When GDP Misleads: Inferring Living Standards from the Value of a Statistical Life

Philip Trammell — Stanford University

Charles I. Jones* — Stanford GSB and NBER

February 25, 2026 — Version 0.2

Preliminary and Incomplete

Abstract

Real GDP per person is a widely used proxy for living standards, but it can be a poor welfare measure when new goods or quality improvements matter, when nonmarket goods are significant, and when preferences are nonhomothetic — all of which are true in practice. We propose an alternative that is robust to these concerns: under weak conditions, the growth rate of the value of a statistical life (VSL), together with standard Euler-equation objects, identifies the growth rate of lifetime utility. The intuition is that people routinely trade off consumption against mortality risk, and their willingness to pay for small risk reductions reveals the value of remaining lifetime utility. Implementing this approach for the United States suggests that lifetime utility may have risen by more than a factor of five since 1940, whereas conventional consumption-based calculations using a stable log/CRRA flow utility imply much smaller gains, on the order of a doubling.

*We are grateful to Sebastian Di Tella and Chris Tonetti for helpful discussions.

1. Introduction

The richest person in the world in the 1830s was Nathan Rothschild, whose personal net worth was around 0.6% of British national income (Ferguson, 2008, p.82). Despite this vast fortune, Rothschild died at age 58 in 1836 of an infection that $10 of antibiotics could likely cure today (Economist, 2004). Similarly, the richest people in the world today, such as Bill Gates or Jeff Bezos, presumably have a marginal utility from additional consumption spending that is zero. Nevertheless, their utility still increases when new goods (smartphones or LLMs) are invented. These examples suggest that more consumption of a fixed set of goods eventually hits a marginal utility of zero while the invention of new goods or higher quality goods continues to increase wellbeing.

At a macro level, our conventional measure of living standards is real GDP per person. Of course, it is well-known that real GDP is an imperfect measure of living standards along many dimensions. For example, the welfare gains from new goods and expanded product variety — ranging from household innovations such as air conditioning to modern digital goods such as smartphones and social media — are only imperfectly captured by conventional price indices and therefore by measured real GDP (Bresnahan and Gordon, 1996; Hausman, 1996; Brynjolfsson et al., 2025). Quality improvements are also difficult to measure and can be missed in official price indices and real growth statistics (Griliches, 1961; Boskin et al., 1996). And broader changes related to political and economic freedoms, crime, environmental quality, and climate are largely outside the scope of GDP accounting, motivating complementary welfare measures (Stiglitz, Sen and Fitoussi, 2009; Jones and Klenow, 2016).

This paper makes three points. First, it provides a simple example involving non-homothetic preferences and Baumol's cost disease to illustrate that real GDP can be a profoundly misleading measure of living standards, even beyond the concerns just raised. Imagine a world of two goods: food and string quartets. Food has rapid productivity growth at 5% per year while string quartets has slow productivity growth at 1% per year. Real GDP growth is the weighted average of these two rates, where the weights are the spending shares on the two goods. In the optimal allocation, because of Baumol's cost disease, the spending share on entertainment rises and real GDP growth rates fall toward 1% per year. The incongruity is that if string quartets had never been invented, real GDP growth would be 5% per year forever — creating arbitrarily large discrepancies between the two scenarios — even though utility would be permanently lower.

The main contribution of the paper is to propose a solution that addresses all the concerns laid out so far: under surprisingly weak conditions, the growth rate of the value of a statistical life (VSL) can be used to infer the growth rate of lifetime utility. The intuition is that individuals make decisions that involve taking risks to their lives. Their willingness to take such risks reveals what they expect about their lifetime utility — including all the changing economic and social circumstances that occur over time.

The math involved is straightforward. The VSL equals the remaining lifetime utility divided by the marginal utility of consumption spending (to convert utils into dollars). If we could somehow measure the marginal utility of consumption spending, we could recover a measure of lifetime utility. While we don't know how to do this in levels, we can do it in growth rates. In particular, the standard Euler equation for consumption sets the growth rate of marginal utility equal to the difference between the rate of time preference (ρ) and the interest rate (i_t). This leads to our key result. Letting g denote growth rates, we have

$$g_t^{utility} = g_t^{VSL} + \rho - i_t$$

Intuitively, the growth rate of lifetime utility is the growth rate of the VSL adjusted for the rate at which the marginal utility of consumption spending declines.

Finally, we implement this approach empirically. Standard measures of living standards based on the log of real per capita consumption increase by a factor of around 2 since 1940 — and preferences with more curvature would show an even smaller increase. In contrast, the VSL-implied utility measures increase by a factor of 5 to 7 in our baseline scenarios. The exact numbers are sensitive to measures of all three terms on the right-hand side of the formula, so this calculation is just suggestive. But the paper's insight is that a relatively straightforward set of measurements can be surprisingly informative about how living standards have changed over time in a way that is not subject to the numerous concerns that plague real GDP.

Moreover, the empirical results and the Rothschild and Gates examples at the start of the paper highlight a different view of growth. Passing exponential increases in real GDP — or even explosive increases — through a stable, bounded utility function (such as $u(c) = \frac{c^{1-\theta}}{1-\theta}$ with $\theta > 1$) is standard practice in the growth literature. But this practice implies that even explosive growth cannot increase flow utility in the long run. In contrast, the VSL-implied utility measures suggest that living standards have changed enormously. A natural explanation is that new goods and higher quality versions of existing goods shift the upper bound on utility, leading to much larger increases in utility than the stable and bounded utility function approach suggests.

Related literature. A large tradition in index-number theory emphasizes that real GDP and standard price/quantity indices admit clean welfare interpretations only under restrictive conditions, most famously homotheticity (Samuelson and Swamy, 1974). Closely related work on Divisia and chain indices clarifies why the connection to welfare breaks once expenditure shares move endogenously with income: chain indices can become path dependent and need not track a representative consumer's cost-of-living or welfare (Hulten, 1973; Diewert, 1978; Reinsdorf, 1998; Oulton, 2008). Our first contribution builds on this insight by highlighting a sharp implication for long-run growth comparisons: when the set of goods changes and demand shifts systematically across sectors, measured real GDP can move in the opposite direction from utility, and the gap can become arbitrarily large.[1]

Other papers build on the index literature and seek to provide improved measures of welfare in environments with income effects and preference instability. Baqaee and Burstein (2023) consider an equivalent variation measure of welfare, evaluated at final preferences, and show how this can be calculated if one has knowledge of the input-output production structure as well as elasticities of substitution in both production and consumption. Jaravel and Lashkari (2024) have a similar money-metric welfare criterion and show how cross-sectional data on prices and quantities for consumers with heterogeneous incomes can be used to measure the income effects and estimate growth in consumer welfare. While making valuable progress, these approaches remain vulnerable to the new goods problem: no amount of income in his last weeks of life could make up for Nathan Rothschild's lack of access to antibiotics (i.e. for their infinite price in the base period).

A broader literature seeks to discipline welfare and cost-of-living measurement with revealed preference and other approaches such as Engel-curve restrictions (Feenstra and Reinsdorf, 2000; Blundell, Browning and Crawford, 2003; Redding and Weinstein,

[1] Oberfield (2023) highlights how nonhomothetic preferences interacting with directed technical change can lead measured real income growth to provide a misleading guide to welfare across groups.

2020). Our perspective is complementary: rather than refining price and quantity indices directly, we seek a welfare-growth measure that is robust to changing goods, qualities, and tastes by leveraging cardinal utility and mortality-risk tradeoffs.

2. A Fundamental Problem with Real GDP

To begin, we make a point that is well-known among economists who study index number theory. However, the point is often overlooked, and we think a simple example highlights some profound implications for measuring living standards.

Suppose flow utility is the sum of the utility from each of N_t goods:

$$U_t = \sum_{i=1}^{N_t} u_i(c_{it}) \tag{1}$$

where $u_i(c) \geq 0$, $u_i(0) = 0$, $u_i(\infty) = \bar{u}_i$, and $u_i'(c) > 0$.

In this setup, none of the goods are necessary in the sense that if a good has not yet been invented (or if it simply is not consumed), U_t can still be positive. Also, consuming infinite amounts of a particular good only delivers bounded utility. In the long run, progress in living standards must come from inventing new goods or higher qualities of existing goods, which here correspond to new goods.

Our view is that this setup captures some key features of modern economies. Nontrivial increases in living standards do not come simply from us consuming more quantities of the corn and wheat that existed back in 1870. There are sharp diminishing returns to consuming more and more corn on a given day. Instead, progress comes from inventing antibiotics and smartphones.

