

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>

November 1, 2004 — Version 2.0

Health care extends life. Over the past half century, Americans spent a rising share of total economic resources on health and 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: Is the growth of health spending 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 account that emerges is that the marginal utility of non-health consumption diminishes faster than the marginal utility of health spending. As a result, the composition of total spending shifts toward health. The health share continues to grow as long as income grows. In projections based on our parameter estimates, the health share reaches 33 percent by the middle of the century.

*We are grateful to David Cutler, Amy Finkelstein, Ron Lee, Tomas Philipson, David Romer, and participants at numerous seminars and NBER meetings for helpful comments. Jones thanks the Center for Economic Demography and Aging at Berkeley for financial support.

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 decision is the division of total resources between health care and non-health consumption. Utility depends on quantity of life (life expectancy) and quality of life (consumption). People value health spending because it allows them to live longer and to enjoy better lives. In our analysis, the rise in the health share occurs because health is a superior good with an income elasticity well above one. Income grows at an approximately constant rate. The marginal utilities of both consumption and health spending fall. But saturation occurs faster in consumption than in health spending. As people grow 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.

We take an alternative approach. We examine the allocation of resources that maximizes social welfare in our model. We abstract from the complicated institutions that shape spending in the United States and ask a more basic question: from a social welfare standpoint, how much should the nation 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.

The recent health literature has emphasized the importance of technological change as an explanation for the rising health share—for example, see Newhouse (1992). According to this explanation, the invention of new and expensive medical technologies causes health spending to rise over time.

Although the development of new technologies unquestionably plays a role in the rise of health spending, the technological explanation is incomplete for at least two reasons. First, expensive health technologies do not need to be used just because they are invented. We investigate whether the social payoff associated with the use of new technologies is in line with the cost. Second, the invention of the new technologies is itself endogenous: Why is the U.S. investing so much in order to invent these expensive technologies? By focusing explicitly on the social value of extending life and how this value changes over time, we shed light on these questions.

We begin by documenting the facts about aggregate 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.

Our approach is closest in spirit to the theoretical papers of Grossman (1972) and Ehrlich and Chuma (1990), who consider the optimal choice of consumption and health spending in the presence of a quality-quantity tradeoff. Our work is also related to a large literature on the value of life and the willingness of people to pay to reduce mortality risk. Classic references


include Schelling (1968) and Usher (1973). Arthur (1981), Shepard and Zeckhauser (1984), Murphy and Topel (2003), and Ehrlich and Yin (2004) are more recent examples that include simulations of the willingness to pay to reduce mortality risk and calculations of the value of life. Nordhaus (2003) and Becker, Philipson and Soares (2004) conclude that increases in longevity have been roughly as important to welfare as increases in non-health consumption, both for the United States and for the world as a whole.

We build on this literature in two ways. First and foremost, the focus of our paper is on understanding the determinants of the aggregate health share. The existing literature focuses on individual-level spending and willingness to pay to reduce mortality. Second, we consider a broader class of preferences for longevity and consumption. Many earlier papers specialize for their numerical results to constant relative risk aversion utility, with an elasticity of marginal utility that is between zero and one. In part, this restriction occurs because these papers do not consider a constant term in flow utility. As we show below, careful attention to the constant is crucial to understanding the rising health share. In particular, when a constant is included, standard utility functions that exhibit a rapidly declining marginal utility of consumption are admissible. This is the key to the rising health share in the model.

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

FIGURE 1. The Health Share in the United States




Note: The numerator of the health share 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 Section 5.

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 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

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.

remaining years of life (which depends on unknown future mortality rates), but is life expectancy for a hypothetical 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. We assume for now that period utility depends only on consumption; in the next section, we will introduce a quality-of-life term associated with health. 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 representative individual receives 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, a health production function governs the individual's state of health:

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

The social planner chooses consumption and health spending to maximize the utility of the individual in (1) subject to the resource constraint (2) and the production function for health status (3). That is, the optimal allocation solves

$$\max_{c,h} f(h)u(c) \quad s.t. \quad c + h = y. \quad (4)$$

The optimal allocation equates the ratio of health spending to consumption to the ratio of the elasticities of the health production function and the flow utility function. Letting $s \equiv h/y$, the optimum is

$$\frac{s}{1-s} = \frac{\eta_h}{\eta_c}, \quad (5)$$

where $\eta_h \equiv f'(h)\frac{h}{x}$, and $\eta_c \equiv u'(c)\frac{c}{u}$.

