Economics Letters 74 (2002) 347–351

economics letters

www.elsevier.com/locate/econbase

A verification of the expected utility calibration theorem

William S. Neilson[a,*], Harold Winter[b]

[a]*Department of Economics, Texas A&M University, College Station, TX 77843, USA*
[b]*Department of Economics, Ohio University, Athens, OH 45701, USA*

Received 3 April 2001; accepted 27 July 2001

Abstract

We investigate whether a single utility function can fit data from both small- and large-scale risks. Utility functions consistent with the wage–fatality risk tradeoff are less risk averse than those consistent with portfolio choice data, as predicted by the expected utility calibration theorem.

Keywords: Risk aversion; Expected utility; Calibration theorem; Value of life

JEL classification: D81

1. Introduction

Over the past 30 years, expected utility theory has come under attack on a number of different fronts, and these attacks have led to changes in the ways that expected utility is used, and even whether it is appropriate to use at all. First, many researchers, including Kahneman and Tversky (1979), showed that the traditional carrier-of-value in the utility function, the individual's final wealth position, is not appropriate; utility should be defined over changes in wealth. Second, Kahneman and Tversky also showed that utility functions are not everywhere risk averse; they are risk averse over gains but risk seeking over losses. Third, a large number of studies show that the independence axiom is violated, and so a more general model should be used.[1] Most recently, Rabin (2000) showed that expected utility suffers from a calibration problem — if the decision maker is averse to a single,

*Corresponding author. Tel.: +1-979-845-7352; fax: +1-979-862-8483.
E-mail addresses: w-neilson@tamu.edu (W.S. Neilson), winter@ohiou.edu (H. Winter).

[1]For a survey, see Starmer (2000).

0165-1765/02/$ – see front matter
PII: S0165-1765(01)00570-5

arbitrarily small risk at every wealth level, he must necessarily be ridiculously averse to large risks.[2] For example, if an expected utility maximizer refuses a 50:50 lose $10/gain $10.10 bet at every wealth level, he must necessarily refuse *any* gamble with a 50% chance of losing $1000 (Rabin and Thaler, 2001).[3] Neilson (2001) extends this result to rank-dependent utility preferences, showing that the result is not an artifact of second-order risk aversion, as defined in Segal and Spivak (1990). Instead, the calibration result stems from the use of a single utility function to govern choices in all settings. The calibration argument is purely theoretical, though, and the purpose of this paper is to determine if existing empirical evidence across two literatures — value of life studies and portfolio choice studies — lend support to Rabin's calibration problem.

Workers regularly accept small wage increases to compensate for a small increase in the probability of dying at work, and there is a large literature that uses the estimates of these wage differentials to calculate a *statistical* value of life.[4] However, these wage differentials can also be interpreted as providing evidence about how averse individuals are to large-scale risks, i.e. risks with potentially large gains or losses. At the other extreme, portfolio choice studies show how risk averse individuals are over smaller-scale risks. If individuals are expected utility maximizers with a single utility function that is used in all circumstances, the two sets of evidence should be compatible. If, however, Rabin's calibration problem holds, the wage premium evidence should be compatible with a less risk averse utility function than the portfolio choice evidence does. This is, in fact, what we find.

2. Empirical calibration attempts

A number of studies have attempted to calibrate the utility function for moderate risks. In particular, there is a large literature that fits constant relative risk averse (CRRA) utility functions to consumption data. A summary of the estimated coefficients is presented in Table 1.[5] All of these studies estimate constant relative risk averse utility functions, but with a variety of data sets, time periods, and methodologies.[6] The results of these studies are found in Table 1, where c denotes the coefficient of relative risk aversion, $c(x) = -xu''(x)/u'(x)$.

Now turn attention to the compensating wage premium. Consider the case of a risk averse, expected

---

[2]Epstein (1992, pp. 11–12) also points out this problem. He shows, through an example, that while risk premia associated with a small actuarially fair gamble might be reasonable and small, those associated with moderate or large stakes gambles are unreasonably large.

[3]The gist of Rabin's argument is as follows. A dislike of a 50:50 lose $10/gain $10.10 bet at every wealth level places restrictions on the concavity of the utility function. From any initial wealth level w, the utility drop from losing $10 must outweigh the utility gain from winning $10.10. Aggregating yields that the utility decreases from losses must accumulate quickly, while utility increases from gains can never get very large. Rabin's analysis yields that the utility decrease from a loss of $1000 is larger than the utility increase from *any* gain, and hence the result.

[4]For a survey, see Viscusi (1993).

[5]The approach of all of these studies is to use a consumption–saving model to predict a relationship between the risk and return on assets and intertemporal consumption behavior. The studies all use consumption data from several periods, look at the risk and return numbers implied by stock market data (NYSE), and then recover the underlying utility parameters that make the consumption data consistent with the risk–return data.

