I Z A

IZA DP No. 4508

Policy Relevant Heterogeneity in the Value of Statistical Life: New Evidence from Panel Data Quantile Regressions

Thomas J. Kniesner
W. Kip Viscusi
James P. Ziliak

October 2009

Policy Relevant Heterogeneity in the Value of Statistical Life: New Evidence from Panel Data Quantile Regressions

Thomas J. Kniesner
Syracuse University
and IZA

W. Kip Viscusi
Vanderbilt University

James P. Ziliak
University of Kentucky

ABSTRACT

Policy Relevant Heterogeneity in the Value of Statistical Life: New Evidence from Panel Data Quantile Regressions[*]

We examine differences in the value of statistical life (VSL) across potential wage levels in panel data using quantile regressions with intercept heterogeneity. Latent heterogeneity is econometrically important and affects the estimated VSL. Our findings indicate that a reasonable average cost per expected life saved cut-off for health and safety regulations is $7 million to $8 million per life saved, but the VSL varies considerably across the labor force. Our results reconcile the previous discrepancies between hedonic VSL estimates and the values implied by theories linked to the coefficient of relative risk aversion. Because the VSL varies elastically with income, regulatory agencies should regularly update the VSL used in benefit assessments, increasing the VSL proportionally with changes in income over time.

JEL Classification: C23, I10, J17, J28, K00

Keywords: value of statistical life, VSL, quantile regression, panel data, fixed effects, PSID, fatality risk, CFOI

Corresponding author:

Thomas J. Kniesner
Department of Economics
Syracuse University
426 Eggers Hall
Syracuse, NY 13244
USA
E-mail: tkniesne@maxwell.syr.edu

The underlying economic benefits principle for policy assessments is that the benefit value is society's willingness to pay for the benefit. The standard process of calculating the benefits of policies that reduce fatality risks involves multiplying the average risk reduction for the affected population by an average willingness to pay for the risk reduction based on the value of statistical life (VSL) revealed in labor market decisions.[1] Whether the VSL from the labor market appropriately transfers as a benefits measure to other contexts depends on whether the population whose risks are being reduced has preferences for risk reduction that are similar to those in the labor market studies being used as the reference point.[2] There are many sources of heterogeneity, including those stemming from age and other demographic factors. Our research presented here examines the economic and policy importance of heterogeneity in VSL based on income and the closely related matter of the person's risk-taking behavior.

To elaborate, we examine the VSL distribution across the population using quantile regression estimates from panel data on workers' wages. Our approach yields estimates of the VSL for different segments of the wage distribution, which are associated with different fatality risk levels and income levels. Because safety is a normal good, workers with higher income levels will have higher VSL levels and will tend to choose jobs with lower risk levels. Our empirical research explores the heterogeneity of VSL across different segments of the wage distribution and the associated income levels of people in the segments to infer the income elasticity of VSL. We consequently provide evidence on the relationship of VSL to income levels and to fatality risk levels, each of which have been matters of policy concern.

---

[1] The benefit assessment procedure is consistent with that summarized by the U.S. Office of Management and Budget (2003), OMB Circular A-4, Regulatory Analysis. Although altruistic concerns may be relevant, typically the focus is on the willingness to pay by those whose risk level has been reduced by the policy.
[2] There may also be differences in the kinds of deaths involved and the associated morbidity effects, which may affect valuations as well (Viscusi 2009).

The U.S. Environmental Protection Agency (EPA) has long been concerned with the relationship of the VSL to the risk taking activities of the agency's target populations. Officials at EPA often express doubts as to whether VSL estimates for workers who have self selected into risky jobs are pertinent to the protected population for involuntary environmental risks. The voluntary/involuntary risk distinction has also been included in proposed Federal legislation as a factor that should influence the selection of the VSL.[3] Although there have been many pleas for using higher VSL amounts for involuntary risks, we know of no policymakers who have advocated using lower VSL amounts for people who have chosen to incur large risks, such as cigarette smokers or workers in extremely hazardous jobs. There is a clear asymmetry in the willingness of policymakers to draw distinctions, as increases in the VSL for involuntary risks are viewed with favor, but decreases in VSL are not advocated for people who face large risks voluntarily.[4]

The counterpart of the risk-VSL relationship is the risk-income relationship. The U.S. Department of Transportation (DOT) provides for consideration of income in setting the VSL for benefit assessment, as it permits the use of an income elasticity adjustment of 0.55 based on the meta analysis of Viscusi and Aldy (2003).[5] Agencies within DOT, notably the Federal Aviation Administration, have long maintained that the higher income levels of those protected by their regulations imply that a higher VSL level should be used for their regulations than for transportation policies generally.[6] Senate Bill S. 3564 also recognizes the role of income

---

[3] The distinction arises in Section 4A of the bill S. 3564 introduced by Senator Barbara Boxer in 2008, 110[th] Congress, 2d Session, the "Restoring the Value of Every American in Environmental Decisions Act."

[4] Such asymmetries in valuation are similar to a wide variety of behavioral anomalies, such as loss aversion in prospect theory and a variety of models of reference dependence.

[5] U.S. Department of Transportation, Office of the Assistant Secretary for Transportation Policy, Revised Departmental Guidance: Treatment of the Value of Preventing Fatalities and Injuries in Preparing Economic Analyses, August 2, 2005.

[6] Graham (2008) suggests that when policies have clearly different target groups, such as air travel versus bus travel, then income elasticities might be taken into account. Similarly, the Federal Aviation Administration sought to use a

adjustments, stating that the VSL amount must be increased annually to reflect changes in income.[7] The proposed legislation does not permit decreasing VSL if income levels decline, which incomes did during the recession that began shortly after the legislation was proposed in 2008. Our estimates will indicate whether an income-based adjustment to VSL is warranted and to what extent.

