

May 2014 ■ RFF DP 14-13

Heterogeneity in the Value of Life

Joseph E. Aldy and Seamus J. Smyth

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org



Heterogeneity in the Value of Life

Joseph E. Aldy and Seamus J. Smyth*

May 2014

Abstract

We develop a numerical life-cycle model – with choice over consumption and leisure, stochastic mortality and labor income processes, and calibrated to U.S. data – to characterize willingness to pay (WTP) for mortality risk reduction. Our theoretical framework can explain many empirical findings in this literature, including an inverted-U life-cycle WTP and an order of magnitude difference in prime-aged adults WTP. By endogenizing leisure and employing multiple income measures, we reconcile the literature's large variation in estimated income elasticities. By accounting for gender- and race-specific stochastic mortality and income processes, we explain the literature's black-white and female-male differences.

JEL codes: J17, D91, Q51

Keywords: value of statistical life, mortality risk reduction, income elasticity

*Aldy is an Assistant Professor, Harvard Kennedy School; Visiting Fellow, Resources for the Future; and Faculty Research Fellow, National Bureau of Economic Research. Smyth is an Economist at Caxton Associates. Address correspondence to joseph_aldy@hks.harvard.edu, 617-496-7213, and 79 JFK Street, Mailbox 57, Harvard Kennedy School, Cambridge, MA 02138. We have benefitted from comments from David Cutler, Jeff Liebman, Aleh Tsyvinski, and Kip Viscusi as well as seminar participants at Central Florida, Duke/NC State/RTI, Harvard, Yale, the NBER Summer Institute, and the ASSA-AEA meetings.

1 Introduction

From purchasing safe automobiles, to taking vitamins, to exercising at the gym, individuals make decisions everyday that reveal their willingness to pay (WTP) to reduce mortality risk. Government agencies implement a variety of measures, from regulating air pollution, to improving food safety, to vaccination programs, that reduce mortality risk. What are the economic benefits of these mortality risk reduction efforts, and how do they vary across the U.S. population?

To address this question, we develop a life cycle model of the consumption and labor supply choice problem facing an individual. Numerically solving this problem under uncertainty in life expectancy and wage income, we assess how much wealth an individual would forego for a marginal reduction in mortality risk. We calibrate the model to labor market, wealth, and longevity data for the United States, and then run a series of “simulated laboratory experiments.” Each experiment begins with an identical set of 10,000 age-20 agents who make annual consumption and leisure decisions in response to persistent and temporary wage shocks, and agents advance to the next year if they do not face a stochastic mortality shock, until age 100 when all surviving agents die. Through these experiments, we address how the willingness to pay to reduce mortality risk varies with age, income, wealth, and other individual characteristics correlated with longevity expectations, including gender and ethnicity.

Through our simulated lab experiments, we find the following about the heterogeneity in the value of life. The mean willingness to pay to reduce a marginal reduction in mortality is about one-third larger than the median, and the 95th percentile WTP is about 12 times the 5th percentile WTP for a cross-section of age-45 agents. These results reflect the dispersion in incomes that evolves in each experimnt and enables us to investigate the income elasticity for the willingness to pay to reduce mortality risk. We find that the elasticity

of WTP with respect to the permanent component of wages is about 1.7, while the elasticity of WTP with respect to realized income is about 0.6. This average elasticity across all working agents masks heterogeneity with respect to age: age-60 agents have elasticity estimates less than 1/2 and less than 1/4 the estimates of age-30 agents for these two measures, respectively. Across most of the income distribution, the WTP to reduce mortality risk follows an inverted-U shape over the life cycle, with the agents bearing the most adverse productivity shocks (and hence lowest incomes) revealing a steadily declining WTP with age. White agents' WTP to reduce mortality risk is significantly higher than black agents' WTP, with a ratio of about 2 at their life-cycle peaks, and we decompose the drivers of this white-black WTP gap in terms of different labor income processes and life expectancy. Finally, we show that men have higher WTP than women, also as a result of higher wages, until about age 70 when longer female life expectancy dominates the differences in accumulated wealth driven by wage processes.

This analysis makes several contributions to the literature. First, our numerical approach advances the theory and simulation literature focused on valuing mortality risk reduction. Unlike earlier papers that only specified consumption as a choice variable in the life cycle problem (Shepard and Zeckhauser [1984], Rosen [1988], Moore and Viscusi [1990], Johannson [1996], Hall and Jones [2007]), our approach allows for the individual to choose optimal consumption and leisure over the life cycle that permits us to better match our model to observed consumption and leisure patterns and to specify willingness to pay for risk reduction as a trade-off among consumption, leisure, and risk. In addition, ours is the first paper in this literature to explicitly account for uncertainty in the income process. Modeling uninsurable idiosyncratic labor income risk yields consumption and saving behavior consistent with the empirical evidence

on precautionary saving among the young and wealth accumulation typically starting in middle age. This has important results for how willingness to pay varies both over the life cycle and across the cross-section of the population.

Second, our numerical simulations explain some of the observed variations in the value of a statistical life (VSL) estimates in recent empirical work.¹ In revealed preference and stated preference studies, Aldy and Viscusi [2008], Viscusi and Aldy [2007], Kniesner et al. [2006], Cameron and DeShazo [2013], and Krupnick [2007] find that the VSL generally takes an inverted-U shape over the life cycle. ² Costa and Kahn [2004] estimate an income elasticity for the VSL of about 1.5 and Kniesner et al. [2010] estimate income elasticities greater than 1.0 from hedonic labor market analyses, although meta-analyses by Viscusi and Aldy [2003], Mrozek and Taylor [2002], and Liu and Liu [1997] produce income elasticities in the 0.4 and 0.6 range and Cameron and DeShazo [2013] estimate an income elasticity of about 0.7 based on a stated preference survey. Viscusi [2004] and Leeth and Ruser [2003] estimate larger VSLs for whites than for blacks in the labor market. Viscusi and Hersch [2008] find that male workers have larger VSLs than female workers. Our numerical model can provide a theoretical underpinning for these results based on the variation in lifetime income, income shocks, and life expectancy across various demographic groups.

Third, our focus on heterogeneity across the population serves as a comple-

¹The value of a statistical life is simply the aggregation of a population's willingness to pay for a given risk reduction such that it equals the total amount that population would pay to reduce one statistical mortality. For example, if each member of a population of 10,000 is willing to pay \$500 to reduce a 1 in 10,000 risk, then the value of a statistical life for this population is \$5 million. Thus, VSL represents a population-expressed measure that aggregates individuals' willingness to pay to reduce mortality risk. We use both terms in this paper, and in our simulated laboratory experiments, the willingness to pay to reduce mortality risk by one given agent could be considered identical to the VSL for a population of agents that experience identical wage and mortality shocks as this specific agent that represents this population in our experiment.

²Some studies that evaluate the age-VSL relationship do not find this result. For example, Smith et al. [2004] show that the VSL may increase with age. Blomquist and O'Connor [2011] find that VSL follows an inverted-U for adults (aged 30-92 in their sample), but may be higher for children. Early hedonic wage models that simply included the interaction of age and occupational risk typically found a statistically significant, negative interaction effect Viscusi and Aldy [2003].

ment to several papers that abstract from the choice under uncertainty characteristic problem and yield national aggregates for the value of a statistical life. For example, [Hall and Jones, 2007] and Murphy and Topel [2006] provide population aggregate estimates for VSL and use these to assess the benefits of increasing life expectancy through health care investments. Our approach produces a similar life-cycle shape for age-specific mean VSLs as the Murphy and Topel model, but accounting for gender-specific wage processes produces a different result for gender-specific VSLs than they found in a model that only accounts for differences in life expectancy across gender.

