

The Value of Life and the Rise in Health Spending

Robert E. Hall

Hoover Institution and Department of Economics, Stanford University and NBER

E-mail: rehall@stanford.edu

<http://stanford.edu/~rehall>

and

Charles I. Jones

Department of Economics, U.C. Berkeley and NBER

E-mail: chad@econ.berkeley.edu

<http://elsa.berkeley.edu/~chad>

July 30, 2004 — Version 0.75

Preliminary — Please do not quote without permission

Health care extends life. Over the past half century, Americans have spent a rising share of total economic resources on health and have enjoyed substantially longer lives as a result. Debate on health policy often focuses on limiting the growth of health spending. We investigate an issue central to this debate: can we understand the growth of health spending as the rational response to changing economic conditions—notably the growth of income per person? We estimate parameters of the technology that relates health spending to improved health, measured as increased longevity. We also estimate parameters of social preferences about longevity and the consumption of non-health goods and services. The story of rising health spending that emerges is that the diminishing marginal utility of non-health consumption combined with a rising value of life causes the nation to move up the marginal-cost schedule of life extension. In principle, this mechanism can produce a health share that eventually approaches 100 percent if incomes continue to grow. Quantitatively, simulations based on our parameter estimates project that the trending health share will reach about 33 percent by the middle of the century.

1. INTRODUCTION

The United States devotes a rising share of its total resources to health care. The share was 5.2 percent in 1950, 9.4 percent in 1975, and 15.4 percent in 2000. Over the same period, health has improved. The life expectancy of an American born in 1950 was 68.2 years, of one born in 1975, 72.6 years, and of one born in 2000, 76.9 years.

Why has this health share been rising, and what is the likely time path of the health share for the rest of the century? We present a framework for answering these questions. In the model, the key allocative decision is the division of total resources between consumption and health care. Utility depends on quantity of life (life expectancy) and quality of life (consumption). Health spending is valued because it allows people to live longer and better lives. In our analysis, the rise in the health share occurs because of an income effect. As consumption increases with income, the marginal utility of consumption falls rapidly while the value of extending life rises. As agents in our model get richer and richer, consumption rises but they devote an increasing share of resources to health care.

Many of the important questions related to health involve the institutional arrangements that govern its financing—especially Medicare and employer-provided health insurance. One approach would be to introduce these institutions into our model and to examine the allocation of resources that results.

Instead, we take an alternative approach. We examine the allocation of resources that maximizes social welfare in our model. This allows us to abstract from the complicated institutions that shape spending in the United States and to ask a more basic question: from a social welfare standpoint, how much should we spend on health care, and what is the time path of optimal health spending? We look at these issues from two points of view, first under the hypothesis that historical levels of health care were optimal

and second under the hypothesis that they were not. In the second case, we make progress by drawing on the results of a large body of existing research on the value of a statistical life.

We start by documenting the aggregate facts for health spending and life expectancy, the two key variables in our model. We then present a simple stylized model that makes some extreme assumptions but that delivers our basic results. From this foundation, we consider a richer and more realistic framework and develop a full dynamic model of health spending. The remainder of the paper estimates the parameters of the model and shows a number of simulation results.


Related contributions that we will discuss in more detail in future drafts include Grossman (1972), Usher (1973), Shepard and Zeckhauser (1984), Ehrlich and Chuma (1990), and Ehrlich (2000).

2. BASIC FACTS

We will be concerned with the allocation of total resources to health and other uses. We believe that the most appropriate measure of total resources is consumption plus government purchases of goods and services. That is, we treat investment and net imports as intermediate products. Similarly, we measure spending on health as the delivery of health services to the public and do not include investment in medical facilities. Thus we differ conceptually (but hardly at all quantitatively) from other measures that include investment in both the numerator and denominator. When we speak of consumption of goods and services, we include government purchases of non-health goods and services.


Figure 1 shows the fraction of total spending devoted to health care, according to the U.S. National Income and Product Accounts. The numerator is consumption of health services plus government purchases of health services and the denominator is consumption plus total government

FIGURE 1. The Health Share in the United States



Note: The health share is calculated as follows. The numerator is consumption of health services plus government purchases of health services and the denominator is consumption plus total government purchases of goods and services. For further information on sources, see the data section of this paper.

FIGURE 2. Life Expectancy in the United States



Note: Life expectancy at birth data are from Table 12 of National Vital Statistics Report Volume 51, Number 3 "United States Life Tables, 2000", December 19, 2002. Center for Disease Control.

purchases of goods and services. The fraction has a sharp upward trend, but growth is irregular. In particular, the fraction grew rapidly in the early 1990s, flattened in the late 1990s, and resumed growth after 2000.

Figure 2 shows life expectancy at birth for the United States. Following the tradition in demography, this life expectancy measure is not expected remaining years of life (which depends on unknown future mortality rates), but rather represents life expectancy for an individual who faces the cross-section of mortality rates from a given year.

Life expectancy has grown about 1.7 years per decade. It shows no sign of slowing over the 50 years reported in the figure. In the first half of the 20th century, however, life expectancy grew at about twice this rate, so a longer

times series would show some curvature. Whether life expectancy rises linearly or less-than-linearly over time is somewhat open to debate in the demography literature. Oeppen and Vaupel (2002) document a remarkable linear increase in the upper tail of female life expectancy dating back to 1840. See Lee (2003) for an overview of this debate.

3. BASIC MODEL

We begin with a model based on the simple but unrealistic assumption that mortality is the same in all age groups. We also assume that preferences are unchanging over time and income and productivity are constant. This model sets the stage for our full model where we incorporate age-specific mortality and productivity growth. As we will show in Section 4, the stark assumptions we make in this section lead the full dynamic model to collapse to the simple static problem considered here.

The economy consists of a collection of people of different ages who are otherwise identical, allowing us to focus on a representative person. Let x denote the person's state of health, which we will call *health status*. The mortality rate of an individual is the inverse of her health status, $1/x$. Since people of all ages face this same mortality rate, x is also equal to life expectancy. For simplicity at this stage, we assume zero time preference.

Expected lifetime utility for the representative individual is

$$U(c, x) = \int_0^\infty e^{-(1/x)t} u(c) dt = xu(c). \quad (1)$$

That is, lifetime utility is the present value of her per-period utility $u(c)$ discounted for mortality at rate $1/x$. In this stationary environment, consumption is constant so that expected utility is the number of years an individual expects to live multiplied by per-period utility. Here and throughout the paper, we normalize utility after death at zero.

Rosen (1988) pointed out the following important implication of a specification of utility involving life expectancy: When lifetime utility is per-period utility, u , multiplied by life expectancy, the level of u matters a great deal. In many other settings, adding a constant to u has no effect on consumer choice. Here, adding a constant raises the value the consumer places on longevity relative to consumption of goods. Negative utility also creates an anomaly—indifference curves have the wrong curvature and the first-order conditions do not maximize utility. As long as u is positive, preferences are well behaved.

Rosen also discussed the following issue: If the elasticity of utility rises above one for low values of consumption—as it can for the preferences we estimate in this paper—mortality becomes a good rather than a bad. A consumer would achieve a higher expected utility by accepting higher mortality and the correspondingly higher level of later consumption. Thus one cannot take expected utility for a given mortality rate as an indicator of the welfare of an individual who can choose a lower rate. This issue does not arise in our work, because we consider explicit optimization over the mortality rate. An opportunity for improvement of the type Rosen identified would mean that we had not maximized expected utility.

