REEXAMINING UTILITY FUNCTIONS FOR NORMATIVE RISK ANALYSIS:

CARA UTILITY AND GAIN AVERSION

by

Robert A. Collins
May 16, 2006

REEXAMINING UTILITY FUNCTIONS FOR NORMATIVE RISK ANALYSIS: CARA UTILITY AND GAIN AVERSION

The expected utility hypothesis has lost credibility as a descriptive theory over the past three decades. Experimental data has shown considerable support for Machina's (1982) assertion that the independence axiom is too restrictive to reflect the actual choices of individuals. Following Kahneman and Tversky (1979), it has been clear that individual choices may also be systematically altered non-rational influences like framing effects. In addition, the loss aversion hypothesis from prospect theory has considerable empirical support and is also not consistent with the expected utility hypothesis. Rabin (2000) suggests that loss aversion may even be a better explanation for modest-scale risk aversion by individuals than expected utility theory. All in all, it seems safe to say that few economists would insist that the expected utility hypothesis is an adequate model to describe how individuals actually behave.

In spite of the developments in positive economics, there is still potential for legitimate normative application of the expected utility hypothesis. For example, business inventory decisions are not likely to be affected by framing. It is also acceptable to assume in the context of purely rational business choices that the independence axiom holds. In situations where the axioms actually apply, the standard for evaluation of risky prospects is still the expected utility hypotheses. However, practical application of expected utility requires a utility function with appropriate properties so that risky prospects are ranked consistently with preferences. It is frequently claimed that a utility function is satisfactory for this purpose if it everywhere reflects non-satiety and risk aversion, which is often interpreted as a positive first derivative and a negative second. Despite the ubiquity of this claim, these restrictions on the signs of the derivatives are not sufficient for producing consistent rankings of risky prospects.

Most normative studies use either the negative exponential (CARA) or power (CRRA) utility function. Even when these functions have a positive slope and concavity, neither satisfies the axiomatic requirements for consistent ranking of risky prospects because neither is bounded from both above and below. While Arrow (1971 page 69) and Samuelson (1977 page 35) both demonstrated convincingly that utility functions must be bounded from both above and below to provide consistent rankings of risky prospects, it seems that this requirement has been almost universally ignored. While discussions of the necessity of boundedness have appeared in the literature for some time, the issue appears to not yet be resolved. Bernoulli (1738) proposed the log utility function to eliminate the so-called St. Petersburg paradox, but Menger (1934) showed that with simple changes in the game, any utility function that did not approach a finite upper bound would produce paradoxical results. While the power function is not bounded from above, many economists believe that gambles with infinite moments are not likely to occur in practical settings, and so it is reasonable to use it in practice even if it technically does not meet the axiomatic requirements for consistent ranking of risky prospects. The negative exponential utility function does have an upper bound eliminating the possibility of a paradox from what Samuelson calls a "Mengerian super game", but the fact that it is bounded only from above causes it to have an even more undesirable characteristic, gain aversion.

This note demonstrates that the commonly used CARA utility function is unsuitable for economic analysis because it does not satisfy the requirement of non-satiety, in a stochastic sense. When the probability of winning becomes sufficiently large, a CARA decision-maker will reduce the size of their bet as the potential gain from winning the gamble gets larger, other things equal. Therefore, the CARA decision-maker exhibits "gain aversion" and avoids a bet with a high probability of a large potential gain as if it were "risk". This is unacceptable as a premise for any economic model. Rejecting a larger potential gain for a smaller one at the same odds implies preferring less to more which violates the non-satiety axiom.

While previously discussed flaws in utility functions arise because of some unimaginable situation like a Mengerian super game, the CARA gain aversion problem arose in a practical real-world setting. In a risk management model for inventory decisions with a CARA utility function, it was observed that when inventory purchases were almost certain to lead to profits, the model gave the bizarre result that inventory levels should be reduced as amount for potential profit increased. After many hours of checking for computational mistakes etc, it eventually became clear that the expected utility model with the CARA utility function was simply producing the wrong answer. The reason is that "gain aversion" is an inherent problem with the CARA utility function.

The proof of the CARA gain aversion assertion is straightforward and requires only elementary mathematics. Suppose the decision-maker is faced with a risky prospect and must decide if they wish to place a bet, and if so how big. The choice for the size of the bet is $X \geq 0$ and the utility of the outcome of the bet (w) is CARA, $U(w) = 1 - e^{-\gamma(w)}$, $\gamma > 0$. The agent either loses their bet, or wins a multiple (G) of their bet:

Outcome(w)	Probability	
-X	p	where $0 < p < 1$
GX	(1 - p)	where $G > 0$

The expected utility of the gamble is a function of the risk aversion parameter γ, the probability of losing p, and the multiple of the bet received for a win G. With CARA utility, the expected utility hypothesis suggests the optimal choice of a bet will maximize:

$$E\left[U(w\{X\})\right] = p\left[1 - e^{-\gamma(-X)}\right] + (1 - p)\left[1 - e^{-\gamma GX}\right].$$

Differentiating w.r.t the choice variable,

$$\frac{dE[U]}{dX} = -p\gamma e^{\gamma X} + (1-p)\gamma G e^{-\gamma G X}.$$

Therefore, the first-order condition for a maximum is:

$$pe^{\gamma X^*} = (1-p)Ge^{-\gamma G X^*}$$

$$\frac{p}{(1-p)G} = e^{-\gamma X^*(1+G)}$$

$$X^* = \frac{\ln\left[\frac{p}{(1-p)G}\right]}{-\gamma(1+G)} = \frac{\ln\left[\frac{(1-p)G}{p}\right]}{\gamma(1+G)}.$$

For a risk averse agent to play the game, the expected value must be greater than zero, which implies G > p/(1-p) or (1-p)G/p >1. Therefore, the numerator of the first-order condition must positive.

