Journal of Risk and Uncertainty, 3: 403–413(1990)

Discounting Statistical Lives

JOHN K. HOROWITZ
Department of Agriculture and Resource Economics, University of Maryland, College Park, Maryland 20742

RICHARD T. CARSON*
Department of Economics, D-008, University of California, San Diego, LaJolla, CA 92093

Key words: discount rate, statistical life, discrete choice

Abstract

Benefit–cost analysis of government projects that reduce health risks over an extended period of time requires an estimate of the *value of a future life*. This in turn requires a discount rate. We suggest and carry out a method to estimate the discount rate using observations on discrete choices between projects with different time horizons. This method is implemented in a survey context. For our primary example, the estimated median discount rate is close to the market rate. A substantial proportion of the sample is estimated to have quite low discount rates. We provide some evidence that discount rates may differ for different types of risks.

Economists have come to accept the idea that placing a dollar value on life is unavoidable in choosing among projects that reduce mortality risks. A related and possibly more controversial question, however, has now surfaced: What is the present value of a risk reduction that occurs in the future? If the value of future risk reductions is to be discounted, what discount rate should be used?

This question arises, for instance, in decisions about when to undertake the cleanup of a toxic waste site, how quickly to phase out the use of a dangerous pesticide, or what kind of sewage treatment plant to build.[1] Environmental and consumer groups have argued that a zero discount rate should be used for public policies that reduce risk. Proponents of this vew are concerned about issues like environmental quality, market failure, and the welfare of future generations. They often believe that programs that reduce long-term health risks address these issues. Thus, they favor a low discount rate for statistical lives since this promotes government action to reduce long-term health risks. The Office of Management and Budget, on the other hand, has adopted the position that a market rate

*This research was supported in part by the U.S. Environmental Protection Agency under Cooperative Agreement in Environmental Economics Research CR-813557-01-0. The opinions expressed do not necessarily reflect the views of the U.S. Environmental Protection Agency. We would like to thank Paul Portney for bringing the importance of discounting statistical lives to our attention, and John Conlisk, Maureen Cropper, Mark Machina, Robert Mitchell, Peter Navarro, and Walter Oi for helpful comments.

is appropriate.[2] Proponents of this view point out that project funds could be invested at a market rate of interest and the future revenue used to buy $(1 + r)$ amounts of risk reduction in the following period. The practical implication of this viewpoint is that less emphasis is placed on reducing long-term health risks in favor of a greater emphasis on programs that can achieve very quick reductions in health risks, even if those programs may save fewer statistical lives in the long run than an alternative program that produces less immediate results. A third viewpoint is that the characteristics of the risks are important and a different discount rate should be used for evaluating each risk. To look at public attitudes toward health risks and to distinguish between these competing views about the appropriate discount rate requires a method that estimates individuals' discount rates and allows for the possibility that their discount rates for different types of risks are different.

In this article, we demonstrate a method to elicit individuals' discount rates for a particular risk. The intuition behind our approach can be seen in the following example. Consider program A, which reduces a particular risk and saves ten lives in the year 1990, and program B, which saves 16 statistical lives from the same risk in the year 1995. Both programs cost the same and must be initiated in 1990. If an agent prefers program A to program B, then his discount rate must be greater than 10%. If program B is preferred, then the agent's discount rate must be less than 10%.

Our method, which can be based either on surveys or on observations of actual behavior, uses agents' choices between policies that differ in the timing and/or number of the statistical lives that are saved. Agents are not required either to put a dollar value on risk reduction or to formulate a discount rate explicitly, both of which can be fairly difficult tasks. Furthermore, the interval in which the agent's discount rate is estimated can be narrowed by changing either the times when lives are saved by the programs or the number of lives that are saved by the programs and then observing choices between the new set of policies. Alternatively, if the set of possible programs varies randomly across subjects, with changes in either the timing or the number of lives saved, then the distribution of discount rates in the population can be estimated without asking an agent to make more than one binary discrete choice.