The individual is endowed with a constant flow of resources y that can be spent on consumption or health:

$$c + h = y. \quad (2)$$

The economy has no physical capital or foreign trade that permits shifting resources from one period to another.

Finally, the individual's state of health is determined by health spending according to a health production function:

$$x = f(h). \quad (3)$$

The optimal allocation of resources in this economy involves the choice of consumption and health spending that maximizes the utility of the individual in (1) subject to the resource constraint (2) and the production function for health status (3). The first order condition for the optimal allocation is

$$\frac{U(c, x)}{u'(c)} = \frac{x^2}{f'(h)}. \quad (4)$$

This condition equates the marginal benefit of saving a life, $U(c, x)/u'(c)$, to the marginal cost. The marginal cost of saving a life is dh/dm , the amount of output that must be given up to reduce the mortality rate m by one unit. In terms of our model, it is straightforward to show that $dh/dm = x^2/f'(h)$. Intuitively, the mortality rate is $1/x$, so that saving one life requires x^2 people to increase x by one unit, at a cost of $dh/dx = 1/f'(h)$.

Throughout the paper, we will take the health technology to have the form,

$$x = f(h) = A(zh)^\theta, \quad (5)$$

where z is the productivity of health resources in producing health status and A is a constant. The specification implies diminishing returns to adding resources in health care if θ is below one. Note that $f'(h) = \theta x/h$.

Let $s = h/y$ denote the share of resources devoted to health care, and let $L(c, x) \equiv U(c, x)/u'(c)$ denote the value of a life on the preference side, in consumption units. Dividing by $u'(c)$ converts units of utility into units of consumption. Then the health technology implies that the first-order condition is

$$s = \theta \cdot \frac{L(c, x)/x}{y}. \quad (6)$$

The optimal allocation of resources makes the health share proportional to the value of a year of life L/x as a fraction of per capita income. This

first-order condition together with the production function for health status and the definition of the value of life determine the optimal allocation of resources in the model.

Now suppose we ignore the fact that income and life expectancy are taken as constant in this static model and instead consider what happens if productivity grows. Despite this short-cut of using a static model to answer a dynamic question, the results anticipate the findings in our full dynamic model. With the value of life held constant, equation (6) implies that growing y would cause the health share to decline. In other words, to get a rising health share, it must be the case that the value of a year of life, L/x , rises faster than total resources y .

One final manipulation of the first-order condition is also informative. In this basic model, $L/x = u(c)/u'(c)$. Using this fact, we can rewrite equation (6) as

$$\frac{s}{1-s} = \frac{\theta}{\eta_{uc}}, \quad (7)$$

where $\eta_{uc} \equiv \frac{du}{dc} \frac{c}{u}$ is the elasticity of flow utility with respect to consumption. This form of the optimality condition shows two results. First, the health share is a decreasing function of the elasticity of utility with respect to consumption. For a large class of utility functions, this elasticity declines as consumption rises, which will produce a rising health share in the model. For example, this occurs if $u(c)$ is bounded or if $u(c) = \log c$. Second, because $s = 1 - c/y$, this expression shows that the optimal health share is invariant to the technology parameter z . That is, technical change in the production of health leaves the optimal health share unchanged in this simple formulation. Our framework does not put technical change in the health sector at the center of an explanation for the rising health share; we will discuss this more in the conclusion to the paper.

3.1. An Example: Constant-elastic utility

To illustrate some of the principles of the analysis, we calibrate the simple static model to data for the U.S. for 1950 and 2000. In the calibration, we ignore the fact that the planner in both cases would understand that the economy is not static and also that mortality is far from constant by age. We tackle both of those issues in the next section. It will turn out that the static model is similar to the full dynamic model with age-specific mortality.

Suppose that the flow of utility takes the constant-elastic form with an additive constant:

$$u(c) = b + \frac{c^{1-\gamma}}{1-\gamma}. \quad (8)$$

Concavity requires $\gamma > 0$. When $\gamma > 1$, the realistic case, the base level of utility, b , needs to be positive and large enough to ensure that flow utility is always positive, so that the optimum occurs at an interior point. Notice that in this case $u(c)$ is bounded so that the elasticity η_{uc} is decreasing in consumption.

In the constant-elastic case, it is straightforward to show that the value of a year of life, $L(c, x)/x$, satisfies

$$\frac{L(c, x)}{x} = bc^\gamma - \frac{c}{\gamma - 1}. \quad (9)$$

For $\gamma > 1$, the growth rate of the value of a life year approaches γ times the growth rate of consumption from above. Therefore, the value of a year of life will grow faster than consumption if γ is larger than 1. It turns out that this is enough to get the value of a year of life to grow faster than income as well. For example, if the health share were constant, consumption and income would grow at the same rate. But then the value of a year of life would be growing faster than income, producing a rising health share. According to equation (6), this is exactly what we need in order for the model to generate a rising health share.

Now we pick some parameter values for this utility function and the production function for life expectancy. According to the data underlying the figures at the beginning of the paper, resources per capita rose from 0.343 to 1.0 (as an index based in 2000), health spending rose from 0.018 to 0.154 (in the same units), and life expectancy rose from 68.2 to 76.9 years. We assume that health productivity, z , is normalized to one in the year 2000 and grows at a constant exponential rate; as discussed later in the paper, we assume this growth rate is one percent per year. Fitting the life expectancy and health spending data to our health technology, equation (5), we find parameter values $\theta = 0.0453$ and $A = 83.68$.

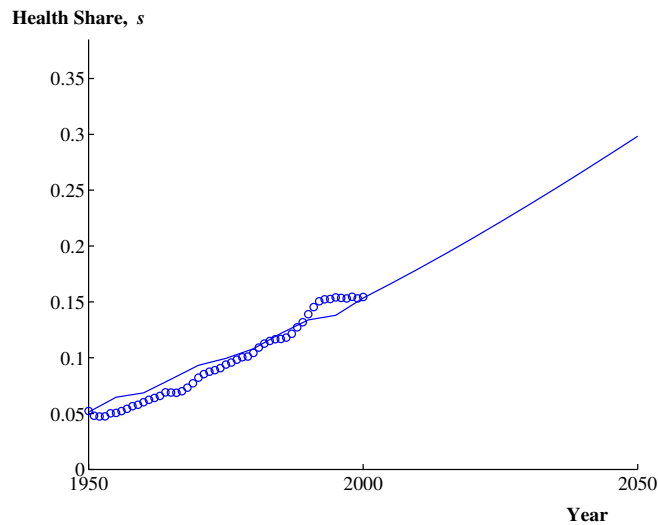
We estimate the curvature parameter, γ , and the base level of utility, b , for the constant-elastic utility function by matching the first-order condition to the data in 1950 and 2000. Using the consumption, health spending, and life expectancy data, we solve the two instances of the first-order condition, for 1950 and 2000, for the two parameters. We find $\gamma = 1.8$ and $b = 6.0$.

Figures 3 and 4 show the results of the model simulation with these parameter values. Because we chose the four parameter values to match the beginning and ending data points for each figure, the simulation results match up nicely with the data. Projecting forward, the basic model suggests a health spending share of nearly 30 percent by the year 2050, when life expectancy reaches 85 years.

