The Journal of Risk and Uncertainty, 31:1; 23–34, 2005

The Value of a Statistical Life and the Coefficient of Relative Risk Aversion

LOUIS KAPLOW
Harvard University and National Bureau of Economic Research, Harvard Law School, Hauser Hall 322, Cambridge, MA 02138, USA

Abstract

Individuals' risk preferences are estimated and employed in a variety of settings, notably including choices in financial, labor, and product markets. Recent work, especially in financial economics, provides estimates of individuals' coefficients of relative risk aversion (R's) in excess of one, and often significantly higher. However, it can be shown that high R's imply equally high values for the income elasticity of the value of a statistical life. Yet estimates of this elasticity, derived from labor and product markets, are in the range of 0.5 to 0.6. Furthermore, it turns out that even an R below one is difficult to reconcile with these elasticity estimates. Thus, there appears to be an important (additional) anomaly involving individuals' risk-taking behavior in different market settings.

Keywords: coefficient of risk aversion, equity premium, risk aversion, value of a statistical life

JEL Classification: D80, G11, G12, I10, J17

Individuals' risk preferences have an important influence on a wide variety of behavior, ranging from portfolio selection to occupational choice. Arrow (1971, pp. 97–98), in his seminal work on behavior under uncertainty, suggested on grounds of boundedness of the utility function that individuals' coefficients of relative risk aversion (R's) should be approximately 1. For a constant-relative-risk-aversion utility function, this implies that utility would depend on the log of income, so that when income doubles, marginal utility is halved. Subsequently, Maitel (1973) offered a range of evidence favoring an estimate of R of approximately 1.5, and other early work supported similarly low values for R.[1]

More recently, a growing body of empirical work, especially in financial economics, provides higher estimates of the value of individuals' R's. Most of this work indicates an R of 2 or more, and some, including that attempting to reconcile the equity risk premium with rational behavior, indicates that individuals' R's may be above 10.[2]

[1]For example, the work surveyed in Mehra and Prescott (1985) in their presentation of the equity premium puzzle finds R's in the range from near 0 to 2.

[2]See, for example, Blake (1996), Brav, Constantinides, and Geczy (2002), Campbell (1996, 2003), Kocherlakota (1996), Mankiw and Zeldes (1991), and Pålsson (1996).

Similarly high estimates are also reported in studies of risk-taking behavior in other markets.[3]

Separate bodies of empirical work study individuals' risk-taking behavior in labor, product, and housing markets, examining compensating wage differentials or individuals' willingness to pay for safer products or for homes in areas with lower environmental hazards. An important strand of this literature measures the value of a statistical life (VSL) based on individuals' tradeoffs between wages or prices on one hand and job or product safety on the other. Researchers have long recognized that VSL depends positively on income. Viscusi and Aldy (2003) survey the large body of pertinent literature and extend it through their own meta-analysis, finding that the income elasticity of VSL is in the range 0.5 to 0.6, with the upper bound of the 95 percent confidence interval falling below 1.0.

It is worth considering whether estimates of individuals' behavior in risky situations are consistent across contexts. A priori, one would suppose that the same individuals, with the same utility functions, should display similar risk-taking behavior in different settings.[4] What does not seem to have been recognized, however, is that *R*'s, estimated by financial economists and others, have particular implications for the income elasticity of VSL, measured in studies of labor and product markets (and conversely).

The relationship between the income elasticity of VSL and *R* is formally derived in Section 1.[5] The fundamental connection thereby identified can be understood intuitively as follows. VSL depends (in significant part) on the marginal utility cost of expenditures to protect one's life. It follows from this that the income elasticity of VSL will depend on how the marginal utility cost of such expenditures varies with income, which is to say on the rate at which the marginal utility of income falls as income rises. An individual's *R* is a measure of this rate.

