Chapter two: **How to Think About Policy**

Tyler Cowen
Department of Economics
George Mason University Fairfax, VA 22030
tcowen@gmu.edu

*This is from a manuscript in progress called **The Growth Imperative: A New Approach to the Theory of Economic Policy**. Comments are most welcome, and references are available on request.

July 3, 2005

Economic growth alleviates human misery and lengthens human lives. Wealthier societies have better living standards, better medicines, and offer greater autonomy, greater fulfillment, and more sources of fun.

As recently as the end of the nineteenth century, life expectancy in Western Europe ran about forty years of age. Polio, tuberculosis, and typhoid were common ailments, even among the rich. Most individuals worked at hard physical labor, and a college or university education was a luxury.

Robert Fogel (2004, pp.8, 9, 34) paints a grim picture of the European past:

"...at that time [eighteenth and early nineteenth centuries] food constituted between 50 and 75 percent of the expenditures of laboring families...however...the energy value of the typical diet in France at the start of the eighteenth century was as low as that of Rwanda in 1965, the most malnourished nation for that year in the tables of the World Bank. England's supply of food per capita exceeded that of France by several hundred calories but was still exceedingly low by current standards. Indeed, as late as 1850, the English availability of calories hardly matched the current Indian level. One implication of these low-level diets needs to be stressed: Even prime-age males had only a meager amount of energy available for work…the average efficiency of the human engine in Britain increased by about 53 percent between 1790 and 1980."

Leisure time has grown as well. Let us define discretionary time as those hours not spent working, eating, or engaging in the bare minimum of personal grooming. We find that "the lifetime discretionary hours spent earning a living have declined by about one-third over the past century...In 1880 four-fifths of discretionary time was spent earning a living. Today, the lion's share (59) percent is spent doing what we like. Moreover, it appears probable that by 2040, close to 75 percent of discretionary time will be spent doing what we like, despite a further substantial increase in discretionary time due to the continuing extension of the life span" (Fogel 2004, p.70).

Imagine a time traveler from the eighteenth century visiting the life of Bill Gates. He would witness television, automobiles, refrigerators, central heating, antibiotics, plentiful food, flush toilets, cell phones, personal computers, and affordable air travel, among other remarkable benefits. But the most impressive features of Gates's life, from a historical point of view, are those shared by most middle-class Americans today. The very existence of an advanced civilization – the product of cumulative economic growth -- confers immense benefits on ordinary citizens.[1]

Even today's poor, in the United States, enjoy a virtually unimaginable standard of living by previous times. In 1995, 41 percent of "poor" households owned their own homes. The average poor home had three bedrooms, one-and-a-half baths, a garage, and a porch or patio. Seventy percent of poor households own a car; 27 percent own two or more cars. Ninety-seven percent of the poor have a color television, and almost half of the poor have two or more color televisions. Sixty-four percent own microwave ovens, half have a stereo system, and over a quarter have an automatic dishwasher. Two-thirds of "poor" households classified as poor have air conditioning. (For purposes of comparison, only 36 percent of the entire U.S. population had air conditioning as recently as the early 1970s.) Today's American poor are more likely to be overweight than are middle-class persons. Most poor children today grow up to be, on average, one inch taller and ten pounds heavier than WWII American GIs.[2]

Just as the present appears remarkable from the vantage point of the past, our future may offer comparable advances in benefits. Continued robust growth might bring greater life expectancies, cures for debilitating diseases, and cognitive enhancements. Millions or billions of people will have much better and longer lives. Many features of modern life might someday seem as backward as we now regard the large number of women who died in childbirth for lack of proper care. Most of all, economic growth limits and mitigates tragedies. It is a simple failure of imagination to believe that human progress has run its course.

---

[1] I am indebted to Don Boudreaux for this way of framing the point.
[2] The U.S. Census provides the data, I am indebted to the summary by Rector (1998).

The economic growth of the wealthier countries benefits the very poor as well. Most generally, the evidence suggests that the bottom quintile of an economy shares proportionally in growth (Dollar and Kraay 2000). Furthermore a richer economy will have a greater capacity to absorb immigrants. Poor people who migrate to rich countries earn much higher incomes, and their children become much richer yet. A typical migrant from rural Mexico to the United States will move from earning $2 a day to $10 an hour. Over time the children of immigrants approach the national average, depending how long they have been in the country and how concerned they are to assimilate. Of course, the richer the wealthier country, the more new immigrants will benefit.

Immigrants also send remittances back home. Total remittances around the world are now about $80 billion a year, about twice the amount of the formal category of foreign aid. Remittances, however, bypass governments and therefore do not encounter comparable problems with waste or corruption. Remittances are now ten times the amount of net private capital flows, after adjusting for profit repatriation and interest payments. To cite one example, Mexicans working in the United States send back home $20 billion every year, circa 2003. This sum is twice the value of Mexico's agricultural exports, and over a third more than tourist revenue. Many Latins have used U.S. remittances to start new businesses or revitalize their communities through infrastructure investments.[3]

Furthermore many migrants return to their home countries, bringing skills and liberal democratic ideas. Software repatriates have helped build India's competitiveness in high-tech industries. Thousands of Asian students have obtained science or engineering degrees from American universities, thereafter returning home to start new businesses. If a country is willing to offer some scope for entrepreneurship, it need not fear a "brain drain." Instead foreign contacts, training, and periods of residency will help promote domestic development.

---

[3] All the figures are from the November/December 2004 issue of Foreign Policy.

The global poor also benefit from new medicines, new global technologies, and research and development efforts. E. Helpman (2004, p.84) summarizes: "the main finding -- that R&D capital stocks of trade partners have a noticeable impact on a country's total factor productivity -- appears to be robust... [consider] a coordinated permanent expansion of R&D investment by 1/2 of GDP in each of twenty-one industrial countries. The U.S. output grows by 15 percent, while Canada's and Italy's output expands by more than 25 percent. On average the output of all the industrial countries rises by 17.5 percent. And importantly, the output of all the less-developed countries rises by 10.6 percent on average. That is, the less-developed countries experience substantial gains from R&D expansion in the industrial countries…"

Cross-regional trade has brought wealth to the poor throughout history. The Greek city-states and the Roman Empire benefited greatly from maritime trade across the Mediterranean; those regions in turn spread growth-enhancing institutions around Europe, Northern Africa, and the Middle East. The commercial revolution of the late Middle Ages and Renaissance reopened many of the trade routes of antiquity; more generally, the European technological breakthrough relied on trade with China and the Islamic world. Trade also played a central role in driving the growth eras of nineteenth century Europe and post World War II Asia. Today most poor countries seek greater access to wealthier Western and Asian markets.[4]

Most generally, look back at how we climbed out of the poverty of the year 1000 A.D. or 5000 B.C. Analogous reasoning applies today. If we are sufficiently forward-looking, sustainable economic growth will be our priority. We will favor growth even if we focus only on special egalitarian concerns about the very poor.

[4] Statistical work on trade, investment, and growth cannot always sort out which variables are endogenous. Nonetheless the available evidence shows definite correlations between openness to trade and growth; see Helpman (2004, p.70-1) for a survey of the literature.

World history offers precedents for the idea of a "great transformation," leading to enormous increases in the quality and quantity of human lives. Our ancestors did not foresee the evolution of humans, the agricultural revolution, the "urban revolution" (Sumeria and Mesopotamia, circa 4000 B.C.), or the Industrial Revolution. Each development, over time, drastically changed the human condition, and eventually very much for the better. The history of economic growth, to some extent, is the history of working out the consequences of such unforeseen transformations (Hanson 2002). It is unlikely that we have seen the last of such revolutions, at least provided that civilization manages to stay afloat.

The importance of the growth rate increases, the further into the future we look. If a country grows at two percent, as opposed to growing at one percent, the difference in welfare in a single year is relatively small. But over time the difference becomes very large. For instance, had America grown one percentage point less per year, between 1870 and 1990, the America of 1990 would be no richer than the Mexico of 1990.[5] At a growth rate of five percent per annum, it takes just over eighty years for a country to move from a per capita income of $500 to a per capita income of $25,000, defining both in terms of constant real dollars. At a growth rate of one percent, such an improvement takes 393 years.

VISUALS HERE

Robert Lucas (1988) put the point succinctly: “the consequences for human welfare involved in questions like these are staggering: once one starts to think about [exponential growth], it is hard to think about anything else.”

Although I am focusing largely on utilitarian values, the benefits of growth are by no means restricted to utility. For instance wealthier societies bring much greater access to the arts and education (Cowen 1998). Economic growth also minimizes "the tyranny of place." Individuals suffer a lack of freedom when they have little or no chance to escape

[5] Cowen (2002b).

the circumstances of their birth. Perhaps they are born poor, into the wrong social class, into a community with little tradition of formal education, or far removed from urban culture. Today, because of wealth, more individuals escape these shackles than ever before. We are more mobile, more able to shape our selves, more able to choose our friends, and more able to weave together different cultural traditions, when constructing our personal narratives.