The basic mechanism of the rising health share is straightforward. Because γ is greater than one, the value of a year of life rises faster than income. Life becomes increasingly valuable relative to consumption for a society with these preferences.

We can calculate exactly how valuable life becomes by computing $L(c, x)$ in the simulation. For the year 2000, the model implies a value of life of \$7.8 million (in 2000 prices). We will discuss the empirical literature on the value of life in a later section, where we will conclude that this value

FIGURE 3. Simulation with the Basic Model: the Health Share

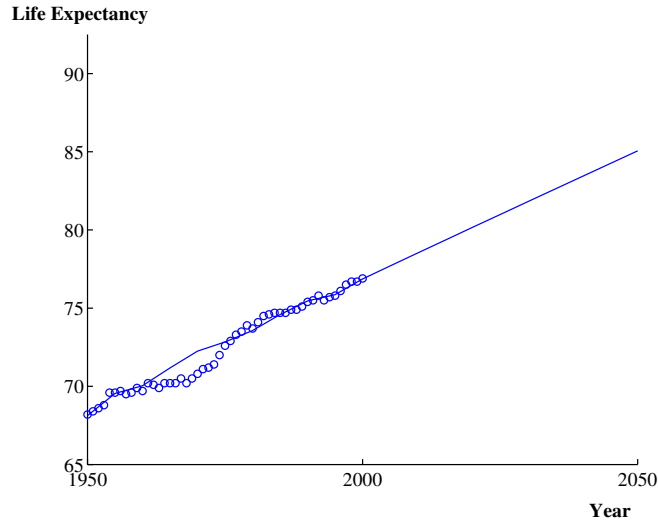


Circles “o” show actual data. The solid line shows the simulated values.

is probably at the upper end of the reasonable range. Taken at face value, the model implies that aggregate health spending in 2000 was at roughly the right level—if we believe the true value of life in 2000 was around \$8 million—or was too high by as much as a factor of three—if instead the true value of life was \$1 million.

The simple model develops intuition, but it falls short on a number of dimensions. First, the model assumes constant total resources and constant health productivity—we have been flouting these assumptions in the simulations by letting these grow. Second, the model assumes a constant mortality rate at all ages. Finally, as we will show, in the full dynamic model with age-specific mortality, the value of life is substantially lower.

FIGURE 4. Simulation with the Basic Model: Life Expectancy



Circles “o” show actual data. The solid line shows the simulated values.

4. THE FULL DYNAMIC MODEL

We turn now to the full dynamic model, allowing age-specific mortality and the associated heterogeneity, as well as growth in total resources and productivity growth in the health sector. An individual of age a in period t has an age-specific state of health, $x_{a,t}$. As in the basic model, the mortality rate for an individual is the inverse of her health status. Therefore, $1 - 1/x_{a,t}$ is the per-period survival probability of an individual with health $x_{a,t}$.

An individual’s state of health is produced by spending on health $h_{a,t}$:

$$x_{a,t} = f_a(z_t h_{a,t}), \quad (10)$$

where z_t is an exogenous productivity factor that converts spending on health $h_{a,t}$ into effective health input. Note that we allow the production function for health to depend on age—mortality varies by age because of

variations in health input and variations in the effectiveness of health inputs in raising health status.

The starting point for our specification of preferences is the flow utility of the individual, $u_{a,t}(c_{a,t})$. Utility depends on both time and age. For simplicity, we assume that the time and age effects are additive, so that

$$u_{a,t}(c_{a,t}) = b_{a,t} + u(c_{a,t}) \quad (11)$$

Here $b_{a,t}$ is the base value of flow utility for a person of age a and $u(c_{a,t})$ is the part that varies with the current quality of life, as measured by the flow of consumption, $c_{a,t}$.

In this environment, we consider the allocation of resources that would be chosen by a social planner who places equal weights on each person alive at a point in time and who discounts future flows of utility at rate β . Let $N_{a,t}$ denote the number of people of age a alive at time t . Then social welfare is

$$\sum_{t=0}^{\infty} \sum_{a=0}^{\infty} N_{a,t} \beta^t (b_{a,t} + u(c_{a,t})). \quad (12)$$

The optimal allocation of resources is a choice of consumption and health spending at each age that maximizes social welfare subject to the production function for health in (10) and subject to a resource constraint we will specify momentarily.

It is convenient to express this problem in the form of a Bellman equation. Let $V_t(N_t)$ denote the social planner's value function when the age distribution of the population is the vector $N_t \equiv (N_{1,t}, N_{2,t}, \dots, N_{a,t}, \dots)$. Then the Bellman equation for the planner's problem is

$$V_t(N_t) = \max_{\{h_{a,t}, c_{a,t}\}} \sum_{a=0}^{\infty} N_{a,t} (b_{a,t} + u(c_{a,t})) + \beta V_{t+1}(N_{t+1}) \quad (13)$$

subject to

$$\sum_{a=0}^{\infty} N_{a,t} (y_t - c_{a,t} - h_{a,t}) = 0, \quad (14)$$

$$N_{a+1,t+1} = \left(1 - \frac{1}{f(h_{a,t})}\right) N_{a,t}, \quad (15)$$

$$N_{0,t} = N_0, \quad (16)$$

$$y_{t+1} = e^{g_y} y_t, \quad (17)$$

$$z_{t+1} = e^{g_z} z_t, \quad (18)$$

and subject to the production function for health status in equation (10). The first constraint is the economy-wide resource constraint. Note that we assume that people of all ages contribute the same flow of resources, y_t . The second is the law of motion for the population. We assume a large enough population so that the number of people aged $a + 1$ next period can be taken equal to the number aged a today multiplied by the survival probability. The third constraint specifies that births are exogenous and constant at N_0 . The final two constraints are the laws of motion for resources and health productivity, which grow at rates g_y and g_z .

Let λ_t denote the Lagrange multiplier on the resource constraint. The optimal allocation satisfies the following first order conditions for all a :

$$u'(c_{a,t}) = \lambda_t, \quad (19)$$

$$\beta \frac{\partial V_{t+1}}{\partial N_{a+1,t+1}} \cdot \frac{f'(h_a)}{x_{a,t}^2} = \lambda_t. \quad (20)$$

That is, the marginal utility of consumption and the marginal utility of health spending are equated across people and to each other at all times. This condition implies that people of all ages have the same consumption c_t at each point in time, but they have different health expenditures $h_{a,t}$ depending on age.

Let $v_{a,t} \equiv \frac{\partial V_t}{\partial N_{a,t}}$ denote the change in social welfare associated with having an additional person of age a alive. That is, $v_{a,t}$ is the social value of life at age a in units of utility. Combining our two first-order conditions,

we get a new version of the optimal allocation condition that we derived before for the static model:

$$\frac{\beta v_{a+1,t+1}}{u'(c_t)} = \frac{x_{a,t}^2}{f'(h_{a,t})}, \quad (21)$$

The optimal allocation sets health spending at each age to equate the marginal benefit of saving a life to its marginal cost. The marginal benefit is the social value of life from the preference side, $\beta v_{a+1,t+1}/u'(c_t)$. The marginal cost is dh/dm , the additional units of output that must be given up to reduce the mortality rate m by one unit to save a life. In terms of our setup with health status x , $dh/dm = x^2/f'(h)$. The mortality rate is $1/x$, so that saving one life requires x^2 people to increase x by one unit, at a cost of $dh/dx = 1/f'(h)$.