To illustrate this relationship between *R* and the income elasticity of VSL, consider again the case in which *R* is constant and equal to 1. When utility thus equals the log of income, marginal utility equals the inverse of income. Therefore, if an individual's income doubles, marginal utility and hence the utility cost of a precautionary expenditure that reduces the risk of death falls to half the prior level, so individuals with twice the income would on

[3]See, for example, Barsky et al. (1997) (survey on willingness to gamble on lifetime income) and Mankiw (1985) (consumption spending). Some studies report lower *R*'s. See, for example, Chetty (2003) (evidence on labor supply implies *R* near 1.0), Choi and Menezes (1992) (discussing wide range of estimates in earlier literature, some below one; article examines willingness to accept gambles and suggests the implausibility of *R*'s below one and the probability that they are much greater than one), and Halek and Eisenhauer (2001) (life insurance purchases; median estimate of 0.888 and mean estimate of 3.735).

[4]The relationship between this proposition and results in the various empirical literatures, however, is not straightforward because different literatures implicitly examine utility as a function of different variables and in different settings. For elaboration and a partial reconciliation of disparate estimates of *R*, see Meyer and Meyer (2004). Note 8, below, discusses how this issue is addressed in the present article.

[5]The model (in particular the extension that allows for a bequest motive) can readily be interpreted as applying as well to risks of injury rather than death. Though there is less empirical work on the subject, estimates of the income elasticity of the value of statistical occupational injuries reported in Viscusi and Evans (1990) are in the range 0.6 to 1.0, similar to the values for the income elasticity of VSL, thereby reinforcing the anomaly discussed in the text.

this account be willing to pay twice as much for a given safety measure. This factor taken alone suggests an income elasticity of VSL equal to 1 when *R* is equal to 1; likewise, when *R* equals 2, 10, or some other value, the income elasticity of VSL should on this account equal 2, 10, or that other value, respectively.

There is, however, another effect of income on VSL: The value of preserving one's life is higher when income is higher, because utility is accordingly higher. This suggests that the income elasticity of VSL should tend to exceed *R*.

The derivation in Section 1 shows that indeed the income elasticity of VSL is a sum of these two components (with a weighting on the latter, as will be discussed). This second effect is rather small when *R* is large because a high *R* implies that the marginal utility of income falls rapidly, so total utility rises little as income increases. Accordingly, for the high *R*'s often reported, the income elasticity of VSL would not be much higher than *R*. By contrast, this latter effect is larger when *R* is low because in that case higher income does imply substantially higher levels of utility. As a consequence, very low values of *R* (well below 1) do not necessarily (or even probably) imply that the income elasticity of VSL is below 1.

Taken together, the results indicate that estimates of the income elasticity of VSL in the empirical literature, which are about half the lowest value apparently obtainable from utility maximization, seem difficult to reconcile with rational behavior. Furthermore, the low estimates of the income elasticity of VSL—tightly clustered in the range of 0.5 to 0.6—seem to provide a theoretical upper bound on *R*, which seems quite inconsistent with the high estimates of *R*—at least 2 and perhaps 10 or more—obtained in work on financial economics and in other literatures.

Section 2 pursues extensions and offers further comments aimed toward possible resolution of the identified inconsistency. First, a bequest motive is introduced. In the simplest case, this extension has no effect on the result; more broadly, the possibility of utility from bequests might help to partially explain low estimates of the income elasticity of VSL, but only in the case in which *R* is substantially less than 1. Next, account is taken of the fact that the concrete illustrations in section 1 used an approximation, namely, that expenditures to avert death do not rise with income. However, this factor also has a negligible impact unless *R* is substantially less than 1. Moreover, the magnitude of the appropriate adjustment in such cases appears to be small.

Further discussions relate the foregoing analysis more directly to the pertinent empirical literatures. Regarding estimates of the income elasticity of VSL, possible biases are considered. One involves risk misperceptions, but the likely direction of any bias only exacerbates the discrepancy identified here. Another concerns moral hazard due to life insurance; although this consideration does suggest that estimates may be downward biased, the likely impact seems far too modest to be important. However, some recent literature that examines long-term changes in income and in VSL does indicate somewhat higher elasticities, which would be consistent with an *R* above (though not greatly above) 1. Literature on *R* also provides some basis for reconciliation. Specifically, alternative formulations of utility functions and other features favor lower estimates of *R*, though still well above 1.