These pluralistic considerations all point toward the same conclusion. The more rapidly growing economy will, at some point, bring much higher levels of human well-being on an ongoing basis.

The Principle of Growth:

"We should make political choices so as to maximize the rate of sustainable economic growth."

The Principle of Growth would return economics back to its roots in Adam Smith. Smith held a straightforward, common-sense approach to political economy. He understood that the benefits of cumulative growth were enormous and should be the focus of economic investigation. It is no surprise that his economics treatise was entitled The Wealth of Nations: An Inquiry into the Nature and Causes. In turn, his earlier Theory of Moral Sentiments argued that discounting the future is not always rational. Taking the two books together, Smith produced a systematic and advanced approach to welfare economics; he is the most important source for these arguments. While Smith remains a famous economist, The Principle of Growth does not dominate formal economic thinking today as it should (see further below).

If we so wish, we can attach "rights constraints" to this recommendation of growth maximization. That is, we might reject growth-maximizing policies to the extent they grossly violate human rights. After all, the introduction already opened the door for a pluralistic rather than narrowly utilitarian approach. Such rights fit into a growth

maximization framework at least as easily as into cost-benefit analysis. A belief in human rights therefore does not provide an argument against (possibly constrained) growth maximization as a normative standard for policy choices.

We can classify non-growth-correlated values into two categories. The first category consists of those values we are willing to trade off against higher growth. To make this more explicit, there is some level of accumulated growth – albeit possibly a very high level -- that will trump this value in importance.

Enough growth, for a long enough period of time, must at some margin defeat the relevance of that value. So for long enough time horizons, the other value becomes irrelevant for our decision-making. We are left with the initial recommendation of growth-maximization. To make this more concrete, we should not accept a lower growth rate to obtain more toothpicks today. Over time the higher-growth economy is likely to deliver a higher quantity of either toothpicks or toothpick-equivalent-value in any case.

To the extent our time horizon shortens, we might need a high growth rate in absolute terms to offset the potential loss of non-growth-correlated values. Assume for instance that the relevant time horizon is five years. We might be willing to give up non-growth-correlated values to achieve a growth rate of seven or eight percent. We may be more reluctant to trade off non-growth-correlated values to boost the growth rate from two to three percent. Yet a growth rate of seven percent might be impossible to achieve. Should we give up non-growth-correlated values for a mere five years of three percent growth?

Of course even a very small growth increment becomes enormous after enough time has passed. If we find seven percent growth a compelling reason to disregard some non-growth-correlated value, we will find a smaller boost in the growth rate a compelling reason as well, but only if our time horizon is sufficiently long.

If you are comparing a one percentage point boost to the growth rate, and starting at parity, you need a time horizon of 110.4 years to establish this 3:1 ratio. If you are

comparing a two percentage point boost in the growth rate you need a time horizon of 55.5 years. For a more ambitious boost of five percentage points, the time horizon must stretch for only 22.5 years. We can see that The Principle of Growth does not require time horizons of millions or billions of years long, much less infinities. These figures make The Principle of Growth, and The Principle of Modified Growth, more plausible.

Keep in mind that the indicated number of years expresses when a 3:1 ratio will be reached. Over time, if the higher growth continues, the 3:1 ratio will be exceeded on an ongoing basis. For growth boosts of one, two, and five percentage points net, we would need 161, 81, and 33 years respectively to reach a quintupling of real income. So the 3:1 requirement is only a temporary milestone on the path toward even greater discrepancies of wealth and (likely) welfare. Asking a policy to deliver a 3:1 wealth ratio is thus quite a stringent rule of thumb for ascertaining the welfare-dominance of one outcome over another.

This analysis, of course, applies only when we are willing to make the trade-offs. The second category consists of those absolute values we are not willing to trade off against growth at all; we can refer to this latter category as absolute rights. For this latter category, by definition, no amount of accumulated future growth will cause us to override these values. We can think of these values as absolute human rights. For instance, it might be argued that we should not torture one thousand innocent people, even if the ongoing rate of growth will be higher for a very long period of time. This argument, by the nature of its construction, is robust to the duration of higher growth. Along these lines, Robert Nozick (1974) has written of rights as "side constraints." These side constraints need not overlap with Nozick's very particular libertarian vision, but they will satisfy his broader notion of rights as restrictions on our choice set.

I will not assess which values (if any) might fall into this second camp, or what absolute human rights might look like. But we do end up with a surprising conclusion. If the time horizon is sufficiently long, the only non-growth values that will matter are the absolute side constraints, or the non-violable human rights.

The resulting moral theory has a clear substantive structure. There is a consequentialist stricture to maximize the rate of sustainable economic growth. Yes we are still pluralists, but if we face a sufficiently long time horizon we can fill in some of the content of that pluralism. Many non-growth-correlated plural values will fall away in practical importance, when compared to the benefits of growth. With a long enough time horizon, those other values will not influence concrete policy decisions. The proper constraints on growth are absolute human rights which provide strictly binding constraints on policy and behavior.

We thus can state a modified Principle of Growth:

The Modified Principle of Growth:

"We should push for sustainable economic growth, but not at the expense of inviolable human rights."

What does growth mean for economic policy?

The more exact implications of growth will depend on empirical investigation. A large and growing empirical literature studies growth using multivariate statistical analysis, applied to pooled data across countries and time. This research examines the correlations between observed rates of economic growth and economic policies or demographic facts. The goal is to start with as few theoretical preconceptions as possible, and see where the data lead us. Typically growth rates are on the left hand side of the regression (to be explained) and various policy instruments and endowment measures are on the right hand side as the explanatory variables. Causality is difficult to infer from the resulting regressions and the t statistics may be hard to interpret, given the difficulty of knowing when the model is properly specified. Nonetheless the gross correlations can help

identify general patterns. Growth hypotheses that "go against" these regressions have some explaining to do.[6]

Overall the empirical results suggest that a stable market order, private property, and the rule of law are strongly correlated with economic growth. Infrastructure-oriented government spending and years of education are positively correlated with growth as well. Non-infrastructure government spending, high and volatile inflation, and regulatory interventions all are negatively correlated with growth. A good distribution of income may benefit growth. For lesser developed nations, colonial origins and the quality of current legal institutions help predict growth. A quick and dirty summary of these results suggests a leading role for capitalistic market institutions, as have driven the growth of the West and parts of Asia. Governments should focus on providing growth-maximizing public goods.[7]

I will not, however, examine these empirical debates. The recommended innovation here is broad and conceptual, rather than empirical or numerically specific. Given our best understanding of growth, what should we do with that knowledge? And what does that knowledge imply for central questions of normative economics? Our methods for estimating the growth-enhancing effects of policies, while they can always use further improvement, do not require significant reformation. Rather we need to take those results more seriously and give them a more central place in our normative toolbox.

Does more wealth really bring greater well-being?

A growing body of literature suggests that additional riches do not make citizens in wealthy countries any happier, at least not above a certain level. Using information taken from questionnaires, once a country has a per capita income of roughly $10,000 a year or

[6] On the robustness of cross-sectional growth regressions in light of critiques, see Sala-i-Martin (1997).
[7] Seminal investigations include Grier and Tullock (1989), Barro (1991), Easterly and Rebelo (1993) and Mahoney (2001). For a survey of time-series tests of similar growth propositions, see Greiner, Semmler, and Gong (2005).

more, the aggregate income-happiness link appears weak. Helliwell (2002, p.28) argues that the curve flattens out at about half of current American per capita income, or roughly the standard of living in contemporary Greece. These results might lead us to wonder whether economic growth is so important after all.[8]

Despite this body of evidence, I nonetheless wish to treat wealth and happiness as comoving in the broad sense. The questionnaire evidence, however intriguing in some regards, does not overturn the basic conclusions of this chapter. Economic growth remains our primary institutional means of improving human well-being.

At most the happiness literature shows that many more changes are irrelevant than we had previously thought. This result would not, however, eliminate the major benefits of economic growth, as experienced over longer periods of time. It might turn out that (if we believe the happiness literature) many "small" changes are irrelevant or nearly irrelevant for happiness. Yet sufficiently large changes still can boost or harm our welfare by significant amounts. If the small changes do not much matter, that is all the more reason to focus on the large changes and thus reason to elevate the importance of infra-marginal welfare economics.