Finally, by taking the derivative of the value function, we find that the social value of life satisfies the recursive equation:

$$v_{a,t} = b_{a,t} + u(c_t) + \beta \left(1 - \frac{1}{x_{a,t}}\right) v_{a+1,t+1} + \lambda_t (y_t - c_t - h_{a,t}). \quad (22)$$

The additional social welfare associated with having an extra person alive at age a is the sum of four terms. The first is the base level of flow utility enjoyed by that person. The second is the adjustment for the quality of life associated with the current consumption level. The third is the expected social welfare associated with having a person of age $a+1$ alive next period, where the expectation employs the survival probability $1 - 1/x_{a,t}$. Finally, the fourth term is the net social resource contribution from a person of age a , her production less her consumption and health spending.

4.1. Relation to the Static Model

It is worth pausing for a moment to relate this full dynamic model to the simple static framework. With constant income y , a time- and age-invariant health production function $f(h)$, $\beta = 1$, and a representative individual,

the Bellman equation can be written as

$$V(y) = \max_{c,h} b + u(c) + (1 - 1/f(h))\beta V(y) \text{ s.t. } c + h = y. \quad (23)$$

Given the stationarity of this environment, it is straightforward to see that the value function is just

$$V(y) = \max_{c,h} (b + u(c)) x(h) \text{ s.t. } c + h = y. \quad (24)$$

This is exactly the static model that we postulated earlier, restated in discrete time.

4.2. Summary of the Dynamic Model

Here we pull together all of the elements of the model:

Functional form for health technology:

$$f_a(h_{a,t}) = A_a (z_t h_{a,t})^{\theta_a}. \quad (25)$$

Functional form for flow utility:

$$b_{a,t} = \frac{c_t^{-(\gamma-1)}}{\gamma-1} \quad (26)$$

Recursion for value of life measured in utility units:

$$v_{a,t} = b_{a,t} - \frac{c_t^{-(\gamma-1)}}{\gamma-1} + c_t^{-\gamma} (y_t - c_t - h_{a,t}) + \beta \left(1 - \frac{1}{x_{a,t}}\right) v_{a+1,t+1}. \quad (27)$$

Value of a statistical life in consumption units = marginal cost of a statistical life:

$$\frac{\beta v_{a+1,t+1}}{c_t^{-\gamma}} = \frac{x_{a,t} h_{a,t}}{\theta_a}. \quad (28)$$

5. DATA

We organize the data into 20 five-year age groups, starting at zero and ending at 99. We consider 11 time periods in the historical period, running from 1950 through 2000.

We obtained data on age-specific mortality rates from Table 35 of National Vital Statistics Report Volume 51, Number 3 *United States Life Tables, 2000*, December 19, 2002, Center for Disease Control. This source reports mortality rates every 10 years, with age breakdowns generally in 10-year intervals. We interpolated by time and age groups to produce estimates for 5-year time intervals and age categories.

We obtained data on age-specific health spending from Meara, White and Cutler (2004). These data are for 1963, 1970, 1977, 1987, 1996, and 2000. We distributed national totals for health spending across age categories, interpolated to our 5-year age categories.

We obtained data on national totals from Table 2.5.5 of the revised National Income and Product Accounts of the Bureau of Economic Analysis, accessed at bea.gov on February 13, 2004 (for private spending) and Table 3.15 of the previous NIPAs, accessed December 2, 2003 (for government spending). Data on government purchases of health services are no longer reported in the accounts.

The empirical counterpart for our measure, y , of total resources is total private consumption plus total government purchases of goods and services, from the sources described above.

6. ESTIMATING THE HEALTH PRODUCTION FUNCTION

We begin by assuming a functional form for the production function of health:

$$x_{a,t} = f_a(z_t h_{a,t}) = A_a (z_t h_{a,t})^{\theta_a}. \quad (29)$$

Notice that we allow the parameters of this production function, A_a and θ_a , to depend on age.

We also need to specify how the productivity of the health technology changes over time. Recall that we assume exponential improvement over time at rate g_z :

$$z_t = z_0 e^{g_z t}, \quad (30)$$

and normalize the level of z in the year 2000 to one.


Our procedure is to use outside evidence on g_z and then to estimate A_a and θ_a from data on health spending and mortality. Evidence on the value of g_z is limited to case studies of particular disorders. These studies estimate price declines of more than 1 percent per year, but are selectively chosen. A key contribution of these studies is to adjust for quality change. Cutler, McClellan, Newhouse and Remler (1998) find that the real quality-adjusted price of treating heart attacks declined by 1.1 percent per year between 1983 and 1994. Shapiro, Shapiro and Wilcox (1999) estimate a real annual rate of decline of 1.5 percent for the price of treating cataracts between 1969 and 1994. Berndt, Bir, Busch, Frank and Normand (2000) find that the price of treating acute phase major depression declined at an annual rate of more than 3 percent between 1991 and 1996. See Jones (2003) for the details underlying these numbers. We will consider a benchmark value of $g_z = .01$ and (in a later draft) consider robustness to $g_z = 0$ and 0.02.

To estimate A_a and θ_a , we take logs of the production function in equation (29) and add an error term:

$$\log x_{a,t} = \log A_a + \theta_a \log(z_t h_{a,t}) + \epsilon_{a,t}. \quad (31)$$

Our identifying assumption is that the error, $\epsilon_{a,t}$, has mean zero and no trend. Accordingly, we use a GMM estimator based on the two orthogonality conditions of zero mean and zero covariance with a linear trend variable.

FIGURE 5. Estimates of the elasticity of health status with respect to health inputs




Note: The height of each bar measures the elasticity, θ_a , in the specification $\log x = \log A_a + \theta_a \log(zh)$. Standard errors are shown at the top of the bars.

Because $h_{a,t}$ is strongly trending, the resulting estimator has small standard errors—the trend instrument is strong rather than weak.

Figure 5 shows the GMM estimates of θ_a , the elasticity of health status, x , with respect to health inputs, zh , by age category. The groups with the largest improvements in health status over the 50-year period, the very young and the middle-aged, have the highest elasticities.

Figure 6 shows the actual and fitted values for two representative age groups. Because the health technology has two parameters for each age—intercept and slope—the equations are successful in matching the level and trend of health status. The same is true in the other age categories.

FIGURE 6. Estimation of the parameters of the health technology



Note: The curving lines show actual health inputs zh on the horizontal axis and health status, x , on the vertical axis, for two age groups, 35-40 and 65-70, for the period 1950 through 2000. The smooth lines shows the estimated functions, $\log x = \log A + \theta \log(zh)$.

From these estimates, we can calculate the marginal cost of saving a life at each age. However, before turning to these calculations, it is helpful to have in mind a summary of the empirical literature on the value of a statistical life (VSL).

6.1. Evidence on the Value of a Statistical Life

In evaluating our results, three dimensions of the VSL literature are relevant. We are interested in (i) the level of the VSL, (ii) the rate at which the VSL changes over time, and (iii) how the VSL varies with age.

Most estimates of the level of the value of a statistical life are obtained by measuring the compensating differential that workers receive in more dangerous jobs. Viscusi and Aldy (2003) provide the most recent survey of this evidence and find estimates of the value of a statistical life that range from \$4 million to \$9 million, in year 2000 prices.