The surprise is that even this simple setup violates the conditions that are necessary for real GDP to be a consistent measure of living standards. Technically, these preferences are nonhomothetic, and homotheticity is a necessary condition for real GDP to be a consistent measure of living standards (Samuelson and Swamy, 1974).[2]

[2] Preferences are homothetic if for any two consumption bundles x and y,

$$x \succsim y \implies \alpha x \succsim \alpha y \quad \forall \alpha > 0.$$

Two features of the preferences in (1) combine to break homotheticity. The first is that preferences saturate and the second is that no good is necessary. To see why, consider the case where $\bar{u}_i$ is the same for all goods and equal to $\bar{u}$. As the consumption of N goods goes to infinity, utility approaches $N\bar{u}$ — it is determined

To make this point concrete, we state the following proposition.

Proposition 1. In the setup just given, real GDP can be profoundly misleading: real GDP can be arbitrarily higher if new goods are *not* invented even though utility — the standard of living — is higher if they are.

The proposition is proved in the next subsection by providing a simple example.

2.1 A Baumol Cost Disease Example

Consider the following example — a special case of the setup above. The economy begins with a single good (f, for "food"). Eventually at some date t^*, a new good (s, for "services" or "string quartets") is invented; its discovery is exogenous in this example. Utility is the sum of the utility from each good: $U_t = u_f(c_t^f) + u_s(c_t^s)$, and the utility from each good is given by a bounded utility function. You only get so much utility from eating a bunch of apples or repeatedly watching the same violin performance on a given day:

$$u_i(c) = \bar{u}_i + \frac{(c + \bar{c}_i)^{1-\theta}}{1-\theta} \quad \text{where} \quad \theta > 1. \tag{2}$$

The Stone-Geary term $\bar{c}_i$ is chosen so that $u_i(0) = 0$.[3] An example of this function is shown in Figure 1.

Each good is produced using labor and there is a fixed supply of labor, normalized to 1:

$$c_t^f = A_t^f \ell_t \quad \text{and} \quad c_t^s = A_t^s (1 - \ell_t)$$

Markets are perfectly competitive and the equilibrium allocation is optimal.

Both productivity terms A_t^i are assumed to grow at constant exponential rates g_A^i. Productivity growth in food (agriculture) is fast, while productivity growth in services

by how many goods are consumed, not by the consumption of any one good. This leads to reversals. Consider two bundles, where x consumes a tiny amount of each of two goods, while y involves 1 unit of only one of the goods. Clearly we can set this up so that $y \succ x$ if the consumption in x is sufficiently small. But now consider blowing up consumption in each bundle by the factor α and let $\alpha \to \infty$. Then $x \succ y$ because the bundle x features two goods while the bundle y features only one.

[3] In particular, $\bar{c}_i = [\bar{u}_i(\theta - 1)]^{\frac{1}{1-\theta}}$.

Figure 1: Utility from consuming a given good

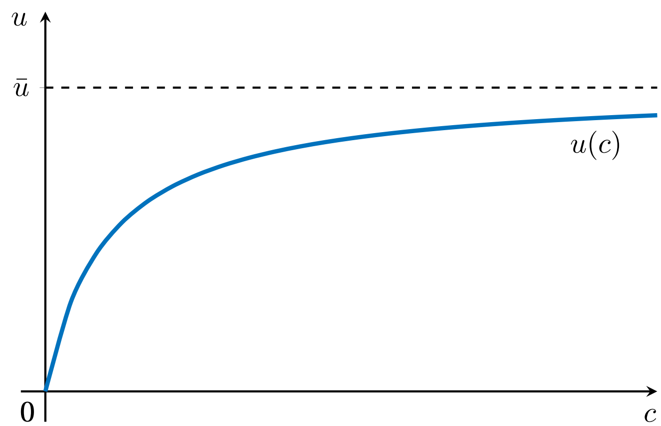

Note: Utility is zero if consumption is zero and approaches $\bar{u}$ as consumption goes to infinity.

(string quartets) is slow: $g_A^f > g_A^s$. For example, suppose productivity growth in making food is 5% per year, while productivity growth in performing string quartets is 1% per year.

Finally, we make a technical assumption about $\bar{c}_s$ so that at date t^* when string quartets are invented, the economy still keeps 100% of its labor in food. This assumption is not necessary for the results that follow, but it ensures that utility does not jump up at date t^*. There is a classic observation that "new goods" are not properly counted in GDP when their consumption jumps discontinuously from zero to a positive quantity, and we are simply turning off this force to illustrate other features of the model.

The optimization problem. The allocation of labor is chosen optimally to maximize utility at each date:

$$\max_{\ell_t} u_f(A_t^f \ell_t) + I_t u_s(A_t^s(1-\ell_t)) \tag{3}$$

Here, I_t is an indicator variable that is 1 after date t^* when string quartets are invented and 0 before t^*.

Solving the model. Before string quartets are invented, the allocation is trivial. The economy puts all of its labor into making food and the growth rate of GDP and consumption is g_A^f. Because utility is bounded and growth is fast, the marginal utility of consumption of food falls quickly.

After date t^*, string quartets have been invented, and the economy gradually shifts its labor toward that good. This occurs because the fast growth of productivity in food drives down the marginal utility in that sector, while the slow growth in string quartets means marginal utility falls more slowly. The model is straightforward to solve analytically, but we omit the details here; see Appendix A.

This basic example exhibits the classic "cost disease" of Baumol (1967) in which the relative price of string quartets rises over time and the economy tilts its labor allocation and spending share away from food and towards string quartets once the new good is invented. As t goes to infinity, the labor allocation shifts entirely to string quartets, approaching 100%, while the high productivity growth rate in food ensures that c_t^f still goes to infinity even faster than c_t^s.

GDP growth and utility. National income accounts in this economy would calculate GDP growth using the chaining method; for a nice illustration, see Whelan (2003). In particular, real GDP growth is a weighted average of the growth rates of the two goods where the weights are the spending shares. Intuitively, before string quartets are invented, the growth rate is just g_A^f. In the distant future when the spending share on string quartets approaches 100%, the growth rate is g_A^s. And for the periods in between, the growth rate is a weighted average of the growth rates of the two goods. This is illustrated in Figure 2a.

Overall utility in this example is also straightforward to calculate. Before date t^*, the economy just follows the $u_f(c_t^f)$ utility function. After date t^*, we get to add utility from string quartets. As time goes on, economic growth means that utility U_t asymptotes to $\bar{u}_f + \bar{u}_s$ as consumption of each good goes to infinity. This is illustrated in Figure 2b.

Key insight. So far, all of this is straightforward and nothing is surprising. The incongruity comes from pondering the following question: Is the invention of string quartets a valuable thing or not?

The fundamental answer is clear from Figure 2b: utility is higher after the service good is invented because this allows the economy to escape the sharp diminishing returns and bounded utility associated with consuming only more and more food.