In the next section, we outline the individual’s life-cycle model. The third section discusses the calibration of the model and the numerical implementation. In the fourth section, we employ the model to assess our base case and provide results from these simulations showing how the willingness to pay to reduce mortality risk varies with age, income, ethnicity, and gender. The final section concludes the paper, and several appendices provide more information on the computational methods and wage profile and variance estimation.

2 Life-Cycle Utility Model

Our model captures several important features that play a role in determining an individual’s willingness to pay to reduce mortality risk, and hence the population-based measure value of a statistical life. First, the model allows idiosyncratic, age-specific mortality that can differ across demographic groups. Second, agents face uninsurable wage shocks whose magnitudes can differ across different types of agents. Third, agents value both consumption and leisure. Fourth, agents choose optimally their savings and allocation of time to market work accounting for both earnings and mortality risk. This type of life-cycle model with idiosyncratic risk is common in the macroeconomic literature, for

examples see French [2005], Nishiyama and Smetters [2005] and Smyth [2006]. These innovations to the typical life-cycle model used in the health and mortality risk valuation literature allow us to conduct simulations that closely represent actual means and distributions for key economic measures over the life cycle, including labor income, consumption, wealth accumulation, labor participation, and mortality.

2.1 Agent’s Utility Maximization Problem

Agents choose their consumption and leisure to maximize the discounted sum of future utility:

$$\max_{c,l} \mathbb{E} \sum_{t=0}^T \beta^t \phi_t u(c_t, l_t) \quad (1)$$

Agents have a constant, time-invariant rate of time preference β . Their probability of survival to the next period, ϕ_t , differs by age. ϕ_t can differ across agents by ethnicity and gender. This age-dependent survival probability, which declines as agents get older, effectively shortens their time horizon and affects decisions on consumption and leisure as well as their values for mortality risk reduction.

Agents face incomplete markets: private markets for insurance against future wage shocks do not exist, but agents can self-insure through savings. All saving is risk-less and pays the market rate of return. This ability to partially self insure, as well as conduct life-cycle savings, provides a very important means of smoothing consumption in the face of idiosyncratic shocks. There are two uses of an agent’s resources: consumption and saving for future consumption. The agent’s resources consist of prior savings, a_{it} , interest on those savings, $R_{it}a_{it}$, after-tax labor market compensation, $w_{it}(1 - \tau^{SS})(1 - l_{it})$ (note that we denote labor taxes as τ^{SS} to represent the social security payroll tax rate), and social security benefits, b_{it}^{SS} . In each one-year period, the agent allocates

one unit of time between leisure, l_{it} , and market work. This yields the agent's budget constraint:

$$c_{it} + a_{i,t+1} = a_{it}(1 + R_t) + w_{it}(1 - \tau^{SS})(1 - l_{it}) + b_{it}^{SS} \quad (2)$$

The base case is a borrowing constraint of zero. A borrowing constraint of zero ensures that no agents die in debt:³

$$a_{t+1} > a_{min} \quad (3)$$

$$a_{min} = 0 \quad (4)$$

We also consider cases where agents can borrow by relaxing the condition in equation 4. In these cases the age-dependent borrowing constraint starts fairly loose and rises to zero as the agent ages. The tightening borrowing constraint prevents agents from dying in debt with certainty.⁴

2.2 Wage Process

We draw from the macro-consumption literature a wage process that has both permanent and transitory shocks to characterize the wage risks agents face. Estimating such a wage process from microeconomic data is consistent with several important empirical features of the U.S. economy, including wages, asset holdings, and consumption [Samwick and Carroll, 1997, Gourinchas and Parker, 2002, Storesletten et al., 2004].

This wage process includes a permanent component of wages and a transitory

³If agents can borrow, some agents who die young will owe money. Setting a borrowing constraint of zero ensures that this case never happens and is an easy solution. A borrowing constraint of zero also helps in mimicking the observed distributions of consumption, income and wealth. In cases where we relax the borrowing constraint we ignore any agents who die in debt as their number is fairly small and offset by those who die holding assets.

⁴If agents were allowed to borrow at older ages they would have incentives to borrow, consume what they borrowed, and die in debt with certainty. Making agents pay back all debts before retirement ensures that this does not occur.

component, which are subject to i.i.d. shocks, that determine agents' realized wages. The permanent component, P_{it} , is characterized by:

$$P_{it} = P_{i,t-1} \left(\frac{M_t}{M_{t-1}} \right) \eta_{it} \quad (5)$$

M_t is the economy-wide mean wage in period t and M_{t-1} is the previous period's mean wage. Their ratio, $\frac{M_t}{M_{t-1}}$, gives the deterministic element of wages over the life cycle. The permanent shock, η_{it} , persists into future periods through P_{it} depending on the previous permanent component of wages, P_{it-1} .

The actual wage received by an agent is:

$$w_{it} = P_{it} \varepsilon_{it} \quad (6)$$

This is composed of the permanent component, P_{it} , and a purely transitory shock, ε_{it} , representing a shock without persistence. The effects of ε_{it} , except as they persist through differential asset holdings, disappear after one period.

All agents retire at 65. Although they can choose to retire (by working zero hours) before age 65, they cannot receive Social Security benefits, b_{it}^{SS} , until they reach 65. The Social Security benefits the agent receives are based on a progressive non-linear transformation of the final permanent wage similar to the translation between wages and benefits in the actual system.⁵

⁵Social security wages are based on final permanent wage in our model rather than some form of AIME (Average Indexed Monthly Earnings). Basing benefits on an AIME-type formula would add an additional state variable to the agent's problem and increase the computational complexity of our exercise while providing little additional value in understanding the heterogeneity in the value of life.

2.3 Recursive Representation

Given the economic environment outlined above, we can represent the individual's problem recursively:

$$V_t(A_{it}, P_{it}, \varepsilon_{it}) = \max_{c, l} u(c_{it}, l_{it}) + \beta \phi_t \mathbb{E}[V_{t+1}(A_{i,t+1}, P_{i,t+1}, \varepsilon_{i,t+1})] \quad (7)$$

subject to

$$c_{it} + a_{i,t+1} = a_{it}(1 + R_t) + w_{it}(1 - \tau^{SS})(1 - l_{it}) + b_{it}^{SS} \quad (8)$$

This recursive representation illustrates how solving the problem simplifies to solving a sequence of many one-period maximization problems. The state vector includes three variables: two continuous (the level of the permanent component of wages and the agent's asset holdings) and one discrete (the transitory shock to income). Appendix A presents details on our numerical methods for solving the individual's problem.

