

Learning to be Risk Averse?

Robert E. Marks

Economics, University of New South Wales,

and University of Melbourne

Email: robert.marks@gmail.com

Abstract—The purpose of this research is to search for the best (highest performing) risk profile of agents who successively choose among risky prospects. An agent’s risk profile is his attitude to perceived risk, which can vary from risk preferring to risk neutral (an expected-value decision maker) to risk averse. We use the Genetic Algorithm to search in the complex stochastic space of repeated lotteries. We find that agents with a CARA utility function learn to possess risk-neutral risk profiles. Since CARA utility functions are wealth-independent, this is not surprising. When agents have wealth-dependent, CRRA utility functions, however, they also learn to possess risk profiles that are about risk neutral (from slightly risk-averse to even slightly risk-preferring), which is surprising.

Keywords—

I. INTRODUCTION

Informally, it is widely held that in an uncertain world, with the possibility of the discontinuity of bankruptcy, the most prudent risk profile is risk aversion. Indeed, “Risk aversion is one of the most basic assumptions underlying economic behavior” [1], perhaps because “a dollar that helps us avoid poverty is more valuable than a dollar that helps us become very rich” [2]. But is risk aversion the best risk profile? Even with bankruptcy as a possibility?

To answer this question, we use two kinds of utility function (the wealth-independent exponential utility function, or Constant Absolute Risk Aversion CARA, and the Constant Relative Risk Aversion CRRA function, which is sensitive to the agent’s level of wealth) and run computer experiments in which each agent chooses among three lotteries, and is then awarded with the outcome of the chosen lottery k .

Repetition of this choice by many agents allows us use a technique from machine learning – the Genetic Algorithm or GA [3] – to search for the best risk profile, where “best” means the highest average payoff when choosing among lotteries.

Modelling the agent’s utility directly allows us to avoid the indirect inference of Szpiro [1], who argues that the evolutionary learning technique of the GA does two things: it allows wealth-maximizing agents to succeed even in highly stochastic environments, and it allows the emergence of risk aversion. Indeed, Szpiro argues that risk aversion is the best risk profile to adopt in such an environment.

II. DECISIONS UNDER UNCERTAINTY AND RISK PROFILES

The von Neumann-Morgenstern formulation of the decision-maker’s attitude to risk is based on the observation that individuals are not always expected-value decision makers. That is, there are situations in which people apparently prefer

a lower certain outcome to the higher expected (or probability-weighted) outcome of an uncertain prospect (where the possible outcomes and their possibly subjective, or Bayesian, probabilities are known).

An example is paying an insurance premium that is greater than the expected loss without insurance. On the other hand, people will sometimes “gamble” by apparently preferring a lower uncertain outcome to a higher sure thing: this is risk-preferring.

We can formalise this by observing that, by definition, the utility of a lottery is its expected utility, or

$$U(L) = \sum p_i U(x_i), \quad (1)$$

where each (discrete) outcome x_i occurs with probability p_i , and $U(x_i)$ is the utility of outcome x_i . It is useful to define the Certainty Equivalent \tilde{x} (or C.E.), which is a certain outcome which has the identical utility as the lottery:

$$U(\tilde{x}) = U(L) = \sum p_i U(x_i) \quad (2)$$

We can use the C.E. to describe the decision-maker’s risk profile [4]. Define the Expected Value \bar{x} of the Lottery as:

$$\bar{x} = \sum p_i x_i. \quad (3)$$

When $\tilde{x} = \bar{x}$, then the decision-maker’s utility function exhibits risk neutrality; when $\tilde{x} < \bar{x}$, then risk aversion; and when $\tilde{x} > \bar{x}$, then risk preferring.

A. Approximating the Certainty Equivalent

Expand utility $U(\cdot)$ about the expected value \bar{x} .

$$U(x_0) \approx U(\bar{x}) + (x_0 - \bar{x})U'(\bar{x}) + \frac{1}{2}(x_0 - \bar{x})^2 U''(\bar{x})$$

The C. E. \tilde{x} of a continuous lottery is obtained by integration over the probability density function (p.d.f.) $f_x(\cdot)$:

$$U(\tilde{x}) = \int dx_0 U(x_0) f_x(x_0)$$

$$\therefore U(\tilde{x}) \approx U(\bar{x}) + 0 + \frac{1}{2}\sigma^2 U''(\bar{x}), \quad (4)$$

where σ^2 is the variance. But, by expansion,

$$U(\tilde{x}) \approx U(\bar{x}) + (\tilde{x} - \bar{x})U'(\bar{x}). \quad (5)$$

