Population and Welfare:
The Greatest Good for the Greatest Number

Peter J. Klenow Charles I. Jones Mark Bils Mohamad Adhami *

August 28, 2023 – Version 0.7

Abstract

Economic growth is typically measured in per capita terms. But social welfare should arguably include the number of people as well as their standard of living. We decompose social welfare growth — measured in consumption-equivalent (CE) units — into contributions from rising population and rising per capita consumption. Because of the diminishing marginal utility of consumption, population growth is scaled up by a value-of-life factor that substantially exceeds one and empirically averages around 3 across countries since 1960. Population increases are therefore a major contributor, and CE welfare growth around the world averages more than 6% per year since 1960, as opposed to 2% per year for consumption growth. Countries such as Mexico and South Africa rise sharply in the growth rankings once population growth is incorporated, whereas China, Germany, and Japan plummet. We show the robustness of these results to incorporating parental time use and fertility decisions using data from the U.S., Mexico, the Netherlands, Japan, South Africa, and South Korea. The effects of falling parental utility from having fewer kids are roughly offset by increases in the "quality" of kids associated with rising time investment per child.

*Klenow: Stanford University and NBER; Jones: Stanford University and NBER; Bils: University of Rochester, Richmond Fed, and NBER; Adhami: Stanford University. We reversed the author order to combat alphabetical inequality, not based on contribution. We thank Kyeongmin Park for his excellent research assistance on the project, and participants at many seminars for comments.

1. Introduction

Economic growth is almost invariably measured in *per capita* terms. The reason for this is clear: we seek to quantify the gains in living standards in the economy, and individual consumption is a key argument of people's utility functions. From the standpoint of social welfare, however, the total population of the economy arguably matters as well: a world with a billion happy people is "better" than a world with a million people if they are identical in every other way. Similarly, a disaster that killed half the world's population would be a tragedy, even if it left consumption per person unchanged.

Consider two hypothetical economies with exactly the same time path of total factor productivity (TFP) so that the production opportunities for these economies are the same. One country chooses to keep population constant, and all the TFP growth shows up as higher per capita consumption. The other chooses to keep per capita consumption constant and use the extra resources to raise its population. The traditional focus on per capita consumption says that the first country is more successful, but this seems odd given that the two countries have exactly the same production possibilities.

This hypothetical is not too far from some real world examples. Between 1960 and 2019, consumption per person increased by a factor of 6 in Japan but by a factor of 3 in Mexico. However, Mexico's population tripled while Japan's population only rose by 30 percent. Which country was more successful?

Or consider economic growth over the past 25,000 years. For the bulk of that period — until the last few centuries — technological improvements like better stone tools, the wheel, agriculture, and cities showed up as higher population rather than higher per capita consumption. But shouldn't that progress count somehow?

When it comes to public investments in nonrival knowledge, shouldn't the number of beneficiaries in the future matter for a social cost-benefit calculation? For example, shouldn't we be willing to spend more to adapt to or mitigate global warming the larger the future population relative to today's?

This paper reconsiders the pace of economic growth over time and across countries using a consumption-equivalent metric based on total utilitarian social welfare.[1] Stated more coarsely, we put "humans" into a Human Development Index.

To see how, consider an economy of N_t identical people with consumption per person c_t. Let the annual flow of individual utility be $u(c_t)$. The aggregate flow of utility is then $N_t \cdot u(c_t)$, a natural version of "the greatest good for the greatest number." For a large set of countries since 1960 and for several groups of countries over longer time periods (including "The West" and "The World"), we calculate consumption-equivalent welfare growth for this metric. It is worth emphasizing the this calculation is for the *flow* of social welfare rather than a present-discounted value that takes into account the welfare of future generations. In this sense, it is more like GDP itself.

A simple version of our calculation just uses aggregate consumption growth rather than growth in consumption per person. This amounts to putting equal weight on the number of people and consumption per person. When welfare is given by $N_t \cdot u(c_t)$, however, the diminishing returns in $u(\cdot)$ implies that constant and equal weighting of consumption growth and population growth is not correct. We show that the growth rate of consumption-equivalent (CE) social welfare is instead given by

$$v(c_t) \cdot g_{Nt} + g_{ct}$$

where g_N and g_c denote population growth and per capita consumption growth, respectively. The weight $v(c)$ on population growth is the value of a year of life measured in years of per capita consumption and is almost always greater than one because of the "consumer surplus" associated with life. In fact, we will show that the appropriate weight is typically much larger, approximately 5 in recent years for the United States, and averaging 2.7 for the world as a whole since 1960. In other words, from a social welfare standpoint, population growth is substantially more important than growth in per capita consumption.

[1] Throughout the paper, we use social welfare as a shorthand for total utilitarian social welfare, also known in the literature as the Benthamite criterion.

In our results, CE-welfare growth across 101 countries averages 6.2% per year between 1960 and 2019, versus average annual per capita consumption growth of 2.1%. Population growth accounts for two-thirds of CE-welfare growth on average for this sample. These numbers are also illustrative of the United States: CE-welfare growth averages 6.5% per year between 1960 and 2019, versus 2.2% for per capita consumption growth.

The growth rate of consumption-equivalent social welfare therefore provides a very different perspective on the success of various countries over time. Mexico, South Africa, and Kenya move from the bottom third of growth rates to the top 60%; Mexico, for example rises from the 36th percentile to the 88th, with CE-welfare growth equal to 8.6% per year. On the flip side, traditionally fast-growing countries like Germany, Japan, and China have much slower CE-welfare growth. Germany grows at just 0.2% per year, at the 11th percentile. Similarly, Japan and China fall to the 32nd and 45th percentiles respectively, below the United States and below the median. Overall, the correlation between CE-welfare growth and per capita consumption growth across 101 countries is 0.51.

Large differences also appear when we study growth over long periods of time. For example, for Western Europe and its Western offshoots (which we shorten to "The West" below), per capita consumption rose by a factor of 20 in the two centuries between 1820 and 2018 while CE-welfare rose by a factor of nearly 3000. Similar numbers apply to the world as a whole since the year 1500.

It is also important to appreciate what we are *not* doing in this paper. We use the marginal rate of substitution between population and per capita consumption to do accounting exercises along a social indifference curve. This paper is about preferences, not technology. There are lots of other interesting questions one might like to answer such as "What would optimal fertility look like?" or "Did the demographic transition reduce social welfare?" Answering these questions requires specifying the production possibilities of the economy, including things like the nonrivalry of ideas, limited resources, human capital, and pollution. Our approach is more narrow. We cannot address these richer questions, but we get by without making any assumptions about the production structure of the economy. Clearly it would be interesting for future research to make assumptions on the production side to consider these richer questions.

The remainder of our paper proceeds as follows. After a brief review of the literature, Section 2 lays out our basic theory and derives an expression for consumption-equivalent social welfare growth. Section 3 applies this framework, first to a broad set of countries over the period 1960 to 2019 and then to much longer time periods for the United States, "The West," and for the world as a whole. As our baseline calculation attributes the welfare from migration entirely to the destination country, Section 4 provides an alternative calculation at the opposite extreme that assigns the welfare to the source country. For most countries, this makes little difference. Section 5 gauges the robustness of these results to incorporating parental time use, fertility decisions, and a quantity-quality tradeoff. Using detailed time use data for the U.S., Netherlands, Japan, South Korea, Mexico, and South Africa, we find that the effects of falling parental utility from having fewer kids is roughly offset by increases in the "quality" of kids associated with rising parental time investments per child. That is, the detailed calculations with rich micro data confirm that the simpler calculations in the first half of the paper are robust. Finally, Section 6 concludes.

Related literature (in progress). Parfit (1984) posed the "repugnant conclusion" challenge to total utilitarianism. This challenge holds that trying to maximize total welfare could result in an undesirable outcome in the form of many people being alive but enjoying very low flow utility per capita.[2] More recently, in Zuber et al. (2021), 29 philosophers offer reasons to entertain total utilitarianism in spite of the repugnant conclusion. MacAskill (2022) provides a detailed and nuanced case for a total utilitarian perspective.

A large literature estimates the value of a life year using various methods. See Viscusi and Aldy (2003) for a classic survey. The U.S. Environmental Protection Agency (2020) uses such estimates in setting safety thresholds. And the World Health Organization (WHO, 2001) uses them to determine the cost effectiveness of spend-

[2] There is also the "sadistic conclusion" challenge to per capita utilitarianism. This notes that trying to maximize per capita utility without regard to the total number of people alive may imply choosing a very low population.

ing to avoid lost life years due to mortality.

Young (2005) analyzes the impact of the AIDS epidemic on deaths and the standard of living in South Africa. He discuses how spending on antiretrovirals could prevent deaths and the resulting implications for the standard of living for survivors.

Golosov, Jones and Tertilt (2007) propose two Pareto efficiency criteria for assessing outcomes when population levels are endogenous due to fertility choices and parents face quantity-quality tradeoffs.

Cordoba (2015) explicitly analyzes how rising longevity of children (and parents) offsets fertility reductions in terms of the growth of parental living standards. He looks at the impact of the quality-quantity tradeoff on welfare in 116 countries from 1970 to 2005.

Jones and Klenow (2016) quantify consumption-equivalent welfare across countries and time. They incorporate consumption, leisure, life expectancy, and inequality. But their focus is entirely on *per capita* welfare.

De la Croix and Doepke (2021) propose a "soul-based utilitarian" social welfare function that postulates a fixed number of souls who can be born or not, or even reborn. They show that this nests various other social welfare functions considered in the literature, including total utilitarianism. Chichilnisky et al. (2020) propose a related, survival-probability weighted social welfare function.

2. The Framework: Aggregate Welfare

To make our point as clearly as possible, consider an economy of N_t identical people, each with consumption per person c_t. Each person gets flow utility $u(c_t)$. The total flow of welfare enjoyed by this economy is then

$$W(N_t, c_t) = N_t \cdot u(c_t).$$

Notice that this is also a standard utilitarian social welfare function.

Without loss of generality, the value of death is normalized to zero; we also assume that nonexistence is equivalent to death.[3] For life to be valuable, it must then be that $u(c) > 0$. In addition, we make the standard assumptions that u exhibits diminishing marginal utility: $u'(c) > 0$, and $u''(c) < 0$.

There are multiple ways to derive the growth rate of consumption-equivalent welfare. For example, with discrete time, consider the value of λ such that $W(N_t, c_t) = W(N_{t-1}, \lambda c_{t-1})$. We prefer continuous time because, as the time difference shrinks to zero, the compensating variation will equal the equivalent variation and we do not need to make this distinction. To begin, include λ_t as an adjustment to consumption so that $W_t = N_t \cdot u(\lambda_t c_t)$ and totally differentiate:

$$\begin{aligned} dW_t &= dN_t\, u(\cdot) + N_t u'(\cdot)\left[c_t d\lambda_t + \lambda_t dc_t\right] \\ \Rightarrow \frac{dW_t}{W_t} &= \frac{dN_t}{N_t} + \frac{u'(\lambda_t c_t)\lambda_t c_t}{u(\lambda_t c_t)}\left[\frac{d\lambda_t}{\lambda_t} + \frac{dc_t}{c_t}\right] \end{aligned}$$

To get the consumption-equivalent measure, we solve for the growth rate of λ_t that keeps us at the original level of welfare so that $dW_t = 0$ and we evaluate at the initial level of welfare with $\lambda = 1$:

$$\boxed{\underbrace{g_{\lambda t} \equiv -\frac{d\lambda_t}{\lambda_t}}_{\text{CE-Welfare growth}} = \underbrace{\frac{u(c_t)}{u'(c_t)c_t}}_{\equiv v(c_t)} \cdot \frac{dN_t}{N_t} + \frac{dc_t}{c_t}} \tag{1}$$

This puts population growth in consumption-equivalent units, essentially using the slope of the indifference curve. The weight on population growth is $v(c_t) \equiv \frac{u(c_t)}{u'(c_t)c_t}$. That is, $v(c_t)$ is the value of having one more person live for one period, $u(c_t)$; dividing by the marginal utility of consumption puts this in consumption units, and then taking the ratio to consumption puts it in units of "years of per capita consumption." A percentage point of population growth is worth $v(c_t)$ percentage points of consumption growth.

