



Tyler Cowen



TYLER COWEN

Appendix A

Some Optional Mathematics and Remarks on a Few Metrics


Appendix B

Animal Welfare and Derek Parfit's Repugnant Conclusion

References

References

. . .



Chapter one: Introduction

Three Philosophical Starting Points

When it comes to the future of our world, we have lost our way in a fundamental manner and not just on a few details. We must return to

principles, but we do not always have good principles to guide us. We have wandered from ideals of a society based on prosperity and the rights and liberties of the individual, and yet we do not know how to return to those ideals.

That sounds so simple: prosperity and individual liberty. Who could be opposed? In the abstract, few people speak up against those values. But in practice we turn away from them all the time. We pursue many other ends which we should ignore or reject and that is one regard in which we have lost our way. We instead need a tougher, more dedicated, and indeed a more stubborn attachment to prosperity and freedom than is currently the case. When you see what this means in practice, you may wince at some of the implications and you may be put off by the moral absolutism it will require. Yet these goals—strictly rather than loosely pursued—are of world-historic importance for civilization and if we adhere to them they will bring an enormous amount of good into our world.

But how do we know which goods we should be pursuing, and how do we weigh one value against another? How should we make decisions when moral values clash? These rather corny questions of freshman bull sessions, presented and chewed over around the dorm, remain of vital import.

Before considering how to make such trade-offs, here is some background on my underlying philosophical stance and what I am bringing to the table.

I treat questions of right and wrong as having correct answers, at least in principle. We should admit significant grey areas, but right and wrong are a kind of “natural fact,” as many philosophers would say. To put it more bluntly, there is an objective right and wrong. Relativism is a non-starter and most people are not sincere in their relativist pronouncements in any case. At some gut level they think they know right and wrong; if you doubt this watch them lecture their kids or better yet criticize their colleagues.



“I will of course yield to whomever agrees with me.”

That said, I am not going to spend time on what the concepts of right and wrong “really mean,” whether they come from God, or whether we always have compelling reason to act in a moral way. I will not consider what is called meta-ethics, namely the study of the underlying nature of ethical judgments. I’m going to assume that right and wrong are concepts which make fundamental sense. Even if that’s not exactly your view, perhaps you could slot many of my arguments to your favored alternative stance.

It is often very difficult—in concrete cases—to discover which particular course of action is right and which is wrong. The skeptic is underappreciated, especially in an age of polarized politics when each side is convinced it is right and the other is unacceptably wrong. Science is our main path to knowledge and yet so often science tells us we don’t know. That is all the more true for social science, and perhaps macroeconomics stands at the summit of our epistemic limitations. So, to consider the realm of politics, we should not engage in the sport of building a coalition of like-minded individuals, defeating the competing coalitions, and then implementing what we already know to be best. That’s a popular approach, and it makes us feel good about ourselves and our own supposed superiority, but it is unjustified. We need to be more modest than that when it comes to what we can possibly know.

Philosophers David Hume and William James both understood the smallness of the individual human mind, compared to the vast expanse of nature and society, and they emphasized the irrationalities of the human mind, compared to the daily problems which are put before us. If we are building principles for politics, we need approaches which are relatively robust to human error, robust to the rampant human tendency for self-

deception, and which can transcend our own tendencies for excessive “us vs. them” thinking.




“To err is human, to really mess up takes planning.”

Yet, at the same time, we need doctrines *we can actually believe in* and which provide a foundation for a political and social order. A fine-tuned philosophic doctrine which no one accepts or ever could accept won’t be of much use. Reconciling the need to accommodate both skepticism and belief is one of the trickiest tasks for a philosophy. If we are indeed skeptics of a sort, how can we end up believing in anything of real import? To give this question another framing, what are the roles of reason and the roles of faith in how we move forward?

Next, I hold **pluralism** as a core moral intuition. What’s good about an individual human life can’t be boiled down to any single value. It’s not all about beauty, all about justice, or all about happiness. More plausibly, pluralist theories postulate a variety of relevant values, including human well-being, justice, fairness, beauty, the artistic peaks of human achievement, the quality of mercy, and the many different and indeed sometimes contrasting kinds of happiness. Life is complicated! That means no single value is a

“trump card” which overwhelms all other values in all instances and thus there is a fundamental messiness as to the nature of the good. At first that recognition of messiness may seem inconsistent with an attachment to rigid ideals of prosperity and liberty, and that reconciliation will prove to be a central issue.

Sometimes my fellow economists argue that “satisfying people’s preferences” is the only value which matters—because in their view it encapsulates all other relevant values—but that approach doesn’t work. It is not sufficiently pluralistic, as it also matters whether our overall society fits standards of justice, beauty, and other values from the plural canon. “What we want” does not suffice to define the good. Furthermore often we must judge people’s preferences by invoking other values, external to those preferences. To give an extreme example, when we condemn a wife-beater, must we really calculate whether the suffering of the victim exceeds the pleasure of the hitter? I think not. Furthermore, if individuals are poorly informed, confused, or downright inconsistent—as nearly all of us are—the notion of “what we want” isn’t always so clear. So while I am an economist, and I will use a lot of economic arguments, I won’t always be siding with the normative approach of my discipline, which puts too much emphasis on satisfying preferences at the expense of other ethical values. We need more room for justice and beauty.



I sometimes call myself a “two-thirds utilitarian,” since I look first to human well-being when analyzing policy choices. If a policy harms human well-being, on net, it has a high hurdle to overcome. If “doing the right thing” does not create a better world in terms of well-being, on a repeated basis, we should start wondering whether our conception of “the right thing” makes sense. That said, human well-being is not an absolute priority and thus the half-in-jest reference to the two-thirds weighting for utility. We sometimes ought to do that which is truly just, even if it is painful for many people involved. We should not take away one of your kidneys by brute force simply because you can do without it and someone else needs one. We should not end civilization to do what is just, but justice sometimes trumps utility. And justice does not reduce to what makes us happy or to what satisfies our preferences.[1]

In short, my philosophical starting points are:

1. “Right and wrong” are very real concepts which should possess great force.
2. We should be skeptical about the powers of the individual human mind.
3. Human life is complex and it offers many different goods, not just one trump value.

Given these beginnings, we now must turn to actual choice and I see some key questions for the individual and also for the collective. Why do we prefer one choice over another? To what extent do we have good reasons for such preferences? Exactly which choices should we make?

[1] Smart and Williams (1973), Scheffler (1982), and Singer (1993) offer some standard treatments of consequentialism—the evaluation of choices in terms of their consequences, a philosophical doctrine which includes utilitarianism as one variant. Pettit (1997) offers one good introduction to consequentialist reasoning and why it is persuasive.

. . .



⑥

Six Critical Issues

*Time, Aggregation,
Rules, Uncertainty,
Rights, and Morality*

Six Critical Issues


To make progress on these queries, I will consider six critical issues, each of which can help us resolve clashes of value:

1. Time

How should we weight the interests of the present against the more distant future? This relates to a more metaphysical question: do we have reasons to weight the present more heavily simply because it is the here and now? Does the economic approach to time discounting—which suggests the future declines in moral importance as time passes, and in rough proportion to market interest rates—apply? (I will say no.) This is a key question for deciding how much we should commit to making the future a better place.

2. Aggregation

Aggregation refers to how we resolve disagreements and how we decide that the wishes of one individual should take precedence over the wishes of another. At a very simple level, if your daughter wants Japanese anime on the DVD player and your son wants a Disney cartoon, whose desire should prevail? At the social level, if John wants greater income equality and Cecilia does not, on what grounds might we elevate one preference over another? There are well-known results from economics, for instance the Arrow and Sen Impossibility Theorems, which suggest aggregation problems are very difficult to solve and perhaps intractable altogether. Some commentators read these theorems as showing that we cannot rationally decide who should get his or her way when people disagree. Yet we must resolve clashing preferences every day and we cannot in fact retreat into an operative nihilism. I'll tip my hand and suggest that I'm optimistic about resolving this problem in a workable manner.



“I get so tired of jousting with that guy.”

3. Rules

The idea of rules and general principles for our choices and also for politics is a compelling one. But what does it mean to adhere to such rules? We'd like to think that rules have independent power and force for their own sake, but that view is difficult to defend under fire. After all, it seems that virtually all rules have exceptions. Sometimes you should lie, if only to save innocent victims from their persecutors, to cite one classic example from moral philosophy. When we decide when and whether to break a given rule, we're back to judging individual cases, which is what we tried to get away from in the first place.

Philosophers pose similar questions with a different language. They debate the doctrine of rule utilitarianism, which suggests that we should choose the rules to maximize social utility. Similarly a broader and more pluralist approach—called rule consequentialism—looks for the rules to maximize a broader notion of good consequences.[1] Yet does the doctrine of rule utilitarianism (or rule consequentialism) collapse into act utilitarianism? That is, are we not always second-guessing our adherence to the supposed rule? Should we just break the rule, when the individual circumstances of the case dictate we should do so?

Under one common view, rules are a mere fiction or a phony trick, albeit a useful one, but holding no independent force in our moral reckoning. Can we generate a coherent morality in which we should respect rules and principles for their own sake? Can we make a fundamental choice to think in terms of rules and principles per se? Might we even obsess over rules? Or are we caught in the trap of always worrying about the exceptions and thus we end up back with rules as a useful fiction?

I'm going to speak up for rules.

4. Radical uncertainty

I'm a skeptic, but I'm also a skeptic who has a "can do" temperament and who realizes how paralyzing skepticism can be. To put it bluntly, it is extraordinarily difficult to predict the distant future. I'm not just talking about the difficulty of constructing good theories in the social sciences and testing those theories against the data. There's a broader inability to trace definitive chains of cause and effect in human affairs.



I'm still bugged by some pretty common problems from science fiction and speculative fiction. To take a simple example, even our smallest moves can set off extensive chain reactions with far-reaching repercussions. Imagine if Hitler's father—or how about Caesar's father?—had arrived one second later

to the marriage bed. A different sperm and egg would have come together and the whole subsequent course of human history probably would have been different. When I was a kid, I read a comic book story about a team of researchers going back in time to observe the dinosaurs, but also trying not to disturb the past. By mistake, their ship crushes one leaf and the entire subsequent course of human history is changed as a result and the American Indians end up conquering Europe. If there is going to be a future version of Hitler, my actions writing this work—and your role in reading it—are helping to give rise to his later conception and birth.

Given such long-run uncertainty, how can we pretend to assess good and bad consequences as a product of our actions? How can we make any decision at all without being morally paralyzed and totally uncertain about what we are doing?

Usually people skirt over this problem by saying we need to move on and do the best we can. There is much truth to that, but the question is still how our epistemic modesty should shape what we are actually going to do.

5. How is it that we can believe in rights?

The notion of rights is often put forward in philosophic discussion and in the world it exercises considerable sway, including in international law. But within philosophical circles, the foundations for rights are often viewed as shaky or, even when a notion of rights is accepted, rights are seen as relying too much on pure intuitionism. I'm not going to derive rights deductively from scratch, but I do believe in (nearly) absolute human rights. I will put forward a doctrine of "rights without embarrassment." That's not nearly as strong as a proof or derivation of rights, but some key elements of ethical reasoning point our attention in the direction of objectively valid human rights and indeed their near sacred character.[2]

As you might expect from someone who is speaking up for rules, I'm going to speak up for rights too.