Figure 2: Baumol Cost Disease Example: GDP growth, utility, and real GDP

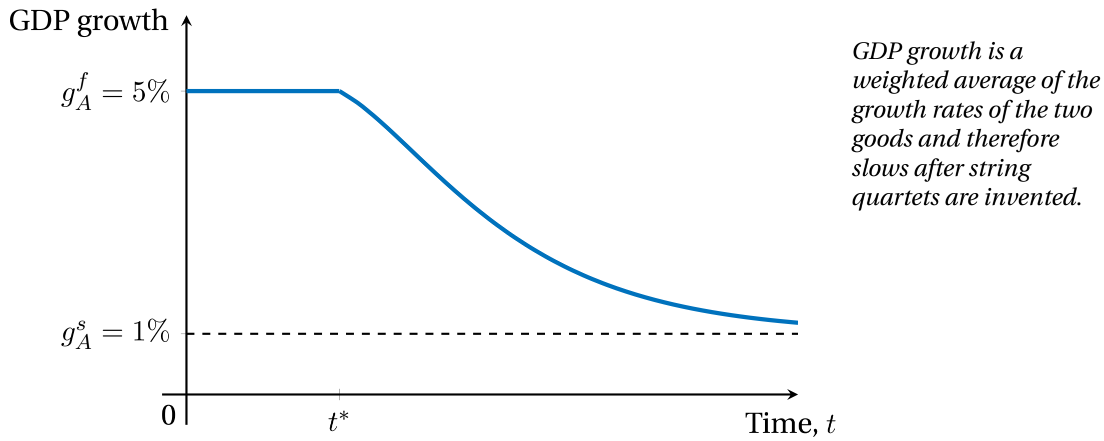

(a) GDP growth

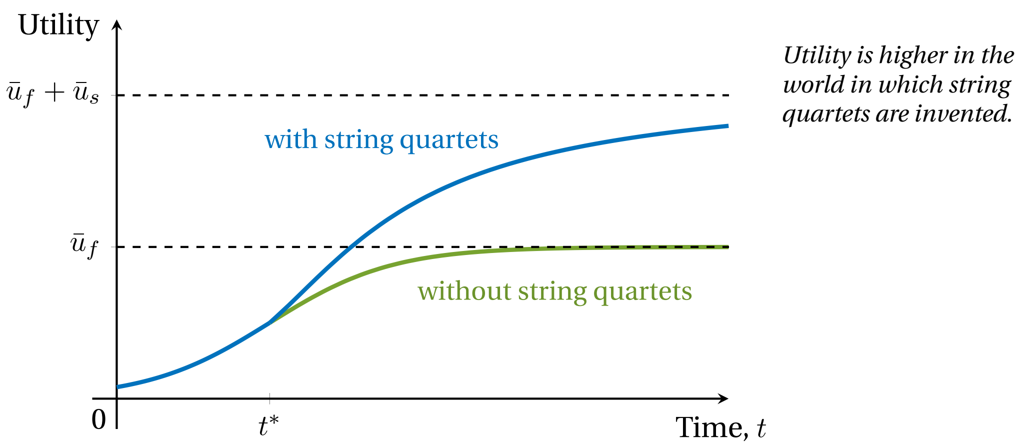

(b) Utility

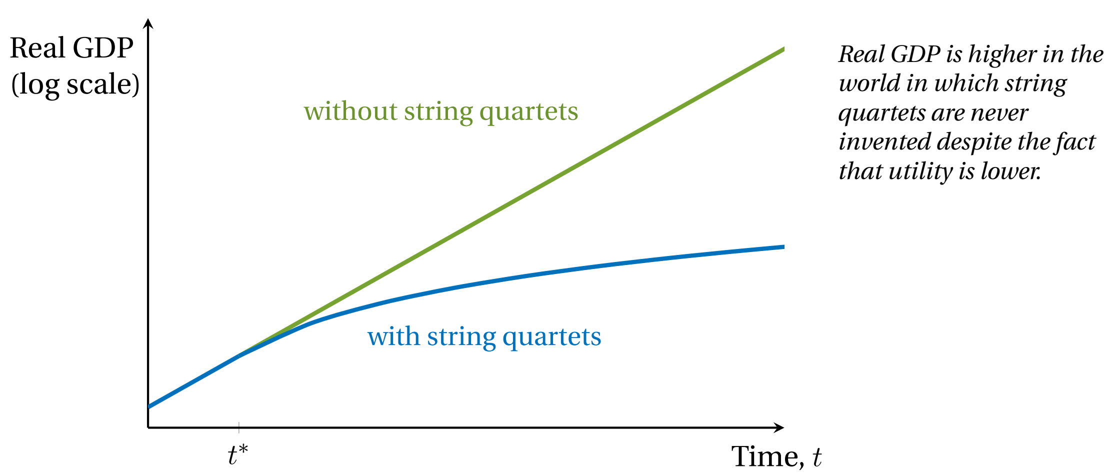

(c) Real GDP

However, consider the answer that comes from national income accounting. GDP growth is given by Figure 2a. Growth slows after string quartets are invented, eventually falling from, say, 5% to just 1%. In particular, *had string quartets never been invented, GDP growth would have been 5% per year forever*. Since real GDP is just the cumulation of GDP growth, this means that the level of real GDP is much lower in the future relative to the path in which string quartets had never been invented, as shown in Figure 2c.

As time goes on, the gap between the two paths of real GDP gets arbitrarily large. The message from real GDP is that we should never invent string quartets and continue to enjoy the rapid 5% growth that comes from consuming only food. However, this perspective is entirely misleading. Utility is clearly higher in the world in which string quartets are invented.

2.2 Summary

Index number theory long ago pointed out that with nonhomothetic preferences, real GDP can be a misleading guide to welfare. The simple example in this section highlights how misleading it can be. Modern economies have experienced a structural transformation away from fast-growing agriculture and toward slow-growing services. A natural implication of this structural transformation is that *measured GDP would likely be substantially higher if we had never made the structural transformation.* But this says much more about the measure of GDP than it says about living standards.

While this basic point has been known since at least the 1970s, it is not deeply internalized in the growth literature (though there are important exceptions that we discussed in the literature review). Most likely, this is because the topic of economic growth itself is extremely important, and there is seemingly no better alternative: a flawed measure may be better than none at all. In the next section, however, we propose a better measure.

3. Measuring Utility using the Value of a Statistical Life

The point of this section is that the willingness of individuals to trade off consumption for mortality risk is inherently informative about their expectations of lifetime utility. This tradeoff naturally incorporates the changing set of goods that people consume over time, changes in the quality of goods, and even changes in social norms — anything that affects utility. We now lay out the argument in detail.

Consider the following decision problem for a person who faces a time path of mortality rates δ_t and has to choose the consumption of a range of goods at each point in time to maximize lifetime utility. Given a path of prices $\{p_{it}\}_{i=1}^{\infty}$, the wage w_t, the interest rate $\tilde{i}_t$, and the mortality rates δ_t, the agent solves

$$\begin{aligned}
\max_{\{c_{it}\}} \quad & W_0 = \int_0^{\infty} e^{-\rho t} S_t U(c_{1t}, c_{2t}, \ldots, c_{nt}, \ldots; t)\, dt \quad \text{subject to} \\
& c_t^e = \sum_{i=1}^{\infty} p_{it} c_{it} \\
& \dot{a}_t = \tilde{i}_t a_t + w_t - c_t^e, \quad a_0 \text{ given} \\
& \dot{S}_t = -\delta_t S_t, \quad S_0 = 1
\end{aligned}$$

S_t is the probability of surviving to time t, a_t is the agent's wealth, and c_t^e denotes total consumption spending on the potentially changing set of goods. $U(\cdot)$ is continuous, concave, and differentiable.

Notice that we allow the static part of the utility function $U(\cdot)$ to be nonseparable and to be defined over an arbitrary and potentially infinite number of goods. Moreover, this static utility function can depend on time t: the set of goods can change, their qualities can change, and even preferences can change. With respect to $U(\cdot)$, we assume that none of the goods are essential in the sense that if a good has not yet been invented (or if it is not consumed), U_t can still be positive as long as some other goods are consumed in sufficient quantity. We incorporate new goods by allowing the prices of goods that have not yet been discovered to be infinite so that those goods have $c_{it} = 0$. This nests the additively separable case considered in Section 2, but is obviously much more general.

The fact that the date t flow utility can depend on time t in an arbitrary way allows for a great deal of richness in what determines utility. We already noted a changing set of goods and changes in the quality of goods. But the approach allows for even more — things like pollution externalities, the utility from living in a free society, and potential costs associated with social media. Anything that people take into account in trading off consumption today versus the risk of losing their lifetime utility is incorporated.

Static problem. We solve the problem in two parts. First, we consider the static problem of allocating a given amount of total spending c_t^e across the goods that exist at date t:

$$U_t(c_t^e;\{p_{it}\}) \;=\; \max_{\{c_{it}\}} \; U(c_{1t}, c_{2t}, \ldots, c_{nt}, \ldots; t) \;\; \text{subject to}$$
$$c_t^e = \sum_{i=1}^{\infty} p_{it} c_{it}$$

Overloading our notation even more, let $U_t(c_t^e) \equiv U_t(c_t^e;\{p_{it}\})$ be the indirect utility function from the static problem, i.e. the optimal utility from spending c_t^e dollars at prices $\{p_{it}\}$.

Dynamic problem. We now solve the dynamic problem. The dynamic problem is to choose the time path of total consumption spending c_t^e to maximize lifetime utility given a deterministic path of the wage and interest rate:

$$\max_{\{c_t^e\}} \quad W_0 = \int_0^{\infty} e^{-\rho t} S_t U_t(c_t^e)\, dt \;\; \text{subject to} \tag{4}$$
$$\dot{a}_t = \tilde{i}_t a_t + w_t - c_t^e, \quad a_0 \text{ given}$$
$$\dot{S}_t = -\delta_t S_t, \quad S_0 = 1$$

The first order conditions for this problem lead to a standard Euler equation:

$$g_{muc,t} = \rho + \delta_t - \tilde{i}_t \tag{5}$$

where *"muc"* denotes the marginal utility of consumption spending $U_t'(c_t^e)$ and $g_{muc,t} \equiv \frac{d \log U_t'(c_t^e)}{dt}$ is the growth rate of the marginal utility of consumption spending. The first order condition is the familiar Euler equation that says that the marginal utility of consumption spending evolves exponentially at rate $\rho + \delta_t - \tilde{i}_t$.