2.4 Defining the Value of a Statistical Life

Using the above representation of an agent's problem, we define a discrete time analogue to the more commonly used continuous time definition of mortality risk reduction valuation. Consider, at period t , a small increase in the probability of surviving to the next period $\partial\phi_t$. The probability of survival is now $\phi_t + \partial\phi_t$ instead of ϕ_t . All future survival probabilities remain unchanged so the function $\mathbb{E}[V_{t+1}(A_{i,t+1}, P_{i,t+1}, \varepsilon_{i,t+1})]$ also remains unchanged. Since agents change their savings behavior in response to the change in survival probabilities in that one period, the assets carried forward will be different and the evaluated expectation will also differ. This allows us to compare two distinct value functions representing agents' utility maximization problem over consumption

and leisure. $V_t(A_{it}, P_{it}, \varepsilon_{it})$ is the baseline value function and $V_t^*(A_{it}, P_{it}, \varepsilon_{it})$ reflects the case of the 1-period increase in agents' survival probabilities. The change in the mortality risk the agents face represents the only difference in parameters.

To evaluate the value of the mortality risk reduction we consider how much agents would pay for this increase in their survival probability. Consider the change in wealth, ∂A_{it} , such that

$$V_t(A_{it} + \partial A_{it}, P_{it}, \varepsilon_{it}) = V_t^*(A_{it}, P_{it}, \varepsilon_{it}) \quad (9)$$

This increment ∂A_{it} is the amount of additional wealth that makes the person indifferent between living longer in expectation and living the same length of time as before with that additional wealth. The value of a statistical life typically refers to the aggregate willingness to pay of a sub-population to reduce the mortality risk borne by this sub-population equivalent to one statistical fatality. We construct a VSL from this wealth-risk trade-off by assuming a sub-population of individuals of size $\frac{1}{\partial \phi_t}$ identical to agent i and aggregating the ∂A_{it} over this sub-population. This yields a VSL of:

$$VSL_{it} = \frac{\partial A_{it}}{\partial \phi_t} \quad (10)$$

Note that since our population is heterogeneous in all of A_{it} , P_{it} and ε_{it} the VSLs are also heterogeneous across the population. We report VSLs throughout this paper to facilitate comparison with the literature.

This corresponds to the definition of the VSL in simple one-period models and most revealed preference studies [Viscusi and Aldy, 2003, Johannesson et al., 1997]. For example, hedonic wage models in which the wage is regressed on occupational job mortality risk and other controls yield VSL estimates by

differentiating the estimated wage equation with respect to annual job mortality risk. The VSL is the change in the wage with respect to the change in risk and is given by the estimated coefficient for job mortality risk (for linear specifications and the coefficient scaled by the wage for log-linear specifications) aggregated over a population. This definition corresponds to our one-period change in wealth divided by the one-period increment in the survival probability expressed above.

3 Model Calibration

The parametrization of our model draws on a variety of sources. We employ some parameters from the literature estimated through simulated method of moments and other life-cycle based microeconomic procedures. We estimate the parameters characterizing the wage process with data from the Panel Study of Income Dynamics (PSID) and we employ survival probability estimates from the 2002 U.S. Life Tables.

3.1 Agent's Utility

Based on previous research by, among others, French [2005], ? and Nishiyama and Smetters [2005] we assume that consumption and leisure are non-separable. The period utility function is:

$$u(c, l) = k + \frac{(cl^v)^{1-\gamma}}{1-\gamma} \tag{11}$$

When $v > 1$, consumption and work are complements. At lower levels of leisure the marginal utility of consumption increases. Agents who work a lot also consume a lot. γ is the coefficient of relative risk aversion over the combined utility from consumption and leisure. We assume $\gamma = 2$ and $v = 2$. A value of

2 for γ is quite commonly used in the literature. A value of 2 for v implies that most working agents spend about $\frac{1}{3}$ of their time allotment on market work. These parameters align closely with those estimated for this utility function by French [2005], who finds that this non-separable specification better fits the empirical data than an additively-separable alternative.

We include a constant k in our utility function, which has no effects on decisions regarding savings, consumption, and leisure. Excluding a constant, however, yields negative values of utility for all positive values of consumption and leisure with a coefficient of risk aversion in excess of 1. This implies that the standard normalization that dead agents receive zero utility would result in agents preferring death to being alive. Other recent papers have employed such a constant, motivated by a minimum level of utility from consumption and leisure that makes an individual prefer life to death. For example, ? define a $z = u(c, l)$ which is aggregate of consumption and leisure. They further define z_0 as the minimum level of utility for which the agent prefers life to death. Utility is then $u(z) = \frac{z^{1-\gamma} - z_0^{1-\gamma}}{1-\gamma}$ and extracting the z_0 yields $u(z) = \frac{z^{1-\gamma}}{1-\gamma} - \frac{z_0^{1-\gamma}}{1-\gamma}$. The term $-\frac{z_0^{1-\gamma}}{1-\gamma}$ is positive when $\gamma > 1$, as we assume and is commonly assumed. Our k corresponds to ?'s $-\frac{z_0^{1-\gamma}}{1-\gamma}$. Hall and Jones (2007) also assume a constant in their model, and they elaborate in detail the need to do so in order to generate WTP to reduce mortality risk values with a standard utility function like what we employ here.

We choose k to match the peak VSL estimated in empirical studies. Aldy and Viscusi [2008] have recently estimated workers' life cycle VSL patterns based on labor market hedonic wage analyses over 1993-2000. They estimate a peak VSL of approximately \$8 million and we calibrate k to match that peak. Note that this calibration only implies that the peak estimated in our simulations will be about \$8 million on average and has no bearing on the life-cycle VSL

pattern. The results that follow are robust to different values of k . Changing k acts as a shifter to VSL, moving it up or down, but has little effect on either distributions or patterns of VSL.

3.2 Life-Cycle Pattern and Variance of Wages

We estimate the mean life-cycle pattern of hourly labor compensation using worker data compiled in the PSID over the 1968-1993 period. For the entire population of workers, we find that labor incomes increase rapidly for workers in their 20s and peak and then decline for workers in their 50s. We have also estimated this deterministic life-cycle wage pattern for various sub-populations by race and gender. We draw from the work of Samwick and Carroll to estimate the permanent and transitory shocks to labor income. Table 2 contains the estimates of the variances for the total sample and various subgroups of interest. Appendix B provides more details on the construction of the life-cycle wage profile and the permanent and transitory variances.

3.3 Survival Probabilities

We use in our model survival probabilities from the 2002 U.S. Life Tables constructed by the National Center for Health Statistics [Arias, 2004]. The Life Tables provide us with the age-conditional probability of surviving to the next year for the whole population, by male and female populations, and by black and white populations. Agents are born into our economic environment at age 20 and can live to be at most 100 at which age they die with certainty. Since less than 1 percent of agents survive to 100, truncating life at that age has no significant effects on the results. Also note that the U.S. Life Tables do not report age-specific survival probabilities for those older than 100.

3.4 Social Security Benefits Function

Our approximation to Social Security replaces income based on an agent's final wage. It replaces a minimum of approximately \$7,000 and increases for agents with higher wages until reaching the maximum of about \$16,000. Consistent with the progressive nature of the statutory benefits structure, the replacement rate increases fastest near the bottom of the wage distribution and then levels off as wages increase. These features of Social Security are also important in driving the distributional dynamics of wealth. As shown by Hubbard et al. [1995], Social Security and other government insurance programs drive those with low lifetime earnings to save little. This increases the dispersion of wealth in the economy. In addition, as numerous authors have pointed out, Social Security is a powerful risk-sharing tool.

