1

How Pricing Lives Saves Lives

The Challenge of Valuing Mortality Risks

"We can't do that. That's immoral." This was the reaction I got in 1980 when I suggested to a prominent Occupational Safety and Health Administration (OSHA) official that the agency monetize the reduced risks of death from job safety regulations using labor market estimates of workers' valuation of fatality risks. The values I advocated were based on the extra amounts workers are paid for each expected workplace death. Early studies often referred to these figures as the value of life, but the terminology used to describe this approach has evolved to be the "value of a statistical life," which is both more accurate and somewhat less inflammatory.

The idea of monetizing the benefit of reduced worker fatality risks was not controversial. OSHA and other agencies had routinely attached dollar values to the expected lives that would be saved by regulations. In doing so, they followed the general approach patterned after that used by the courts in wrongful death cases, in which they equated the benefit of reduced risks of death to the value of lost earnings and medical expenses. This formulation enabled agencies to generate a mortality risk reduction benefit number and to be able to point to the use of a similar approach by the courts, thus providing

some evidence of its reasonableness. However, there is a fundamental disconnect between these values and the core principle underlying benefit assessment, which is that benefit values for government policies should reflect society's willingness to pay for the benefit. An instructive way to ascertain these values is to examine the revealed preferences based on risk-taking behavior. The estimates of the value of a statistical life make such a connection by using the value that the workers themselves place on risks of death.

The benefit valuation amounts per expected fatality prevented that I suggested to OSHA were quite substantial, far in excess of the figures OSHA was using to value reductions in mortality risks. So the net effect of adopting my numbers would be to make lives considerably more valuable than the agency's practice at that time, a consequence that hardly seemed "immoral." I was also personally disappointed in their negative reaction, having just published the book edition of my doctoral dissertation on job safety,[1] as well as related articles advocating the value of a statistical life approach. But in addition to my personal stake in the methodology, I viewed valuing risks of death based on how workers themselves valued these risks as the only correct economic procedure for producing sensible policy valuations.

The idea of conceptualizing the valuation task in terms of statistical lives had a firm economic basis and was due to Nobel laureate Thomas Schelling.[2] However, given the state of empirical methods at the time of his original analysis, he was not optimistic about the prospects of assessing this value based on either surveys or empirical work: "The main problem is that people have difficulty knowing what it is worth to themselves, cannot easily answer questions about it, and may object to being asked. Market evidence is unlikely to reveal much." Fortunately, advances over the past half-century in available data, statistical methods, and survey procedures have enabled us to develop meaningful estimates of the value of a statistical life.

At the time of my original suggestion that OSHA adopt this methodology, I was on leave from my position at the Northwestern University economics department to the government, where I served as the Deputy Director of the President's Council on Wage and Price

Stability. This agency managed the pay-price guidelines and was responsible for oversight of all proposed federal regulations for the Carter administration. Under the Carter administration's Executive Order No. 12044, agencies were required to assess the benefits and costs of all major rules and to show that the chosen regulatory approach was the most cost-effective alternative. Although agencies were required to analyze the benefits and costs of policies and ideally should monetize these benefits and costs, the executive order did not require agencies to demonstrate formally that the benefits of the proposed regulations exceeded the costs. That requirement arrived in 1981 with the Reagan administration, which moved the regulatory oversight staff to the US Office of Management and Budget (OMB), Office of Information and Regulatory Affairs. Nevertheless, even without a formal benefit-cost test, there were political pressures within the Carter administration to strike a sensible balance between calculated benefits and costs, in part because of the perceived potential inflationary impact of costly but ineffective regulations.

The shift of the regulatory oversight group to OMB by President Reagan in 1981 was accompanied by other policy changes. As part of its economic reform agenda, the Reagan administration had ramped up the emphasis on bringing regulatory costs under control. The Carter administration had initiated a variety of deregulation initiatives with respect to airlines, trucking, and banking. Although the economic rationales for risk and environmental regulation are quite different than for economic regulations, as is the potential justification for deregulation as well, the Reagan administration sought to extend the deregulation concept to these newer social regulations. As part of this effort, the Reagan administration's Executive Order No. 12291 imposed a benefit-cost test for proposed major regulations. Agencies must show that the benefits exceeded the costs before being permitted to issue the regulation. This benefit-cost analysis requirement has remained in place through all subsequent administrations. The Reagan administration also strengthened the regulatory oversight process. The requirement that agencies obtain OMB approval before launching major regulatory initiatives replaced the advisory White House reviews under Carter. The Reagan

administration also targeted many regulations for elimination, such as a host of recent safety regulations pertaining to the auto industry.

