The Effect of Life Expectancy on Risk Aversion

Vincent van der Goes A.E. Eiben

Vrije Universiteit, Amsterdam
vgoes@feweb.vu.nl, gusz@cs.vu.nl

Abstract

In this paper we investigate the effect of life expectancy on risk attitude in an artificial multi-agent world, which is a variant of the sugarscape [2]. Our basic approach is evolutionary in that we encode risk attitude as a gene within the agents that undergoes variation and selection and perform simulations to see the evolved (i.e., presumed optimal) values of this gene/attitude under different conditions. We observe that agents that have much to loose (i.e., their "eternal life") evolve a more risk averse type of behavior than agents with a limited life span.

1 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.

The classical economic *expected utility* model [1], that will be used in this paper, is based on the assumption of rational behavior. Modern, more general descriptive models allow for a certain degree of irrationality, based on empirical evidence [4]. Computational techniques such as multi-agent systems open up new avenues for research on behavior under uncertainty and risk. They allow us to integrate insights from multiple disciplines, such as economics, psychology and even biology.[8]

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.

To answer this question we set up a simple system of agents with the basic properties of a species: they breed and depend on the consumption of a single resource, sugar. It is a variant of the *sugarscape* system in [2]. Next, our goal is to find out what kind of behavior would be optimal in this environment when an agent would be offered a gamble, with chances to lose or gain sugar. If our hypothesis is correct, this optimal behavior would likely vary significantly when the life duration of the agents is modified.

Since expected utility is generally accepted as the normative model of behavior under risk [5], we expect optimal behavior to take this form. Under expected utility, the behavior under risk of an agent is completely defined by its utility function. In the case of our model, utility is a function of the total amount of sugar that the agent owns. We measure the utility function by introducing a choice involving risk into our model. There will be the possibility of gaining more resource, but also a risk of death. The agents decide in correspondence to an inherited genetic code. The resulting evolution will be reiterated many times. A comparison will be made between agents with no age limit and inherently short-lived agents.

Since this evolution takes place in a strongly simplified environment and the agents have complete information about possible outcomes and the probabilities of those outcomes, we assume that evolution will result in an equilibrium that is at least close to rationality. That is, we assume that

Assumption 1 *The expected utility theory is a suitable descriptive model for the average evolved behavior for the agents at the end of each evolution run of our multi-agent system.*

The paper is set up as follows. In section 2 we introduce the economical theory that will be needed. Section 3 specifies the multi-agent system. This is a variant of the sugarscape model from [2]. The details of the experiment are given in section 4.1 and the results are summarized and analyzed in section 4.2. Finally, section 5 gives a summary and conclusions.

2 Decision Under Risk in Economics

2.1 Expected Utility Maximization

In economics, increasingly complex models have been developed for human *behavior under risk* and *uncertainty*. These are two different contexts. Under risk, all the possible outcomes and the probabilities of each are perfectly known, while this is not necessarily the case for decision under uncertainty. For the scope of this project, we will only consider decision under risk. Accordingly, agents will know all the probabilities that are involved.

Modern descriptive models can be used to model irrational behavior (see, e.g., [4]). The most flexible model that is still considered to be fully rational is the model of expected utility [1]. This model is built on the concept of utility. Informally speaking, utility is a measure of the happiness of an agent under given conditions. Utility reveals itself in and is defined by preferences. If an agent prefers state A over state B, denoted as $A \succ B$, his utility level in state A is by definition greater than in state B. So if we measure all preference relationships of an agent for all possible states X of the world, we can construct a corresponding utility function $U(X)$. Assuming that the agent has complete and consistent preferences, that is.

However, such a utility function is not unique. First, adding a constant to it would not change any of the preferences. Secondly, multiplying utility with a positive constant does not change preferences either. So when defining a utility function, we are free to choose the origin and also free to choose a unit length.

Suppose an agent has to choose one action out of n possible actions A_i. Action A_i has m_i possible outcomes X_{ij} $(i = 1, 2, ..., n; j = 1, 2, ..., m_i)$. Each outcome has a known probability p_{ij} of occuring after action A_i. The expected utility model predicts that the agent will choose the action that maximizes expected utility. Expected utility is the weighted average of the utility values of the possible outcomes, as defined in equation 1.

$$EU(A_i) = \sum_{j=1}^{m_i} p_{ij} U(X_{ij}) \tag{1}$$

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, we have a complete knowledge of its behavior under risk.