[3] The philosophy literature notes that alternative assumptions could matter. For example, imagine using our free normalization to set pre-existence to zero. Perhaps everyone goes to heaven after dying, or perhaps hell, or perhaps some probabilistic combination of those possibilities. Absent evidence, assuming the same value before and after life is natural, in our view. As Mark Twain remarked, "I do not fear death. I had been dead for billions and billions of years before I was born and had not suffered the slightest inconvenience from it."

A couple of brief examples are helpful for intuition. Notice that $v(c)$ is the inverse of the elasticity of utility with respect to consumption. If $u(c) = c^\alpha$, then $v(c) = 1/\alpha$. With linear utility ($\alpha = 1$), then $v(c) = 1$ and the value of a year of life equals per capita consumption. If $\alpha = 1/2$, then $v(c) = 2$. More generally, the sharper is diminishing marginal utility — the lower is α — the higher is $v(c)$. In general, one would expect $v(c) > 1$, capturing the "consumer surplus" associated with diminishing marginal utility.

Measuring $v(c_t)$ in the United States in 2006. The key weight $v(c)$ is the value of a year of life, measured in dollars, as a ratio to consumption per person:

$$v(c) \equiv \frac{u(c)}{u'(c) \cdot c} = \frac{\text{VSLY}}{c} \approx \frac{\text{VSL}/e_{40}}{c} \approx \frac{\$7,400,000/40}{\$38,000} = \frac{\$185,000}{\$38,000} \approx 4.87$$

The U.S. Environmental Protection Agency (2020) uses a value of life (VSL) equal to \$7.4m in 2006 prices. Given that a middle-aged American had a remaining life expectancy of around 40 years in 2006, this corresponds to a VSLY of around \$185,000. This is very much in line with the value of a life year used in many other studies. Cutler, Ghosh, Messer, Raghunathan, Rosen and Stewart (2022) use a round number of \$100,000, based on Viscusi and Aldy (2003) and consider robustness to twice and half that value.

Consumption per person in the United States in 2006 was \$38,000, including both private consumption and government consumption, which implies a value of $v(c_{usa,2006}) \approx 4.87$. That is, a year of life in 2006 in the United States is valued at approximately 5 years worth of consumption per person. This is our baseline.

Functional form of flow utility. To determine the value of $v(c_t)$ in other years and in other countries — and more generally to do our accounting — we would ideally have similar VSL estimates from many different countries and time periods. There is not a lot of well-identified evidence of this kind, so we take a different approach. In particular, we specify a functional form for flow utility and use that to calculate $v(c)$ at different levels of consumption. Our benchmark, which we relax in a robustness check, is

$$u(c) = \bar{u} + \log c.$$

With this functional form, the value of a year of life is given by

$$v(c) \equiv \frac{u(c)}{u'(c) \cdot c} = u(c) = \bar{u} + \log c. \tag{2}$$

Both of these equations make clear the importance of the constant term $\bar{u}$. To calibrate its value, we choose consumption units such that $c_{usa,2006} = 1$, which means that $v(c_{usa,2006}) = \bar{u} = 4.87$.[4]

The other interesting thing to note about $v(c)$ from equation (2) is that it is not constant. In particular, $v(c)$ increases with the log of consumption: as living standards increase, life becomes increasingly valuable, even relative to consumption.

Using data from the National Income and Product Accounts back to 1929 and from Barro and Ursua (2008) before that, Figure 1 shows the implied value of life $v(c)$ for the United States over time. As in equation (2), this value rises linearly over time, reflecting the exponential growth in consumption. The value is slightly below 2 in 1820 and rises to nearly 5 by 2019.

Figure 2 shows the values of $v(c)$ for a cross-section of some of the most populous countries in the world in 2019 using the Penn World Table 10.0. Interestingly, the range of values is similar in the world's cross section for 2019 as for the United States back to 1820, ranging from a low of just under 2 for Ethiopia to the high of 5 for the United States. The average value across 101 countries in 2019 is 2.7. Kremer, Luby, Maertens, Tan and Wiecek (2023) cite and use a World Health Organization thresholds of valuing a life year of between one and three times per capita GDP. This implies roughly 1.5 to 4.5 times per capita consumption, which is remarkably close to the range across our countries.

[4] If consumption is sufficiently low, then flow utility could turn negative. This issue is discussed extensively in Rosen (1988), who noted that individuals with low consumption would become risk-loving and take gambles between death and a higher level of consumption in order to convexify utility. In our calculations, flow utility turns out to be positive in every year and country.

Figure 1: $v(c)$ over time in the United States

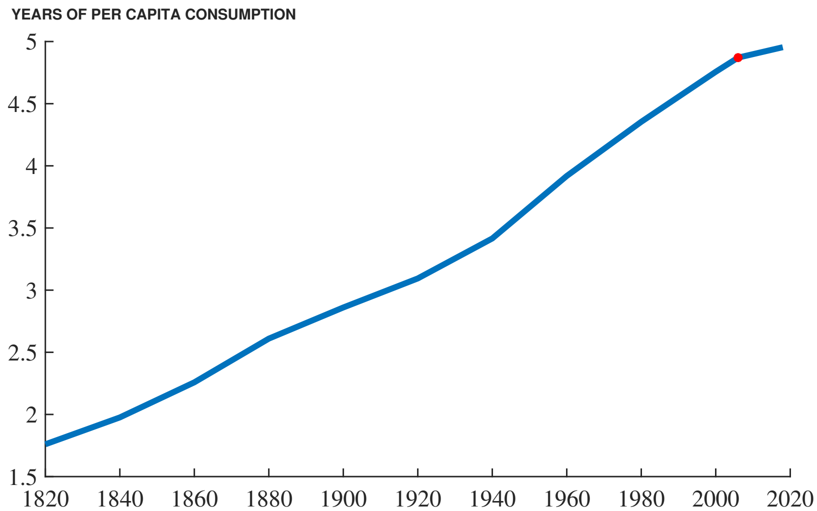

Note: $v(c)$ computed using data from the U.S. National Income and Product Accounts back to 1929 and from Barro and Ursua (2008) before that, assuming log utility.

Figure 2: $v(c)$ across countries in 2019

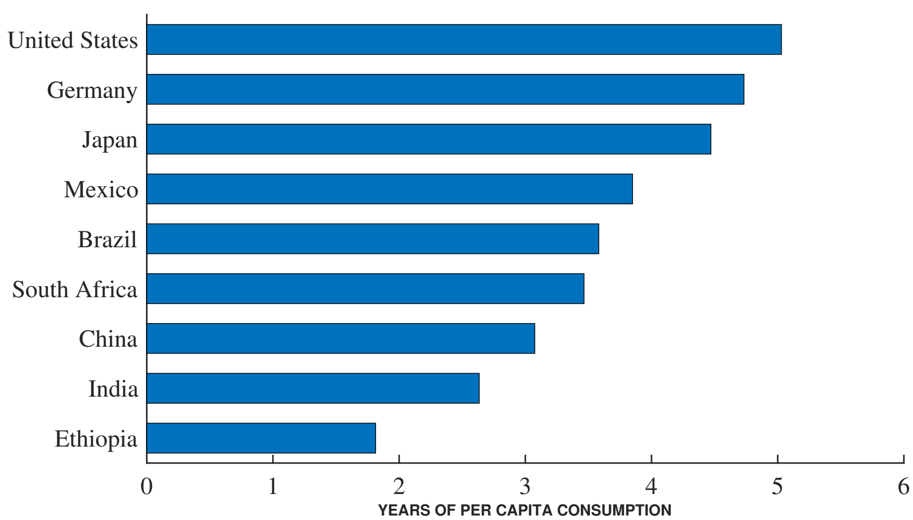

Note: $v(c)$ computed using data from the Penn World Tables 10.0 assuming log utility.

Summary. Motivated by this analysis, we can now summarize our framework. Consumption-equivalent social welfare growth, g_λ, is the sum of per capita consumption growth and population growth, where population growth is scaled by the value of a year of life, $v(c)$:

$$g_\lambda = v(c) \cdot g_N + g_c. \tag{3}$$

Because $v(c)$ is a number around 2 to 5, population growth gets a much higher weight than consumption growth. The remainder of this paper applies this equation empirically, first to a broad cross-section of countries from 1960 to 2019 and then to several groups over countries over much longer time periods.

When implementing this calculation, we always use annual data and then average the result over longer time periods. When annual data are not available, for example in looking at data back to the 1800s or 1500s, we interpolate between the observations assuming a constant consumption growth rate and then proceed as if we have annual data. This strikes us as the best way to treat the data given the changing $v(c_t)$ over time. It is closest to our continuous-time derivation and allows us to avoid the usual "CV" versus "EV" discrepancy.

3. Results: Consumption-Equivalent Social Welfare Growth

Data. We draw on several data sources for our calculations. For the period 1960–2019, we use the Penn World Table 10.0, an updated version of Feenstra, Inklaar and Timmer (2015), which gives us a large sample of 101 countries. Consumption is calculated as the sum of private consumption and government consumption. For historical calculations back to the year 1500, we use The Maddison Project data of de Pleijt and van Zanden (2020).

Table 1: Overview of Results from 1960 to 2019

	Baseline	— Robustness —	
	$\bar{u} = 4.87$	$\bar{u} = 4.44$	$\bar{u} = 6.16$
CE-Welfare Growth	6.2%	5.4%	8.5%
Contribution of population	4.1%	3.3%	6.4%
Average value of life $v(c)$	2.7	2.3	4.0
Pop Share of CE-Welfare Growth	66%	63%	73%
Pop Share (if weight by population)	51%	46%	62%
# of countries with pop share $\geq 50\%$	78	69	89

Notes: Average $g_c = 2.1\%$ and average $g_N = 1.8\%$ across 101 countries. Data on consumption per capita and the population is from Penn World Tables 10.0.

3.1 Macro Results for 1960 to 2019

Table 1 summarizes our results for the 101-country sample from the Penn World Table, applying equation (3) annually and taking the average.[5] While growth in consumption per person averages 2.1% per year between 1960 and 2019, CE-welfare growth is substantially higher at 6.2% per year. Growing at 2.1% per year, average living standards double every 35 years. But taking into account population growth as well, social welfare doubles every 12 years in this sample. The 4.1 percentage point difference between consumption growth and social welfare growth is accounted for by the fact that population growth averages 1.8 percent per year and the value of life $v(c)$ over this period has an average value equal to 2.7 years of consumption (covariances mean that the average of the product is not equal to the product of the average here). Across the 101 countries, population growth accounts for 66% of social welfare growth; weighting countries by their population, which means that China plays a large role, the population share of welfare growth falls to 51%.

The last two columns of Table 1 show how these summary results change when we consider alternative calibrations for the value of life. The median value dis-

[5]The PWT has consumption data for 111 countries since 1960, but we are dropping any country labelled as an outlier in any of the sample years.

cussed in the survey of Viscusi and Aldy (2003) is $5 million, while the median across their preferred studies is $7 million, both in 2000 prices.

Using the lower value of life of $5 million implies $\bar{u} = 4.44$. In this case, social welfare growth falls from 6.2% to 5.4% per year, and the population share declines from 66% to 63%. Conversely, using the higher $7 million value of life, which is around $8 million in 2006 prices, leads to $\bar{u} = 6.16$. In this case, social welfare growth rises to 8.5% per year and the population share rises to 73%.

