties of a natural species. They breed and depend on the consumption of a resource. The multi-agent model is inspired by and similar to the Sugarscape [4]. Agents move around a grid to gather their food, sugar. They are in constant need of sugar for consumption and when they don't manage to gather enough of it, they die from starvation. However, if they manage to gather sugar a large amount of sugar in excess of their needs for consumption, they can use it to create offspring.

The system has been modified, to introduce an element of risk. In the process of gathering sugar, the agents are now offered a choice between a safe method of harvesting and

a more risky method. The risky method can yield a greater amount of sugar, but at the risk of gaining no yield at all.

Agents decide on this issue in accordance to a genetically encoded risk aversion parameter. Agents with a lower parameter tend to opt for the safe method of harvest, while agents with a higher parameter tend to use the risky method. The risk parameter is inherited over generations and subject to mutation. Thereby, it can evolve over time. We investigate the evolution of the risk parameter, while varying the maximal age of agents.

This research is an extension of earlier work [12], where the same research question was also addressed with a multi-agent system. However, in this paper we test the same hypothesis for a much larger set of parameter settings and come up with an improved method of measuring risk aversion. In contrast to earlier work, we manage to evolve a coefficient of relative risk aversion, which provides a quantitative measure of behavior under risk.

The remainder of this paper is set up as follows. Section II gives a brief introduction to the underlying theory of risk models. A detailed description of the multi-agent model can be found in section III. Section gives an overview of the experiments we performed and the parameter settings that have been used. The results are in section V and in section VI we conclude and suggest directions for further research.

II. ECONOMIC BACKGROUND: EXPECTED UTILITY

In this section we outline the basics of the theory of expected utility maximization, which is considered as the normative model of choice under risk in economics. The theoretic concepts outlined here will serve as the base of the mental model of our agents. The main advantage of having agents following the expected utility model is not that it realistically models biological evolution in nature, but rather that it provides a simple and elegant mental framework, where the degree of risk aversion can be tuned with a single parameter with a clear meaning.

A. Expected utility maximization

In the field of economics, various formal models have been developed to capture human behavior in situations where outcomes are subject to probabilities. In this context we speak of *behavior under risk*, when the set of possible outcomes of each action is known, as well as their respective probabilities. When the probabilities are unknown beforehand, or when even not all of the possible outcomes are known, we talk about *uncertainty*. This paper deals exclusively with decision under risk. The agents have full knowledge of their prospects whenever they need to make a decision.

Vincent van der Goes is with the Department of Economics and Business Administration, VU University Amsterdam, The Netherlands; email: vgoes@feweb.vu.nl

A.E. Eiben is with the Department of Computational Sciences, VU University Amsterdam, The Netherlands; email: gusz@few.vu.nl

Agents in our model will behave according to the model of *expected utility maximization*. This is considered to be the rational behavioral model [2] [8], even though it is widely considered as outdated for the purpose of describing human psychology (see, e.g., [7], [1]).

Consider a case where an agent is facing probabilities of gains or losses of a single resource. This resource could be money, as is often the case in economics, but it could also be another resource. Within the scope of this work, it will be sugar.

Further suppose that the probabilities of specific gains or losses depend on the behavior of the agent. The agent can choose between several different *actions* A_i . After the agent has made a decision, performing the action can result in several different *outcomes*, depending on random chance. Each outcome is represented by an amount of resource x that the agent gains. When x is negative, it is a loss rather than a gain. The possible values of x after taking action A_i are x_{ij} ($i = 1, 2, \dots, n; j = 1, 2, \dots, m_i$). The probabilities of each of those outcomes to occur, given that the agent has chosen action A_i , is p_{ij} .

The theory of expected utility assumes a rational agent, having given preferences regarding the possible outcomes. The preferences are assumed to satisfy a number of conditions, such as consistency and completeness. It is also assumed that the agent has full knowledge of x_{ij} and p_{ij} . Under these assumptions, expected utility theory explains how the agent should rank the possible actions in order to optimize his prospects [2].

The expected utility model is based on the concept of a *utility function*. Informally speaking, utility measures the happiness of an agent at a given wealth level. The utility function maps every wealth level w to a utility level $U(w)$. The utility value corresponding to a loss or gain of x , given an initial wealth level w_0 , is then $U(w_0 + x)$.