6. Common sense morality

Common sense morality suggests that we should work hard, take care of our families, and live virtuous but self-centered lives, while giving to charity at the margins and helping out others on a periodic basis. Utilitarian philosophy, on the other hand, appears to suggest an extreme degree of self-sacrifice. Why should a mother tend to her baby when she could sell it and send the proceeds to save a greater number of babies in Haiti? Shouldn't anyone with the

training of a doctor be obliged to move to sub-Saharan Africa to save the maximum number of lives? What percentage of your income do *you* give to charity? Given the existence of extreme poverty, shouldn't it be at least fifty percent? If you belong to the upper middle class, how about giving away eighty percent of your income? You don't really need cable or satellite TV, or perhaps you should eat beans with freshly ground cumin instead of meat. The bank might let you borrow to give away even more than you are earning. How terrible is personal bankruptcy anyway, if you have saved seven lives in the meantime?



Is eating that ice cream cone so important? Common sense morality implies it's OK to enjoy that chocolate but utilitarianism suggests maybe not. British philosopher Henry Sidgwick was obsessed with whether the recommendations of utilitarianism and common sense morality can be made broadly compatible; later British philosopher Derek Parfit picked up on the

same theme. And if those two approaches are not compatible, which ethical perspective should we prefer at which margin and why?

The **six** questions stated above are all tough ones. Yet I will put forward and defend a controversial claim: all these difficult problems are more of a whole than we used to think. These problems can indeed be resolved and by a relatively small number of intellectual and philosophical “moves.” By no means do these moves make life easy for normative evaluation because they will create *new* problems. But I will be suggesting that we trade in our old problems for these new problems. I will be suggesting that, upon reasoned examination, we *must* trade in our old problems for these new ones.

[1] For two recent discussions of what consequentialism means, see Hurley (2009) and also Brown (2011).

[2] For one recent look at rights, which also surveys many of the major questions, see Griffin (2008). Wenar (2013) also surveys some aspects of the current debate, in addition to his original contributions. On intuitionism, see Huemer (2005).

. . .

Tough Questions and Faith in the Future




Tough Questions and Faith in the Future

The fundamental philosophical moves I have in mind are twofold.

First, I do not take the productive powers of economies for granted. Production could be much greater than it is today and our lives could be more splendid. Or if we make some big mistakes production could be much less and we could all be much poorer. This simple observation helps us put the idea of production at the center of moral theory, as without production value is problematic. For all of her failings, Ayn Rand is the one writer who has best understood the importance of production for moral theory, a point which she expressed enthusiastically at great length, albeit with numerous unfortunate caricatures. It is the work of capital, labor, and natural resources—driven by the creative individual mind—which undergird the achievements of our civilization. Whether or not you agree with all of Rand's political views, do not take the existence of wealth for granted.

Second, I will seek to revise some of our intuitions about “moral distance.” Which individuals should have a pull on our choices and which should have less influence? I will argue for instance that the individuals who will live in the future should be less distant from us, in moral terms, than many people believe. Their interests should have greater sway over our calculations and that means we should invest more in the future. Even though it is sometimes hard for us to imagine the consequences of our actions for future people, especially people from the more distant future, their moral import remains high. I thus will be asking for a greater faith in the future. I am not asking for a faith at the *expense* of reason, but still it will take an attitude very much akin to faith to consistently think so far ahead in our calculations. It is no accident that religious people on average have much higher rates of fertility, or that they engage in so many long-term business and charitable projects for the future, as suggested long ago by Max Weber.

When considering the questions outlined above, I will focus on clashing arguments and the substantive bottom line. I do not devote much time to building consensus on familiar material, surveying what everyone has said on a particular topic, or other niceties. I do not retread familiar ground, offering some “suggestive remarks” on the tough problems at the very end. I do not “argue by elimination” by focusing on the weaknesses in other views and downplaying the weaknesses of my own. Instead I seek to start with the tough questions, try to pinpoint the hardest points of dispute, and spend the rest of the exposition trying to pick up the pieces. That is what I like to read and thus what I am trying to write.



This is where my internal mental debate stands. I hope you enjoy it. If you are the kind of reader I want, no matter how hard I push, you will feel I have not pushed hard enough on the tough questions.

I also hope that you will respond by taking a stronger stand on behalf of the ideals of freedom and prosperity. I hope that you will join more firmly in the cause of making our civilization stronger, more durable, and also making our civilization a more wondrous place. I am suggesting that we need a radical reawakening. And this reawakening, it turns out, will prove a new and compelling way of reaffirming our own power as individuals.

With all that in mind, let's proceed.

. . .



ON VALUE, THE CRUSONIA PLANT, WEALTH PLUS &

Chapter two: Wealth makes the world go round

On Value, the Crusonia Plant, and “Wealth Plus”

here does value come from? How is value created and maintained and indeed augmented over time? These simple questions sound like

W clichés but they are central to ethics.

Let's consider the idea of systems, networks, norms, or policies which create systematically increasing value over time, without apparent end. Milton Friedman used to argue that there is no such thing as a free lunch but at some level of the analysis this has to be false. The universe exists and who had to pay for it? (Even better, under some cosmological hypotheses the universe arose out of quantum fluctuations and it is forever giving rise to offshoot baby universes, all "for free.") That seems to be at least one free lunch—a big one—so maybe there are others. What can we usefully think of as a free lunch when it comes to social decision-making? After all, somehow planet earth got from a bunch of trilobites to trillions of dollars in gdp and a Louvre full of beautiful paintings.

Since free lunches aren't always easy to find, we should think about where free lunches might be hiding or why some of those free lunches are less than evident. In particular we might uncover hidden gains if we more closely consider the dimension of time. Maybe some of our choices could release a steady gush of benefits, but we don't always see those benefits as clearly as we ought to. As Adam Smith noted long ago in the eighteenth century, we visualize many future events very poorly and with a deficit of proper imagination.



In economics, there is at least one (hypothetical) example of a free lunch over time. Economist Frank Knight wrote of the Crusonia plant, a mythical, automatically growing crop which creates more output each period. If you lay the seeds the plant just grows and you don't have to water it or tend to it. Imagine for instance an apple tree, which each year yields some apples. The tree also produces apple seeds. The apple seeds germinate and there is a steady and indeed growing supply of new apples and also of new apple trees, albeit based on some sun and some rain. A Crusonia plant, measured in terms of its ability to produce apples, might grow five percent each year on net. At the same time, it looks like a modest apple tree, and it does not appear to resolve key ethical and political questions.

The Crusonia plant may sound unrealistic or a bit silly, but it's a useful example for pinpointing the nature of our quest. The Crusonia plant is an example of a free lunch—at least a free lunch of apples—once you have obtained it. By the way, if you're wondering why it is called the Crusonia plant, Crusonia was the name of Robinson Crusoe's island. In the Daniel Defoe novel, Robinson Crusoe's island had a lot of trees which yielded fruit for nothing, without requiring labor from Crusoe or anyone else and thus Crusoe enjoyed some free lunches, just as we can too.



Frank Knight postulated the Crusonia plant to make some technical points about the theory of capital and investment, but those debates have since passed. In contrast, I see the Crusonia plant is an entry point for resolving aggregation problems. A Crusonia plant would be better than a plant which dies after one month and leaves no successors, even if this short-lived plant were quite lovely or brilliant. We could compare two plants in terms of various qualities, such as their color or their scent, but after a while the unceasing free yield of the Crusonia plant has to prove better. At some point the sheer accretion of value, from the ongoing growth of the Crusonia plant, dominates the comparison between the two plants. We thus have a principle of both ethics and prudence: when in doubt choose the Crusonia plant.

When it comes to making tough decisions, we should try to identify which elements in the choice set resemble a Crusonia plant. If we could find some choices or policies which give rise to the equivalent of Crusonia plants, namely ongoing and self-sustaining surges in value, the case for those choices

would be compelling. Furthermore, if it turned out that Crusonia plants were more common than is at first sight apparent, aggregation problems would be eased more generally.

So for a social setting, what might count as analogous to a Crusonia plant? Look for social processes which are ongoing, self-sustaining and which create rising value over time. The natural candidate for such a process is economic growth or some modified version of that concept. If sustainably positive-sum institutions exist, there may be Crusonia plants all over. As we'll see, standard definitions of economic growth do not fully qualify as true Crusonia plants, in part because they ignore environmental sustainability and in part because they do not adequately value leisure time. Nonetheless if we think about economic growth a little more broadly, we will have a relevant Crusonia plant for making decisions.



"I'm knitting my son a scarf."

Economists cite the concept of gross domestic product (gdp) to refer to the total value of goods and services produced over some period of time, usually a year or a quarter. The rate of economic growth is then the rate at which gdp

increases. As I will use the concept, maximizing the long run rate of economic growth refers to gross domestic product as it should be understood properly and not as it is currently measured by most governments. “Wealth Plus,” if I may use that term to refer to the accumulated gains from growth, accounts for leisure time, household production (valuable activities you do at home for free, whether mending socks or using Facebook), and environmental amenities, among other adjustments. Current gdp statistics have a bias towards what can be measured easily and relatively precisely, rather than focusing on what contributes to human welfare.

With this in mind, I will define the concept of Wealth Plus:

Wealth Plus:

“The total amount of value produced over some time period. This includes traditional measures of economic value, as would be found in gdp statistics, but also measures of leisure time, household production, and environmental amenities, as summed up in a relevant measure of wealth.”

In this context, maximizing Wealth Plus does not mean that everyone should work as much as possible. A fourteen-hour workday might maximize measured gdp in the short run but it would be less propitious over time, once we take the value of leisure into account, not to mention labor burnout. Still, this standard is going to value a strong work ethic.[1]

Maximizing Wealth Plus also does not mean destroying the natural environment. It’s now well understood that environmental problems can lower or destroy economic growth through feedback effects. We therefore should protect the environment enough to preserve and indeed extend economic growth for the more distant future.




Most generally, we wish to maintain higher growth over time, and not just for a single year or for some other shorter period of time. 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 economically and politically. 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 and the economy has gone backwards in many ways. The Shah's forced modernization did not in fact maximize true economic growth and more cautious policies likely would have been better.[2]

The concept of sustainability, as embedded in the Crusonia plant idea, thus focuses our attention on the prerequisites for a durable civilization. And although most people live in a peaceful and relatively prosperous time, we should not take durability for granted. If we look at the broader historical record, economic growth is hardly the rule and civilizations show 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, various dynasties and republics of China, four periods in Africa, three in India, two in Japan, 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.[3]

While the exact numerical estimates depend on how we define the concept of civilization, and how we pin down start and end points, the more general point of fragility stands.[4] Human beings can and indeed do experience significant and ongoing losses of their prosperity and freedom.



We might wonder whether we would maximize the relevant pluralist values by existing at a very modest population and economic level for a very long time, “living in harmony with nature,” so to speak. Think of Tolkien’s quaint hobbits. But poorer societies from the past have collapsed repeatedly through military weakness, eco-catastrophe, famine, tyranny, and natural disasters, among other factors.

Kee in mind that the wealthier tyrant will conquer or at least disrupt the noble savage. Even if in principle the life of the noble savage were best, no society following this path will, on its own, keep its autonomy in the longer run. Given the previous path of human development, someone will have tanks and nuclear weapons, whether we like it or not. It is important that the more benevolent societies be both richer and more technologically advanced, and again we see the relevance of sustainable economic growth.