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 income grows. The short-cut of using a static model to answer a dynamic question anticipates the findings of our full dynamic model quite well.

The response of the health share to rising income depends on the movements of the two elasticities in equation (5). The crux of our argument

is that the consumption elasticity falls relative to the health elasticity as income rises, causing the health share to rise. Health is a superior good because satiation occurs more rapidly in non-health consumption.

Why is η_c decreasing in consumption? In most branches of applied economics, only marginal utility matters. For questions of life and death, however, this is not the case. We have normalized the utility associated with death at zero in our framework, and how much a person will pay to live an extra year hinges on the level of utility associated with life. In our application, adding a constant to the flow of utility $u(c)$ has a material effect—it permits the elasticity of utility to vary with consumption.

Thus our approach is to take the standard constant-elastic specification for marginal utility but to add a constant to the level of utility. In this way, we stay close to the approach of many branches of applied economics that make good use of a utility function with constant elasticity for marginal utility. In finance, it has constant relative risk aversion. In dynamic macroeconomics, it has constant elasticity of intertemporal substitution. In the economics of the household, it has constant elasticity of substitution between pairs of goods.

What matters for the choice of health spending, however, is not just the elasticity of marginal utility, but also the elasticity of the flow utility function itself. With the constant term added to a utility function with constant-elastic marginal utility, the utility elasticity declines with consumption for conventional parameter values. The resulting specification is then capable of explaining the rising share of health spending.

We specify flow utility as:

$$u(c) = b + \frac{c^{1-\gamma}}{1-\gamma}. \quad (6)$$

Based on evidence discussed later in the paper, we consider $\gamma > 1$ to be likely. In this case, the base level of utility, b , needs to be positive and large

enough to ensure that flow utility is always positive. The flow of utility $u(c)$ is then bounded because the exponent on consumption is negative. This means the elasticity η_c is decreasing in consumption. More generally, any bounded utility function $u(c)$ will deliver a declining elasticity, as will the unbounded $u(c) = \alpha + \beta \log c$. Thus the key to our explanation of the rising health share — a marginal utility of consumption that falls sufficiently quickly — is obtained by adding a constant to a standard class of utility functions.

An alternative interpretation of the first-order condition is also informative. Let $L(c, x) \equiv U(c, x)/u'(c)$ denote the value of a life in units of output. Then, the optimal allocation of resources can also be characterized as

$$s = \eta_h \cdot \frac{L(c, x)/x}{y}. \quad (7)$$

The optimal health share is proportional to the value of a year of life L/x divided by per-capita income. If the flow of utility is given as in equation (6), it is straightforward to show that the value of a year of life satisfies

$$\frac{L(c, x)}{x} = bc^\gamma - \frac{c}{\gamma - 1}. \quad (8)$$

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 (and income) if γ is larger than 1. According to equation (7), this is one of the key ingredients needed for the model to generate a rising health share.

A rapidly-declining marginal utility of consumption leads to a rising health share provided the health production elasticity η_h does not itself fall too rapidly. For example, if the marginal product of health spending in extending life were to fall to zero — say it was technologically impossible to live beyond the age of 100 — then health spending would cease to

rise at that point. As we discuss later, for the kind of health production functions that match the data, the production elasticity is either constant or declines very gradually, and the declining marginal utility of consumption does indeed dominate, producing a rising health share.

Finally, we can also generalize the utility function to $U(c, x)$ in place of $xu(c)$, so that lifetime satisfaction is not necessarily proportional to the length of the lifetime. The solution for this case is $s/(1-s) = \eta_h \eta_x / \eta_c$, where $\eta_x \equiv U_x x / U$ is the elasticity of utility with respect to life expectancy. Our result, then, is that the health share rises when the consumption elasticity falls faster than the product of the production and life expectancy elasticities. As just one example $U(c, x) = x^\alpha u(c)$ delivers a constant η_x even with sharply diminishing returns to life expectancy (that is, α close to zero), so our main results are unchanged in this case.