[6]One other parameterization is available: Cicchetti and Dubin (1994) estimate coefficients for a hyperbolic absolute risk aversion (HARA) utility function. With the magnitudes of the payoffs used here, their parameterization is essentially the same as a constant relative risk aversion utility function with coefficient $c = 0.63$.

Table 1
Summary of the estimated coefficients

Study	Consumption data source	Estimated risk aversion coefficient
Friend and Blume (1975)	Household wealth from survey data	$c=2$
Grossman and Shiller (1982)	Aggregate consumption	$c=4$
Hansen and Singleton (1983)	Aggregate consumption	$0.07<c<0.62$
Brown and Gibbons (1985)	Aggregate consumption	$0.1<c<7$
Restricted sample		$2.55<c<2.70$
Litzenberger and Ronn (1986)	Aggregate consumption	$c=4.2$

utility maximizing worker who must decide what job to take. Assume that jobs differ only in their fatality rates and their compensating wage premiums, with the available jobs characterized by the wage–risk schedule $\pi(p)$, where p is the fatality rate and $\pi(p)$ is the corresponding wage premium. The worker selects the job which maximizes his expected utility $(1-p)\,u(L+\pi(p))+pv(D)$, where L is considered to be the *subjective* value of life, D is the subjective value of death, u is the state-dependent utility function corresponding to no fatal accident, and v is the state-dependent utility function corresponding to death.[7] The subjective value of life L includes everything that gives him utility if he lives, including future income, but does not include the compensating wage premium.[8] Maximizing with respect to p and rearranging the first-order condition yields:

$$\pi'(p)=\frac{u(L+\pi(p))-v(D)}{(1-p)\,u'(L)+\pi(p)} \tag{1}$$

Given the state-dependent utility functions u and v and values of p, $\pi(p)$, $\pi'(p)$, and D, it is possible to solve (1) for the subjective value of life L. Estimates of the function π and the function u are available from other studies, as noted below, but there have been no attempts to measure the value or utility of death. It is tempting to normalize $v(D)$ to zero, but, unfortunately, this does not work with constant relative risk averse utility functions.[9] Instead, we assign the utility of death an arbitrary small number, so that $v(D)=u(1000)$, which implicitly assumes that the subjective value of life is at least \$1000.

The values of p, $\pi(p)$, and $\pi'(p)$ we use come from the empirical estimates of Moore and Viscusi (1990), although similar results hold with other wage regression estimates. Moore and Viscusi regress log wages on the number of deaths per 100 000 workers and a vector of personal job characteristics. Their sample has a mean wage of \$7.01 per hour and a mean fatality risk of 7.918 deaths per 100 000

[7] The value of a death state is needed in this paper because it allows the calculation of the subjective value of life with constant relative risk averse utility functions; otherwise, the calculation is impossible. State-dependent utility functions are used because there is no reason to think that the utility of wealth when alive is the same as the utility of wealth remaining after death.

[8] Note that L is not just the worker's income, or even the worker's discounted expected future earnings. For the example used here, the average annual salary is \$14 000, but the subjective value of life for a risk neutral worker taking the average job is almost \$5 million (see Table 2 below).

[9] Letting c equal the coefficient of relative risk aversion, the constant relative risk averse utility function takes the form $u(x)=x^{1-c}$ if $0\le c<1$, $u(x)=\log(x)$ if $c=1$, and $u(x)=-x^{-c}$ if $c>1$. The utility function is undefined at zero if $c\ge 1$.

workers per year, and their regression implies that at $p = 7.918 \times 10^{-5}$, $\pi(p) = \$380$ per year and $\pi'(p) = \$4\,962\,000$.

We wish to conduct the following exercise. Suppose that a worker's behavior is consistent with the average behavior of one of the studies in Table 1; that is, suppose that a worker's preferences over relatively small risks exhibit constant relative risk aversion with one of the parameters listed in Table 1. Further suppose that the worker faces the average wage premium – fatality risk tradeoff estimated by Viscusi and Moore. We can then calculate the worker's subjective value of life. If the calculated subjective value of life is too low, in a sense discussed below, we infer that a single constant relative risk averse utility function cannot explain the worker's behavior in both types of choice situation, consistent with Rabin's calibration result.