Table 2 reports the decomposition of growth in consumption-equivalent social welfare for a select sample of countries based on equation (3). To begin, consider the countries with the fastest and slowest growth in the table. Social welfare growth in Mexico averages 8.6% per year since 1960, far exceeding its modest growth in consumption per person of just 1.8% per year. This is for two reasons: population growth averages more than 2% per year and the value of life factor $v(c)$ averages 3.4. Population growth accounts for 79% of social welfare growth in Mexico. At the other extreme is Germany. Its relatively higher growth rate of consumption is barely augmented by its very modest population growth of 0.2% per year even though its value of life factor is 4.0. Population growth accounts for just 22% of social welfare growth in Germany. Figure 3 shows these data graphically, in part to make comparisons with later figures easier.

Figure 4 illustrates the role of $v(c)$ in our calculations. Recall that in our baseline with log utility, the value of life satisfies $v(c) = u(c) = \bar{u} + \log c$; that is, it rises with consumption. While the value of life averages 2.7 in our 101 country sample from 1960–2019, the value for the United States in 2019 is 5 while the value in Ethiopia in 1960 is just 0.4, for example. Instead, Figure 4 treats all life-years as being worth 2.7 years of consumption. While this does not matter for our overall qualitative finding that population growth is more important than consumption growth in our welfare calculations, it does matter for specific countries. Ethiopia, for example, moves from below-average growth to above-average growth, becoming the fastest growing country among our select sample. U.S. growth becomes noticeable slower and below average as life-years are half as valuable.

To show results for a broad set of countries, Figures 5 and 6 show scatterplots of

Table 2: Decomposing Welfare Growth in Select Countries, 1960–2019

	g_λ	g_c	g_N	$v(c)$	$v(c) \cdot g_N$	Pop Share
Mexico	8.6	1.8	2.1	3.4	6.8	79%
Brazil	7.9	3.1	1.8	2.8	4.8	61%
South Africa	7.9	1.4	2.1	3.1	6.4	82%
United States	6.5	2.2	1.0	4.4	4.3	66%
China	5.7	3.8	1.3	1.8	2.0	34%
India	5.3	2.6	1.9	1.6	2.8	52%
Japan	4.9	3.2	0.5	3.8	1.7	34%
Ethiopia	4.4	2.5	2.7	0.7	1.9	44%
Germany	3.8	2.9	0.2	4.0	0.8	22%

Notes: g_λ denotes consumption-equivalent social welfare growth, g_c is the growth rate of per capita consumption, g_N is population growth, $v(c)$ is the value of life year relative to consumption, and the population share is $v(c) \cdot g_N / g_\lambda$. Data on consumption per capita and the population is from Penn World Tables 10.0.

Figure 3: Welfare Growth in Select Countries, 1960–2019

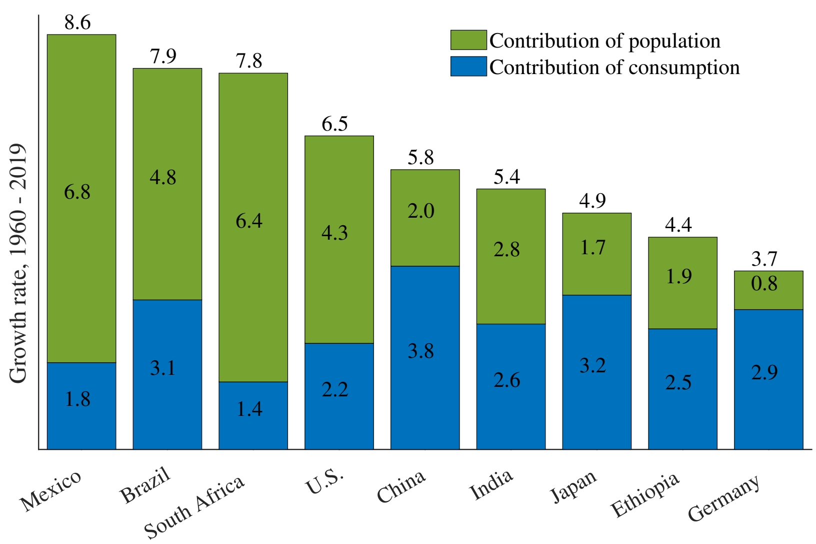

Notes: The numbers in the bars are CE-welfare growth, the percentage point contribution from population growth, and per capita consumption growth. Data are from the Penn World Tables 10.0, an updated version of Feenstra, Inklaar and Timmer (2015).

Figure 4: Welfare Growth in Select Countries for $v(c) = 2.7$

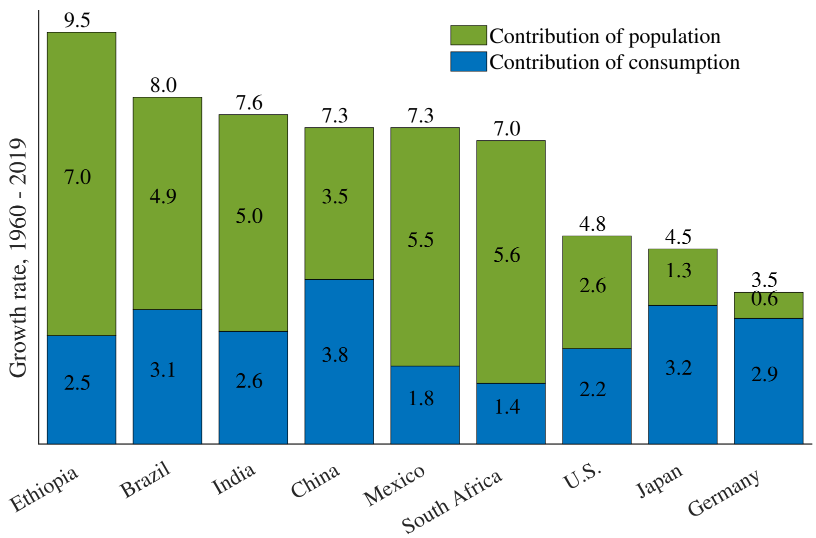

Notes: We impose a constant value of life of $v(c) = 2.7$ for all countries and years. The numbers in the bars are CE-welfare growth, the percentage point contribution from population growth, and per capita consumption growth. Data are from the Penn World Tables 10.0.

CE-welfare growth against consumption growth and population growth. The range of variation in CE-welfare growth is striking. Even the slowest-growing countries have growth rates of CE-welfare between 1960 and 2019 that exceed 2% per year. This constrasts with the negative consumption growth rates observed for a handful of countries. Equally striking is the fact that the fastest-growing countries have average annual growth rates of CE-welfare that exceed 10% per year, versus a maximum of 5% per year for consumption growth.

Neither consumption growth nor population growth are particularly highly correlated with CE-welfare growth. The correlation with consumption growth is 0.51, while the correlation with population growth is 0.29.

Figure 7 provides a different way of illustrating the difference between our CE-welfare growth and standard consumption growth measures by ranking countries from fastest to slowest growing. For example, China, Japan, and Germany are among the fast-growing countries over this period in terms of consumption growth, with

Figure 5: Plot of CE growth against consumption growth, 1960-2019

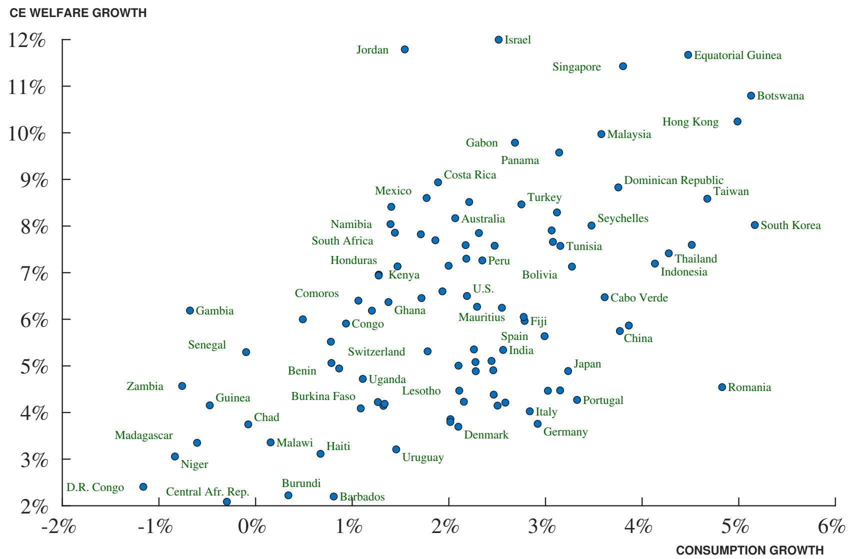

Figure 6: Plot of CE growth against population growth, 1960-2019

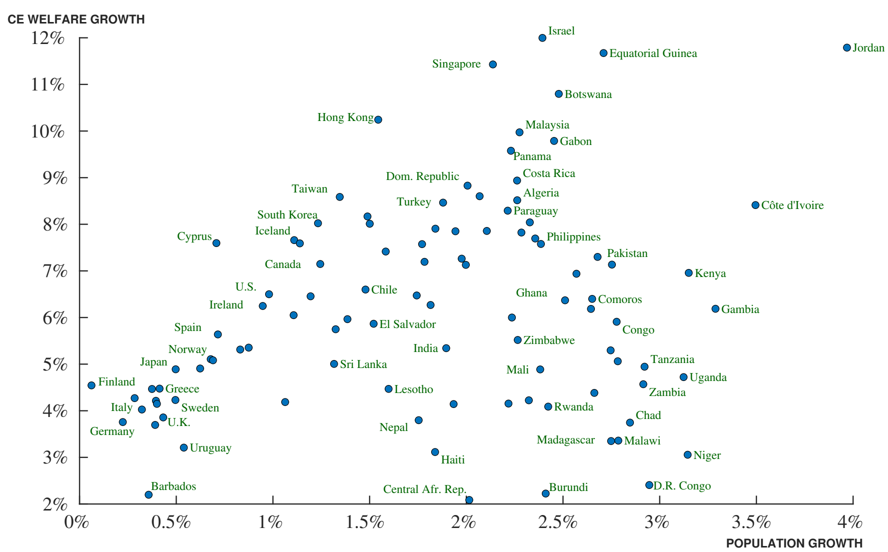

Figure 7: Changing Perspectives on Who is Growing Rapidly

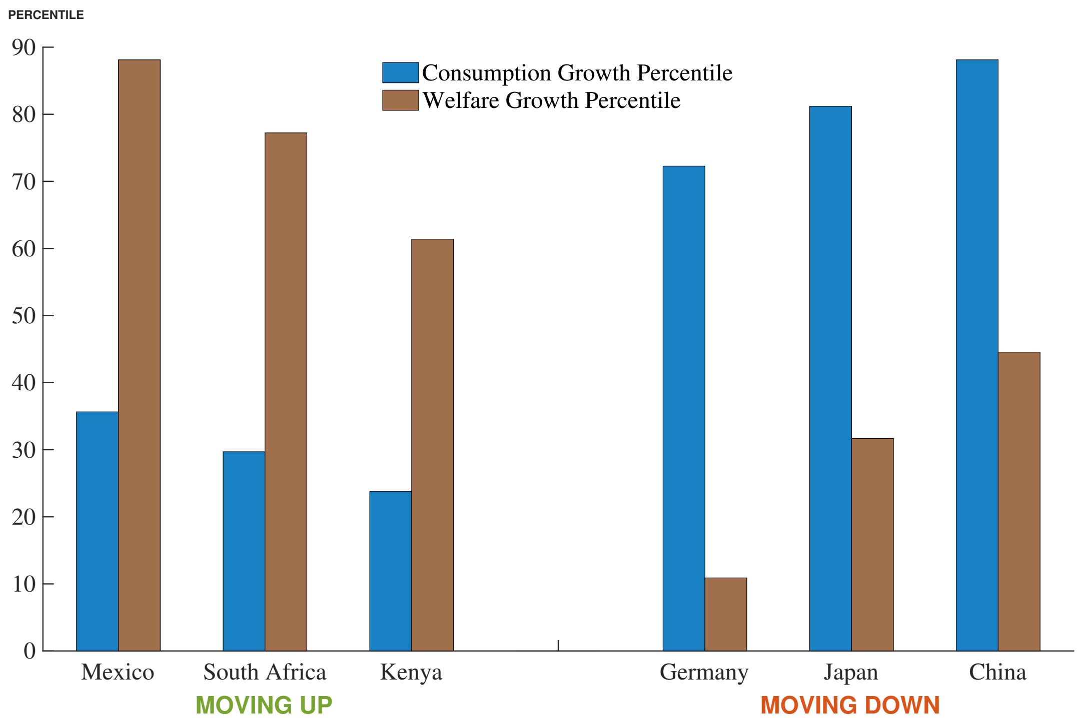

Notes: The chart shows the percentile in the cross-country distribution of growth rates between 1960 and 2019 for a select set of countries. Data is from the Penn World Tables 10.0.

