Searching for Agents' Best Risk Profiles

Robert E. Marks

Economics, UNSW Australia,
Sydney, NSW 2052, Australia
robert.marks@gmail.com
http://www.agsm.edu.au/bobm

Abstract. The purpose of this research is to seek the best (highest performing) risk profiles 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, or even a dual-risk attitude. We use the Genetic Algorithm to search in the complex stochastic space of repeated lotteries. We examine three families of utility (or value) functions: wealth-independent CARA and wealth-dependent CRRA, in which an agent's risk profile is unchanging, and the Dual-Risk-Profile (DRP) functions from Prospect Theory, in which the agent can be risk-averse (for gains) or risk preferring (for losses). Statistical analysis of the simulation results suggests that the best (profit-maximizing) CRRA functions are risk neutral, while the other functions remain slightly risk-averse. The most profitable are slightly risk-averse DRP functions.

Keywords: decision making under risk, comparing utility functions, Prospect Theory, Genetic Algorithms, risk aversion, risk neutrality

1 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.[1] Indeed, "Risk aversion is one of the most basic assumptions underlying economic behavior" [2], perhaps because "a dollar that helps us avoid poverty is more valuable than a dollar that helps us become very rich" [3]. But is risk aversion the best risk profile? Even with bankruptcy as a possibility?

To answer this question, we use three kinds of utility function: the wealth-independent exponential utility function, or Constant Absolute Risk Aversion CARA; the Constant Relative Risk Aversion CRRA function, which is sensitive to the agent's level of wealth; and the DRP functions of Prospect Theory, where an agent's risk profile can vary depending on prospects of losing or gaining. We run computer experiments in which each agent chooses among three lotteries, and is then awarded with the outcome of the chosen lottery k.

[1] An preliminary version of this paper was presented at the IEEE Computational Intelligence for Finance Engineering & Economics 2014, London, March 29 [1].

Repetition of these choices by many agents allows us to use a technique from machine learning – the Genetic Algorithm or GA [4] – to search for the best function from each utility family, 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 [2], 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. We compare the cumulative winnings (fitnesses) of our agents to see whether this is so.

2 Decisions under Risk 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

$$U(L) = \sum p_i U(x_i), \tag{1}$$

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:

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

We can use the C.E. to describe the decision-maker's risk profile [5]. Define the Expected Value $\bar{x}$ of the Lottery as:

$$\bar{x} = \sum p_i x_i. \tag{3}$$

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.

2.1 Approximating the Certainty Equivalent

Expand utility $U(.)$ 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.) $f_x(.)$:

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

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

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

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

Therefore, from (4) and (5),

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

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

2.2 Risk Aversion

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

- linear (say, at a point of inflection, where $U'' = 0$), 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.

3 Utility Functions

We consider three 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);
2. those where the risk preference is a function of the wealth of the decision maker (the Constant Relative Risk Aversion CRRA functions); and
3. those in which the risk profile is a function of the prospect of gaining (risk averse) or losing (risk preferring): the DRP Value Functions from Prospect Theory.

3.1 CARA Utility Functions

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 – constant-absolute-risk-aversion (CARA) functions.

CARA utility functions charactise 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.

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 CARA functions, where utility U is given by

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

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

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

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.

3.2 CRRA Utility Functions

We want utility functions which are *not* wealth-independent, to see whether such functions will result in risk-averse agents doing best.

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

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

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:

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

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

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.

3.3 The Dual-Risk-Profile DRP function from Prospect Theory

From Prospect Theory [8], we model the DRP Value Function, which maps from quantity X to value V with the following two-parameter equations (with $\beta > 0$ and $\delta > 0$):

$$V(X) = \frac{1 - e^{-\beta X}}{1 - e^{-100\beta}}, 0 \leq X \leq 100, \tag{11}$$

$$V(X) = -\delta \frac{1 - e^{\beta X}}{1 - e^{-100\beta}}, -100 \leq X < 0. \tag{12}$$

The parameter $\beta > 0$ models the curvature of the function, and the parameter $\delta > 0$ the asymmetry associated with losses. The DRP function is not wealth independent. This function (in Fig. 1, with $\delta = 1.75$, for prizes between $\pm\$100$) exhibits the S-shaped asymmetry postulated by Kahneman and Tversky [8]. It exhibits risk seeking (loss aversion) when X is negative with respect to the reference point $X = 0$, and risk aversion when X is positive. We use here a linear probability weighting function (hence no weighting for smaller probabilities). As Fig. 1 suggests, as $\delta \to 1$ and $\beta \to 0$, the value function asymptotes to a linear, risk-neutral function (in this case with a slope of 1).