3.5 Simulated Laboratory Experiments

For each case, we simulate the model with 10,000 agents. Agents in a given population start with the same initial labor market compensation at age 20, experience identical deterministic wage profiles, face the same survival probability profile over the life cycle, and have wage shocks drawn from the same distribution. We introduce the change in the probability of survival for that one period and assess the change in wealth for the alternative value function with the initial probability of survival for each agent necessary to make agents indifferent between the two profiles. This generates the age 20 VSLs for the agents. We then move the agents forward to age 21, subject to an age-specific mortality probability, and replicate the experiment. This proceeds on until age 100 when all agents die. These simulations yield age-specific VSLs for the agents in our model runs. Table 3 presents the experiments that we implemented.

4 Willingness to Pay to Reduce Mortality Risk

In contrast to the deterministic simulation models used previously in this literature, our model with uncertainty in the income process allows us to investigate the distribution in VSLs.⁶ Figure 1 shows the life-cycle VSL profiles for the mean, median, and 5th and 95th percentiles of the 10,000 sample of agents calibrated to the total population.⁷ This figure depicts two stark findings. First, the inverted-U shape is common to most agents across the distribution. The inverted-U takes an even more pronounced shape for those at the upper end of the distribution. For agents who experience a series of very negative shocks, however, the pattern of their VSL can exhibit a quite different pattern as it falls over their life cycle, evident by those at the 5th percentile of the distribution. Second, those agents who experience a sequence of positive permanent labor market shocks have much, much higher VSLs. The mean VSL is nearly 33% higher at its peak than the median for this population, reflecting the effect of the high VSLs on the mean. The VSLs for those at the 95th percentile are an order of magnitude larger than the VSLs for those at the 5th percentile. As the first analysis to account for uncertainty in the labor income process, we show substantial heterogeneity in the value individuals place on mortality risk reduction simply as a function of the labor compensation outcomes over their life cycle.⁸

Figure 2 explicitly illustrates this heterogeneity in the VSL at one point in the life cycle, age 40. The distribution of agents' values is very skewed. Although most agents' mortality risk reduction valuation yields a VSL ranging between \$2 million and \$6 million, the long right tail shows the existence of much higher

⁶Note that Hall and Jones (2007) do not include endogenous labor supply or consider wage shocks in their analysis. Murphy and Topel (2006) do not account for labor income shocks in their composite good Z that aggregates consumption and leisure in their model.

⁷Percentiles are over cross-sections at each age.

⁸Our results are broadly consistent with the empirical results of the spread of estimated values of statistical life in Kniesner et al. [2010] using quantile regression of labor market data.

values. This was the same effect that resulted in the divergence between the mean and the median: the relatively small number of high VSL agents pulls the mean above the median. The mean VSL of about \$8 million at age 40 masks the variation over the population. We also find that the variance of VSL over the life-cycle also follows an inverted-U shape, with higher cross-sectional variances for those agents in their 50s than any other age group. The variance in VSLs declines rapidly starting in the mid-60s because agents are forced into retirement at age 65 and are no longer subject to permanent and transitory wage shocks. Since only the top of the income distribution has substantial amounts of wealth to run down, VSLs decline much more for those at the top of the income distribution than for those at the bottom, resulting in the smaller variance starting at 65.⁹

4.1 Heterogeneity by Age

Figure 3 shows the mean age-specific VSL profile for the agents in this base case. The mean VSL for these agents is about \$3 million at age 20, it increases until its peak of \$8 million in the mid-forties, and steadily declines over the rest of the life cycle. The inverted-U shape of the VSL reflects the findings in Shepard and Zeckhauser [1984] for their “Robinson Crusoe” case in which an agent cannot borrow against future income. The inability to insure against idiosyncratic income risk and the non-negative borrowing constraint in our model constrains agents in much the same way as the borrowing constraint in their Robinson Crusoe case. Our result is driven, however, much more by the inability to insure against income shocks than it is by the borrowing constraint – we have also undertaken simulations with looser borrowing constraints and we still find this inverted-U shape because risk-averse agents are not inclined to borrow

⁹This effect is due in large part to the presence of Social Security. Poor agents have little incentive to save since Social Security benefits replace a large fraction of their consumption.

substantially from future income under the possibility of future adverse income shocks. Our VSL shape mimics the life cycle pattern of consumption, an analytic result of Johansson’s [2002] theoretical model. Our inverted-U shape is generally similar to that of Murphy and Topel (2006), although we find a peak a little later in the life-cycle.

In Figure 3, we have plotted life-cycle hedonic wage VSL estimates from Aldy and Viscusi [2008]. Controlling for both age and cohort effects in panel data of workers, they use a minimum distance estimator to flexibly allow for age variation in the wage-occupational mortality risk tradeoff (VSL). Recall that we calibrated the additive constant, k , in the utility function to the maximum VSL in Aldy and Viscusi [2008], but all other model inputs come from different sources. While we expect our simulations, on average, to match their peak VSL, there is nothing in our calibration that forces our model to match their life-cycle pattern.

Our model achieves results strikingly similar to their estimates, in terms of both the shape of the VSL profile and VSL levels. The two curves track almost exactly from ages 25 to 55. At low ages our model underpredicts VSL by about \$0.5 million, and at older ages our model slightly overpredicts relative to the revealed preference results. Our results also appear similar to the empirical VSL estimates over the life cycle in Kniesner et al. [2006].

4.2 Heterogeneity by Income

This dramatic dispersion in the willingness to pay to reduce mortality risk – within and across age cohorts – reflects the dispersion in incomes as workers’ diverge in their productivities in our simulated experiments. Individuals should be willing to pay more for a given mortality risk reduction as their incomes increase. Viscusi and Aldy [2003] estimate an income elasticity for the value of

statistical life on the order of 0.5-0.6 in a meta-analysis of hedonic wage studies, while Mrozek and Taylor [2002] estimate an elasticity of 0.46 and Liu and Liu [1997] estimate an elasticity of 0.53. In contrast, Kaplow [2005] argues that the income elasticity should exceed 1.0 based on theoretical grounds and empirical estimates of the coefficient of relative risk aversion exceeding 1.0. Murphy and Topel (2006) suggest that the income elasticity should be at least 1.3 given the parameters assumed in their model. Consistent with this, Costa and Kahn [2004] estimate an income elasticity of about 1.5 for a series of U.S labor market cross-sections over 1950-1990. With our model we can explicitly address the question of how the value of life varies with income.

We have estimated income elasticities in separate regressions of the natural logarithm of VSL on the natural logarithm of the permanent component of wages and the natural logarithm of realized income. The two measures differ because agents will choose to adjust their labor-leisure time allocation as labor compensation changes: a percentage increase in realized income will typically exceed the percentage increase in the permanent component of wage because of the complementarity between labor and consumption assumed in our model. This is consistent with the common empirical finding of an upward sloping labor supply function. Since accumulated wealth also plays a role in determining the willingness to pay for mortality reductions, we also control for it in some regressions. The simulated data from our model runs generate more than 400,000 observations for the all working agents sample and at least 9,200 observations for the age-specific working agents samples.