Because our method estimates discount rates by observing choices between programs that reduce the same type of risk, we do not rely on differences in risks to define the discount rate. This allows the distribution of the public's discount rates to be different for different risks. Treating risks separately is consistent with previous work by Slovic, Fischhoff, and Lichtenstein (1985), which suggests that people view a particular risk as a bundle of characteristics, such as how voluntary or how dreaded it is. In this light, the public's discount rate for a risk might be seen as one more of its characteristics, but one with fairly large implications for public policy decisions. A hypothesis as to whether the median discount is the same for all risks is then testable. In the same vein, other testable hypotheses are whether the median discount rate for a particular risk is equal to zero, whether it is equal to the market rate, and whether a significant fraction of the population holds discount rates that are consistent with being zero or the market rate.

We implement this method using a sample of students and look at three different types of risk. For our sample, we show that the median discount rate for each of the three types of risks can be robustly estimated. Our very limited work suggests that for some risks the median discount rate may be close to the real market discount rate, while for other risks it may not, and that we can reject the hypothesis that the discount rates for all risks are equal.

1. The value of a statistical life

There is a large literature beginning with Schelling (1968) and Mishan (1971) that examines the value of a statistical life and the willingness to pay for decreases in risk.[3] These works led to utility-based theories of the value of life that included both atemporal models (Jones-Lee, 1976) and intertemporal, life-cycle models in which the timing of risks could be considered (Arthur, 1981; Ehrlich and Chuma, 1987). Only recently has the life-cycle utility model been used to derive a theoretical expression for an individual's discount rate for risks (see Rosen, 1988; Horowitz, 1989).

Virtually all the empirical work on the value of risk reductions has considered risks that occur entirely in the present, e.g., accidents (Jones-Lee, Hammerton, and Phillips, 1985), or has used the atemporal model. The few empirical studies that have measured willingness to pay to avoid (only) future risks, using exposure to carcinogens with a long latency period as the source of risk, have not varied the timing dimension and thus cannot be used to estimate a discount rate directly (Mitchell and Carson, 1986; Smith and Desvousges, 1987).[4]

Empirical estimates of the discount rate for risk reductions, like empirical estimates of values of other nonmarketed goods, can be obtained either by inferring them from market transactions or by eliciting them directly through the use of experiments and surveys (Hausman, 1979; Fuchs, 1982). The indirect market approach was recently adopted by Moore and Viscusi (1988; Viscusi and Moore, 1989) using wages and risks from different occupations.

An experiment or survey for estimating discount rates has a number of strengths and weaknesses relative to inference from actual market transactions. Our survey method's strengths are that it can be used to elicit discount rates for different types of risks, not just those for which close market substitutes exist, and it can also explicitly specify the size of the risks, thereby giving all individuals (and the researcher) a common set of information. The potential weakness of survey methods is, of course, that they are based on responses to hypothetical situations. The method itself, however, is not tied only to an experimental or survey context, but could in principle be used with actual observed choices in a market or voting context.

2. A method for estimating discount rates

In this section, we introduce some mathematical notation and show how estimates of the discount rate for the reduction in statistical lives lost to a particular hazard can be estimated. This is done by showing how simple binary discrete choices between programs that save statistical lives over different time horizons reveal information about the population's underlying distribution of discount rates.

2.1. The relationship between discount rates and discrete choice

The value to an individual of a project that saves statistical lives over a period of time can be expressed as

$$\sum_{t=0}^{T} L(t)V_i(t), \tag{1}$$

where $L(t)$ is the number of statistical lives saved in year t, T is the duration of the project, and $V_i(t)$ is individual i's present value of a statistical life saved in year t.[5] Estimates of the value of a statistical life in the current period, $V_0 = V(0)$, are given by Jones-Lee (1976), Thaler and Rosen (1976), and many others (see Fisher, Chestnut, and Violette, 1989, for a summary). The assumption of exponential discounting yields $V_i(t) = V_0/(1 + \delta_i)^t$, where δ_i is the individual's (constant) discount rate.[6]

To estimate an individual's discount rate, and to estimate the distribution of discount rates in the population, we introduce a discrete choice problem. Consider two projects that save statistical lives. One project saves L_1 lives per year for every year from the present year T, until the project expires. The second project saves L_2 lives per year and has the same horizon T, but does not take effect until year $j > 0$. For individual i under the specification $V_i(t) = V_0/(1 + \delta_i)^t$, the values of the two projects are

$$\text{Value of present project} = \sum_{t=0}^{T} L_1 \frac{V_0}{(1+\delta_i)^t} \tag{2a}$$

and

$$\text{Value of future project} = \sum_{t=j>0}^{T} L_2 \frac{V_0}{(1+\delta_i)^t}. \tag{2b}$$