We use the GA to search the joint plane (β, δ) as the agents (each characterised by a point in the (β, δ) plane) choose the lottery that has the greatest expected value of the three. Each lottery has two known prizes in the interval of $[-\$100, +\$100]$ of known probabilities, p_i. So the agent chooses the lottery k with the highest expected value:

$$U_k = \sum_{i=1}^{2} p_{ki} V(X_{ki}, \beta, \delta).$$

The GA jointly searches for points in (β, δ) that result in high payoffs after the payoff of each chosen lottery is subsequently realised, based on the probabilities of the possible outcomes.

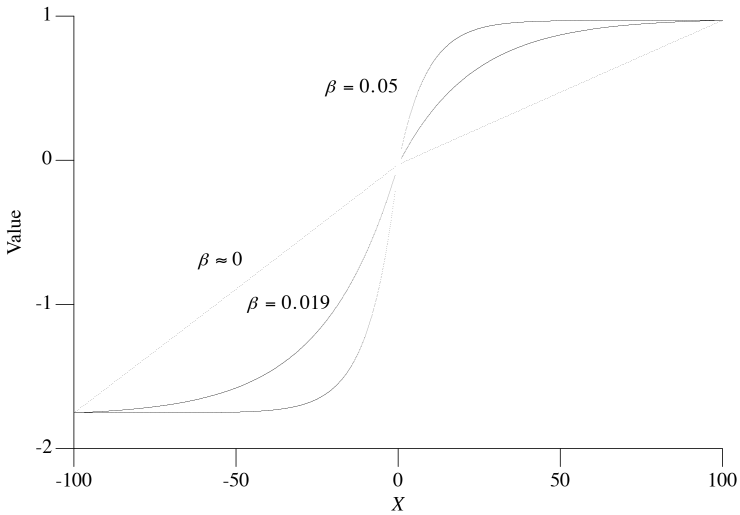

Fig. 1. A Prospect Theory (DRP) Value Function

4 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 \$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 or value function (a function of its γ or ρ/w or (β, δ)), and chooses the lottery k with the highest expected utility. Doing 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 agent (that is, the wealth of the CRRA 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, or (β, δ)):

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

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

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

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

4.1 Searching with the Genetic Algorithm

We use a population of 100 agents, each of which has an 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 [6] 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 [3] asserts that "theory actually predicts virtual risk neutrality." We return to this in the Discussion below.

4.2 Simulations with Utility-Maximizing (or Value-Maximizing) Agents

Using NetLogo [7], we model each agent as a binary string which codes to its risk-aversion coefficient/s, (γ for CARA agents, ρ for CRRA agents, and (β, δ) for DRP agents) in the interval ± 1.048576. The DRP agents search for $0 \leq \beta < 0.21$ and for δ in the interval ± 10.48.

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 (7) and (10), based on its value of γ (respectively, ρ and wealth w), or from (11) and (12), based on its value of β and δ. 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. Because all agents face the same probabilistic lotteries, we can also compare the the three utility families based on their average winnings or fitnesses. The processes are stochastic. For each model we perform a number n of Monte Carlo simulation runs to obtain sufficient data to analyse the results statistically.

4.3 The CARA Results

The on-line simulations [2] show three things clearly:

[2] See `http://www.agsm.edu.au/bobm/teaching/SimSS/NetLogo4-models/RA-CARA-EU-312p.html` for a Java aplet and the Netlogo code.

1. The mean (black) fitness (cumulative winnings) grows quickly to a plateau after 20 generations or so;
2. the mean, maximum, and minimum risk-aversion coefficients γ 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 fitness in that generation collapsing from close to the mean fitness.

These observations 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.

Eye-balling single output plots, however, is not sufficient to reach clear conclusions about the best utility functions. We have preferred 55 independent Monte Carlo runs using the GA to search for better CARA utility functions.