The Triumph of the Value of a Statistical Life Approach: The Hazard Communication Policy Debate

In this era of deregulation, agencies nevertheless continued to develop new regulatory proposals.[3] The most expensive major new initiative proposed in the early years of President Reagan's first term was the OSHA hazard communication regulation, which OSHA proposed in 1982. If this regulation was enacted, for the first time there would be regulatory requirements that firms label dangerous chemicals used in the workplace. Since these chemicals are often considerably more potent than household chemicals, the absence of any such labeling regulation more than a decade after OSHA's establishment was surprising. After having read these proposed chemical labels, workers would be aware of their chemical risk exposures and be able to take appropriate precautions, possibly including the decision to quit and seek safer employment. In addition, the regulation would require that firms maintain material safety data sheets so that if workers were exposed to a dangerous chemical, the medical personnel would be aware of the consequences of the exposure and know how to treat the worker.

OSHA's regulatory impact analysis for the hazard communication standards tallied the prospective costs and expected improvements in worker health and attached a dollar benefit to these health effects. The dominant benefit component for this regulation, as it is for most other health, safety, and environmental regulations, was the value of the mortality risk reduction. However, instead of using the value of a statistical life to monetize these effects, at that time OSHA and other agencies used the value of medical costs and lost earnings, or what they termed the "cost of death." After completing the analysis based on the costs of the deaths prevented by the regulation, OSHA submitted the proposed regulation to the OMB regulatory oversight group for the approval that was required before the regulation could be issued. Although OSHA's evaluation concluded that

the regulation was desirable on balance, its economic assessment was flawed in several respects. Based on the critique of the regulatory impact analysis by the OMB economists, if OSHA had done a proper analysis, the result would have been that costs exceeded the benefits so that the regulation failed a benefit-cost test. As a result, OMB rejected OSHA's regulatory proposal.

OSHA nevertheless wished to pursue the possibility of issuing the regulation. The procedure that the Reagan administration had established for agency appeals was that in the event of a dispute, the regulatory agency could appeal the decision to then Vice President George H. W. Bush. The vice president characterized the disagreement as a technical economics dispute and suggested that an outside expert assess the merits of the competing arguments. I was asked to resolve the dispute after being approved by the Secretary of Labor and OMB. By that time, I had left the government and was then at Duke University, where my research continued to focus on risk regulation issues. The OMB regulatory oversight staff, most of whom were my former colleagues at the Council on Wage and Price Stability, raised a host of criticisms of OSHA's benefit estimates. All of these critiques were well founded. The result of adopting the OMB corrections to the analysis was that the calculated costs of the regulation exceeded the estimated benefits in terms of improved worker safety. In my assessment of the competing agency arguments, I accepted all of OMB's critiques as being sound.

Where my approach differed from that taken by both OSHA and OMB was with respect to how the expected lives saved would be valued. Both OSHA and OMB valued the lives saved based on what OSHA termed the "cost of death," or the present value of the medical costs and lost earnings that would be saved by preventing workers from being killed by chemical exposures. Under this approach, lives would have a value of several hundred thousand dollars, which was not a trivial amount, but was an arbitrary accounting measure that bore little relationship to how workers themselves valued risks of death. The approach I suggested utilized my labor market estimates of how much compensation workers required to face small risks of death. My estimate of the value of a statistical life at that time was

$3 million per expected fatality prevented, or about $7.4 million adjusted for inflation. More recent estimates of the value of a statistical life generally place its value at between $9 million and $11 million. Using this estimate in the regulatory benefits analysis instead of the cost-of-death approach boosted benefits by an order of magnitude. Following a similar approach, I also attached values to the prevention of nonfatal worker injuries, but the driving force in the benefit assessment was the value of the fatalities prevented by the proposed regulation. The result of abandoning the cost-of-death approach was that benefits now exceeded the costs so that the hazard communication regulation would now pass OMB's economic test. President Reagan approved the regulation almost immediately after my report in support of the regulation reached the White House.