In order to rank the possible actions, the agent calculates the expected value of the utility of the outcomes for each action, as in equation 1. The best action is the action producing the highest expected utility.

$$EU(A_i) = \sum_{j=1}^{m_i} p_{ij} U(w_0 + x_{ij}) \quad (1)$$

When there are several actions maximizing expression 1, the agent is *indifferent* between those actions. The theory makes no prediction of the choice of the agent in this case.

An important feature of expected utility theory is that the attitude of an agent toward risk is completely identified with the utility function. If we know the utility function of the agent for every wealth level, we can predict the choice for any given action A_i , x_{ij} and p_{ij} .

However, utility functions are not uniquely defined. Adding a constant to it would not change anything about the behavior of the agent, since the choice of an agent depends only on the relative quality of outcomes. Similarly, multiplying a utility function with a positive constant does not change any preferences either. This implies that given

a utility function, we are free to apply any positive affine transformation, ie, we are free to choose the origin and to choose a unit length.

B. The Standard Gamble Question

When an agent is maximizing expected utility, a convenient way of reconstructing the utility function is by means of the *standard gamble question*.

A standard gamble question confronts an agent with a binary choice. If the agent chooses action A_1 , he will gain y for sure. If he chooses action A_2 , he has a probability of p to gain $z > y$, but a probability of $1 - p$ to gain $x < y$. The outcomes for A_2 can be abbreviated to $z_p x$ (where the probability is tied to the first state).

For values of p close to zero, it is clear that the agent will prefer action A_1 . At the other end of the spectrum, for p close to one, the agent will prefer A_2 . Therefore, there should be some value p^* for which the agent is indifferent between the two actions [2]. This implies that the utility of gaining y equals the expected utility of $z_p x$. In a standard gamble, an agent is given a choice such as above and has to reveal the value of p^* satisfying this condition, as in equation 2.

$$U(w_0 + y) = (1 - p^*)U(w_0 + x) + p^*U(w_0 + z). \quad (2)$$

In the case that $w_0 = 0$, this expression simplifies to

$$U(y) = (1 - p^*)U(x) + p^*U(z). \quad (3)$$

Since we are free to choose our unit length and origin, we may choose $U(x) = 0$ and $U(z) = 1$. It then follows from equation 3 that $U(y) = p^*$.

Now suppose the agent can reveal his choices for any given values of w_0 , x , y and z . In that case the utility function can easily be reconstructed over a given interval $[A, B]$. To this end, we fix w_0 at $w_0 = 0$, $x = A$ and $z = B$. For each value of $y \in [A, B]$, the agent should now reveal the corresponding p^* that makes him indifferent between the safe option y and the gamble $x_p z$. The utility of any wealth level $y \in [A, B]$ is then equal to the probability p^* .

C. Relative Risk Aversion

For utility functions depending on a single parameter, it is possible to quantify the risk attitude of the agent. When the utility function is the identity function $U(x) = x$, the agent is maximizing the expected value of the resource itself. This is called *risk neutral* behavior, because it implies that the agent will take a bet if and only if it has a positive expected value.

When the utility function is concave, each increment in the resource conveys progressively less utility to the agent. As a consequence, the agent will reject a fair bet, implying *risk averse* behavior. If, on the other hand, the utility function is convex, the agent will take some bets that are less than fair, leading to *risk seeking* behavior.

In order to define those types of behavior quantitatively, we can use *relative risk aversion* or *RRA*, as in equation 4.

$$R_u(x) = \frac{-xu''(x)}{u'(x)} \quad (4)$$

It quantifies the local concavity of the utility function, which in turn characterizes the type of behavior. A class of simple utility functions that is often considered in economics is the class of *isoelastic utility functions*. Those are utility functions with a constant relative risk aversion for all values of x . They have the additional property of being insensitive to scale. That is, when an agent maximizing such a utility function has to make a choice involving possible losses or gains of a resource, multiplying all relevant quantities of the resource by a constant factor will have no influence on his decision.

The general functional form of isoelastic utility functions is shown in equation 5.

$$U_\eta(x) = \begin{cases} \frac{x^{1-\eta}}{1-\eta} & \text{for } \eta \neq 1 \\ \log(x) & \text{for } \eta = 1 \end{cases} \quad (5)$$

Where η , the *coefficient of risk aversion*, equals the relative risk aversion $R_{U_\eta}(x)$ for any value of x .