Table 4 presents our estimates of the elasticity of VSL with respect to both the permanent component of wages and the realized labor income of agents. All estimates generate statistically significant positive elasticities for VSL with respect to either realized income or permanent wages. When using income, the

elasticity of VSL is 0.55 but when using productivity this jumps to 1.71. The productivity measure also does a much better job predicting the value an agent places on these mortality risk reductions with an R^2 of 0.935 as opposed to 0.674. The regressions with productivity are thus doing a better job of approximating the true function that relates our state variables to VSL.¹⁰

Columns (C), (D), (E) and (F) add a variety of controls, including level of assets and a cubic in age. Including these controls has only a modest impact on the two VSL-income elasticity measures. The differences between the two measures are statistically significant and are large in economic terms.

Table 5 shows substantial variation in the elasticities by age. Young workers have much higher income elasticities: an age-30 worker has an estimated VSL-income elasticity for the permanent component of wages of 2.01, and this declines to 1.56 for an age-45 worker and still further to 0.81 for an age-60 worker. Likewise, an age-30 worker has a VSL-income elasticity for realized income of 0.82, nearly twice the age-45 estimate of 0.56 and nearly five times the age-60 estimate of 0.15. For young agents the VSL varies much more with changes in either permanent wages or income than for old agents. Since they have many more years of possibly higher consumption (and/or leisure) from their higher income, they should be willing to pay more to maintain this. Young agents also have the ability to stretch each additional dollar out over more years. Since marginal utility of both consumption and leisure declines with more of either, the same amount of extra consumption spread out over a longer time period raises expected utility more. The larger increase in expected utility for an

¹⁰Theoretically, VSL is a four dimensional function of age, assets, the permanent level of productivity and the transitory shock, $VSL_{it} = VSL_{it}(\text{age}, a_{it}, P_{it}, \varepsilon_{it})$. This function can generally be non-linear. In running the regression we are estimating it as a linear function of these state variables. Errors in the linear regression are thus not random error but deviations of the true underlying function from the linear approximation. Adding more controls thus allows us to better approximate the function. The age-specific regressions we perform later are then separately estimating the series of functions $VSL_i = VSL_i(a_i, P_i, \varepsilon_i)$, and age is necessarily dropped.

additional dollar of either permanent wages or realized income drives the larger response of WTP to reduce mortality risk for younger agents.¹¹

The substantial difference between the elasticities with respect to income and permanent wages mirrors the significant variation in estimated VSL-income elasticities in the revealed preference literature. The elasticity with respect to realized income is closer in spirit to the empirical work that employs cross-sectional data to estimate a VSL-income elasticity while the productivity-based VSL elasticity is more akin to estimates that vary across time or over cohorts. Since the cross-sectional studies generally do not take into account leisure in their analysis, they can understate how much a larger wage makes agents willing to pay for extra years of life. The permanent component of wages captures that agents have some stock of future earnings potential that they can take in either consumption or leisure. Since agents value both, an extra dollar of earnings potential is very valuable to the agents. It is important to recognize that a model that does not account for both leisure choice and income shocks (that can generate VSL heterogeneity) would not be able to illustrate this effect.¹²

These differences in the VSL elasticity measures for the realized income and the permanent component of wages provides some numerical and theoretical context for the existing literature. Costa and Kahn [2004] found an income elasticity with a series of U.S. cross-sections over 40 years of 1.5, only slightly less than our 1.67 estimate for the permanent component of wages VSL elasticity

¹¹Assets have the opposite pattern. An extra dollar of assets has a large impact on the VSL for old agents than young ones. Since older agents don't have future labor income which is affected by changes in their permanent component of wages, their main difference is their asset holdings. Higher asset holdings support a higher future level of consumption making a given mortality change that increase the chances of living to enjoy that consumption more valuable.

¹²Evans and Smith [2010] derive a richer model of individual choice over income and risk than the model in Kaplow [2005], and they find how exogenous income shocks affect labor supply decisions and hence the income elasticity of the value of a statistical life. While they do not present estimates of how the income elasticity varies with shocks, they present indirect evidence of how shocks affect labor supply and the coefficient of relative risk aversion in an empirical assessment using the Health and Retirement Study. Our results are generally consistent with the implications of their framework.

(Table 4, column F). Likewise, Kaplow [2005] and ? suggest that the income elasticity should exceed unity. Across cohorts, we should expect such an elastic response of VSL to income changes. A variety of meta-analyses of labor market VSL estimates – summarized, replicated, and extended upon by Viscusi and Aldy [2003] – tend to find much lower VSL income elasticities on the order of 0.5 to 0.6. Most of the studies included in these meta-analyses are based on labor market survey data over a relatively short period of time from the mid-1970s through late-1980s. The variation in income in these studies’ samples likely reflects differences within cohorts – such as from sample composition effects by focusing on specific occupations, specific countries, etc. – and less differences across cohorts. The smaller estimated income elasticities by integrating these VSL results in a meta-analysis would appear to square with our estimate of a 0.55 VSL-realized income elasticity.

To illustrate how the elasticity with respect to permanent wages corresponds more closely to how the VSL changes with wages rising over time we undertook the following experiment. We vary the initial wage of agents entering the economy. An increase in the entering wage persists permanently and increases the permanent component of wages for every agent in the economy. Given this secular increase in permanent incomes, VSLs across the population should also rise and we can consider the elasticity of this with respect to the increase in permanent income.

This approach yields an estimate of the VSL-income elasticity of 1.56.¹³ This is comparable to our estimate of the elasticity with respect to the permanent component of wages presented in Table 5. This elasticity of VSL in response to a secular increase in wages for our model economy supports our interpretation

¹³Formally the procedure is very similar to estimating the cross-sectional elasticities above. We consider increasing the initial wage by 10%, 20%, 50%, 75%, 125% and 200%. We then take the log of the population average of VSL. We then run the regression of this on the log of the initial component of permanent wages. The above estimate is the coefficient on initial permanent wages.

above of a difference between measuring an elasticity with respect to wages and realized labor income.

4.3 Heterogeneity by Ethnicity

In the U.S. economy, white and black workers differ in terms of their labor market returns and expected longevity. We assess the extent to which these differences can impact black and white willingness to pay for risk reduction.

To decompose the effects of differences in income profiles and longevity by race, we present four mean life-cycle VSL patterns in Figure 4. Two of the patterns are black and white life-cycle VSL profiles. The other two reflect a modification of a black population characteristic to its white population representation. In one profile, we have modified the black population calibrated model so that it has the white population's deterministic wage profile and labor market income variances while maintaining black population survival probabilities. In the other profile, we have modified the black population calibrated model so that it maintains the black population's deterministic wage profile and labor market income variances while having the white population survival probabilities.

The modest difference in the initial (age 20) VSLs between blacks and whites reflects the longer longevity and higher expected future incomes for whites. White VSLs peak at over \$9 million in their late forties, nearly double the peak for blacks at a similar age, comparable to the Viscusi [2003] findings. Figure 4 clearly shows that the vast majority of the difference in VSL between blacks and whites results from their differences in wages. Improving the survival rates of black agents to the levels of white agents only yields a small increase in the mean VSL, closing less than ten percent of the gap between peaks. The very modest effect of changing survival probabilities at early stages of life reflects

the effect of discounting. A 20-year old worker expecting to live to 73 (black life expectancy conditional on reaching 20) or 78 (white life expectancy) would not value these differences much since they would occur more than 50 years in the future. The figure does show that moving from black survival probabilities to white survival probabilities does increase in absolute and percentage terms the VSL for those at older ages because they have fewer years over which to discount before enjoying the longer life expectancy. Improving the black wage profile to match the white wage profile closes over ninety percent of the gap between the VSLs.