The Genesis of Estimates of the Value of a Statistical Life

Where did this $3 million figure come from? The average annual worker fatality risk at that time was 1/10,000, which is more than double the current level of dangerousness. In return for bearing this risk, workers received an annual wage premium of $300, where this amount was estimated statistically controlling for other aspects of the job and worker characteristics. The result is that for a group of 10,000 workers, on average one of them would be killed on the job in the coming year. The amount of compensation that this group of 10,000 workers would receive for the one expected death is 10,000 × $300, or $3 million. Thus, the value of a statistical life is simply the total amount of compensation required per expected workplace death. The value of a statistical life reflects the values that the workers themselves believe that bearing these risks is worth rather than an accounting measure or an arbitrary number assigned by a government analyst.

But why should any finite value be applied to the expected lives saved? One could treat each expected life saved as having an infinite value. In that case, it would be desirable to expend the entire federal budget on safety measures that would eliminate a small chance of even one expected death. Given the multiple risks that we face

and the limits on our financial resources, this uncompromising approach is infeasible. Consider data for 2014, the most recent year for which comprehensive accident data are available. There were 136,053 accidental deaths in the United States in 2014.[4] If the entire gross domestic product of $17.4 trillion in 2014 were allocated to preventing accidents, it would only be possible to spend an average of $128 million per death to prevent these accidents, leaving nothing left to prevent illnesses or to provide for daily living expenses.

To motivate the reasonableness of using workers' wage-risk tradeoffs as the guide, it is useful to ask people to conceptualize scenarios in which the reader makes similar decisions that do not reflect an unbounded commitment to safety, whether it involves living in a riskier but less expensive neighborhood or driving a car that doesn't have all possible safety enhancements. Most transportation choices and dietary decisions entail at least some risk. We do not plan our lives to minimize all possible risks. Through daily risk-taking decisions, people reveal that they place a finite value on reduced risks to their life. People's unwillingness to display an unbounded commitment to safety is consistent with a myriad of other risk-taking decisions that we make. The labor market estimates undertake a similar comparison in which the tradeoff is based on the extra wages that workers are paid for the additional risks posed by their jobs.

Critiques of My Initial Estimates of the Value of a Statistical Life

My estimates of the value of a statistical life on the order of $3 million came under attack from both extremes. Some critiques suggested that the numbers were too big, while others thought they were too small. The critiques suggesting that they were too small involved appeals to value lives at an infinite amount and were not grounded on any actual empirical estimates. Those suggesting that my estimates were too large were more empirically based. My figure exceeded the present value of lifetime earnings, so how could it be feasible for people to value their lives so highly? Surely people could not afford to pay more for their lives than their resources permitted. The raw

numbers presumably suggested that the valuations were excessive. However, what workers are valuing is not the certainty of death but rather a very small risk of death of 1/10,000 annually. It should also be noted that workers are being compensated for these small risks. They are not paying for greater safety so the budgetary constraint issues are less salient. Moreover, for small risks of death, budgetary issues are not influential whether we are talking about the amount that workers are paid to incur a risk or the amount that they would be willing to pay to reduce their level of risk. The amounts that workers require to face slightly greater risks will equal the amount that they are willing to pay to make their jobs safer to that same degree. Even if workers were paying for the risk reduction, it is not unrealistic to assume that workers would be willing to pay $300 to reduce their annual fatality risk by 1/10,000. This is a quite different matter than assuming that workers have $3 million to buy out of the risk of certain death.

A more sophisticated critique was based on the economics literature at that time. Some other competing estimates of the value of a statistical life pegged the value at well below $1 million. Didn't that lower value imply that my numbers were too high? The disparity in the estimates also gave the false impression that estimates of the value of a statistical life differed widely and were too unreliable to be used for policy analysis. However, the observed differences were quite plausible. Whereas workers in my studies faced risks that were comparable to the US average of 1/10,000 annually, the estimates of under $1 million were based on workers in very high-risk jobs with annual fatality rates of 1/1,000 per year.[5] These divergent estimates are not incompatible. The value of a statistical life is not a universal constant. Rather, it reflects the average rate of tradeoff between wages and risk for particular samples of workers. Those who place a comparatively low value on risk to their lives will tend to gravitate to higher-risk jobs and reveal through their choices a lower wage requirement per unit of risk, which will imply a lower value of a statistical life.