2.2 The Standard Gamble Question

When we assume an agent to act in accordance to the theory of expected utility, we can measure the relative utility assigned to states with the standard gamble question. The standard gamble question confronts an agent with a hypothetical binary choice. If he chooses action A_1, he will enter state Y for sure. If he chooses action A_2, he has a probability of p to enter a better state Z, but also a probability of $1 - p$ to enter a worse state X. So for these states it is known that $Z \succ Y \succ X$. The outcomes for A_2 can be abbreviated to $Z_p X$ (the probability is tied to the first state here).

In this scenario there should be some value of p for which the expected utility of A_1 equals the expected utility of A_2. We then say the agent is indifferent between A_1 and A_2, or $A_1 \sim A_2$. The agent has to answer the question which p satisfies this condition, as depicted in equation 2.

$$Y \sim \begin{cases} Z & \text{with chance } p \\ X & \text{with chance } 1-p \end{cases} \tag{2}$$

From equations 1 and 2 it now follows that

$$U(Y) = (1-p)U(X) + pU(Z). \tag{3}$$

So the standard gamble question can be used to measure the relative utility of states X, Y and Z.

2.3 Utility of a Single Resource

Consider utility as a function of only a single resource, such as monetary wealth. A state then corresponds to owning a specific amount of money. Let x, y and z be the amount of money owned in states X, Y and Z respectively. If an agent can answer all of the standard gamble questions for a fixed highest amount $z = z_1$, a lowest amount $x = x_0$ and all amounts $y \in [x_0, z_1]$, it is easy to construct the utility function of monetary wealth over the domain $[x_0, z_1]$. Since we are free to choose the origin and unit length of the utility function, we may choose $U(x_0)$ to be 0 and $U(z)$ to be 1. It then follows from equation 3 that $U(y), y \in [x_0, z_1]$ is equal to the corresponding outcome of the standard gamble question, p_y.

One special case would be an agent with utility function $U(y) = y$. An agent with this utility function will simply maximize the expected value of money. This is called *risk neutral* behavior. Another special case is when the utility function has the property that $U(y) < y$. In that case the agent will be indifferent between the gamble $z_p x$ and the safe amount y when the expected value of the gamble is smaller than y. Such an agent is called *risk seeking*. Similarly, an agent with a utility function $U(y) > y$ is *risk averse*.

3 The Sugarscape System

3.1 Basic Model

We will consider a multi-agent system that is similar to the Sugarscape model from [2]. It is a multi-agent system that implements agents consuming and depending on a single resource, called sugar. The agents move around to gather the sugar and can store it in their stock. Sugar in stock does not decay. The agents can reproduce. When an agent runs out of sugar, it dies from starvation. Our aim is to keep the model as simple as possible.

The world consists of a square lattice of locations. The world is also connected at the edges, giving it the topology of a torus. A location can either be empty, or it can contain a fixed quantity of β sugar. When a location is empty, it has a probability α to 'grow' a new quantity of β sugar in the next timestep and a probability $1 - \alpha$ to remain empty. We call this growback rule $G_{\alpha,\beta}$.

Initially, all locations are filled with β sugar. The world will also contain a number of agents, which are placed at random positions. Multiple agents can occupy the same location. All agents start with a stock of γ sugar.

At each timestep, the agents carry out a number of action rules. First, all agents move a single step into a random direction, either horizontally, vertically or diagonally. The order in which they do this is random. This is movement rule M.

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. The order in which they act is again randomized, so when there are multiple agents at the same location, one of them gets all the sugar and all the others get nothing. This is harvest rule H.

Now the agents may reproduce. If an agent has gathered at least 2γ units of sugar, a new agent is created at the same location. This costs the parent agent γ sugar. The child agent is endowed with γ sugar. This is reproduction rule R_γ.

Finally, the agents will need to digest sugar to stay alive. This action is defined by metabolism rule E: the stock of the agent is reduced by one unit of sugar.

The order in which the agents act is randomized for each action. The complete rule set is
$(\{G_{\alpha,\beta}\}, \{M, H, R_\gamma, E\})$.

Optionally, the agents may also have a maximal age. In that case, the agent dies when it has existed for that number of timesteps.

3.2 Introducing Risk

The agents in our model will get the opportunity to engage in a risky activity, called *the lottery*. The lottery will be offered randomly and only on rare occasions, to prevent a significant effect on the economy. The lottery takes the form of the standard gamble question in equation 2. There are three possible outcomes: X, Y and Z. In state X, the agent loses all sugar and dies. In state Y the stock of the agent is reset to a fixed amount of y units of sugar, no matter how much sugar the agent had before. Likewise, in state Z the stock is reset to z units, with $z \geq y \geq x = 0$. The probability of winning the high amount z in the lottery will vary over time.