The correlation between γ and Fitness in these MC simulations (where Fitness is the winnings averaged across 100 agents, each of which chooses the expected "best" of three lotteries 1000 times per generation, for 200 generations) is 0.7148. This suggests that the larger the value of γ, the higher the value of Fitness. The p-value for γ against a null hypothesis of $H_0 : \mu_\gamma = 0$ is 0.0006, which provides a very strong presumption against the null, that is, although it is close to the risk-neutrality of $\gamma = 0$, the CARA function does not converge to exact risk neutrality.[3] The final means of γ in the 55 runs are significantly positive, suggesting weak risk-aversion.

The wealth-independent CARA utility function precludes bankruptcy. What of utility function that does not exclude this possibility?

4.4 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.[4] We have performed 109 independent Monte Carlo runs using the GA to search for better CRRA utility functions.

The correlation between ρ and Fitness is 0.0907. This suggests that ρ and Fitness are not correlated. The p-value for ρ against a null hypothesis of $H_0 : \mu_\rho = 0$ is 0.2996, which provides no presumption against the null, that is, the data suggest that the mean $\mu_\rho = 0$, or the CRRA function converges to risk neutrality, despite our prior expectations for this function.

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

[3] For the detailed statistics of this and the other four analyses, see the Appendix in the paper at `http://www.agsm.edu.au/bobm/papers/singapore14.pdf`.

[4] See `http://www.agsm.edu.au/bobm/teaching/SimSS/NetLogo4-models/DRA-CRRA-EU-revCD-312p.html` for a Java aplet and the NetLogo code.

4.5 The Dual-Risk-Profile Results

We considered first the marginal results: holding β constant at zero, asking what values of δ emerge as conditionally best, then, holding δ constant at unity, asking what values of β emerge as conditionally best. That is, first considering a kinked function, possibly linear, and then a symmetric DRP function, possibly linear.

We performed 55 independent Monte Carlo runs using the GA to search for better δ, while holding $\beta = 0$. The correlation between δ and Fitness in these MC simulations (where Fitness is the winnings averaged across the 100 agents, each of which chooses the expected "best" of three lotteries 1000 times per generation, for 200 generations) is 0.5602. This suggests that the larger the value of δ, the higher the value of Fitness. The p-value for δ against a null hypothesis, $H_0 : \mu_\delta = 1$, is effectively zero (< 0.00001), which provides a very strong presumption against the null, that is, the data suggest that the mean $\mu_\delta \neq 1$.

We then performed 50 independent Monte Carlo runs using the GA to search for better β, while holding $\delta = 1$. The correlation between β and Fitness in these MC simulations is 0.0941. This suggests that β and Fitness are not correlated. The p-value for β against a null hypothesis, $H_0 : \mu_\beta = 0$ is 0.0012, which provides a strong presumption against the null, that is, the data suggest that the mean $\mu_\beta \neq 0$.

Then we undertook a search in the (β, δ) plane. We performed 54 independent Monte Carlo runs using the GA to search for better β and δ jointly. Fig. 2 shows the result with the two variables' final means across these runs, where $\beta = 0.007186$ and $\delta = 1.2598$: almost risk neutral (in which an expected-value decision maker becomes an expected-outcome decision maker).[5]

The correlation between δ and β (which is -0.0015) suggests that there is little if any trade-off between the two values in maximizing Fitness.[6] This is not surprising: from the two marginal explorations, we see that with fixed $\beta = 0$, the best mean $\delta = 1.4658$ (compared to 1.2598 when the search is in the (β, δ) plane), while with fixed $\delta = 1$, the best mean $\beta = 0.003846$ (compared with 0.007186 when the search is in the joint plane). That is, fixing δ constrains the Fitness more than fixing β.

4.6 Comparing Models

This can be seen in another way. All five sets of Monte Carlo simulations are searching the same space: given the known prizes (between \$100 and \$100) and known probabilities, choose the expected "best" lottery. This means we can

[5] See `http://www.agsm.edu.au/bobm/teaching/SimSS/NetLogo4-models/RA-PTh-EV-both.html` for a Java aplet and the NetLogo code of a DRP Value Function model.

[6] Another formalization of Prospect Theory Value Functions is to model separate curvature parameters for gains ($X > 0$) and losses ($X < 0$), but this does not really capture the full asymmetry between gains and losses that our model includes. And $\delta \neq 1$ implies a different slope (given $\beta = 0$) for losses.