Notwithstanding such critiques and the political sensitivity of the task of valuing risks to life, other agencies also adopted the value

of a statistical life, or the VSL. While it may be the case that agencies were swayed by the compelling economic logic of the VSL approach, it is also likely that the fact that the VSL enabled agencies to boost the calculated economic benefit estimates by a factor of ten contributes to its attractiveness. The mortality risk reduction benefit valuation approach using the VSL is now the standard practice throughout the US government as well as in many other countries. The largest benefit component of all US federal regulations is the monetized value of the statistical lives that will be saved by the regulation.[6] The VSL has become the most important parameter driving the attractiveness of government regulations generally as it plays the central role in the evaluation of health, safety, and environmental regulations.[7]

Why Monetizing the Effects Matters

The monetization of the reduced mortality risks through application of the value of a statistical life is instrumental in the assessment of the benefits by the government, as it enables these effects to be put in the same terms as other economic impacts. The benefit-cost analysis procedure that lies at the heart of regulatory analyses involves a comparison of the benefits and the costs and a judgment that the benefits exceed the costs. To make such a comparison, at some point all effects must be put in comparable units, at least implicitly. Cost figures are dollar amounts that appear to be real economic consequences. Indeed, the regulatory oversight efforts in the Ford and Carter administrations were motivated primarily by a concern with the economic burdens arising from the inflationary effects of regulatory costs, not a concern with benefits or benefit-cost balance. Monetization of risk reduction benefits puts these effects on the same footing as the cost numbers, making clear that they are just as real economic effects as are regulatory costs.

Instead of assigning an explicit finite value to the expected fatality risks, another possible benefits approach is not to monetize the fatality reduction benefits at all. Rather, one can simply note that a policy will lead to a reduction in the risk of death for some

population and that these expected risk reductions are valuable but are not monetized effects. One might advocate this nonmonetized approach based on the rationale that monetizing the expected lives saved in effect denigrates them by treating lives as a commodity.[8] The danger of this attempt to skirt around the use of monetary values is that effects that are not monetized may be treated as being less consequential than the purportedly real economic costs and benefits that have a monetized market value.

In an earlier study of the use of benefit-cost analysis to value water-resource projects, I found that government agencies placed almost exclusive emphasis on the monetized effects.[9] In that era, environmental consequences received qualitative discussion but were largely set aside in favor of emphasis on the series of tangible economic benefits that were monetized. As a result, somewhat oddly the traditional evaluation of water-resource projects placed negligible emphasis on nonmonetized ecological consequences, and instead focused on monetized impacts such as the value of the irrigation water that would be sold to farmers and the value of increased municipal and industrial water supplies. Monetized benefits such as the irrigation water counted, but qualitative environmental consequences did not. If the risk reduction effects are not monetized, the greater likelihood is that they will be treated as having zero or negligible value, not that they will be viewed as being more consequential than if the value of a statistical life numbers were not used.

Failing to monetize the effects also limits the ability of benefit-cost analysis to provide a comprehensive index of the policy's attractiveness. By converting benefits into monetary terms, it is possible to compare the benefits and costs to assess the net attractiveness of the policy. Ultimately, any policy decision will implicitly make such a comparison even if the benefits are not monetized, so the failure to monetize risks disguises the hard choices being made but does not avoid the task of setting an implicit value on expected lives saved. Approval of a policy that costs $100 million and expects to save fifty lives implies that the government values lives at least at $2 million. Similarly, failure to move forward with a policy that costs $100 million but expects to save only ten lives indicates that the government

values expected lives saved at less than $10 million. In instances in which there are multiple attributes that are not monetized, such as lives, cancer cases, and disabling injuries, it becomes more difficult to impute the specific implicit value of a statistical life that is incorporated in the policy decision. But whether the valuation is undertaken implicitly or explicitly, there is still some valuation judgment being made.

Exploiting the Potential of the Value of a Statistical Life

While the idea of pricing lives may seem to be morally reprehensible, the practice of valuing lives and risks to life arises either explicitly or implicitly in a variety of contexts in government policies, corporate decisions, and compensation awarded by the courts. The principal recurring theme of this book is that US institutions and practices throughout the world have consistently undervalued lives. Properly confronting issues regarding pricing lives may offend moral sensitivities but will ultimately lead to more protective policies than those that emerge when the monetization of risks to life is suppressed as being too grisly an enterprise to contemplate, much less implement in actual risk decisions.