The agent can either choose to enter state Y for sure, or take the gamble $Z_{p_t} X$.

Now our goal is to find the chance p^* for which agents are indifferent between the two available options. More precisely, we want to know the indifference value p^* leading to optimal agent behavior, optimal in the sense of providing the best selective advantage.

The idea is to measure p^* through a coevolution of the agents and the lottery. The agents have a new genotype, defined in equation 4. The genotype of agent i $(i = 1..n)$, $\mathcal{G}_i$, contains a single real-valued allele π_i.

$$\mathcal{G}_i = \pi_i \tag{4}$$

π_i is the answer of agent i to the standard gamble question with states X, Y and Z. The agent will prefer the gamble Z_pX over Y if and only if $p_t > \pi_i$.

To ensure a good selective pressure on the genotype, we adopt the following update rule U for p_t:

$$p_t = \frac{1}{n}\sum_{i=1}^{n}\pi_i \tag{5}$$

So why would the values of π_i converge toward p^*? Suppose that at timestep t the average value of π_i is below p^*. Update rule U will then set p_t to the average of those values. If an agent with a value of π_i above the average is offered the lottery, it will refuse the gamble and take the safe amount y. Since p_t is below p^*, this is a good decision. An agent with a value of π_i below the average, on the other hand, will take the gamble $x_{p_t}z$. Therefore, the former agents have a selective advantage over the latter agents. Over time, the average value of π_i is likely to rise.

Similarly, when the average value of π_i is above p^*, agents with a value of π_i have a selective advantage and the average value of π_i is likely to fall.

However, when the average value of π_i is close to or equal to p^*, it is likely to diverge because of *genetic drift* [9] [3]. So even though π_i is on average likely to converge toward p^*, there is also noise.

Each turn, every agent will take the lottery with chance ϵ. This is a new action rule for the agents. We call it lottery rule $L_{\epsilon,x,y,z}$. The lottery rule is applied after the harvest rule H, but before the reproduction rule R_γ.

Evolution of the gene π works only through mutation, there is no crossover operator. Inheritance of the gene π is defined as follows:

$$\pi_{child} = \max(\min(\pi_{parent} + N(0, \sigma), 1), 0) \tag{6}$$

The stepsize σ in equation 6 becomes a new parameter of the reproduction rule, which now becomes $R_{\gamma,\sigma}$. The new rule set of the sugarscape system is $(\{\, G_{\alpha,\beta}\}, \{\, M,\ H,\ L_{\epsilon,x,y,z},\ R_\gamma, C\}, \{U\})$.

4 Experiments

4.1 Experimental Settings

In order to measure the utility curve and the associated risk attitude of the agents, the multi-agent system described in section 3 has been run multiple times, varying the sure amount of sugar y between a lowest amount of $x_0 = 0$ to a highest amount of $z_1 = 100$ units of sugar, taking incremental steps of 5 units. The amount of sugar that could be won in the gamble, z_1, was held fixed at an amount of 100 units. The punishment for losing the gamble also remained fixed at losing all sugar and subsequent death, as x_0 was held fixed at 0 units. As explained in section 2.3, varying y this way allows us to measure the function $U(y)$ in the domain $[x, z]$. The upper bound z was chosen to be 100 because an agent in the basic system from section 3.1 can not have more than $2\gamma + \beta = 90$ sugar in stock (the agent would immediately reproduce before harvesting any new sugar).

To get a good approximation of p^*, we ran the multi-agent system for a period of 20000 timesteps and then collected the average value of π at the end of the run. The run was repeated 100 times for each value of y.

The experiment decribed so far was first conducted for short-lived agents, with a fixed maximal lifetime of 40 timesteps. Next, it was repeated for long-lived agents with no age limit.