China at around the 90th percentile. Slow population growth in these countries moves them sharply down in the distribution of CE-welfare growth, with Germany falling to just the 10th percentile and China falling to the 44th percentile.

In constrast, a number of countries with slow consumption growth move up sharply in the distribution. Mexico rises from the 35th percentile to the 88th, and Kenya rises from the 23rd percentile to the 61st.

3.2 Growth Rates in Sub-Periods

Figure 8 considers average annual growth starting in a later period, 1980 to 2019. Two things stand out in this figure. First is the negative correlation between consumption growth and CE-welfare growth for the first six countries in the chart. China

Figure 8: Welfare Growth in Select Countries, 1980–2019

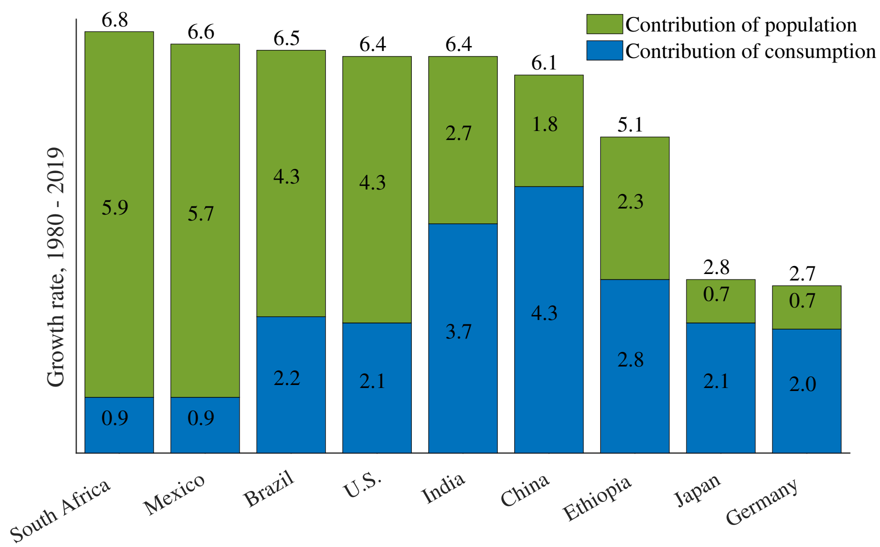

Notes: The numbers in the bars are CE-welfare growth, the percentage point contribution from population growth, and per capita consumption growth. Data are from the Penn World Tables 10.0.

Figure 9: Welfare Growth in Select Countries, 2000–2019

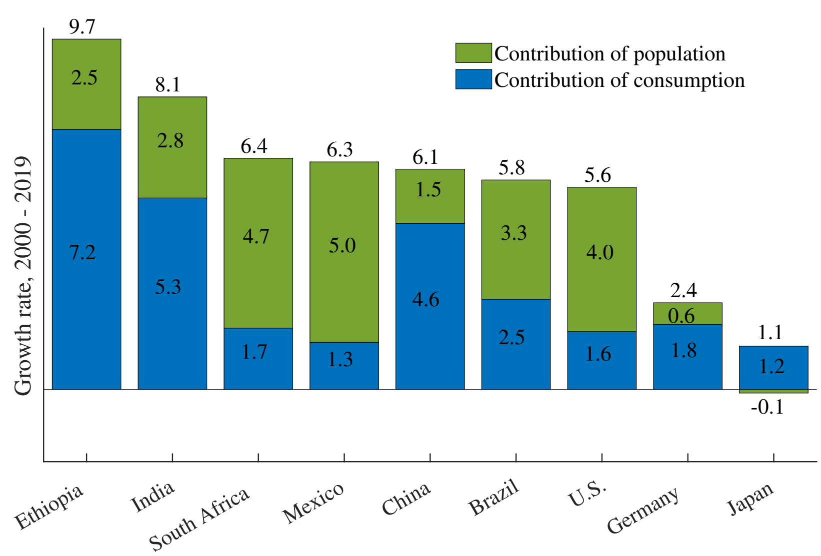

Notes: The numbers in the bars are CE-welfare growth, the percentage point contribution from population growth, and per capita consumption growth. Data are from the Penn World Tables 10.0.

has the fastest growth in per capita consumption among these six countries but the slowest growth in CE-welfare, falling behind the U.S., India, and Brazil, with all around 6% per year. The second feature of the graph that stands out is once again the very slow CE-welfare growth in Germany and Japan, driven by their very low rates of population growth.

Many of these trends are reinforced in even more recent time periods, as shown in Figure 9 for 2000–2019. CE-welfare growth in India rises to 8.1% per year, in China falls to 6.1% per year, and in Japan falls to just 1.1% per year, dragged down by a small negative population growth rate.

Another way of looking at these facts is to focus on specific countries over time. Figure 10 shows CE-welfare growth in Japan every decade since the 1960s. The well-known slowdown in Japanese growth is evident in the blue bars, which show consumption growth. But this slowdown is reinforced by declining rates of population growth. Overall CE-welfare growth slows from 9.7% per year in the 1960s to -0.3% per year in the 2010s. For this most recent decade, a negative population growth rate of -0.15% per year — when scaled up by $v(c)$ — more than offsets the modest consumption growth rate of 0.4%.

Figure 11 shows growth in China since the 1960s. Population growth in China (not shown) slows from 2.3% per year in the 1960s to just 0.5% per year in the 2010s. However, the rising value of life $v(c_t)$ to some extent offsets this decline: the contribution of population growth to CE-welfare growth falls from just 2.2% per year in the 1960s to 1.5% per year in the 2010s. CE-welfare growth has slowed since the 1990s, but the decline is modest, from 7.0% per year to 5.7% per year.

In contrast, the bulk of CE-welfare growth in Sub-Saharan Africa since the 1960s has been due to population growth, as shown in Figure 12. Population growth was relatively stable at just over 2.5% per year during the entire period, and the population term accounts for around 4pp of CE-welfare growth in Sub-Saharan Africa each decade. Consumption growth rose in the 2000s and 2010s, leading to a robust CE-welfare growth rate of more than 8% in the 2010s.

Figure 10: Growth in Japan by Decade

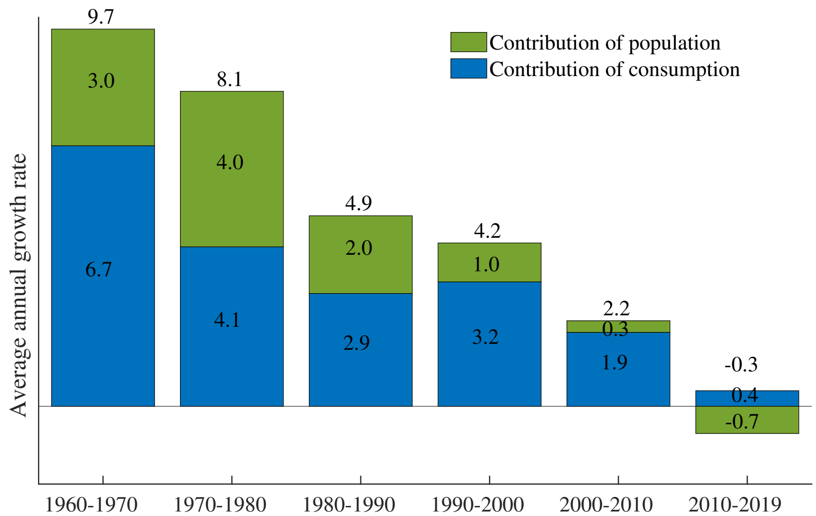

Notes: The numbers in the bars are CE-welfare growth, the percentage point contribution from population growth, and per capita consumption growth. Data are from the Penn World Tables 10.0.

Figure 11: Growth in China by Decade

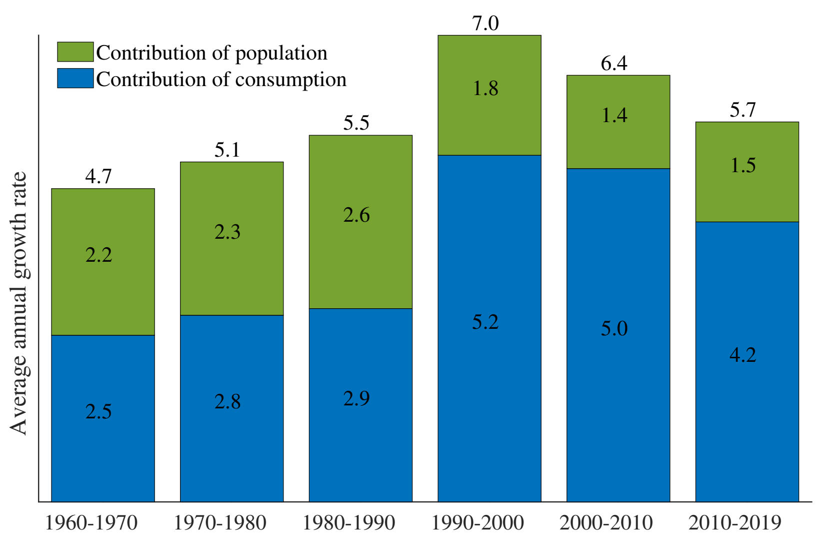

Notes: The numbers in the bars are CE-welfare growth, the percentage point contribution from population growth, and per capita consumption growth. Data are from the Penn World Tables 10.0.