Because in a quantile wage regression the VSLs vary with the potential wage ($\hat{w}$), it admits the possibility that VSL varies positively with income levels and permits the calculation of income elasticities for the different quantile groups. When income elasticities of VSL differ by income levels, it opens the door for Ramsey pricing of safety programs similar to that used for other health-related programs such as hospital pricing, where the decision is health with distributional consequences (Harris 1977).

Estimates of the income elasticity are also useful in exploring whether the model underlying the hedonic equilibrium equation yields results consistent with economic theory. Kaplow (2005), among others, has noted the connection between risk aversion and VSL and has emphasized that VSL must be income elastic for the empirical results to be consistent with estimates based on the estimated coefficient of relative risk aversion in the consumption literature. Based on existing estimates in the literature, there is currently a disconnect between the theory and empirical evidence. The VSL has an estimated income elasticity less than 1.0, whereas risk aversion estimates would imply a VSL that is income elastic (Kniesner and Ziliak 2002; Viscusi and Aldy 2003; Aldy and Viscusi 2008). Our income elasticity estimates establish

---

higher value of life than for other transportation policies because of the higher income of air travelers. See Viscusi (1993), which is based on his report to the FAA. Based on this report, the agency sought to justifying using a higher VSL in its policy analyses.

[7] More specifically, VSL must be increased at least once a year to reflect "the average annual total compensation of individuals, including income and benefits."

the consistency between the risk preferences in theoretical models and in the VSL literature that has eluded economists up till now.

The relationship between VSL and income has also been a prominent component of broader commentaries on benefit-cost analysis that have challenged the willingness-to-pay principle for benefit assessment. According to the critiques, VSL amounts should be adjusted for people's ability to pay and should not be limited by a person's current resources. Given the positive income elasticity of VSL, such a proposal without an upper bound on the reference income amount is not well-defined. The ability to pay approach also is inconsistent with basing policies on the actual preferences of those being protected.

Our examination of the variation of the VSL with wage levels also has implications for the relation of VSL to risk levels. Consider a standard hedonic wage model. For any given firm, the offer curve of wages for different levels of risk will be governed by the requirements that workers be paid their marginal product and profits equal zero. Because safety is costly to provide, and increasingly so as the level of safety is increased, the wage offer curve $w(\pi)$ for the firm as a function of the risk level π will satisfy $dw/d\pi > 0$ and $d^2w/d\pi^2 < 0$. The market offer curve $W(\pi)$ is the outer envelope of the individual firms' offer curves and is generally assumed to have similar shape with $dW/d\pi > 0$ and $d^2W/d\pi^2 < 0$.[8] Workers choose various points along the market offer curve, with the local rate of tradeoff simultaneously reflecting the wage-risk tradeoff along the worker's constant expected utility locus and the wage-risk tradeoff on the firm's iso-profit curve.

The hedonic labor market equation $W(\pi)$ estimates the locus of such choices. If all workers chose the same job, individual preferences would be determining the market tradeoff

---

[8] Should there be any downward dip in the market offer curve, no worker would ever choose such a job that is dominated by available jobs offering higher wage for lower risk.

rate rather than the curvature of the market opportunities locus. But with people sorted along different points of the opportunities locus, the shape of the opportunities will be influential. Four empirical studies have examined the curvature of the market equilibrium wage-risk tradeoffs.[9] Each of the studies found that a worker's wages are positively related to the level of risk, but at a diminishing rate. Thus, the VSL varies with the level of risk and declines as the risk level increases, as workers most willing to bear risk sort themselves into the riskiest pursuits. The VSL-risk relationship is also consistent with the hypothesized theoretical shape of the market opportunities locus. Our quantile regression analysis estimates the VSL levels at different junctures of the wage offer curve, generating information on the relationship of VSL to the fatality risk and worker income levels.

Section 1 outlines our econometric approach. The estimation procedure for our panel data uses a new extension of the quantile regression method that recognizes the role of latent person-specific heterogeneity, which Kniesner et al. (2008) have found to be critical in VSL studies. For the estimation, we construct a data set consisting of five waves of worker data from the Panel Study of Income Dynamics to which we match highly refined fatality risk data using the Census of Fatal Occupational Injuries. The empirical estimates in Section 3 indicate substantial variation in the VSL estimates across the quantile ranges, with a median value of about $7.5 million. In addition to providing VSL estimates at different levels of fatality risk and wage rates, our findings have fundamental policy ramifications as well. Section 3 uses the findings to show how a social welfare function based on willingness to pay will be more protective from a safety standpoint than the implications of a social welfare function in which all individuals' risk reductions receive the same weight. In Section 4 we show that the VSL is highly elastic with

---

[9] The studies, which are reviewed in Viscusi and Aldy (2003), are Viscusi (1981), Olson (1981), Dorsey and Walzer (1983), and Leigh and Folsom (1984).

respect to income. Our income-elastic VSL result using within sample estimates of the elasticity of the VSL with respect to family income is more in line with theoretical models, such as that of Kaplow (2005), than previous estimates based on comparisons across hedonic wage studies. Section 5 places our results in the context of the VSL literature.

1. Econometric Framework

In a previous paper we established that individual heterogeneity in the form of a latent time invariant intercept was crucially important to estimating VSL in a linear panel data regression model, much more so than endogeneity of fatal injury risk (Kniesner et al. 2008). It was also the case that how one controlled for heterogeneity (fixed effects versus a mis-specified exogenous random effects model) was far less important to the estimates than whether one controls for differential intercepts at all.