Ashenfelter and Greenstone (2004) provide an alternative approach to estimating the VSL. Their research design exploits the fact that states took differential advantage of the relaxation of federal mandatory speed limits that occurred in 1987. They find that a much lower number of 1.5 million dollars (in 1997 prices) represents an upper bound on the VSL, suggesting that various problems including omitted variable bias and selection problems account for the higher estimates in the labor market literature.

How does the value of life change over time? Recall that a rising value of life is crucial in this model to understanding the rising health share. Unfortunately, there is relatively little empirical evidence on changes in the value of life over time.

Costa and Kahn (2003) appear to provide the first estimates from a consistent set of data on changes in the value of life in the United States. They use decennial census data from 1940 to 1980 and estimate the value of a statistical life in 1980 of 5.5 million dollars (in 1990 dollars). Moreover,

they find that this value has been rising over time at a rate equal to between 1.5 and 1.7 times the growth rate of per capita GDP. Hammitt, Liu and Liu (2000) made a similar study for Taiwan, combining a time series of cross-sections, and they estimate an elasticity of the value of a statistical life with respect to per capita GDP of between 2 and 3. Because life expectancy itself grows relatively slowly, these studies therefore support the key requirement in this paper that the value of a year of life as a ratio to per capita income is rising over time, and provide an estimate of how rapidly the rise occurs.

A different approach to estimating changes in the value of life finds the opposite result, however. In addition to surveying the existing literature that estimates the value of life at a point in time, Viscusi and Aldy (2003) also conduct a “meta-analysis” to estimate the elasticity of this value with respect to income. Looking across some 60 studies from 10 countries, they regress the average value of life estimates from each study on a measure of average income from each study and obtain an estimate of the elasticity of the dollar value of life with respect to income of about 0.5 or 0.6, with a 95 percent confidence interval that is typically about 0.2 to 0.8. This appears to be consistent with several other estimates from different “meta-analysis” studies that are also summarized by Viscusi and Aldy.

Some additional insight on this issue comes from looking back at our model. Recall that equation (9) in the simple model suggests that the value of life as a ratio to life expectancy is roughly proportional to consumption raised to the power γ . That is, in units of output, the value of a year of life grows with c^γ . One way of thinking about γ is that it is the inverse of the intertemporal elasticity of substitution, which recent empirical work estimates to be less than one. This suggests that $\gamma > 1$, and in fact the values that Costa and Kahn (about 1.6) and Hammitt, Liu, and Liu (about 2 or 3) find accord well with this interpretation. Kaplow (2003) puzzles over

the low income elasticity estimates from the meta-analysis literature for a similar reason; the recent empirical work by Costa and Kahn and Hammitt, Liu, and Liu helps to resolve this puzzle, we think.

Finally, we turn to evidence on variation in the value of a statistical life by age. Aldy and Viscusi (2003) summarize the existing empirical literature, which primarily consists of contingent valuation studies. They go on to provide their own age-specific estimates using the hedonic wage regression approach. Qualitatively, they support the contingent valuation literature in finding an “inverted-U” shape for VSL by age. Quantitatively, their main finding is that the value of life for a 60-year old is about 5.5 million dollars, while the value of life for 30 to 40-year olds is about 2.5 to 3.0 million dollars, a gradient of about 1/2 across these age groups.

To summarize, we take the following stylized facts from the VSL literature. First, there is substantial uncertainty regarding the level of the VSL: it could be as low as 1.25 million in the late 1980s, but could range much higher to numbers like 5 million or more. These numbers are plausibly interpreted as the value of life at some average age, which we will take to be the 35 to 39-year olds. Second, recent estimates suggest that the VSL grows over time, at a rate something like 1.6 or 2 times the growth rate of income. Finally, it appears that the VSL varies with age in an inverted-U pattern, with a relatively gentle slope, falling by about 1/2 between the ages of 35 and 60.

6.2. The Marginal Cost of Saving a Life

Our estimates of the health production function allow us to calculate the marginal cost of saving a life, given the observed allocation of resources. Recall that this marginal cost is $x^2/f'(h)$. (See, for example, the discussion surrounding equation (21).) With our functional form for the health technology, $f'(h) = \theta x/h$, so the marginal cost of saving a life is hx/θ . It

TABLE 1.
The Marginal Cost of Saving a Life (thousands of 2000 dollars)

Age	1950	1970	1990	2000	Growth Rate
0-5	10	60	310	430	7.7
10-15	140	480	2,060	4,690	7.1
20-25	410	1,310	2,580	3,240	4.1
30-35	310	760	1,460	2,900	4.5
40-45	120	240	780	1,370	4.9
50-55	60	130	490	810	5.4
60-65	30	130	390	680	5.9
70-75	30	100	420	590	6.2
80-85	30	120	470	560	6.1
90-95	30	160	520	570	5.6

Note: The middle columns of the table report the value of a statistical life for various age groups. The estimates are calculated as the marginal cost of extending life, hx/θ , using the estimates of θ given in Figure 6 and using actual data on health spending and mortality by age. Standard errors for these values based on the standard errors of θ_a are small. The "Growth Rate" column reports the average annual growth rate between 1950 and 2000.

is important to realize that this calculation only involves the health production function. For this part of the paper, the preference side of the model is irrelevant.

Table 1 shows this marginal cost of saving a life for various age groups. We can interpret these results in terms of the three findings from the empirical VSL literature. First, the marginal cost of saving the life of a 40-year old in the year 2000 was about 1.4 million dollars. This lines up reasonably well with lower end of the level estimates from the literature, for example, the numbers from Ashenfelter and Greenstone (2004). This would suggest that health spending was at approximately the right level as a whole for this age group in 2000. Alternatively, of course, if one believes the higher estimates of the VSL from the literature, the calculation from Table 1 would suggest that health spending for this group was too low.

Next, the final column of the table shows the growth rate for the marginal cost of saving a life. These growth rates are high, on the order of 5 percent per year or more. By comparison, the empirical VSL literature finds significantly lower growth rates. Taking income growth to be about 1.8 percent per year, for example, the income elasticity from Costa and Kahn (2003) of about 1.6 suggests that the VSL grows at a rate of $1.8 \times 1.6 = 2.9$ percent per year. This implies that the value of life in 1950 or 1960 would have been much higher than the marginal cost of saving a life. Therefore, we may have been spending too little on health prior to the most recent decade, even taking the level of the VSL from the lower end of the estimates.

Finally, in terms of variation by age, the general pattern of the marginal cost of saving a life matches the inverted-U shape found in the VSL literature. Quantitatively, the slope of the relation appears to be slightly steeper than what that literature finds. Another way of saying that is this: if the value of life at age 60 is really as high as half of that at age 30, then we may be spending too little on health at higher ages. Similarly, it appears that the marginal cost of saving a life at very young ages is so low that additional spending at those ages might be worthwhile.

7. ESTIMATING THE PREFERENCE PARAMETERS

We present results for two approaches. The first takes the observed levels of health spending as optimal and estimates the preference parameters. The second estimates preference parameters from the evidence in the empirical VSL literature; it implies that health spending was inefficiently low until recently.

7.1. Estimation Using the First-Order Condition

Our model contains the following preference parameters: the discount factor β , the base levels of flow utility $b_{a,t}$, and the curvature parameter γ . We will assume a period length of 5 years.