To give a few examples of how large changes matter, most life catastrophes create significant misery. Very sick individuals have less autonomy, experience more pain, and face the stress of dealing with their condition. The death of a child or close family member has a strongly detrimental affect on happiness for most individuals, often persisting for many years. Torture, extreme stress, rape and severe physical pain also produce depression, trauma, and persisting unhappiness. Individuals who have been through wars, revolutions, and collapses of civil order typically experience recurring

[8] See Argyle (1999), Oswald (1997), and Myers (2000). Wealthy countries, when they become wealthier over time, do not become happier in the aggregate. In some cases (e.g., the United States 1946-91) greater wealth is correlated with lower levels of self-reported happiness; see Dieter (1984), Blanchflower and Oswald (2000), Diener and Oishi (2000), Myers (2000), Kenny (1999), Lane (1998), Frey and Stutzer (2000,) and Easterlin (1995). On the United States, see (Frey and Stutzer 2002a, pp.76-7).

flashbacks, nightmares, irritability, depression, alcoholism, troubled relationships, and inability to concentrate. Well-functioning capitalist democracies tend to minimize these problems.[9]

It is sometimes questioned whether even extreme catastrophes make people less happy. Individuals who experience severe disabilities or physical handicaps do to some extent adjust their expectations. Often these victims compare themselves to individuals who are even worse off than they are, or they lower their aspirations in life. The loss in happiness is not as great as a naïve perspective might expect. Nonetheless victims of catastrophe still report lower levels of happiness than do comparable healthy individuals. The happiness difference is most distinct, the greater the extent of the disability. At the very least, a significant percentage of victims experience an ongoing "core of distress" for many years.[10]

People cope least successfully when the catastrophe or malady is ongoing and involves an ongoing deterioration of condition. Most of the counterintuitive results come when the bad event has a “once and for all” nature, such as a one-time physical handicap. In these cases many people recover their initial level of self-reported happiness. But individuals remain subjectively badly off when they suffer from progressive or degenerative problems.[11] So to the extent that a poorer society brings an ongoing worsening of conditions for many individuals, the associated human suffering will be greater. Once again, we are led back to significant benefits from ongoing economic growth.

---

[9] On the link between catastrophes and unhappiness, see Dyregrov (1990), Lehman et.al. (1987), Sanders (1980), Weiss (1987), Frederick and Loewenstein (1999), Lehman, Wortman, and Williams (1987), Archer (2001), and Wortman et.al. (1992). For evidence on the difficulty of recovering from rape, see Meyer and Taylor (1986) and Wirtz and Harrell (1987).

[10] On coping, see Brickman, Coates, and Janoff-Bulman (1978), Bulman and Wortman (1977), Kessler, Price, and Wortman (1985), Meyer and Taylor (1987), and Wortman and Silver (1987). On the "core of distress" idea see Frey and Stutzer (2002a, p.56), Wirtz and Harrell (1987), and Stroebe et.al. (2001).

[11] See (Frederick and Loewenstein 1999).

The observed flat-lining of the happiness-wealth relationship may in part reflect framing. The literature usually focuses on aspiration or treadmill effects. Under this view, you get more but you start expecting more as well, or aspiring to more. The greater wealth therefore translates into less happiness than might have been expected. But this is not the only adjustment occasioned by growing wealth. The wealthy also recalibrate how they should respond to questions about their happiness. If happiness itself is subject to framing effects, surely talk about happiness is subject to framing effects as well. The wealthy develop higher standards for reporting when they are "happy" or "very happy." If you are a millionaire living next door to a billionaire, you might be less likely to report that you are ecstatically well-off. This does not mean that you spend your entire time envying the billionaire or suffering because of your lower relative status.

So let us assume that both framing effects – concerning happiness and talk about happiness – operate at the same time. This will imply that even a constant measured level of reported happiness implies growing real happiness over time. Life improvements do usually make us happier, while both our expectations and our reporting standards adjust upwards. This is the most likely interpretation of the aggregate data. Furthermore it is supported by direct observation. Most individuals strive to earn higher incomes, even after they have experienced the strength of "aspiration" and "treadmill" effects.

Note that within a country wealthier people report unambiguously higher levels of happiness, on average, than do poorer people (Dieter 1984). This result has not been challenged seriously. Now to some extent this result may reflect a zero-sum relative status effect. The wealthier people feel better, but their gains make the poor feel worse off. Nonetheless it is unlikely that the entire effect boils down to a zero-sum game. Wealthier lives are easier and happier in absolute terms in numerous ways, as discussed above. I do not envy your kidney operation or your pacemaker very much if at all, even if I may be less able to afford those same benefits in my old age.

Furthermore, some of the apparent "zero-sum" element will be a framing effect for "talk about happiness," rather than happiness itself. If I buy a Mercedes, my polled neighbor may express greater dissatisfaction with his Volkswagen. That same neighbor, if he had a Lada in Moscow, circa 1978, might express a very high level of satisfaction on a questionnaire. Nonetheless in absolute terms he still likely prefers having the Volkswagen in contemporary America. So my neighbor may envy my new car, but when I buy a Mercedes the gainer's gain still may exceed the loser's loss.

Note that if the world as a whole became much wealthier, the "point of flattening" would shift out to a higher income level over time. In a hypothetical future, people without access to limb regeneration and daily supersonic transport would feel deprived and thus less happy, relative to their neighbors. The standard of living found in contemporary Greece – even if comparable in happiness to the current U.S. -- would not match up to what people will likely enjoy in this wealthier society of the future. Or go back to the past. The happiness-wealth curve curve may have had a flat range in the Stone Age, but the entire range of that curve would not make for a very satisfying existence today. So a flat range of the curve, at any point in time, does not break the happiness-wealth link. Economic growth still can shift the curve up over time.

The happiness literature also takes a limited view of what well-being consists of. The contemporary empirical literature on happiness starts with the operational definition of whether an honest, self-aware person would report himself or herself as being happy, if so asked. Even if this accurately captures one notion of happiness, it is not the only relevant variable concerning well-being.

For instance a wealthier economy probably gives us more "fleeting" happiness experiences, or at least greater chances to trade-off long and short-term sources of happiness. Recent research (Kahneman, et.al. 2004) looks at the allocation of time during the day and classifies events according to how much (temporary) happiness they produce. It turns out that intimate relations, time spent with friends, and television, all appear to make people happier in this sense. Working and commuting make people less

happy. A wealthier economy will offer greater options for structuring these choices, again noting that there may be trade-offs between long- and short-term happiness. Wealthier economies, on average, are associated with higher levels of leisure time, although they accommodate workaholic preferences as well.

Often context effects matter for temporary happiness. An individual will admit to being happier if he has recently found a dime, or if his soccer team won rather than lost (Schwarz and Strack 1999). These sources of happiness will likely be systematically more potent in the wealthier society. A diverse commercial economy offers more sources of temporary stimulations and more short-term turns of good fortune. This means more new gadgets, more fun videos, and more serendipitous encounters with new people. Of course a new gadget may make you happy for only a limited period of time. But if the stream of new gadgets is steady, your boost in happiness can be quite real.

Most generally, we must ask which institutional structures give people the greatest opportunities to structure their lives to achieve their preferred forms of happiness or well-being. Some persons may seek temporary stimulations, others may want to feel fulfilled at the end of their lives, and others may seek to maximize the quality of their modal day. Some will seek happiness through out-competing their peers for status, while others will look inward. Again, greater opportunities and freedoms will likely favor the wealthier society in these regards. Well-being is not a single variable to be maximized; rather individuals prefer to structure the kinds of well-being or happiness they can achieve.

Finally, even if we accept the "flat-line" empirical result as valid, the questions are posed to individuals in normal life circumstances. The answers will not pick up the ability of wealthier economies to postpone or mitigate extreme tragedies, whether in the wealthier or poorer parts of our world. For instance the happiness measures, by their nature, do not pick up the benefits of greater life expectancy. The dead and incapacitated cannot complain about their situation, at least not in questionnaire form. If an immigrant, or a child of immigrants, fills out the form, there is no comparison with a pre-immigration state of affairs. By its very nature, happiness research draws upon a fixed pool of people

in relatively normal circumstances. This will limit its ability to measure some of the largest welfare changes brought by economic growth. Happiness research, whatever its positive uses, is poorly suited to underpin a welfare economics of tragedy.

Aggregation issues

The benefits of growth help us address aggregation problems more generally. Growth maximization, of course, will not make everyone better off. At the very least, some individuals living today would be better off with higher levels of consumption, and lower rates of future growth. And even if growth makes many people better off in the longer run, it is unlikely to make each and every person better off.

Moving forward on this aggregation problem requires some version of interpersonal welfare comparisons. That is, at some level of the analysis, through some method or another, we must assert that the benefits to one group of people outweigh the losses to another. Cost-benefit analysis, of course, invokes a "potential compensation" principle toward this end. If the winners could in principle compensate the losers, as measured in material wealth, cost-benefit analysis recommends the policy, at least subject to distributional caveats. This is not commonly considered an interpersonal utility comparison, but in practice it functions as one. In most real world cases, the compensation is never paid or even seriously contemplated. So we are judging one set of gains as socially "worth more" than a set of losses on the other side of the scale.