In continuing our exploration of the heterogeneity of regression parameters, we examine slope differences using the concept of quantile regression. It is only recently that estimators have appeared for quantile regression in a panel context that also permit latent person-specific heterogeneity (Koenker 2004; Lamarche 2006). Interquartile differences in estimated fatality risk marginal effects and the associated VSLs capture distributional issues such as asymmetry not evident in mean regression. Even a simple comparison of the mean versus median VSL is instructive for safety policy where the VSL is a benefit comparison point for evaluating life-saving programs with different cost levels. In particular, using the median program benefit as a cutoff value ensures that a majority of the affected population will benefit from the program, which need not be the case using the mean of an asymmetric benefit distribution as a cost-effectiveness cutoff.

To fix notation let our regression model be

$$\ln w_{it} = x_{it}'\beta + \alpha_i + u_{it},\tag{1}$$

where i indexes the worker, t indexes time, and the vector of covariates includes fatal injury risk (π) and the usual demographic characteristics of the worker (z). The α_i's are the time-invariant worker-specific effects, and u_{it} is the usual random error term. In (1) there are common marginal effects.

1.1 Quantile Regression Approach

In a quantile regression model one has $\beta(\tau_j)$ where τ_j indexes the quantiles of the potential regression outcomes. For tractability, when T is finite the researcher usually begins by assuming that the worker-specific effect, α_i, is a so-called pure location shift that applies (is common) to all the conditional quantiles of the regression outcomes. In our estimation we follow the innovative regression model developed in Lamarche (2006). He begins by noting that a shrinkage estimator wherein a tuning parameter (call it λ) controls the degree of inter-person intercept differences is a way to limit the variability of the multiplicity of marginal effect estimates in a quantile regression. The tuning parameter ranges across the spectrum of complete to no heterogeneity and can either be estimated or fixed ex ante. The idea is the optimal shrinkage of the differentials toward a common intercept. We minimize the trace of the covariance matrix when estimating the tuning parameter as well as provide estimates of several extreme cases, such as complete versus no intercept heterogeneity.

Lamarche's (2006) estimator finds the

$$\arg\min(\beta,\alpha)\sum_{j=1}^{J}\ \sum_{t=1}^{T}\ \sum_{i=1}^{N}\omega_{\tau_j}\rho_{\tau_j}(\ln w_{it} - x_{it}'\beta(\tau_j) - \alpha_i) + \lambda\sum_{i=1}^{N}|\alpha_i|,\tag{2}$$

where ω_{τ_j} is the relative weight of the j^{th} quantile and $\rho_{\tau_j}(u) = u(\tau_j - I(u \leq 0))$ is the quantile loss function.

Here the tuning parameter ($\lambda = \sigma_u^2 / \sigma_\alpha^2$) regulates the influence on the quantiles of the estimated worker effects. In the case where $\lambda = 0$ the fixed effects estimator emerges while for the case where $\lambda > 0$ a penalized (shrinkage) estimator with fixed effects appears (Lamarche 2006). As noted, we explore two ways of valuing λ, fixing it ex ante, as the estimator $\hat{\beta}(\tau, \lambda)$ is asymptotically unbiased for all $\lambda > 0$, and finding its trace minimizing estimated value, which is a form of feasible generalized least squares (GLS) analogous to penalized least squares for panel data (Lamarche 2006). Finally, the pattern of $\hat{\beta}(\tau_j)$ describes how the resulting marginal impact of fatality risk varies across potential wage outcomes.

We also use our estimates of $\hat{\beta}_1(\tau_j)$ to construct estimates of the value of a statistical life and note how the associated VSLs vary across potential wage outcomes. Accounting for the fact that fatality risk is per 100,000 workers and that the typical work-year is about 2,000 hours, the estimated value of a statistical life for a quantile of the potential wage distribution is

$$VSL(\tau_j) = \left[(\frac{\partial \hat{w}}{\partial \pi}(\tau_j) = \hat{\beta}_1(\tau_j) \times w(\tau_j)) \times 2000 \times 100,000 \right]. \tag{3}$$

To summarize, equations (1) and (3) describe our organizing framework: (1) is the hedonic market wage locus and (3) is the associated value of statistical life (VSL), which depends on the estimated parameters of (1) via the marginal effect of fatal injury risk (π) and is non-constant in the context of a quantile regression estimator (2). There are two kinds of heterogeneity that come into play. One is econometric heterogeneity whereby the wage equation intercepts vary with the person indicator (i). The other is economic heterogeneity whereby there

is curvature of the hedonic locus ($\hat{\beta}$ varies with τ) to reflect both latent worker and firm differences in risk tolerance and cost functions.

1.2 PSID Data

The main body of our data comes from the 1993–2001 waves of the Panel Study of Income Dynamics (PSID), which provides individual-level data on wages, industry and occupation, and demographics. The PSID survey has followed a core set of households since 1968 plus newly formed households as members of the original core have split off into new families.

The sample we use consists of male heads of household ages 18–65 who are in the random Survey Research Center (SRC) portion of the PSID, and thus excludes the oversample of the poor in the Survey of Economic Opportunity (SEO) and the Latino sub-sample. The male heads in our regressions (i) worked for hourly or salary pay at some point in the previous calendar year, (ii) were not permanently disabled or institutionalized, (iii) were not in agriculture or the armed forces, (iv) had a real hourly wage greater than $2 per hour and less than $100 per hour, and (v) had no missing data on wages, education, region, industry, and occupation. All real wages and VSL estimates are in terms of 2001 dollars based on the personal consumption expenditure deflator.

Beginning in 1997 the PSID moved to every-other-year interviewing. For consistent spacing of survey response we use data from the 1993, 1995, 1997, 1999, and 2001 waves. The use of every-other-year responses will be one of many mechanisms to reduce the influence of measurement error in our estimated VSL. We do not require individuals to be present for the entire sample period; we have an unbalanced panel where we take missing values as random

events.[10] Our sample filters yielded 2,036 men and 6,625 person-years. About 40 percent of the men were present for all five waves (nine years); another 25 percent were present for at least four waves.