For these two projects we can define an implicit *rate of return* or *equilibrating discount rate* as the discount rate an individual must have if he is indifferent between the two projects. We denote this equilibrating discount rate as δ^*. It is the level of δ such that the value of the present policy is equal to the value of the future policy,

$$\sum_{t=0}^{T} L_1 \frac{V_0}{(1+\delta^*)^t} = \sum_{t=j}^{T} L_2 \frac{V_0}{(1+\delta^*)^t}. \tag{3}$$

In our experiment each participant is asked which of the two projects he prefers, given L_1, L_2, j, and T. A participant who prefers the present policy must have a discount rate δ_1 greater than δ^*, since $\delta_1 > \delta^*$ is equivalent to $\Sigma_{t=0} L_1/(1 + \delta_1)^t > \Sigma_{t=j} L_2/(1 + \delta_1)^t$. A participant who prefers the future policy must have a δ_1 smaller than δ^*. Note that under our specification V_0 can be canceled from the equation. Preference depends only on the relationship between δ^* and δ_1. This allows us to separate discounting behavior from questions about the value of a statistical life.

The parameter δ^* provides either an upper or a lower bound on the individual's discount rate, given his or her choice of a rate greater than or less than δ^*. We exploit this feature in the experiment's design by giving each participant a different pair of projects to choose between—namely, each person's future project used a different value of L_2. Because one-to-one mapping exists between each L_2 and δ^* (given L_1, j, and T), this random assignment of L_2 is equivalent to assigning a different δ^* (say δ_i^*) to each participant. We observe which of the two projects is preferred for the given value of L_2 and thus for the given δ^*. We then estimate a relationship between δ^* and the proportion of the individuals assigned that δ^* who were in favor of the future project.[7] This relationship is an estimate of the cumulative distribution function for discount rates in the population.

2.2. *Estimating the relationship between choice and* δ_i

Suppose that individual discount rates are distributed in the population according to $g(\delta_i) = \mu + h(X_i, \gamma) + \epsilon_i$, where the function $g(\cdot)$ is increasing; the function $h(\cdot)$ is centered around zero with arguments X_i, a vector of individual taste variables, and γ, a vector of parameters; ϵ is a zero-mean disturbance term with variance σ^2; and μ is a location parameter, such as the mean or median, of the distribution of discount rates.

Following the argument presented in the preceding section, we assume that an individual chooses the future policy if his δ_i is less than his assigned δ_i^* or, equivalently, if $g(\delta_i)$ is less than $g(\delta_i^*)$. The probability that a randomly selected individual who faces a choice at δ_i^* chooses the future policy ($CHOICE_i = 1$) is

$$PROB(CHOICE_i = 1 | \delta_i^*) = PROB(g(\delta_i^*) > g(\delta_i)) = PROB(g(\delta_i^*) - \mu - h(X_i,\gamma) > \epsilon_i). \quad (4)$$

This probability P can be estimated in a straightforward manner once we make an assumption about the distribution of ϵ (the subscript i will be suppressed from here on whenever no possibility for confusion arises). For example, if ϵ is distributed normally with standard deviation $\sigma = 1/\beta$, then this probability is

$$P = PROB(CHOICE) = 1 | \delta^*) = \int_{-\infty}^{\beta[g(\delta^*) - \mu - h(X,\gamma)]} \frac{1}{(2\pi)^{1/2}} EXP\left[-\frac{t^2}{2}\right] dt$$
$$= F[\beta g(\delta^*) - \beta\mu - \beta h(X,\gamma)]. \quad (5)$$

where $F(\cdot)$ is the standard normal cumulative distribution function.

We can estimate the parameters in (5) through maximum likelihood methods. This leads to a standard probit regression. A typical feature of probit models is that only the product $\beta\mu$ (i.e., μ/σ) can be estimated, and not the individual parameters β and μ. In the formulation of our model, though, separate estimates of both β and μ are possible as can be seen by examining the likelihood function. Because the treatment variable $g(\delta^*)$ enters linearly with respect to μ, we can estimate β as the coefficient on $g(\delta^*)$ and can estimate μ separately by dividing $\beta\mu$ by β (Cameron and James, 1987). We also note that random assignment of the δ_i^* implies that functions of δ_i^* are independent of $h(X_i,\gamma)$. This means that β and μ can be estimated from either the conditional distribution ($h(\cdot)$ included) or the unconditional distribution ($h(\cdot)$ excluded).

