

Calibrating Methods for Decision Making Under Uncertainty

Robert E. Marks¹

Economics, University of New South Wales Sydney, Sydney NSW 2052, Australia
robert.marks@gmail.com,
WWW home page: <http://www.agsm.edu.au/bobm>

Abstract. This paper uses simulation (written in R) to compare six methods of decision making under uncertainty: the agent must choose one of eight lotteries where the six possible (randomly chosen) outcomes and their probabilities are known for each lottery. Will risk-averse or risk-preferring or other methods result in the highest mean payoff after the uncertainty is resolved and the outcomes known? Methods include max-max, max-min, Laplace, Expected Value, CARA, CRRA, and modified Kahneman-Tversky. The benchmark is Clairvoyance, where the lotteries' outcomes are known in advance; this is possible with simulation. The findings indicate that the highest mean payoff occurs with risk neutrality, contrary to common opinion.

Keywords: decision making, uncertainty, utility functions, simulation, Clairvoyance, risk neutrality

1 Introduction

This is not a descriptive paper. It does not attempt to answer the positive question of how people make decisions under uncertainty. Instead, it attempts to answer the normative question of how best to make decisions under uncertainty. How best to choose among lotteries.

We must first define “best” and “uncertainty”. By “best” we mean decisions that result in the highest payoffs, where the payoffs are the sum of the prizes won across a series of lotteries. The experimental set-up is that each period the agent is presented with eight lotteries, each with six possible known outcomes or prizes (chosen in the range $\pm\$10$). No uncertainty about possible payoffs. But there is uncertainty in each lottery about which payoff or prize will occur. The best information the agent has are the probabilities of the six possible prizes or payoffs in each lottery. Each lottery has six possible payoffs, but the values of these payoffs and their probabilities vary across the eight distinct lotteries. Choosing among these is what we mean by “decision making under uncertainty.”

2 Decision Making under Uncertainty

We model agents as possessing various approaches to this problem.

- A simple approach (the Laplace method) is to ignore any information about the probabilities of payoffs and instead just choose the lottery with the highest average or mean payoff, by calculating the mean of each lottery’s six possible payoffs.
- Another method (modelling an optimistic agent) is to choose the lottery with the highest possible best payoff, the max-max method.
- Modelling a pessimistic agent, another method is to choose the lottery with the highest possible worst payoff, the max-min method. Neither of these methods uses the known probabilities, or even five of the six payoffs.
- A fourth method is to use the known probabilities to choose the lottery with the highest expected payoff, weighting each possible payoff by the probability of its occurring, the Expected Value method.
- Three different families of utility functions.

2.1 Clairvoyance

The so-called Clairvoyant decision maker [1] knows the realisation of any uncertainty, so long as this requires no judgment by the Clairvoyant, and the realisation does not depend on any future action of the Clairvoyant. Here, with simulation of probabilistic outcomes, we can model a Clairvoyant who knows the realised outcome (among the six random possibilities) of each of the eight lotteries, while other decision makers remain ignorant of this. We simulate each outcome as occurring with its (known) probability: only one realised outcome per lottery. The Clairvoyant chooses the lottery with the highest realised outcome of the eight.

We can say something of this: if A_1, \dots, A_n are i.i.d. uniform on $(0,1)$, then $M_n = \max(A_1, \dots, A_n)$ has the expectation of $\frac{n}{n+1}$. Here, $n = 6$ and the expected maximum outcome for any lottery must be $\frac{6}{7} \times 20 - 10 = \7.14 .¹ But the realisation of any lottery is in general less than its maximum outcome, and its simulated realised outcome is generated from the weighted random probability distribution of the six possible outcomes. The Clairvoyant is faced by eight lotteries, and chooses the lottery with the highest simulated *realised* outcome (which the Clairvoyant knows). It turns out (from the simulation) that the expected maximum of these eight realised outcomes is $\$7.788$.² This is the best on average that any decision maker can achieve, given our experimental platform. It is our benchmark.