The demographic controls in the model include years of formal education, a quadratic in age, dummy variables for state of residence, dummy indicators for region of country, race, union status, marital status, and one-digit occupation.

1.3 Fatality Risk Data

We used the fatality rate for the worker's two-digit industry by one-digit occupation group. We distinguished 720 industry-occupation groups using a breakdown of 72 two-digit SIC code industries and the 10 one-digit occupational groups. After constructing codes for two-digit industry by one-digit occupation in the PSID we then matched each worker to the relevant industry-occupation fatality risk. Our worker fatality risk variable uses proprietary U.S. Bureau of Labor Statistics data from the Census of Fatal Occupational Injuries (CFOI) for 1992–2002.[11]

The CFOI provides the most comprehensive inventory to date of all work-related fatalities in a given year. The CFOI data come from reports by the Occupational Safety and Health Administration, workers' compensation reports, death certificates, and medical examiner reports. To be classified as a work-related injury the decedent must have been employed at the time of the fatal event and engaged in legal work activity that required the worker be present at the site of the fatal incident. In each case the BLS verified the work status of the decedent with two or more of the above source documents or with a follow-up questionnaire in conjunction with a source document.

---

[10] Ziliak and Kniesner (1998) showed that when there is non-random attrition fixed effects models should remove it along with the other time-invariant factors.

[11] Restricted access to CFOI fatality data was obtained via an agreement with the U.S. Bureau of Labor Statistics. Our variable construction procedure follows that in Viscusi (2004), which describes the properties of the 720 industry-occupation breakdown in greater detail.

Because our fatality risk variable is by industry and by occupation, it provided a much more pertinent measure of the risk associated with a particular job than a more broadly based index, such as the industry risk alone, which is the most widely used job risk variable. Taking into account the occupation as well as the industry as we did here substantially reduces the measurement error in the fatality risk variable.

The importance of the industry-occupation structure of our risk variable is especially great within the context of a panel data analysis. By using a fatality risk variable that varies over time and is defined for 720 industry-occupation groups, we greatly expanded the observed variance in workers' job risks across different periods.

Our focal measure of the fatal injury risk rate used the number of fatalities in each industry-occupation cell. The denominator is the number of employees for that industry-occupation group in survey year t. Our measure of the fatality risk is time-varying because of changes in both the numerator and the denominator.[12]

We expect less reporting error in the industry information than in the occupation information, so our annual measure should have less measurement error than if the worker's occupation were the basis for matching (Mellow and Sider 1983; Black and Kniesner 2003). To reduce the influence of large swings in fatality risk further, we dropped person-years where the percentage change in fatality risk exceeded a positive 300 percent or negative 75 percent. The sample mean fatality risk for the annual measure is 6.4/100,000.

2. Quantile Hedonic Wage Equation Estimates: Marginal Effects and VSL

Table 1 presents quantile regression estimates of the fatal injury risk for no latent heterogeneity and varying degrees of possible heterogeneity ex ante as reflected in the specified

---

[12] We used the bi-annual employment averages from the U.S. Bureau of Labor Statistics, Current Population Survey, unpublished table, Table 6, Employed Persons by Detailed Industry and Occupation for 1993–2001.

value of λ. The regression for $\lambda = 1$ minimized the trace of the variance-covariance matrix. Note that the most curvature of the hedonic wage function (the biggest inter-quantile differences in the estimated marginal effect) appears in the regression without latent heterogeneity and that the least curvature appears when the most heterogeneity is specified ex ante. We rely on the estimated marginal effects from the regression selected by the econometric model, which is where $\hat{\lambda} = 1$, and $\hat{\beta}(\tau = 0.90) / \hat{\beta}(\tau = 0.10) = 1.5$.

Figures 1 and 2 present two different calculations of VSL. Figure 1 uses the quantiles' actual wage rate averages so that the VSL computation using (3) reflects inter-quantile differences in both marginal effects and potential wages. Figure 2 standardizes the VSL by the median wage rate. Again focusing on the VSLs for the regression with the heterogeneity weight selected by the econometric model ($\lambda = 1.0$), the $VSL(\tau = 0.90) / VSL(\tau = 0.10) \cong 5$ when both sources of variation are included versus $VSL(\tau = 0.90) / VSL(\tau = 0.10) \cong 1.5$ based on differences in marginal effects alone. The effect of the wage in the calculation is much more important than the effect of the differences in estimated marginal effects of fatal injury risk. Of course, in both cases the median VSL is about $7 million to $8 million and is a possible cutoff for regulatory policy decisions in that regulations costing no more than this benefit at least half the affected population.

The essence of our quantile regression approach is that the estimated marginal effects differ by potential wage (and therefore income) so that there is (more) nonlinear curvature potential. The policy relevance of our results includes how they may pertain to updating of the VSL for income growth when used in regulatory agency decisions. The implication of how the curvature is reduced when controlling for individual intercepts is analogous to the issue of poolability in linear panel models. As is well-known, if there is a positive marginal effect of a

covariate and intercepts vary positively with the covariate, then an equation that pools the data and forces a common intercept will estimate a steeper regression line than the flatter (set of) true regression lines that have intercept heterogeneity. Analogously, we found that the curvature is reduced dramatically (the marginal effect of fatality is more constant across quantiles) when we allowed for intercept heterogeneity. As in the case of the linear panel models we have estimated in earlier papers, the researcher needs to allow for intercept heterogeneity in panel quantile models of hedonic wage equilibrium with fatal injury risk.

There is yet another dimension of the curvature we found in the estimated marginal effects and the associated VSL. Recently, Chetty (2009) fleshed out the concept of a sufficient statistic for policy analysis. The idea is that a deep structural parameter, while generally desirable for policy evaluation, may be econometrically difficult to obtain because of endogeneity of the constraint set and the implausibility of instruments coupled with the inability or unwillingness to wait for nature or the NSF to produce a suitable experiment or random controlled trial. An alternative research approach is to use economic theory to locate reduced form parameters that are sufficient to infer welfare effect levels and directions, so-called sufficient statistics.