The other parameters of the multi-agent system were kept equal between all runs. They are summarized in table 1. When initializing an evolution run, the world was completely filled with sugar. A population of 100 agents was created on random locations, with random genes and a stock containing $\gamma = 40$ units of

parameter	**value**
world size	32
α in $G_{\alpha,\beta}$	0.03
β in $G_{\alpha,\beta}$	10
γ in $R_{\gamma,\sigma}$	40
σ in $R_{\gamma,\sigma}$	0.02
ϵ in $L_{\epsilon,x,y,z}$	0.01
x in $L_{\epsilon,x,y,z}$	0
y in $L_{\epsilon,x,y,z}$	$0, 5, .., 100$
z in $L_{\epsilon,x,y,z}$	100
maximal age	40,∞
time limit	20000 time steps
initial agents	100
initial sugar of agents	$\gamma = 40$
initial π	uniform random from $[0, 1]$
initial agent location	uniform random

Table 1: Parameter settings of the multi-agent system used in the experiment.

sugar. The maximal age of the short-lived agents was chosen to be 40, because a much shorter life span could lead to population crashes. The mutation strength σ was manually tuned to stabilize the evolution of π well within the time limit.

4.2 Results and Discussion

Figure 1 and figure 2 show the experimental results, pointwise averaged over all 100 repetitions of the experiment for each separate amount of sugar y. It also shows the standard deviations to indicate the spread in the results. The averages of both experiments have been plotted together in figure 3, with the addition of the risk-neutral curve as a visual aid.

From the figures it appears that the long-lived agents evolved a more risk-averse strategy than the short-lived agents. The average evolved values of π (averaged over the populations and then averaged over all runs) are higher for all values of y.

In order to test if this is a statistically significant effect, we compared the results point by point with a two-tailed Student's T-test for each value of y. Each of those tests yields a p-value smaller than 0.00001, except at $y = 80$, where the p-value is 0.0002. This is extremely significant evidence of difference for the entire length of the curve. We conclude that even within a very simple multi-agent system, life expectancy can have a notable influence on the evolution of risk attitude.

The spread in the results tells us something about the validity of assumption 1. In the case of the long-lived agents we see a very small spread. However, the short-lived agents show a greater spread in evolved behaviour. This can be interpreted as a greater tolerance for suboptimal genotypes in the case of short-lived agents, even though it is not clear what causes this difference.

A notable feature is that both utility curves show a bend around $y = 80$. This may be related to the fact that an agent in our experimental settings requires $2\gamma = 80$ sugar in stock to reproduce. The immediate opportunity to produce offspring could explain why acquiring exactly this amount seems to be worth taking additional risk.

One peculiarity in the utility curve is that it starts above 0 and ends below 1. On theoretical grounds it is obvious that p^* equals 0 for $y = 0$ and 1 for $y = 100$. However, at these extremes there is a problem with the evolution. Mutation rule 6 allows π to become greater than 0, but it does not allow a mutation below 0. Therefore, when $y = 0$ there will be a bias toward evolution of π values greater than 0, in violation of assumption 1. Likewise, the assumption is also violated at $y = 100$.

Still, as figure 3 shows, the endpoints of the curve are reasonably close to their theoretical values in spite of these violations.

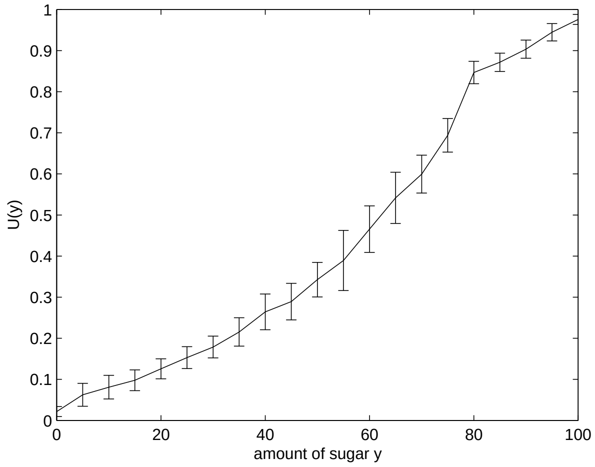

Figure 1: Experimental results for short-lived agents (maximal age 40 generations) with standard deviations.

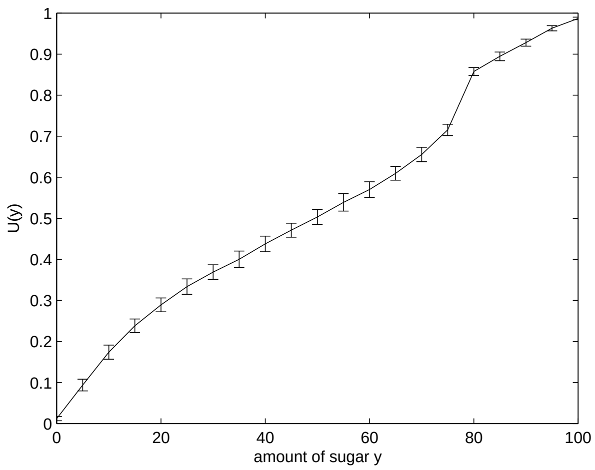

Figure 2: Experimental results for long-lived agents (no age limit) with standard deviations.