Therefore, from (4) and (5),

$$\tilde{x} - \bar{x} \approx \frac{1}{2}\sigma^2 \frac{U''(\bar{x})}{U'(\bar{x})}$$

$$\therefore \tilde{x} \approx \bar{x} + \frac{1}{2}\sigma^2 \frac{U''(\bar{x})}{U'(\bar{x})} \quad (6)$$

B. Risk Aversion

Risk aversion is not indicated by the slope of the utility curve: it's the *curvature*: if the utility curve is locally –

- linear (say, at a point of inflection), then the decision maker is locally risk neutral;
- concave (its slope is decreasing – Diminishing Marginal Utility), then the decision maker is locally risk averse;
- convex (its slope is increasing), then the decision maker is locally risk preferring.

III. UTILITY FUNCTIONS

We consider two types of utility function:

- 1) those which exhibit constant risk preference across all outcomes (so-called wealth-independent utility functions, or Constant Absolute Risk Aversion CARA functions), and
- 2) those where the risk preference is a function of the wealth of the decision maker (the Constant Relative Risk Aversion CRRA functions).

A. Wealth Independence

If an increase of all outcomes in a lottery by an equal amount Δ increases the C.E. of the lottery by Δ , then the decision maker exhibits wealth independence:

$$U(\tilde{x} + \Delta) = U(L') = \sum p_i U(x_i + \Delta).$$

Acceptance of this property restricts possible utility functions to be linear (risk neutral) or exponential, or constant-absolute-risk-aversion (CARA) functions.

CARA utility functions characterise risk preference by a single number, the *risk aversion coefficient*, γ . Since CARA utility functions are wealth-independent, any aversion to bankruptcy is thus precluded, by definition. Whether a decision maker exhibits a wealth-independent utility function is an empirical question.

B. CARA Utility Functions

When utility is linear in outcomes, the decision maker is risk-neutral, across all outcomes, but such a simple constant-risk-profile utility function is of no further interest. Instead, we consider the exponential constant absolute risk averse (CARA) functions, where utility U is given by

$$U(x) = 1 - e^{-\gamma x}, \quad (7)$$

where $U(0) = 0$ and $U(\infty) = 1$, and where γ is the *risk aversion coefficient*:

$$\gamma = -\frac{U''(x)}{U'(x)}. \quad (8)$$

1) *Risk Aversion with Exponential Utility*: From (6) and (8), for exponential utility,

$$\tilde{x} \approx \bar{x} - \frac{1}{2}\sigma^2\gamma$$

which indicates that when $\gamma = 0$, then $\tilde{x} \approx \bar{x}$ (risk neutrality), when $\gamma > 0$, then $\tilde{x} < \bar{x}$ (risk averse), and when $\gamma < 0$, then $\tilde{x} > \bar{x}$ (risk preferring), with positive variance.

C. CRRA Utility Functions

We want a utility function which is *not* wealth-independent, to see whether that will result in risk-averse agents doing best.

The Arrow-Pratt measure of relative risk aversion (RRA) ρ is defined as

$$\rho(w) = -w \frac{U''(w)}{U'(w)} = w\gamma \quad (9)$$

This introduces wealth w into the agent's risk preferences, so that lower wealth can be associated with higher risk aversion. The risk aversion coefficient γ is as in (8).

The Constant Elasticity of Substitution (CES) utility function:

$$U(w) = \frac{w^{1-\rho}}{1-\rho}, \quad (10)$$

with positive wealth, $w > 0$, exhibits constant relative risk aversion CRRA, as in (9).

1) *Risk Aversion with CES Utility*: In the CRRA simulations, we use the cumulative sum of the realisations of payoffs won (or lost, if negative) in previous lotteries chosen by the agent plus the possible payoff in this lottery as the wealth w in (10). Each agent codes for ρ .

From (6), the C.E. with CES utility is approximated by

$$\tilde{x} \approx \bar{x} - \frac{1}{2} \frac{\rho}{w} \sigma^2.$$

Iff $\frac{1}{2} \frac{\rho}{w} \sigma^2 > 0$ (or $\rho/w > 0$), then then C.E. $\tilde{x} <$ the expected mean \bar{x} , and the decision maker is risk averse.