The simple model develops intuition, but it falls short on a number of dimensions. Most importantly, the model assumes constant total resources and constant health productivity. This means it is inappropriate to use this model to study how a growing income leads to a rising health share, the comparative static results notwithstanding. Still, the basic intuition for a rising health share emerges clearly. The health share rises over time as income grows if the joy associated with living an extra year does not diminish as quickly as the marginal utility of consumption.

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. This model also incorporates a quality-of-life component associated with health spending.

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 $h_{a,t}$:

$$x_{a,t} = f_a(z_t h_{a,t}), \quad (9)$$

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, $u_{a,t}(c_{a,t}, x_{a,t})$. In addition to depending on consumption, flow utility depends on health status, $x_{a,t}$. Spending on health therefore affects utility in two ways, by increasing the quantity of life through a mortality reduction and by increasing the quality of life.

For reasons that will become clear in the empirical section, we also allow flow utility to depend on both time and age. For simplicity, we assume the time and age effects are additive, so that

$$u_{a,t}(c_{a,t}, x_{a,t}) = b_{a,t} + u(c_{a,t}, x_{a,t}) \quad (10)$$

Here $b_{a,t}$ is the base value of flow utility for a person of age a and $u(c_{a,t}, x_{a,t})$ is the part that varies with the current consumption and health status. Furthermore, we assume the invariant part of the utility function takes the following form:

$$u(c_{a,t}, x_{a,t}) = \frac{c_{a,t}^{1-\gamma}}{1-\gamma} + \alpha \frac{x_{a,t}^{1-\sigma}}{1-\sigma}, \quad (11)$$

where γ , α , and σ are all positive. The first part of this function is the standard constant-elastic specification for consumption. We assume further

that health status and consumption are additively separable in utility and that quality of life is a constant-elasticity function of health status.

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

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

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 (9) 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

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

subject to

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

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

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

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

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

and subject to the production function for health status in equation (9). 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 :

$$u_c(c_{a,t}, x_{a,t}) = \lambda_t, \quad (19)$$

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

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 together with the additive separability of flow utility 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 the two first-order conditions, we get:

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

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 sum of two terms. The first is the social value of life from the preference

side, $\beta v_{a+1,t+1}/u_c$. The second is the additional quality of life enjoyed by people as a result of the increase in health status.

Marginal cost is the squared value of x divided by the marginal product of health spending h in producing changes in x . Because the reason for the presence of the square of x may not be obvious, we will spell out the details. In brief, it arises from translating changes in health status x into mortality $1/x$.

The marginal cost of saving a statistical life is dh/dm , where dh is the increase in resources devoted to health care and dm is the reduction in the mortality rate. For example, if reducing the mortality rate by .001 costs \$2000, then saving a statistical life requires $1/.001 = 1000$ people to undertake this change, at a cost of \$2 million. Our model contains health status x as an intermediate variable, so it is useful to write the marginal cost as $\frac{dh}{dm} = \frac{dh/dx}{dm/dx}$. Since health status is defined as inverse mortality, $m = 1/x$ so that $dm = dx/x^2$. In the previous paragraph, we required $1/dm$ people to reduce their mortality rate by dm to save a life. Equivalently, setting $dx = 1$, we require x^2 people to increase their health status by one unit in order to save a statistical life. Since the cost of increasing x is $dh/dx = 1/f'(x)$, the marginal cost of saving a life is therefore $x^2/f'(x)$.

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

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

The additional social welfare associated with having an extra person alive at age a is the sum of three terms. The first is the level of flow utility enjoyed by that person. The second 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 last 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 flow utility function that depends only on consumption, the Bellman equation for a representative agent can be written as

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

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

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

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

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:

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

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 :

$$z_t = z_0 e^{g_z t}, \quad (26)$$

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 consider a benchmark value of $g_z = .01$. This value implies that technical change in the health sector has proceeded at a rate that is one percentage point faster than technical change in the rest of the economy.


Because health spending has grown at an annual rate of about 4.5 percent, our assumption of $g_z = .01$ amounts to assuming that increased health spending accounts for just over 80 percent of the average decline in mortality between 1950 and 2000. We considered increasing g_z to capture a trend in other determinants of mortality, but we could not find evidence of any other determinants that might have trended consistently over this period. One natural candidate is changes in smoking behavior, but this is probably better captured as first increasing and then decreasing over our sample period and therefore would not explain much of the overall trend.