Furthermore primitive warfare appears to have been at least as frequent, bloody, and arbitrary in its violent effects as modern warfare.[5] 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 comfort. In other words, we probably need to move forward rather than seek a static, quiet existence, yet our path requires a tightrope act, balancing progress and stability along the way.

We already can see that three key questions should be elevated in their political and also philosophical importance, namely:

1. What can we do to boost the rate of economic growth?
2. What can we do to make our civilization more stable?, and
3. How should we best deal with environmental problems?

The first of these is more commonly considered a “right wing” or perhaps libertarian concern, the second is most commonly a “conservative” preoccupation, and the third is, especially in the United States, most commonly associated with “left wing” perspectives. Yet these questions should be central, rather than peripheral, for everybody. We can see right away how the political spectrum should be reshaped. Politics should be about finding the best means to achieve these ends, rather than disputes about the importance of these ends.

[1] In terms of providing operational guidance in calculating Wealth Plus, the two best efforts I know are Jones and Klenow (2010) and Becker, Philipson, and Soares (2005). For a good piece which puts production and productivity at a central place in moral theory, see Stanczyk (2012). See also numerous works by the philosopher David Schmidtz.

[2] For a formal look at the concept of sustainability, see Heal (1998, chapter one).

[3] Rees (2003) and especially Posner (2004) both focus on issues of civilizational collapse. S.E. Finer (1999, 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.

[4] 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.”

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

. . .

HOW GOOD IS GROWTH ANYWAY?



How Good is Growth Anyway?

The history of economic growth indicates that, with some qualifications, it alleviates human misery, improves human happiness and opportunity, and it lengthens human lives. Wealthier societies have better living standards, better medicines, and offer greater autonomy, greater fulfillment, and more sources of fun. While measured wealth does not exactly correspond to Wealth Plus, in the past these two concepts have moved together pretty closely, especially across the range of outcomes we have observed (as opposed to hypothetical thought experiments and counterfactuals).

We often forget how overwhelmingly positive the effects of economic growth have been. Economist Russ Roberts reports that he frequently polls journalists and asks them how much economic growth there has been since 1900. By Russ's account the typical answer is that the standard of living has gone up by around fifty percent. In reality, the U.S. standard of living has gone up by a factor of five to seven—estimated conservatively—and possibly much more,

depending on which techniques we use for measuring prices and the values of outputs over time, a highly inexact science.



The data show just how much living standards have gone up. In 1900 for instance almost half of all U.S. households (0.49) had more than one person per room and almost one-quarter (0.23) over 3.5 persons per sleeping room. Slightly less than one-quarter (0.24) of all U.S. households had running water, eighteen percent had refrigerators, twelve percent had gas or electric light, and today the figures for all of these stand at 99 percent or higher. Back then only five percent of households had telephones and none of them had radio or TV. The high school graduation rate was only about six percent, and most jobs were physically arduous and to high rates of disability or even death. In the mid-nineteenth century a typical worker might have put in

somewhere between 2800 and 3300 hours a year, but now this is closer to 1400 to 2000 hours a year (Huberman and Minns 2007).[1]

Until recently in history, polio, tuberculosis, and typhoid were common ailments, even among the rich. In the United States, earlier Presidents George Washington, James Monroe, Andrew Jackson, Abraham Lincoln, Ulysses S. Grant and James A. Garfield all caught malaria during their lives. Antibiotics and vaccines have existed for only a tiny fraction of human history, and it is no coincidence they exist in the wealthiest time period humanity has seen. There is also a strong and consistent relationship between wealth and rates of infant mortality; small children do best when they are born into the wealthier countries, and that is because wealth supplies the resources to take better care of them.



As recently as the end of the nineteenth century, life expectancy in Western Europe ran about forty years of age. Economist 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.”

In earlier times most individuals worked at hard physical labor and a college or university education, or even a high school education, was a luxury. Leisure time has risen with economic growth. Again citing Robert Fogel (2004, p.70), 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.”

The splendors of the modern world are not just frivolous baubles but they are important sources of human comfort and well-being. 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 citizens of wealthy countries today. My smart phone is as good as his. The very existence of an advanced civilization—the product of cumulative economic growth—confers immense benefits on ordinary citizens, including their abilities to educate and entertain themselves or to choose one life path over another.[2]

Evolution of Phone








The economic growth of the wealthier countries benefits the very poor as well, even though this is sometimes with considerable lags. The distribution of wealth changes over time, and not all growth trickles down, but as an overall historical average the bottom quintile of an economy shares in growth.[3] You can see this by comparing the bottom quintile in say the United States and in India or Mexico.

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 usually become richer yet. A typical migrant from rural Mexico to the United States will move from earning \$2 a day to \$10 an hour or more. Over time the children of immigrants approach the national average, depending how long they have been in the country and how keen they are to assimilate. And the richer the receiving country, usually the more new immigrants will benefit. Central American immigrants to the United States do better than Central American immigrants to Mexico or Nepalese

immigrants to India. Immigrants also send remittances back home and at a rate which far exceeds governmental foreign aid. In all of these ways the actual upward mobility of the United States far exceeds what the usual numbers indicate, because published mobility numbers do not typically include a comparison with pre-immigration outcomes.

Many migrants eventually return to their home countries, bringing skills and liberal democratic ideas. Software repatriates have helped build India's competitiveness in high-tech industries. Many thousands of Asian students have obtained science or engineering degrees from Western universities, thereafter returning home to start new businesses. If a country offers scope for entrepreneurship, it will more likely experience a "brain gain" than suffer from a "brain drain." Foreign contacts, training, and periods of residency all help promote domestic development.



The global poor also benefit from new medicines, new global technologies, and research and development efforts, all of which are the product of wealth. Elhanan 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..."

Although these historical processes have often embodied unfairness and long lags, of decades or more, economic growth nonetheless has brought wealth to the poor and elevated their status. The Greek city-states and the Roman Empire benefited 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 and eventually human beings started to climb out of the Malthusian trap of very low per capita incomes at subsistence. The wealth of the West helped enable the export miracles of the East Asian economies. Today most poor countries seek greater access to wealthier Western and Asian markets and flourish if they can achieve it.[4]

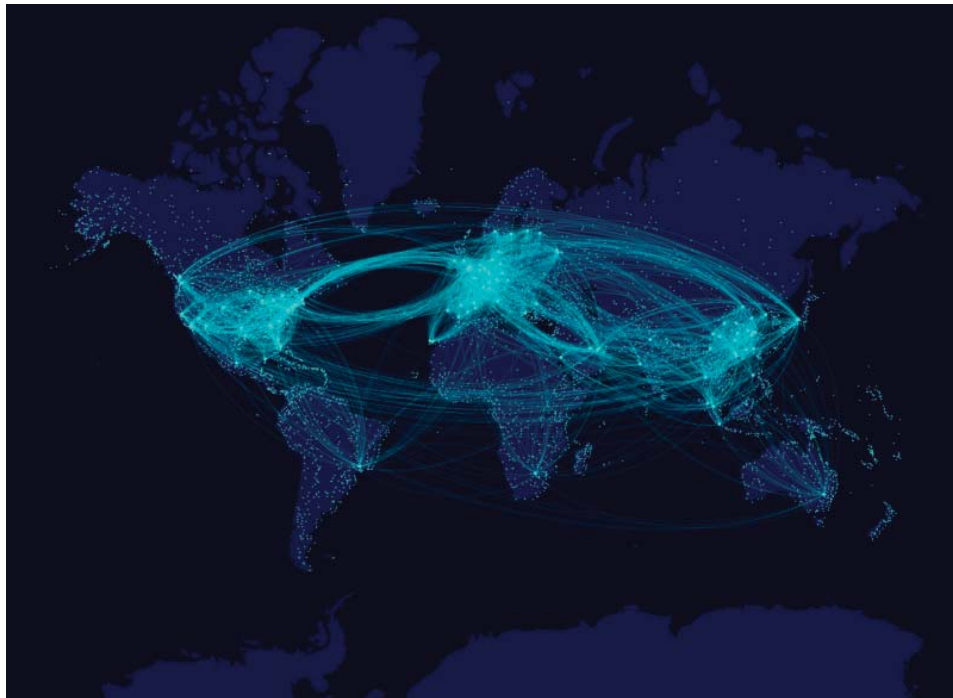


For all the recent increases in inequality within individual nations, *global* inequality has declined over the last few decades, in large part because of growth in China and India. The growth in emerging nations in large part has in turn been driven by earlier growth in the West and in East Asia. China for instance engaged in “catch up” growth based on adopting Western technologies and also by exporting to the wealthier nations. China has gone from being a quite poor nation to a “middle income” nation with a sizable middle and upper class. Although it belies a lot of the recent media coverage, which focuses only on “within nation” magnitudes, recent world history has been an extraordinarily egalitarian time. Most of all, it is a story of how global economic growth helps the poor. There has been a squeezing of the middle class in the wealthier nations, in part because of increasing global competition. Still, we have seen economic growth, aggregate wealth, and global income equality all rising together over the last twenty-five years. Most citizens in East Asia, South Asia, and Latin America have seen significant gains in their living standards, and much of this has been a trickle-down effect from the earlier growth of the wealthier countries. Much of Africa is now following suit as well, in part boosted by China’s demand for raw materials and also by the spread of modern technologies such as affordable cell phones.[5]

Sometimes extended periods of growth do not bring full or fair benefits for the poor or lower classes, for instance during the early phase of the British industrial revolution in the late eighteenth century. Still, even given these transitional troubles, the historical record suggests it was better for Britain to push ahead with economic growth, as this eventually drove the greatest boost in living standards the world has seen. To be sure, there probably were better policies at the time which would have distributed the benefits of growth more widely (e.g., fewer wars and Poor Law reform and free trade for the British). But taking misguided policies as given, it was better than Britain pursued economic growth rather than turning its back on the idea, even though significant real wage gains for the working class often did not arrive until the 1840s.


Nobel Laureate Amartya Sen has promoted the idea of “capabilities” as, if not quite a substitute for economic growth, an alternative focus. Sen points out that our positive opportunities in life often matter more than the amount of cash in our bank accounts. He also notes that some parts of the world, such as the state of Kerala in India, have relatively good health and education indicators, even though their per capita incomes are relatively low. Sen’s points are well taken but they do not put a fundamental dent in the relevance of wealth, or as I am calling it here Wealth Plus. The benefits from “capabilities,” insofar as they are important, such as health benefits, count in Wealth Plus even if they are not represented properly in current gdp

measures. In other words, Kerala is wealthier than some limited statistical measures imply. Wealth and good social outcomes still are strongly correlated on average and this correlation is stronger over longer time horizons. For instance, if Kerala does not grow much in more narrow economic terms, it is unlikely to look so impressive in its social indicators fifty or one hundred years from now. Even today Kerala manages as well as it does in large part because so many Keralans take jobs in wealthier countries, especially in the Gulf States, and send money back home. And compared to other Indian states, Kerala has a clearly above-average measure of wealth and also above-average consumption expenditures, as picked up by traditional statistics.[6]



The truth is that economic growth is the only permanent path out of squalor. Economic growth is how the Western world climbed out of the poverty of the year 1000 A.D. or 5000 B.C., it is how much of East Asia became remarkably prosperous, and it is how our living standards will improve in the future. Just as the present appears remarkable from the vantage point of the past, the future, at least provided growth continues, will offer comparable advances, including perhaps greater life expectancies, cures for debilitating diseases, and cognitive enhancements. 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. I have myself written of “the great stagnation” as a growth slowdown which overtook the Western world, starting in about 1973. It is a failure of imagination, however, to believe that human progress has run its

course. The more plausible view is that progress is unevenly bunched, we have been in a slow period as of late, various new developments are percolating, and we should do our best to help them along. Whether we like it or not, economic growth and technological progress do not always come in steady doses.