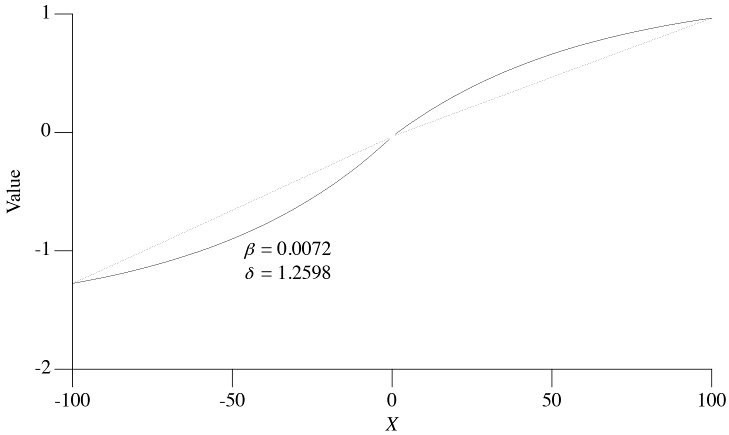

Fig. 2. The Best Prospect Theory (DRP) Value Function

compare the Fitnesses (dollar winnings) across the simulation runs. This is shown in Table 1.

Table 1: Mean Fitnesses of the Five Models

Model	Mean Fitness ($)
CARA	37,650
CRRA	29,403
DRP with $\beta = 0$	37,666
DRP with $\delta = 1$	38,879
DRP joint β, δ	37,721

It is clear that CRRA, the only model whose "best" parameter reflects risk-neutrality, performs worst at maximizing Fitness, while the best model is the DRP model from Prospect Theory with symmetric ($\delta = 1$) loss-averting and gain-preferring (its utility is convex for losses and concave for gains). This is strange: a constrained optimization outperforming an unconstrained optimisation, when the constrained value is available to the unconstrained. The GA search of the joint model could/should find that Fitness is higher when $\delta \approx 1$ but hasn't. This suggests that the runs be lengthened, perhaps because the joint search in (β, δ) is hard. Indeed, from Table 1, the apex of the hill of optimal fitness is quite flat. It is likely that this anomaly suggests that the GA optimisations have been prematurely terminated; would longer runs change our risk profile results? Further work will tell.

At any rate, Table 1 shows that CARA, DRP with $\beta = 0$, and DRP with joint β, δ search are very close in terms of best Fitness.

5 Discussion

Like the GA simulations of Szpiro [2], we find that the best-performing CARA agents are risk-averse, not risk-neutral. 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"), while our models allow any CARA risk profile to emerge, we argue that our results are more general than Szpiro's.

Rabin [3] suggests a reason why risk-neutral functions will not do better than risk-averse functions, 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.

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 [9]), we are interested in seeing what is the best (i.e. most profitable) risk profile for agents faced with risky choices.

We find that for wealth-independent CARA utility functions (exponential) agents do not learn to become risk-neutral decision makers in order to maximise their returns when choosing among risky propositions. But for wealth-dependent CRRA utility functions (CES) our agents do learn to be risk neutral, despite the possibility (even if small) of bankruptcy, or the loss of all accumulated wealth.

Rabin [3] argues that *loss aversion* [8], rather than risk aversion, is a better (i.e. more realistic) explanation of how people actually behave when faced with risky decisions. This is captured in our DRP Value Function.

An analytical study of Prospect Theory Value Functions [10] 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 [8] – the so-called dual risk attitude, with utility convex for losses and concave for gains – the agent eventually learns to make risk-neutral choices. Their result appears consistent with our results for the CRRA model, 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 ours, at least for CRRA utility.

A simulation study [11] examines the survival dynamics of investors with different risk preferences in an agent-based, multi-asset, artificial stock market and finds that investors' survival is closely related to their risk preferences. Examining eight possible risk profiles, the paper finds that only CRRA investors with relative risk aversion coefficients close to unity (log-utility agents) survive in the long run (up to 500 simulations). This does not appear consistent with our results.

Our last finding is obtained by comparing the mean fitnesses (accumulated winnings) of the five models. We find that a symmetric DRP model (with $\delta = 1$) does better than any of the other models, while the only model which learns to be risk-neutral (the CRRA model) does worst.