3 Three Utility Functions

The remaining methods map the known possible payoffs to “utilities,” where the utilities are monotone (but not in general linear) in the dollar amounts of the

¹ The lottery outcomes fall randomly in the range $\pm\$10$; see Section 4.

² With 48 outcomes, the expected maximum outcome across the eight lotteries is $\$9.59$; the expected maximum of the eight simulated realised outcomes is 81.2% of this maximum.

possible payoffs. These methods vary in how the utilities are mapped from the payoffs.

By definition, the utility of a lottery L 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 .

Risk aversion is 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.

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; see equation (2) below);
2. those where the risk preference is a function of the wealth of the decision maker (the Constant Relative Risk Aversion CRRA functions; see equation (5) below); 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. See equations (6) and (7) below.

Since the utility functions are monotone transformations of the possible payoffs, it would be pointless to consider the max-max, max-min, or Laplace methods using utilities instead of payoff values.

3.1 Constant Absolute Risk Aversion, CARA

Using CARA, utility U of payoff x is given by

$$U(x) = 1 - e^{-\gamma x}, \quad (2)$$

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

$$\gamma = -\frac{U''(x)}{U'(x)}. \quad (3)$$

When γ is positive, the function exhibits risk aversion; when γ is negative, risk preferring; and when γ is zero, risk neutrality, which is identical with the Expected Value method.

3.2 Constant Relative Risk Aversion, CRRA

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

$$\rho(w) = -w \frac{U''(w)}{U'(w)} = w\gamma. \quad (4)$$

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 (3).

The Constant Elasticity of Substitution (CES) utility function:

$$U(w) = \frac{w^{1-\rho}}{1-\rho}, \quad (5)$$

with positive wealth, $w > 0$, exhibits constant relative risk aversion CRRA, as in (4). 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 (5). It can be shown that 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 [2], we model the DRP Value Function, which maps from quantity x to value V with the following two-parameter equations (with $\beta > 0$ and $\delta > 0$):

$$V(x) = \frac{1 - e^{-\beta x}}{1 - e^{-100\beta}}, 0 \leq x \leq 100, \quad (6)$$

$$V(x) = -\delta \frac{1 - e^{\beta x}}{1 - e^{-100\beta}}, -100 \leq x < 0. \quad (7)$$

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.³ Three DRP functions in Fig. 1 (with three values of β , and $\delta = 1.75$, for prizes between $\pm\$100$) exhibit the S-shaped asymmetry postulated by Kahneman and Tversky [2]. The DRP function 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 $\delta \rightarrow 1$ and $\beta \rightarrow 0$, the value function asymptotes to a linear, risk-neutral function (in this case with a slope of 1).

³ This does not require that we include wealth w in the ranking of the lotteries, as in CRRA case; instead we choose a reference point at the current level of wealth, and consider the prospective gains and losses of the eight lotteries.




Fig. 1. A Prospect Theory (DRP) Value Function ([3])

4 The Experiments, by Simulation

The experimental set-up is to generate eight lotteries, each with six possible outcomes, each outcome with its own probability of occurrence. The outcomes are chosen from a uniform distribution between +\$10 and -\$10; the probabilities are chosen at random so they add to unity for each lottery. The agent has complete information about the outcomes and their probabilities. Then the agent chooses the “best” lottery, based on the method of choice.

The actual realisation of one of the six possibilities from the chosen lottery is simulated, using the generated probabilities: a payoff with a probability of $0.x$ will be realised on average with a frequency of $100x\%$. The realisation of outcome in the chosen lottery is the agent’s score (in dollars, say). In each iteration, payoff realisations are derived for each of the eight lotteries.

Agents are presented with n iterations of the proceeding choice, and each iteration generates new lotteries with new possible payoffs and new probabilities of the payoffs. The mean payoff over these n choices is the score of the specific decision method being tested.⁴

General opinion is that firms, at least, are better served by slightly risk-averse behaviour. Too risk averse and attractive prospects are ignored (“nothing ventured, nothing gained”), but too risk preferring is the same as gambling, with the risk of losing heavily. What do our simulations tell us about the best method of decision making under uncertainty?