Over the past three decades, the VSL has emerged as the most influential economic parameter in the evaluation of federal regulatory policies. All federal agencies have adopted the methodology to ascertain the benefits of different regulatory efforts and the appropriate levels of stringency of these regulations. My examination of the role of the VSL explores a variety of issues that arise in policy contexts in which the use of the VSL approach has become well established. Should, for example, reduced risks of life be valued at the same level for everyone? Should income levels matter? Are reduced risks to the lives of older people less valuable than the lives of children, who have much greater life expectancy? How should we conceptualize such issues pertaining to the heterogeneity of the VSL and tailor the appropriate policy approach? Are there any groups in the workforce, such as Mexican immigrants, who appear to be left out of this ideally functioning economic market, and what are the

ramifications of this shortcoming in market performance? These are not always easy questions to answer, but available research addresses many of these issues.

The task of setting a price on life arises in many other contexts as well. My focus is on pricing lives more generally, not simply on valuing statistical lives for government policies that are directed at reducing mortality risks. Assigning either an explicit or implicit price for lives is an integral aspect of safety decisions by firms and is a routine component of court awards in wrongful death cases. Does the VSL provide more meaningful guidance in these situations? At one extreme, might it serve as an all-purpose number for valuing lives? Or should the application of the VSL in the courtroom be more focused on particular situations?

The most direct extension of the use of the VSL is to corporate decisions regarding safety. Private companies must decide on the safety features of their products, and in doing so, they are valuing risks to life just as are government agencies. While companies have not generally adopted the VSL approach, should they use the VSL estimates in their safety decisions? The failure of firms to utilize the VSL may seem to be puzzling. After all, the VSL numbers are generated by decisions made in markets. The VSL estimates from the labor market reflect the joint influence of decisions by firms to offer wages for jobs at different levels of risk and the decisions by workers to accept such jobs. The economic mechanism works analogously for products, as the VSL reflected in product choices simultaneously reflects the cost to the firm of increasing product safety as well as how much consumers themselves are willing to spend for safer products. Why have companies failed to adopt the VSL, which serves as a measure of how much their consumers will value additional safety features? Application of the VSL creates tremendous potential for providing a sounder basis for corporate risk decisions that will also lead to greater levels of safety.

But there are pitfalls companies face that might account for their failure to adopt the VSL. Using the VSL may expose corporations to greater liability threats than if they didn't undertake a thorough risk analysis. Just as government agencies periodically come under

fire for valuing lives, corporations may be subject to similar and even more costly critiques. However, unlike government agencies, corporations are not shielded from potential liability costs. Much like government agencies before the adoption of the VSL, corporations have a history of shortchanging safety. Establishing a legal structure that enables companies to undertake responsible risk analyses is essential. Applying the VSL to corporate behavior as well as to government policies establishes a consistent and more protective safety framework across societal institutions.

In addition to the prospective safety decisions by government agencies and corporations, there are also situations in which mortality risks must be valued retrospectively. Government regulators levy fines on firms when workers are killed on the job as well as penalties on firms that generate product and environmental hazards that lead to fatalities. Courts make damages awards to companies in wrongful death cases ranging from routine automobile accident cases to mass tort actions, such as for pharmaceutical products. Both the regulatory sanctions and the court awards entail the pricing of identified lives after the fact. In neither instance is the VSL the standard approach now used to value these deaths. As I will indicate, the VSL does have a pivotal role to play for government regulatory sanctions but a more nuanced role in the courts, for which the principal concern is addressing the economic losses incurred by the accident victims. In these instances as well as in valuing prospective risks, the VSL establishes safety guideposts that will lead to more protective outcomes than current approaches.

Risk Awareness and Beliefs

The application of the VSL estimates appears to offer enormous potential dividends in a variety of contents, but is there any reason to believe the numbers are credible? Do workers even understand the risks posed by their jobs? This is an important question, since the VSL estimates are based on economic frameworks that assume that workers are aware of the risks they face. An assumption of fully informed workers seems unreasonable for hazardous exposures that

cannot be readily monitored and that lead to illnesses not captured in accident statistics. However, the VSL estimates are based on acute fatality risks generally caused by accidents rather than deferred risks associated with occupational illnesses. Basing the VSL approach on a framework that presupposes awareness of the risk is reasonable for the types of hazards captured in labor market studies, which are the more visible risks, such as motor vehicle hazards and risks of fires and falls, rather than less apparent risks, such as the cancer risks of asbestos.

