Risk Neutral is Best for Risky Decision Making

Robert E. Marks [1]

Economics, UNSW Australia, Sydney,
and University of Melbourne, Australia
Email: robert.marks@gmail.com

Revised 2014-11-11

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) function from Prospect Theory, in which the agent can be risk averse (for gains) or risk preferring (for losses). Analysis of the simulation results reveals that the best function for risky decision-making is the risk-neutral linear function.

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. Indeed, "Risk aversion is one of the most basic assumptions underlying economic behavior" (Szpiro 1997), perhaps because "a dollar that helps us avoid poverty is more valuable than a dollar that helps us become very rich" (Rabin 2000). 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.

Repetition of these choices by many agents allows us use a technique from machine learning — the Genetic Algorithm or GA (Holland 1992) — to search for the best function from each 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 (1997), who argues that the evolutionary learning technique of

---

1. Earlier versions of this paper were presented at the IEEE Computational Intelligence for Finance Engineering & Economics 2014, London, March 29, (Marks 2014a) and the Asia Pacific Symposium on Intelligent and Evolutionary Systems (IES'2014), Singapore, November (Marks 2014b).

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 (fitness) of our agents to conclude that risk neutrality is the best profile.

2. 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). On the other hand, some 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 (Howard 1968). 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 $\bar{x}$.

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

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

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

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

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),

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

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

The ratio U''/U' is proportional to the curvature of the function, which refelcts the risk profile modelled.

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 functions from Prospect Theory.

3.1 CARA Utility Functions

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

$$U(\tilde{x} + \Delta) = U(L') = \sum p_i U(x_i + \Delta). \tag{7}$$

Acceptance of this property restricts possible utility functions to be linear (risk neutral) or exponential, constant-absolute-risk-aversion (CARA) functions (Howard 1968).

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 human 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 here we instead consider the exponential constant absolute risk averse (CARA) functions, where utility U is given by

$$U(x) = 1 - e^{-\gamma x}, \tag{8}$$

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

$$\gamma = -\frac{U''(x)}{U'(x)}\,. \tag{9}$$

From (6) and (9), for exponential utility,

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

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 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

$$\rho(w) = -w \frac{U''(w)}{U'(w)} = w\gamma \tag{10}$$

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 (9).

We use the Constant Elasticity of Substitution (CES) utility function,

$$U(w) = \frac{w^{1-\rho}}{1-\rho}, \tag{11}$$

with positive wealth, $w > 0$, which exhibits constant relative risk aversion CRRA, as in (10).

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 (11). Each agent codes for ρ.

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

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

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, $U(w) \approx \log(w)$.[2] With $w > 0$ and $\rho < 0$, it is equivalent to risk preferring. With $\rho = 0$, the CRRA function is risk neutral.

3.3 The DRP fuunctions from Prospect Theory

From Prospect Theory (Kahneman and Tversky 1979) we model the DRP Value Function, which maps from quantity X to value V with the following two-parameter equations (with $\beta > 0$ and $\delta > 0$):

$$V = \frac{1 - e^{-\beta X}}{1 - e^{-100\beta}}, \quad 0 \le X \le 100 \tag{12}$$

$$V = -\delta \frac{1 - e^{\beta X}}{1 - e^{-100\beta}}, \quad -100 \le X \le 0. \tag{13}$$

The parameter $\beta > 0$ models the curvature of the function, and the parameter $\delta \ge 1$, the asymmetry associated with losses. The DRP function is not wealth-independent.[3]

---

2. Since we are only interested in the ranking of the three lotteries, when $\rho = 1$ we ensure that the argments of the logs are positive.