“At least we are consistently inconsistent.”

World history offers various 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. For that matter the East Asian revolution in economic growth was not widely anticipated. Each development, over time, dramatically 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. It is unlikely that we have seen the last of such revolutions, at least provided that civilization manages to stay afloat.

A critical point is that 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.[7] 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, defined in terms of constant real dollars. At a growth rate of one percent, such an improvement takes 393 years.

Robert E. Lucas (1988, p.5), Nobel Laureate in economics, 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 on utilitarian values, the benefits of growth are by no means restricted to utility in the narrow sense. For instance wealthier societies bring much greater access to the arts and education. Economic growth also minimizes “the tyranny of place.” Individuals suffer a lack of freedom when they have little or no chance to escape 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 diverse cultural traditions when constructing our personal narratives. Benjamin M. Friedman, in his brilliant *The Moral Consequences of Economic Growth*, shows just how many of the virtues of the modern world depend on higher and indeed growing levels of wealth.[8]

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—and other plural values—on an ongoing basis. If some set of choices or policies give us a higher rate of economic growth, those same choices or policies are akin to a Crusonia plant.

[1] These figures are drawn from the Bureau of the Census, via an unpublished paper by Bradford DeLong, “Lecture 2: Slow Growth and Poverty in the North Atlantic, 1800–1870.” On work hours, see Huberman and Minns (2007).

[2] I am indebted to Don Boudreaux for this framing of the point.

[3] See for instance Dollar and Kraay (2000).

[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.

[5] On the decline in global inequality, see for instance Lakner and Milanovic (2014).

[6] See for instance Venkatraman (2009) and Tsai (2006).

[7] Cowen (2003).

[8] See Friedman (2006). On the connection between the arts and economic growth, see Cowen (1998). For a recent defense of gdp maximization, see Oulton (2012). For a recent philosophic defense of the importance of economic growth based on pluralistic considerations, see Moller (2011). Stanczyk (2012) considers the importance of “productive justice,” namely making sure the output gets produced in the first place.



DOES ECONOMIC GROWTH MAKE US HAPPIER?

Does Economic Growth Make Us Happier?

We all like to speculate where in the world life is best or where the people are happiest. Is it in the richest countries, the countries where the people have the best psychological attitudes, or is there some other clue for answering this question?

Some recent research suggests that wealth boosts happiness and that is true for a great variety of people, including for the relatively wealthy who are already meeting their basic needs. For instance economists Betsey Stevenson and Justin Wolfers, in the most comprehensive study of the income-happiness link to date, find that the relationship between measured well-being and income is roughly linear-log, which implies income boosts happiness even at higher levels of earnings. The comprehensive study by Nobel Laureate economist Angus Deaton finds similar results, namely that extra income brings extra happiness, even in relatively wealthy settings.[1]

An older body of literature suggests that additional riches do not make citizens in wealthy countries any happier, at least not above a certain level of wealth. The core evidence here is taken from questionnaires which ask people how happy they are. Once a country has a per capita income of roughly \$10,000 a year or more, the aggregate income-happiness link appears weak to many observers. Some commentators argue that the curve flattens out at about half of current American per capita income. These results cast doubt upon whether economic growth does in fact yield ongoing benefits in terms of happiness.[2]

Despite this evidence, I see wealth and happiness as comoving in the broad sense, again subject to a properly sophisticated understanding of wealth.



The observation of a nearly flat happiness-wealth relationship reflects more about the nature of language than about happiness. To give an example, if you ask the people of Kenya how happy they are with their health, you get a pretty high rate of reported satisfaction, not so different from the rate in the healthier countries and in fact higher than the reported rate of satisfaction from the United States. The correct conclusion is not that Kenyan hospitals have hidden virtues, or that malaria is absent in Kenya, but rather that Kenyans have recalibrated their use of language to reflect what they can reasonably expect from their daily experiences. In similar fashion, people in less happy situations or less happy societies often attach less ambitious meanings to the claim that they are happy. Evidence based on questionnaires therefore will underrate the happiness of people in wealthier countries.[3]

The literature on happiness often focuses on aspiration or treadmill effects. Under this view, you get more but you start expecting more as well, or you start aspiring to more. The greater wealth therefore translates into less happiness than might have been expected at first. But it is unlikely that treadmill effects “eat up” all of the happiness gains from greater wealth. Along the lines of the Kenya example, growing wealth also causes people to recalibrate their language and 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 (if anything it could be easier to recalibrate your language than to recalibrate your happiness expectations). 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 even though your day to day existence is pretty sweet. The failure to issue a totally glowing report does not mean that you spend your entire time envying the billionaire or suffering because of your somewhat lower relative status; you still can lord it over plenty of other people, if you do desire. In the meantime, the presence of your neighbor may be creating higher standards for how you use the terms “happy” or “very happy.”

That means even a constant level of reported happiness implies growing real happiness over time, because the “happy” word is taking on more ambitious meanings as society accumulates more wealth and richer experiences. Life improvements do usually make us happier, while both our expectations for happiness and our reporting standards for “being happy”—our use of language—adjust upwards.



The belief that greater wealth correlates with greater happiness is supported by direct observation. Many individuals strive to earn higher incomes, even after they have experienced the strength of “aspiration” and “treadmill” effects. It’s not that they’re all being tricked, but rather they know at a gut level that money helps them achieve valuable ends, or that it helps their families. Individuals value those goals for their families, even if the channel of happiness is a quite indirect one which is not always reflected in their daily moods or their moment-to-moment self-reports.

It is also the case that within a country wealthier people report unambiguously higher levels of happiness, on average, than do poorer people. [4] For all the talk about how some happiness studies present a revisionist view of material wealth, this result has not been challenged and it pretty decisively demonstrates that, at least on average, wealth brings more happiness. To some extent the greater reported happiness of the wealthy may reflect a zero-sum relative status effect, namely that the wealthier people feel better but their possessions make the poor feel worse off. Nonetheless it is unlikely that the entire gains from wealth, or even most of the gains, dissipate in zero-sum games. Wealthier lives are easier and happier in absolute terms in numerous ways, as discussed above, and this is reflected by looking at where most immigrants wish to migrate, namely to the wealthier countries.

If a wealthy man buys a Mercedes, his polled neighbor may express greater dissatisfaction with his Volkswagen. That same neighbor, if he had a Lada in Moscow, circa 1978, might express a high or at least decent level of

satisfaction on a happiness questionnaire. Nonetheless in absolute terms he still likely very much prefers having the Volkswagen in contemporary America to having the Lada in 1978 Moscow. The Volkswagen is a pretty good car and the Lada broke down a lot and was hard to start. So the neighbor might envy the wealthy man's new car, but still the happiness gains from wealth, and from better cars, do not dissipate through envy. Better cars really do make us better off. Or put it this way: it is better to envy your neighbor's car than to envy his horse and buggy. Envyng his supersonic transport would be better yet.




On top of all of those considerations, happiness isn't a single, simple variable which can be measured unambiguously. Happiness means a lot of different things to different people. 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 typical day. Some will seek happiness through

the process of out-competing their peers for status, while others will look inward for contemplative delights. Most likely we seek some mix of these ends but with varying emphases and weights. Wealthier societies offer greater opportunities and freedoms to pursue preferred concepts of happiness, even if this privilege does not always show up in the measurement of a single, aggregate number.

The happiness literature takes a relatively limited view of human well-being. Usually 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, and perhaps it does, it is only one component of human well-being. A wealthier economy will offer greater options for structuring choices of work vs. friends, thrills vs. long-term satisfaction, enjoying children vs. a life of theater and travel, and so on.


A wealthier economy also gives us more “fleeting” happiness experiences. An individual will admit to being happier if he has recently found a dime, or if his soccer team won a big game. These sources of happiness will likely be more frequent and more consistent 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 fun and interesting new people. That sounds a bit superficial and indeed it is, but it is yet another reason why economic growth will boost happiness in its more complex and plural forms.[5]



At most, the happiness literature shows that many changes in individual conditions are irrelevant for our well-being, due to habituation and expectation effects. This conclusion would not, however, eliminate the major benefits of economic growth. Even if many “small” changes in income are irrelevant or nearly irrelevant for happiness, sufficiently large changes in life circumstances still may boost or harm our welfare.

To give 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 truly damages happiness for most individuals and those effects can persist 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 very often experience recurring flashbacks, nightmares, irritability, depression, alcoholism, troubled relationships, and inability to concentrate. However psychologically troubled our modern, wealthy societies may seem, poverty is not the solution to these problems and in fact it makes them worse.[6]

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 and a significant percentage of victims experience an ongoing “core of distress” for years.[7]



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 or at least their reports would indicate as such. But individuals remain subjectively badly off when they suffer from progressive or degenerative problems.[8] So to the extent that a poorer society brings an ongoing worsening of conditions for many individuals, the associated human suffering will be greater and this does indeed represent a true loss of human well-being. Once again, there are significant benefits from ongoing economic growth.

Extra wealth also serves as a cushion against very bad events, if only against later declines in wealth. Ten or fifteen years ago, it was common to

hear the claim that once a nation reaches the level of material wealth found in Greece, happiness more or less flat-lines, and indeed this was more or less where the flat-lining point seemed to be. Circa 2016, with Greece embroiled in an ongoing economic crisis for years, and no one uses the Greek example any more to make points about the flat-lining of the happiness-income relationship. The country lost almost a quarter of its economic output over the crisis, unemployment has run over twenty percent, there have been riots in the streets, a neo-Nazi party was elected to the legislature, and at times basic medicines have been unavailable. Some additional cushions of wealth, prior to the crisis, would have helped the country a good deal and perhaps would have prevented the troubles altogether, by easing debt repayment.



Finally, even if we accept the “flat-line” empirical result on happiness and wealth to be valid, the questions about self-reported happiness 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. 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 from the grave, at least not for researcher questionnaires. Life expectancy rises with wealth, but questionnaires and self-reports miss this benefit. Or if an immigrant, or a child of immigrants, fills out a happiness questionnaire, there is no comparison with a pre-immigration state of affairs, either as it was or as it might have been. By its 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 benefits brought by economic growth and also by change more generally. If we want to be around

to even have the option of answering happiness questionnaires, wealth is extremely important.

[1] See for instance Stevenson and Wolfers (2008, 2013), Sacks, Stevenson, and Wolfers (2010), and Deaton (2007).

[2] On the flattening of the curve, see for instance (e.g., Helliwell 2002, p.28). Other relevant sources are Argyle (1999), Oswald (1997), and Myers (2000). Wealthy countries, when they become wealthier over time, do not in general 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). For a criticism of the income-happiness from a time use perspective, see Kahneman et.al. (2006).

[3] See Deaton (2007).

[4] See for instance Dieter (1984), among many other sources.

[5] See Schwartz and Strack (1999) on some of these dilemmas, and Levinson (2013) offers some interesting observations on the role that “projection” (assuming excess permanence of current events) may play in these evaluations.