Figure 12: Growth in Sub-Saharan Africa by Decade

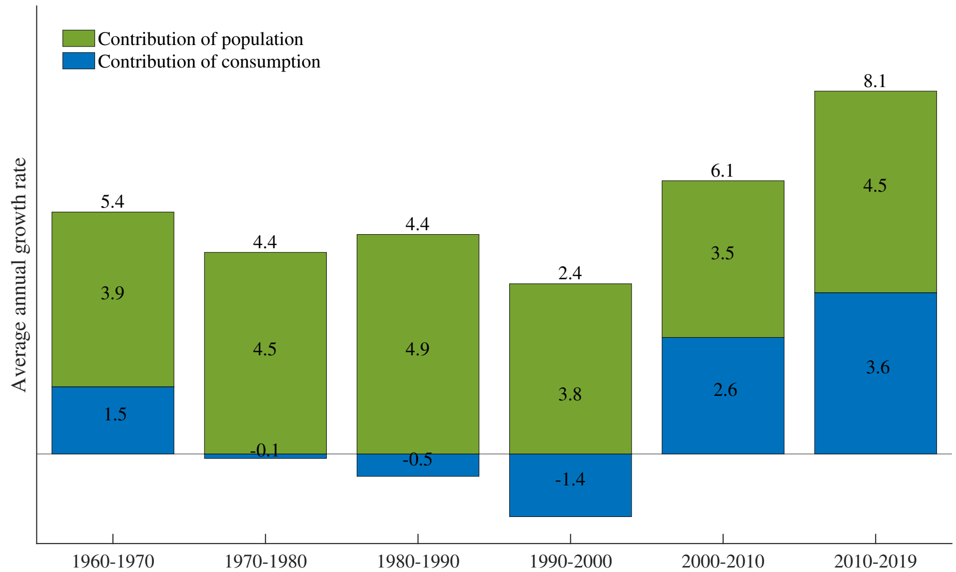

Notes: The numbers in the bars are CE-welfare growth, the percentage point contribution from population growth, and per capita consumption growth. Data are from the Penn World Tables 10.0.

3.3 Long-Run Growth in the U.S., the West, and the World

Using data from the Maddison Project from de Pleijt and van Zanden (2020), we can study CE-welfare growth back further in time. To begin, Figure 13 shows the level of various outcomes in the United States, with the values in 1820 normalized to 1.0. For example, consumption per capita increased by a fctor of 24 in the U.S. in the 200 years since 1820, implying an average growth rate of 1.6% per year. In comparison, aggregate consumption rose by a factor of 801 over this same period, reflecting a 33-fold increase in population. As we have emphasized in the paper, however, CE-welfare puts a larger weight on population. This is reflected in the graph by the fact that CE-welfare increased by a factor of around 600,000(!) in the U.S. since 1820, an average growth rate of 6.7% per year. A population growth rate that averaged 1.8% per year accounted for just over three-quarters of all CE-welfare growth in the United States over this period.

Figure 14 breaks the period since 1820 into twenty-year intervals so the time

Figure 13: U.S. cumulative growth, 1820–2018

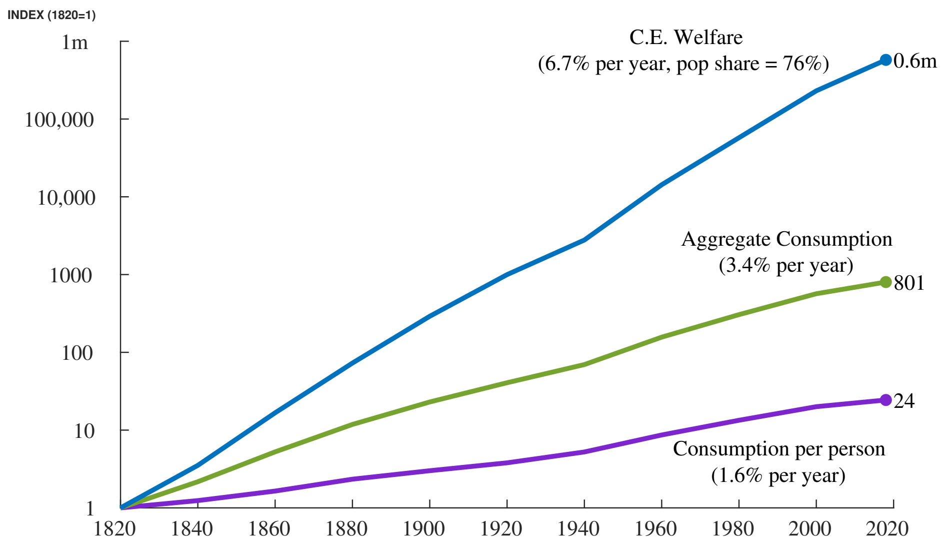

Note: Data the U.S. National Income and Product Accounts back to 1929 and from Barro and Ursua (2008) between 1820 and 1929.

series of U.S. growth is apparent. Interesting, while there are long-term swings in the growth rate, there is no overall trend. Growth was especially fast in decades of high immigration like the 1840–1880 period and 1940–1960. The period since 2000 is the slowest period of growth since 1820, tied only by the 1920–1940 period that encompassed the Great Depression and restricted immigration.

An obvious concern about the U.S. numbers is that migration among countries is a zero-sum game. The U.S. population has grown rapidly during the past 200 years in part because of migrants from the rest of the world. In Section 4 below, we adjust for migration on a country-by-country basis. For now, instead, we report results over long time periods for larger collections of countries, first for "The West" and then for "The World."[6]

Figure 15 shows cumulative growth since 1820 for "The West." Over the past 200 years, consumption per capita rises by a factor of 20 and aggregate consumption rises by a factor of 114, while CE-welfare rises by a factor of nearly 3000. This is much

[6] "The West" is Western Europe plus the "Western Offshoots" in de Pleijt and van Zanden (2020).

Figure 14: Growth in the United States, 1820–2018

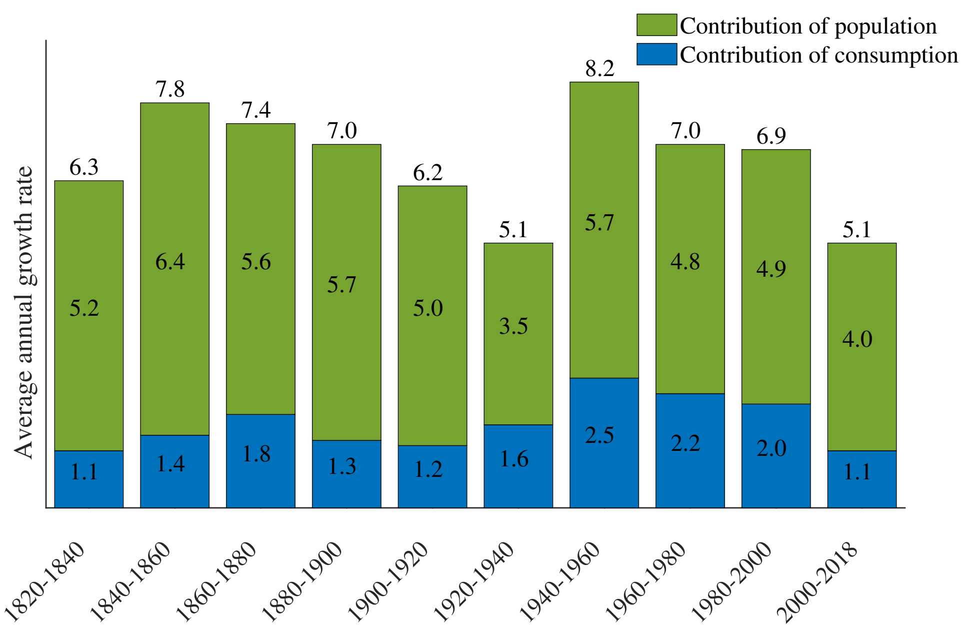

Note: Data the U.S. National Income and Product Accounts back to 1929 and from Barro and Ursua (2008) between 1820 and 1929.

Figure 15: Cumulative Growth in "The West," 1820–2018

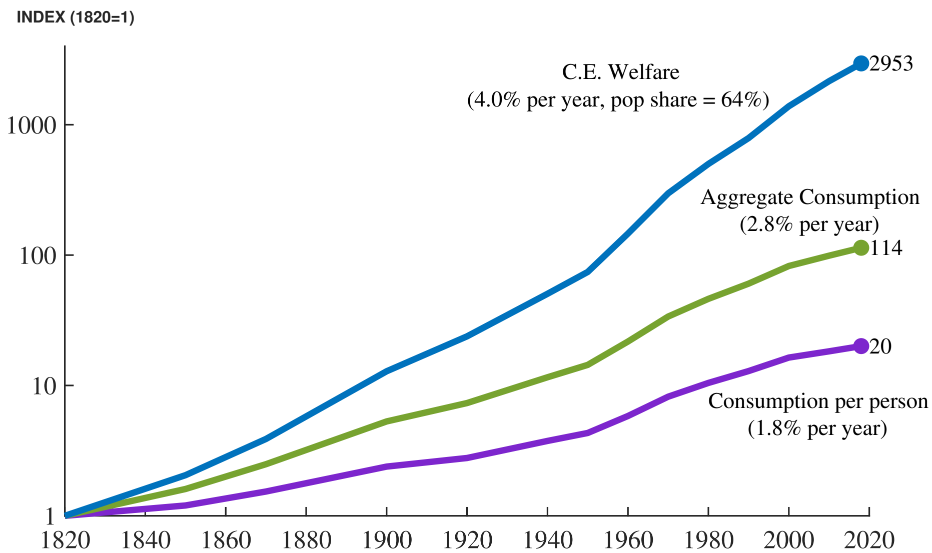

Notes: "The West" is the sum of Western Europe and "Western Offshoots" in the Maddison Project data of de Pleijt and van Zanden (2020). We estimate consumption as 0.8 times per capita GDP.

lower than the factor of 600,000 we saw for the U.S., largely reflecting the zero-sum nature of migration. Nevertheless, CE-welfare growth averages 4.0% per year over these two centuries, and population growth accounts for 64% of the overall gain.

Figure 16 shows the time series of CE-welfare growth for the past two centuries. While migration masked the "hump shape" for the United States, the hump is much more apparent for "The West" taken as a whole. CE-welfare growth peaks in the 1940–1970 period at 5.9% per year and falls to 4.7% per year by 1990–2018.

Finally, Figure 17 shows the gain in CE-welfare for the world as a whole since 1500. Over more than five centuries, consumption per person rises by a factor of 20, for an average annual growth rate of 0.6% per year. Population growth of 0.5% per year boosts the growth rate of aggregate consumption to 1.1%, which cumulates to a factor of 316 since 1500. But the gain in CE-welfare is more than 10-fold larger: it rose by a factor of 3700 since 1500, at an average annual growth rate of 1.6p̂er year. Gains in population account for 61% of this growth.

Figure 18 shows the time series of these growth rates. Between 1500 and 1850,

Figure 16: West CE growth over the long run, 1820-2018

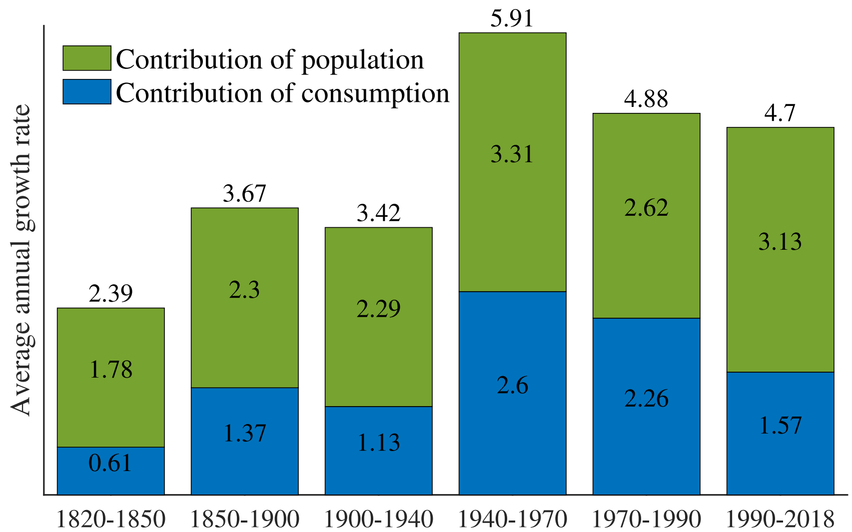

Notes: "The West" is the sum of Western Europe and "Western Offshoots" in the Maddison Project data of de Pleijt and van Zanden (2020). We estimate consumption as 0.8 times per capita GDP for this figure.