The evidence that workers' risk beliefs are plausible is quite strong.[10] While the economic theory does not require that every worker can identify the job risks with pinpoint accuracy, workers' subjective risk perceptions do follow the expected patterns. Workers are more likely to believe their jobs expose them to dangerous or unhealthy conditions if they work in industries with higher reported accident rates. Whereas only 24 percent of workers in low-risk industries view their jobs as dangerous, all workers in the highest-risk industries believe that their jobs are dangerous. The most prominent hazards cited by workers are inherently dangerous materials, inherently hazardous equipment, inherently hazardous procedures, dangerous exposure to dust, and transportation hazards. A particularly striking result is that for my survey of hundreds of chemical workers described below, workers' perceptions of the accident rate in their jobs were equal to the objective measure of the average injury risk in the chemical industry.

The sources of information about risks are quite diverse. Many hazards, such as the dangers of high-rise construction work, are readily apparent. Moreover, over two-fifths of all workplace fatalities are related to motor vehicles, which pose quite familiar risks. Workers also learn through their on-the-job experiences. Workers who have experienced an on-the-job injury on their current job are more likely to regard their jobs as being dangerous, as 71 percent of those who have experienced such an injury view their jobs as dangerous. These patterns of risk beliefs are consistent with the notion that people learn about risks in a sensible way. If the wages paid for these perceived risks are inadequate, they may quit. My estimates have

found that one-third of all manufacturing quit rates are attributable to this learning response to job risks. Moreover, after quitting these jobs, workers switch to positions for which they receive appropriate risk compensation in line with the VSL estimates.

How do estimates of the wage premiums workers receive for subjectively assessed risks compare with economists' estimates derived using objective risk measures in the statistical analysis of compensation for risk? There are in fact strong parallels in terms of the wage-premium effects implied by objective measures of the risk and subjective worker assessments. The wage-premium estimates for risk that I obtained using fatality risk measures conditional on workers perceiving that their jobs exposed them to dangerous or unhealthy conditions were the same as the values obtained with empirical estimates based on objective industry risk measures. Also, if workers are completely ignorant of the hazards, there will be no market mechanism that generates wage premiums for risks, so the existence of observed wage differentials provides additional evidence of some underlying risk awareness. Situations in which there are gaps in the provision of such risk compensation, such as for immigrant workers in dangerous jobs, serve as a red flag that there is something amiss and more vigilant government regulation might be beneficial.

Linkage to Hazard Communication

It is particularly noteworthy that the OSHA hazard communication regulation that led to the adoption of VSL is closely linked to a linchpin of the economic theory used to derive the VSL estimates, which is that workers will demand compensation for the risks they face. If workers are unaware of the risks, as is possible when facing dimly understood chemical hazards, they will not require extra wages to face these hazards. Communicating information about the risks through a labeling effort is consequently a pivotal mechanism for fostering the kind of market response that underlies the VSL approach.

But what effect would such labeling regulations actually have? The regulatory impact analyses of the proposed hazard communication

standard by OSHA and OMB were based largely on hypothesized possible impacts and plausible conjectures. The 1980s was a period in which there was substantial interest in the comparatively novel informational regulations, but little was known about their efficacy. Other than cigarette warnings (in place in the United States since 1966), warnings for prescription drugs, and warnings for a short list of highly dangerous chemicals such as hydrochloric acid, there was much less reliance on informational approaches than is the norm today. OSHA's hazard communication regulation was an innovative policy, but it was not the only informational effort in that era. Other informational initiatives followed. Congress passed legislation requiring new rotating cigarette warnings in the mid-1980s, and in that decade the US Environmental Protection Agency (EPA) also embarked on various right-to-know efforts for chemical exposures and toxic environmental releases. However, despite the flurry of policy activity, agencies did not have a firm grasp of the efficacy of the warnings policies. In many instances, policymakers regarded informational efforts as a holding action until the pertinent policy response became evident or political obstacles to more stringent regulation could be overcome.

At the time of my evaluation of the OSHA hazard communication regulatory proposal, there was no firm empirical basis that could be used to ascertain the likely efficacy of the regulation. In collaboration with a chemical industry official, I designed a survey with an experimental structure to explore how workers would respond to the warnings.[11] This survey also illuminated how workers perceive and respond to risks, providing a detailed exploration of the formation of workers' risk beliefs and how these beliefs translate into the wages they require for the job. The survey involved showing different warnings to hundreds of workers at several major chemical plants. The products ranged from very dangerous chemicals, such as asbestos, to fairly innocuous materials, such as sodium bicarbonate (household baking soda). If the worker had to work with the warning on the chemical label instead of the chemicals at the worker's current job, would that affect the assessment of the risks of the job? The results were remarkably strong. Before being shown any warnings, workers were asked to

assess the injury risks of their jobs. The average response was exactly equal to the reported injury rate in the chemical industry, which is a number that they were not told. These remarkably accurate risk beliefs were consistent with workers being cognizant of acute accident risks, as the theory underlying the VSL assumes. The study told the workers that the new chemicals that had the warning label would replace those chemicals currently used in their jobs. Giving workers information about potentially dangerous chemicals boosted their risk beliefs for chemicals that were more dangerous than the chemicals that the workers were currently exposed to on their jobs.