Section 3 concludes by restating what appears to be an important anomaly involving individuals' risk-taking behavior in different market settings, one that appears to exist even if one accepts the lower estimates of *R* that researchers find most plausible. Of course, it is hardly novel to suggest that individuals' behavior under uncertainty may not be fully rational and, in particular, may exhibit inconsistencies across different contexts. Nevertheless, the new anomaly identified here seems significant, and its exploration may contribute to the range of efforts that seek to improve our understanding of risk-taking behavior.

1. Analysis

Individuals are assumed to choose expenditures x that reduce the probability $p(x)$ of death, $p'(x) < 0$ and $p''(x) > 0$, to maximize their expected utility

$$U(x) = (1 - p(x))u(y - x), \tag{1}$$

where y is exogenous income and u is a concave function of consumption, $c = y - x$. (The possibility of a bequest motive is considered in Section 2.1.) An individual's first-order condition regarding the choice of x is[6]

$$p' = -\frac{(1 - p)u'}{u}, \tag{2}$$

where $u' = du/dc$.

Expression (2) can readily be transformed into an expression for VSL.

$$VSL = -\frac{1}{p'} = \frac{u}{(1 - p)u'}. \tag{3}$$

The numerator on the right side is simply the utility benefit of saving one's life, and the denominator is the expected marginal utility cost (in units of utility per dollar), so the ratio gives the value of a statistical life (in dollars).

Clearly, because u depends on consumption c, so does VSL. Let η denote the elasticity of VSL with respect to c, that is,

$$\eta = \frac{dVSL}{dc}\frac{c}{VSL}. \tag{4}$$

To analyze this elasticity, we first find

$$\frac{dVSL}{dc} = \frac{(1 - p)u'^2 - u[(1 - p)u'' - u'p'x_c]}{(1 - p)^2u'^2}, \tag{5}$$

where $x_c = dx/dc$ indicates how individuals adjust their investment in reducing the probability of death as their consumption rises.[7] Combining expressions (3)–(5),

$$\eta = c\frac{(1 - p)u'^2 - u[(1 - p)u'' - u'p'x_c]}{(1 - p)uu'} = \frac{cu'}{u} - \frac{cu''}{u'} + \frac{cp'x_c}{1 - p}. \tag{6}$$

[6] The second-order condition holds globally, justifying later substitutions using this first-order condition.

[7] It may be more intuitive to consider how x varies with income, y. It is straightforward to show that $x_c = x_y/(1 - x_y)$.

In the third term on the right side of (6), one can substitute for p' using the first-order condition (2), simplify, and combine with the first term to yield

$$\eta = \frac{cu'}{u}(1 - x_c) - \frac{cu''}{u'}. \tag{7}$$

Examining expression (7), the first component of the first term is the elasticity of u with respect to consumption c, which will be denoted η_{uc}. And the second term is R (of u with respect to c). Accordingly, we can rewrite expression (7) as

$$\eta = \eta_{uc}(1 - x_c) + R. \tag{8}$$

The two terms on the right side of expression (8) correspond to the two effects identified in the introduction. First, as income and therefore consumption rises, the utility value of saving one's life rises at a rate indicated by η_{uc}, the elasticity of u with respect to c. This effect is weighted by $1 - x_c$ because some of one's increased income is spent on safety and this expenditure does not contribute directly to utility. Second, as income and thus consumption rises, the cost, measured in utility, of safety expenditures falls at a rate indicated by R, which is (the negative of) the elasticity of the marginal utility of consumption with respect to consumption.[8]