The difference in wages between blacks and whites appears to drive the differences in VSL that have been observed in the empirical literature. Research by Viscusi [2003] based on hedonic wage models shows that black workers have a value of statistical life equal to about one-half that of white workers. Thus, our analysis suggests that this white-black gap could be explained by differences in labor income and life expectancy, as opposed to unobserved differences in attitudes toward risk.

4.4 Heterogeneity by Gender

Men and women differ in both their labor market compensation (mean and variances) and their life expectancy. Based on the 2002 Life Tables, women should expect to live 5.5 years longer than men. This should increase the VSL of women relative to men. Countering this effect, however, is the higher labor market compensation men experience. We assess the life-cycle VSLs for men and women to discern which of these two effects dominate.

In our simulations, female VSLs peak in their late forties, about five years later than the male VSL peak, reflecting the effects of differences in life expectancy. At their peaks, men have VSLs of about \$8 million, more than 33%

greater than women's VSLs of about \$6 million at their peak. Averaged over the working years of 20-65, the ratio of female-to-male VSLs in our simulations comes to about 0.87, which is modestly larger than the female-to-male VSL ratio of 0.75 in the labor market hedonic analysis by Viscusi and Hersch [2008]. To decompose the effects of differences in income profiles and longevity by gender, we present four mean life-cycle VSL patterns in Figure 5. Two of the patterns are male and female life-cycle VSL profiles. The other two reflect a modification of a female population characteristic to its male population representation. In one profile, we have modified the female population calibrated model so that it has the male population's survival probabilities while maintaining the female deterministic wage profile and labor market income variances. In the other profile, we have modified the female population calibrated model so that it has the male population's deterministic wage profile and labor market income variances while maintaining the female survival probabilities.

The figure illustrates that the vast majority of the differences in VSL between males and females, as in the black-white comparison, derives from differing wage processes. Giving females the shorter lives of males results in only a modest reduction in their VSL. Since wages for females increase at a slower rate than those of males, their VSL also rises less steeply and has a longer plateau than the male VSL. Giving females the wages of males, however makes a substantial difference. With their longer life expectancies, their VSL now surpasses that of males at all ages. The mortality difference makes little difference in the shape with VSLs as males' and females' VSLs have almost the exact same shape once they have the same pattern of wages.

Our results differ from work by Murphy and Topel (2006), which shows female VSLs exceeding male VSLs in their simulation model. This reflects their assumption that males and females only differ according to survival probabili-

ties. Thus, their analysis is effectively similar to a comparison of the life-cycle VSL patterns in Figure 5 for male wages and male life expectancy to male wages and female life expectancy, or a comparison for female wages and female life expectancy and female wages and male life expectancy. Since we do not assume the same income profile for men and women in these experiments, we find a different male-female VSL relationship.

5 Conclusion

We have developed a numerical life-cycle model that is novel to the mortality risk reduction literature by accounting for choice over consumption and leisure and stochastic wage and mortality processes. Our calibrated model matches several important regularities in the U.S. economy, including life cycle consumption patterns, precautionary saving among young adults, wealth accumulation starting in middle age, and labor market compensation and participation outcomes. The realism of the model and the accounting for uncertainty allows us to explore important questions on the value of mortality risk reductions that existing theoretical and simulation models cannot address, and our simulated laboratory experiments serve as a bridge between the theoretical and empirical literatures.

We find that the value agents' place on reductions in mortality risk (or their population-equivalent VSL) varies substantially with income and assets. Accounting for empirically estimated life cycle wage patterns and idiosyncratic labor market compensation shocks yields a wide range of VSL estimates both over the life cycle and across the population. The mean, median, 95th percentile, and variances of the VSL follow an inverted-U shape over the life cycle, peaking in agents' mid-40s – at about the same point in the life cycle as the peak in consumption. Agents who experience very negative labor market shocks have a VSL that declines over their life cycle, as evident by the 5th percentile of

the distribution. The mean VSL substantially exceeds the median as the VSL distribution is non-normal with an extremely fat, long upper end tail. The peak VSLs for agents at the 95th percentile are an order of magnitude larger than the peak VSLs for those at the 5th percentile of the distribution. The inverted-U shape is driven by the idiosyncratic labor market shocks that we assume, consistent with real-world practice, are uninsurable (except through savings-based self insurance). Accounting for these labor market shocks, extending beyond the current VSL literature, allows us to characterize the entire age-specific VSL distribution and not just a single, deterministic series of values.

The elasticity of the VSL with respect to income differs substantially between the wage and realized income. Estimating the elasticity with respect to realized income ignores the substantial value that agents place on leisure. We find that the VSL-income elasticity for all working agents based on the permanent component of the wage is on the order of 1.5, three times greater than the estimate for realized income. We believe that this difference, only evident in such a model with leisure choice and a stochastic labor market compensation process over the life cycle, explains the difference in VSL-income elasticities in the revealed preference literature: the permanent wage results are consistent with cohort-based income elasticities and the realized income results are consistent with cross-sectional results. We also show that the VSL-income elasticity declines significantly with age for both income measures.

We have also assessed how the VSL varies over the life cycle by race and gender. We find that differences in the life cycle wage process and mortality profile drive the substantial differences in VSLs between blacks and whites and males and females. Assuming no preference heterogeneity across demographic groups, we show that virtually all of the black-white gap in recent revealed preference research can be explained by labor market compensation. Although

the longer female life expectancy should increase female VSLs relative to males' VSLs, the higher labor market compensation for men more than offsets this effect.

Our results may have several implications for future research. First, the dispersed and skewed distribution of agents' VSLs indicates that empirical estimates based on assumptions of normality might be biased. Second, these results suggest the possibility of explaining some other observed risk anomalies based on the agents' position in wage and asset distributions. Third, our estimates of an elasticity in response to changes in wages substantially in excess of one implies that as an economy becomes richer more resources might be devoted to health care and other measures to reduce mortality risk.

The substantial heterogeneity in the VSL over the life cycle and across the population could have several implications for the application of VSL estimates to mortality risk reduction policy proposals. First, the life cycle variation in the value of statistical life in this study does not support current EPA practice of a constant VSL irrespective of age or the past FDA practice of a constant value of a statistical life-year irrespective of age. The inverted-U life cycle pattern we find implies that EPA practice overestimates the benefits of mortality risk reduction for the elderly and underestimates the benefits of risk reduction for prime-aged adults and FDA practice overestimates the benefits of mortality risk reduction for young adults. Second, the VSL-income elasticity results can inform assessments of long-term mortality risk reduction policies, such as reducing UV-B exposure through the phase-out of chlorofluorocarbons. Such policies would reduce mortality risk for future generations that would be expected to have higher lifetime incomes and would be willing to pay more for risk reduction. Third, the heterogeneity across the population suggests that a one-size fits all VSL would not appropriately reflect how much individuals would be willing to

pay for risk reduction if they had such an opportunity to do so in a competitive market. For proposed policies that may affect the mortality risk profile for a specific demographic group, population-specific VSLs may be appropriate. Given the negative response to the so-called “senior discount” used by EPA in its 2002 evaluation of the Clear Skies Initiative, this may raise political and ethical concerns (see discussion in Aldy and Viscus [2007]). Such concerns may be entirely legitimate for those contexts in which at least some of the differences in lifetime incomes and wage profiles, and hence estimated VSLs, reflect gender- or ethnic-specific discrimination in labor markets (see discussion in Viscusi [2010]).