Within this paper we will restrict attention to the case $\eta < 1$. Values of η greater than one correspond to extreme risk aversion. For such utility functions, $U(0) = -\infty$. This implies that the agent would avoid a possible loss of all wealth at any cost, ie, would even decline a bet where an extremely large amount of wealth can be gained at a vanishingly small risk of losing everything. This type of behavior would not be realistic within our setting.

For values of $\eta < 1$, equation 5 can be simplified. Recall that utility functions can be freely transformed by positive affine transformations. Since $1 - \eta$ is a positive constant for $\eta < 1$, the utility function can be multiplied by this number.

In order to simplify the equation further, we use $\rho = 1 - \eta$ as our measure of risk aversion, rather than η itself. Restricting to the case $\eta < 1$, multiplying by $1 - \eta$ and substituting η for $1 - \rho$ produces the Cobb-Douglas class of utility functions, see equation 6.

$$U(x) = x^\rho, \rho \in \mathbb{R}^+ \quad (6)$$

Where ρ is a risk aversion parameter. When $\rho = 1$, the agent is risk neutral. For $\rho < 1$ the agent is risk averse, while for $\rho > 1$ the agent is risk loving. Cobb-Douglas utility functions represent a convenient behavioral model for the scope of this research, since it can model both risk averse and risk loving behavior. Another advantage is that the only parameter that is allowed to vary between agents, risk aversion, is exactly the quantity we are trying to estimate.

III. THE MULTI-AGENT MODEL

We investigate the influence of an exogenously given fixed maximum lifetime on the evolution of risk attitude. To this end, we consider a multi-agent system loosely based on the

Sugarscape model [4], although it differs in some details. It features agents consuming a single resource, called *sugar*. The agents are in constant need of sugar in order to survive. However, sugar is a scarce good in this system, because it grows in limited quantities. Thereby the system supports a limited population of agents, where they have to compete for survival.

A. Birth, Death and Consumption

The environment of the agents consists of a square lattice of $N \times N$ locations. Each location can be occupied by zero or more agents. The world is connected at the edges, giving it the topology of a torus. That is, when an agent walks off one edge it reappears at the opposite edge. A location can be in two different states. Either it is empty, or contains sugar. When a location is empty, it has a probability α to 'grow' new sugar in the next timestep and a probability $1 - \alpha$ to remain empty. A location containing sugar will remain in that state until an agent harvests the sugar.

Agents can move around the world and harvest sugar when they find it. Once harvested, the sugar is stored in the stock of the agent. The current amount of sugar held by agent i is denoted as s_i . Agents can keep an unlimited amount of sugar in stock for indefinite time, but need to consume one unit of sugar from stock at each timestep. When the agents has no sugar in stock left, it dies from starvation.

Initially, all locations are empty. An initial number of N_0 agents are created at random positions. Multiple agents can share the same location. All agents start with a stock containing $\gamma/2$ sugar.

During each timestep, all agents carry out a number of action rules. Each of those action rules is carried out simultaneously by all agents. In cases where this would lead to conflicts, the agents get to act in a random order.

First, all agents move a single step into a random direction. The direction can either be horizontally, vertically or diagonally.

Next, all agents get to harvest sugar. If the new location contains sugar, the agent picks it up and adds it to its own stock. As soon as any agent harvests the sugar, the location reverts to the empty state. As a result, when multiple agents share the same location with sugar, only the agent acting first gets to harvest it.

The agent has to choose one of two possible methods of harvesting. It can either choose a safe method, which has a guaranteed but low yield of y units of sugar. Alternatively, it can choose a risky method, which could either yield a higher amount of z sugar, or it could fail and yield a lower amount of x sugar. The resulting amount of sugar is added to the stock s_i .

Since harvesting will never yield more than z sugar, the total available amount of sugar in the system is kept limited. Once an agent has harvested sugar, all sugar at the location disappears, regardless of the amount of sugar the agent has managed to extract.