This change in risk beliefs also had the expected effects on workers' attitudes toward their jobs. If workers were not offered higher wages for greatly increased risks, they indicated that they would be unwilling to take the job again or they would be willing to quit their position. In addition, the wage rate that they required to face greater risks from the hypothetical chemical exposures were in line with existing labor market estimates of wage premiums for objective measures of injury risks. This study consequently documented the reasonableness of worker risk beliefs and the broad tradeoff that workers must make in terms of the discrete decision to work on a particular job based on whether the wage and job-risk package made the job as attractive as alternative positions.

The chemical worker study addressed many key aspects of individuals' responses to hazard warnings, particularly those related to risk beliefs and wage responses, but did not focus on the pivotal regulatory issue, which was whether there would be an increase in precautionary behaviors that the worker would take when using the chemicals. Whether warnings would also foster safety-related behaviors also was a particular concern for the EPA, which had in place a series of regulations regarding chemicals and pesticides used by consumers and workers. The agency wanted to ascertain whether it was sufficient to utilize warning labels for dangerous chemicals and pesticides to promote safe handling of risky chemicals rather than banning these chemicals or greatly restricting their use, such as requiring that users be certified pesticide applicators. The tradeoffs in this instance involved risk-averting choices the individual

makes, conditional on use of the product, rather than the discrete decision to use a product to take a job. Would the consumer take appropriate precautions such as wearing rubber gloves and mixing herbicides with water at the recommended levels? For that agency I directed a series of studies of how best to design chemical and pesticide warnings.[12] In this household-chemical situation, the tradeoff was not in terms of wages and risk. Risk and precautionary efforts were often the focus, as consumers had to choose the precautions they would take with the product. In other studies that we undertook, consumers could pay a higher price for a safer product that required fewer precautions during use or that was safer generally, which is more analogous to the wage-job risk situation. Thus, there are two dimensions by which warnings could exert their influence. They could affect a consumer's discrete decision to use a product at all, in much the same way as hazard warnings could affect whether the worker would take a job, or they could influence how consumers choose to use the product. We found that properly designed warning labels could be effective in influencing each of these dimensions.[13]

The underlying principle for understanding such informational efforts is that their function is intrinsically related to the economic principles pertaining to the VSL. The tradeoffs embodied in decentralized decisions by consumers and workers involve decisions to engage in potentially risky activities, the precautions they take in such activities, and the other aspects of the activity that govern its overall attractiveness, such as the job's wage rate and the product's price. Informational efforts seek to make these decisions function in a sound way. Similarly, the use of the VSL serves as a mechanism for replicating the guidance for how policies would be structured if they were designed based on the patterns of behavior reflected in informed consumer and worker decisions.

The Organization of This Book

My exploration of the issues pertaining to pricing lives begins with the fundamentals. Chapter 2 considers the key components of government practices. What is the most reliable evidence for the VSL,

and how have agencies incorporated these values in their policy assessments? The dominant approach in the United States has been to rely on labor market evidence, but survey approaches also have a role to play to address risks such as cancer and in countries such as the United Kingdom where labor market studies are less reliable. US estimates of the VSL based on labor market evidence are greater than those using other techniques, such as asking people in surveys how much they value risks.

The basic elements of the VSL involving setting a monetary value on the worth of risks to life should extend beyond government policies to corporate risk decisions, such as whether more costly safety devices should be added to the product design. Corporations' failure to adopt the VSL and the causes of their reluctance to undertake meaningful risk analyses are the subject of chapter 3. Historically, companies have been vilified for undertaking risk analyses that strive to balance costs and safety, but this impediment to responsible risk analyses could be addressed by providing appropriate legal protections for corporations undertaking responsible risk analysis. The attendant liability costs in the absence of such protections have led to the suppression of meaningful risk analyses, as exemplified in the policy practices at General Motors (GM) with respect to defective ignition switches. This case study of corporate neglect dramatizes the dangers of shunting explicit risk-cost tradeoff issues to the side. Banishing explicit discussion of the underlying risk-cost concerns led to over one hundred deaths. Chapter 4 takes an in-depth look at the GM ignition switch defect experience and the potential transformation of corporate practices that would be achieved through the use of the VSL estimates. This case study also illuminates the shortfall in government policies that fail to incorporate the VSL in setting regulatory sanctions for corporations' safety infractions.