References

- J.E. Aldy and W.K. Viscus. Age differences in the value of statistical life: Revealed preference evidence. *Review of Environmental Economics and Policy*, 1(2):241–260, 2007.
- J.E. Aldy and W.K. Viscusi. Adjusting the value of statistical life for age and cohort effects. *Review of Economics and Statistics*, 90(3):573–581, 2008.
- E. Arias. United states life tables, 2002. *National Vital Statistics Report*, 53(6), 2004.
- M. Dickie Blomquist, G.C. and R.M. O’Conor. Willingness to pay for improving fatality risks and asthma symptoms: Values for children and adults of all ages. *Resource and Energy Economics*, 33(2):410–425, 2011.
- T.A. Cameron and J.R. DeShazo. Demand for health risk reductions. *Journal of Environmental Economics and Management*, 65:87–109, 2013.
- Dora Costa and M. E. Kahn. Changes in the value of life, 1940-1980. *Journal of Risk and Uncertainty*, 29(2):159–180, 2004.
- M.F. Evans and V.K. Smith. Measuring how risk tradeoffs adjust with income. *Journal of Risk and Uncertainty*, 40(1):33–55, 2010.
- Eric French. The effects of health, wealth and wages on labor supply and retirement behavior. *Review of Economic Studies*, 72(2):395–427, 2005.
- Pierre-Olivier Gourinchas and Jonathan Parker. Consumption over the lifecycle. *Econometrica*, 70(1):47–89, 2002.
- Robert E. Hall and Charles I. Jones. The value of life and the rise in health spending. *Quarterly Journal of Economics*, 122(1):39–72, 2007.

- R. Glenn Hubbard, Jonathan Skinner, and Stephen P. Zeldes. Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399, 1995.
- M. Johannesson, P. O. Johannsson, and K. G. Lofgren. On the value of changes in life expectancy. *Journal of Risk and Uncertainty*, (15):221–239, 1997.
- P. O. Johannsson. On the value of changes in life expectancy. *Journal of Health Economics*, (15):105–113, 1996.
- P. O. Johannsson. On the definition and age-dependency of the value of a statistical life. *Journal of Risk and Uncertainty*, 25(3):251–263, 2002.
- L. Kaplow. The value of a statistical life and the coefficient of relative risk aversion. *Journal of Risk and Uncertainty*, 31(1):23–34, 2005.
- T. J. Kniesner, W. K. Viscusi, and J. P. Ziliak. Life-cycle consumption and the age adjusted value of life. *Contributions to Economic Analysis and Policy*, 5(1), 2006.
- T. J. Kniesner, W. K. Viscusi, and J. P. Ziliak. Policy relevant heterogeneity in the value of statistical life: New evidence from panel data quantile regressions. *Journal of Risk and Uncertainty*, 40(1):15–31, 2010.
- A. Krupnick. Mortality-risk valuation and age: Stated preference evidence. *Review of Environmental Economics and Policy*, 1(2):261–282, 2007.
- J. D. Leeth and J. Ruser. Compensating wage differentials for fatal and nonfatal injury risk by gender and race. *Journal of Risk and Uncertainty*, 27(3):257–277, 2003.
- J. Hammit Liu, J.-T. and J.-L. Liu. Estimating hedonic wage function and value of life in a developing country. *Economic Letters*, 57:353–358, 1997.

- J. C. Mesa. An object-oriented class library for nonlinear optimization. *Sandia Technical Report*, (SAND94-8225), 3 1994.
- M. J. Moore and W. K. Viscusi. Models for estimating discount rates for long term health risks using labor market data. *Journal of Risk and Uncertainty*, (3):381–401, 1990.
- J.R. Mrozek and L.O. Taylor. What determines the value of life? a meta-analysis. *Journal of Policy Analysis and Management*, 21(2):253–270, 2002.
- K. M. Murphy and R. H. Topel. The value of health and longevity. *Journal of Political Economy*, 114(5):871–904, 2006.
- Shinichi Nishiyama and Kent Smetters. Consumption taxes, risk sharing and economic efficiency. *Journal of Political Economy*, 113(5):1088–1115, 2005.
- Sherwin Rosen. The value of changes in life expectancy. *Journal of Risk and Uncertainty*, 1:285–304, 1988.
- Andrew Samwick and Christopher Carroll. The nature of precautionary wealth. *Journal of Monetary Economics*, (40):41–72, 1997.
- Philip J. Schneider and David H. Eberly. *Geometric Tools For Computer Graphics*. The Morgan Kaufmann Series in Computer Graphics and Geometric Modeling. Morgan Kaufmann Publishers, 2003.
- D. S. Shepard and R. J. Zeckhauser. Survival versus consumption. *Management Science*, 30(4):423–439, 1984.
- V. K. Smith, M. F. Evans, H. Kim, and D. H. Taylor. Do the 'near' elderly value mortality risks differently. *Review of Economics and Statistics*, 86(1): 423–429, 2004.

- Seamus J. Smyth. A balancing act: Optimal non-linear taxation in oligonomies. *Harvard University Working Paper*, 5 2006.
- Kjetil Storesletten, Chris Telmer, and Amir Yaron. Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51:609–633, 2004.
- W. K. Viscusi. Racial differences in the value of statistical life. *Journal of Risk and Uncertainty*, 27(3):239–256, 2003.
- W. K. Viscusi. The value of life: Estimates with risks by occupation and industry. *Economic Inquiry*, 42(1):29–48, 2004.
- W. K. Viscusi and J. E. Aldy. The value of a statistical life: A critical review of market estimates throughout the world. *Journal of Risk and Uncertainty*, 27(2):5–76, 2003.
- W. K. Viscusi and J. E. Aldy. Labor market estimates of the senior discount of the value of statistical life. *Journal of Environmental Economics and Management*, 53(3):377–392, 2007.
- W.K. Viscusi. Policy challenges of the heterogeneity of the value of statistical life. *Foundations and Trends in Microeconomics*, 6(2):99–172, 2010.
- W.K. Viscusi and J. Hersch. The mortality cost to smokers. *Journal of Health Economics*, 27:943–958, 2008.

Tables

Table 1: Summary of Parameter Values, Sources and Main Economic Interpretation

Parameter	Value	Source	Parameter Function
β	0.965	Gourinchas and Parker [2002], French [2005]	Rate of time preference
γ	2.0	Gourinchas and Parker [2002], French [2005]	Risk aversion
v	2.0	French [2005], Nishiyama and Smetters [2005]	Trade-off between leisure and consumption
ϕ_i	varies	Arias [2004]	Survival probability
R	5%	Business cycle literature	Interest Rate
k	8.0	Match VSL peak from Aldy and Viscusi [2008]	Preference for living over dying
Income process	varies	Author's estimation	Earnings pattern over life cycle

Table 2: Permanent and Transitory Variances of Wage Shocks

	Permanent (σ_η^2)	Transitory (σ_ε^2)
Full Sample	.0070	.026
Males	.0067	.028
Females	.0071	.021
Whites	.0067	.027
Blacks	.0062	.026