We estimated the parameters of interest under the specification $g(\delta_i) = \delta_i$, which implies that the unconditional distribution of discount rates in the population is $\delta_i = \delta_m + \nu_i$, where δ_m is the unconditional mean (and median) discount rate and $\nu_i = \epsilon_i + h(X_i,\gamma)$. In Horowitz and Carson (1988), we also estimate δ_m using the specification $g(\delta_i) = \ln(\delta_i)$.[8]

3. Experimental design

To estimate individuals' discount rates for statistical lives, we conducted an experiment using the discrete choice setup described above. Seventy-five undergraduates who were enrolled in an upper-division environmental economics class at the University of California, San Diego took part.

The experiment portrayed scenarios for three types of risk. In each scenario there was a choice between two policies, one that saved a given number of lives starting immediately and an alternative policy that saved a possibly different number of lives but that started later. Only one scenario will be detailed here. The policies described in this scenario were improvements in the design of airplanes or airports that reduced the risk of fatal airplane collisions.[9]

The air travel scenario question is as follows:

> Airplane accidents result in a number of deaths each year, mostly from landings and takeoffs. One option to reduce the number of deaths is to require that a safety feature be placed on the existing fleet of airplanes. This feature will save an expected 20 lives per year over the next 15 years (the life of the safety feature). However, the government could instead require that a new radar system be installed at the airport. Construction of the radar system would take about 5 years and would be financed by the airlines. The radar will save about (L_2) lives per year over the next 10 years (year 5 to year 15).
>
> The costs of these two alternatives are the same. But because these costs are large, the government will not require both actions to be taken and it must make a choice about which one to use. Which action should the government require?
>
> _____ Airlines should be required to install safety equipment immediately.
>
> _____ Airlines should be required to construct the radar system, to be ready in 5 years.

The first policy presented in this question involves the immediate installation of a safety feature on airplanes that is expected to prevent $L_1 = 20$ deaths per year for the next $T = 15$ years, beginning immediately, $j = 0$. The second, "future" policy involves the construction of a radar system that will be completed $j = 5$ years from the present and will last for $T - j = 10$ years. This policy is predicated to prevent L_2 deaths per year, with the value for L_2 randomly assigned to each participant. Between one and seven participants were assigned each value of L_2. L_2 ranged from 29 to 54. Because of the unique relationship between L_2 and δ^*, we can refer to each possible L_2 value by its δ^*. The equilibrating discount rate when $L_2 = 32$, for example, is $\delta^* = 0.036$, or 3.6%. The equilibrating discount rates that we randomly assigned in the air travel scenario ranged from −1% to 20% with a mean of 6.6% and a median of 7.5%.[10]

4. Results

Our data for the airplane safety experiment consist of each respondent's choice ($CHOICE_i \epsilon \{0,1\}$) and the δ^* assigned (through L_2). They also include the total number of participants who were assigned each value of δ^* and the proportion, p, of those that favored each policy.

Values of p are plotted against δ^* in figure 1. Results from the probit estimation are reported in (6) with the asymptotic t-statistics given in parentheses:

$$\Phi(\hat{p}) = \underset{(-1.71)}{-0.458} + \underset{(2.97)}{10.082}\,\delta^*. \qquad (6)$$

Here $\Phi(\cdot)$ is the inverse of the standard normal cumulative probability function. Predicted probabilities from (6) are graphed in figure 1.