We have explored a variety of parametric restrictions on the base utility, $b_{a,t}$. These include making it a constant for all ages and years, making it vary by age, and giving it a trend over time. The evidence in favor of age effects is strong. There is evidence of trends in base utility, but not at the same rate for different age groups. We have not found a useful parametric restriction—candidates such as a set of age effects and a set of time effects result in sufficiently large residuals that the other parameters take on improbable values.

Accordingly, we treat the values of $b_{a,t}$ as parameters themselves, without imposing any restriction. Because there is one of these parameters for each data point, estimation is a matter of solving for the values, not minimizing a GMM norm or other criterion. Further, this means there are not enough equations to estimate the other two parameters, β and γ . We use outside evidence on these parameters before solving for the values of $b_{a,t}$.

With respect to the discount rate, we use the following approach: the Euler equation for consumption with constant-elastic preferences is


$$\frac{c_t}{c_{t-1}} = [\beta(1+r)]^\sigma \quad (32)$$

where σ is the intertemporal elasticity of substitution (see, for example, Hall (1988)). We take its value to be 0.5, in line with a substantial amount of research on this subject over the past two decades. r is the real return to saving, which we take to have the value 0.05—see Hall (2003) and the research cited there. The value of β that reconciles the observed growth rate of consumption per person over the sample period 1950-2000 of 2.08 percent per year is 0.992, or, for the 5-year intervals in the model, 0.963.

For the utility curvature parameter γ , we look to other circumstances where curvature affects choice. First is intertemporal choice, as just discussed. The intertemporal elasticity of substitution, σ , is the reciprocal of the elasticity of marginal utility with respect to consumption. If marginal utility is quite elastic, diminishing marginal utility of future consumption inhibits trading off future for current consumption and the intertemporal elasticity is low. If we make an analogy between valuing future consumption for the purposes of choosing its level and valuing the enjoyment of future life for the purpose of choosing health spending and thus choosing mortality, then we would relate our curvature parameter γ to the reciprocal of σ , suggesting a value of 2. This value is also not too different from the value of 1.8 we calculated in section 3.2.

The second dimension of choice is over uncertain outcomes. Risk-averse consumers will apply a concave utility function to those outcomes. If the utility function has a constant elasticity, the elasticity of marginal utility is the coefficient of relative risk aversion. Lucas (1994) discusses the evidence on this parameter and concludes that 2 is a reasonable value. Again, if we make an analogy between choosing among risky outcomes and evaluating the benefit of the value of future life, we would relate our curvature parameter γ to the coefficient of relative risk aversion, once more suggesting a value of 2.

We should also note that the relatively high values of the base value of life that we find have the implication that our results are not at all sensitive to the value of the curvature parameter γ . Most of the motive to invest in longevity comes from the base value and relatively little from the part of utility that depends on consumption.

FIGURE 7. Estimates of base flow utility, $b_{a,t}$ 

Note: Each line shows the cross section of base levels of utility in a period. The periods cover 5 years each from 1950 through 2000.

Given the data and the values of θ_a , β , and γ , we first calculate the implied value of life in utility units from equation (28):

$$v_{a,t} = \frac{x_{a-1,t-1} h_{a-1,t-1} c_{t-1}^{-\gamma}}{\beta \theta_{a-1}}. \quad (33)$$

and then recover the base levels of utility from a rearranged version of equation (27):

$$b_{a,t} = v_{a,t} + \frac{c_t^{-(\gamma-1)}}{\gamma-1} - c_t^{-\gamma} (y_t - c_t - h_{a,t}) - \beta \left(1 - \frac{1}{x_{a,t}} \right) v_{a+1,t+1}. \quad (34)$$

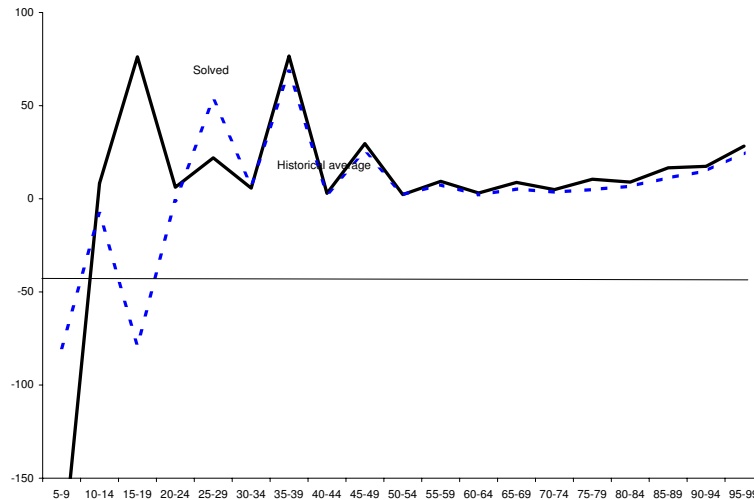
Figure 7.1 shows the results of the calculations. Each line portrays the base level of utility for every age group in a particular year, for the 11

years at 5-year intervals from 1950 through 2000. The lines share a common pattern—negative flow utility in the youngest group and usually in the second-youngest group, and also negative flow utility for teenagers. Negativity of flow utility does not contradict any principles of the model. The motivation for continuing to live is to capture next period's value of life. Negative flow utility marks a difficult period of life that people choose to live through so that they can enjoy later periods with positive utility. For older people, flow utility stabilizes at a common, lower positive level over all periods. Flow utility rises somewhat for the very elderly.

We could also interpret the solved values for $b_{a,t}$ as the residuals from the first-order condition in a model with a constant b . Economically, they arise because the marginal cost of saving a life—the right side of equation (28), with values shown in Table 1—varies considerably more across ages than the value of life on the preference side would in the absence of variation in b . That is, with a constant b across ages, the value of life on the preference side—the left side of equation (28)—turns out to be relatively flat across ages. For example, consider the marginal cost of saving a life reported in Table 1. The only way the model can make sense of the fact that we spend so little on health care for children from 0 to 4 and so much those between 4 and 9 is by having a substantially lower b for the younger group.

These calculations provide estimates of the base level of utility during the historical period. For our projections for the next 50 years, we need future values of the base utility parameters. For this purpose, we make use of additional information, namely the level of the value of life in utility units from equation (33) in the last historical year, 2000. This level information is not used in the calculation of the historical values of $b_{a,t}$ from equation (34), which is in difference form. To make use of the level information, we hypothesize that $b_{a,t}$ will not change over the future from its values in 2000. This hypothesis makes sense, because there is no systematic trend in

FIGURE 8. Comparison of historical average base utilities and values inferred from values of life in 2000



Note: The dashed line is the average of all of the lines in Figure 7.1 and the solid line is the set of base values of utility inferred from the value of life in 2000

the historical values in Figure 7.1. Then we proceed in the following way: When we solve the model for the years 2000 through 2095, we treat b_a as a set of unknowns to solve and then require that the model solution match the value of life in 2000.

Figure 8 compares the results of the two approaches. The solid line infers the future values of b_a from the year-2000 values of life and the dashed line is the average of all of the historical lines in Figure 7.1. Except for the more erratic values for the younger groups, the match is quite good.

7.2. Matching the Value of Life Estimates

For this approach, we drop the assumption that the observed data are generated by maximizing social welfare given our estimated health technology. Instead, we take the age-specific spending data and the consumption data

as given and compute VSL at each age, $\beta v_{a+1,t+1}/u'(c_t)$, from these data. For future values of health spending by age, we project the existing data forward at a constant growth rate.