⁴ See the R[4] code at <http://www.agsm.edu.au/bobm/papers/riskmethods.r>

5 Results

Table 1 presents the mean results of 10,000 iterations (independent samples) of the eight lottery/six prize experimental platform, with results for:

1. The benchmark Clairvoyant method
2. the Expected Value method
3. the Laplace method
4. the max-max method
5. the max-min method
6. random choice among the eight lotteries

Table 1. Mean payoffs by method.

Method	Payoff (\$)	% Clairvoyant	% EV
Clairvoyant	7.7880	100	
Expected Value	3.8718	49.7143	100
Laplace	3.3599	43.1425	86.7800
max-max	1.3917	17.8702	35.9500
max-min	2.4279	31.1752	62.7100
Random	0.0216	0	0

The Clairvoyant would have won \$7.79 with perfect foresight. The other methods, of course, cannot see the future, which is the essence of decision making under uncertainty. Expected Value (the risk-neutral decision maker) is second, with 49.7% of the Clairvoyant's score; Laplace is third, with 43.1%. Surprisingly, the (pessimist's) max-min, at 31.2%, is almost twice as good as the (optimist's) max-max, at 17.9%. Unsurprisingly, choosing among the eight lotteries randomly is worst, with effectively a zero mean payoff (of 2.16 cents, or 0.56% of EV).

Table 2 presents the mean results of 10,000 iterations of the CARA method with different values of the risk-aversion coefficient γ : the results show that the best decisions are made when $\gamma \approx 0$, that is when the method is risk neutral and approximates the Expected Value method.

Table 3 present the mean results of 10,000 iterations of the CRRA method with different values of the RRA parameter ρ and reveals that with a CRRA decision maker, again the best profile (the value of ρ that results in the highest expected payoff) is close to zero. That is, as with the CARA method, there is in this set-up no advantage to being risk averse or risk preferring (even a little): the best profile is risk neutrality, as reflected in the Expected Value method. Note that the logarithmic utility method (with $\rho = 1.0$) performs at only 98.88% of the Expected Value method.

Table 4 presents the mean results of 10,000 iterations of twelve DRP functions, combinations of three values of δ and four values of β . The results are

Table 2. CARA mean payoffs, varying γ .

gamma γ	Payoff (\$)	% Clairvoyant	% EV
-0.2000	3.4714	44.5739	89.6600
-0.1600	3.6111	46.3670	93.2669
-0.1200	3.7005	47.5160	95.5782
-0.0800	3.8196	49.0448	98.6532
-0.0400	3.8582	49.5405	99.6503
1×10^{-4}	3.8718	49.7151	100.0016
0.0400	3.8330	49.2163	98.9982
0.0800	3.7840	48.5873	97.7330
0.1200	3.7290	47.8818	96.3138
0.1600	3.6534	46.9111	94.3613
0.2000	3.5615	45.7301	91.9858

Table 3. CRRA, mean payoffs, varying ρ .

rho ρ	Payoff (\$)	% Clairvoyant	% EV
-2.5000	3.7570	48.2408	97.0360
-2.0000	3.8114	48.9391	98.4407
-1.5000	3.8350	49.2426	99.0512
-1.0000	3.8490	49.4222	99.4124
-0.5000	3.8665	49.6475	99.8656
1×10^{-4}	3.8718	49.7143	100
0.5000	3.8577	49.5343	99.6379
1.0000	3.8284	49.1581	98.8812
1.5000	3.8056	48.8655	98.2926
2.0000	3.7773	48.5012	97.5598
2.5000	3.7521	48.1780	96.9098

Table 4. DRP, % of EV, varying δ and β .

beta β	$\delta = 1.001$	$\delta = 1.2$	$\delta = 1.4$
0.0010	100	99.8069	99.4421
0.1000	99.5989	98.5836	98.6017
0.2000	98.2308	97.9890	97.2238
0.4000	96.9848	95.9122	95.2202

the percentages of the EV method. From the mean result for Random in Table 1 (which is +0.56% of EV), we can conclude that the errors in Table 4 are about 1.12% ($\pm 0.56\%$) of EV. Again we see that risk-neutral behaviour, here with $\delta \rightarrow 1$ and $\beta \rightarrow 0$, is the best method for choosing among risky lotteries.