After harvesting comes the reproduction rule. If an agent has gathered at least γ units of sugar, it will reproduce. A

new child agent is created at the same location. The parent loses half of the sugar it has in stock and transfers it to the stock of the child. After reproduction, s_{child} equals s_{parent} .

The next action rule is metabolism. All agents digest a quantity of δ sugar, which is subtracted from s_i .

The final action rule in each timestep is death. If $s_i \leq 0$, agent i dies and is removed. Additionally, agents have a maximal age \mathcal{A}_{max} . If the agent has been alive for this many timesteps, it also dies during this phase.

All of the action rules of the system and their related system parameters are summarized in table I. The table also lists symbolic representations for each of the action rules. The total rule set of the system is $\{\{G_\alpha\}, \{M, H, R_\gamma, E_\delta, D_{\mathcal{A}_{max}}\}\}$, where the world action rules are $\{G_\alpha\}$ and the agent action rules are $\{M, H, R_\gamma, E_\delta, D_{\mathcal{A}_{max}}\}$.

TABLE I
ACTION RULES OF THE MULTI-AGENT SYSTEM

Order	Description	Symbol	Parameters
	Sugar Growback	G_α	α
1	Movement	M	-
2	Harvest Sugar	H	-
3	Reproduction	R_γ	γ
4	Metabolism	E_δ	δ
5	Death	$D_{\mathcal{A}_{max}}$	\mathcal{A}_{max}

B. Evolution of Decision Rules

Agents in our model get a choice between two methods of harvesting sugar. The choice takes the form of the standard gamble question in equation ???. There are three possible resulting sugar yields from harvesting: x , y and z , with $x < y < z$. If the agent chooses the safe method of harvesting, he gains y sugar for sure. When the agent chooses the risky method, he has gains z sugar with probability p and x sugar otherwise.

The quantities x , y and z and the probability p are revealed to the agent beforehand. The agent now calculates the expected utility of both options. Utility is taken as a function of the total amount of sugar that the agent will have in stock after adding x , y or z . The agents have a Cobb-Douglas utility function, as specified in equation 6. The risk aversion parameter ρ differs between agents and is genetically encoded. The resulting calculations for agent i are given in equations 7 and 8.

$$EU_{safe}(s_i, \rho_i) = (s_i + y)^{\rho_i} \quad (7)$$

$$EU_{risk}(s_i, \rho_i) = p(s_i + x)^{\rho_i} + (1 - p)(s_i + z)^{\rho_i} \quad (8)$$

If $EU_{safe} > EU_{risk}$, the agent chooses the safe harvest method and if $EU_{risk} > EU_{safe}$, he chooses the risky method. In the unlikely event that $EU_{safe} = EU_{risk}$, the agent is indifferent and chooses either option with 50% probability.

In the encoding of ρ_i there is no distinction between genotype and phenotype. Each agent i has a genetic code consisting of the single allele ρ_i . At reproduction, the genetic code of the child agent is inherited from the parent agent and subjected to a mutation operator known as *lognormal mutation*, see equation 9.

$$\rho_{child} = \rho_{parent} \cdot e^\zeta \quad (9)$$

$$\zeta \sim N(0, \sigma) \quad (10)$$

Where σ in equation 10 is a constant mutation parameter of the model.

C. Parameters of the Harvest Process

Recall that the harvest decision takes the form of a standard gamble question. As explained in section II-B, the shape of a utility function can be recovered over an interval by keeping $w_0 = s_i$, x and z fixed and letting y increase from x to z .

In order to simulate the evolution of a utility function that works well over the entire domain of interest, we adopt a similar approach. Unfortunately, s_i depends on the history of the agent and can't be kept fixed. Nonetheless, in the multi-agent system we keep the amounts x and z in the harvest gamble fixed, while varying y . The safe yield y is drawn randomly with uniform distribution from $\langle x, z \rangle$, implying $x < y < z$. This ensures that the utility function is evaluated at many different points.

The probability parameter p is also varied. There is a problem with the choice of p that merits careful consideration. If p were simply a function of y , for example, it would be possible that expected utility of one of the two methods of harvesting would be higher than the other for the entire population for any value of y . In this scenario, all agents would prefer the same method of harvesting. In that case, no agent would have any selective advantage over any of the others. Therefore, there would be no selective pressure on risk aversion at all.