We keep $\gamma = 2$ as before. We assume that the base value of utility takes the form

$$b_{a,t} = e^{g_b t} \tilde{b} \quad (35)$$

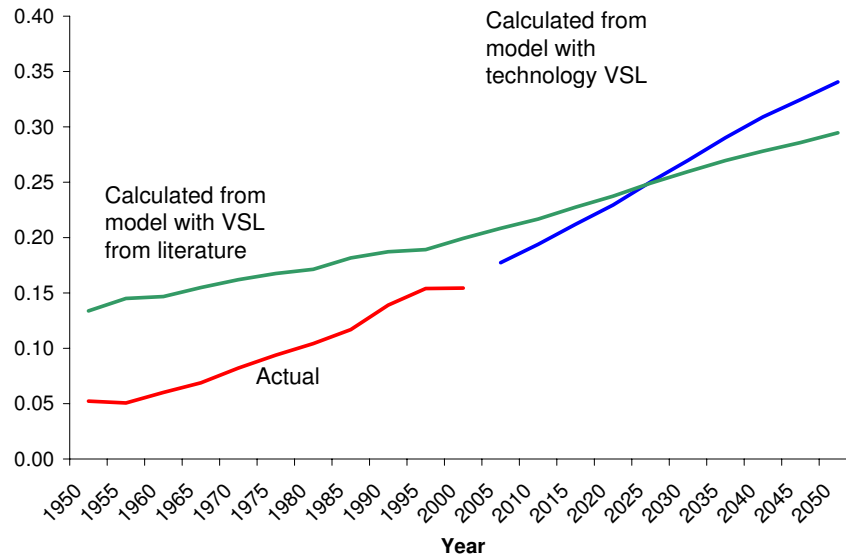
and choose the level \tilde{b} and growth rate g_b to match the estimates from the VSL literature. We start with the estimate from Ashenfelter and Greenstone of a value of life for 35–39 year olds of \$1.25 million in 1987. We project this back to 1950 and forward to 2000, using a growth rate of $1.6 \times 2.31 = 3.70$ percent per year, based on the Costa and Kahn income elasticity. The resulting estimates are $\tilde{b} = 16.85$ (the level in 2000) and $g_b = -0.735$ percent per year.

An alternative approach, which yields results similar to what we will see in the next section, is to assume that b is constant over time and across ages, and choose γ to match the change in the value of life. This yields a value of γ that is approximately 1.8. [to be explored in a later draft].

8. SOLVING THE MODEL

We now solve the model over the sample period 1950 through 2000 and also project the economy out to the year 2050. For the historical period, we take resources per person, y , at its actual value. For the projections, we use the historical growth rate for the sample period, 2.31 percent per year. During the historical period, we use the actual age distribution of the population in the model solution. For the projections, we get the age distribution in year t by applying the model's mortality rate from the previous year to that year's population by age. At the moment, we project births to be a constant 4 million per year and ignore immigration; in the future,

FIGURE 9. Simulation Results: The Health Share of Spending



The line at the lower left shows actual data for the health share. The upward-sloping line for the period 2000-2050 shows the projected health share based on the full model where VSL is inferred from the health technology, as in Table 1. The gently sloping line for the period 1950-2050 shows the hypothetical historical and projected share for preferences inferred from the VSL literature.

we should improve on these assumptions. The details for the numerical solution of the model are discussed in the Appendix.

Figure 9 shows the calculated share of health spending over the period 1952 through 2100. For the historical period, the figure also shows the actual share.

The rising health share observed in the data is a robust feature of the optimal allocation of resources in the health model. For both approaches, we find that the optimal allocation of resources involves a rising health share. The key force at work in the model behind this result is that the marginal utility of consumption falls rapidly. Intuitively, spending an extra

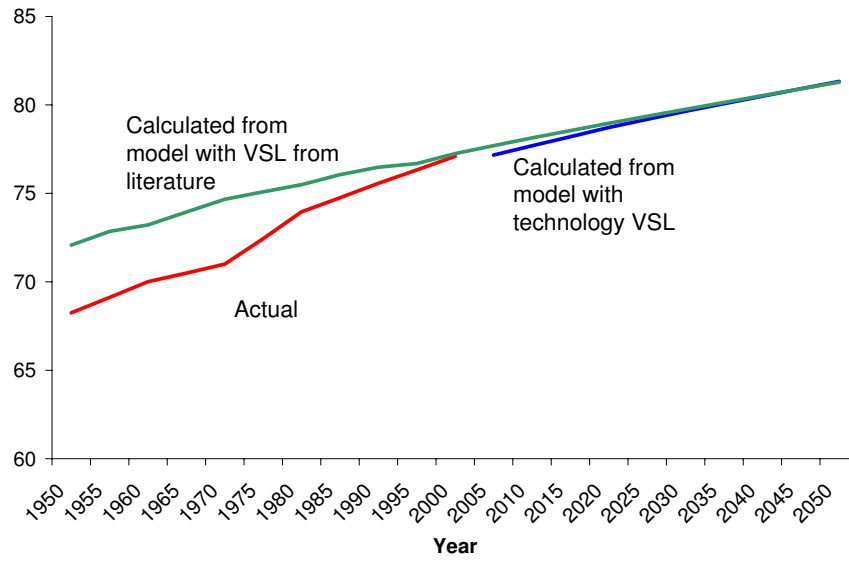
dollar on consumption yields a gain of $u'(c)$, which goes to zero quickly, while spending an extra dollar on health yields a gain of $\beta v f'(h)/x^2$. This term also goes to zero, of course, but if the value of life is growing over time, the decline is slower.

It is important to be clear that this occurs as consumption is rising in all years. Consumption grows, but at a rate that is slower than the rate of income growth. As we get richer and richer, the most valuable thing we have to purchase is more time to live.

The figure shows a substantial difference between projected health shares for the two approaches. Our first approach would match the actual health share between 1950 and 2000 exactly. The projection based on that approach, labeled “technology VSL” in the figure, implies a rapidly growing health share in the future, reaching 34 percent in 2050. The second approach, based on the VSL estimates in the literature, produces a much flatter health share. It suggests underspending on health for the last 50 years. The optimal health share rises gradually in the future, reaching only 28 percent in 2050.

Figure 10 shows the actual and projected levels of life expectancy at birth. For the projection based on inferring the VSL from actual spending and our estimated health technology, the projected path does not grow quite as fast as the historical life expectancy. The slowdown is the result of decline of the elasticity of health status with respect to spending, θ_a , at high ages. It is also interesting to notice how similar the predicted life expectancies are in the year 2050, despite the relatively large differences in health spending observed in Figure 9. The main reason for the similarity is that the pattern of health spending for the projection based on outside estimates of the VSL is more efficient at producing life expectancy than is the one based on the historical pattern of VSL inferred from the age distribution of health spending. In particular, spending on the youngest group of children is

FIGURE 10. Simulation Results: Life Expectancy at Birth



See notes to Figure 9. Life expectancy is calculated using the cross-section distribution of mortality rates at each point in time.

more than four times higher in the projection based on outside estimates. In that projection, we assume equal flow value of life for all age groups, whereas we infer a much lower flow value for young children based on actual spending patterns.

The slight discontinuity in the rise of life expectancy in the figure is the result of an under-prediction of health status in 2000 in our health technology, concentrated in the three youngest age groups.