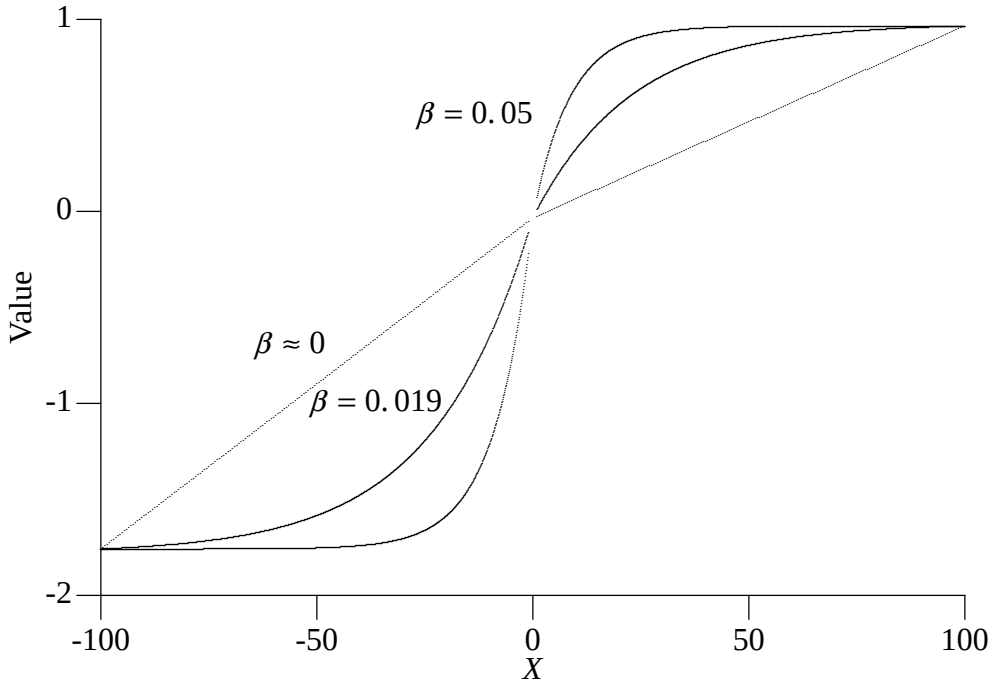

Figure 1: A DRP Function (δ = 1.75).

This function (see Figure 1, with δ = 1.75, for prizes between ± \$100) exhibits the S-shaped asymmetric function postulated by Kahneman and Tversky (1979). 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 Figure 1 demonstrates, as $\delta \rightarrow 1$ and $\beta \rightarrow 0$, the value function asymptotes to a linear, risk-neutral function (in this case with a slope of 1).

4. The Simulations

We employ a numerical simulation method: the Genetic Algorithm (GA). The GA is an optimisation technique used here to search the space of risk-profile parameters of each family of utility functions for those parameters that result in the highest payoff in the process of choosing risky lotteries, as described below. Our prior was that the best parameters would reflect slight risk aversion (at least in the case of CARA and CRRA functions).

The GA also reports the value of the maximand for each set of parameters, which allows the experimenter to compare the relative performances of the best sets of parameters for the CARA, CRRA, and DRP functions. This provides an important insight for our search (below).

---

3. That is, it does not satisfy (7). This does not require that we include wealth w in the ranking of the lotteries, as in the CRRA case; instead we choose a reference point at the current level of wealth, and consider the possible gains and losses of the three lotteries.

Each lottery is randomly constructed: the two payoffs ("prizes") are uniformly chosen in the interval [–$100, +$100], and the probability is chosen uniformly from [0,1]. (Each lottery has, of course, a single degree of freedom for probability.) Each agent chooses the lottery with the highest expected utility of the three from (1) and (8), based on its value of γ (respectively, ρ and wealth w), or from (12) and (13), based on its values of β and δ. To do this, agents know the prizes and probabilities of all three lotteries.

Then the actual (simulated) outcome of the chosen lottery is randomly realised, using its probability. The winnings of the CARA agent or the DRP agent (respectively, the wealth of the CRRA agent) are incremented accordingly. This becomes the agent's "fitness" in the GA search. Each agent successively chooses 1000 lotteries.

4.1 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 1000 choices among the lotteries. Using NetLogo (Wilensky 1999), we use an implementation of the GA (Gilbert 2004) to search for the best risk profile, using each agent's cumulative winnings as its "fitness.".

We model each agent as a binary string which codes to its risk-aversion coefficient (γ, for CARA agents, ρ, for CRRA agents) in the interval ±1.048576. The DRP agents search for $0 \leq \beta < 0.21$ and for δ in the interval ± 10.48. We select the best-performing agents after the 1000 choices to be the "parents" of the next generation of agents, generated by "crossover" (exploiting the genetic information already present in the populations of agents) and "mutation" (exploring the solution space by generating new genetic information) of the chromosomes of the pairs of parents. The "sexual" process of selecting the best-performing agents in the population and then pairing them off to produce the next generation of agents results in a new population inheriting the better characteristics (here, risk profiles) of its parents' generation.