[6] On the link between catastrophes and unhappiness, see Dyregrov (1990), Lehman et.al. (1987), 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).

[7] On coping, see Brickman, Coates, and Janoff-Bulman (1978), Bulman and Wortman (1977), Kessler, Price, and Wortman (1985), 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).

[8] See Frederick and Loewenstein (1999).

work, it is usually assumed that consumption tomorrow is a relatively close substitute for consumption today.

. . .



Chapter five: What about redistribution?


On Moral Obligation and Redistribution

Why eat an ice cream cone when someone in Malawi is starving? That's one version of an age-old question in morals, yet it remains difficult to answer.

Under any moral theory which counts the interests of people in a more or less cosmopolitan manner, our personal obligations toward the poor appear strong. For instance several billion people live on less than \$2 a day. Last year millions of children died of preventable diseases, such as diarrhea, or they

experienced stunted growth and development. It's easy enough to feel that we should all be attending to such problems with more resources and life energy and indeed we should. The difficult question, however, is how far such obligations extend and whether such obligations should overwhelm us from pursuing our more personal or more individualistic goals. In contrast to utilitarianism, common-sense morality typically suggests that we do have the right to pursue our own lives and life plans.

Under some accounts of our obligations to others, every individual is obliged to work years 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, if not their entire careers. Many more of us would have to become doctors or nurses in the first place, unless of course you would be more effective as a Wall Street wizard and then a wealthy philanthropist. Maybe those options sound meritorious to you, but how far are you willing to go? Are you willing to treat the interests of other people on a par with your own or those of your family and friends? If you buy into the standard utilitarian logic of beneficence, a mother might have to abandon or sell her baby to raise the resources to allow her to send food to the babies of others. At this point most people balk and search for some moral principle which limits our obligations to the very poor.



"Wait, I'm YOUR therapist? I thought you were MY therapist."


One problem is that the needs of the suffering are so enormous that few able or wealthy individuals would be able to carry out individual life projects of their own choosing. They would instead become a kind of utility slave, serving only the interests of others and giving themselves just enough food and fuel to keep going. The result is that utilitarianism, or for that matter many forms of

consequentialism, often is seen as an excessively demanding moral philosophy. People fall into two camps: those who reject utilitarianism for its extreme and unacceptable implications, and those like the early Peter Singer, who trumpet the call for greater sacrifice and pursue the utilitarian logic to a consistent extreme.

As I've stated, I am a pluralist rather than a simple utilitarian. Still, utility is a central and important value, so I would like to confront these dilemmas and consider the scope of our obligations to the poor.

I don't, however, wish to focus on government vs. private charity. To be sure, that is an important practical issue, but I'll instead focus on the broader conceptual question of comparing growth to redistribution—public or private sector—as ways of helping the poor. When framed in this manner, we'll see there are some strong and strict limits on our obligations to redistribute wealth, even if we accept a full utilitarian framework. In other words, I'll end up siding with common-sense morality, which allows us (mostly) to pursue individual life goals.

It can be argued very plausibly and I think correctly that we are obliged to help the poor more than we are doing now. But the correct approach to our cosmopolitan obligations does not lead to 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. Our strongest obligation is to contribute to sustainable economic growth and to support the general spread of civilization, rather than to engage in massive charitable redistribution in the narrower sense. In the longer run, greater economic growth, and a more stable civilization, will help the poor most of all.



The point can be expressed as follows: 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, and sometimes redistribution of a different kind, namely growth-enhancing redistribution. (By no means do all of today's government programs actually redistribute to the poor, much less boost economic growth.) It will not, however, suggest that

a utilitarian or consequentialist approach is obliged to redistribute most of national income to the very poor. Nor are productive individuals typically morally required to spend most of their years serving the very poor.

There nonetheless remains a good case for some amount of redistribution. For instance a well-constructed welfare state can give the poor greater access to education and nutrition. The supported 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. Furthermore welfare systems make many higher income individuals feel good about their state and induce their political support. Many people want to have states whose benevolence they can feel good about. This framing of political benevolence contributes to state legitimacy and thus to public order. [1]

These arguments provide good reasons to support some investments in a welfare state—and yes we should describe them as investments—and thus to support some wealth redistribution. Furthermore they suggest an appropriate nature and scope for such redistribution, namely that we try to enhance sustainable economic growth.

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 often combined with static, insider-oriented labor market regulations. Furthermore if everyone approaches government looking for grants of money and resources, basic mechanisms of governance can break down, leading to rent-seeking, corruption and fiscal bloat. Alternatively, excess or poorly conceived welfare expenditures may create urban cultures of dependency and crime, which endanger social order. The empirical literature suggests that non-infrastructure government spending is correlated positively with lower growth rates, with the caveat that those results are measuring traditional gdp rather than some notion such as Wealth Plus, as discussed earlier.[2]



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 because they resent being taxed to pay for them. Yet immigration is by far the most effective anti-poverty program that has been discovered. So even if a specified set of welfare expenditures brings some growth benefits, alternative investments may do more for human welfare.

We can thus see some 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, on a one-time basis for the year. But if current yearly income were divided equally, average income would then fall rapidly, due to incentive effects and also the inability of many countries,

including the United States, to pay off its current debts. Over time the world's poor would fall into a deeper state of misery than they currently experience. The poor countries no longer would benefit much 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, most of all supporting immigration and also producing new technologies with global reach, such as cell phones or new methods for boosting agricultural productivity.[3]



[1] 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).

[2] 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.

[3] On utilitarian obligations, Scarre (1996, chapter VIII) offers a good survey. In addition to Smart and Williams (1973), see Rand (1967), Scheffler

(1982), Wolf (1982), Railton (1984), and Nagel (1986) for other critiques of extreme utilitarian obligations.

. . .

An aerial night photograph of a city skyline, featuring a prominent skyscraper with a red-tipped spire. The city lights are visible against the dark sky, and the water of a bay or harbor is visible in the background.

GROWTH RATES & LONG TIME HORIZONS

Growth Rates and Long Time Horizons

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 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 our chips.

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 different reasons. Utilitarianism tells us we should work, save, and innovate to serve the purposes of others, including future generations. Common sense morality tells us that we should work and save to take care of our families and because we value our own lives. These two perspectives remain different in their methods and their justifications. 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. After all, I favor pluralism, rather than utilitarianism or common sense morality per se. So we do not have to bring those two perspectives into complete accord. Instead, for the purposes of constructing a workable pluralism, it may suffice to know that two of the “kits in our toolbox” point in broadly compatible directions.[1]

Common sense morality also suggests that public and private ethical codes should differ in their advice and it may be possible to generate or justify this kind of recommendation as well. For instance a mother typically is considered justified in giving preference to her own baby, rather than tending to the babies of others. At the same time, governments should adopt a more impersonal perspective, taking into account the interests of all citizens (and to some extent residents) under their jurisdictions.[2] When allocating resources, governments should not favor one particular baby over another, at least not for the same reasons that a mother will favor *her* baby.



**“I’m bringing litigation against everyone
who is responsible for my overly demanding socialization.”**

It has remained an unanswered question why morality should be split in this fashion, namely between our private and public obligations. 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 some tools to defend such a bifurcation. This division of responsibilities stands a very good chance of boosting the long-term rate of sustainable economic growth. It’s pretty easy to see why parents do best when attending to their family priorities, and why governments do best when instituting a legal regime with some degree of impartiality and rule of law. The resulting proscribed behavior of both private citizens and governments then would spring from a common principle of growth maximization. The resulting principles would be roughly compatible with common sense morality and the bifurcation of private and public sector responsibilities. Again, I’m not arguing that this argument provides an exact justification of the details of common-sense morality as found in the world, with all of its quirks and custom-dependent variations. I’m just noting that the gap between common-sense morality and utility maximization, properly construed, is much smaller than it might at first appear.

We can test these hypotheses by looking for cases when a utilitarian or consequentialist should favor large-scale redistribution toward the very poor. For instance, the case for redistribution would be stronger if the world were going to end in the near future. If the time horizon is extremely short, the benefits of continued higher growth will be choked off and the scope for compounding over time would be correspondingly limited. The immediate returns to charity therefore should weigh more heavily in the decision calculus. To present this point in its starkest form, imagine that the world were set to end tomorrow. There would be little point in maximizing the growth rate and arguably we should just throw a party and consume what we can, after feeding the hungry that is.

Alternatively, the real return on investment might be permanently negative or zero. In this case compounding of returns would not operate, a long time horizon would not lead to a Crusonia plant, and there would again be greater reason to redistribute wealth. A high degree of redistribution also makes sense in a lot of “lifeboat” settings, where say a group of desperate individuals are afloat in the open sea, and wondering how to share the cans of spam. These examples typically have an implicit assumption of a zero or negative rate of return on investment, or perhaps investment simply isn’t possible in such an artificial setting. Again, there isn’t a Crusonia plant to run after.



“I’m rewriting my book as an epic novel of the sea.”

The attitude of *historical pessimism* thus is one of the most important critiques of my arguments. If historical pessimism is true, as suggested by many old school conservatives, expected rates of return are negative, there are no long-lasting Crusonia plants, and my arguments, even if they hold up logically, do

not very much apply. Historical pessimism is therefore much more than a mood or an attitude, it would shape our substantive views and our practical choice significantly, if it were true. That said, for the purposes of Crusonia plant arguments, positive economic growth need only be possible *with some probability*, an argument whose underlying structure recurs throughout these chapters. That positive probability means that the growth calculus will dominate our estimates of the expected returns from different choices. We therefore can reject the final practical stance of the historical pessimists, even while recognizing that they often get the better of the optimists in the substantive arguments about the future course of the world.

The current political opinions of social scientists do not always match up to the conclusions which I am suggesting. 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 improving the environment, all positions associated with the political left overall, at least in the context of the United States. They see a lower discount rate as providing support for all of these policies. Yet they also tend to favor redistribution, even when such policies conflict with economic growth. In this sense the political left does not have a consistent attitude about the importance of the future. Many thinkers on the right suffer from the opposite inconsistency. They often favor market-based discount rates, which are relatively high, but when the topic is redistribution they worry much more about the longer-term consequences.[3]

In contrast, I see a deep concern for the distant future as cutting across most current political spectrums. 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, because they require resource reallocations, namely many workers must try to find new jobs. Furthermore, market economies often 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 the United States gained ground on the Soviet Union, rather than allowing convergence.

[1] Hurley (2009) raises the question of whether consequentialism, even if it stipulates that some set of actions is good, can derive a strong obligation to perform those actions. In the synthetic view outlined here, that sense of obligation can come from common-sense morality, if need be.

[2] On this distinction, see Goodin (1995).

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



Who Should Sacrifice? And When?

Under more normal circumstances, with a long time horizon, 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.

To cite a simple example, Martin Luther King brought much good to the world, with respect to both justice and long-term economic growth. It is likely King did the right thing, rather than playing golf all day for fun, even though he lost his life in doing so. The same can be said for Gandhi. Nonetheless such obligations to sacrifice cannot be universal or near-universal. If we all went around sacrificing to an extreme degree, 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 cases our obligations should be viewed at a collective level. The framework does not pin down a uniquely correct course of behavior for each individual involved, and so it's not morally clear *which* individual is obliged to make the sacrifice. What if there is an innocent girl drowning in the lake, and any one of us could jump in and save her?