Here the sufficient-statistic issue relates to the curvature of the hedonic locus. Theoretically, the hedonic wage equation that we estimate is the locus of tangencies of workers' indifference curves and firms' iso-profit curves or the set of market equilibrium outcomes, and, as such, need be upward sloping as both sets of curves slope upward but can have any shape depending on the degree of heterogeneity of firms versus workers. We found not only the needed upward slope, but also a slope that declines with the potential wage so that the hedonic locus we estimated reflects not only the indifference curve shape but also the heterogeneity of preferences across workers at different income levels. This is sufficient to infer that differences in the

aversion to a work-related death by workers dominate cost structure differences of firms. The result is also sufficient to conclude that the compensation a worker needs to accept a non-infinitesimal additional risk is at least as great as the amount indicated by the slope of the hedonic locus (the locus would be an upper bound if inter-firm heterogeneity were instead dominant). Thus, the heterogeneity that we find in the marginal effects via quantile regression is a sufficient statistic, so to speak, for using the set of estimated VSL numbers in Figures 1 and 2 as thresholds for regulatory costs of a life-saving regulation. On average, or at any potential wage quantile, our results imply cost-effectiveness of a regulation that is less costly.

3. Social Welfare Implications of Heterogeneity in Our VSL Estimates

As we have noted, the standard benefits measure in benefit-cost analyses is the average societal willingness to pay for the benefit. The theoretical foundation for the benefit calculation is based on a social welfare function v that is additive in the utilities of the individuals in the citizenry. In addition, whatever distributional weights are applied are constructed so that the marginal social benefit of increasing the wealth (y_i) of person i is equalized across the population, or that $\partial v / \partial y_i$ is the same for all i. Thus, the benefit value is the sum of the different individuals' willingness to pay for the benefit, where each person's willingness to pay amount is accorded the same weight.

In the case of reduced mortality risks, consider the benefit associated with small reductions in mortality risk that will eliminate m expected deaths for population n, where the fatality risk probability (π) being reduced is m/n. Let VSL_i be the value of statistical life for person i. Then the value of the benefit is

$$\sum_{i=1}^{n} VSL_i (m/n) = m[(1/n)\sum_{i=1}^{n} VSL_i] = m\overline{VSL_i} \,, \tag{4}$$

where $\overline{VSL_i}$ denotes the mean VSL for the population. The mean VSL across the population

multiplied by the expected number of deaths that will be prevented yields the benefit amount.

Recent theoretical research by Manski (2009) implies that the arithmetic mean VSL

would be the cutoff for program economic efficacy, even if VSL is heterogeneous across people.

The argument is that the government should buy the identical amount of safety for everyone,

which is the average amount that individuals would demand because a common amount is

welfare maximizing in a probabilistic setting, much like a certain income dominates income

variability under insurance. Manski's theoretical conclusion is in contrast to the social welfare

function discussion in Baker et al. (2008), where the implication is that the harmonic mean of

VSL may be a reasonable alternative cost threshold for regulatory policy if differences in

willingness to pay are ignored.

Although the formulation in (4) is the standard approach in the literature, one could

hypothesize other forms for the social welfare function. There has long been discussion of the

merits of using the willingness-to-pay measure as the reference point. Some non-economists

have suggested that ability to pay should not be a factor, but the critics do not specify how the

social welfare function should be modified to achieve the purported equity. Can one achieve the

desired equitable results by endowing people with income they do not have? Will the net value

of endowments still be within the societal budget constraints? How much should this endowment

be and how should it be distributed across the population?

In contrast to the ill-defined critique of conventional social welfare functions, Baker et al.

(2008) explore an approach that is both well defined and incorporates a concept of equity. Rather

than constructing the social welfare function to equalize the marginal social value of wealth

across individuals, they suggest a model in which the social welfare function equalizes the

marginal social welfare value of the risk reduction. So, suppose that irrespective of differences across the population in a person's willingness to pay for a given reduced mortality risk π_i, that the distributional weights in the social welfare function are constructed so that rather than $\partial v / \partial y_i$ being equalized it is $\partial v / \partial \pi_i$ that is equalized across the population. In the alternative formulation of the social welfare function, the key benefit parameter is not the arithmetic mean VSL across the population but rather the harmonic mean VSL.[13] Although the harmonic mean framework is unconventional, it shows the fundamental role of VSL in a variety of social welfare contexts, and we will use our quantile estimates to explore both the arithmetic mean and harmonic mean values.

The standard hedonic wage models do not yield estimates of either the arithmetic mean VSL or the harmonic mean VSL. Rather, they generate the coefficients that minimize the sum of squared residuals when fitting a regression to the set of observed wage-fatality risk combinations observed in the labor market. For our data, the conventional VSL estimate from a hedonic wage equation that will serve as the reference point is $7 million to $8 million (Kniesner et al. 2008). Thus the conventional regression VSL estimate is similar to our median quantile estimate of $7.55 million.

By using the results implied by the quantile wage estimates, we are able to explore the shape of the distribution of the individual VSL_is across the population. For purposes of our research we pool the quantile VSL estimates from the focal set of quantile estimates (0.1, 0.25, 0.50, 0.75, 0.90) and the additional quantile values (0.13, 0.38, 0.63, 0.83, 0.95). For concreteness, we will assume that the population distribution of VSL_i values is the same as that

---

[13] For policies financed by taxes, there is also an adjustment term involving the covariance of individual i's tax rate and the inverse of VSL_i. If that covariance is zero, then the appropriate unit benefit value is simply the harmonic mean VSL. When policies are funded by a lump sum tax the benefits measure equals the harmonic mean VSL. More generally, the harmonic mean will be a lower bound on the benefits measure. See Baker et al. (2008).

for our sample of workers. It is straightforward to match the VSL levels to particular affected

populations based on wage rate differences using the empirical results above. As Figure 1

indicates, there is substantial heterogeneity; the VSL rises at an increasing rate as one moves to

the upper quantiles of the wage distribution. The discussion below will focus on the estimates for

the preferred model, which is the random effects model for which $\lambda = 1.0$.