With $w > 0$, $\rho > 0$ is equivalent to risk aversion. With $w > 0$ and $\rho = 1$, the CES function becomes the (risk-averse) logarithmic utility function, $U(w) \approx \log(w)$. With $w > 0$ and $\rho < 0$, it is equivalent to risk preferring.

IV. THE SIMULATIONS

Each lottery is randomly constructed: the two payoffs (“prizes”) are randomly chosen in the interval between – and + MAP, (where the Maximum Absolute Prize, MAP, is usually 100); and the probability is also chosen randomly. (Each lottery has, of course, a single degree of freedom for probability). Each agent calculates the expected utility of each of the three lotteries, using its utility function (a function of its γ or ρ/w), and chooses the lottery k with the highest expected utility. To do this, agents know the prizes and probabilities of all three lotteries.

Then the actual (simulated) outcome of the chosen lottery k is randomly realised, using its probability. The winnings of the Constant Absolute Risk Aversion agent (respectively,

the wealth of the Constant Relative Risk Aversion agent) is incremented accordingly. Each agent successively chooses 1000 lotteries.

Calculate the three expected utilities for lotteries X, Y, and Z, functions of γ (or ρ and w):

$$U(X) = p_x U(x_1) + (1 - p_x) U(x_2)$$

$$U(Y) = p_y U(y_1) + (1 - p_y) U(y_2)$$

$$U(Z) = p_z U(z_1) + (1 - p_z) U(z_2)$$

Choose the lottery I with the highest expected utility. Win (or lose) whichever prize (i_1 or i_2) is realised in that lottery, based in the lottery's probability p_i .

A. Searching with the Genetic Algorithm

We use a population of 100 agents, each of which has a average winnings or a cumulative level of wealth, based on its risk profile and the successive outcomes of its choices among the lotteries. The GA's mutation rate is controllable by the simulator, on-screen.

We use an implementation [5] of the GA to search for the best risk profile. That is, we select the best-performing agents to be the "parents" of the next generation of agents, which is generated by "crossover" and "mutation" of the chromosomes of the pairs of parents. Each of the new generation of agents chooses the lottery k with highest expected utility a thousand times. Again, the best are selected to be the parents of the next generation.

We use the GA simulation in this search as an empirical alternative to solving for the best (highest performing) risk profile analytically. Note that Rabin [2] asserts that "theory actually predicts virtual risk neutrality." We return to this paper in the Discussion below.

B. Simulations with Utility-Maximizing Agents

Using NetLogo [6], we model each agent as a binary string which codes to its risk-aversion coefficient, γ , for CARA agents (respectively, ρ , for CRRA agents) in the interval ± 1.048576 .

Each lottery is a two-prize lottery, where each prize is chosen from a uniform distribution, between $-$ and $+$ MAP (Maximum Absolute Prize), where MAP can be set up to 100 by the simulator, and the single probability is chosen randomly from uniform $[0,1]$.

Each agent chooses the lottery k with the highest expected utility from (1) and (7), based on its value of γ (respectively, ρ and wealth w). Then a realised outcome is calculated for that lottery, based on its probability.

Each agent faces 1000 lottery choices, and the cumulative winnings that agent's "fitness" for the GA.

C. The CARA Results

The windows in Fig. 1 captured from the NetLogo simulations¹ show three things clearly:

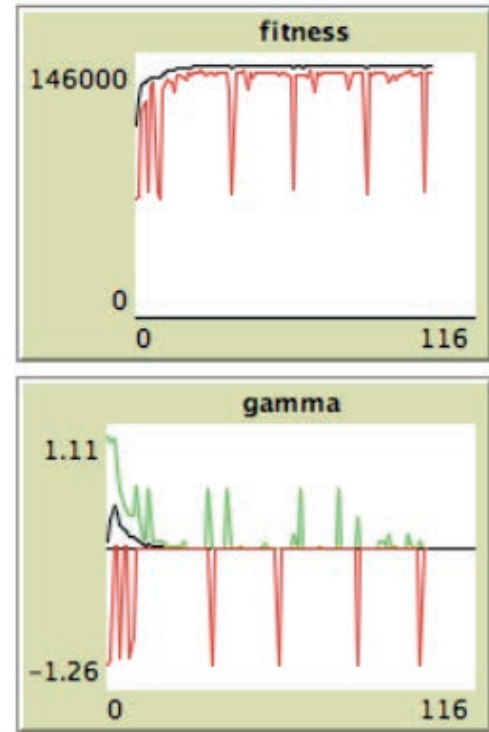


Fig. 1. CARA Simulation Results

- 1) The mean (black) fitness (cumulative winnings) grows quickly to a plateau after 20 generations (along the x axis) or so;
- 2) the mean, maximum, and minimum risk-aversion coefficients γ (respectively, black, green, red) converge to close to zero (risk neutrality) over the same period, and
- 3) Any γ deviation from zero up (more risk-averse) or down (more risk-preferring) leads to the minimum (red) fitness in that generation collapsing from close to the mean fitness.