In order to avoid such a situation, p is manipulated in a manner designed to optimize selective pressure. The idea is that when the agents start to evolve risk loving strategies, the probability of winning the high amount z is proportionally decreased to test whether the agents are willing to take even more risk. Conversely, when agents become risk averse, the probability of winning is increased to test if the agents will now evolve even more risk aversion.

We gauge the dominating risk attitude of the agents in the current population by calculating the average value of ρ_i at the beginning of each timestep, as in equation 11.

$$\bar{\rho}_t = \frac{1}{n} \sum_{i=1}^n \rho_i \quad (11)$$

Now we choose p such that an agent with a risk aversion parameter $\rho_i > \bar{\rho}_t$ will always choose the risky harvest method, while an agent with $\rho_i < \bar{\rho}_t$ will choose the safe option. This creates a high selective pressure. For example,

consider the case where the agents could on average increase their fitness significantly by becoming more risk averse. It is then ensured that a part of the population will choose the safe method of harvesting, giving them a significant selective advantage over the rest.

This condition can be restated as follows. When the risk aversion gene of an agent would be equal to the population average, he should be indifferent between the safe harvest method and the risky harvest method. This is formally specified in equation 12

$$EU_{safe}(s_i, \bar{\rho}_t) = EU_{risk}(s_i, \bar{\rho}_t) \quad (12)$$

Solving equation 12 yields the following function for p :

$$p(s_i, y, \bar{\rho}_t) = \frac{(s_i + y)^{\bar{\rho}_t} - (s_i + x)^{\bar{\rho}_t}}{(s_i + z)^{\bar{\rho}_t} - (s_i + x)^{\bar{\rho}_t}} \quad (13)$$

In summary, x and z are kept constant, y is drawn uniform randomly from $\langle x, z \rangle$ and p is calculated from s_i , y and $\bar{\rho}_t$ as in equation 13.

IV. EXPERIMENTAL SETUP

We performed a series of experiments with the multi-agent system of section III. In each experiment, the model was run a total of 100 times. The simulations were all aborted after $T = 20000$ timesteps. At the end of each run, we calculated and stored the average risk aversion in the population at the beginning of the final timestep $\bar{\rho}_T$. Those averages were averaged over the simulation runs.

The entire experiment of 100 system runs has been repeated for several parameter settings. We explored the behavior of the system by systematically varying some parameters of interest and keeping all other parameters equal.

The first parameter that was varied was the maximum lifetime \mathcal{A}_{max} , since the main question is whether this has any influence on the result. The maximum lifetime was varied between the values of 40, 80, 160, 320 and ∞ . An infinite lifetime means agents do not have a maximum lifetime at all and is equivalent to $\mathcal{A}_{max} = T$.

Another parameter which is likely to play a major role in the evolution of behavior under risk is the sugar threshold for reproduction γ . The reason why we believe so, is that winning a gamble while harvesting could push the available amount of sugar beyond this threshold. In that case the agent would be able to reproduce immediately, which could provide a strong incentive to risk the gamble.

In order to gain insight in the influence of γ , it was also varied between experiments. It was consecutively held at the values of 40, 80 and 160. Every combination of these two parameters has been tried, resulting in a total of 5×3 experiments of 100 runs each.

The stepsize parameter σ for evolution and the run length T were manually tuned for good performance. There are two criteria that we tried to meet in tuning those parameters. First of all, σ should not be too large, in order to keep the effect of genetic drift limited. Secondly, T should be sufficiently large for the evolution to reach a long-term equilibrium. We

found that the values of $\sigma = 0.01$ and $T = 20000$ timesteps met both criteria reasonably well.

A detailed list of all the parameter settings that have been used in the experiments can be found in table II.

TABLE II
EXPERIMENT DESCRIPTION TABLE FOR THE TESTS

Experiment Details	
Simulation Length T	20000 Timesteps
Number of Repeats	100
World Size	50×50 Locations
Sugar Growth (α)	0.01
Harvest Choice Details	
x	0
z	100
y	random uniform over $\langle x, z \rangle$
p	$p(s_i, y, \bar{\rho}_t)$
Agent Details	
Reproduction Threshold γ	40, 80, 160
Metabolism δ	1
Maximal Age \mathcal{A}_{max}	40, 80, 160, 320, ∞
Initial Agents	100
Initial Sugar $s_{i,0}$	$\gamma/2$
Initial Agent Location	Uniform Random
Evolution Details	
Genetic Encoding	Direct Representation
Mutation Operator	Lognormal Mutation
Mutation Parameter σ	0.01
Initial σ	1
Crossover Operator	none

V. EXPERIMENTAL RESULTS

Table III shows the experimental results, averaged over all 100 repetitions of the experiment. It also shows the standard deviations to indicate the spread in the results.