We use the GA simulation in this search as a numerical alternative to solving for the best (highest performing) risk profile analytically. Note that Rabin (2000) asserts that "theory actually predicts virtual risk neutrality." We return to this paper in the Discussion below.

Each agent faces 1000 lottery choices, and its cumulative winnings is that agent's "fitness" for the GA. The processes are stochastic. For each model we perform a number of Monte Carlo simulation runs to obtain sufficient data to analyse the results.

Unlike the GA simulations of Szpiro (1997), we find that the best-performing 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"), while our models allow any risk profile to emerge, we argue that they are more general than Szpiro's.

5. The Results

5.1 The CARA Results

The on-line NetLogo simulations[4] show three things clearly:

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 suggest 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 performed 55 independent Monte Carlo runs using the GA to search for better CARA utility functions. Appendix A.1 presents the results in Table 2. The data suggest that the CARA function has not (yet) converged to risk neutrality. (We reject the null.) The wealth-independent CARA utility function precludes bankruptcy. What of a utility function that does not exclude this possibility?

5.2 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 (11) that exhibit CRRA.

The results are surprising:[5] We have performed 109 independent Monte Carlo runs using the GA to search for better CRRA utility functions. Appendix A.2 presents the results in Table 3. The data suggest that the CRRA function has converged to risk neutrality. (We do not reject the null.)

5.3 The DRP Results

We use the GA to search the joint plane (β, δ) as the agents (each characterised by a point on the (β, δ) plane) choose the one of three lotteries that has the greatest expected value. 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

---

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

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

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

We considered first the marginal results: holding β constant at zero, asking what values of δ emerged as conditionally best, then, holding δ constant at unity, asking what values of β emerged as conditionally best. That is, first considering a kinked function possibly linear, and then a symmetric dual-risk-profile function possibly linear. The results are presented in Appendix A.3 in Table 4. The data suggest that the constrained DRP function has not (yet) converged to risk neutrality. (We reject the null.)

We then performed 50 independent Monte Carlo runs using the GA to search for better β, while holding $\delta = 1$. The data in Table 5 suggest that the constrained DRP function has not (yet) converged to risk neutrality. (We reject the null.)

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. 3 shows the result with the two variables' means, where β = 0.007186 and δ = 1.2598: almost, but not yet, risk neutral (in which an expected-value decision maker becomes an expected-outcome decision maker).[6] The data in Table 6 suggest the unconstrained DRP function has not (yet) converged to risk neutrality. (We reject the null.)

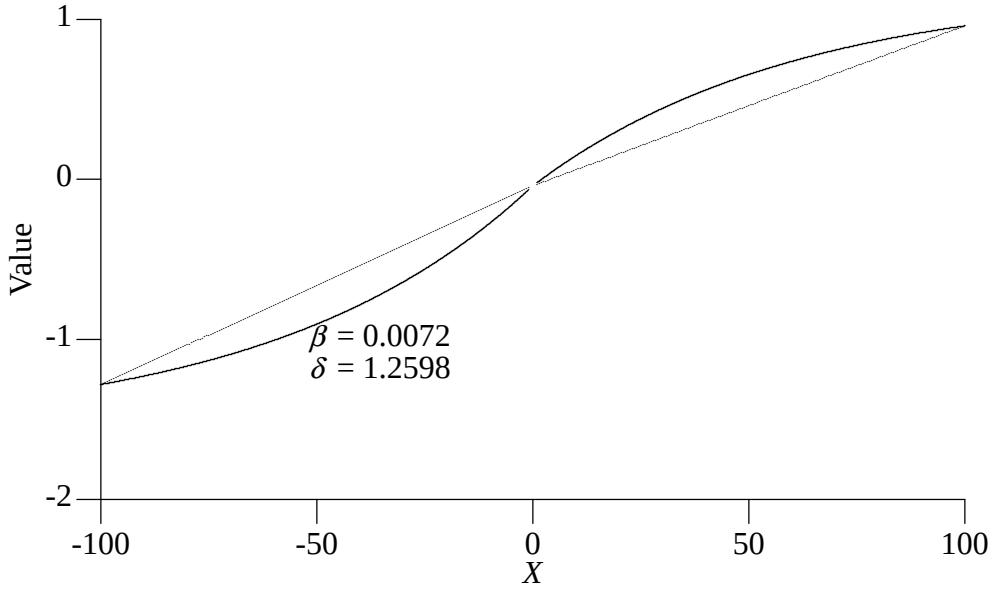

Figure 2: The best DRP function (so far).