In these cases the question "What should I do?" allows for considerable latitude in answers and the scope of my individual obligation, as a group member, may be indeterminate. It's good enough if I make the sacrifice but it's also fine if someone else does instead. Similarly, it doesn't have to be *my* group which protests government injustice, because many other groups could do this as well. This becomes a problem of game theory, and as we know from game theory, very often the implied obligation for a single individual or group is indeterminate.



Imagine a game with payoffs, so that it is better if someone sacrifices, to achieve a socially valuable end, but it is worse if everyone sacrifices or tries to sacrifice to achieve that end. The structure of this problem is common to many

questions of morality and individual obligation, including the problem of global poverty. Some people should make sacrifices to help out but, because we must keep economically advanced civilization up and running, not everyone should make such sacrifices. Arguably this is the paradigmatic payoff structure to address questions related to global poverty, sacrifice, and obligation.

A utilitarian standard, in its simplest form, suggests that the “least cost supplier” should make such a sacrifice. That’s already moving us yet another bit closer to common-sense morality, namely by stipulating that the sacrificers should be the people most inclined to do so. That would include saints, moral saints, and dedicated agents of social change, as well as individuals who, for whatever other reasons, don’t find the required sacrifices to be so very daunting.

If several potential sacrificers face the same cost and can produce the same expected social benefits from sacrifice, it can be said that one of them (or more broadly, some subset of them) should sacrifice. There is then no fact of the matter, in utilitarian terms, who should jump into the water to rescue the child. In these cases it is morally indeterminate who should sacrifice and who should not or who should move to Malawi to administer the vaccinations to the children. That’s less oppressive, and less destructive of the integrity of individual life plans, than the typical utilitarian scenarios of enslaved Western doctors, toiling away their lives to suit purposes which are not their own.



When confronted with global poverty, or other forms of calamity, many (most?) people feel something like: “well, *someone* ought to do something.” At the same time, they do not obviously feel that this “someone” ought to be themselves. Such intuitions may seem incoherent or perhaps even selfishly irresponsible. But they match up to one of the realities of this game, as specified above. Yes, *someone* ought to do *something*, but the game itself does not make it clear who that someone is. The result is under-provision of the public good and that many people continue to make the assumption that, largely for selfish reasons, this “someone” is going to be someone else. The point is not that the observed response is socially optimal, only that the best solution to the game corresponds to some underlying features of common sense morality.[1]

In any case, individuals are more likely to sacrifice too little than too much, and too few individuals are willing to sacrifice much at all. So we can look to a specific recommendation of another kind, even if it cannot always be calculated *which* individual should perform a given sacrifice. 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, if it turned out to be best, all things considered, that they be the ones who should step forward. We ought to honor and reward such sacrifices more, to increase their likelihood, and again this is not so far from common-sense morality.

[1] Alternatively, a sacrificer could be specified by lot. Or we could look to game theory. We could think of morality as proscribing that individuals should play what economists call “randomized Nash strategies,” which will mean some probability of selfishness and some probability of sacrifice for each individual, as if we are rolling a die. A correctly specified randomization strategy will bring about the right amount of sacrifice—on average, at least—but again it won’t have everybody sacrificing. The proper randomized strategy will be the one which maximizes sustainable, expected global economic growth.

. . .



SHOULD MONEY BE REDISTRIBUTED TO THE RICH?

Should Money Be Redistributed to the Rich?

In many cases, purely utilitarian prescriptions will have morally counterintuitive implications, but running counter to the usual fears of enslaved American doctors serving poor Africans. 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. Indeed it is a common complaint in the literature on inequality that “the rich get richer,” while the “poor get poorer,” or at least more or less stay put. If this portrait is to be believed, it implies that the rich earn higher returns on their accumulated wealth, as indeed has been stressed by the French economist Thomas Piketty. If there is a trickle-down effect from the wealth of the wealthy, combined with a zero rate of discount, it is easy to generate scenarios where utilitarianism recommends redistribution to the wealthy.



For instance let's say—for purposes of argument—that the wealthy earn eight percent on their holdings, annually and on average, and the poor earn one percent. If a fifth of the gains to the wealthy, over time, trickle down to the poor, the poor are better off if the wealthy command more resources. They

will get a fifth of the eight percent, or 1.6 percent, rather than the one percent they earn on their own. Usually this trickle down won't come right away, but over time the rich will build more factories, buy more products, hire more servants, fund more research and development, push for more immigration, and so on. Sooner or later a lot of the poor will benefit. If we have a deep concern for the distant future, it matters less if a lot of these benefits come later.



The implications of this redistribution to the rich 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 consider the appropriate broader stretch of time, but they will appear anti-egalitarian by many common metrics, which focus on the short run only or which focus on a single country only, rather than taking an appropriately longer, broader, and more global perspective.

I am not suggesting that a good pluralist theory will, all things considered, endorse major, systematic redistribution toward the wealthy or the

talented. For one thing, this may be one case where a “rights constraint” limits the core recommendation of growth maximization. Maybe it’s just not right for Bill Gates to seize resources from a poor bricklayer, no matter how good a manager, investor, and philanthropist he may be. An alternative worry, also limiting redistribution towards the wealthy, is that a sufficiently unequal distribution of wealth may lead to lower growth through a number of channels discussed above.



What’s important, however, is how this reframing shifts the burden of proof by examining the implications of a very low discount rate. Direct, short-term redistribution to today’s poor is no longer the default option for a moral theory that emphasizes individual well-being. Instead, in many cases utilitarianism has to work to avoid the conclusion of redistributing more resources to the wealthy. Once again, it is possible to have a moral theory

which focuses on good consequences without requiring everyone to give up eighty percent of their income or to work as a doctor in Africa.[1]

Finally, we need to think carefully about where the most significant gains of the past have come from, and we should emphasize the extension of those gains rather than redistribution *per se*. Arguably the most important gifts of the past generation to the current generation come from wise investments, a belief in rules of just conduct, good political institutions, and good values, among other related historical factors. Growth-enhancing institutions do require hard work, but that investment is a positive-sum rather than a zero-sum game across the generations. Again, the moral intuition which results is the idea of strengthening good rules which are conducive to future economic growth, properly understood, and that is again not so far from common sense morality.

[1] For a discussion of the issues surrounding our obligations to save, see Rawls (1999, p.252).

. . .

A photograph of a hand holding a single autumn leaf with yellow and red hues. The leaf is positioned in the center, with the text 'OUR OBLIGATIONS TO THE ELDERLY' overlaid in white, bold, sans-serif capital letters. The background is a soft-focus field of many other autumn leaves in various shades of yellow, orange, and red.

OUR OBLIGATIONS TO THE ELDERLY

Our Obligations to the Elderly

Given the limits on our obligations to the poor, we will have comparable limits on our obligations to the elderly, using the same logic. In fact we can think of the elderly as individuals who are poor in one particular dimension, namely in their future human capital. The elderly are more likely

to die soon than are the young. And while we should do a good deal to help the elderly, the logic of sustainable growth places limits on these obligations too.

There is a more general question of how much a consequentialist or utilitarian standard should value the lives of the elderly. I recall giving a job interview in 1986 and being asked by my interviewer, the economist Julius Margolis, “why don’t we value human lives at replacement cost?” I was caught off guard and I didn’t have a good answer for him. Yet his challenge sounded so wrong to me. Do not the elderly deserve more respect than that? Do they not experience special lives, the value of which cannot be captured by metaphors of replacement? I’ve been thinking about that question for a long time because it challenges a lot of our moral intuitions.

Let’s start with a simple example. If a house is worth \$1 million in the marketplace but it can be built anew at a replacement cost of \$500,000, the correct value of the house is \$500,000, at least *provided we actually take the actions to replace it*. To make that concrete, we should not spend \$800,000 to save the house from destruction when we can replace it with an expenditure of \$500,000 on a perfect copy instead. We should spend only up to \$500,000.



So let’s say a human life is worth \$4 million, as estimated by standard economic “willingness to pay to reduce risk” measures, but we can create another human life for about \$10,000, say by subsidizing births or by saving another human life in a more economical manner elsewhere. It seems that birth subsidies are going to be cheaper than spending \$4 million to save a life. So how much should we spend to save or preserve that first human life?

Should we spend \$4 million? Or is this human life worth only \$10,000 or in other words its replacement cost?

Virtually all of us believe that \$10,000—it could be even less—is not in general the correct figure for valuing a life, including the life of an elderly person. For one thing, the life we (possibly) lose and the new life we would create are not the same person and they do not express the same individuality. So we would not be doing a strict replacement of what has been lost. In particular we may have special obligations to older individuals for their roles in raising us, building our nation, and perhaps defending us in earlier wars. Furthermore failing to save the first life and investing to help or create another life are not usually causally related. Letting an elderly person die, rather than spending to save that life, does not in fact automatically activate higher birth subsidies or other life-saving measures elsewhere in the economy.

So no, it is unlikely that measures for replacement cost are the correct way of valuing human lives. Still, there is some replaceability at the margin and replaceability should not be irrelevant for how we think about the values of lives. If one life disappears and another is added, the new life does make up for some of the value lost, at least in utility terms. Coming at this from the other side, losing an irreplaceable civilization is a much greater tragedy than losing a civilization in a way which allows for the birth of a new and different one in its place. Replaceability therefore seems to count for something, even if we do not agree for how much.



Most practically, the additional wealth from economizing on life-saving expenditures does lead people to buy safer cars, to take less risky jobs, and so on, so at some margin we will save some number of other lives by investing less in direct life preservation for the elderly. We don't know whether an increment of wealth saved will in fact rescue or preserve other human lives, but *there is some chance this might be the case*. And this possibility lowers the value of spending a lot of money to extend a human life, it does not raise that value. So, in a contemporary setting, a human life probably should be valued at less than \$4 million or whatever other sum the willingness to pay method, or some other utilitarian calculation, is going to serve up.

We probably do not know the exact correct valuation of an individual life, but we do know that the possibility of commensurability, the pull of the more distant future, the ongoing replenishment of human civilization, and the value of investing in future lives—when considered all together—exert some downward pressure on how much we should invest to extend the lives of the elderly today. My arguments therefore suggest lower estimates for the value of lives, including older lives, than will be derived from most other plausible frameworks, if only because replacement and replenishment of the civilizational flow are *considered as one factor among many*. Replacement and replenishment should not be taken as the final word, but yes they do exert downward pressure on our value of life estimates.



To put it most concretely, today in the United States we are spending too much on the elderly and not enough on the young. And given that the elderly are the ones who vote with greatest frequency, and the young often do not or cannot, that mistake should hardly come as a surprise. Governments should focus on investment, but in the United States at least government spending on investment is falling; currently government investment is estimated at about 3.6% of gdp, compared to a postwar average of about 5%; broadly similar patterns can be found in many of the European budgets.[1]

Unfortunately, when government spending needs to be limited or cut, very often investment is the first area to go and entitlements for the elderly remain intact. In this regard, I am suggesting some significant revisions to current trends.

[1] See Harding, McGregor and Muller (2013).

. . .



One Final Building Block: The Nature of Economic Growth

Now let's look at one final building block for deciding the appropriate scope of redistribution, namely the nature of economic growth.