These observations clearly show that CARA agents perform best (in terms of their lottery winnings) who are closest to risk neutral ($\gamma = 0$). Too risk averse, and they forgo fair lotteries; too risk preferring and they choose too many risky lotteries.

Despite our prior belief, the CARA agents do not learn to be risk averse, but to be risk neutral. Is this because the wealth-independent CARA utility function precludes bankruptcy?

D. The CRRA Results

We could, of course, put a floor on agent wealth, below which is oblivion, but better to use a utility formulation that is not wealth independent and repeat the search. We use the CES utility functions (10) that exhibits CRRA.

The results are surprising (see Fig. 2, for one simulation run):² the CRRA agents do not learn to be risk averse, but are very close to risk neutral.

¹See <http://www.agsm.edu.au/bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-3l2p.html> for a Java applet and the NetLogo code.

²See <http://www.agsm.edu.au/bobm/teaching/SimSS/NetLogo4-models/DRA-CRRA-EU-revCD-3l2p.html> for a Java applet and the NetLogo code.




Fig. 2. CRRA Simulation Results

Remember: $\gamma = \frac{\rho}{w}$, so dividing the ρ values by the high w values attained implies corresponding minute values of γ here.

E. Changing the Fitness Function

For the simulations in 4.1.(CARA) and 4.2 (CRRA), the GA's fitness function (or performance measure) is the arithmetic mean of the agent's cumulative winnings (or losses) in the 1000 lottery choices:

$$fitness = \frac{1}{1000} \sum_t L_{kt}$$

where L_{kt} is the realisation of the the highest-expected-utility lottery k , chosen at period t .

We now use the geometric mean:

$$fitness = \prod_t (L_{kt})^{0.001}$$

The motivation behind this is that the log of the geometric mean of the expected value of a lottery equals its expected utility with a logarithmic utility function.

Does risk-aversion emerge with this new fitness? ³ This would require $\gamma > 0$ (for CARA) and $\rho > 0$ (for CRRA). No: whether we use arithmetic or geometric means for our fitness function, the agents – whether CARA or CRRA – appear to converge (if slowly) on risk-neutral decision making. ⁴

³See <http://www.agsm.edu.au/bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-GM-revC.html> for a Java applet and the NetLogo code of the CARA model with geometric mean.

⁴See <http://www.agsm.edu.au/bobm/teaching/SimSS/NetLogo4-models/DRA-CRRA-EU-GM-revH-312p.html> for a Java applet and the NetLogo code of the CRRA model with geometric mean.

V. DISCUSSION

Unlike the GA simulations of Szpiro [1], we find that the best-performing CARA agents are risk-neutral, not risk averse. Because of the indirect way in which Szpiro modelled the risk profiles of his agents (unlike a referee's suggestion, footnote 3, Szpiro's model "only distinguishes between risk-averse automata and all others"), explanation of the contradictory results is not easy, but since our models allow any risk profile to emerge, we argue that they are more general than Szpiro's.

Should we be surprised that risk neutrality does better than risk aversion in CARA utility functions? Rabin [2] suggests a reason why not, at least for small-stakes lotteries. He argues that von Neumann-Morgenstern expected-utility theory is inappropriate for reconciling actual human behaviour as revealed in risk attitudes over large stakes and small stakes. If there is risk aversion for small stakes, then expected-utility theory predicts wildly unrealistic risk aversion when the decision maker is faced with large stakes. Or risk aversion for large stakes must be accompanied by virtual risk neutrality for small stakes.

Rabin [2] argues that *loss aversion* [7], rather than risk aversion, is a better (i.e. more realistic) explanation of how people actually behave when faced with risky decisions. This suggests possibilities for further simulations, although "loss aversion" suggests a prior conclusion.