I do not wish to reject the potential compensation principle out of hand. But the principle does not command the loyalty of many philosophers or even of many specialists in normative welfare economics. It uses only a paucity of information about individual well-being, it does not build in distributional concerns upfront, and the principle does not guarantee transitive rankings.[12] Most generally, if some people gain and others lose, why

[12] See Chipman (1974) on some of the technical problems. Ronald Dworkin (1980) presents the most effective normative philosophical critique of the potential compensation principle.

should measured total wealth be the relevant tiebreaker? Should not other values stand a chance of serving as tiebreakers? Must we not consider a broader bundle of pluralist values?

Nonetheless the potential compensation principle contains a significant core of truth. A policy that brings overwhelming costs but few offsetting benefits would have a hard time overcoming the relevant hurdle. Therefore the potential compensation principle is most plausible when the surplus of benefits far exceeds the measured costs. And we therein find a means of resurrecting the potential compensation principle in restricted form. If the well-being gains to the future are significant and ongoing, those gains should far outweigh one-time costs to the present, as expressed by the Overtaking Criterion.[13]

Our concern for the future expands the number of cases where an overwhelming preponderance of benefits lie in the same direction. As argued above, a sustainable increase in economic growth will boost many plural values in the medium and long runs. To be sure, some people will be worse off, and some values, in the short to medium run, will not be favored. In these regards aggregation problems do not disappear. Nonetheless the competing options do not generally offer a deadlock of roughly equivalent values and interests on each side of the scale. The higher growth alternative will, at some point in the future, offer a clear and ongoing preponderance of plural values in its favor. These plural values will include the distribution of large benefits to very poor individuals. This predominance of values stands on one side of the scale whether we consider ordinal or cardinal measures of welfare, or some weighted average of the two. In other words, if the relevant time horizon is long enough, an overtaking criterion will kick in.

This approach tries to represent a minimum set of distributional commitments. Many recommendations of traditional cost-benefit analysis are vulnerable to the distinction

[13] Furthermore the grossness of the real income comparisons should significantly limit the scope for intransitivity and Scitovsky double-switching problems, as outlined by Chipman (1974).

between ordinal and cardinal measures of welfare. Applied cost-benefit analysis, in its most practical form, recommends that we pursue wealth maximization. Yet for small changes the question remains to what extent wealth translates into welfare. According to numerous philosophers, a dollar for a rich man may bring less "social value" than a dollar to a poor man. Many cost-benefit analysts attempt to address this problem by applying "distributional weights" to the dollar values created. I have no particular objections to this procedure, but I would not wish to rest critical policy judgments on it. The distributional weights require that we can fine-tune the scales with some accuracy. The (Modified) Principle of Growth, in contrast, looks for a long run where one set of values decisively swamps another in magnitude. Insofar as we can find such improvements, their value will be relatively robust to particular distributional theories and commitments.

This approach to the aggregation problem coincides with common-sense morality. Not everyone can be happy all of the time, but we nonetheless should choose an option that makes a strong preponderance of people much better off. We will never resolve aggregation problems by producing some new algorithm or voting rule that magically resolves all conflicts. If John and Sally favor different policy options, there is always some irreducible clash of interests. We instead should look for decision rules that will put a preponderance of values on one side of the decision scale, and then select that option. We find such rules by considering a broader class of affected individuals, namely the more distant future. The longer our time horizon, the more likely we can find a preponderance of values pointing in a particular direction, in this case toward growth-enhancing policies.

Choosing the pro-growth alternative addresses intra-self aggregation problems as well. For instance individual preferences do not always reflect individual interests. Observed preferences often appear to be irrational, transitive, spiteful, or otherwise morally dubious or at least questionable. Given these facts, it is often asked why the concept of preference satisfaction should stand at the center of an economic approach to welfare. Arguably satisfying preferences does not always make people happier or make the world a better place.

Focusing on the macroeconomics of growth sidesteps these dilemmas. We might doubt that the marginal fast food cheeseburger is worth $4.89 for me, all things considered. Perhaps the offer is manipulating my evolutionarily-programmed desire for more fat, to the detriment of my health. It is more difficult to doubt that living in a much wealthier society is good for me, all things considered. When we consider the macroeconomics of growth, the relevant comparisons become quite gross after the passage of enough time. A given individual is likely better off living an extra twenty years, receiving anesthesia at the dentist, enjoying plentiful foodstuffs, more years of education, and not losing any children to premature illness. Similarly, people one hundred years from now will be much better off if the rate of sustainable growth stays high. At some point these cumulated benefits will be sufficiently gross and sufficiently obvious to be robust to particular instances of irrational or misguided preferences.

We thus see that Arrow's Impossibility Theorem does not present an insuperable obstacle to recommending the pursuit of growth. First, it is well-known that the use of interpersonal utility information defuses the Arrow problem (Sen 1984). Since I am willing to make interpersonal utility comparisons across much richer and poorer societies, we can generate rank orderings without resorting to postulates of dictatorship. Second, growth maximization does not fit Arrow's definition of a social welfare function. One of Arrow's assumptions, the independence of irrelevant alternatives, implies that we can consult only ordinal preference information in constructing a ranking across social outcomes. The Principle of Growth goes beyond mere preferences by looking also to information about wealth and constraints. That is, if we specify "growth maximization," we also must lay out some trading possibilities. We are starting with the broad outlines of the status quo and evaluating policies as they would interact with current market institutions. This will include information about marginal rates of substitution, wealth endowments, and the scope of transactions costs, among other variables. We can think of the policy recommendation as a social welfare function, but as drawing upon more information than ordinal preference rankings alone. In this sense most real world institutions, including the market, or the operation of political pressure groups, fail to

meet Arrow's definition of a social welfare function. They include or rely upon too many sources of information, thus showing that the Arrow problem bites for only a limited class of normative comparisons.

Growth maximization vs. cost-benefit analysis

The importance of the growth rate is hardly news, but rarely is the point made fully explicit for the theory of economic policy. Instead of performing a cost-benefit analysis, the recommended alternative procedure is to ask whether a given policy is likely to increase or decrease the rate of economic growth.

Over a sufficiently long time horizon, the growth effects of the policy will overwhelm the static allocation effects as measured by cost-benefit analysis. Of course the postulate of deep concern for the distant future is central to this judgment. As discussed above, the growth approach sidesteps aggregation problems in a way that standard cost-benefit analysis does not. We can identify cases where a strong preponderance of plural values lie in one direction. Our policy recommendations therefore become more certain, the net benefits of good policies are more decisive, and our recommendations are more robust to the possibility of static measurement errors.

Under the traditional cost-benefit view, we should not in general maximize the rate of economic growth. At the growth maximum there will always exist some increase in current consumption that will lower growth yet improve current welfare. In this view future periods will be wealthier, and thus maximizing the utility path of growth will imply pulling some additional resources toward the present. Marginal utilities are equalized over time, taking into account a declining marginal utility of wealth. The recommended rate of growth will stand below the maximum possible rate of growth, due to the growing wealth of the world over time (see equation (5) from the last chapter).

In contrast, the Overtaking Criterion (Equation (3)) implies a more future-oriented perspective. The more rapidly growing economy will, after some point, have

systematically higher levels of cardinal well-being; we should not sacrifice these future benefits for some current increment of consumption.

To view this point in terms of comparative statics, assume we are not at the growth-maximizing point and that the current generation cuts back on consumption and increases saving by some discrete amount. More will be saved and invested and the growth rate will rise. There will be a loss today but a sustained increase in well-being and other plural values over the remainder of the time horizon. If we accept a deep concern for the distant future, we will find this trade-off to be attractive. Furthermore such a trade-off is at our disposal unless we already stand at the growth maximum.[14]

Of course we can perform the same experiment in reverse. Assume we start at the growth-maximizing point. Why would we wish to leave a growth maximum with an increase in current consumption? We would reap a discrete benefit now, but future generations will be worse off for a long time to come. At first the differences in welfare will be relatively small (at least assuming that the consumption increase is small). But at some sufficiently distant point in the future, the difference in welfare will be very large. For a long enough time horizon, the final cost will be a series of generations who are much poorer than they otherwise might have been. Some set of generations will have lives with many more tragedies, relative to what might have been. The difference in welfare would be analogous to putting several or more generations of the modern world back into eighteenth century conditions.

If the time horizon is extremely short, the benefits of continued higher growth will be choked off and will tend to be small in nature. Even if we hold a deep concern for the distant future, perhaps there is no distant future for humanity to care about. To present this point in its starkest form, imagine that the world were set to end tomorrow. There

[14] In some economic models, having more capital does not always increase the growth rate. In some so-called "golden rule" models (Sidrauski 1967) the costs of maintaining extra capital can exceed its rate of return; in these cases the proper recommendation is to decrease the rate of investment. It is not in general argued, however, that these models are policy-relevant.

would be little point in maximizing the growth rate, and arguably we should just throw a party and consume what we can. Even if we could boost growth in the interim hours, the payoff would be small and not very durable. The case for growth maximization therefore is strongest, the longer the time horizon we are willing to consider.