At a very general level, we can see beneficial policies as falling into one of two categories. First, such policies may yield some benefit in once-and-for-all fashion. Imagine increasing the power of all light bulbs for one year. Second, the new benefits may be ongoing and self-augmenting. Imagine science policies which speed the rate at which better light bulbs (or other innovations) are discovered. Such policies would permanently increase the rate of economic growth, or in other words they would count as a Crusonia plant, namely a self-generating and self-refreshing source of ongoing value.



If we are examining a policy change, or an act of redistribution, it is important to know whether it involves an upfront, once-and-for-all benefit (or cost), or a systematic boost (or decline) in the growth rate over time. This somewhat arcane distinction, drawn from the economic literature on growth, is of great importance for adjudicating potential social changes.

To consider a simple 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). Many of these storms may come only after some time, but a greater concern for the

future means that we must pay heed to these consequences. More generally, many environmental problems hurt the prospects for long-run sustained growth. I am suggesting that such problems are especially important.

At the same time, counterintuitively, a concern for the distant future will militate against some environmental investments. In contrast to the ongoing storms, some of the costs of climate change appear to be “one-time” in nature, such as the costs of relocating coastal and inland settlements. The induced relocation might count as a rights violation of some sort, but from a consequentialist point of view these costs decline in importance precisely because they are one-time in nature. Once they are incurred, and the relocation is accomplished, from an economic point of view the story is over. When the rate of intergenerational discount is sufficiently low, maximizing the growth rate takes 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 compounding growth. So if we are approaching climate change as a serious problem, we should pay greater heed to the storms and lesser heed to the relocation costs, compared to what other frameworks will suggest. Most commonly, these differing kinds of costs are often lumped together without explicit attempt at differentiation, but that is a mistake.



They showed up for the cruise a little late.

It's very important whether a specified policy choice will affect the overall growth rate or not. The details of this distinction are complicated, however, because economic models provide differing accounts of which changes alter growth rates, as opposed to bringing one-off changes or improvements.

The most prominent economic approach to growth, the Solow model, is named after MIT economist and Nobel Laureate Robert Solow, who laid out the basics of the model in the 1950s. The Solow model, in its most basic form, 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.[1] In this model the primary way to increase ongoing growth is to induce a higher rate of technological innovation. Indeed many empirical tests have shown embodied technological progress to be a significant factor behind U.S. economic growth and many of these tests are run within the framework of Solow's work.


The Solow model helps us understand the phenomenon of “catch-up” growth, as has been so significant in East Asia. In the model the rate of return on capital diminishes as the capital stock increases. So when there is not much capital, the returns to investing and thus the incentives to invest are high, at least provided the labor force is of good quality and institutions will protect the property rights of investors. Given this approach, poorer countries should be expected to catch up to richer countries, as they borrow new technologies and increase their capital stocks to implement a lot of new and lucrative opportunities. Chinese economic growth, starting in the 1980s, had very much this catch-up flavor.



The Solow model also implies economies should recover quickly from one-time negative shocks, such as earthquakes or destructive wars. Although the capital stock has fallen from the destruction, the rate of return on capital is now higher, precisely *because* the capital stock has fallen. Additional savings should make up the gap and, over time, restore the economy to its previous growth path. That's another example of this more general phenomenon of catch-up.

What's surprising in the model is how many apparent improvements do not ultimately contribute to the long-run rate of growth. For instance a boost in savings and investment is seen as contributing to "transition growth paths" but not to "steady state growth" in the long-run. In the model, more savings means the catch-up happens more rapidly but it won't raise long-run growth. Alternatively, a decrease in wealth lowers the base on which growth occurs, and it slows down catch-up, but it also doesn't lower long-run growth. To use a biological metaphor, consider a lobster. If an arm is lopped off another arm grows to replace it. In the long run the lobster has the same capabilities, even if it is worse off along the transition path, namely while it is re-growing the appendage. In economic terms the mechanism runs as follows. The decline in the capital stock, as might result from a war or earthquake, raises the rate of return on capital, which induces more savings, which restores a higher capital stock. In the long run, the temporarily higher savings rate makes up for the destroyed resources. The very rapid recovery of Japan and Germany after World War II represents this mechanism in operation. In these cases, according to the theory, the rate of growth will remain lower only if the negative shock somehow permanently reduces the rate of technological progress.[2]

That's the Solow model, but I'm not arguing it is always the best model for understanding our world or for judging redistribution. It is simply one possibility.



[1] See Solow (1956, 1957); Romer (2000) provides a more recent summary.

[2] That 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), Lucas (1988) and others stress the role of human capital in boosting or maintaining the growth rate.



GROWTH BEGETS MORE GROWTH

Growth Begets More Growth

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 rates should continue to increase over time. Improvements beget further improvements and negative

events are likewise cumulative, thus the moniker “increasing returns.”

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. The first idea spreads, begets subsequent ideas, and so growth increases. Or there’s another way to look at increasing returns. Larger markets generate stronger incentives for ideas production, if only because the innovators can sell their product to a larger market (e.g., it would not be worth inventing the iPhone for New Zealand customers alone). That means large economies can grow more rapidly than small economies. The more the economy grows, the greater the incentive for subsequent new ideas and that in turn reinforces the incentive for growth. New ideas will lead to more growth, which in turns encourages more new ideas, and so on.




Increasing returns models are most commonly associated with the economist Paul Romer but they can be traced back to Adam Smith and the very beginnings of economics as a systematic object of study. 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 such as the United States, can drive an increasing returns to scale process.

To put these growth theories in my current terminology, increasing returns models suggest there are lots and lots of Crusonia plants, whereas Solow models indicate we have Crusonia plants only in policies which very directly and very specifically raise the rate of idea generation.[1] The increasing returns model that virtually any gain in resources can be translated into higher growth in the long run, rather than washing out in the adjustment process.

Under increasing returns models, a one-time negative shock harms the long-run rate of growth. Intuitively, we can think of the increasing returns concept as suggesting that resources multiply themselves at increasing rates. The larger the economy, the faster it will grow. Rather than losing the arm of a lobster, we have lost a colony of fertile rabbits or alternatively a Crusonia plant. Even if the rabbit 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 or limit each and every possible negative shock. The Solow model suggests a picture of greater resilience, since catch-up effects prevent each and every mistake from compounding over time into a larger collapse.

Increasing returns growth models will make us more wary about non-growth-enhancing wealth redistribution than will the Solow growth model. In the Solow growth model, the costs of the redistribution might be "once-and-for-all," rather than lowering the long-term rate of growth. We can make up for our temporary losses and eventually get back to where we ought to be. In contrast, under the increasing returns model any setback makes the economy smaller and thus limits future rates of growth, with significant implications for the standard of living in the distant future.


This distinction shows yet another way how traditional political debates should be redrawn. Individuals who believe in increasing returns models should be much more skeptical of non-growth-enhancing redistribution than individuals who believe in the Solow catch-up model.



More generally, it is a central question whether the logic of the Solow model or the logic of the increasing returns model holds. Even if you don't buy into all the details of these models are stated, the two core options are still that one-time costs do matter a lot in the long run or that they don't matter a lot in the long run. I'm treating the models as stand-ins for those two broader views. Maybe you have some other favorite growth model but we must in some way return to the key question of whether gains and losses compound or whether they dwindle away into longer-run insignificance.[2]

Whatever your exact view on the Solow and increasing returns models, the logic of the increasing returns model will likely carry significant weight in our final evaluation. In many cases our best answer, given current knowledge, suggests that a given cost brings some probability of an ongoing growth effect (as with increasing returns) and some probability of a one-and-for-all adjustment cost, followed by catch-up (as in the Solow model). In our expected value calculations, this will operate as an expected impact on the long-term rate of economic growth. Therefore we should incorporate the logic of the increasing returns model into how we evaluate social changes, even if the increasing returns model is not our single best current theory of economic

growth. In expected value terms, most of our social choices have an impact upon future rates of economic growth. *Crusonia plants are everywhere*, again in expected value terms. We are making decisions about Crusonia plants all the time.



Finally, both the Solow and Romer models emphasize ideas as the wellspring of economic growth. New ideas, of course, are the product of human reason and it is Aristotle who defined man as the rational animal. A preoccupation with pursuing growth—or some modified version of the growth ideal—thus means a preoccupation with ideas, a preoccupation with cultivating human reason, and a preoccupation with the notion of that man should realize, perfect, and extend his nature as a generator of powerful ideas which can change the world. Cringe all you wish, but on this point I'll send another

credit along to Ayn Rand, who stresses this point even more than do most philosophic rationalists. If we are pursuing self-sustaining and self-generating bundles of plural values, one way or another we are paying homage to the power of human reason.

[1] 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*.

[2] 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. The neo-institutional models have very real merits in explanatory terms but for normative purposes they do not remove the importance of taking a stance on the long-run growth effects of a policy change. On neo-institutionalist approaches, see Douglass North's work on American and European economic history (North 1981, North and Thomas 1976); see also Olson (1984) and Bates et.al. (1998) and Acemoglu and Johnson (2004).

. . .

SCIENCE FICTION, HOPE AND THE EPISTEMIC CRITIQUE



Chapter six: Must uncertainty paralyze us?

Science Fiction, Hope and the Epistemic Critique

When I was a kid I loved science fiction stories. I loved to think and speculate about how things could be totally different from how they were. There was one story in particular which intrigued me and I encountered its premise in a few different books, for that matter comic books. It still appears in popular culture, and if you don't like science fiction try Gwyneth Paltrow's 1998 movie *Sliding Doors*, which asks how much a single life can be altered by the simple act of missing a train.

It's pretty simple. If you can somehow manage to go back in time, and alter one small event, the entire history of the world can change. One extra sneeze from one caveman, millennia ago, probably would overturn everything we know. Ray Bradbury's short story "A Sound of Thunder," published in 1952, is

one of the early sources for this idea. It seems a little crazy, but the more you think about it the more it seems to hold.



The key point is that small changes usually turn into big changes. What if John bends down to pick up a banana peel? If nothing else, this action will likely affect the identities of all his future children, if only by changing the timing of future sex acts by a slight amount, or by reconfiguring the position of John's testicles and thus altering his sperm production. A different set of people born into the future will, most or all of the time, cause the world to take a different path. Even if most people "don't much matter" for broader aggregate outcomes, it sure seems that some do, such as Jesus or Hitler or Lenin or Chairman Mao. For instance, without Hitler Nazism probably would not have succeeded or had the same impact on the world stage. The Second World War, as we know it, would not have happened, nor would have the Holocaust occurred. Virtually every country's subsequent history would have been different and we would end up with a quite different world history for the rest of humanity's time on earth.

That's what gives so much power to the people who go back in time, in these stories at least. They can't predict or control how they will change the future, but they can indeed change the future by tapping on the hinges of some extreme contingencies.



In some cases we might think that aggregate outcomes are stable to small perturbations. The long-run random effects will sometimes cancel out or offset each other; Tolstoy went so far as to suggest that the “great men” of history had no impact at all, citing Napoleon as someone whose acts were later undone by his successors. Or some physical systems may be stable with regard to small perturbations; for instance getting through one traffic light may just mean you wait longer at the next one.


That said, we cannot count on all such perturbations to cancel out or diminish in importance, if only because even a small act reshuffles the entire future

genetic history of humanity. Even on a prosaic level, Napoleon's acts changed the course of life in Germany, which underwent a liberal intellectual revolution because of the French invasion and then modernized, or Egypt, which received the printing press, a large dose of liberal ideas, and then finally started to resent European interference with local affairs, a feeling which to this day has not vanished. The histories of these regions were changed irrevocably, as were the histories of the Jews who were liberated under Napoleon's rule.