6. See http://www.agsm.edu.au/bobm/teaching/SimSS/NetLogo4-models/RA-PTh-EV-both.html for a Java applet and the NetLogo code of a dual-risk Value Function model.

5.4 Comparing Models

As remarked above, all 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 compare the Fitnesses (dollar winnings) across the simulation runs. This is shown in Table 1.

Model	Mean Fitness ($)
CARA	37,650
CRRA	29,403
DRP with $\beta = 0$	37,666
DRP with $\delta = 1$	38,814
DRP joint β, δ	37,721
Linear	39,192
Clairvoyant	50,251

Table 1: Mean Fitnesses of the Models

We include two further models in Table 1: Linear (or risk-neutral) and Clairvoyant.

Clairvoyant is included as benchmark of best possible performance in chosing among the three lotteries. It answers the question: what is the upper limit to agents' performance? If agents could foresee the future, and hence know which three of the six possible prizes would eventuate, then they could choose the lottery which would result in the highest payoff. Although such clairvoyance is impossible in the real world, in our experimental world we build a model in which the realisations of the three lotteries have occurred *before* the agent chooses. The expected value of such choices over the 1000 decisions could be calculated analytically, but Monte Carlo simulations finesse such calculations.[7]

As reported in Table 8, the mean payout across the 50 MC runs, each of 1000 choices, is $50,251, or a mean of $50.25 per choice. Since random choice across the lotteries as constructed would result in a mean of zero, Clairvoyence is valuable. The best of our models results in 77.4% of the Clairvoyant benchmark.

It is clear from Table 1 that CRRA, the only model whose "best" parameter results in risk-neutrality, performs worst in maximizing Fitness, while the best model is the Dual-Risk-Profile 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 optimization, when the constraint 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.

That the GA process in the joint search of the DRP function did not find $\delta = 1$, with its higher fitness, suggests short-circuiting the search and simulating

---

7. The expected maximum value of the choice among the three lotteries, as constructed, can be shown analytically to be $(¾ \times 200) - 100 \geq \$50$.

with a risk-neutral, linear function (that is, DRP with $\delta = 1$ and $\beta = 0$): if the resulting mean fitness is significantly greater than 38,879 then the linear function dominates. As seen in Table 1, the Linear function results in a higher mean fitness than any of the other functions (bar Clairvoyant). The difference is statistically significant.[8]

I am convinced that this shows that the best risk profile for making risky decisions, such as choosing lotteries of known outcomes and probabilities, is the risk-neutral, linear function. In this case its performance is 78.05% of the benchmark Clairvoyant result.[9] The results with the Linear model show clearly that the experiments with the GA and the three families of functions have not (yet) converged, although they are progressing. A comparison of the Fitnesses in Table 1 (with the anomaly that the constrained model of Table 5 below had a higher Fitness then the unconstrained model of Table 6) reveal this lack of convergence and suggest the short cut of jumping to the Linear model. That this mdoel outperforms all the others (bar Clairvoyant) vindicates the short cut and results in the author's conviction that Linear (risk-neutral) is best.

6. Discussion

In the three centuries since the Bernoulli cousins posed and solved the St. Petersberg Paradox (Bernoulli 1738), the idea that a non-linear function of money (or winnings), not simply the amount of money itself, can be used to model decision making under risk, where prospects and probabilities are known, has been been developed. If the Bernoullis saw a concave function, or utility function, as a way of resolving the issue of an infinitely valued lottery, more recently behavioural economists have argued that using non-linear utility functions when faced with risky decisions is a realistic model of human behaviour.

A linear utility function implies zero marginal utility of money or income: an extra dollar is worth an extra dollar no matter what one's wealth. Not a realistic view of human decision making perhaps, but this paper is seeking the optimal utility function in risky decision making, not the most realistic for human decision making.

Should we be surprised that risk neutrality does better than risk aversion? Rabin (2000) 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.

---

8. Null hypothesis: the means of the Linear model and DRP with $\delta = 1$ are identical: but p-value ≈ 0 (Wilcoxen two-tailed), so reject the null.