Figure 17: Cumulative Growth in "The World," 1500–2018

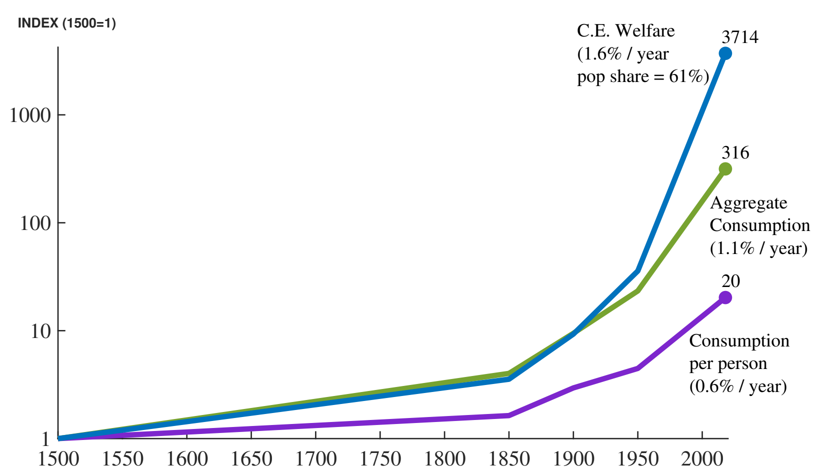

Note: Data from The Maddison Project data of de Pleijt and van Zanden (2020). We estimate consumption as 0.8 times per capita GDP for this figure.

Figure 18: World CE growth over the long run, 1500-2018

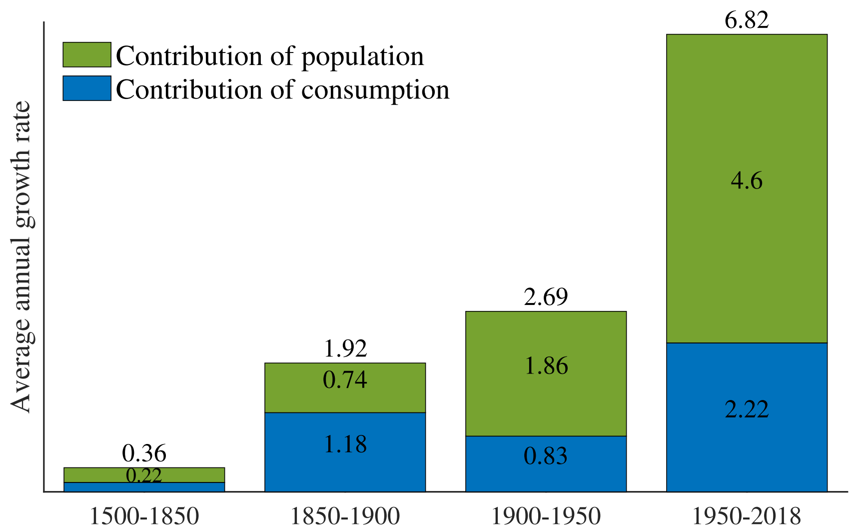

Note: Data from The Maddison Project data of de Pleijt and van Zanden (2020). We estimate consumption as 0.8 times per capita GDP for this figure.

CE-welfare growth for the world as a whole averaged just 0.36% per year. Growth accelerated over the next 150 years, reaching 6.8% per year for the world between 1950 and 2018. During this latter period, per capita consumption growth was 2.2% per year, with the value of population growth contributing more than twice as much as 4.6% per year and accounting for 67% of world CE-welfare growth.

4. Taking Migration into Account

Our calculations to this point credit countries for the growth in the number and standard of living of its resident populations. This makes no distinction based on where the individuals were born and consequently assigns the contribution of migrants to their destination country. Taking the other extreme, one might instead attribute people to the country in which they are born. Compared to our baseline calculation, we can add flow utility for out-migrants and subtract flow utility from in-migrants:[7]

$$W_{it} = N_{it} \cdot u(c_{it}) + \sum_{j \neq i} N_{i \to j,t} \cdot u(c_{jt}) - \sum_{j \neq i} N_{j \to i,t} \cdot u(c_{it})$$

where $N_{i \to j,t}$ is the population born in country i and living in country j in year t and $N_{j \to i,t}$ is the population born in country j living in country i in year t.

Growth in country welfare adjusted for migration is then

$$\begin{aligned} g_{\lambda_{it}} = \;& v(c_{it}) \cdot g_{N_{it}} \;+\; g_{c_{it}} \\ &+ \sum_{j \neq i} \frac{N_{i \to j,t}}{N_{it}} \cdot \frac{u(c_{jt})}{u(c_{it})} \left(v(c_{it}) \cdot g_{N_{i \to j,t}} + \frac{v(c_{it})}{v(c_{jt})} \cdot g_{c_{jt}} \right) \\ &- \sum_{j \neq i} \frac{N_{j \to i,t}}{N_{it}} \left(v(c_{it}) \cdot g_{N_{j \to i,t}} + g_{c_{it}} \right) \end{aligned}$$

The first term is our baseline, which credits all immigrants to the *destination* country. The second term adds in growth from out-migrants, and the third term subtracts growth from in-migrants. This adjusted measure therefore credits migrants to the *source* country.

To implement this migration adjustment, we use data from the World Bank's Global Bilateral Migration Database. This contains data on what share of each country's resident population in select years (1960, 1970, 1980, 1990, and 2000) was born in each country of origin. We have the necessary data to make our migration adjustment for 81 countries.

Figure 19 plots migration-adjusted welfare growth vs. our baseline welfare growth for the 81 countries from 1960 to 2000. The points hug the 45 degree line as results with and without the migration adjustment are highly correlated at 0.92. The adjustments for individual countries can be sizable, though they still don't alter the important role of population growth. Figure 20 shows countries in which in-migration

[7] We could also explore intermediate cases such as giving countries credit for the higher consumption enjoyed by in-migrants from poorer countries.

Figure 19: Baseline vs. Migration-Adjusted CEW growth

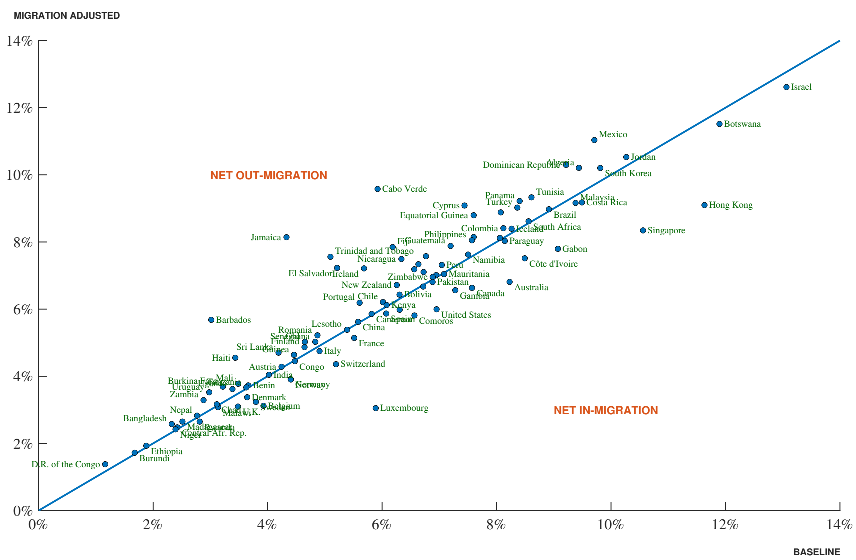

arguably overstates welfare growth, including the United States, Hong Kong, Australia, Germany, France and the U.K. And Figure 21 illustrates that out-migration potentially understates welfare growth for Ireland, Mexico, Turkey, Portugal, Algeria, and the Philippines.

5. Beyond Consumption

One limitation of our simple approach above is that it only incorporates consumption in flow utility. For example, it ignores utility from leisure. Where parents are concerned, it also ignores fertility decisions and time investments in children that can yield flow utility. In particular, over time parents may have increasingly made a "quantity-quality" tradeoff, choosing fewer kids, but investing more in each (Barro and Becker, 1989a). In the process they may have also substituted away from childcare towards leisure.

We therefore extend our framework in this section to incorporate parental fer-

Figure 20: Countries for which in-migration biases our baseline upward

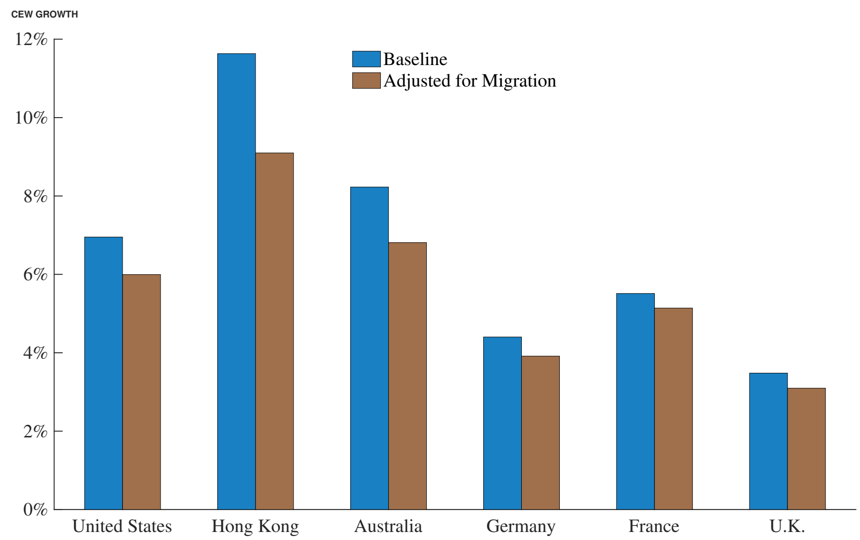

Figure 21: Countries for which out-migration biases our baseline downward

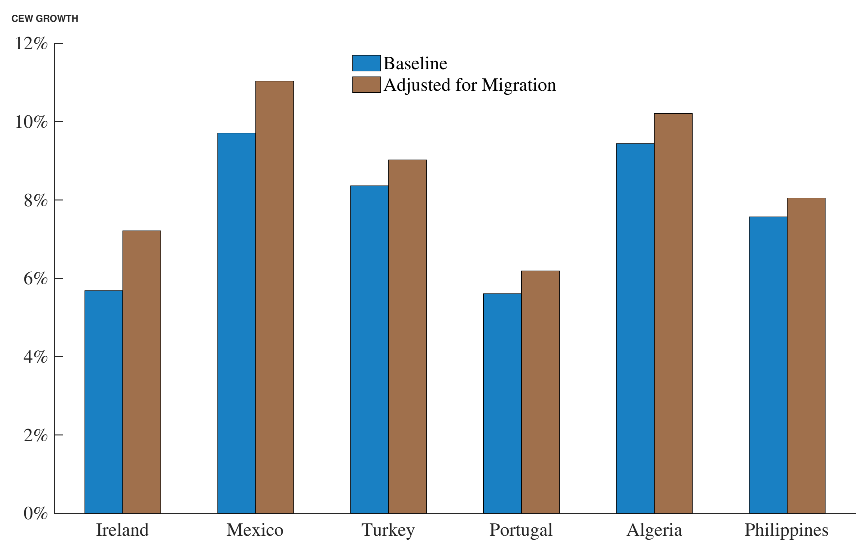

tility decisions that trade off altruism toward their kids (their consumption and human capital) with their own parental consumption and leisure time.