The calculated subjective value of life for different values of the coefficient of relative risk aversion are shown in Table 2.[10] The table shows that if a slightly risk averse individual takes the average job, his behavior is consistent with a high subjective value of life. For example, the subjective value of life for a risk averse worker with constant relative risk aversion parameter 0.5 is about $2.5 million. On the other hand, if an individual whose constant relative risk aversion coefficient is greater than 2.6 takes the average job, his behavior is consistent with extremely low subjective values of life. For example, if the relative risk aversion parameter is 4.2, as estimated by Litzenberger and Ronn (1986), the subjective value of life is only $6400.[11]

These results are consistent with Rabin's calibration theorem. According to the calibration theorem, if an individual is averse to small risks at every wealth level, he must be ridiculously averse to large risks. Because the estimate of the wage differential is taken as given, if a more risk-averse worker accepts the same wage differential for the same job as a less risk-averse worker, the calculated value of life must be lower for the more risk averse worker. Thus, a high degree of risk aversion leads to a low calculated subjective value of life, which is what is found in Table 2. For worker's wage–fatality risk choices to be consistent with them treating their lives as valuable, the utility function used to

Table 2
Calculated subjective value of life for different values of the coefficient of relative risk aversion

Coefficient of relative risk aversion (c)	Subjective value of life ($)
0	4 958 700
0.5	2 528 900
1.0	748 600
2.6	13 500
4.2	6400
7.0	3300

[10]The numbers in Table 2 do not represent an effort to construct a value of life estimate that can be used in policy analysis. The proper value of life estimate for policy analysis is the *statistical* value, which is approximately $5 million. The *subjective* value of life estimates in Table 2 are calculated to see whether the utility functions used to generate them are 'reasonable' or not. For moderately high values of c, the subjective value of life estimates are sufficiently low to reject the specification of the preference function. Note that with $c = 0$, the statistical and subjective value of life estimates are equal, as the empirical value of life literature implicitly assumes that workers are risk-neutral.

[11]The figures in Table 2 are robust to changes in D in the following sense. If D is increased by some amount ΔD, the numbers in Table 2 also increase by approximately ΔD.

evaluate large risks must be less risk averse than the utility function used to evaluate small or moderate ones. Further evidence of this fact is provided by Viscusi and Moore (1989), who use a hedonic wage regression to estimate the rate of time preference and the degree of relative risk aversion for workers facing job fatality risks. They estimate a relative risk aversion coefficient of $c=0.7$, which is below most of the estimates from portfolio choice data in Table 1. Thus, a comparison of the wage–fatality risk data and the portfolio choice data provides empirical validation of the expected utility calibration theorem.

3. Conclusion

This paper lends further support to the contention that an expected utility preference representation with a single utility function is unable to describe choices over both small- or moderate-stakes risks and large-stakes risks. We show that the coefficients of constant relative risk aversion compatible with wage–fatality risk premium data are smaller than the coefficients compatible with portfolio choice data, suggesting that utility functions used to evaluate large-stakes risks are less risk averse than those used to evaluate small- or moderate-stakes risks. This result is compatible with the theoretical analysis presented by Rabin (2000).

References

Brown, D.P., Gibbons, M.R., 1985. A simple econometric approach for utility-based asset pricing models. Journal of Finance 40, 359–381.

Cicchetti, C.J., Dubin, J.A., 1994. A microeconometric analysis of risk aversion and the decision to self-insure. Journal of Political Economy 102, 169–186.

Epstein, L.G., 1992. Behavior under risk: Recent developments in theory and applications. In: Laffont, J.J. (Ed.), Advances in Economic Theory. Sixth World Congress, Vol. II. Cambridge University Press, Cambridge.

Friend, I., Blume, M.E., 1975. The demand for risky assets. American Economic Review 65, 900–922.

Grossman, S.J., Shiller, R.J., 1982. Consumption correlatedness and risk measurement in economics with non-traded assets and heterogeneous information. Journal of Financial Economics 10, 195–210.

Hansen, L.P., Singleton, K.J., 1983. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. Journal of Political Economy 91, 249–265.

Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica 47, 263–291.

Litzenberger, R.H., Ronn, E.I., 1986. A utility-based model of common stock price movements. Journal of Finance 41, 67–92.

Moore, M.J., Viscusi, W.K., 1990. Compensation mechanisms for job risks. Princeton University Press, Princeton, N.J.

Neilson, W.S., 2001. Calibration results for rank-dependent expected utility. Economics Bulletin 4, 1–5.

Rabin, M., 2000. Risk aversion and expected-utility theory: a calibration theorem. Econometrica 68, 1281–1292.

Rabin, M., Thaler, R.H., 2001. Anomalies: Risk aversion. Journal of Economic Perspectives 15, 219–232.

Segal, U., Spivak, A., 1990. First order versus second order risk aversion. Journal of Economic Theory 51, 111–125.

Starmer, C., 2000. Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. Journal of Economic Literature 38, 332–382.

Viscusi, W.K., 1993. The value of risks to life and health. Journal of Economic Literature 31, 1212–1946.

Viscusi, W.K., Moore, M.J., 1989. Rates of time preference and valuations of the duration of life. Journal of Public Economics 38, 297–317.