To calculate the value of VSL implied by the quantile results, we fit a quadratic function

to the quantile estimates.[14] The arithmetic mean VSL across the entire wage distribution equals

$10.7 million, which exceeds both the median value as well as the mean from a standard hedonic

wage regression. The mean VSL here would be increasingly closer to the standard mean hedonic

VSL if the number of quantiles used for the illustrative calculation were increased. For our

purposes what is of primary interest is the relative magnitude of the mean VSL and the harmonic

mean VSL, where the arithmetic and harmonic means are both calculated in a consistent way

across the distribution of estimates.

The non-standard social welfare formulation in which there is an equal marginal social

value placed on risk reductions for all citizens irrespective of their individual willingness to pay

is not necessarily more protective, and indeed the opposite may be the case. The harmonic mean

VSL that emerges as the key parameter in the equalized social valuations of risk reduction

models must always be less than the arithmetic mean for VSL amounts, which are by definition

never negative and not constant across the population. In the special case in which policies are

financed by a lump-sum tax that is identical for all individuals and one adopts the risk value

---

[14] Let q indicate the quantile percentile VSL, then VSL = 5.578 (1.215) – 15.453 (7.373) q + 38.454 (8.181) q^2, $R^2 =$ 0.97, where the numbers in parentheses are robust standard errors. For the semi-log version of the model, ln (VSL) = 1.142 (0.063) + 1.525 (0.374) q + 0.737 (0.383) q^2, $R^2 = 0.99$. Based on the linear model, the arithmetic mean VSL is $10.7 million, and the harmonic mean VSL is $7.2 million, and with the semi-log model the arithmetic mean is $10.7 million, and the harmonic mean is $7.0 million. The fitted predicted values for the semi-log specification have more attractive properties in that the VSL is always increasing as the quantile percentage rises so those estimates will be the focus of the discussion.

equalization social welfare function approach, the pertinent VSL for benefit assessment is the harmonic mean. For our quantile estimates, the harmonic mean is $7.0 million, which is about two-thirds of the size of the arithmetic mean calculated using the same quantile estimates and about equal to the median VSL. The implications of the standard social welfare approach in which willingness to pay valuations receive the same marginal social value are larger estimated benefits and greater levels of risk protection than would be generated by a social welfare function that equalizes the social welfare weight placed on different marginal risk reductions. Put somewhat differently, the sometimes criticized emphasis on individual willingness to pay will lead to more stringent regulations than an approach divorced from willingness to pay coupled with an assumption that risk reductions of a given magnitude have the same value across society irrespective of one's financial resources.

4. Additional Policy Relevant Heterogeneity: New Estimates of Income Elasticity of VSL

Quantile estimates of the VSL can be used in conjunction with family income information in the PSID to calculate the income elasticity of the VSL,

$$\eta(\tau_j)_{VSL,y} = (\partial VSL(\tau_j) / \partial y(\tau_j))(y(\tau_j) / VSL(\tau_j)), \tag{5}$$

which Table 2 presents.

The VSL levels exhibit particularly sharp increases for the 0.75 and 0.90 quantiles. The two high-wage quantiles also are coupled with substantial decreases in the fatality risk level so that we observe both a positive income elasticity of VSL as well as a negative elasticity of the VSL with respect to average fatality rates. We estimate the elasticity of the estimated VSL at different quantiles with respect to the real family income levels at the quantiles.[15] Based on a linear regression of VSL on the quantile-specific real family income, the overall income

---

[15] Note that our estimates do in fact construct an income elasticity based on real family income levels, whereas what the previous literature refers to as an income elasticity is actually an elasticity with respect to worker wage levels.

elasticity of VSL across the quantiles is 1.44.[16] The elasticity amount varies depending on the

quantile, with a high value of 2.24 at the low-wage 0.10 quantile to a low elasticity value of 1.23

at the high-wage 0.90 quantile.[17] The estimates are similar for a specification that also includes a

quadratic real family income term, with an overall elasticity of 1.32 and an elasticity range from

2.06 at low wage levels to 1.16 at high wage levels. In each case, higher income levels have the

most pronounced effect on the VSL for the low-wage workers, and the elasticity estimates

throughout all quantile values are above 1.0.

The very high income elasticity of VSL at low wage levels may account for why

countries at less advanced stages of development enact very few health, safety, or environmental

regulations until per capita income levels in those countries rise to an adequate level. The

perceived economic benefits of risk regulation will be extremely low if income levels are low

because the VSL declines at very low income levels at a rate that is more than proportional to the

decrease in income.

Our estimates of a large income elasticity of VSL are consistent with the simple

theoretical models that have been developed. Kaplow (2005), for example, develops a model of

VSL in which there is no bequest motive and the person picks the level of precautionary

expenditures.[18] In his model, the coefficient of relative risk aversion provides a lower bound on

the consumption elasticity of VSL, which is similar to but not identical to the income elasticity.

With recent estimates of the coefficient of relative risk aversion being around 2 based on the

---

[16] The linear equation used to estimate income elasticity is VSL = -5499606 (773449) + 259.3 (9.6), $R^2 = 0.998$.
[17] The quantile ranges are based on the quantile values at points estimated in the regressions. One would, for example, expect the elasticity value to be even greater at the 0.01 quantile.
[18] Other studies of this relationship along similar lines include Eeckhoudt and Hammitt (2001) and Evans and Smith (2010). The latter paper shows that with more complex and more realistic models the relationship between the income elasticity of VSL and the coefficient of relative risk aversion becomes less clear cut.

labor supply analysis of Chetty (2006) and the consumption analysis of Kniesner and Ziliak (2002), one would expect the VSL to be income elastic, which is what the results above indicate.