But we do not appeal to empirical evidence or even to prior beliefs of what sort of risk profile is best. Whereas there has been much research into reconciling actual human decision making with theory (see [8]), we are interested in seeing what is the best (i.e. most profitable) risk profile for agents faced with risky choices.

And we find that for wealth-independent CARA utility functions (exponential) agents learn to become risk-neutral decision makers in order to maximise their returns when choosing among risky propositions. This is different from the risk-averse agents that Szpiro [1] observed. But for wealth-dependent CRRA utility functions (CES) our agents often do learn to be slightly risk averse, as expected, but not always.

An analytical study [9] posits an adaptive process for decision-making under risk such that, despite people being seen to be risk averse over gains and risk seekers over losses with respect to the current reference point [7] – the so-called dual risk attitude, the agent eventually learns to make risk-neutral choices. Their result appears consistent with our results, although the learning in their model is not that of the GA, but rather agents observing how their choices result in systemic undershooting (or overshooting) of their targets, which then results in more realistic targets and choices. Their lotteries are symmetrical (for tractability), unlike ours. Our results suggest that their results might generalise to asymmetric lotteries, such as our.

VI. CONCLUSION

Using a demonstrative agent-based model – which demonstrates principles, rather than tracking historical phenomena – we have used the Genetic Algorithm to search the complex, stochastic space of decision making under uncertainty, in which agents successively choose among three (asymmetric) lotteries with randomly allocated probabilities and outcomes

(two per lottery), in order to maximize their expected utilities. The GA searches for the best-performing utility function, whether CARA (or wealth-independent) or CRRA (when wealth, and hence bankruptcy, matters).

Despite our prior belief that a risk-averse agent does best in these circumstances, we find that both CARA and CRRA utilities converge (if slowly, in the latter case) to risk neutrality. This is consistent with analytical work that proves that with symmetric lotteries, and agents with dual risk attitude, risk-neutral decisions are the eventual outcome of agents adjusting their aspirations and targets in response to the realisations of their choices. Further work will reveal the robustness of our results.

ACKNOWLEDGMENT

I should like to thank Simon Grant, Luis Izquierdo, the participants of the Complex Systems Research Summer School 2007 at Charles Sturt University, the participants at the 26th Australasian Economic Theory Workshop 2008 at Bond University, Jasmina Arifovic, James Andreoni, Seth Tisue, Marco LiCalzi, and Nigel Gilbert for his implementation of the Genetic Algorithm in NetLogo. Two anonymous reviewers have also been helpful.

Note: Java applets of the simulation models and the NetLogo code is available online, together with graphical output of the simulation results, as referenced in the four footnotes above.

These models will also generate real-time results, including graphs of their performance, when your computer's Java security allows. Moreover, you can explore the impact of the GA mutation rate on the simulation evolution.

REFERENCES

- [1] G. G. Szpiro, "The emergence of risk aversion," *Complexity*, vol. 2, pp. 31-39, 1997.
- [2] M. Rabin, "Risk aversion and expected-utility theory: a calibration theorem," *Econometrica*, vol. 68, pp. 1281-1292, 2000.
- [3] J. H. Holland, *Adaptation in Natural and Artificial Systems*, 2nd ed. Cambridge, MA: M. I. T. Press, 1992.
- [4] R. A. Howard, "The foundations of decision analysis," *IEEE Trans. on Systems Science and Cybernetics*, vol. ssc-4, pp. 211-219, 1968.
- [5] N. Gilbert, "Axelrod's Iterated prisoners' dilemma tournament," in *Simulation for the Social Scientist*, N. Gilbert and K. G. Troitzsch, 2nd ed. Maidenhead, England: Open University Press, 2005. [Online]. Available: <http://cress.soc.surrey.ac.uk/s4ss/code/NetLogo/axelrod-ipd-ga.html>
- [6] U. Wilensky, NetLogo. Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston, IL. (Version 4), 1999. [Online]. Available: <http://ccl.northwestern.edu/netlogo>.
- [7] D. Kahneman, and A. Tversky, "Prospect theory: an analysis of decision under risk," *Econometrica*, vol. 47, pp. 263-291, 1979.
- [8] W. B. Arthur, "Designing economic agents that act like human agents: A behavioral approach to bounded rationality," *American Economic Review Papers & Proceedings*, vol. 81, pp. 353-360, 1991.
- [9] S. DellaVigna and M. LiCalzi, "Learning to make risk neutral choices in a symmetric world," *Mathematical Social Sciences*, vol. 41, pp. 19-37, 2001.