A question arises as to what happens to the assets of a person when they die. As Blanchard (1985) notes, in a general equilibrium model with an actuarially fair annuity market, the mortality rate drops out of the Euler equation. Individuals effectively get an extra return on their assets equal to the mortality rate: $\tilde{i}_t = i_t + \delta_t$, where i_t is the risk-free interest rate not including the annuity (e.g. the Tbill rate). Making this substitution in (5), we are left with the traditional Euler equation:

$$g_{muc,t} = \rho - i_t \tag{6}$$

There is of course a question of whether or not people actually have access to actuarially fair annuity markets. We will use the traditional Euler equation in (6) as our baseline case — it turns out that this is the conservative approach. But we also consider the case in which the effective discount rate ρ is augmented by the mortality rate in a robustness check.

The Value of a Statistical Life. The value of a statistical life (VSL) is the willingness to pay, per unit of risk, for a small reduction in mortality risk. What follows is a standard analysis in the tradition of Arthur (1981) and Rosen (1988).

Just as we reused the U_t notation to denote the indirect utility function, let's do the same thing with W_0. That is, let $W_0(a_0, \{p_{it}, w_t, i_t, \delta_t\})$ denote the optimal lifetime utility that solves the dynamic problem.

To define a well-behaved "marginal change in mortality risk today," it is convenient to augment (4) slightly and work with the cumulative hazard $\Delta_t \equiv \Delta_0 + \int_0^t \delta_s \, ds$ where $\Delta_0 = 0$ in the baseline problem and $S_t = \exp(-\Delta_t)$. Taking the derivative of W_0 in (4) with respect to Δ_0 (using the envelope theorem) gives

$$\frac{\partial W_0}{\partial \Delta_0} = -W_0.$$

Now consider a change in initial assets a_0 and mortality Δ_0 that keeps the individual at the original lifetime utility W_0 (i.e. so that $dW_0 = 0$):

$$\begin{aligned} \frac{\partial W_0}{\partial a_0} \cdot da_0 + \frac{\partial W_0}{\partial \Delta_0} \cdot d\Delta_0 = 0 \\ \Rightarrow \quad \left.\frac{da_0}{d\Delta_0}\right|_{W_0} = -\frac{\partial W_0 / \partial \Delta_0}{\partial W_0 / \partial a_0}. \end{aligned}$$

From the dynamic optimization problem, $\frac{\partial W_0}{\partial a_0} = S_0 U_0'(c_0^e) = U_0'(c_0^e)$ since $S_0 = 1$, and from the argument above, $\frac{\partial W_0}{\partial \Delta_0} = -W_0$.

Finally, we've derived these expressions for time 0, but they clearly are valid for any time t that represents "today." Collecting these expressions gives

$$VSL_t \equiv \left.\frac{da_t}{d\Delta_t}\right|_{W_t} = \frac{W_t}{U_t'(c_t^e)}. \tag{7}$$

The value of a statistical life, VSL_t, is the willingness to pay, per unit of mortality risk, for a small reduction in that risk. For example, if we were considering a reduction in this mortality risk by 0.001, then roughly speaking, we would need to do that 1000 times to save one statistical life. So $0.001 \times 1000 = 1$ and VSL_t is the willingness to pay to save one statistical life.

Rewriting (7) gives a simple expression for lifetime utility:

$$W_t = VSL_t \cdot U_t'(c_t^e) \tag{8}$$

Lifetime utility is the product of the VSL and the marginal utility of consumption spending. Marginal utility is the exchange rate that converts dollars into utils.

Notice that if we had some way to measure the marginal utility of consumption spending, we could use the level of the VSL to measure the level of lifetime utility. Unfortunately, we do not know of a good way to measure the marginal utility of consumption spending. However, we can observe the growth rate of the marginal utility of consumption spending from the Euler equation. This in turn allows us to measure the growth rate of lifetime utility, as we see next.

The Growth Rate of Utility. Taking logs and derivatives of (8) and using the Euler equation (6), we have the following key result:

Proposition 2 (*The Growth Rate of Lifetime Utility*)**.** The growth rate of lifetime utility is the growth rate of the value of a statistical life adjusted for the rate at which the marginal utility of consumption spending declines. The growth rate of marginal utility comes from the standard Euler equation, which depends on the gap between the discount rate and the interest rate:

$$\begin{aligned} g_t^W &= g_t^{VSL} + g_{muc,t} \\ &= g_t^{VSL} + \rho - i_t \end{aligned}$$

The proposition provides a way to measure the growth rate of lifetime utility using a few basic economic objects. Notice that this result is completely robust to the problems of new goods and quality improvements that typically plague the measurement of real GDP. The insight is that individuals trading off consumption versus mortality is informative about their expected lifetime utility. From equation (8), if we could somehow measure the marginal utility of consumption spending at date t, we could use the *level* of the *VSL* to measure the level of lifetime utility. This is difficult — we do not typically observe the marginal utility of spending. However, the Euler equation allows us to observe the growth rate of the marginal utility of consumption spending.

The VSL tells us the value of future lifetime utility in dollars. In fact, if we want a measure of utility in dollars, then the VSL is the only thing we need, which is a tempting prospect. But there is a key problem. The VSL in nominal dollars is not especially useful, and the same problems with real GDP make a "real" dollar conversion problematic. The use of a utility function tells us how everything should be added up and converted into the relevant units. Fortunately, it is only the marginal utility of consumption spending that we need to measure, rather than the entire structure of how utility depends on an expanding range of goods with potentially changing qualities.

Individuals vs. cohorts. One important limitation of the result is that, though it allows the utility function to take essentially any form, it does not resolve the need for a representative agent. In particular, the result requires marginal utility to evolve in the same way for individuals as for cohorts. To illustrate this issue, suppose that, due to imperfect intergenerational altruism, marginal utility in consumption is systematically lower for the old than for the young. This rapid individual-level decline in the muc will then be reflected in a high interest rate, and the growth rate of lifetime utility from one cohort to the next will be estimated to be lower than it truly is.

Since consumption does in fact increase strongly with age, it seems likely that even our high estimates of g_t^W are biased downward in this respect.

4. Measuring Utility: Empirical Application

4.1 Data

To calculate the growth rate of lifetime utility in Proposition 2, we need data on the growth rate of the value of a statistical life, the rate of time preference, and the interest rate.

Value of a Statistical Life. There is a large literature on the value of a statistical life, surveyed by Aldy and Viscusi (2003) and Kniesner and Viscusi (2019). Most of this literature estimates the VSL at a single point in time. A prominent study that uses repeated cross-sections over a long period of time is Costa and Kahn (2004). That paper uses U.S. census data every decade from 1940 to 1980 to estimate the VSL at each point in time. Our first approach to the VSL (called the "Costa-Kahn" approach below) uses these estimates and merges them with the most recent estimate of $13.7 million for 2024 from U.S. Department of Transportation (2025).

As an alternative, one can combine the cross-section estimates at different dates from various papers in order to get a time series for the VSL. A common approach is to use collection of VSL estimates to obtain an income elasticity of the VSL and then use that income elasticity to extrapolate the VSL to different points in time. In fact, this is the basis of the 2024 estimate by the U.S. Department of Transportation (2025). Costa and Kahn (2004) find an income elasticity ranging from 1.5 to 1.7. U.S. Department of Transportation (2021) reviews the evidence on the income elasticity from a number of highly-regarded studies and reports elasticities ranging from 0.5 to 1.6. They adopt a value of 1.0 in updating the VSL to 2024. The survey by Kniesner and Viscusi (2019) reports a similar range. Based on this evidence, our second approach to the VSL assumes it is proportional to (nominal) GDP per person raised to a power of 0.7, 1.0, and 1.3. Our GDP and consumption data are from the National Income and Product Accounts (NIPA) of the Bureau of Economic Analysis.

Interest rate. For the interest rate, the Euler equation approach suggests we need a measure of the short-term risk-free rate. A conventional choice is the 3-month Treasury yield. However, following Krishnamurthy and Vissing-Jorgensen (2012), it is widely appreciated that short-term Treasuries also have special liquidity privileges that allow them to pay an even lower rate. Estimates of this "convenience yield" range from 0.4%

to 0.7% or possibly even 2%, though this last number is for foreign investors (Mota, 2023; Jiang et al., 2019). We therefore measure the risk-free rate as the annualized 3-month Treasury yield plus a 1% convenience yield, and consider robustness to setting the convenience yield to 0% or 2%.