Our VSL income elasticity estimates are somewhat higher than those estimated for the valuation of job injuries. Estimates of the income elasticity of job injuries based on the estimated functional forms of utility functions imply lower elasticities than those indicated by the coefficient of relative risk aversion. For two different models estimated in Viscusi and Evans (1990), the estimated income elasticity of the value of statistical job injuries was 1.10 and 0.67. They also noted that the elasticity values exceed the estimated income elasticity for health insurance, which typically are in the range of 0.5 and below.

The greatest inconsistency with previous empirical findings is that our estimate of an income-elastic VSL is inconsistent with the implications of VSL meta analyses in the literature. The meta analyses have estimated the elasticity of VSL with respect to the average worker wage rate in the study, which the authors generally refer to as the income elasticity. Using a large sample of hedonic wage studies, Viscusi and Aldy (2003) take the models and authors' methods used in a series of previous meta analyses and find an income elasticity of VSL of 0.51-0.61. The specifications developed by Viscusi and Aldy (2003) likewise yielded income elasticity estimates ranging from 0.5-0.6. The U.S. Dept. of Transportation (2008) has selected the midpoint of their range 0.55 as its VSL income elasticity figure for policy analysis. It is especially noteworthy that none of the 95% confidence intervals for the VSL income elasticities reported in Viscusi and Aldy (2003) include estimates as high as 1.0.

The inconsistency between the implications of the meta analyses and our quantile results may be another example of the importance of accounting for within sample heterogeneity in VSL. Unlike the meta analyses across studies, our research focuses on the within sample

elasticity. Our quantile results demonstrate that there is considerable heterogeneity in the VSL as evidenced by the substantial curvature in the relationship between the VSL and the worker wage rate. In contrast, the meta analyses are based on pooling the single VSL estimates from different studies, none of which capture the within sample heterogeneity that we do here. Because of the nonlinear relationship between the VSL and income levels, meta analyses of the VSL-wage relationship do not capture the strong curvature in the VSL-income relationship and will tend to mute the estimated income elasticity.

5. Conclusion

We have noted the dependence of VSL on demographic factors such as age, race, sex, and health habits in the current literature (Viscusi and Aldy 2003; Aldy and Viscusi 2008). Potential wealth and consumption also produces interpersonal differences (Kniesner, Viscusi, and Ziliak 2006). Somewhat more controversial is heterogeneity in VSL according to how one dies (Viscusi 2009). The heterogeneity issue is policy relevant as targeted (more efficient) regulatory cutoffs are the objective.

We find substantial heterogeneity in the marginal effects of fatal injury risk on wages using panel data and a quantile regression model that admits latent intercept heterogeneity. Inter-quantile differences in wages further magnify the differences in VSL so that it varies by a factor of 4.0 in our preferred regression specification. Although the extreme differences affect the arithmetic mean VSL, it does not affect the median, which is also close to the harmonic mean VSL. We noted reasons why a potential cut-off for regulatory policy might be the median, which guarantees a majority benefit, or the harmonic mean, which may maximize social welfare. Both the median and the harmonic mean we find are $7 million to $8 million so that the regulatory cutoff is robust to the harmonic mean and median regulatory cost cutoff justifications.

Finally, we examine the income elasticity of VSL, and find that it is income elastic in our data. The income-elastic VSL we find is important for two reasons, one related to removing a seeming inconsistency in two strands of empirical research and the other related to updating regulatory policy decisions for societal income growth or targeting cutoffs according to the affected population. An income-elastic VSL is required for it to be consistent with the estimated risk aversion parameters in the literature. Our results are the first we know of to be consistent with the risk preferences estimates elsewhere. Moreover, even if one does not find it politically feasible to use regulatory cost cutoffs that vary with the incomes of the target population, at a minimum our results support increasing a common cutoff over time proportionately with society's income growth.

References

Aldy, J. E. & Viscusi, W.K. (2008). Adjusting the value of a statistical life for age and cohort effects. *Review of Economics and Statistics, 90*(3), 573-581.

Baker, R., Chilton, S., Jones-Lee, M., & Metcalf, H. (2008). Valuing lives equally: Defensible premise or unwarranted compromise? *Journal of Risk and Uncertainty, 36*(2), 125-138.

Black, D.A. & Kniesner, T.J. (2003). On the measurement of job risk in hedonic wage models. *Journal of Risk and Uncertainty, 27*(3), 205-220.

Chetty, R. (2006). A new method of estimating risk aversion. *The American Economic Review, 96*(5), 1821-1834.

Chetty, R. (2009). The simple economics of salience and taxation. *National Bureau of Economic Research Working Paper Series, No. 15246*.

Dorsey, S. & Walzer, N. (1983). Workers' compensation, job hazards, and wages. *Industrial and Labor Relations Review, 36*(4), 642-654.

Eeckhoudt, L. R. & Hammitt, J. K. (2001). Background risks and the value of a statistical life. *Journal of Risk and Uncertainty, 23*(3), 261-279.

Evans, M. F. & Smith, V. K. (2010). Measuring how risk tradeoffs adjust with income. *Journal of Risk and Uncertainty*, 40(1).

Graham, J. D. (2008). Saving lives through administrative law and economics. *University of Pennsylvania Law Review, 157(2)*, 395-540.

Harris, J. E. (1979). Pricing rules for hospitals. *The Bell Journal of Economics, 10*(1), 224-243.

Kaplow, L. (2005). The value of a statistical life and the coefficient of relative risk aversion. *Journal of Risk and Uncertainty, 31*(1), 23-34.