5.1 Framework

Suppose total flow welfare takes the form:

$$W(N_t^p,\, N_t^k,\, c_t^p,\, l_t,\, c_t^k,\, h_t^k,\, b_t) \;=\; N_t^p \cdot u(c_t^p,\, l_t,\, c_t^k,\, h_t^k,\, b_t) + N_t^k \cdot \widetilde{u}(c_t^k)$$

where N^p is the number of adults ("p" for parents), N^k is the number of children ("k" for kids), b is number of children per adult, c^p is adult consumption, l is adult leisure, c^k is consumption per child, and h^k is human capital per child. Total population satisfies $N = N^p + N^k = (1 + b) \cdot N^p$.

Aggregate flow welfare is the sum of all parents' flow welfare (from their own consumption, their own leisure, their kids' consumption during childhood, their kids' human capital, and the number of kids per parent) and all kids' flow welfare. The fact that the consumption, human capital, and number of kids affect parental utility is reminiscent of Barro and Becker (1989b) and Farhi and Werning (2007); something like this is necessary to explain why parents have kids and invest resources in them.

We make kids' flow welfare a function of their consumption only. We have in mind that their leisure is fixed at one so is suppressed and that the kids will enjoy the benefits of their human capital in the form of higher consumption when they are themselves adults.

To calculate consumption-equivalent welfare growth, we ask by what factor λ_t one would have to scale up both parents' and kids' consumption at t to match the flow utility at $t + dt$ given the changing numbers of parents and kids and changing per capita variables:

$$[W(N_t^p,\, N_t^k, \lambda_t \cdot c_t^p,\, l_t,\, \lambda_t \cdot c_t^k,\, h_t^k,\, b_t) = W(N_{t+dt}^p\,,\, N_{t+dt}^k\,,\, c_{t+dt}^p\,,\, l_{t+dt}\,,\, c_{t+dt}^k\,,\, h_{t+dt}^k\,,\, b_{t+dt})$$

To solve for λ_t, consider adjusted social welfare:

$$W(\lambda_t) = N_t^p \cdot u\left(\lambda_t c_t^p,\ l_t, \lambda_t\, c_t^k,\ h_t^k,\ b_t\right) + N_t^k \cdot \widetilde{u}\left(\lambda_t c_t^k\right).$$

We then set $dW/W = 0$ and solve for growth in consumption-equivalent welfare $g_{\lambda t} \equiv -d\lambda_t/\lambda_t$ around the initial level of welfare with $\lambda = 1$:

$$\begin{aligned} g_{\lambda_t} = \kappa_t \Bigg[& \omega_t^p \left(\frac{dN_t^P}{N_t^P} + \frac{u_{c_t^p} c_t^p}{U_t} \frac{dc_t^p}{c_t^p} + \frac{u_{lt} l_t}{U_t} \frac{dl_t}{l_t} + \frac{u_{c^k t} c_t^k}{U_t} \frac{dc_t^k}{c_t^k} + \frac{u_{h^k t} h_t^k}{U_t} \frac{dh_t^k}{h_t^k} + \frac{u_{bt} b_t}{U_t} \frac{db_t}{b_t} \right) \\ & + \omega_t^k \left(\frac{dN_t^K}{N_t^K} + \frac{\tilde{u}'(c_t^k) c_t^k}{\tilde{u}(c_t^k)} \frac{dc_t^k}{c_t^k} \right) \Bigg] \end{aligned}$$

where ω_t^p and ω_t^k are the total welfare shares of parents and kids in year t, and κ_t puts this expression in consumption equivalent units.

In our framework, parental decisions are privately optimal. Once we pick functional forms and calibrate parameters, optimality conditions from parents' utility maximization problem will allow us to express the weights on each growth rate above in terms of observables. Specifically, we assume that parents at each date solve the following problem:

$$\begin{aligned} \max_{c^p,\, l,\, c^k,\, h^k,\, b} \quad & u(c_t^p,\ l_t,\ c_t^k,\ h_t^k,\ b_t) \\ \text{subject to:} \quad & c_t + b_t \cdot c_t^k \leq w_t \cdot h_t \cdot l_{c_t} \\ & h_t^k = f(h_t \cdot e_t) \\ & \text{and } \ l_{c_t} + l_t + b_t \cdot e_t \leq 1 \end{aligned}$$

where w is the real wage per unit of human capital, h is parental human capital, h^k is kids' human capital, l_c are parental hours worked, and e is parental time investment per child. Parents spend their earnings on their own consumption and their kids' consumption. Kids' human capital is an increasing function of their parents' human capital and their parents' time investment in them. And parents have a unit of time to spend on work, leisure, and time with their kids.

To make progress, we assume these specific functional forms for parents' and kids' flow utility, respectively:

- **Assumption 1:** $u(c_t^p,\ l_t,\ c_t^k,\ h_t^k,\ b_t) = \log(c_t^p) + \alpha b_t^\theta \cdot \log(c_t^k) + g(l_t, h_t^k, b_t)$
- **Assumption 2:** $\widetilde{u}(c^k) = \bar{u}_k + \log(c_t^k)$

In Assumption 1, $\alpha > 0$ and $\theta > 0$ are parameters governing parental altruism towards their kids. In the special case where $\alpha = 1$ and $\theta = 1$ parents are total utilitarians with respect to their own family. The literature often considers cases with $\alpha < 1$ and $\theta < 1$ (Doepke and Tertilt, 2016).

With these functional forms, growth in consumption-equivalent welfare is:

$$
\begin{aligned}
g_{\lambda_t} &= \pi_t^p \cdot v\left(c_t^p, c_t^k, \vec{x}_t\right) \cdot \frac{dN_t^p}{N_t^p} \ +\ \pi_t^k \cdot \tilde{v}(c_t^k) \cdot \frac{dN_t^k}{N_t^k} && \text{Population} \\
&+ \pi_t^p \cdot \frac{dc_t^p}{c_t^p} \ +\ (1-\pi_t^p) \cdot \frac{dc_t^k}{c_t^k} && \text{Consumption} \\
&+ \pi_t^p \cdot \frac{u_{lt}\ l_t}{u_{ct}\ c_t} \cdot \frac{dl_t}{l_t} && \text{Leisure} \\
&+ \pi_t^p \cdot \frac{u_{bt}\ b_t}{u_{ct}\ c_t} \cdot \frac{db_t}{b_t} && \text{Quantity of kids} \\
&+ \pi_t^p \cdot \frac{u_{h^k t}\ h_t^k}{u_{ct}\ c_t} \cdot \frac{dh_t^k}{h_t^k} && \text{Quality of kids}
\end{aligned}
$$

where

$$
\pi_t^p = \frac{N_t^p}{(1+\alpha b_t^\theta)N_t^p + N_t^k} \quad ; \quad \pi_t^k = \frac{N_t^k}{(1+\alpha b_t^\theta)N_t^p + N_t^k}
$$

$$
v\left(c_t^p, c_t^k, \vec{x}_t\right) = v\left(c_t^p, l_t, c_t^k, h_t^k, b_t\right) = \frac{u\left(c_t^p, l_t, c_t^k, h_t^k, b_t\right)}{u_c\left(c_t^p, l_t, c_t^k, h_t^k, b_t\right) \cdot c_t^p} \quad ; \quad \tilde{v}(c_t^k) = \frac{\tilde{u}(c_t^k)}{\tilde{u}'(c_t^k) \cdot c_t^k}
$$

The first line in the CEW growth expression is the new version of the "population growth" term. This population term differs from the simple $g_N \cdot v(c)$ specification in previous sections for several reasons. First, parents' value of a year of life v and kids' value of a year of life $\widetilde{v}$ may differ. Second, the value of a year of life for parents depends on not only their own consumption but also on their kids' consumption, their own leisure, their own fertility, and their kids' human capital. Third, we have a scaling factor out in front which is less than one. The intuition for this is that the λ factor enters three times rather than just twice: once for the parents, once for the kids, and then once because the parents themselves care about their kids' consumption (the αb_t^θ term). A smaller increase in consumption is then required to value the additional people.

The remaining lines in the CEW growth expression are the new version of the "per capita growth" term. It now includes growth in leisure, kids' human capital, and fertility along with growth in consumption per parent and per kid. Note that the weight on parent terms π_p is *less than* the share of parents in the population, and the corresponding weight on kids' consumption growth $1 - \pi_p$ exceed the share of kids in the population. This likewise reflects parental altruism, which results in "double counting" (upweighting) the growth of kids' consumption. This point was emphasized by Caplin and Leahy (2004) and Farhi and Werning (2007).

Illustrative example. A special case of this growth accounting is helpful for intuition. Suppose $\alpha = 1$ and $\theta = 1$, which implies $dc^k/c^k = dc^p/c^p$, and evaluate at a point where the value of a year of life happens to be the same for parents and kids, $\widetilde{v}(c_t^k) = v(c_t^p, \vec{x}_t) = v(c_t)$. Then CEW growth becomes

$$
\begin{aligned}
g_{\lambda_t} \;=\; & \frac{dc_t}{c_t} + \frac{N_t^p + N_t^k}{N_t^p + 2N_t^k} \cdot v(c_t) \cdot \frac{dN_t}{N_t} \\
& + \frac{N_t^p}{N_t^p + 2N_t^k} \cdot \left(\frac{u_{lt} l_t}{u_{ct} c_t} \cdot \frac{dl_t}{l_t} + \frac{u_{bt} b_t}{u_{ct} c_t} \cdot \frac{db_t}{b_t} + \frac{u_{h^k t} h_t^k}{u_{ct} c_t} \cdot \frac{dh_t^k}{h_t^k} \right)
\end{aligned}
$$

Note the 2 appearing in the denominators on all terms other than consumption growth. The double counting of kids' consumption (their own utility and their parents' utility from it) downweights all non-consumption terms: we need to scale up consumption by a smaller amount to match the value of more people, leisure, etc.

5.2 Implementation

We use parents' first order conditions to map weights in the growth accounting to observables. Specifically,

$$l_t: \quad \frac{u_{lt} l_t}{u_{ct} c_t} = \frac{w_t h_t l_t}{c_t} \tag{4}$$

$$b_t: \quad \frac{u_{bt} b_t}{u_{ct} c_t} = b_t \frac{(c_t^k + w_t h_t e_t)}{c_t} \tag{5}$$

$$h_t^k: \quad \frac{u_{h^k t} h_t^k}{u_{ct} c_t} = b_t \frac{1}{\eta_t} \frac{w_t h_t e_t}{c_t}, \text{ where } \eta_t = \frac{f'(h_t e_t) h_t e_t}{f(h_t e_t)} \tag{6}$$

Equation (4) says that that the weight on leisure growth should be tied to the marginal rate of substitution between consumption and leisure, which in turn equals earnings relative to consumption. Equation (5) says the weight on fertility growth is connected to the marginal rate of substitution between fertility and consumption. The latter can be assessed using total spending on kids (including foregone earnings due to time spent investing in kids' human capital) relative to adult consumption. Equation (6) indicates that the weight on human capital growth is related to the marginal rate of substitution between human capital and consumption, which equals implicit spending on kids' human capital relative to adult consumption.

Calibration. The weight given to a child's human capital growth partly reflects the elasticity of a child's human capital with respect to parental input: $\eta_t = \frac{f'(h_t e_t) h_t e_t}{f(h_t e_t)}$. We first impose a constant η. To calibrate η we exploit that h_t^k's elasticity with respect to $h_t e_t$ is the same as for h_t alone. We base that elasticity on Mincer-equation estimates by Lee, Roys and Seshadri (2015), who include schooling of a respondent's parents, as well as the respondent's, as predictors of the respondent's wage. Assuming that (i) the respondent's schooling coefficient proxies for the impact of parental schooling on the parents' own human capital, and (ii) that parents' choice of e_t is orthogonal to their schooling, then η is identified by the estimated impact of parental schooling on the respondent's wage relative to the impact of their own. This ratio, summing the impacts of both parents' schooling, equals 0.24 (=.0142/.0591).