9. As the number n of lotteries grows, the expected value of the Clairvoyant model approaches the maximum possible outcome ($n/(n+1) \rightarrow 1$).

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 Arthur 1991), we are interested in seeing what is the best (i.e. most profitable) risk profile for agents faced with risky choices.

Rabin (2000) argues that *loss aversion* (Kahneman and Tversky 1979), 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 function, which nonetheless converges on risk neutrality. An analytical study of Prospect Theory DRP Value Functions (DellaVigna and LiCalzi 2001) 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 (Kahneman and Tversky 1979), the agent eventually learns to make risk-neutral choices. Their result is 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 ours.

A simulation study (Chen and Huang 2008) 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).

7. 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, among CARA (or wealth-independent), CRRA (when wealth, and hence bankruptcy, matters), or for the best-performing Value Function, which exhibits the Dual Risk Profile 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.

Against our prior belief that a risk-averse agent does best in these circumstances, we find that only linear, risk-neutral functions perform best. Our findings are therefore consistent with analytical work that proves that with symmetric lotteries, and agents with DRP, risk-neutral decisions are the eventual outcome of agents adjusting their aspirations and targets in response to the realisations of their choices.

Simulations, of course, can not prove necessity, only sufficiency (Marks

2012), 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, converge to risk neutral (linear). Nonetheless, the results in Table 1, I believe, are convincing proof that linear, risk-neutral functions dominate risky decision making. These results therefore disprove the common belief that a small amount of risk aversion is best in a risky world.

Acknowledgments

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, Robert Kohn, Ron Howard, 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 are available online, together with graphical output of the simulation results, as referenced in the 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.

Bibliography

[1] W.B. Arthur, (1991) "Designing economic agents that act like human agents: A behavioral approach to bounded rationality," *American Economic Review Papers & Proceedings*, vol. 81, pp. 353–360.

[2] Daniel Bernoulli, (1738); trans. by Louise Sommer, "Exposition of a New Theory on the Measurement of Risk". *Econometrica* 22 (1): 22–36. (January 1954).

[3] S.-H. Chen and Y.-C. Huang, (2008) "Risk preference, forecasting accuracy and survival dynamics: simulation based on a multi-asset agent-based artificial stock market," *Journal of Economic Behavior and Organizations,* 67(3–4): 702–717.

[4] S. DellaVigna and M. LiCalzi, (2001) "Learning to make risk neutral choices in a symmetric world," *Mathematical Social Sciences*, vol. 41, pp. 19–37.

[5] N. Gilbert, (2005) "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. [Online]. Available:

http://cress.soc.surrey.ac.uk/s4ss/code/NetLogo/axelrod-ipd-ga.html

[6] J.H. Holland, (1992) *Adaptation in Natural and Artificial Systems*, 2nd ed. Cambridge, MA: M.I.T. Press.

[7] R.A. Howard, (1968) "The foundations of decision analysis," *IEEE Trans. on Systems Science and Cybernetics*, vol. ssc-4, pp. 211–219.

[8] D. Kahneman and A. Tversky, (1979) "Prospect theory: an analysis of decision under risk," *Econometrica* vol. 47, pp. 263–291.

[9] R. E. Marks, (2012) "Analysis and synthesis: multi-agent systems in the social sciences", *The Knowledge Engineering Review* 27(2): 123–136.

[10] R.E. Marks, (2014a) "Learning to be risk averse?" In *Proceedings of the 2014 IEEE Computational Intelligence for Finance Engineering & Economics (CIFEr)*, London, March 28–29, ed. by A. Serguieva, D. Maringer, V. Palade, and R.J. Almeida, IEEE Computational Intelligence Society, pp. 1075–1079.

[11] R.E. Marks, (2014b) "Searching for Agents' Best Risk Profiles," In *Proceedings of the 18th Asia Pacific Symposium on Intelligent and Evolutionary Systems (IES'2014)*, Singapore, November.

[12] M. Rabin, (2000) "Risk aversion and expected-utility theory: a calibration theorem," *Econometrica,* vol. 68, pp. 1281–1292.

[13] G.G. Szpiro, (1997) "The emergence of risk aversion," *Complexity* vol. 2, pp. 31–39.