For the rate of time preference, we use a conventional value of 1%.

Nominal vs. "Real" Data. An issue we have skirted around until now is "real" vs. nominal data. An important consequence of the first part of the paper is that "real" GDP — and hence the deflator that is implied — can be highly misleading. Crucially, the approach laid out in Proposition 2 does not require any "real" measures; nominal data is all that is needed.

This can be seen by looking back at Proposition 2. Notice that the two variables involving prices are the growth rate of the VSL and the interest rate. Because they are netted out, any common inflation adjustment drops from the calculation.

In what follows, we begin by doing our analysis with purely *nominal* data. This is intended to emphasize the point that implementing our approach does not require any "real" measures. However, we will also show the results using a CPI deflator to convert the VSL and the interest rate, in part because seeing the "real" growth rate of the VSL is helpful given that we traditionally take this approach so often (Bureau of Labor Statistics, 2024).

4.2 Baseline Results

Table 1 reports the growth rate of lifetime utility according to Proposition 2 for the Costa-Kahn VSL scenario. This table uses the baseline values for the rate of time preference ($\rho = 1\%$) and the convenience yield (1%).

First, notice that the average annual growth rate of lifetime utility is the same whether we use nominal or "real" data, as just discussed. Between 1940 and 2024, lifetime utility grew at an average rate of 2.3% per year. For comparison, "real" per capita GDP over this period grew at 2.2% per year — and passing this through a stable concave utility function would lower this rate further. (We have mixed feelings about making this comparison. After all, the point of the first part of the paper is that "real" GDP is not a useful measure. However, because it is something we are all familiar with, it

Table 1: Utility Growth

Period	g_t^W	g_t^{VSL}	$g_{muc,t}$	Discount Rate	Interest Rate
			Nominal Data (Current Dollars)		
1940–1950	10.1	10.6	-0.5	1	1.5
1950–1960	1.1	3.2	-2.1	1	3.1
1960–1970	7.4	11.5	-4.2	1	5.2
1970–1980	3.0	9.6	-6.5	1	7.5
1980–2024	-0.5	3.4	-3.9	1	4.9
1940–2024	**2.3**	**6.0**	**-3.7**	**1**	**4.7**
			"Real" Data (2024 Prices)		
1940–1950	10.1	5.2	4.9	1	-3.9
1950–1960	1.1	1.1	0.0	1	1.0
1960–1970	7.4	8.8	-1.4	1	2.4
1970–1980	3.0	2.1	1.0	1	0.0
1980–2024	-0.5	0.4	-0.9	1	1.9
1940–2024	**2.3**	**2.2**	**0.1**	**1**	**0.9**

Note: Average annual rates (percent). The growth rate of lifetime utility, g_t^W, is computed as in Proposition 2: $g_t^W = g_t^{VSL} + g_{muc,t} = g_t^{VSL} + \rho - i_t$. The first panel is based on nominal values for the VSL in the Costa-Kahn approach and the interest rate, while the second panel deflates these variables to 2024 prices using the CPI from Bureau of Labor Statistics (2024). A 1% convenience yield and 1% rate of time preference are assumed.

is sometimes helpful to present the comparison: to what extent does our approach change what you thought about the growth of living standards?)

Second, consider where this 2.3% rate comes from by looking at the last line in the table. The "real" VSL grew at 2.2% per year while the marginal utility of consumption was nearly constant, rising at 0.1% per year. Clearly something must have shifted flow utility upward in order for the marginal utility to rise while consumption spending is increasing. A declining mortality rate is one candidate. The flow utility function shifting up because of new and higher quality goods is another.

The growth rate of the marginal utility of consumption is itself the difference between the rate of time preference and the interest rate. Over the 1940 to 2024 period, the real 3-month Treasury yield averaged -0.1% per year; adding a convenience yield of 1% only gets it up to 0.9%, which is slightly below the rate of time preference.

Finally, the time series pattern of the growth rates of lifetime utility are themselves interesting. The growth rate in the 1940s and 1960s was more than 7% per year. In contrast, the growth rate between 1980 and 2024 was actually negative, at -0.5% per year. An important caveat is that this period mixes the Costa-Kahn VSL numbers with the number from the Department of Transportation, so measurement error could be important in this period. We return to this point below.

Figure 3 cumulates these growth rates to show an index of the *level* of lifetime utility over time. That is, the graph shows the factor by which lifetime utility has increased since 1940. For the baseline scenario with the 1% convenience yield, lifetime utility increased by a factor of 8.7 between 1940 and 1980 and by a factor of 6.9 between 1940 and 2024. Measuring the risk-free rate as the Tbill rate with no convenience yield leads lifetime utility to increase by a factor of 16.1 over the entire period. Even with a high convenience yield of 2%, lifetime utility rises by a factor of 3.0.

By way of comparison, the red line at the bottom of Figure 3 shows the level of lifetime utility calculated in a traditional way. That is, we report $W_t = LE_t(\bar{u} + \log c_t)$ where LE is life expectancy, c_t is consumption per capita measured in "real" chained 2017 dollars, and $\bar{u}$ is measured so that utility is zero at \$1000 of consumption.[4] One subtlety is that we do not observe the future of consumption. Therefore, we follow Jones and Klenow (2016) in computing an "as if" lifetime utility measure that summarizes the

[4] If utility hits zero at even lower levels of consumption, then W_t rises by even less.

Figure 3: Utility Levels for the Costa-Kahn VSL Scenario

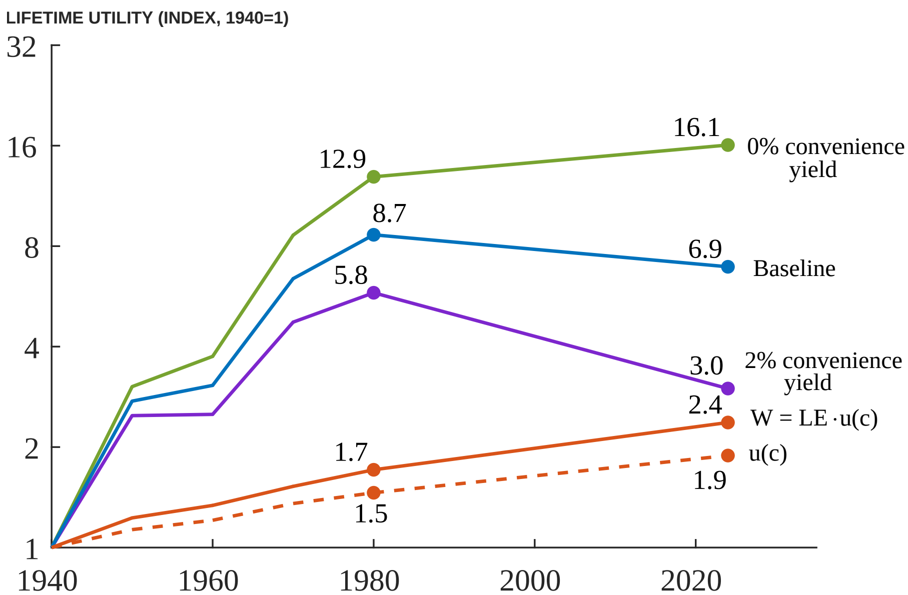

Note: The solid and dashed red lines correspond to traditional measures of utility using "real" per capita consumption. The dashed line shows $u(c) = \log(c/1000)$ where c is measured in chained 2017 dollars and preferences are such that flow utility would be zero at $1000 of consumption. The solid red line multiplies this base life expectancy at birth to get a traditional lifetime utility measure.

data at date t: it is as if the individual lived her entire life in year t (with no growth and no discounting).

The point of the red line is that even incorporating the increases in life expectancy over time, a traditional measure with log consumption grows much more slowly than the measure based on the *VSL*. Between 1940 and 2024, the traditional measure rises only by a factor of 2.4, or at an average annual rate of just 1.04%. In contrast, the Costa-Kahn VSL measure rises by a factor of 6.9, or at an average annual rate of 2.3%. A natural candidate explanation for this difference is the point made at the beginning of the paper: utility is sharply bounded for an existing set of goods, but the expansion in the quality and quantity of goods raises the upper bound and generate much larger increases in utility than the traditional measure with a stable utility function implies.