One generalizing approach might classify policies in terms of the size and time distribution of their benefits. Policies with small and temporary benefits would call for a relative balance closer toward cost-benefit analysis. Other policies – with larger benefits and more durable effects on long-term growth -- would give greater weight to infra-marginal methods. Rather than competing with marginal welfare economics, infra-marginal welfare economics could coexist with it. For the time being, I will consider policies as bringing large and durable changes rather than small and temporary changes. In practical terms, we know that the correct policy will involve greater emphasis on growth maximization than is suggested by current methods of cost-benefit analysis. In theoretical terms, the relative scope of marginal and infra-marginal welfare economics remains open, noting that I will return to this question in the next chapter.

Sustainability

Maximizing the long run rate of growth refers to gross domestic product as properly understood, and not as currently measured by governments. "True GDP," if I may use that term, accounts for leisure time, household production, and environmental amenities. Current GDP statistics have a bias towards what can be measured, rather than what contributes to human welfare. For this reason, "maximizing the rate of growth" does not mean that everyone should work the maximum number of hours in a day. An eighteen-hour workday might maximize measured GDP but certainly would not maximize true GDP over time, once we take the value of leisure into account, not to mention the possibility of labor burnout. (That being said, leisure and physical capital need not be treated symmetrically; a significant portion of capital typically is bequeathed to the next generation, but leisure serves as consumption only.)

Furthermore we wish to maintain higher growth over time, and not just for a single year. Maximizing the sustainable rate of economic growth does not imply pursuing immediate growth at the expense of all other values. Policies that seek growth at breakneck speed are frequently unstable in both economic and political terms. The Shah of Iran, for instance, tried to bring his country into the modern world very rapidly. Growth rates were high for a while but in the longer run could not be maintained. Since the revolution, Iran has done poorly, often with negative rates of growth. The Shah's forced modernization did not in fact maximize economic growth, and a more cautious set of policies likely would have been better.[15]

To give another example, many scientists believe that global warming will increase the number of virulent and persistent storms. (This is illustrative, I am not seeking to debate the relevant facts.) These storms may come only in the long run, but a greater concern for the future means that we must pay greater heed to these consequences. More generally, many environmental problems hurt the prospects for long-run growth, especially once we include suitable measurements for environmental amenities into true GDP.

Sustainability also focuses our attention on the prerequisites for a durable civilization. The Western world does not appear to be on the verge of collapse, so policy analysts and economists often neglect threats to the basic stability of our civilization. Nonetheless we should not take the continuing existence of public order for granted.[16]

Economic growth is the exception not the rule, and civilizations show historical fragility. Michael Shermer (2002) has compiled an informal database on civilizational survival. He catalogued sixty civilizations, including Sumeria, Mesopotamia, Babylonia, the eight dynasties of Egypt, six civilizations of Greece, the Roman Republic and Empire, nine dynasties and two republics of China, four periods in Africa, three in India, two in Japan,

[15] Chapter one already has discussed the Overtaking Criterion. For a formal look at the concept of sustainability, see Heal (1998, chapter one).
[16] Rees (2003) and especially Posner (2004) are two notable exceptions to the tendency to dismiss talk of civilizational collapse.

six in Central and South America, and six in modern Europe and America. He finds that the average civilization endured 402.6 years. He also finds that decline comes more rapidly over time. Since the fall of Rome average duration of a civilization has been only 304.5 years.[17]

While the numerical features of these estimates depend on how we define the concept of civilization, the more general point of fragility stands.[18] Human beings can experience significant and ongoing losses of their prosperity and freedom.

I do not intend to analyze which causes are most likely to threaten civilization (see Posner 2004). Should we invest more resources in spotting and preventing comet and asteroid impacts? Or are infectious diseases the real danger? How about the slow erosion of liberty and initiative through high tax rates, or perhaps moral collapse as

[17] S.E. Finer (1997, pp.30-34) provides an alternative look at civilizational survival, under the heading "Total Life-Spans." He defines a civilization in grosser terms than does Shermer, so for him ancient Egypt is one civilization, not eight. Civilizations have correspondingly longer lifespans. Some of the longer-lived civilizations are 2,820 years (Egypt), 2,133 years (China), and 1,962 (the Byzantine empire). The Venetian Republic lasted 1,112 years. The shorter examples include the Achaemenian Persian Empire (220 years) and the Sassanian Persian Empire (427 years). Finer also develops the more finely grained category of civilizational breakdowns, which occur more frequently. A breakdown is the "disintegration of a previously united state" (p.32), in contrast to the more severe examples of total civilizational collapse. Egyptian breakdowns come after varying periods of 675, 184, 206, 215, and 1,238 years. For Chinese breakdowns they come after periods of 400, 500, 442, 360, 326, 69, and 936 years. Assyrian breakdowns come after periods of 157, 82, 38, 143, and 312 years. Arnold Toynbee, in his classic A Study of History, classifies world history into twenty-six civilizations. By his count, sixteen of these civilizations no longer exist. Samuel Huntington's The Clash of Civilizations (1996, pp.44-45), citing Matthew Melko (1969), refers to twelve major civilizations, seven of which have perished.

[18] Samuel Huntington (1996, chapter two) cites a variety of definitions of “civilization,” including the concepts of “settled, urban, and literate” (p.40), and “the broadest cultural entity” (p.43). Fernández-Armesto (2001, pp.16, 20), referring to the work of Kenneth Clark, designates a civilization as a society with the confidence to build for the future. Clark (1969, p.1) noted that while he did not know exactly what civilization was, he could recognize it when he saw it. Matthew Melko (1969, p.113) remarks that "when a civilization is operating effectively, it is likely to grow."

feared by the conservatives? Rather than debate these questions, the relevant point is that all civilizational dangers are now a greater problem, given our deep concern for the distant future. With a five percent rate of discount, it matters relatively little whether the world lasts two hundred or two thousand years. But if we accept the arguments of chapter one it matters a great deal.

Focusing on sustainable growth does not require an obsessive concern with each and every danger to civilization. For instance truly massive investments to preserve the environment would damage seriously the long-run rate of growth and thus would not be recommended. When we look at particular investments, a very low discount rate often tends to favor the environment. When we look at overall patterns of investment, however, we are again led to favor maximizing the sustainable growth rate of true GDP, which leads to reasonable limits on environmental protection.

More generally, a concern with stability does not imply technological stasis. The literature on environmentalism sometimes presents the "precautionary principle" (see Morris 2001). The precautionary principle has been interpreted in differing ways, but in practice it usually means that we should not adopt a new technology, or new policy, unless we can prove it will bring no harm. This emphasis on fragility is useful, but often the principle is misapplied. The precautionary principle involves a status quo bias and requires that change be justified. In reality change is inevitable, and a status quo that sits still has poor survival properties. Unless we take some chances, we will perish or be conquered. Modernity and new technologies give us some chance of navigating successfully and creating a plenitude of rich plural values for many generations to come.

We might wonder whether we would maximize the relevant human values by existing at a modest population and economic level for a very long time, perhaps "living in harmony with nature," so to speak. But poorer societies from the past have collapsed repeatedly through military weakness, eco-catastrophe, and natural disasters. Furthermore primitive warfare appears to have been at least as frequent, bloody, and arbitrary as modern

warfare.[19] Earlier societies were neither idyllic nor peaceful. So returning to the past, or attempting to throttle economic growth, does not guarantee the future prospects of civilization, much less any degree of comfort.

The issue of military defense further boosts the case for modernity. The tyrant will conquer the noble savage. Even if in principle the life of the noble savage were best, no society following this path will, on its own, manage to keep its freedom for very long. So the real alternative is not the noble savage, but rather the life of the savage under a conquering tyrant, which hardly appears attractive. We must accept the idea of advanced technologies and try to manage our risks as best as we can. Given the previous path of human development, someone will have those technologies, whether we like it or not. It is important that the more benevolent societies be both richer and more technologically advanced.[20]

We should instead adopt a more favorable opinion of science and perhaps subsidies to science as well. Science and the experimental method have been central to the growth of the Western world, the Industrial Revolution, and many subsequent innovations. In earlier times, the superior economic performances of China and the Islamic world were correlated with their scientific advances. While growth boosts science, science causes subsequent growth as well. Both the Solow and increasing returns models (see below) assign a decisive role to ideas and new technologies in stimulating economic growth. Note also that science may help us address environmental and other sustainability problems at lower cost than would otherwise be the case.