6 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 risk, 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, among CARA (or wealth-independent), CRRA (when wealth, and hence bankruptcy, matters), or for the best-performing Value Function, which exhibits the DRP of Prospect Theory, although we use the same parameter β to describe the curvature of both risk averse (gains) and risk preferring (losses), which is a restriction that could be relaxed with further study.

Consistent with our prior belief that a risk-averse agent does best in these circumstances, we find that only one of our three models – CRRA – converges to risk neutrality. Our findings are therefore only partly 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. But the other two models – CARA and the DRP from Prospect Theory – converge on (slightly) risk-averse parameters, when we search using accumulated winnings as the Fitness.

Comparing the mean accumulated winnings across our models, we find that the best performing model is a symmetric DRP model. This might prove of use to future simulators.

Simulations, of course, can not prove necessity, only sufficiency [12], so our results for each of the three functions – CARA, CRRA, and DRP Value Functions – are existence proofs only: the best (highest performing) functions, in choosing among lotteries of known prizes and known probabilities, do not generally tend to risk neutral (linear). The results suggest relaxing the assumption of known probabilities might also be of interest. These results therefore tend to confirm the common knowledge that a small amount of risk aversion is best in a risky world.

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, the participants at the 2014 IEEE CIFEr meetings in London, Jasmina Arifovic, James Andreoni, Seth Tisue, Marco LiCalzi, Shu-Heng Chen, Arthur Ramer, David Midgley, and Nigel Gilbert for his implementation of the Genetic Algorithm in NetLogo. Several anonymous reviewers have also been helpful.

Note: Java aplets of the simulation models and the NetLogo code are available online, together with graphical output of the simulation results, as referenced in the three footnotes above. These models will also generate real-time results, including graphs of their performance, when one's computer's Java security allows. Moreover, one can explore the impact of the GA mutation rate on the simulation evolution.

References

1. Marks, R. E.: Learning to be Risk Averse? In: Serguieva, A., Maringer, D., Palade, V., Almeida, R.J. (eds.) Proc. of the 2014 IEEE Computational Intelligence for Finance Engineering & Economics (CIFEr). London, March 28-29, IEEE Computational Intelligence Society, pp. 1075-1079 (2014)
2. Szpiro, G. G.: The Emergence of Risk Aversion. Complexity 2, 31-39 (1997)
3. Rabin, M.: Risk Aversion and Expected-Utility Theory: a Calibration Theorem. Econometrica 68, 1281-1292 (2000)
4. Holland, J. H.: Adaptation in Natural and Artificial Systems. 2nd ed. MIT Press, Cambridge (1992)
5. Howard, R. A.: The Foundations of Decision Analysis. IEEE Trans. on Systems Science and Cybernetics ssc-4, 211-219 (1968)
6. Gilbert, N.: Axelrod's Iterated Prisoners' Dilemma Tournament. In: N. Gilbert, K. G. Troitzsch. Simulation for the Social Scientist. 2nd ed. Open University Press, Maidenhead, England (2005) `http://cress.soc.surrey.ac.uk/s4ss/code/NetLogo/axelrod-ipd-ga.html`
7. Wilensky, U.: NetLogo. Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston, IL. (Version 4) (1999) `http://ccl.northwestern.edu/netlogo`
8. Kahneman, D., Tversky, A.: Prospect Theory: an Analysis of Decision Under Risk. Econometrica. 47, 263-291 (1979)
9. Arthur, A. B.: Designing Economic Agents that Act like Human Agents A Behavioral Approach to Bounded Rationality. American Economic Review Papers & Proceedings. 81, 353-360 (1991)
10. DellaVigna, S., LiCalzi, M.: Learning to Make Risk Neutral Choices in a Symmetric World. Mathematical Social Sciences. 41, 19-37 (2001)
11. Chen, S.-H., Huang, Y. C.: Risk Preference, Forecasting Accuracy and Survival Dynamics: Simulation Based on a Multi-Asset Agent-Based Artificial Stock Market. Journal of Economic Behavior and Organization. 67(3-4), 702-717 (2008)
12. Marks, R. E.: Analysis and Synthesis: Multi-Agent Systems in the Social Sciences. The Knowledge Engineering Review. 27(2), 123-136 (2012)