In one case, where $\mathcal{A}_{max} = 40$ and $\gamma = 160$, we did not get a result at all. The reason is that the system was no longer stable under this parameter setting, resulting in frequent population crashes. The most likely explanation is that agents were living too short to gather enough sugar to reach their reproduction threshold, thus failing to reproduce.

The standard deviations in $\bar{\rho}_T$ give an indication of the stability of the evolutionary process and the effect of genetic drift at $t = T$. It turns out that the signal-to-noise ratio is good enough to draw meaningful conclusions about the behavior of the system.

The results clearly show that γ plays an important role in the evolution of risk aversion in this multi-agent system. When the agents need a larger amount γ of sugar to be able to reproduce, they tend to develop more risk-loving strategies.

For each of the three values of γ , the average value of $\bar{\rho}_T$ over the runs decreases strictly with \mathcal{A}_{max} . Not all of the differences between those average values are statistically significant. However, it can be shown that for these values of γ the overall negative correlation of $\bar{\rho}_T$ with \mathcal{A}_{max} is statistically significant at high confidence levels.

For $\gamma = 40$, a T-test indicates that the average value of $\bar{\rho}_T$ is significantly greater for $\mathcal{A}_{max} = 40$ than for $\mathcal{A}_{max} = 80$ at the 99% confidence level. The latter value is in turn significantly greater than for $\mathcal{A}_{max} = \infty$, again at the 99% confidence level.

For $\gamma = 80$ we can say exactly the same: $\bar{\rho}_{T, \mathcal{A}_{max}=40} > \bar{\rho}_{T, \mathcal{A}_{max}=80} > \bar{\rho}_{T, \mathcal{A}_{max}=\infty}$. These inequalities are established with statistical significance at the 99% confidence level.

For $\gamma = 160$, it likewise holds that $\bar{\rho}_{T, \mathcal{A}_{max}=80} > \bar{\rho}_{T, \mathcal{A}_{max}=160} > \bar{\rho}_{T, \mathcal{A}_{max}=\infty}$ with statistical significance at the 99% confidence level.

Summarizing, we have no case where the result increases with \mathcal{A}_{max} , while the average value of $\bar{\rho}_T$ over the test runs decreases with \mathcal{A}_{max} for each of the three values of γ in our test suite.

These results are fully in line with our hypothesis that agents with a longer maximum lifetime are likely to evolve more risk aversion.

TABLE III

EVOLVED RISK PARAMETER $\bar{\rho}_T$, AVERAGED OVER 100 RUNS EACH.

\mathcal{A}_{max}	$\gamma = 40$	$\gamma = 80$	$\gamma = 160$
40	1.34 ± 0.10	1.87 ± 0.13	–
80	1.22 ± 0.084	1.76 ± 0.10	2.73 ± 0.15
160	1.21 ± 0.085	1.60 ± 0.088	2.52 ± 0.097
320	1.20 ± 0.086	1.56 ± 0.80	2.10 ± 0.089
∞	1.19 ± 0.081	1.54 ± 0.070	1.84 ± 0.076

VI. CONCLUSIONS

We investigated the effect of an exogenously given life expectancy on the evolution of risk attitude in a multi-agent system. The multi-agent system was a variant of the sugarscape world [4]. It models a species consuming and depending on a single resource, sugar, living in a niche of limited carrying capacity.

We introduced risk into the model, taking the form of a standard gamble question. Agents were offered a choice between two methods of harvesting sugar, a safe and a risky method. The safe method offered a fixed reward, while the risky method had a probability of yielding a higher amount, but also a risk of yielding no sugar. The agents decided by maximizing their expected utility. The functional form of the utility curve was fixed up to a single degree of freedom. The one parameter that was allowed to vary between agents corresponds to relative risk aversion [2].