While many governmental investments in hi-tech have failed, science subsidies need not be counterproductive per se. In addition to the atom bomb, science subsidies of varying

[19] Diamond (2005) focuses on how ecological catastrophes have destroyed or damaged civilizations of the more distant past. On primitive warfare, see Keeley (1996) and LeBlanc (2003).

[20] Trying to maximize the longevity of human existence, at the complete expense of quality of life, also may run afoul of Derek Parfit's "Repugnant Conclusion," as discussed in chapter one.

kinds have hastened the computer, the airplane, the Internet, and penicillin, among other benefits. Governmental prizes for scientific advances were common in the seventeenth and eighteenth centuries, and many famous scientists have received government patronage (Sobel 1996). It is true that past governments often have hindered science, but a science-supporting government is not a contradiction in terms. It is unlikely that laissez-faire will bring us an optimum degree of scientific development. Individuals respond to incentives, and it is within the power of government to increase the returns to doing science. So while not all science subsidies succeed, we should upgrade the importance of achieving effective science subsidies.

Subsidies to pure science, if done appropriately, are likely to have a positive long-run payoff, even if they bring costly mistakes and bureaucratic boondoggles in their wake. Recall that the United States beat totalitarian societies to the atomic bomb by a relatively small number of years, and only because of government involvement. The benefits of this single success probably far outweigh all the other costs we have experienced from science subsidies to date. When it comes to science, government does not in general adopt a longer time horizon than do markets. Often government responds to short-term electoral incentives and special interest pressures for pork, privilege, and political control. Politicians are notorious for running deficits or leaving problems in the hands of their successors. We can, however, identify some areas where political incentives lead government to (sometimes unintentionally) favor a longer time horizon than do markets. Science subsidies are in case in point. Often the subsidy is driven by the desire to delivery rents to key political constituencies, or by the (not always benevolent) desire for greater military power. Despite these mixed motives, and in part because of them, subsidies to science can help humanity invest more in its future.

Once-and-for-all changes vs. growth rate changes

Beneficial policies may fall into one of three categories. First, they may increase some benefit in once-and-for-all fashion. Imagine increasing the power of all light bulbs for one year. Second, they may yield a fixed benefit for an ongoing number of future time

periods. Imagine discovering a light bulb that burns longer. Third, they may lead to a permanent increase in the rate of economic growth. Imagine a new laboratory that speeds the rate at which better light bulbs are discovered.

Gains of the first kind, which do not stretch into time, become in relative terms much less important than gains of the other two kinds (the same can be said for costs). Gains of the third kind are now more important, in relative terms, than gains of the second kind. The temporally distant exponentially increasing gains are not discounted away at such a high rate.

It is a question of special importance which categories best describe the costs and benefits of various policy options. Many policy issues will depend on precisely this question. A given cost might involve an up front, once-and-for-all burden, or instead a systematic decline in the growth rate over time. We should reorient policy research to center on this familiar but undervalued issue.

Counterintuitively, a concern for the distant future sometimes will militate against environmental investments. For instance many of the costs of global warming appear to be "one-time" in nature, such as the costs of relocating coastal and inland settlements. In long-run equilibrium, transition costs aside, it is no worse and arguably better for the world to have a warmer climate (we spend more money on warming space than on cooling it). At the same time stopping or limiting global warming might lower permanently the rate of economic growth. When the rate of intergenerational discount is sufficiently low, maximizing the growth rate tends to take priority over avoiding one-time expenditures and one-time adjustments. Even if those one-time expenditures are large, we will earn back that value over time, due to the logic of investment compounding. So insofar as we view global warming as a serious problem, we must have in mind its long-term growth-reducing effects, such as might occur through a higher rate of storms.

Economic models provide differing accounts of which changes are likely to affect the growth rate. The most prominent approach, the Solow model, postulates a stripped-down economy-wide production function based on constant returns to scale. National output is the result of capital inputs, labor inputs, and technological progress, which renders both capital and labor more effective (see Solow 1956, 1957; Romer 2000 provides a more recent summary). In this model the primary way to increase growth is to induce a higher rate of technological innovation. Along these lines, many empirical tests have shown embodied technological progress to be the driving force behind U.S. economic growth.

In the Solow model the rate of return on capital diminishes as the capital stock increases, and the rate of capital accumulation responds to this rate of return. Given these assumptions, poorer countries should be expected to catch up to richer countries, as they borrow new technologies and increase their capital stocks. Furthermore economies should be able to recover quickly from one-time shocks, such as earthquakes. Although the capital stock has fallen, the rate of return on capital is now higher. Additional savings should make up the gap and restore the economy to its previous growth path.

The Solow model makes a clear distinction between the level of real income and subsequent rates of growth. Many traditional factors – such as a boost in savings and investment – are seen as contributing to "transition growth paths" but not to "steady state growth" in long-run equilibrium. Alternatively, a decrease in wealth lowers the base on which growth occurs, but it has no necessary implications for the succeeding rate of growth. To use a biological metaphor, consider a lobster. If an arm is lopped off another arm grows rapidly to replace it. In the long run the lobster is not much worse off, even if it never quite replaces its original weight. In economic terms the mechanism runs as follows. The decline in the capital stock raises the rate of return on capital, which induces more savings, which tends to restore a higher capital stock. In the long run, an increase in the savings rate makes up, over time, for the "destroyed" resources. The very rapid recovery of some economies after wars or major natural disasters might represent this mechanism in operation. The rate of growth will remain permanently lower only if the negative shock somehow reduces the rate of technological progress. And since

capital has a diminishing marginal product in the model, a higher rate of savings will boost the absolute level of income, and transitional growth, but not the ongoing rate of growth. Growth would proceed from a higher base but not at higher rates in the long run.[21]

In contrast to the Solow model, increasing returns models suggest that growth begets more growth. In this view larger economies should grow more rapidly than smaller economies, and growth patterns should be serially correlated over time. Ideas – and their non-rival nature -- often are cited as the fundamental source of increasing returns. Once an idea has been generated, it can be used many times by many different people at very low marginal cost. Larger markets generate stronger incentives for ideas production and thus face a comparative growth advantage. New ideas will lead to more growth, which in turns encourages more new ideas, and so on.[22]

Increasing returns models are sometimes traced back to Adam Smith. In Smith's implicit model, a larger market size supports a greater division of labor, which in turn makes the economy more productive. In other models greater openness to trade, or a common market area, can drive an increasing returns to scale process. Along these lines, legal and regulatory standardization may help economies grow, as has been one rationale for the European Union.

Under increasing returns models, a one-time negative shock more likely has serious negative effects on the long-run rate of growth. Intuitively, we can think of the increasing returns concept as suggesting that resources multiply themselves at increasingly rapid rates. The larger the economy, the faster it will grow. Rather than

[21] That being said, some later modifications of the Solow model allow for the rates of savings and investment to be correlated with economic growth in a more general manner (see Temple 1999, pp.139-140). Extensions by Uzawa (1965) and Lucas (1988) stress the role of human capital – not just physical capital -- in boosting or maintaining the growth rate.

[22] On increasing returns models, see (Romer 1986, 1990). On the Solow model vs. the increasing returns model, see the 1994 symposium in Journal of Economic Perspectives.

losing the arm of a lobster, we have lost a colony of fertile rabbits. Even if the colony is small at first, it has the potential to become much larger with time. So the increasing returns model implies that we must take great care to avoid each and every negative shock that comes along.

The Solow model suggests a picture of greater resilience. In general the Solow model will give us greater latitude to worry about the present, whereas the increasing returns model imposes a stricter discipline. Many more events will matter greatly for the distant future. To the extent the increasing returns model is true, it is harder to justify the pursuit of non-growth-related values. The Solow model suggests a lower final cost for indulging in these values. When debating one-time costs, it is therefore a central question whether the Solow model, the increasing returns model, or some other approach comes closed to capturing the real world.

That being said, the logic of the increasing returns model will likely carry considerable weight in our final evaluation. In many cases our best answer, given current knowledge, will suggest that a given cost brings some probability of an ongoing growth effect and some probability of a one-and-for-all adjustment cost. In our expected value calculations, this will operate as an expected impact on the long-term rate of economic growth, of course discounted for the uncertainty. Therefore we must face the logic of the increasing returns model, even if it is not our best current forecast of what drives economic growth.[23]

[23] Neo-institutional approaches are less formal than either the Solow or increasing returns models. They point to the importance of property rights, well-functioning institutions, trust, the rule of law, and properly-aligned microeconomic incentives. Nonetheless these views do not typically specify which policy changes cause permanent boosts in the growth rate, as opposed to once-and-for-all changes. To the extent the neo-institutionalist approaches have positive value in explaining economic growth, this suggests they be given a new focus. On neo-institutionalist approaches, see Douglass North's work on American and European economic history (North 1981, North and Thomas 1976) helped found the neo-institutionalist tradition; see also Olson (1984) and Bates et.al. (1998). Acemoglu and Johnson (2004) survey the literature in this area.