[14] U. Wilensky, (1999) NetLogo. Center for Connected Learning and Computer-Based Modeling. Northwestern University, Evanston, IL. (Version 4). [Online]. Available: http://ccl.northwestern.edu/netlogo.

Appendix A

A.1 CARA simulation results for γ, (H_0: $\mu_\gamma = 0$):

Gamma		Fitness ($)	
Min.:	−0.0338341	Min.:	35122
1st Qu.:	−0.0001213	1st Qu.:	37099
Median:	0.0082934	Median:	37848
Mean:	0.0063766	Mean:	37650
3rd Qu.:	0.0140342	3rd Qu.:	38340
Max.:	0.0328237	Max.:	39260
Std. Dev.:	0.01383	Std. Dev.:	899
z-score:	3.4173		
2-t *p*-val.:	0.00063		

$n = 55$; and reject the null

Table 2: Searching for optimal γ in CARA models.

A.2 CRRA simulation results for ρ, (H_0: $\mu_\rho = 0$):

Rho		Fitness ($)	
Min.:	−412.82	Min.:	26060
1st Qu.:	−101.46	1st Qu.:	28447
Median:	21.60	Median:	29472
Mean:	19.21	Mean:	29403
3rd Qu.:	167.90	3rd Qu.:	30342
Max.:	552.21	Max.:	34073
Std. Dev.:	193.36	Std. Dev.:	1513
z-score:	1.0373		
2-t *p*-val.:	0.2996		

$n = 109$; and do not reject the null

Table 3: Searching for optimal ρ in CRRA models.

A.3 Dual-risk-profile function for δ and β:

Marginal, with $\beta = 0$, (H_0: $\mu_\delta = 1$):

Delta		Fitness ($)	
Min.:	0.8077	Min.:	35156
1st Qu.:	1.1808	1st Qu.:	37224
Median:	1.4444	Median:	37666
Mean:	1.4658	Mean:	37666
3rd Qu.:	1.7110	3rd Qu.:	38516
Max.:	2.3576	Max.:	39154
Std. Dev.:	0.36986	Std. Dev.:	1076
z-score:	9.7547		
2-t *p*-val.:	<0.00001		

$n = 55$; and reject the null

Table 4: Searching for optimal δ in constrained DRP models.

Marginal with $\delta = 1$, $(H_0: \mu_\beta = 0)$:

Beta		Fitness ($)	
Min.:	−0.022032	Min.:	38076
1st Qu.:	−0.003355	1st Qu.:	38712
Median:	0.007424	Median:	38852
Mean:	0.004678	Mean:	38814
3rd Qu.:	0.011231	3rd Qu.:	38958
Max.:	0.038922	Max.:	39293
Std. Dev.:	0.0099341	Std. Dev.:	232
z-score:	3.241949		
2-t *p*-val.:	0.0012		

$n = 102$; and reject the null

Table 5: Searching for optimal β in constrained DRP models.

Searching for both δ and β, separate null hypotheses: $H_0: \mu_\delta = 1$ and $H_0: \mu_\beta = 0$. :

Delta		Beta		Fitness ($)	
Min.:	0.8882	Min.:	−0.020642	Min.:	35979
1st Qu.:	1.0665	1st Qu.:	−0.002544	1st Qu.:	37143
Median:	1.1913	Median:	0.007776	Median:	37754
Mean:	1.2598	Mean:	0.007186	Mean:	37721
3rd Qu.:	1.4031	3rd Qu.:	0.017648	3rd Qu.:	38394
Max.:	2.0548	Max.:	0.026858	Max.:	38896
Std. Dev.:	0.26844	Std. Dev.:	0.012353	Std. Dev.:	751
z-score:	7.11062	*z*-score:	4.27443		
2-t *p*-val.:	< 0.00001	2-t *p*-val.:	< 0.00001		

$n = 54$; and reject both nulls

Table 6: Jointly searching for optimal δ and γ in DRP models.

Linear Model	
Min.:	38839
1st Qu.:	39110
Median	39195
Mean	39192
3rd Qu.:	39291
Max.:	39585
Std. Dev.:	150

$n = 102$

Table 7: The Linear, Risk-Neutral Model

Clairvoyant Model	
Min.:	50008
1st Qu.:	50158
Median:	50247
Mean:	50251
3rd Qu.:	50334
Max.:	50555
Std. Dev.:	119
$n = 50$	

Table 8: The Clairvoyant Model