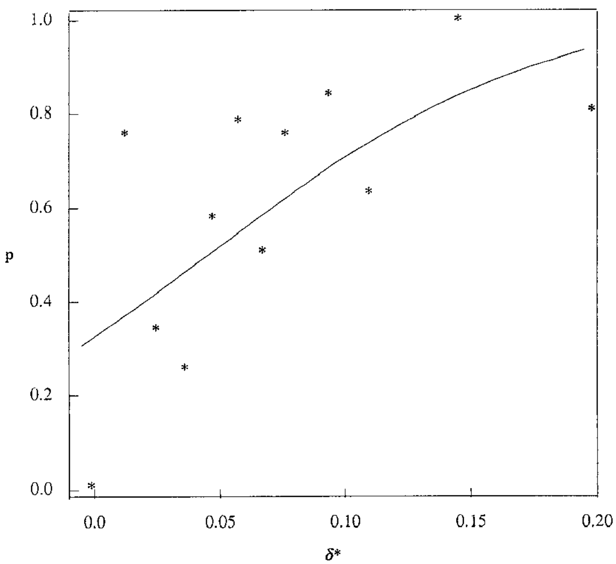

Figure 1. Proportion of individuals choosing the future project as a function of δ^* (air travel scenario). Each point represents the proportion of choices in favor of the future policy of the individuals who were assigned that value of δ^*. The solid line is derived from the relationship estimated in (6): $\Phi(\hat{p}) = -0.458 + 10.982\,\delta^*$.

With δ^* as the independent variable, the estimated median (and by the symmetry of the normal distribution, the estimated mean) discount rate is 4.54% [δ_m = 0.0454(2.91);$\hat{\sigma}_v$ = 0.099(2.46)]. Our null hypothesis was $\delta_m = 0$, which implies that the median individual does not discount future reductions in risk; in other words, future lives are equal in value to present lives. We reject this hypothesis in favor of the alternative that the median discount rate is not equal to zero with $t = 2.91$. A positive median discount rate, however, does not imply positive individual discount rates for all individuals. Under the assumption that the errors are normally distributed, approximately 32% of the population does implicitly espouse a discount rate of zero or less.

The next hypothesis we test is that the discount rate is equal to the market rate, r. A simple measure of the real rate of return is the difference between the nominal rate of return on 25-year treasury bonds and the rate of inflation in the consumer price index at the time of the experiment. This gives a value for r of 5.16%. Using this value, the t-statistic for the hypothesis $\delta_m = r$ is 0.40, and thus we do not reject the hypothesis that the discount rate is equal to the market rate. Although other definitions of the real rate could be used for this test, it is clear that we would not reject the null for a wide range of plausible values for r since the 95% confidence interval for the median (and mean) discount rate in the air travel scenario includes rates from 1.42% to 7.66%.

Looking at the scenarios for the other two types of risks (see Horowitz and Carson, 1988, for details), we uncovered median discount rates of 4.66% for worker safety improvements and 12.8% for traffic safety improvements. These discount rates are significantly greater than zero ($t = 2.44$ and $t = 5.09$), and the second is significantly different from the market interest rate when $r = 5.16\%$ ($t = 3.04$). These results suggest that different discount rates may be appropriate for different risks since we can also statistically reject that the traffic risk has the same discount as the other two risks. The different rates could also have arisen because different combinations of L_2, j, and T were used in each scenario if individuals are not exponential discounters. We were not able to test this hypothesis using only our three scenarios. We also have not examined whether our estimates of the discount rates are sensitive to the particular way in which the questions were framed. The estimation of individual discount rates (as opposed to population median rates) from our data can be achieved by parameterizing the $h(X_i,\gamma)$ function in (4). In contrast to estimating the median discount rate, which can be done nonparametrically and without reference to an individual's characteristics, this generally requires one to make the fairly strong functional form and distributional assumptions typical of the labor market studies of risk. Some empirical results are provided in Horowitz and Carson (1988).

5. Conclusions

This article has proposed an approach for examining the question: What discount rate should be used for future reductions in mortality risks? The distribution of individuals' discount rates was estimated. The approach we used is based on a relationship between an individual's discrete choice between two policies that save statistical lives and the discount rate at which the individual would be indifferent between the two policies. This relationship can be used under fairly general assumptions to estimate median discount rates and, under stronger assumptions, to estimate individual discount rates as well.

Our experimental method is versatile and simple and appears to work well. Our sample of students is probably adequate for making the claim that the median discount rate for risks is greater than zero, but it is less likely to be so for making claims about what discount rate policymakers should use in any particular situation. The next step in addressing this question is to apply the method using a large random sample of the national population. Our method allows this to be accomplished at a fairly low cost.