4.3 Alternative VSL Measures

As discussed in Section 4.1, an alternative approach to the VSL is to assume it is a constant elasticity function of income. We consider three different elasticities: 0.7, 1.0, and 1.3. The results are shown in Figure 4. With an income elasticity of 1.0 — a baseline value that is commonly used in the literature and by the U.S. Department of Transportation (2021) — lifetime utility increases by a factor of 5.1 between 1940 and 2024, or at an average annual rate of 1.9%. However, the rise in lifetime utility is sensitive to the income elasticity. With an elasticity of 0.7, lifetime utility increases by a factor of just 1.3. However, with an elasticity of 1.3 — also a value that has been advocated in the literature (Costa and Kahn, 2004; Hammitt et al., 2000) — the increase in lifetime utility is stunningly large at a factor of 21.1.

4.4 Alternative Interest Rates

Figure 5 shows the sensitivity of the results to different interest rates (in this case, for the GDP-based VSL with an income elasticity of 1.0). We start with the Tbill interest rate and successively add 1%, 2%, 3%, and 4% to get higher interest rates. The higher rates could be from convenience yields but could alternatively capture other reasons why the Tbill rate is not the appropriate measure of the risk-free rate.

The takeaway from the figure is that the measurement of interest rates matters. Depending on how large the spread is over the Tbill rate, utility in 2024 can be as much

Figure 4: Utility when VSL = Income$^{\eta}$

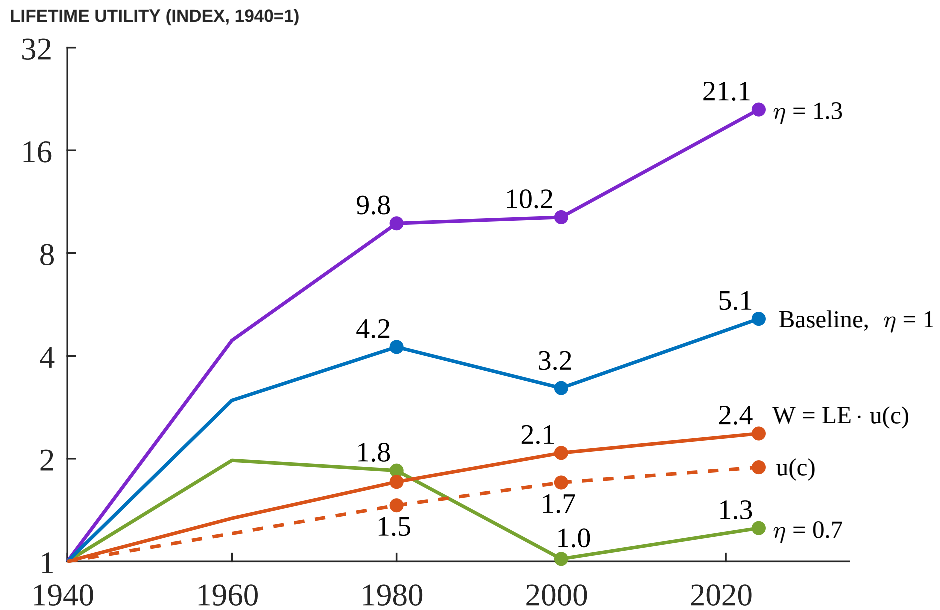

Note: These results are computed assuming the VSL is a constant elasticity function income, measured here as nominal GDP per person raised to the power η. Also see the note to Figure 3 for the "traditional" measures (the red lines).

Figure 5: Utility Levels for Different Interest Rates

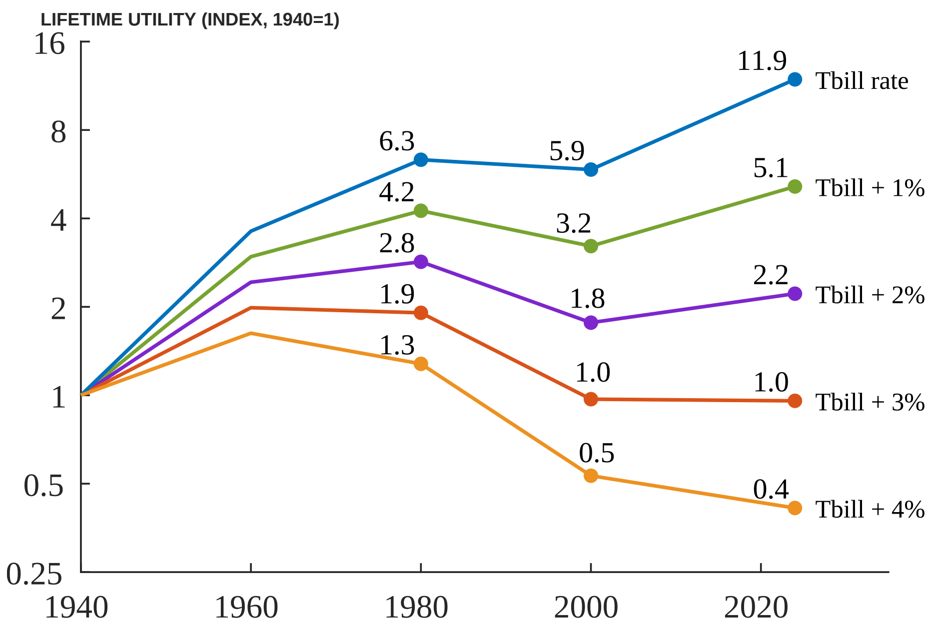

Note: Results for different interest rates. This figure uses the GDP-based VSL with an income elasticity of $\eta = 1.0$. Higher interest rates — especially 3% or more above the Tbill rate — change the overall picture of how living standards have changed over time.

as 11.9 times higher than in 1940 — if we just use the Tbill rate itself — or could have fallen in half since 1940 if the correct interest rate is 4% above the Tbill rate. Perhaps surprisingly, if the interest rate is 2% or more above the Tbill rate, lifetime utility is lower in 2024 than in 1980. Even though our approach leads to a relatively simple formula for the growth rate of lifetime utility, it is clearly of first-order importance to measure the terms in Proposition 2 correctly.

4.5 Robustness: No Annuity Markets

As noted earlier in comparing equations (5) and (6), the presence of perfect annuity markets leads the mortality rate to drop out of the Euler equation, as in Blanchard (1985). Because age-adjusted mortality rates are above 1% per year, this difference matters.[5] With no annuity markets, the presence of the mortality rate in the Euler equation adds around 1pp per year to the growth rate of lifetime utility.

For example, recall that with the Costa-Kahn VSL measures, our baseline scenario showed a rise in lifetime utility between 1940 and 2024 equal to a factor of 6.9. If we were to add in the mortality rate, the rise in lifetime utility would be a factor of 17.3 instead.

Similarly, for the baseline case where we assume an income elasticity of the VSL to GDP equal to 1.0, we found that lifetime utility increased by a factor of 5.1 between 1940 and 2024. If we were to add in the mortality rate, the rise in lifetime utility would be a factor of 12.8 instead.

4.6 Robustness: Risk

Our baseline analysis assumes no risk in the marginal utility of consumption spending. When shocks to marginal utility are mean-independent of the VSL, our lifetime utility estimates remain valid as *expectations* of (expected) lifetime utility across realizations of the *muc*.

In practice, however, the VSL and marginal utility probably negatively covary: both mechanically (since the VSL is the ratio of lifetime utility to marginal utility) and because high marginal utility tends to accompany low consumption and thus low fu-

[5] For mortality, we use the age-adjusted mortality rate from the National Center for Health Statistics (Bastian et al., 2020).

ture welfare. This negative covariance implies that our baseline estimates overstate expected lifetime utility.

Unfortunately, estimating this covariance from consumption data would require specifying a utility function: precisely what we wish to avoid. A more consistent approach may therefore be to use estimates from the finance literature of the unconditional distribution of the *muc*, and use these to generate maximum likelihood estimates of lifetime utility given the observed VSL. Because the *muc* is typically assumed to be right-skewed (in particular, lognormal) and its variance is typically estimated in the finance literature to be extremely high, its mode is far below its mean. Expected marginal utility determines the interest rate, so maximum likelihood estimates of lifetime utility calculated in this way are again far lower than the lifetime utility estimates implied by our risk-free setting.

Both approaches are discussed in Appendix B.

5. Concluding Remarks

Standard measures of real GDP growth can be extremely misleading guides to welfare when preferences are nonhomothetic, when new goods and quality improvements are important, and when nonmarket goods affect welfare. Of course all of these issues are fundamental features of the economy.

This paper shows that under surprisingly weak conditions, the growth rate of the value of a statistical life can be used to infer the growth rate of lifetime utility. This measure naturally incorporates new goods, quality improvements, and even nonmarket goods such as pollution externalities, crime, and the value of living in a free society. The intuition is that individuals make decisions that involve taking risks to their lives, and their willingness to take such risks reveals what they expect about their future lifetime utility.