To assist in interpreting expression (8) quantitatively, it is helpful to make a crude approximation, namely, by assuming that x_c is close to zero. (This approximation is discussed further in Section 2.2.) This allows us to write the income elasticity of VSL as

$$\eta \approx \eta_{uc} + R. \tag{9}$$

For concreteness, it is also useful to consider constant relative risk aversion utility functions. To begin, consider the case in which $R = 1$, that is, $u = \ln c$. For this case, we have

$$\eta \approx \frac{1}{\ln c} + 1. \tag{10}$$

As long as utility is positive, which implies that $c > 1$, the first term is positive. As c becomes large, the term approaches zero. For example, for $c =$ 10,000, $1/(\ln c) \approx 0.11$, and for $c =$ 100,000, $1/(\ln c) \approx 0.09$; thus, $\eta \approx 1.1$. In this simple illustration, our approximation of the

[8] A question arises whether the pertinent utility elasticity and risk-aversion coefficients should be determined with respect to income rather than consumption, and whether one should be examining the traits of u rather than U. The correct choices depend on the implicit time sequence of decisions and realizations concerning, say, job safety and investments and, when examining the different empirical literatures, which model is implicit in the behavior under examination. The version in the text corresponds, for example, to precautions embedded in job choice (as in much of the empirical literature on VSL, in which income depends on a wage that is net of any compensating differential) and subsequent (post-realization of accident outcomes) investments. Kaplow (2003a) considers different combinations (using y and U rather than c and u). In any event, as long as x_c (or, relatedly, x_y) is not very large, the choices make little practical difference.

elasticity of VSL with respect to income is substantially in excess of—about double—the typical magnitudes obtained in the literature of 0.5 to 0.6.[9]

To generalize, consider other constant R utility functions of the form $u = (c^{1-\alpha} - 1)/(1-\alpha)$, where $R = \alpha$ (and $\alpha \neq 1$). Now we have

$$\eta \approx \frac{c^{1-\alpha}}{c^{1-\alpha} - 1}(1-\alpha) + \alpha. \tag{11}$$

For low levels of R, specifically, for $\alpha < 1$, the first component of the first term in (11) approaches 1 (from above) as c becomes large. Therefore, our approximation for η also approaches 1 (from above). Hence, the result—that η is (perhaps slightly) greater than 1.0—is not appreciably different from that in the case in which $R = 1$. (To reinforce the intuition about why low values of R do not imply that $\eta < 1$, consider the case of risk neutrality, in which $\alpha = 0$. The second term in (11) equals zero, but the first term now exceeds 1 (even if only slightly). As consumption rises, the marginal utility cost of expenditures to increase safety does not fall at all, but the utility value of the life one might save is rising roughly proportionately in this case.)

Finally—and of particular interest given the high empirical estimates of R—consider the case in which $\alpha > 1$. For c above 1, $c^{1-\alpha} - 1 < 0$, and $1 - \alpha < 0$; hence, the first term is positive. Further analysis reveals that, as c becomes large, the first term approaches zero (from above).[10] As a consequence, our approximation for η is (perhaps slightly) greater than the value of α. Thus, if one accepts moderate estimates for R, say 2, or much larger estimates, the inconsistency with the estimated value of the income elasticity of VSL (0.5 to 0.6) is striking.

2. Extensions and discussion

2.1. Bequest motive

The analysis in Section 1 assumes that individuals obtain no utility if they die. To incorporate a bequest motive, one can revise expression (1) for expected utility as follows:

$$U(x) = (1 - p(x))u_L(y - x) + p(x)u_D(y - x), \tag{1'}$$

where u_L and u_D are concave state-contingent functions of consumption in the states in which the individual lives and dies, respectively.[11]

It is useful to consider the case in which, for some $\lambda \in [0, 1)$, $u_D = \lambda u_L$. That is, utility in state D, if positive, is taken to be proportional to (and lower than) utility in state

[9] As discussed in Kaplow (2003b), depending on the proper choice of units (an empirical question), much lower values of c might be appropriate to consider, in which case η would be larger.

[10] It is helpful to multiply the numerator and denominator of the first term in (11) by $c^{\alpha-1}$ and to rearrange terms to yield $(\alpha - 1)/(c^{\alpha-1} - 1)$. For $\alpha > 1$, as c becomes large, $c^{\alpha-1} - 1$ approaches infinity, so the first term in (11) approaches zero from above.

[11] This model is similar to that in Viscusi (1994).