To estimate A_a and θ_a , we take logs of the production function in equation (25) and add an error term:

$$\log x_{a,t} = \log A_a + \theta_a \log(z_t h_{a,t}) + \epsilon_{a,t}. \quad (27)$$

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 3. 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 $\log x = \log A_a + \theta_a \log(zh)$. The ranges at the top of the bars indicate \pm two standard errors.

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

Figure 3 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. The fact that the estimates of θ_a generally decline with age, particularly at the older ages, constitutes an additional source of diminishing returns to health spending as life expectancy rises.

Figure 4 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 4. 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 (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 (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 (8) 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 30 to 40-year old is about \$5.5 million while the value of life for a 60-year old is about \$2.5 to \$3.0 million, a gradient of about 1/2 across these age groups. Given their emphasis on the labor market, however, Aldy and Viscusi do not comment on the value of life for children.

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, from the discussion surrounding equation (21), that this marginal cost is $x^2/f'(h)$. With our functional form for the health tech-

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

Age	1950	1980	2000	Per Year of Life Saved 2000	Growth Rate 1950–2000
0-5	10	120	430	6	7.7
10-15	140	970	4,690	73	7.1
20-25	410	1,270	3,240	59	4.1
30-35	310	1,100	2,900	64	4.5
40-45	120	520	1,370	38	4.9
50-55	60	250	810	30	5.4
60-65	30	220	680	36	5.9
70-75	30	210	590	51	6.2
80-85	30	250	560	93	6.1
90-95	30	290	570	267	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 4 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 “Per Year of Life Saved” column divides the cost of saving a life by life expectancy at that age. The “Growth Rate” column reports the average annual growth rate between 1950 and 2000.

nology, $f'(h) = \theta x/h$, so the marginal cost of saving a life is hx/θ . It 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. 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.

The second-to-last column of the table provides an alternative view of the cost of the marginal cost of saving a life by putting it on a “per year of life saved” basis. That is, it shows the cost of saving a statistical life in the year 2000, divided by life expectancy at each age. For example, the marginal cost of saving an extra year of life at age 50 is about \$30,000. Interestingly, the cost of saving a life year in the youngest age category is only about \$6,000, while the cost for saving a life year for the oldest ages rises to well above \$100,000. The marginal cost of saving a life at very young ages is so low that additional spending at those ages might be worthwhile.

Finally, the last 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 2 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 $2 \times 1.6 = 3.2$ 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, the U.S. 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.

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 the end of the 20th century.

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 $b_{a,t}$, the consumption parameter γ , the quality-of-life parameter σ , and the weighting parameter α . For the moment, we consider the case where health status does not affect flow utility so that $\alpha = 0$. We will reintroduce this quality-of-life consideration shortly.

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

$$\frac{c_t}{c_{t-1}} = [\beta(1+r)]^\sigma \quad (28)$$


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.

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.

Given the data and the values of θ_a , β , and γ , we first calculate the implied value of life from equation (21) and then recover the base levels of utility from a rearranged version of equation (22). Figure 5 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 though 2000. The lines share a common pattern—negative flow

FIGURE 5. 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.

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 (21), 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 (21)—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 on those between 5 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 (21) 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 (22), 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 the historical values in Figure 5. Then we proceed in the following way:

FIGURE 6. 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 5 and the solid line is the set of base values of utility inferred from the value of life in 2000

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 6 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 5. Except for the more erratic values for the younger groups, the match is quite good.

7.2. Matching the Earlier Value of Life Estimates

As an alternative approach to estimating the preference parameters, 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

the value of life at each age, $\beta v_{a+1,t+1}/u_c$, from these data. For future values of health spending by age, we project the existing data forward at a constant growth rate. Until the year 2020, this growth rate is the average across the age-specific spending growth rates. After 2020 we assume spending grows at the rate of income growth. The rate must slow at some point; otherwise the health share rises above one. Our results are similar if we delay the date of the slowdown to 2050.