Table 3: Simulated Laboratory Experiments

Case	Wage Process	Mortality Risk Profile
Base	full sample PSID	all population life table
Black	black sample PSID	black population life table
White	white sample PSID	white population life table
Black-White	black sample PSID	white population life table
White-Black	white sample PSID	black population life table
Female	female sample PSID	female population life table
Male	male sample PSID	male population life table
Female-Male	female sample PSID	male population life table
Male-Female	male sample PSID	female population life table

Table 4: Elasticity of VSL

	(A)	(B)	(C)	(D)	(E)	(F)
$\bar{w}(1-l)$.554 (.0006)		.523 (.0006)		.552 (.0006)	
\bar{w}		1.709 (.0007)		1.656 (.0007)		1.672 (.0006)
a			.106 (.0004)	.041 (.0002)	.1708 (.0008)	.100 (.0003)
Other controls					age cubic	age cubic
R^2	0.674	0.935	0.7142	0.941	0.756	0.9619

Regressions are run on simulated data. There are 408,002 observations as regressions are run for all working age (under 65) agents. The coefficients in the first two rows of the table report the elasticities of VSL with respect to realized income and the permanent component of income respectively. a is the asset holdings of an agent. Some cases control for age using a cubic. All coefficients are significant in these cases. All regressions also contain a constant which is not reported.

Table 5: Elasticity of VSL at Different Ages

	All Working	Age 30	Age 45	Age 60
Elasticity of VSL w/r Permanent Component of Wages	1.672 (.0003)	2.011 (0.001)	1.562 (.002)	.814 (.003)
Elasticity of VSL w/r Realized Income	0.552 (.0006)	.819 (0.005)	.555 (.004)	.147 (.002)
Controls	Wealth, Age Cubic	Wealth	Wealth	Wealth

Regressions are run on simulated data. Each coefficient reported in this table represents the results for a separate regression. Income and productivity are not included in the same regressions. For the cases with all agents there are 408,006 observations. For other cases the number of observations ranges from 9,200 observations to 9,997 observations. Standard errors are in parentheses below the reported elasticities. The R^2 for the “All Working” are presented in Table 4 while those for the permanent component of wages range from 0.993 – 0.998 and those for income from 0.775 – .967.

Figures

Figure 1: Mean, Median, 5th and 95th Percentiles of VSL Over Life Cycle




Figure 2: Distribution of VSL at Age Forty




Figure 3: VSL Over Life Cycle



Figure 4: Breakdown of Sources of Differences in VSL between Blacks and Whites





Figure 5: Breakdown of Sources of Differences in VSL between Males and Females



A Solving the Life-Cycle Problem

Solving the individual's life-cycle problem requires numerical methods. The problem includes three state variables (1) the amount of assets (2) the permanent component of wages and (3) the transitory shock to wages. The particular form of our problem makes it impractical to make some common dimensionality reducing state reductions¹⁴. Lacking these simplifications, we solve the more complicated problem with many state variables.

We represent the agent's problem recursively as:

$$V_t(A_{it}, P_{it}, \varepsilon_{it}) = \max_{c, l} u(c_{it}, l_{it}) + \beta \phi_t \mathbb{E}[V_{t+1}(A_{i,t+1}, P_{i,t+1}, \varepsilon_{i,t+1})] \quad (12)$$

subject to:

$$c_{it} + a_{i,t+1} = a_{it}(1 + R_t) + w_{it}(1 - \tau_{SS})(1 - l_{it}) \quad (13)$$

For each value of the discrete state, the transitory wage shock, there is a grid over assets and wages. Solutions generally use 3 transitory shocks. The grid over assets and wages has 20 points for both assets and wages. Adding more grid points does not significantly change any of the decisions or aggregate quantities of the model. Since the decision rules feature high levels of curvature at low levels of assets and wages and much lower curvature at higher levels we use a non-equally spaced grid. The density of the grid increases at lower levels of both wages and assets to better account for the rapid changes in the value function.

Shape-preserving splines approximate the value function between grid points.

¹⁴An example is the common normalization of using the homothecy of the value function to reduce the state variable to $\frac{A+P}{P}$, the ratio of assets plus permanent income to permanent income. Since we have both a permanent and transitory shock combined with an endogenous labor supply decision we lose this convenient normalization.

The value function inherits the monotonicity underlying the utility function. Agents always prefer more assets and higher wages. This implies that approximation methods which preserve this are preferred. We employ two-dimensional Akima shape-preserving splines, which are relatively fast and efficient [Schneider and Eberly, 2003]. Importantly these splines provide information not only about the function but also its derivative. This is necessary to use the much faster Newton’s method based optimization algorithms described below.

To solve the individual’s problem at each of these grid points we use the OPT++ library of Mesa [1994]. This package solves general non-linear programming problems with arbitrary constraints.

Starting in the last period the solution is simple, do not work and consume everything. We then know the the value function V_{80} (corresponding to age-100 agents) at our grid points. The splines described above approximate the value function over the entire area. Given V_{80} , we then solve for V_{79} taking V_{80} as given and construct splines again. From this the algorithm iterates back to the beginning of an agent’s economic life at age 0 (age-20 agents).

When solving for the optimum in the value function we also get decision rules for consumption and leisure. Since these are again on the grid the same spline procedure approximates them.

B Life-Cycle Wage Estimation

We estimate the life-cycle pattern of labor market compensation from Panel Study of Income Dynamics (PSID) data over 1968-1993. We limit the sample to heads of household between 20 and 65 years of age working at least 1,000 hours per year. We also exclude heads of household with top-coded income. This results in a sample of 107,381 observations for 11,040 households. Labor market compensation measures are converted to constant year 2000 dollars using

the CPI-Urban deflator. We then convert labor market incomes to hourly rates and take the natural logarithm of these values.

We estimate the deterministic life cycle component of labor market compensation by specifying the log of wages as a quintic function of age. This quintic version yields very similar life cycle labor market compensation profiles as more flexible specifications, such as those with age-specific dummy variables, but this more parsimonious parametric form allows for more efficient numerical simulation. Wages increase steeply for young agents between 20 and 30. Wages peak at around age fifty and decline thereafter. We have also estimated such age quintic specifications by gender and by race to complement the results for the total population and facilitate our analysis of the heterogeneity of the value of life across demographic characteristics.

To characterize the variance in labor market compensation as the effect of both permanent and transitory shocks, we employ the variance decomposition method developed by Samwick and Carroll [1997]. For our sample, we first remove the predictable component of labor market compensation. We regress the log of the hourly wage on demographic characteristics, age, education, industry, occupation, interactions of these terms, and year dummy variables. We normalize a worker's wage by constructing the ratio of the actual wage to the predicted wage from this regression. For each worker in a given year we difference labor market incomes by subtracting previous years' incomes from the current year's income:

$$\Delta y_{td} = y_{t+d} - y_t \quad \forall d = 1, \dots, D \quad (14)$$

Having removed the predictable component of labor market income, the square of these differences should represent the variance in labor market income

over time. Samwick and Carroll [1997] show how this can be decomposed into permanent and transitory shocks to income by OLS regression on a worker-by-worker basis of these differences on the time length associated with the difference, d , and the constant 2:

$$\Delta y_{itd}^2 = d\sigma_{\eta}^2 + 2\sigma_{\varepsilon}^2 \tag{15}$$

The coefficient estimates yield the permanent and transitory income shocks. Following Samwick and Carroll, we exclude observations with d of 1 or 2 to ensure that our results do not suffer from serial correlation.