L. It is straightforward to demonstrate that, under this assumption, the expression for η is identical to that given in expression (8), restoring our original result.[12] The intuition for this conclusion is that η depends only on how pertinent terms reflecting differences between the two states change as consumption rises; as long as an individual's utilities in the two states are in fixed proportions, the manner in which the relevant differences change with consumption is the same regardless of the particular proportion (including the case in which the proportion is zero, that is, the original case in which there is no bequest motive).

There is, however, no a priori basis for the particular assumption of fixed proportions, and empirical evidence on bequest motives has not been able to pin down even the functional form, much less the pertinent magnitudes, regarding the utility of bequests.[13] It does seem plausible, however, to conjecture that the utility value of consumption in the death state, at least for some individuals, is not subject to as rapid diminishing returns as is utility in the life state. Consider, for example, cases in which the bequest motive is altruistic and involves children, of which there are many, or charities. In this case, there would be a somewhat lower value of η, but as long as we suppose that the marginal utility of income in the life state remains higher than in the death state, the value of R continues to place a lower bound on η.[14] Thus, it is conceivable that bequest motives could help to explain why estimates of η are slightly lower than suggested in Section 1, but unless R is substantially less than 1, it still seems difficult to explain estimates of η in the range 0.5 to 0.6.

2.2. *Assumption that $x_c \approx 0$*

In moving from expression (8) to the approximation in expression (9), the assumption is that x_c roughly equals zero when in fact it is positive—which is intuitive (expenditures on reducing the risk of death should rise with income) and would be straightforward to demonstrate. Thus, the discussion in the latter part of Section 1 overstates the value of η. However, for values of R exceeding one, the entire first term was ignored (on the ground that it approaches zero, from above, in any event[15]), so this approximation is of little consequence in such cases. Likewise, for R equal to one, this term is small, pushing the value of η from one to a slightly higher level.

[12] To see why this must be true, observe that, in the case in question, expression (1′) can be rewritten as

$$U(x) = [1 - (1 - \lambda)p(x)]u_L(y - x).$$

Comparing this expression to expression (1), it is apparent that the only difference is a change of notation: u_L takes the place of u, and $(1 - \lambda)p(x)$, which could be written as $\pi(x)$, takes the place of $p(x)$. For a complete derivation of a version of the more general case, with a demonstration that in this special case the expression for η reduces to that in the original model without bequests, see Kaplow (2003a).

[13] See, for example, Davies (1996), Masson and Pestieau (1997), and Stark (1995).

[14] Formally, in this more general case with bequests, one can show that

$$\eta = c\frac{u'_L - u'_D}{u_L - u_D}(1 - x_c) - c\frac{(1-p)u''_L + pu''_D}{(1-p)u'_L + pu'_D}.$$

Observe that the second term is no longer R for u, but a sort of weighted average (expectation) for this coefficient. (It equals the R for U if x is taken to be fixed.)

[15] If $x_c > 1$, this term would instead approach zero from below.

By contrast, when R is significantly below one, it does seem possible to generate a value of η that is significantly less than one—and thus closer to the range of empirical estimates—if in addition the value of x_c is very large. For example, suppose that $\alpha = 0.5$ and, moreover, that $x_c = 0.5$ (i.e., that of all increases in income, half as much is spent reducing the risk of death as is spent on ordinary consumption). In expression (11), the first term, with $1 - \alpha$, which now equals 0.5, will be weighted by $1 - x_c$, equal to 0.5 in this case, to produce a value of only 0.25. Furthermore, the first component of the first term, as noted previously, will approach 1 for high values of c. Thus, the total value of η in this example would be near 0.75.

Of course, the empirical literature on R's does not support such low estimates of α. Moreover, empirical evidence most closely related to x_c, that concerning how health expenditures rise with income, does not suggest high values of x_c. Estimates reported in Viscusi (1994) indicate values on the order of 0.1 or perhaps smaller. In sum, the approximation made in moving from expression (8) to expression (9) does not appear to contribute significantly to the main result herein.