The CARA agent exhibits the usual risk aversion choosing a smaller bet as the risk aversion parameter gets bigger:

$$\frac{\partial X^*}{\partial \gamma} = \frac{-(1+G)\ln\left[\frac{(1-p)G}{p}\right]}{\left[\gamma(1+G)\right]^2} = \frac{-\ln\left[\frac{(1-p)G}{p}\right]}{\gamma^2(1+G)} < 0$$

The CARA decision-maker's response to an increase in the probability of a loss is also as expected:

$$\frac{\partial X^*}{\partial p} = \frac{\left[\frac{p}{(1-p)G}\right]\frac{-Gp-(1-p)G}{p^2}\gamma(1+G)}{\left[\gamma(1+G)\right]^2} = \frac{\left[\frac{p}{(1-p)G}\right]\frac{-1}{p^2}}{\gamma(1+G)} = \frac{\frac{-1}{(1-p)Gp}}{\gamma(1+G)} < 0$$

However, the partial effect of an increase in the potential gain from the gamble with risk aversion and the probability of a loss held constant produces an absurd result for the CARA agent.

$$\frac{\partial X^*}{\partial G} = \frac{\left[\frac{p}{(1-p)G}\right]\frac{(1-p)p}{p^2}\gamma(1+G) - \gamma \ln\left[\frac{(1-p)G}{p}\right]}{\left[\gamma(1+G\right]^2} = \frac{\frac{\gamma(1+G)}{G} - \gamma \ln\left[\frac{(1-p)G}{p}\right]}{\left[\gamma(1+G)\right]^2}$$

Since the denominator is squared, the sign of the derivative depends on the sign of the numerator.

$$\frac{\partial X^*}{\partial G} < 0 \Leftrightarrow \frac{(1+G)}{G} < \ln\left[\frac{(1-p)G}{p}\right] \Leftrightarrow \frac{1}{G} e^{\frac{1+G}{G}} < \frac{(1-p)}{p}$$

As the probability of a loss (p) gets small, the right hand side gets large. Applying l'Hôpital's rule we see that as G gets large, the left-hand-side approaches zero. Therefore, when the probability of a large gain gets big, the CARA decision-maker will reduce their bet as the potential gain from the bet gets larger. While it is clear from limiting values that this perverse result can occur, the values required for G and (1-p) are not especially large. For example, G = 1 and a 0.9 probability of winning is sufficient. With a 90% chance of winning, the CARA agent will choose smaller and smaller bets as the size of the potential gain increases. It is clear this result arises because the utility function has an upper bound on utility but no lower bound. Note that this result does not involve any assumption about the risk aversion parameter.

This result makes the CARA utility function unsuitable for risk analysis. While it is implausible that the CARA utility function could have been so widely used for so many decades with such a serious flaw, it seems to be so. In addition, it shows that the issue of boundedness of the utility function is important even for practical risk-management settings and not just for "super games". Since none of the utility functions commonly used pass the test of being bounded from both above and below, it appears that even normative risk analysis needs some new utility functions. As Samuelson (1977) pointed out (p36):

> "Those who are tempted to insist upon both an upper and lower bound for utility must sacrifice the property of utility concavity and general risk aversion everywhere. They gain the property that, for any defined probability distribution of ... outcomes, a transitive ordering is well defined."

An alternative way to state this 29 years later might be: Those who are tempted to insist upon both an upper and lower bound for utility must sacrifice the property of utility concavity and general risk aversion everywhere. They could gain S-shaped utility functions more like those supported by the data of experimental economics and also give well-defined transitive orderings of prospects. Since paradoxical results can occur even in practical settings when utility functions don't meet boundedness requirements, perhaps it is time to carefully reevaluate how the expected utility hypothesis is applied even in the context of normative economics.

REFERENCES

Arrow, K.J. (1971): *Essays in the Theory of Risk-Bearing*. Chicago, IL: Markham Publishing Company.

Bernoulli, D. (1738): "Specimen Theoriae Novae de Mensura Sortis" Translated by L. Sommer (1954) as "Exposition of a New Theory on the Measurement of Risk," *Econometrica,* 22,23-36.

Kahneman, D., and A. Tversky (1979): "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica,* 47, 263-91.

Machina, M..J. (1982): "Expected Utility Theory without the Independence Axiom," *Econometrica,* 50, 277-323.

Rabin, M. (2000): "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica,* 68, 1281-92.

Samuelson, P.A.(1977): "St. Petersburg Paradoxes: Defanged, Dissected, and Historically Described," *Journal of Economic Literature,* 15, 24-55.