While the basic components of the VSL approach suffice for providing guidance in most instances, many modifications of the approach may be desirable. Chief among the potential refinements in the use of the VSL estimates is recognition of possible heterogeneity in these values. Chapter 5 examines differences by age, as there is substantial variation across society in remaining life expectancy.

Should risks to the lives of people with very short life expectancy receive the same value as lives of the very young? Should any such quantity adjustments guide our thinking with respect to setting the price on reducing risk, setting levels of compensation, or levying penalties after fatalities have occurred, and are the appropriate age adjustments the same in each context? If we are striving for a more equitable approach, does that mean that all lives are valued equally or that each year of life should have the same value? While posing such questions raises potentially troubling moral issues, framing the answer to such queries to inquire how much the person values risks to his or her own life mollifies many of these concerns. Moreover, the fact that the VSL does not plummet with age and people place a considerable value on risks to each year of life makes the VSL a more protective way to address these issues than some current practices.

Basing the VSL on people's valuation of the risk might serve to disadvantage those who have lower income levels, both within countries and across the world. Chapter 6 considers differences by income, and chapter 7 examines related risk equity issues generally. Whose lives should be saved, and what is the most insightful way to conceptualize what we mean by equity when dealing with risks? There are potential differences in the VSL estimates along these and other dimensions. In some cases, the differences are not as great as one might think. In many other cases, the differences are not critical because government policies, environmental pollution, and consumer products generate risks that often cut across a broad swath of the population so that reliance on average valuations is reasonable. Nevertheless, there are some circumstances in which heterogeneity in valuations matters and should be taken into account. The pertinent VSL that should be used in very low-income nations is less than what is used in the United States. Nevertheless, use of the US VSL after appropriate adjustments for income differences across countries will lead to greater valuation of mortality risks than the practices in those countries or in the practices advocated by international bodies such as the Organisation for Economic Co-operation and Development (OECD), an international organization of 35 countries, most of which are advanced nations.

That this is the case indicates the systematic undervaluation of life throughout the world.

Although the basic components in the VSL literature are clear, there are several ongoing areas of refinements and professional debates, which are considered in chapter 8. An ongoing issue is how illnesses with substantial morbidity components should be addressed. Labor market studies provide insight into the valuation of acute accident risks, but how can we derive valuations for mortality risks such as cancer and other dread diseases? Should cancer risks be accorded a premium, and if so, how much? That chapter also considers issues such as whether the VSL estimates are distorted by publication selection biases in terms of what values are selected to be submitted to journals and published in the literature.

Our understanding of the role of morbidity effects and related complications also has ramifications for how the VSL can best be incorporated in reform of judicial practices with respect to determinations of liability and setting of damages amounts, which is the focus of chapter 9. The VSL does have a legitimate role to play in each instance but in a manner not currently used by the courts. The variations with respect to considerations such as age and income arise in judicial contexts as well, but these are not new issues for the courts. Indeed, the courts treat matters such as income levels and age in a starker fashion than current VSL practices.

Despite the potential controversies and misconceptions that the VSL generates, basing risk determinations on the risk valuations of those exposed to the risk serves as the most meaningful and responsible guide. Refinements in the VSL levels used for policy will remain a worthwhile endeavor. However, the most pressing tasks are to extend the use of the VSL to additional domains and to strike an efficient balance between risk and cost for a wide range of societal decisions. The VSL can serve as the guidepost for establishing appropriate safety levels as well as setting the level of financial sanctions needed to promote safety. Tremendous opportunities remain for incorporating the VSL as an integral part of society's efforts to address health, safety, and environmental risks. Corporations, government agencies, and the courts all should give the VSL a more

prominent role. Previously, government agencies sought to disguise their valuation of lives saved by relabeling it the "cost of death." Corporations have sought to suppress such decisions altogether as being off-limits. The main consequence of these failures to value fatality risks properly is that safety levels fall short and lives are being lost.