2.3. *Empirical estimates of η*

The foregoing analysis indicates that, as a matter of rational behavior, we would expect η to be at least 1 and, if empirical estimates of R are to be believed, possibly far greater. This may be juxtaposed against the empirical evidence bearing on η. Notably, Viscusi and Aldy (2003, p. 40) state: "Based on the approximately 50 wage-risk studies from 10 countries, we can conclude from these results that the income elasticity for the value of a statistical life is less than 1.0. Across a number of specifications with our data, our point estimates of the income elasticity range between about 0.5 and 0.6. Note that in none of our specifications did the income elasticity's 95 percent confidence interval upper bound exceed 1.0."[16]

Despite the apparent confidence in these estimates as an econometric matter, perhaps there are aspects of the underlying behavior that produce a substantial downward bias. One possibility is that individuals are imperfectly informed, greatly underestimating the pertinent risks, or that they behave myopically. There is, however, a difficulty with such an explanation: It implies that estimates of VSL itself are biased downward but carries no immediate implication for estimates of the elasticity of VSL with respect to income. Moreover, if one had to make a conjecture, it seems plausible that such biases would become less severe as income rises. That is, individuals who are better informed, better processors of information, and in particular less myopic and thus are more inclined to invest in their own human capital would tend to have higher rather than lower incomes. If so, as income rises, the degree of downward bias in estimated VSL would fall, which would produce empirical estimates of the income elasticity of VSL that were biased *upward* relative to values that reflect rational, fully informed decisionmaking.

[16]More precisely: "For the OLS specifications, the income elasticity varies from 0.49 to 0.60. The 95 percent confidence intervals never range below 0.2 and never exceed 0.95. For the robust regression specifications, the income elasticity varies from 0.46 to 0.48. The 95 percent confidence intervals never fall below 0.15 and never exceed 0.78."

Other possibilities might also be explored. For example, life insurance—by making the state involving death more attractive (or, put more gently, less unattractive)—will produce moral hazard with regard to market decisions, such as occupational choices and product purchases (if insurance premiums do not fully reflect such choices). This too would cause estimates of VSL drawn from observed behavior to be biased downward. In this instance, however, it seems more plausible to imagine that the effect rises with income, if the relative extent (not absolute amount) of insurance coverage rises with income.[17] Then, empirical estimates of the income elasticity of VSL would be biased downward. Even if this is true, it is difficult to imagine a sufficient downward bias to reconcile estimates of η with the high measured values of R.[18]

Perhaps additional explanations could be adduced. Of course, biases in measuring VSL and its elasticity with respect to income are important in their own right. However, the attempt to reconcile such estimates with evidence on high values for *R*—and with a model of rational behavior that strongly suggests that, regardless of R, η should be at least equal to 1—provides further motivation for exploring the issue.

Recently, researchers have begun to exploit an additional source of information on the income elasticity of VSL, examining how VSL—along with income—changes over time. Using this approach, elasticities above 1 are found by Costa and Kahn (2004) for the United States and by Hammitt, Liu, and Liu (2000) for Taiwan, which is consistent with the lower estimates of *R* mentioned in the next subsection. Why these estimates of the income elasticity of VSL differ so much from the consistently lower estimates in prior work has yet to be explored.

2.4. Empirical estimates of R

As noted in the introduction, empirical estimates of *R* have been derived in the financial economics literature—notably in attempts to explain the equity premium puzzle—and in various other literatures. In much work during the past two decades, estimates of *R* exceed 2 and a number of them exceed 10. Some literature has suggested various adjustments to individuals' utility functions and has introduced other features in attempts to explain equity premiums with what many view as more plausible estimates of *R*, namely, single-digit estimates, often closer to 2 rather than substantially higher values.[19] If, indeed, *R* is close to 2, the extent of the discrepancy with estimates of the income elasticity of VSL is greatly reduced, but hardly eliminated. And if those income elasticity estimates are not substantially off, perhaps work in financial economics and other fields should attempt to reconcile their empirical evidence with a value of *R* closer to 1.