The experimental method proposed is especially useful because it does not require estimating the value of a statistical life. There are at least two situations in which the discount rate is likely to be a more important factor in policy evaluation than the value of statistical life. First, in many cases, the timing dimension is one of the key aspects of regulation over which government agencies have jurisdiction. This situation arises, for instance, when a specific level of risk reduction has been mandated by legislation but the actual mechanisms to be used to meet this target are not specified. What often remains for the regulatory agency to decide is what technological changes should be required (a decision that simultaneously determines how quickly the technologies will be installed, since some technologies require a longer time to be put in operation or to become commercially available) or at what time a given technology should be adopted. Such decisions involve questions of *when* rather than *whether* to undertake a particular risk-reducing project.

The second situation arises when several risks compete for the regulator's attention and those risks have quite different time profiles. The choice is then which risk to reduce first. One good example is in the regulation of drinking water quality. The U.S. Environmental Protection Agency may be forced to choose between mandating reduced risks of contamination from active biological agents, which poses a risk of immediate illness and mortality, and reduced trihalomethane contamination, which pose a risk of future mortality from cancer. The limited financial resources of the municipalities involved generally prevent the adoption of both projects, even when both appear warranted. This tradeoff too can be seen as a tradeoff in the timing of the risks, rather than a tradeoff of risk for money.

Notes

1. To make this example more concrete, consider, for example, a coastal city that has been undertaking primary treatment of its sewage and that is considering constructing a secondary sewage treatment facility. Secondary sewage treatment would reduce the morbidity and mortality risks to swimmers and consumers of seafood. The city could build one particular type of secondary sewage treatment plant that would begin operation almost immediately, or it could build another type of plant that would be more effective in treating sewage and reducing risks but that would require a much longer time to construct. Whether the construction delay is warranted by the greater risk reduction depends in part on the discount rate for mortality risks. This scenario describes in part the situation faced by the city of Los Angeles in the late 1970s with respect to sewage emptied into Santa Monica Bay.
2. Conflict between these two claims was highlighted during the 1987 nomination of Douglas Ginsburg to the U.S. Supreme Court.
3. Changes in the probability of mortality multiplied by the number of individuals affected gives the number of statistical lives saved by a given project. Statistical lives are a common measure of policy effectiveness in reducing risk when small risks are borne by a large group of people rather than large risks by a few specific individuals (Bailey, 1980).
4. Studies using future risks (Mitchell and Carson, 1986) as well as studies that have used actuarial risks, which combine both present and future risks (Thaler and Rosen, 1976), have tended to estimate lower values for a statistical life than studies in which only immediate risks are considered. This relationship suggests that future risks are discounted.
5. This expression can be derived from a life-cycle model of preferences over personal risks. See Horowitz and Carson (1988).
6. When L is constant over time, the discount rate is, for an individual of a given age, one minus the marginal rate of substitution between risk reductions at times t and $t + 1$, i.e., $V(t)/V(t + 1) - 1$. This definition is different from Moore and Viscusi's (1988) discount rate definition, which is based on how the value of a current risk reduction changes with age.
7. This procedure is analogous to the dose–response procedure used in bioassays (Finney, 1978). In our model, δ^* is the stimulus and the choice between the projects is the response.
8. Note that while the assumptions that ϵ is normally distributed and that $g(\cdot)$ is either linear or logarithmic are convenient for estimation, they are not necessary. Assuming $g(\cdot)$ is known, estimates of the median (and the other central quantiles of the distribution of δ_i) will be relatively robust against reasonable alternative distributions for ϵ. On the other hand, if we assume the distribution of the error terms is known, then it is possible to estimate $g(\cdot)$ and $h(\cdot)$ in (4) semi-parametrically using, for instance, smoothing splines or generalized additive models (Hastie and Tibshirani, 1986). If one wants to avoid making assumptions about the functional forms of $g(\cdot)$ and $h(\cdot)$ and the distribution of ϵ, then a completely nonparametric approach such as the ACE algorithm proposed by Breiman and Friedman (1985) can be used to estimate the entire cumulative conditional or unconditional distribution function for δ_i.
9. The air travel risk scenario was loosely modeled after a well-known midair collision over San Diego between a Pacific Southwest Airlines plane and a small private plane. The other two scenarios described traffic and worker safety policies. The traffic risk scenario involved the choice between putting up a traffic signal or completely remodeling the traffic intersection. This scenario was loosely modeled after a controversy over a major intersection near campus which had taken place several years earlier. The worker safety scenario involved vapors from toxic spills and involved the choice between installing technology that was currently available or waiting for a much improved technology that would be available at a later date. This scenario was intended to reflect the type of choice which often faces EPA and OSHA in this area.
10. A negative equilibrating discount rate can be assigned by making $L_1(T) > L_2(T - j)$.