Another critical question is which changes will affect long-run growth within the relevant form of the Solow model. For instance will a new government labor regulation have a negative effect on innovation and thus long-run growth?

Under one view, wages will rise and at the margin employers will hire less labor. Arguably the rate of discovery is likely to fall. To provide a simple example, to the extent it is harder or more costly to hire labor, fewer research laboratories will be opened each period or laboratories will be smaller in size. Now in the Solow model, the induced lowering of the growth rate might be small. For instance the rise in labor costs might not have a large current affect on output. But even if the current growth effect is small, the difference in terms of national income will compound over time. Over a long enough time horizon, real income will be much lower, relative to a world with a weaker regulation.

Under another view, this regulation leads to a one-time adjustment cost but no ongoing decline in long-run growth. We might think of the economy as investing in a one-time and complete adjustment to the new change in circumstances. Perhaps a one-time adjustment on the lab floor will make up for the relevant deficiencies. Most likely we face a probabilistic chance that the cost will bring an ongoing change in the growth rate. We are again led to be extremely wary of policies that may cause deviations from the economy's maximum expected long-run growth path.

What about redistribution?

Increasing returns growth models will make us more wary about redistribution than will the Solow growth model. In the Solow growth model, many economic costs are "once-and-for-all," rather than lowering the long-term rate of growth. Arguably some welfare state costs fall into this category; for instance individuals who go on welfare might not be potential contributors to technological progress and thus to economic growth. Under the increasing returns model, in contrast, any deadweight loss makes the economy smaller

and thus limits future rates of growth, with significant implications for the very distant future.

This is yet another way in which traditional political debates should be redrawn. Individuals who believe in increasing returns models should be much more skeptical of welfare states than those who believe in the Solow model.

Most generally, we should redistribute only up to the point which maximizes the rate of sustainable economic growth. This may mean more redistribution than we currently undertake, or perhaps redistribution of a much different kind, namely growth-enhancing redistribution. (It is debatable how much today's government programs in fact redistribute to the poor at all.) It will not, however, suggest that a utilitarian or consequentialist approach is obliged to redistribute most of national income to the very poor.

To cite some pluses from redistribution, a welfare state can give the poor greater access to education and nutrition. These individuals not only enjoy a higher quality of life, but they produce goods and services, they contribute to tax revenues, and they are less likely to end up as a destructive social force. Other growth-enhancing benefits of redistribution are political in nature. Welfare payments sometimes "buy" the loyalties of special interest groups, thereby inducing them to support public order. Some of the poor will be less desperate and will feel less desperate as well. Those groups receive a financial stake in the system and a socially-sanctioned legitimacy for their claims. More generally, welfare systems make many higher income individuals feel good about their state, and increase levels of political support. Many people want to have states whose benevolence they can feel good about. This benevolence contributes to state legitimacy and thus to public order. [24]

[24] See for instance Alesina and Rodrik (1994) and Persson and Tabellini (1994). For a survey of the growing literature on how income distribution can affect growth, see Greiner, Semmler, and Gong (2005, pp.132-133).

These arguments provide good reasons to support some investments in a welfare state. Furthermore they suggest an appropriate nature and scope for such redistributive activities, namely that we target redistribution in growth-enhancing directions.

Beyond some point a sufficiently generous welfare state limits the rate of growth. It withdraws some individuals from the labor force, weakens productive incentives, necessitates higher tax rates, and is usually combined with static, insider-oriented labor market regulations. Furthermore if everyone approaches government looking for a handout, basic mechanisms of governance can break down, leading to rent-seeking, corruption and fiscal bloat. Alternatively, welfare may create urban cultures of dependency and crime, which endanger social order. As noted above, the empirical literature suggests that non-infrastructure government spending is correlated positively with lower growth rates.[25]

More subtly, high levels of welfare make it harder for wealthy countries to afford large numbers of poor immigrants from around the world. Many immigrants increase government revenue in the short run, but many, especially the poorer ones, do not. They require resettlement assistance, emergency medical care, extra police and public works expenditures, or they otherwise tax the resources of the state. The more we spend on domestic welfare, the less we can spend on absorbing immigrants. In public choice terms, a larger welfare state will make society less willing to take in many immigrants. As noted above, immigrants can boost living standards in both their old and new countries. Our true concern is global growth, and value-maximizing immigration provides a significant boost to this variable. So even if a specified set of welfare expenditures brings some growth benefits, alternative investment opportunities may be superior.

[25] See, for instance, Barro (1991). Goodin, Headey, Muffels, and Dirven (1999) argue that a democratic social welfare state does not lower the rate of economic growth, but they use only two data points, the Netherlands and the United States. See also Lindert (2004). He argues that higher welfare spending tends to be packaged with other growth-enhancing policies, such as low taxation on capital income. He does not show that higher spending at Western European levels is itself good for economic growth.

We can thus see obvious limits to the common utilitarian or consequentialist prescription to redistribute a massive share of global wealth. It is true that sending a large chunk of American GDP to Africa would raise African welfare in the short run. But if current total income were divided equally, world per capita income would be about $3500.[26] This average would then fall rapidly, due to incentive effects. Civilization as we know it could not survive and the world's poor would fall into a deeper state of misery than they currently experience. The poor countries no longer would benefit from their interactions with the previously richer countries. So rather than redistributing most wealth, we would reap greater utilitarian benefits by investing it in high-return activities.[27]

A sufficiently long time horizon will favor growth over redistribution even if we are counting only the interests of the very poor in the social welfare function. The benefits of radical redistribution are one-time in nature. We can try to equalize all wealth today, but we would not be able to draw on comparable resources for the next generation. Such a widespread collective redistribution would lead rapidly to negative economic growth. In contrast the benefits of economic growth will compound over time. It is common to scorn the phrase "trickle down economics," but in fact a steady and ongoing flow of benefits is exactly what we are looking to achieve. A flood is better than a trickle, but a lasting trickle is better than eating our cake today and cashing in all of our chips.

Utilitarian slaves?

As noted above, utilitarianism often is seen as an excessively demanding philosophy. Even if our government is not obliged to engage in widespread redistribution, our personal obligations toward the poor may remain strong. For instance the demands of the suffering are so enormous that few able or wealthy individuals would be able to carry out individual life projects. We can imagine, for instance, that every individual is obliged to

[26] See www.geocities.com/combusem/WORLDGDP.HTM.
[27] On utilitarian obligations, Scarre (1996, chapter VIII) offers a good survey. In addition to Williams (1973), see Rand (1967), Scheffler (1982), Wolf (1982), Railton (1984), and Nagel (1986) for other critiques of extreme utilitarian obligations.

work for charity, or to send most of his or her income to the poor in India. Wealthy doctors should spend large parts of their careers in African villages. Many more of us would have to become doctors or nurses. A mother might have to abandon or sell her baby to send food to the babies of others, and so on.

At the individual level, we may well be obliged to help the poor more than we are doing. But we do not run the risk of personal enslavement or massive redistribution of our personal wealth. Most of us should work hard, be creative, be loyal to our civilization, build healthy institutions, save for the future, contribute to an atmosphere of social trust, be critical when necessary, and love our families. We are now obliged to contribute to sustainable economic growth rather than to engage in massive charitable redistribution. Of course this implies we should engage in growth-enhancing charity when possible.

These stipulated individual obligations are not so far from common-sense morality. To be sure, we have not bridged the gap between utilitarian reasoning and common sense morality. Even when utilitarianism and common sense recommend the same courses of behavior, they do so for very different reasons. Utilitarianism tells us we should work and save to serve the purposes of others, in this case future generations. Common sense morality tells us that we should work and save to take care of our families and because we own our lives. These two perspectives remain different. Nonetheless to the extent the practical conclusions converge, we can think of utilitarian and common sense modes of reasoning as two parts of some broader moral picture. We should be pluralists, rather than advocates of utilitarianism or common sense morality per se. So we do not have to bring those two perspectives into complete accord. Instead it may suffice to know that two of the "kits in our toolbox" point in broadly compatible directions. We do not yet have The Best Ethical Theory, but our quest is no longer so torn between two warring accounts of what we should do. By emphasizing our deep concern for the distant future, we come closer to reconciling utilitarianism and common-sense morality.

Common sense morality also suggests that public and private codes should differ in their advice. For instance a mother might be justified in giving preference to her own baby,

rather than tending to the babies of others. At the same time, perhaps the government should behave as some approximation of an impersonal welfare-maximizer, taking into account the interests of all citizens (Goodin 1995). When allocating resources, governments should not favor one particular baby over another. It has remained an open and unanswered question why morality should be split in this fashion. After all, why do moral obligations change, simply because an individual is labeled as acting privately rather than within the context of a public institution? But we now have the tools to defend a bifurcation of this kind. Such a division of responsibilities stands a good chance of maximizing the long-term rate of sustainable economic growth. Proscribed behavior of both private citizens and government then would spring from a common principle of growth maximization, and the resulting principles would be roughly compatible with common sense morality.