[17] Indirect evidence appears in Bernheim et al. (2003), which indicates that the extent of financial vulnerability from the death of a spouse is lower (in part due to life insurance coverage) for individuals in higher income groups.

[18] Considering that estimates of VSL are in the range of $4 to $9 million (Viscusi and Aldy, 2003, reporting on U.S. labor market studies) whereas individuals at the income level of those in such studies typically have life insurance of no more than a few hundred thousand dollars, the moral hazard effect is almost certainly small; hence, the extent to which it might change with income could not plausibly be very significant.

[19] See, for example, Benartzi and Thaler (1995), Brav, Constantinides, and Geczy (2002), Campbell (2003), Constantinides (1990), Epstein and Zin (1990), McGrattan and Prescott (2003), and Shrikhande (1997). For a skeptical view of the success of most such attempts, see Kocherlakota (1996).

Another possibility is that R is not constant at different income levels and the measures of R in the different literatures involve individuals who, on average, are at different levels of income. Specifically, estimates of the income elasticity of VSL tend to be derived from data on ordinary workers who probably have lower incomes than individuals who would be marginal investors. However, there is no strong evidence that R rises with income or wealth; it may even decrease, which would add to the difficulty of reconciling these disparate results.[20]

3. Conclusion

This paper is motivated by the view that theoretical and empirical work on individuals' risk-taking behavior ought to be reconciled across contexts. It turns out that individuals' R's—typically estimated and employed in the study of behavior in financial and some other markets—have direct implications for the income elasticity of VSL—typically examined in the study of behavior in labor and product markets—and conversely. In the model analyzed here, it is shown that the income elasticity of VSL should, roughly, be as high as R (and plausibly higher, at least equal to 1, if R is below 1). Yet estimates of the income elasticity of VSL across a range of studies seem to cluster tightly in the range 0.5 to 0.6 whereas estimates of R tend to be at least 2, with some estimates exceeding 10. This discrepancy suggests the existence of a previously unrecognized (additional) anomaly concerning individuals' risk-taking behavior in different market settings.

This paper raises more questions than it answers. It may be that individuals in fact behave quite inconsistently in different contexts, in which case the causes of this divergence should be explored. Alternatively, the estimates in one of these literatures may be off substantially or may implicitly be measuring somewhat different phenomena. Finally, the present model may not adequately capture some important factor that might explain the gap. In any case, there seems to be value in researchers paying greater attention to the models and empirical results concerning individuals' risk-taking behavior across different contexts.

Acknowledgments

I am grateful to John Campbell, Jack Meyer, Steven Shavell, Andrei Shleifer, Kip Viscusi, and the referees for comments and the John M. Olin Center for Law, Economics, and Business at Harvard University for financial support.

References

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

Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro. (1997). "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics* 112, 537–579.

[20]Pålsson (1996) finds no significant correlation between income or wealth and R and also surveys the mixed evidence in prior literature which tends, perhaps, to favor a view that R does not vary greatly with income. By contrast, Blake (1996) finds that R's are much higher for poorer investors than for rich investors.

Benartzi, Shlomo and Richard H. Thaler. (1995). "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics* 110, 73–92.

Bernheim, B. Douglas, Katherine Grace Carman, Jagadeesh Gokhale, and Laurence J. Kotlikoff. (2003). "Are Life Insurance Holdings Related to Financial Vulnerabilities?" *Economic Inquiry* 41, 531–554.

Blake, David. (1996). "Efficiency, Risk Aversion and Portfolio Insurance: An Analysis of Financial Asset Portfolios Held by Investors in the United Kingdom," *Economic Journal* 106, 1175–1192.

Brav, Alon, George M. Constantinides, and Christopher C. Geczy. (2002). "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," *Journal of Political Economy* 110, 793–824.

Campbell, John Y. (1996). "Understanding Risk and Return," *Journal of Political Economy* 104, 298–345.

Campbell, John Y. (2003). "Consumption-Based Asset Pricing." In G.M. Constantinides, M. Harris, and R.M. Stulz (eds.), *Handbook of the Economics of Finance: Financial Markets and Asset Pricing*, Amsterdam: Elsevier, vol. 1B, pp. 803–887.