We then estimate a constant and common value $b_{a,t} = b$ and the curvature parameter γ to match some 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. By matching the value of life for this age group in 1950 and 2000, we obtain $b = 13.90$ and $\gamma = 1.584$ for the case where health status does not affect flow utility (i.e. $\alpha = 0$). Finally, we recalibrate the time discount factor β to an interest rate of 5 percent based on this new value of γ .

7.3. The Quality-of-Life Parameters

We know of no empirical literature that allows us to determine values for the quality-of-life parameters σ and α . We therefore calibrate these parameter values in the following way. Consider the following question: What fraction of consumption would a 60 year old be willing to give up in order to have the quality of life implied by the health status of a 35 year old? Similarly, what fraction of consumption would an 85 year old be willing to give up in order to have the quality of life implied by the health status of a 60 year old?

In our model, the answer to this question is given by the value of τ that satisfies

$$\frac{c^{1-\gamma}}{1-\gamma} + \alpha \frac{x^{1-\sigma}}{1-\sigma} = \frac{((1-\tau)c)^{1-\gamma}}{1-\gamma} + \alpha \frac{\tilde{x}^{1-\sigma}}{1-\sigma}, \quad (29)$$

where we assume that \tilde{x} is the new health status “purchased” by giving up consumption.


To calibrate our parameters, we assume the 85 year old would be willing to give up 60 percent of her consumption to have the health status of a 60 year old, and the 60 year old would give up 20 percent of her consumption to have the health status of a 35 year old. At a baseline value of $\gamma = 2$, these data points yield parameter values of $\sigma = 1.64$ and $\alpha = 2.05$. As a check, these parameter values imply that a 35 year old would give up just under 6 percent of her consumption to have the health status of a 20 year old.

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. We solve the model using both of our approaches to calibrating the preferences parameters (the $b_{a,t}$ and γ) and using two approaches to the quality of life ($\alpha = 0$ and $\alpha > 0$). When we recalibrate our other preference parameters to the $\alpha > 0$ case, we find in the first instance values of $b_{a,t}$ that look very similar to those shown in Figure 5. In the second instance where we match the literature’s value of life estimates, we obtain $b = 14.516$ and $\gamma = 1.575$, values very close to the original case of $\alpha = 0$.

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,

FIGURE 7. Simulation Results: The Health Share of Spending



Circles “o” show actual data for the health share. The upward-sloping lines for the period 2005-2050 show the projected health share based on the full model where the VSL is inferred from the health technology, as in Table 1. The gently sloping lines for the period 1950-2050 show the hypothetical historical and projected share for preferences inferred from the VSL literature. Within these two approaches, the upper line corresponds to the case that includes a quality-of-life term ($\alpha > 0$), while the lower line does not ($\alpha = 0$).

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. We project births to be a constant 4 million per year and ignore immigration. The details for the numerical solution of the model are discussed in the Appendix.

Figure 7 shows the calculated share of health spending over the period 1950 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 the U.S. gets richer and richer, the most valuable thing people can 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 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 27 percent in 2050.


Figure 8 examines the micro data underlying the health share. This figure shows actual and simulated health spending by age, for 1950, 2000, and 2050 for our second approach (in the first approach, actual and simulated spending are equated by construction). A comparison of the results for the year 2000, shows that actual and optimal spending are fairly similar for most ages, with two exceptions. Optimal health spending on the youngest age group is substantially higher than actual spending: given the high mortality rate in this group, the marginal benefit of health spending is very high, as was shown earlier. Similarly, while optimal health spending generally rises until age 80, it declines after that point. It is worth noting in this respect

FIGURE 8. Health Spending by Age



Note: Circles denote actual data and solid lines show simulation results for the case in which b and γ are chosen to match estimates from the VSL literature.

FIGURE 9. Simulation Results: Life Expectancy at Birth



See notes to Figure 7. Life expectancy is calculated using the cross-section distribution of mortality rates at each point in time.

that the underlying micro data we use for health spending group all ages above 75 together, so we do not know what the actual pattern of spending looks like above the age of 75.

Figure 9 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 historical life expectancy. The slowdown is the result of the decline in 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 7. The reason is again related to the relatively sharp diminishing returns associated with health spending.