6 Discussion

Whereas there has been much research into reconciling actual human decision making with theory [5], we are interested in seeing what is the best (i.e. most profitable) risk profile for agents faced with risky choices. Rabin [6] argues that loss aversion [2] rather than risk aversion, is a more realistic explanation of how people actually behave when faced with risky decisions. This is captured in our DRP function, which nonetheless favours risk neutrality as a method.

An analytical study of Prospect Theory DRP Value Functions [7] 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 [2], the agent eventually learns to make risk-neutral choices. Their result is consistent with our results.

A simulation study [8] 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 is not what we found (see Table 3 with $\rho = 1$).

Our results here are consistent with earlier work on this topic ([3], [9]) in which we used machine learning (the Genetic Algorithm) to search for agents' best risk profiles in decision making under uncertainty. Our earlier work was in response to [10], which also used machine learning in this search, and which wrongly concluded that risk aversion was the best profile.

7 Conclusion

As economists strive to obtain answers to questions that are not always amenable to calculus-based results, the use of simulation is growing, and answers are being obtained. This paper exemplifies this: the question of which decision-making method gives the highest payoff in cases of uncertainty (where the possible payoffs and their probabilities are known) is not, in general, amenable to closed-form solution. The answer is strongly that risk-neutral methods are best, as exemplified by the Expected Value method. We believe that exploration of other experiments in decision making under uncertainty (with complete information) will confirm the generality of this conclusion. Will relaxing our assumptions of complete information about possible outcomes and their probabilities result in different conclusions? This awaits further research.

8 Acknowledgments

The author thanks Professor Shu-Heng Chen for his encouragement, and discussants of the previous papers ([3], [9]) in this research program. An anonymous reviewer's comments have improved the paper.

References

1. R. A. Howard, "The foundations of decision analysis," *IEEE Trans. on Systems Science and Cybernetics*, vol. ssc-4, pp. 211-219, 1968.
2. D. Kahneman, and A. Tversky, "Prospect theory: an analysis of decision under risk," *Econometrica*, vol. 47, pp. 263-291, 1979.
3. R.E. Marks, "Searching for agents' best risk profiles", In the *Proceedings of the 18th Asia Pacific Symposium on Intelligent and Evolutionary Systems (IES'2014)*, Chapter 24, Volume 1, ed. by H. Handa, M. Ishibuchi, Y.-S. Ong, and K.-C. Tan. In the Series: *Proceedings in Adaptation, Learning and Optimization*, Vol. 1. Springer, pp. 297-309, 2015. <http://www.agsm.edu.au/bobm/papers/marksIES2014.pdf>
4. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. 2013. <http://www.R-project.org/>
5. 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.
6. M. Rabin, "Risk aversion and expected-utility theory: a calibration theorem," *Econometrica*, vol. 68, pp. 1281-1292, 2000.
7. S. DellaVigna and M. LiCalzi, "Learning to make risk neutral choices in a symmetric world," *Mathematical Social Sciences*, vol. 41, pp. 19-37, 2001.
8. S.-H. Chen, and Y.-C. Huang, "Risk preference, forecasting accuracy and survival dynamics: simulation based on a multi-asset agent-based artificial stock market," *Journal of Economic Behavior and Organizations*, vol. 67(3-4), pp. 702-717, 2008.
9. R. E. Marks, "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, 2015.
10. G. G. Szpiro, "The emergence of risk aversion," *Complexity*, vol. 2, pp. 31-39, 1997.