Kniesner, T. J. & Ziliak, J. P. (2002). Tax reform and automatic stabilization. *The American Economic Review, 92*(3), 590-612.

Kniesner, T. J., Viscusi, W. K., Woock, C., & Ziliak, J. P. (2008). The value of a statistical life: Evidence from panel data. Syracuse, NY, Syracuse University.
http://gatton.uky.edu/Faculty/ziliak/KVWZ_VSL.pdf

Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis, 91*(1), 74-89.

Lamarche, C. (2006). Robust penalized quantile regression estimation for panel data. Norman, Oklahoma, University of Oklahoma. Department of Economics. http://faculty-staff.ou.edu/L/Carlos.E.Lamarche-1/rpan.html

Leigh, J. P. & Folsom, R. N. (1984). Estimates of the value of accident avoidance at the job depend on the concavity of the equalizing differences curve. *Quarterly Review of Economics and Business, 24*(1), 55-66.

Manski, C. F. (2009). When consensus choice dominates individualism: Jensen's inequality and collective decisions under uncertainty. *National Bureau of Economic Research Working Paper Series, No. 15172* .

Mellow, W. & Sider, H. (1983). Accuracy of response in labor market surveys: evidence and implications. *Journal of Labor Economics, 1*(4), 331-344.

Olson, C. A. (1981). An analysis of wage differentials received by workers on dangerous jobs. *The Journal of Human Resources, 16*(2), 167-185.

U.S. Bureau of Labor Statistics. (2009). *Employed persons by detailed industry and occupation for 1993-2001*, Current Population Survey.

U.S. Department of Transportation, Office of the Assistant Secretary for Transportation Policy. (2005). Revised departmental guidance: Treatment of the value of preventing fatalities and injuries in preparing economic analyses. Washington, DC.

U.S. Office of Management and Budget. (2003). *OMB CIRCULAR A-4, Regulatory Analysis* (Rep. No. A-4). Washington, DC.

Viscusi, W. K. (1981). Occupational safety and health regulation: Its impact and policy alternatives. In J.P. Crecine (Ed.), *Research in Public Policy Analysis and Management*. Volume 2, pp. 281-299. Greenwich, CT: JAI Press.

Viscusi, W. K. & Evans, W. N. (1990). Utility functions that depend on health status: estimates and economic implications. *The American Economic Review, 80*(3), 353-374.

Viscusi, W.K. (1993). The value of risks to life and health. *Journal of Economic Literature*, 31, 1912-1946.

Viscusi, W. K. & Aldy, J. E. (2003). The value of a statistical life: A critical review of market estimates throughout the world. *Journal of Risk and Uncertainty, 27*(1), 5-76.

Viscusi, W. K. (2004). The value of life: Estimates with risks by occupation and industry. *Economic Inquiry, 42*(1), 29-48.

Viscusi, W. K. (2009). Valuing risks of death from terrorism and natural disasters. *Journal of Risk and Uncertainty, 38*(3), 191-213.

Ziliak, J. P. (1998). Does the choice of consumption measure matter? An application to the permanent-income hypothesis. *Journal of Monetary Economics, 41*(1), 201-216.

Table 1. Quantile Estimates of Labor-Market Fatality Risk Coefficient with and without Unobserved Heterogeneity

Quantile	No Unobserved Heterogeneity	Unobserved Heterogeneity (λ = Tuning Parameter)			
		$\lambda = 2.0$	$\lambda = 1.0$	$\lambda = 0.5$	$\lambda = 0.01$
0.10	0.0025	0.0026	0.0020	0.0013	0.0010
	(0.0010)	(0.0012)	(0.0012)	(0.0011)	(0.0008)
0.25	0.0022	0.0023	0.0021	0.0015	0.0007
	(0.0009)	(0.0013)	(0.0011)	(0.0010)	(0.0008)
0.50	0.0049	0.0044	0.0022	0.0016	0.0007
	(0.0010)	(0.0012)	(0.0011)	(0.0010)	(0.0008)
0.75	0.0062	0.0046	0.0029	0.0016	0.0007
	(0.0010)	(0.0012)	(0.0010)	(0.0011)	(0.0008)
0.90	0.0075	0.0048	0.0030	0.0016	0.0008
	(0.0015)	(0.0014)	(0.0011)	(0.0011)	(0.0009)

Note: Standard errors are in parentheses, calculated from 500 bootstrap replications. The panel quantile bootstrap standard errors jointly sample the residual and regressors. All models control for a quadratic in age, indicators for education, race, union status, marital status, and region of country, as well as year and state fixed effects. The marginal effects are approximately the proportionate change in the wage for a 1/100,000 change in fatality risk.

Table 2. Panel Quantile Estimates of the Income Elasticity of VSL

Quantile	Hourly Wage Rate ($2001)	Fatality Risk (per 100,000)	Average Family Income ($2001)	VSL ($2001)	Income Elasticity
0.10	8.66	7.83	33,583	3,464,000	2.24
0.25	11.71	8.20	43,478	4,918,200	2.04
0.50	17.16	7.68	57,768	7,550,400	1.76
0.75	25.09	6.07	75,762	14,552,200	1.20
0.90	36.74	3.69	117,792	22,044,000	1.23

Note: From the panel quantile model in Table 1 with tuning parameter $\lambda = 1.0$. All dollar amounts are denoted in real 2001 terms based on the personal consumption expenditure deflator. The elasticity is calculated by regressing the VSL at each quantile on average family income at each quantile, and multiplying the coefficient on family income from the regression by the ratio of family income to the VSL at that point. The mean income elasticity of VSL across quantiles is 1.44.

Figure 1: Quantile Estimates of VSL with and without Unobserved Heterogeneity

Figure 2. Quantile Estimates of the VSL Evaluated at the Median Wage and Quantile -Specific Fatality Parameter