Chetty, Raj. (2003). "A New Method of Estimating Risk Aversion." NBER Working Paper 9988.

Choi, E. Kwan and Carmen F. Menezes. (1992). "Is Relative Risk Aversion Greater than One?" *International Review of Economics and Finance* 1, 43–54.

Constantinides, George M. (1990). "Habit Formation: A Resolution of the Equity Premium Puzzle," *Journal of Political Economy* 98, 519–543.

Costa, Dora L. and Matthew E. Kahn. (2004). "Changes in the Value of Life, 1940–1980," *Journal of Risk and Uncertainty* 29, 159–180.

Davies, James B. (1996). "Explaining Intergenerational Transfers." In P. L. Menchik (ed.), *Household and Family Economics*, pp. 47–82. Boston: Kluwer Academic Publishers.

Epstein, Larry G. and Stanley E. Zin. (1990). "First-Order' Risk Aversion and the Equity Premium Puzzle," *Journal of Monetary Economics* 26, 387–407.

Halek, Martin and Joseph G. Eisenhauer. (2001). "Demography of Risk Aversion," *Journal of Risk and Insurance* 68, 1–24.

Hammitt, James K., Jin-Tan Liu, and Jin-Long Liu. (2000). "Survival is a Luxury Good: The Increasing Value of a Statistical Life." (Unpublished manuscript.)

Kaplow, Louis. (2003a). "The Value of a Statistical Life and the Coefficient of Relative Risk Aversion." NBER Working Paper 9852.

Kaplow, Louis. (2003b). "Concavity of Utility, Concavity of Welfare, and Redistribution of Income." NBER Working Paper 10005.

Kocherlakota, Narayana R. (1996). "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature* 34, 42–71.

Maitel, Shlomo. (1973). "Public Goods and Income Distribution: Some Further Results," *Econometrica* 41, 561–568.

Mankiw, N. Gregory. (1985). "Consumer Durables and the Real Interest Rate," *Review of Economics and Statistics* 67, 353–362.

Mankiw, N. Gregory and Stephen P. Zeldes. (1991). "The Consumption of Stockholders and Nonstockholders," *Journal of Financial Economics* 29, 97–112.

Masson, André, and Pierre Pestieau. (1997). "Bequest Motives and Models of Inheritance: A Survey of the Literature." In Guido Erreygers and Toon Vandevelde (eds.), *Is Inheritance Legitimate? Ethical and Economic Aspects of Wealth Transfers*, pp. 54–88. Berlin: Springer-Verlag.

McGrattan, Ellen R. and Edward C. Prescott. (2003). "Average Debt and Equity Returns: Puzzling?" *American Economic Review Papers and Proceedings* 93, 392–397.

Mehra, Rajnish and Edward C. Prescott. (1985). "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15, 145–161.

Meyer, Donald J. and Jack Meyer. (2004). "Relative Risk Aversion: What Do We Know?" (Unpublished manuscript.)

Pålsson, Anne-Marie. (1996). "Does the Degree of Relative Risk Aversion Vary with Household Characteristics?" *Journal of Economic Psychology* 17, 771–787.

Shrikhande, Milind M. (1997). "Non-addictive Habit Formation and the Equity Premium Puzzle," *European Financial Management* 3, 293–319.

Stark, Oded. (1995). *Altruism and Beyond: An Economic Analysis of Transfers and Exchanges within Families and Groups*. Cambridge: Cambridge University Press.

Viscusi, W. Kip. (1994). "Mortality Effects of Regulatory Costs and Policy Evaluation Criteria," *Rand Journal of Economics* 25, 94–109.

Viscusi, W. Kip and Joseph E. Aldy. (2003). "The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World,"*Journal of Risk and Uncertainty* 29, 5–76.

Viscusi, W. Kip and William N. Evans. (1990). "Utility Functions that Depend on Health Status: Estimates and Economic Implications," *American Economic Review* 80, 353–374.