The formula for the growth rate of lifetime utility is simple and elegant: it equals the growth rate of the value of a statistical life adjusted by the growth rate of the marginal utility of consumption spending. From the usual Euler equation logic, the growth rate of muc_t equals the difference between the rate of time preference and the risk-free interest rate. In principle, many of these objects show up throughout economics and we ought to have good empirical measures of them. In practice, the measures are not as precise as we'd like. The growth rate of the VSL is obviously a very important object. But even knowing what value to use for the rate of time preference and the interest rate is not straightforward: each percentage point you add or subtract changes the growth rate of lifetime utility one for one.

Our baseline calculations indicate that lifetime utility increased enormously between 1940 and 2024. For example, using the Costa and Kahn (2004) VSL estimates through 1980 and the U.S. Department of Transportation (2025) value of life for 2024, lifetime utility increased by a factor of 6.9 between 1940 and 2024, or at an average annual rate of 2.3%. Alternatively, assuming an income elasticity of the VSL equal to 1.0 implies that lifetime utility increased by a factor of 5.1 since 1940, or at an average annual rate of 1.9%. Notice that both of these measures are substantially larger than what one would infer from a utility function over log "real" consumption, which roughly doubles between 1940 and 2024. But if the appropriate interest rate is several percentage points higher or the VSL is substantially mismeasured, these numbers could be far off. Our main contribution is a formula that is much simpler than any attempt to start from national income figures and correct for the massive problems posed by new goods, quality improvements, and changing social norms. But future work will be needed to pin down the measurement of the key terms in the formula.

Why the large difference between our baseline approach and the traditional national accounts measures? One answer is the one we began with — with nonhomothetic preferences, the traditional approach of using real GDP per capita can be extremely misleading. In that sense, there is no reason why the two approaches should be comparable.

A related answer is suggested by the Bill Gates / smartphone example from the introduction. Bill Gates surely has a marginal utility of additional consumption — for a given set of consumption goods — that is zero. But even the richest people in the world benefit from new goods and better goods. The traditional approach of computing utility with a stable utility function misses these evidently quite substantial gains.

Finally, note that the "lifetime utility" our approach tracks might more precisely be called the expected utility assigned to life over that assigned to death. It is possible that the VSL has grown rapidly over time not only because our lives have rapidly grown more enjoyable, relative to marginal consumption, but also because we have grown more averse to death for other reasons: a loss of belief in a heavenly afterlife, for instance, or an increase in the desire to stay alive for altruistic reasons.

Future research could extend our approach in several directions. For example, one could measure heterogeneity in the value of life and how it has changed over time — and therefore how much lifetime utility has changed — across different groups of people, such as by income or education level. Another dimension we spent some time exploring is incorporating risk into the analysis. This inevitably involves measuring the variance of the marginal utility of consumption. This variance is widely studied in finance, but the large and well-known disconnects between macro data and financial markets (such as those behind the equity premium puzzle and risk-free rate puzzle) make it difficult to incorporate risk with any precision.

A. Appendix: Solving the Food/String Quartets Example

The allocation of labor is chosen optimally to maximize utility at each date:

$$\max_{\ell_t} u_f(A_t^f \ell_t) + I_t u_s(A_t^s(1-\ell_t))$$

Here, I_t is an indicator variable that is 1 after date t^* when string quartets have been invented and 0 before t^*.

The first order condition for this problem implies[6]

$$\ell_t^* = \begin{cases} 1 & \text{if } t < t^* \\ \frac{\Omega_t(A_t^s+\bar{c}_s)-\bar{c}_f}{A_t^f+\Omega_t A_t^s} & \text{if } t \geq t^* \quad \text{where} \quad \Omega_t \equiv \left(\frac{A_t^f}{A_t^s}\right)^{1/\theta} \end{cases}$$

As $t \to \infty$, the Stone-Geary terms $\bar{c}_s$ and $\bar{c}_f$ become irrelevant. Therefore

$$\ell_t^* \to \frac{1}{\left(\frac{A_t^f}{A_t^s}\right)^{1-\frac{1}{\theta}} + 1} \to 0$$

since A_t^f grows faster than A_t^s and $\theta > 1$.

[6] Recall that we choose $\bar{c}_s$ so that $\ell_{t^*} = 1$.

In a competitive equilibrium that decentralizes this allocation, prices are equal to marginal costs. If the wage is w_t, then $p_t^f = w_t/A_t^f$ and $p_t^s = w_t/A_t^s$. Spending on food and services are

$$p_t^f c_t^f = \left(\frac{w_t}{A_t^f}\right)(A_t^f \ell_t) = w_t \ell_t$$
$$p_t^s c_t^s = \left(\frac{w_t}{A_t^s}\right)(A_t^s(1-\ell_t)) = w_t(1-\ell_t)$$

Therefore, the spending shares on food and services are simply equal to the employment shares: $s_t^f = \ell_t^*$ and $s_t^s = 1 - \ell_t^*$. The spending share on food starts at 100%, equals 100% at the moment string quartets are invented, and then eventually declines to 0% as t goes to infinity.

Chained GDP. To compute the growth rate of real GDP, we follow the national income accounts and use chaining. That is,

$$g_t^{chain} \equiv s_t^f g_{c,t}^f + s_t^s g_{c,t}^s$$

Recall that $c_t^f = A_t^f \ell_t$ and $c_t^s = A_t^s(1-\ell_t)$. Also, the spending shares are equal to the employment shares. Therefore

$$\begin{aligned} g_t^{chain} &= \ell_t \left(g_A^f + \frac{\dot{\ell}_t}{\ell_t}\right) + (1-\ell_t)\left(g_A^s - \frac{\dot{\ell}_t}{1-\ell_t}\right) \\ &= \ell_t g_A^f + (1-\ell_t) g_A^s \end{aligned}$$

Chain-weighted GDP growth turns out to be equal to the employment-weighted average of the TFP growth rates of the two goods.

B. Risk – Thoughts in Progress

Mean-independent risk. The analysis of Section 3 only applies exactly if investors face no risk in their marginal utility of consumption spending, so that, across two periods 0 and 1, with a risk-free interest rate of i_1,

$$U_1'(c_1^e) = e^{(\rho - i_1)} U_0'(c_0^e).$$

In fact, people face shocks to their marginal utility of consumption spending. To understand where this introduces a complication, note first that these shocks do not interfere with our primary analysis *if the realization of U_1' is mean-independent of V_1*. The equation $W_1 = V_1 U_1'$ is an identity, so

$$\mathbb{E}[W_1|V_1] = \mathbb{E}[V_1 U_1'|V_1] = V_1 \mathbb{E}[U_1'|V_1]$$

(where all expectations condition on the observed interest rate i_1), and in the mean-independent case

$$\mathbb{E}[U_1'|V_1] = \mathbb{E}[U_1'] = e^{\rho - i_1} U_0'.$$

The lifetime utility estimates reported in Figure 3,

$$\hat{W}_t = \frac{V_t}{V_0} e^{\rho t - \sum_{s=1}^{t} i_s} W_0$$

(where 1940 is time 0 and W_0 is normalized to 1), are then precisely the expectations of lifetime utility after conditioning on the observed interest rate and VSL series. (The growth rates of W_t reported in Table 1 are the growth rates of these expectations of W_t, rather than expected growth rates of W_t.)

Mean-dependent risk. More generally, defining the expectation $\mathbb{E}[U_1'|V_1]$, and as a result $\mathbb{E}[W_1|V_1]$, requires introducing a joint prior over U_1' and V_1 given the interest rate series. We are not aware of a principled way to choose such a prior, but on any plausible choice, V_1 negatively covaries with U_1', so $\mathbb{E}[U_1'|V_1]$ decreases in V_1 and mean-independence is rejected. This is for two reasons. First, because $V_1 = W_1/U_1'$, higher realizations of U_1' lower V_1 mechanically, fixing W_1. Second, high marginal utility of consumption spending presumably tends to accompany low "real" consumption, and thus low subsequent lifetime utility W_1 (though as reflected by our general utility function, this connection is not necessary).

Given this negative covariance, $V_1\mathbb{E}[U_1']$ exceeds W_1 in expectation:

$$\begin{aligned}\mathbb{E}[V_1\mathbb{E}[U_1']] &= \mathbb{E}[V_1]\mathbb{E}[U_1'] \\ &> \mathbb{E}[V_1]\mathbb{E}[U_1'] + \text{Cov}(V_1, U_1') = \mathbb{E}[V_1U_1'] = \mathbb{E}[W_1].\end{aligned}$$

That is, the estimates of lifetime utility W_t reported in Figure 3, as a multiple of W_{1940}, exceed "expectations" of lifetime utility by the unknown $\text{Cov}(V_t, U_t')$.