References

Arthur, W. B. (1981). "The Economics of Risks to Life," *American Economic Review* 71, 54-64.

Bailey, Martin J. (1980). *Reducing Risks to Life: Measurement of the Benefits*. Washington, DC: American Enterprise Institute.

Breiman, Leo and Jerome Friedman. (1985). "Estimating Optimal Transformations for Multiple Regression and Correlation," *Journal of the American Statistical Association* 80, 580–598.

Cameron, Trudy A. and Michelle D. James. (1987). "An Efficient Method for Closed-Ended Contingent Valuation Surveys," *Review of Economics and Statistics* 64, 269–276.

Ehrlich, Isaac and Hiroyuki Chuma. (1987). "The Demand for Life: Theory and Applications." In G. Radnitzky and P. Bernholz (eds.), *Economic Imperialism*. New York: Paragon House.

Finney, David. (1978). *Statistical Methods in Biological Assay*, 3rd edition. New York: McMillan.

Fisher, Ann, Lauraine Chestnut, and Daniel Violette. (1989). "The Value of Reducing Risks of Death: A Note on New Evidence," *Journal of Policy Analysis and Management* 8, 80–100.

Fuchs, Victor. (1982). "Time Preference and Health: An Exploratory Study." In Victor Fuchs (ed.), *Economic Aspects of Health*. Chicago: University of Chicago Press.

Hastie, Trevor and Robert Tibshirani. (1986). "Generalized Additive Models," *Statistical Science* I, 297–309.

Hausman, Jerry A. (1979). "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables," *The Bell Journal of Economics* 10, 33–54.

Horowitz, John and Richard T. Carson. (1988). "Discounting Statistical Lives," Discussion Paper 88-16, Department of Economics, University of California, San Diego.

Horowitz, John. (1989). "An Economic Theory of Latency," Working Paper 89-29, University of Maryland.

Jones-Lee, M. W. (1976). *The Value of Life*. Chicago: University of Chicago Press.

Jones-Lee, M. W., M. Hammerton, and P. Phillips. (1985). "The Value of Safety: Results of a National Sample Survey," *The Economic Journal* 95, 49–73.

Mishan, Ezra J. (1971). "Evaluation of Life and Limb: A Theoretical Approach," *Journal of Political Economy* 79, 687–705.

Mitchell, Robert Cameron and Richard T. Carson. (1986). "Valuing Drinking Water Risk Reductions Using the Contingent Valuation Method: A Methodological Study of Risks from THMs and *Giardia*." Report to the U.S. Environmental Protection Agency.

Moore, Michael J. and W. Kip Viscusi. (1988). "The Quantity-Adjusted Value of Life," *Economic Inquiry* 26, 369–388.

Rosen, Sherwin. (1988). "The Value of Changes in Life Expectancy," *Journal of Risk and Uncertainty* 1, 285–304.

Schelling, Thomas C. (1968). "The Life You Save May Be Your Own." In Samuel B. Chase (ed.), *Problems in Public Expenditure Analysis*. Washington, DC: Brookings Institution.

Smith, V. Kerry and William H. Desvousges. (1987). "An Empirical Analysis of the Economic Value of Risk Changes," *Journal of Political Economy* 95, 89–118.

Slovic, P., B. Fischhoff, and S. Lichtenstein. (1985). "Characterizing Perceived Risk." In C. Hohenemser, R.W. Kates, and J. Kasperson (eds.), *Perilous Progress: Managing the Hazards of Technology*. London: Westview Press.

Thaler, Richard and Sherwin Rosen. (1976). "The Value of Saving a Life: Evidence from the Labor Market." In N. Terleckyj (ed.), *Household Production and Consumption*. New York: Columbia University Press.

Viscusi, W. Kip and Michael J. Moore. (1989). "Rates of Time Preference and Valuations of the Duration of Life," *Journal of Public Economics* 38, 297–317.