9. QUALITY OF LIFE

We are currently at work incorporating quality-of-life considerations into the model, replacing the flow of utility $u(c)$ with $u(c, x)$. The flow of utility will depend not only on current consumption, but also on current health

status x . In future versions of the paper, our model will likely include these quality-of-life effects from the beginning (we have worked out the model, but we haven't completed the empirical estimation for this case).

10. CONCLUSION

Our research is still at an early stage. Nevertheless, several key findings are already emerging. The most important of these is that a relatively standard economic model yields a strong prediction for the health share. Provided the marginal utility of consumption falls sufficiently rapidly — as it does for a common specification of preferences featuring a constant intertemporal elasticity of substitution less than one — the optimal allocation of resources involves a health share that rises over time. It is important to realize that this occurs as consumption continues rises; it is just that consumption grows slightly more slowly than income. The intuition for this result is that life is extremely valuable, and as we get richer and richer, the most valuable channel for our spending may be to purchase more time to live. This fundamental mechanism in the model receives support from conventional estimates of the intertemporal elasticity of substitution and from the recent work of Costa and Kahn (2003) and Hammitt et al. (2000), who document that the value of a statistical life appears to be rising faster than income over time.

Given the recent trend in the literature of emphasizing the importance of technological change as an explanation for the rising health share (for example, see Newhouse (1992)), readers may be puzzled that this phenomenon has barely made an appearance here.

There are two keys to reconciling our paper with the technological explanation. First and foremost, the technological explanation only works if preferences accommodate the rising health share. For example, the discovery of new drugs and medical procedures can lead the optimal health share

to rise from 5 to 15 percent of income only if spending 15 percent on health is optimal from the preference side. Both sides of Marshall's scissors must be in favor of high health spending.

Second, it should be noted that it is far from clear that technical change in the health sector promotes rising spending. For example, increases in the productivity of the health sector, z , in our model actually work to reduce the health share: for a given value of life, it becomes cheaper to purchase the effective health input that delivers the right life expectancy. This is not to say that one cannot make the technological explanation work; Jones (2003) is one paper that does this. Rather, it is to point out that (a) one needs a story of preferences no matter what happens on the technology side, and (b) the effect of technological change on the health share is not obvious from a theoretical standpoint.

APPENDIX: SOLVING THE FULL MODEL

We solve the model with age-specific mortality numerically using the following procedure.

1. Begin with a guess for per capita consumption at each date.
2. Guess a terminal value for the value function at each age in the final year of the simulation, 2095; we calculate v from the hypothesis that future variables have the same value as has been projected for 2095.
3. Iterate backwards in time using the first order condition to determine health spending at each age and each point in time.
4. Use this time series for health spending by age to get health status, x and hence age-specific mortality. Iterate forward on the population equation to get the size of the population at each age and date.
5. Finally, for each period, the distribution of the population, the guess for per capita consumption, and the health spending level imply a total quantity

of resources used, while the level of per capita income implies a total quantity of resources available. Let d be the $T \times 1$ vector of discrepancies in the resource constraint.

6. Use a nonlinear equation solver (fsolve in Matlab) to find the time series for consumption that makes the discrepancy vector equal to zero.

We use a similar procedure when we solve the complete model starting in 2000 and match the observed values of v in 2000 from the technology estimates. We start with guesses for both c and b and calculate discrepancies in both material balance and in the values of v in 2000. The nonlinear equation solver then finds values for both c and b .

A complete solution takes about 5 seconds using a Pentium M processor.

REFERENCES

- Aldy, Joseph E. and W. Kip Viscusi, "Age Variations in Workers' Value of Statistical Life," December 2003. NBER Working Paper 10199.
- Ashenfelter, Orley and Michael Greenstone, "Using Mandated Speed Limits to Measure the Value of a Statistical Life," *Journal of Political Economy*, February 2004, 112 (1), S226–S267. Part 2.
- Berndt, Ernst R., Anupa Bir, Susan H. Busch, Richard G. Frank, and Shaon-Lise T. Normand, "The Medical Treatment of Depression, 1991–1996: Productive Inefficiency, Expected Outcome Variations, and Price Indexes," 2000. NBER Working Paper No. 7816.
- Costa, Dora and Matthew Kahn, "Changes in the Value of Life, 1940–1980," May 2003. MIT mimeo.
- Cutler, David M., Mark B. McClellan, Joseph P. Newhouse, and Dahlia Remler, "Are Medical Prices Declining? Evidence from Heart Attack Treatments," *Quarterly Journal of Economics*, November 1998, 113 (4), 991–1024.
- Ehrlich, Isaac, "Uncertain Lifetime, Life Protection, and the Value of Life Saving," *Journal of Health Economics*, May 2000, 19 (3), 341–367.
- and Hiroyuki Chuma, "A Model of the Demand for Longevity and the Value of Life Extension," *Journal of Political Economy*, August 1990, 98 (4), 761–782.
- Grossman, Michael, "On the Concept of Health Capital and the Demand for Health," *Journal of Political Economy*, March/April 1972, 80 (2), 223–255.

- Hall, Robert E., "Intertemporal Substitution in Consumption," *Journal of Political Economy*, April 1988, 96 (2), 339–357.
- , "Corporate Earnings Track the Competitive Benchmark," November 2003. Stanford.edu/rehall.
- Hammitt, James K., Jin-Tan Liu, and Jin-Long Liu, "Survival is a Luxury Good: The Increasing Value of a Statistical Life," July 2000. Harvard University mimeo.
- Jones, Charles I., "Why Have Health Expenditures as a Share of GDP Risen So Much?," July 2003. U.C. Berkeley mimeo.
- Kaplow, Louis, "The Value of a Statistical Life and the Coefficient of Relative Risk Aversion," July 2003. NBER Working Paper 9852.
- Lee, Ronald D., "The Demographic Transition: Three Centuries of Fundamental Change," *Journal of Economic Perspectives*, Fall 2003, 17 (4), 167–190.
- Lucas, Deborah, "Asset Pricing with Undiversifiable Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle," *Journal of Monetary Economics*, 1994, 34 (3), 325–342.
- Meara, Ellen, Chapin White, and David M. Cutler, "Trends in Medical Spending by Age, 1963–2000," *Health Affairs*, July/August 2004, 23 (4), 176–183.
- Newhouse, Joseph P., "Medical Care Costs: How Much Welfare Loss?," *Journal of Economic Perspectives*, Summer 1992, 6 (3), 3–21.
- Oeppen, Jim and James W. Vaupel, "Broken Limits to Life Expectancy," *Science*, May 2002, 296 (5570), 1029–1031.
- Rosen, Sherwin, "The Value of Changes in Life Expectancy," *Journal of Risk and Uncertainty*, 1988, 1, 285–304.
- Shapiro, Irving, Matthew D. Shapiro, and David W. Wilcox, "Measuring the Value of Cataract Surgery," 1999. University of Michigan mimeo.
- Shepard, Donald S. and Richard J. Zeckhauser, "Survival versus Consumption," *Management Science*, 1984, 30, 423–439.
- Usher, Daniel, "An Imputation to the Measure of Economic Growth for Changes in Life Expectancy," in M. Moss, ed., *The Measurement of Economic and Social Performance*, New York: National Bureau of Economic Research, 1973.
- Viscusi, W. Kip and Joseph E. Aldy, "The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World," *Journal of Risk and Uncertainty*, 2003, 27, 5–76.