Maximum likelihoods and confidence intervals. Though we cannot produce an expectation for W_t given V_t without a prior, we can produce a maximum likelihood estimate and a confidence interval, given a standard assumption about the *unconditional* distribution of U_t'.

Suppose that

$$\log U_t' \sim N(\mu_t, \sigma^2 t),$$

where σ^2 must be at least 0.1 to match financial statistics (Hansen and Jagannathan, 1991)[7] and, given σ^2, μ_t is chosen so that $\mathbb{E}[U_t'] = e^{\rho t - \sum_{s=1}^{t} i_s}$:

$$\mu_t = \left(\rho - \frac{\sigma^2}{2}\right)t - \sum_{s=1}^{t} i_s.$$

Since $\log V_t = \log W_t - \log U_t'$,

$$V_t \,\big|\, W_t = w \quad \sim \quad \text{Lognormal}\big(\log w - \mu_t,\ \sigma^2 t\big).$$

This yields a maximum likelihood estimate and a $1-\alpha$ confidence interval for W_t:

$$\begin{aligned}\hat{W}_t^{\text{MLE}} &= V_t e^{\mu_t}, \\ \text{CI}_{1-\alpha}(W_t) &= [V_t e^{\mu_t - z\sigma\sqrt{t}},\ V_t e^{\mu_t + z\sigma\sqrt{t}}], \quad z = z_{1-\alpha/2}.\end{aligned}$$

Since the expectation of U_t' is $e^{\mu_t + t\sigma^2/2}$, $\hat{W}_t^{\text{MLE}}$ is less than our lifetime utility series

[7] Hansen and Jagannathan prove a lower bound: $\sigma^2 \geq \log\sqrt{1+\text{SR}^2}$, where SR is the maximum Sharpe ratio across available investments. Since estimates of SR range from 0.3 to 0.5, this lower bound ranges from 0.086 to 0.22. As shown below, even a value of $\sigma^2 = 0.1$ greatly lowers the maximum likelihood estimates of lifetime utility, and greatly widens the associated confidence intervals.

$V_t\mathbb{E}[U_t']$ by a factor of $e^{-t\sigma^2/2}$. For example, if $\sigma^2 = 0.1$, the MLE series can be produced from our unadjusted lifetime utility series by multiplying the latter by $e^{-1/2} \approx 0.6$ each decade.

This downward adjustment is due entirely to the asymmetry of the lognormal distribution, not to the covariance of V and U'. If we posited a symmetric distribution for U', the highest-density value of U' would be $\mathbb{E}[U']$, so we would have $\hat{W}_t^{\text{MLE}} = V_t\mathbb{E}[U_t']$.

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.

Baqaee, David Rezza and Ariel Burstein, "Welfare and Output with Income Effects and Taste Shocks," *The Quarterly Journal of Economics*, 2023, *138* (2), 769–834.

Bastian, B., Vera B. Tejada, and E. Arias et al., "Mortality Trends in the United States, 1900–2018," 2020. National Center for Health Statistics, https://www.cdc.gov/nchs/data-visualization/mortality-trends/index.htm.

Baumol, William J., "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review*, June 1967, *57*, 415–426.

Blanchard, Olivier J., "Debts, Deficits, and Finite Horizons," *Journal of Political Economy*, 1985, *93* (2), 223–247.

Blundell, Richard W., Martin Browning, and Ian A. Crawford, "Nonparametric Engel Curves and Revealed Preference," *Econometrica*, 2003, *71* (1), 205–240.

Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale W. Jorgenson, "Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index," Technical Report, Advisory Commission to Study the Consumer Price Index December 1996.

Bresnahan, Timothy F. and Robert J. Gordon, *The Economics of New Goods* Studies in Income and Wealth, Chicago, IL: University of Chicago Press, 1996.

Brynjolfsson, Erik, Avinash Collis, W. Erwin Diewert, Felix Eggers, and Kevin J. Fox, "GDP-B: Accounting for the Value of New and Free Goods," *American Economic Journal: Macroeconomics*, 2025, *17* (4), 312–344.

Bureau of Labor Statistics, "Historical Consumer Price Index for All Urban Consumers (CPI-U)," 2024. Accessed January 28, 2026.

Costa, Dora L. and Matthew E. Kahn, "Changes in the Value of Life, 1940–1980," *Journal of Risk and Uncertainty*, September 2004, *29* (2), 159–180.

Diewert, W. Erwin, "Superlative Index Numbers and Consistency in Aggregation," *Econometrica*, 1978, *46* (4), 883–900.

Economist, The, "The Health of Nations," *The Economist*, July 17 2004.

Feenstra, Robert C. and Marshall B. Reinsdorf, "An Exact Price Index for the Almost Ideal Demand System," *Economics Letters*, 2000, *66* (2), 159–162.

Ferguson, Niall, *The Ascent of Money: A Financial History of the World*, Penguin Press, 2008.

Griliches, Zvi, "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change," in Price Statistics Review Committee, ed., *The Price Statistics of the Federal Government*, Cambridge, MA: National Bureau of Economic Research, 1961, pp. 173–196.

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.

Hausman, Jerry A., "Valuation of New Goods under Perfect and Imperfect Competition," in Timothy F. Bresnahan and Robert J. Gordon, eds., *The Economics of New Goods*, Chicago, IL: University of Chicago Press, 1996, pp. 207–248.

Hulten, Charles R., "Divisia Index Numbers," *Econometrica*, 1973, *41* (6), 1017–1025.

Jaravel, Xavier and Danial Lashkari, "Measuring Growth in Consumer Welfare with Income-Dependent Preferences: Nonparametric Methods and Estimates for the United States," *The Quarterly Journal of Economics*, 2024, *139* (1), 477–532.

Jiang, Zhengyang, Arvind Krishnamurthy, and Hanno Lustig, "Foreign Safe Asset Demand and the Dollar Exchange Rate," *The Journal of Finance*, 2019, *76* (3), 1049–1089.

Jones, Charles I. and Peter J. Klenow, "Beyond GDP? Welfare across Countries and Time," *American Economic Review*, September 2016, *106* (9), 2426–57.

Kniesner, Thomas J. and W. Kip Viscusi, "The Value of a Statistical Lifeunlocked," *Oxford Research Encyclopedia of Economics and Finance*, 2019.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen, "The Aggregate Demand for Treasury Debt," *Journal of Political Economy*, 2012, *120* (2), 233–267.

Mota, Lira, "The Corporate Supply of (Quasi) Safe Assets," 2023. unpublished manuscript, SSRN 3732444.

Oberfield, Ezra, "Inequality and Measured Growth," Working Paper 31096, National Bureau of Economic Research March 2023.

Oulton, Nicholas, "Chain Indices of the Cost-of-Living and the Path-Dependence Problem: An Empirical Solution," *Journal of Econometrics*, May 2008, *144* (1), 306–324.

Redding, Stephen J. and David E. Weinstein, "Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences," *The Quarterly Journal of Economics*, 2020, *135*, 503–560.

Reinsdorf, Marshall B., "Divisia Indexes and the Representative Consumer Problem," Technical Report, Ottawa Group (Price Index Working Group) 1998. Conference paper / working paper.

Rosen, Sherwin, "The Value of Changes in Life Expectancy," *Journal of Risk and Uncertainty*, 1988, *1*, 285–304.

Samuelson, P. A. and S. Swamy, "Invariant Economic Index Numbers and Canonical Duality: Survey and Synthesis," *The American Economic Review*, 1974, *64* (4), 566–593.

Stiglitz, Joseph E., Amartya Sen, and Jean-Paul Fitoussi, "Report by the Commission on the Measurement of Economic Performance and Social Progress," Technical Report, Commission on the Measurement of Economic Performance and Social Progress 2009.

U.S. Department of Transportation, "Departmental Guidance: Treatment of the Value of Preventing Fatalities and Injuries in Preparing Economic Analyses," Technical Report March 2021. https://www.transportation.gov/sites/dot.gov/files/2021-03/DOT%20VSL%20Guidance%20-%202021%20Update.pdf.

_ , "Departmental Guidance on Valuation of a Statistical Life in Economic Analysis," Technical Report 2025. Accessed January 10, 2025.

Whelan, Karl, "A Two-Sector Approach to Modeling U.S. NIPA Data," *Journal of Money, Credit, and Banking*, August 2003, *35* (4), 627–656.