9. CONCLUSION

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 widely-used specification of preferences featuring a constant intertemporal elasticity of substitution less than one—the optimal health share rises over time. The rising health share occurs as consumption continues to rise, but consumption grows more slowly than income. The intuition for this result is that life is valuable, and as people gets richer, the most valuable channel for spending is to purchase additional years of life. People become saturated in non-health consumption, driving its marginal utility to low levels.

This fundamental mechanism in the model is supported empirically in a number of different ways. First, as discussed earlier, it is consistent with conventional estimates of the intertemporal elasticity of substitution. Second, the mechanism predicts that the value of a statistical life should rise faster than income over time; Costa and Kahn (2003) and Hammitt et al. (2000) find this to be the case. Cross-country evidence also suggests that health spending rises more than one-for-one with income; this evidence is summarized by Gerdtham and Jonsson (2000).

One source of evidence that runs counter to our prediction is the micro evidence on health spending and income. At the individual level within the United States, for example, income elasticities appear to be substantially less than one, as discussed by Newhouse (1992). A serious problem with this existing evidence, however, is that health insurance limits the choices facing individuals, potentially explaining the absence of income effects. Our model makes a strong prediction that if one looks hard enough and

carefully enough, one ought to be able to see income effects in the micro data. Future empirical work will be needed to judge this prediction. A suggestive informal piece of evidence is that exercise seems to be a luxury good: among people with sedentary jobs, high wage people seem to spend more time exercising than low wage people, despite the higher opportunity cost of their time.

As we mentioned in the introduction, the recent health literature has emphasized the importance of technological change as an explanation for the rising health share. We believe that our approach is complementary to the technological view. For example, one could imagine that the production function for health status is $x_t = f(h_t)$ if $f(h_t) < \bar{x}_t$ but is limited to \bar{x}_t at higher levels of spending. The discovery of new medical technologies then shifts out \bar{x}_t over time. Implicitly, then, we are assuming that the medical technology shifts out rapidly enough so that this technology constraint is not binding. In this respect, Jones (2003) considers the alternative extreme where the technology constraint always binds.

What our model shows is that the technological explanation only works if preferences accommodate the rising health share. The discovery of new drugs and medical procedures can lead the optimal health share to rise from 5 to 15 percent only if spending at least 15 percent on health is optimal from the preference side. Both blades of Marshall's scissors must be in favor of high health spending. Our model explains why this might be the case. Moreover, the technological view provides little in the way of guidance as to the future of the health share. In contrast, our approach helps us to understand how the health share will evolve regardless of whether or not the technological constraint continues to bind.

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.
- Arthur, W. B., "The Economics of Risk to Life," *American Economic Review*, March 1981, 71 (1), 54–64.
- 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.
- Becker, Gary S., Tomas J. Philipson, and Rodrigo R. Soares, "The Quantity and Quality of Life and the Evolution of World Inequality," *American Economic Review*, forthcoming 2004, pp. xxx–xxx.
- 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 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.
- and Yong Yin, "Explaining Diversities in Age-Specific Life Expectancies and Values of Life Saving: A Numerical Analysis," September 2004. NBER Working Paper 10759.
- Gerdtham, Ulf-G. and Bengt Jonsson, "International Comparisons of Health Expenditure: Theory, Data and Econometric Analysis," in Anthony J. Culyer and Joseph P. Newhouse, eds., *Handbook of Health Economics*, North Holland, 2000.
- 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.
- Murphy, Kevin M. and Robert Topel, "The Economic Value of Medical Research." In *Measuring the Gains from Medical Research: An Economic Approach* (Murphy and Topel, eds 2003).
- and —, eds, *Measuring the Gains from Medical Research: An Economic Approach*, Chicago: University of Chicago Press, 2003.
- Newhouse, Joseph P., "Medical Care Costs: How Much Welfare Loss?," *Journal of Economic Perspectives*, Summer 1992, 6 (3), 3–21.
- Nordhaus, William D., "The Health of Nations: The Contribution of Improved Health to Living Standards." In Murphy and Topel, eds (2003).
- 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.
- Schelling, Thomas C., "The Life You Save May Be Your Own," in Jr. Samuel B. Chase, ed., *Problems in Public Expenditure Analysis*, Washington D.C.: Brookings Institution, 1968, pp. 127–161.
- 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.