To calibrate the parameters governing parental altruism towards their kids, α and θ, we rely on a USDA study (Lino, 2011) of spending on kids versus parents within households. Note that, under Assumptions 1 and 2, the first-order conditions from the parents' utility maximization problem imply:

$$\frac{c_t^k}{c_t^p} = \alpha b_t^{\theta - 1}.$$

For example, for $\theta = 1$ and $\alpha = 1$ parents equate each kid's consumption to their own. From Lino (2011), households with two parents and two children, for whom $b = 1$, spend approximately two-thirds as much on the children as the parents. From this we calibrate $\alpha = 2/3$. By contrast, two-parent household with one child spend somewhat more *per* child; those with three children spend somewhat less. These patterns are consistent with a value for θ of about 0.8. We treat this as a baseline while considering robustness to $\theta = .6$ and $\theta = 1$

As in previous sections, we target $v(c_t^p, c_t^k, \vec{x}_t) = 4.87$ for the US in 2006. To calibrate $\tilde{v}(c_t^k)$, we assume the value of a year of life to be the same for a child and an adult in the U.S. in 2006.[8] We consider robustness to $\tilde{v}(c_t^k)/v(c_t^p, c_t^k, \vec{x}_t) = 0.8$ or 1.2 for the US in 2006.

We employ welfare accounting to benchmark other countries' levels for $v(c_t^p, c_t^k, \vec{x}_t)$ and $\tilde{v}(c_t^k)$. We chain welfare in the country with the second-highest level of per capita consumption in 2006, the Netherlands, to that with the highest, the United States, based on their differences in consumption, leisure, number of children, and children's human capital. In the same way, we proceed to link the third richest, Japan, to the Netherlands, and so forth. We then chain $v(c_t^p, c_t^k, \vec{x}_t)$ and $\tilde{v}(c_t^k)$ through time within countries to reflect the growth rates in each of their arguments.[9]

Data. As in previous sections, consumption and total population are from the Penn World Table 10.0. The total number of children (0-19 years old) is from the World Bank. We combine data on total hours worked (Penn World Table) and on working age population (World Bank) to calculate hours worked per adult. We mea-

[8] Given our additively separable preferences, this implies equal utility flows in 2006 in the U.S.

[9] In linking welfare through time within countries we use Tornqvist weights to value the factors. (E.g, growth on leisure from 2006 to 2007 is weighted by the average of relative time allocated to leisure in 2006 and 2007). In linking two countries in 2006, we use the average of their weights on an argument in that year to weight the country differences. (E.g, the percent leisure difference between the U.S. and the Netherlands in 2006 is weighted by the average of relative time allocated to leisure in 2006 in the U.S. and in the Netherlands.)

sure parental time investments in kids using data on childcare from time use surveys. Leisure is then the residual after subtracting hours worked and total childcare from waking time, which we set equal to 16. Finally, to obtain growth in human capital, we assume an even split of real wage growth between human capital and real wage per unit of human capital.

The most stringent requirement is the availability of micro data from consistent time use surveys. Such data were available for the following country-years: United States (2003-2019), Netherlands (1975-2006), Japan (1991-2016), South Korea (1999-2019), Mexico (2006-2019), and South Africa (2000-2010).

5.3 Results

Table 3 presents our calculations of consumption-equivalent welfare (CEW) growth based on macro data alongside our new micro data based calculations for the same six countries. We adjust the period of the macro calculation to match the years for which we have micro data.

Table 3 shows modest net effects on total CEW growth from our added "per capita" terms (last three columns). The exception is Mexico, where leisure is falling and the rise in quality of kids is not large enough to offset their falling quantity.

The gap between the macro and micro CEW growth is due to the smaller population term in the micro results. This is the point we emphasized earlier: taking into account parental altruism toward their kids leads to double counting kids' consumption, so that a smaller increase in consumption is required to value the additional people. Table 3 shows that, quantitatively, this adjustment in the population term is modest, and population is still an important contributor to CEW growth.

Table 4 gives the *share* of CEW growth due to population in each of the six countries. We first note that this fraction is fairly similar between our macro and micro calculations. The Table also reports the effects of considering higher or smaller values of θ (the parameter governing diminishing returns in utility from having kids) or v_k relative to v_p for the US in 2006 (respectively, kid's and adult's value of a year of life relative to consumption). The baseline value of θ is 0.8 and the larger and smaller

Table 3: CEW Growth: Macro versus Micro Calcualtions

	MACRO			MICRO					
	CEW growth	pop term	cons term	CEW growth	pop term	cons term	leisure term	quality term	quantity term
USA	5.4	3.9	1.5	4.8	3.2	1.5	0.1	0.2	−0.3
NLD	4.5	2.5	2.1	4.0	2.0	2.1	0	0.4	−0.4
JPN	2.3	0.4	1.9	1.8	0.1	1.9	0	0.2	−0.4
KOR	4.4	1.7	2.6	3.7	0.9	2.6	0.6	0.4	−0.8
MEX	6.5	4.9	1.6	3.8	3.3	1.6	−0.3	0.1	−0.8
ZAF	6.8	4.3	2.6	5.9	3.1	2.6	1	0.3	−1

Notes: 'MACRO' results are based on the framework presented in Section 2, while the 'MICRO' results are based on the augmented framework presented in Section 5. CEW denotes consumption-equivalent welfare growth, decomposed in subsequent columns to show contribution of the different terms. The period is 2003-2019 for the United States, 1975-2006 for Netherlands, 1991-2016 for Japan, 1999-2019 for Korea, 2006-2019 for Mexico, and 2000-2010 for South Africa. Data sources are the Penn World Table 10.0 for population, consumption, and hours worked, time use surveys for fertility ("quality"), World Bank data on population for the number of kids per adult ("quantity").

values are 1.0 and 0.6. The baseline value of v_k/v_p is 1, and the larger and smaller ratios are 1.2 and 0.8. The share of growth due to population growth changes only modestly when we entertain these different parameter values.

Table 4: Share of population growth in CEW growth: Macro versus Micro

	MACRO	MICRO				
		Baseline	Robustness			
			Larger θ	Smaller θ	Larger v_k	Smaller v_k
USA	72%	68%	69%	66%	68%	67%
NLD	54%	50%	52%	48%	49%	52%
JPN	16%	5%	8%	3%	−9%	16%
KOR	40%	24%	27%	21%	15%	32%
MEX	76%	88%	90%	85%	87%	88%
ZAF	63%	53%	55%	51%	51%	55%

Notes: CEW denotes consumption-equivalent social welfare growth. The share of growth due to population growth is the ratio of the population terms to overall CEW growth. For data sources and years see the notes to Table 3. The baseline value of θ is 0.8 and the larger and smaller values are 1.0 and 0.6. The baseline value of v_k/v_p is 1, and the larger and smaller ratios are 1.2 and 0.8.

6. Conclusion

We calculated consumption-equivalent welfare growth based on total utility, including population growth, for many countries and time periods. We consistently find that population growth contributed between 1/2 to 2/3 of growth in country welfare. Because consumption runs into diminishing returns, each additional point of population growth is worth about five percentage points of per capita consumption growth in rich countries. Even taking into account developing countries a percent of population growth is worth 2.7 percentage points of per capita consumption growth over our sample of countries from 1960 to 2019.

Countries with slow population growth — such as China, Japan, and Germany — plummet in our growth rankings. In contrast, middle-income countries exhibiting above-average population growth, such as Mexico, Brazil, and South Africa, move up in our growth rankings.

We found our results to be robust to adjusting for migration and incorporating parent utility from children and privately optimal fertility choices. Crediting migration entirely to source countries has modest net effects in most countries, and does not alter our conclusion that at least half of growth in welfare comes from population growth. Similarly, taking into account intergenerational utility has modest net effects because leisure exhibits little trend and the "quality" of kids is rising to offset the falling quantity of kids.

References

Barro, Robert J and Gary S Becker, "Fertility choice in a model of economic growth," *Econometrica: journal of the Econometric Society*, 1989, pp. 481–501.

Barro, Robert J. and Gary S. Becker, "Fertility Choice in a Model of Economic Growth," *Econometrica*, March 1989, *57* (2), 481–501.

Barro, Robert J and Jose F Ursua, "Consumption disasters in the twentieth century," *American Economic Review*, 2008, *98* (2), 58–63.

Caplin, Andrew and John Leahy, "The social discount rate," *Journal of political Economy*, 2004, *112* (6), 1257–1268.

Chichilnisky, Graciela, Peter J Hammond, and Nicholas Stern, "Fundamental Utilitarianism and Intergenerational Equity with Extinction Discounting," *Social Choice and Welfare*, 2020, *54* (2-3), 397–427.

Cordoba, Juan Carlos, "Children, dynastic altruism and the wealth of Nations," *Review of Economic Dynamics*, 2015, *18* (4), 774–791.

Cutler, David M., Kaushik Ghosh, Kassandra L. Messer, Trivellore Raghunathan, Allison B. Rosen, and Susan T. Stewart, "A Satellite Account for Health in the United States," *American Economic Review*, February 2022, *112* (2), 494–533.

De la Croix, David and Matthias Doepke, "A soul's view of the optimal population problem," *Mathematical Social Sciences*, 2021, *112*, 98–108.

de Pleijt, Alexandra and Jan Luiten van Zanden, "The Maddison Project," 2020.

Doepke, Matthias and Michèle Tertilt, "Families in macroeconomics," in "Handbook of macroeconomics," Vol. 2, Elsevier, 2016, pp. 1789–1891.

Farhi, Emmanuel and Ivan Werning, "Inequality and Social Discounting," *Journal of Political Economy*, 2007, *115* (3), 365–402.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, 2015, *105* (10), 3150–82.

Golosov, Mikhail, Larry E Jones, and Michele Tertilt, "Efficiency with endogenous population growth," *Econometrica*, 2007, *75* (4), 1039–1071.

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

Kremer, Michael, Stephen P Luby, Ricardo Maertens, Brandon Tan, and Witold Wiecek, "Water Treatment And Child Mortality: A Meta-Analysis And Cost-effectiveness Analysis," Technical Report, National Bureau of Economic Research 2023.

Lee, Sang Yoon Tim, Nicolas Roys, and Ananth Seshadri, "The causal effect of parents' education on children's earnings," *Manuscript. https://www. ssc. wisc. edu/˜ aseshadr/WorkingPapers/spillover. pdf [accessed 26 October 2021]*, 2015.

Lino, Mark, "Expenditures on children by families," Technical Report, U.S. Department of Agriculture, Center for Nutrition Policy and Promotion, Miscellaneous Publication No. 1528-2011 2011.

MacAskill, William, *What We Owe the Future*, Basic books, 2022.

Parfit, Derek, "Reasons and Persons Oxford University Press," in "Oxford" 1984.

Rosen, Sherwin, "The value of changes in life expectancy," *Journal of Risk and uncertainty*, 1988, *1* (3), 285–304.

U.S. Environmental Protection Agency, "Mortality Risk Valuation," Technical Report 2020. Accessed April 20, 2020.

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* (1), 5–76.

WHO, "Commisssion on Macroeconomics and Health Investing in Health for Economic Development," 2001.

Young, Alwyn, "The gift of the dying: The tragedy of AIDS and the welfare of future African generations," *The Quarterly Journal of Economics*, 2005, *120* (2), 423–466.

Zuber, Stéphane, Nikhil Venkatesh, Torbjörn Tännsjö, Christian Tarsney, H Orri Stefánsson, Katie Steele, Dean Spears, Jeff Sebo, Marcus Pivato, Toby Ord et al., "What should we agree on about the repugnant conclusion?," *Utilitas*, 2021, *33* (4), 379–383.