We do see circumstances under which a utilitarian should favor large-scale redistribution toward the very poor. Perhaps, for whatever reason, the world was going to end in the near future. Redistribution then would stand a greater chance of being favored in utilitarian terms. The scope for compounding over time would be correspondingly limited and the immediate returns to charity would weigh more heavily in the decision calculus. Alternatively, the real return on investment might be permanently negative or zero. In this case compounding would not operate and we again would see greater reason to redistribute wealth.

Under more normal circumstances, a utilitarian or consequentialist framework still may recommend that some individuals sacrifice significant parts of their lives, or risk such sacrifices, for a greater social good. Martin Luther King brought much good to the world, with respect to both justice and economic growth. Nonetheless such obligations to sacrifice cannot be universal or near-universal. If we all went around sacrificing, there would be no civilization left to advance. As we saw before, we should reject collective sacrificial recommendations that will lower the rate of sustainable economic growth.

In many instances our obligations are collective and we cannot pin down a unique course of individual behavior. In these cases the question "What should I do?" allows for considerable latitude. For instance some subset of a group might have the capacity to act collectively and solve some social problem. But the scope of my individual obligation may be indeterminate. It will not be clear, even on the level of pure theory, who in the group is required to act how much to improve the future. Why should one coalition be obliged to sacrifice and not another? The structure of this problem is common to many questions of morality and individual obligation. If we put aside philosophical conundrums about encountering a drowning man in the wilderness, we see that most problems admit of multiple sacrificial solutions.

So when a sacrifice is called for, the question arises exactly who is obliged to make the sacrifice. A utilitarian standard, in its simplest form, suggests only that the “least cost supplier” should make such a sacrifice. If several individuals face the same cost and can offer the same benefits, it can only be said that one of them should sacrifice. To make this concrete, if several rescuers have equal facility in rescuing a drowning child, there is no fact of the matter, in utilitarian terms, who should jump into the water. In these cases a sacrificer could be specified by lot. Or we could take a game-theoretic route. We could think of morality as prescribing that individuals should play randomized Nash strategies with some probability of selfishness and some probability of sacrifice. The resulting outcome likely will bring some individuals who sacrifice and some who do not. On average this might generate about an appropriate amount of sacrifice.

Whether such a game-theoretic approach portrays every relevant nuance of the problem, we can see a simple truth. In most real world situations, individuals are more likely to sacrifice too little than too much. So we can look to a definite recommendation of another kind. We should strengthen our consciences, and strengthen social norms, to increase the probability that the appropriate individuals would be willing to make a needed sacrifice. Again this recommendation is in line with common-sense morality. We are not compelled to endorse the idea of universal or even near-universal sacrifice, but we can favor exhortation for more sacrifices. Arguably our society does not

encourage enough self-sacrifice for the more distant future and for other important goals. In that case we ought to honor and reward such sacrifices more, to increase their likelihood.

In some cases utilitarian prescriptions will have morally counterintuitive implications, but running counter to the usual fears of enslaved American doctors serving Africa. Namely utilitarianism may support the transfer of resources from the poor to the rich. A talented entrepreneur, for instance, can probably earn a higher rate of return on invested resources than can a disabled great-grandmother. So we will have some reason, when thinking about the future, to redistribute additional resources to the more productive members of society. The implications will be anti-egalitarian at first, but over a sufficiently long time horizon the poor will benefit increasingly from the high rate of economic growth. The results need not be anti-egalitarian if we take the appropriate broader stretch of time, but they still will appear anti-egalitarian by the usual metrics.

I am not suggesting that a good pluralist theory will, all things considered, necessarily endorse systematic redistribution toward the wealthy or the talented. This may be one case where we impose a "rights constraint" on the core recommendation of growth maximization. Nonetheless the example shows how a very low discount rate shifts the burden of proof. Direct, short-term redistribution to today's poor is no longer the default option for an impersonal moral theory that emphasizes individual well-being. Furthermore pure utilitarianism can be anti-egalitarian, at least in the short run, and even in the medium run, in its implications.

Obligations to save and invest?

In addition to obligations to redistribute, we might also consider our obligations to save. John Rawls (1999, p.252) raised the fear that a zero rate of discount will demand unreasonably high levels of savings from the current generation. We can imagine similar objections being leveled against very low rates of discount, or a deep concern with the

distant future. Nonetheless a concern with sustainability may weaken our obligations to considerable degree.

Rawls's concerns are not very specific, but consider a real world example. Arguably a very low rate of discount sometimes may recommend behavior similar to what we saw in postwar Japan. In that case a generation worked extremely hard for the future and saved a good deal. This may be objectionable on Rawlsian grounds (the generation that works hard is poorer than generations to come), but it is not necessarily unacceptable in consequentialist terms. The investment brought an extraordinary return. Furthermore, the Japanese work ethic of that time was sustainable, in part, because it was voluntary rather than coercively enforced. Since each hard worker wanted to provide so much for the future, the investment was compatible with relatively free and thus growth-enhancing institutions. The workers enjoyed the satisfaction that their children would have better lives. So we can look approvingly on the Japanese example, without wishing to coerce all other citizenries to do the same.

To see the potentially positive-sum nature of growth more clearly, consider bequests. Bequests do have a potential zero-sum element (should I spend a dollar or leave it for my heirs?), but those same bequests are precisely not what accounts for a high rate of ongoing growth. Bequests are a one-time wealth transfer, which do not raise the rate of growth permanently or perhaps not even in the short run. (To refer to the example at hand, the postwar Japanese generation did not leave enormous yen-based bequests to its children, nor have commentators cited such bequests as a reason for the Japanese miracle.) So if a generation matures in a society with a high rate of growth, we should not look to bequests as the cause of this good fortune. Instead the older generation has given the greater gift of growth-conducive institutions. So the key transfer was not based on total abstinence from consumption. Instead it came from wise investments, a belief in rules of just conduct, good political institutions, good values, and so on. Growth-enhancing institutions do require hard work, but that investment is a positive-sum rather than a zero-sum game across the generations.

Concluding and summary remarks

My informal polling over the years suggests that many advocates of greater state spending – especially non-economists -- like the idea of a very low discount rate. Many of these individuals would like our government to devote more resources to education, to infrastructure, and to improve the environment. They see a lower discount rate as providing support for all of these policies. More generally, these individuals believe we are not caring enough about the future. Very low rates of discount therefore serve as a "left-wing" view in most cases. Similarly, support for market-based discount rates often comes from centrist or more "right-wing" views.[28]

In contrast, I see a deep concern for the distant future as cutting across the current political spectrum. A greater orientation toward the future is likely to increase the desirability of policies favoring a market economy, economic growth, and technological innovation. Furthermore some of the arguments for these institutions may require a deep concern for the more distant future. For instance positive rates of discount usually imply that we should grant considerable importance to the alleviation of immediate suffering. Market liberalizations, whatever long-run virtues they may have, sometimes increase immediate suffering. Furthermore, market economies tend to invest their surpluses in long-run growth, rather than redistributing it to the immediately suffering poor.

Market economies and market reforms look best, the greater the weight we place on the relatively distant future. A free society is better today than a corrupt and totalitarian alternative. But one hundred years from now, the difference in human welfare, and other relevant values, will prove far more pronounced. Over time America gained ground on the Soviet Union, rather than allowing convergence.

To close, let us sum up some of the core conclusions of this chapter:

[28] For a “left-wing” view of discounting, see, for instance, Solow (1974). Beckerman (1996) offers a market-oriented view, critical of zero discounting.

1. Our deep concern for the distant future has concrete policy implications.
2. Welfare economics should focus on empirical questions of growth and place less emphasis on traditional cost-benefit analysis. The exact scope and extent of this recommendation, however, remains an open question.
3. Economic research should focus much more on which policy changes involve once-and-for-all changes in wealth, and which have a long-term impact on the rate of growth.
4. The proper role of government is to support growth-enhancing public goods.
5. We should pay greater heed to the long-term safety of our civilization.
6. We should care most about those environmental problems that will impact the long run rate of "true GDP" growth. One-time losses and adjustment costs are less important.
7. Even if we are strict utilitarians, our collective obligations to the very poor are more limited than is commonly believed. Given a sufficiently long time horizon, economic growth is the best means of improving the lot of the poor. Our strongest obligation is to adopt growth-maximizing institutions.
8. These answers do not stand in gross practical conflict with the conclusions of common-sense morality.

With those recommendations in place, let us turn to some remaining critical questions for our enterprise.