The risk aversion parameter was treated as the genetic code of the agent. By observing many evolutions of the gene, we were able to measure the preferred degree of risk aversion of the agents under given parameter settings of the model. The experiments have been repeated for various parameter setting. We found that long-lived agents consistently and clearly evolve a more risk averse type of behavior than short-lived agents, for all of the parameter settings in our test suite.

This finding is consistent with earlier work [12], lending independent support to the hypothesis that a longer maximum lifetime will encourage the evolution of more risk averse strategies for survival. In contrast to this earlier study, we established that there is an effect for finite differences in maximum lifetime and under a range of relevant parameter settings of the system. Additionally, we improved on the earlier results by quantifying risk aversion in the form of a coefficient of relative risk aversion.

It should be noted, however, that in both this research and in [12] the agents were allowed to reproduce during their entire lifecycle. The effect of a limited fertile period remains to be investigated. Other features that could be explored in future work include the effects of sexual reproduction, a crossover operator and multiple resources. In the context of human behavior one would also need to consider trade and social factors such as altruism.

There are several real-world cases where our hypothesis could be tested or used as an explanatory factor. As one example, in most animal species, including humans, females have a longer average life expectancy than males, while males are often more inclined to engage in risky activities [13] [9]. More generally, in relation to animal species it is consistent with *life history theory* [3] [11]. Life history is a mathematical framework relating patterns of behavior of animals to effects of natural selection on the key characteristics that define the life course, such as maturation, reproduction and life expectation.

In the field of economics, an evolutionary influence of life expectation could serve as a partial explanation for some findings of *prospect theory*. Prospect theory has widely been accepted as a better descriptive model for human behavior under risk than expected utility theory [7], [1]. An important feature of prospect theory is *framing*: people tend toward risk loving behavior when facing losses, while being much more risk averse when facing possible gains. Of particular interest in relation to our research, it has been found that people are stating to have especially risk-loving tendencies in matters of life and death [5].

Earlier attempts to explain framing from a perspective of biological evolution include [10]. In this research, a model from risk-sensitive *optimal foraging theory* was developed. It was argued that framing has evolved as an adaptation to optimize fitness in a hunter-gatherer society.

We propose the evolutionary influence of life expectancy on risk attitude as an additional explanatory factor for framing effects, especially for situations where life and death are at stake.

ACKNOWLEDGMENT

The authors would like to thank Cees Withagen for reviewing an early draft of the paper.

REFERENCES

- [1] C. Athena Aktipis and Robert O. Kurzban. Is homo economicus extinct?: Vernon smith, daniel Kahneman and the evolutionary perspective. *Advances in Austrian Economics*, 7(1):135–153, 2004.

- [2] K.J. Arrow. *Essays in the Theory of Risk-Bearing*. Markham, 1971.
- [3] E.L. Charnov. *Life History Invariants*. Oxford University Press, 1993.
- [4] J.M. Epstein and R. Axtell. *Growing Artificial Societies: social sciences from the bottom up*. Brookings Institution Press, 1996.
- [5] N.S. Fagley and P.M. Miller. Framing effects and arenas of choice: Your money or your life? *Organizational behavior and human decision processes*, 71(3):355–373, 1997.
- [6] N. Gilbert and K. Troitzsch. *Simulation for the social scientist*. Open University Press, 1999.
- [7] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 4(2):263–291, 1979.
- [8] R.L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Academic Press, 1976.
- [9] D.J. Kruger. An evolutionary life-history framework for understanding sex differences in human mortality rates. *Human Nature*, 17(1):74–97, 2006.
- [10] J.H. Fowler R. McDermott and O. Smirnov. On the evolutionary origin of prospect theory preferences. *International Studies Quarterly*, 70(2):335–350, 2008.
- [11] D. Roff. *The Evolution of Life Histories: Theory and Analysis*. New York: Chapman and Hall, 1992.
- [12] V. van der Goes and A.E. Eiben. The effect of life expectancy on risk aversion. In *Proceedings of the 21st Benelux Conference on Artificial Intelligence*, 2009.
- [13] M. Wilson and M. Daly. Competitiveness, risk taking, and violence: The young male syndrome. *Ethology and Sociobiology*, 6(1):59–73, 1985.