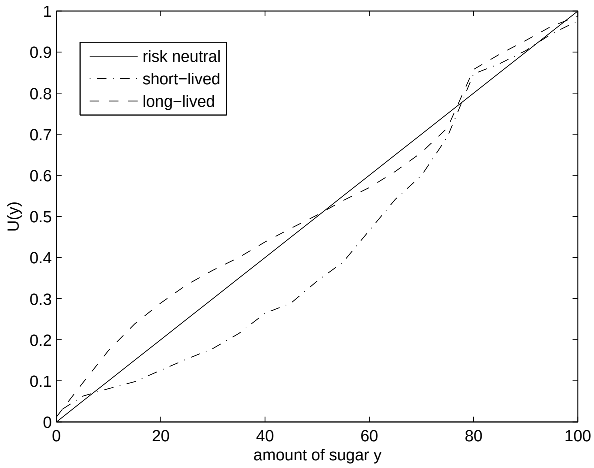

Figure 3: Utility as a function of sugar wealth. The risk neutral curve has been added as a visual aid.

5 Conclusions and outlook

We investigated the effect of life expectancy on the evolution of risk attitude in a multi-agent system. The multi-agent system was a variant of the sugarscape world [2]. It models a species consuming and depending on a single resource.

We introduced a lottery into the model, taking the form of a standard gamble question. The agents had to choose between a smaller safe amount, or a higher amount with a risk of death. The behavior of the agents was encoded with a gene. By observing many evolutions of the gene and under the assumption of evolution leading to optimal behavior which is accurately described by the theory of expected utility maximization, we were able to measure the optimal utility function of agents in our model. Since we assume expected utility maximization, knowing the utility function over a given domain allows us to predict the optimal choice given any decision under risk as long as the outcomes fall within the domain.

Introducing these kinds of lotteries into a multi-agent system to measure utility is a novel technique, with the potential for broader application. For example, in [2], page 98 a utility or welfare function is needed to derive trade rules for agents. The utility function used in that case was postulated in an ad-hoc fashion, while the standard gamble technique used here could be used to derive one.

The experiments have been repeated twice, once for agents with no upper limit on age and once for agents with a shortly limited lifespan. The observed utility functions clearly show a more risk-averse attitude in the first group. This is in agreement with our hypothesis that long-lived agents are likely to exhibit more risk aversion than short-lived agents.

However, it is not yet clear how certain parameters of the algorithm could have influenced our results. In particular, the effect of the probability ϵ an agent will be subjected to a lottery deserves further investigation, as well as the effects of varying the mutation strength σ. In addition, the possibility of adding recombination could be considered.

There are many possible applications of our hypothesis. 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 [7] [6]. However, it is not easy to separate cause and effect when looking for empirical evidence of our hypothesis, since taking more risk will obviously affect life expectancy in turn.

6 Acknowledgements

Many thanks to Cees Withagen for reviewing and providing directions that have lead to this research. We'd like to thank Peter Wakker for reviewing an early version of this document. Further, we'd like to thank Maarten Oosten and Jelle Goeman for their helpful suggestions.

References

[1] K.J. Arrow. *Essays in the Theory of Risk-Bearing*. Markham, 1971.

[2] J.M. Epstein and R. Axtell. *Growing Artificial Societies: social sciences from the bottom up*. Brookings Institution Press, 1996.

[3] K.A. De Jong. *An Analysis of the Behaviour of a Class of Genetic Algorithms, PhD thesis*. University of Michigan, 1975.

[4] D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, 4(2):263–291, 1979.

[5] R.L. Keeney and H. Raiffa. *Decisions with Multiple Ojectives: Preferences and Value Tradeoffs*. Academic Press, 1976.

[6] D.J. Kruger. An evolutionary life-history framework for understanding sex differences in human mortality rates. *Human Nature*, 17(1):74–97, 2006.

[7] M Daly M. Wilson. Competitiveness, risk taking, and violence: The young male syndrome. *Ethology and Sociobiology*, 6(1):59–73, 1985.

[8] Liviu Panait and Sean Luke. Cooperative multi-agent learning: The state of the art. *Autonomous Agents and Multi-Agent Systems*, 11(3):387–434, 2005.

[9] G.G. Simpson. *The Meaning of Evolution*. New Haven, 1949.