For small changes to translate into large final effects, we need only postulate that some individuals, or some leaders, play a significant role on the global stage. Even if most individuals do not matter, or most small changes wash out, some of the small changes today will alter future identities, once we look a generation or two into the future. So the argument requires only that a very small number of personal identities matter for the course of history. If Hitler's great-great-grandmother had bent down to pick one more daisy, thereby postponing her next act of intercourse ever so slightly, many of the effects of that delay might have washed out but Europe today would be a very different place.

Following the philosophers, I refer to this as the epistemic problem.[1] The epistemic problem isn't really about time travel at all, even though science fiction constructs allow us to visualize the problem in an especially vivid manner. The real issue is that we don't know today whether our *current* actions will in fact give rise to a better future, even when it appears they will. If you ponder these time travel conundrums enough, you'll realize that the effects of our current actions are very hard to predict, and that has nothing to do with whether time travel ever comes to pass.

The epistemic critique suggests consequentialism cannot be a useful guide to action because we hardly know anything about long-run consequences. While it is true we can calculate expected value, such a calculation is typically based on a very limited range of information about present consequences or consequences in the near future. Like many philosophers, I wonder if we have the correct moral theory in the first place, if we cannot know ninety percent, or perhaps 99.9 percent, of what is to count toward a good outcome.



Maybe you think it is a good rule to teach teenage drivers not to plough through the yellow light. After all in the United States about 40,000 people die in auto accidents every year. But today, when such a driver stops at a yellow traffic light, rather than accelerating, he likely affects the length of other commutes and thus changes the timing of millions of future conceptions. Subsequent genetic identities will change as well. Come the next generation, these different identities lead to different marriage patterns and thus an entirely new set of individuals in the future. So how can we really tell if our yellow light rule is a good one? Aren't we operating in the dark? If you think about these conundrums for long enough, you wonder how we can ever judge good consequences at all.

Once you start worrying about the epistemic problem, you may fear the onset of an extreme moral nervousness. Virtually every action would appear to have enormous consequences for our future. You would feel that maybe, just maybe, you had caused the painful deaths of millions when you speed through a green light; but cheer up, you probably have saved millions of others as well. All those lives rested upon your decision. Any moment most of us might be doing something that will lead to truly wonderful results, truly terrible results, or most likely a mix of both at the same time. It seems paralyzing. If you internalized that feeling, all of life would be like walking around on eggshells, except that the eggshells are geopolitical changes, possibly worth millions or even billions of future human lives.



**“Next time, don’t start the presentation by asking,
‘Can you tolerate ambiguity?!’”**

I know that some of you may be thinking this argument is just plain, flat out stupid. But keep your patience. I’m not here to defend nihilism or to suggest that ethics should focus on the paradoxes of time travel and the timing of conception for the future Hitler. I’d like to defend a version of common sense morality, but I do want to look a little more closely at why the epistemic critique does not imply hopelessness about trying to make the world a better place. If we have that understanding, we’ll see that some versions of common sense morality work better than do other versions, and we will move closer to those forms of common sense morality. That will have some concrete implications. We’ll also see some new arguments for some of the positions which I’ve already staked out. In the meantime I’m simply suggesting that we should take seriously the problems with the dogmatic assertion that we absolutely *know* we are doing the world good.

By the way, the difficulty of calculating consequences becomes a much more important issue when we maintain a deep concern for the distant future. At sufficiently high rates of discount, most future consequences cease to matter much within a thirty-year horizon or so. But with sufficiently low rates of discount, welfare-relevant consequences continue until the end of the world. We cannot dismiss the importance of the future simply because it is distant from us in time and thus we need to worry about epistemic problems all the more.

So, to proceed, I'm going to step back and consider whether epistemic problems upset the entire consequentialist framework. At the same time, I'm going to revisit some questions from earlier chapters about Crusonia plants. Might the epistemic problem and the importance of Crusonia plants have some underlying connection? I'll be returning to that question, but first I will work through some examples about the radical uncertainty of the future.

Finding cases where consequences clearly should matter

Let us start with a simple example, namely a suicide bomber who seeks to detonate a nuclear device in midtown Manhattan. Obviously we would seek to stop the bomber, or at least we would work to reduce the probability of a detonation. I will treat this example as standing in more generally for choices, decisions, and policies that affect the long-term prospects of our civilization.

If we stop the bomber, we know that in the short run we will save many lives, avoid a massive tragedy, and protect the long-term strength, prosperity, and freedom of the Western world. For this same reason, it should count as a significant benefit if we can lower his chance of succeeding. Reasonable and moral people, regardless of the details of their ethical and meta-ethical stances, should favor stopping the bomber.

No matter how hard we try to stop the bomber, we are not, *a priori*, committed to a very definite view of how effective prevention will turn out in the very long run. After all, stopping the bomber will likely reshuffle future genetic identities, and may imply the birth of a future Hitler or maybe an even worse tyrant, this one armed with nuclear and biological weapons. Even trying to stop the bomber, with no guarantee of success, will remix the future in similar fashion, for reasons I've discussed. Still, there is at least a slight chance (and maybe a very definite chance) that stopping the bomber will favor good consequences, even in the longest of runs. To put it simply, it is difficult to see the violent destruction of Manhattan as on net, in **ex ante** terms, favoring the long-term prospects of the world.



We can imagine scenarios where the destruction of Manhattan works out for the better **ex post**; perhaps, for instance, the explosion leads to a powerful anti-proliferation movement, which turns out to save the entire world in some longer run. Maybe so. But we would not breathe a sigh of relief on hearing the news of the destruction of Manhattan. We would not think that world is now on a path toward future salvation. More likely we would fear subsequent attacks on other major Western cities, the emboldening of terrorists, or broader panics and perhaps an economic collapse. There would be a new and very real gateway toward global chaos and tyranny. The future is remixed radically, but we would not rationally believe that some new positive option has been created to counterbalance the current destruction and the new possible negatives. Furthermore the negatives from the destruction could continue for decades or even centuries. The memory of such an event, and its immediate consequences, would not fade soon.

If this example does not convince you, consider the value of stopping a terrorist attack that would decimate the entire United States. Or consider an attack that would devastate all of Western civilization, or the entire world. What if a cosmological disaster destroyed 99.999 percent of all intelligent life across the universe? At some point we can find a set of consequences so significant that we would be spurred to action, without much epistemic reluctance, even though we would be recognizing the broader uncertainties of

the very long run. Surely at some point the upfront benefit must be large enough to persuade us to pursue it. We can debate “how large” an upfront event is needed to sway us toward making an actual evaluation and recommendation, but a large enough upfront event should suffice.[2]

[1] Lenman (2000) provides one clear statement of the epistemic critique of utilitarianism and cites some precursors. For additional perspectives on the epistemic critique, see also Norcross (1990), Frazier (1994), Howard-Snyder (1997), Dorsey (2012), Mason (2004), Lang (2008), Burch-Brown (2013), and Kolmar and Rommeswinkel (2013). Hayek (1991) can be thought of as offering a version of the epistemic critique as well. Shelley Kagan (1998, p.64) calls this epistemic argument “the most common objection to consequentialism.” William Whewell provided one early statement of the epistemic critique, see the discussion in Mill (1969 [1852]). Cowley, Ambrose, and McCullough (2000), writing as historians, consider a series of “what if?” questions, in the context of military history. Moriarty (2005) considers the implications of the epistemic argument for traditional concepts of desert. My initial crack at thoughts on these topics is Cowen (2006). MacAskill (2014) considers how we ought to maximize across the expected values of possibly conflicting moral theories with possibly conflicting conceptual frameworks.

[2] The economic literature on probability offers a debate on whether we can ever say we have “no idea” about the likelihood of an outcome. Under one view, we can always attach a Bayesian probability, whether explicit or implicit, to various outcomes (Caplan 1999). Even if we are very uncertain, in principle there exist betting odds that we would or would not be willing to take on a given choice. These counterfactual bets help us pin down implicit probability estimates for any imaginable outcome. Under a second view, we simply cannot assign probabilities at all to some events (O’Driscoll and Rizzo 1996). Those events are unique and “radically uncertain” and thus do not fit into the standard categories of probability theory. But even in these cases we still have degrees of uncertainty. I may have “no idea” about my forthcoming birthday surprise, but this uncertainty is not comparable to my “no idea” about intelligent life on other planets. Background social context will give us some expectations, even if we cannot assign definite numbers to probability forecasts. For an earlier look at the argument that we often have some idea about consequences, see Mill (1969 [1852], p.180).



“STOP THE NUCLEAR BOMBER IN
MANHATTAN.
ABOUT THE DOG’S LEG I COULDN’T SAY”



“About the Dog’s Leg I Couldn’t Say”

We therefore can avoid complete paralysis or sheer and absolute agnosticism, *at least for some of our choices*. No matter how high the uncertainty surrounding our estimates of subsequent consequences, we can take some actions to favor good consequences in the short run.

It is only necessary that those short-run good consequences are of sufficiently large and obvious value. We can recognize the subsequent radical uncertainty, but still the upfront benefit can be large enough for us to proceed with some hope—albeit not absolute confidence—that matters will work out for the better in the longer run.



If you look at the epistemic critique, it does not focus on the pursuit of large, upfront benefits. Instead the articles in this philosophic literature often specify very small, indeed “squirrely” benefits. There’s a reason for this—namely in those cases the epistemic critique has greater weight—and so let’s look at these arguments in more detail.

James Lenman, a philosopher and a commentator in this literature, doubts the importance of consequences as an account of moral rightness. Lenman’s arguments are interesting, but I think that, properly understood, they strengthen the case for rules-based, “big picture” thinking about consequences. Let’s first go through Lenman’s arguments and then we’ll return to what the whole mess might mean.[1]

One of Lenman’s central intuitions involves a D-Day example. Imagine we must decide which French beach to invade to fight Hitler. Our decision might have monumental consequences for the future of civilization. Imagine further that we have two potential candidate beaches for the invasion, but no particular military reason to favor one beach or the other. One beach is probably much better than the other, but when making the choice we do not know which one will be better. We do know, however, that if we land at beach A, we will cause inconvenience to a dog. As Lenman constructs the example,

the dog will experience a broken leg, but the inconvenience could be less painful to make the contrast with the human consequences all the more glaring. (How about a slight sprain for the dog?) If we land at beach B, no dog is injured.



“How dare you say I am dogmatic!”


In most plausible moral theories we attach some weight to the suffering of animals. It is bad if the dog on beach A must suffer. Yet Lenman suggests that the fate of the dog provides at most an extremely weak reason to favor landing at beach B rather than beach A. The fate of the dog is tiny relative to what stands at stake in the comparison between the two beaches. According to Lenman, the real comparison, if we knew it, would swamp the factor of the pain of the dog. In the meantime, our ignorance of the true relevant factors should not allow us to elevate one small factor (the dog) in importance, when that factor would otherwise become negligible in the final calculations. In Lenman’s view, our uncertainty about beach B vs. beach A does not cause two those options to cancel, leaving the fate of the dog to tip the scales one way or the other. Instead, he wonders whether we can evaluate actions in terms of their consequences at all. As you can see, this argument is one version of the epistemic critique.

The hard line response, of course, dismisses Lenman’s intuition rather than responding to it. We can imagine the extreme consequentialist yelling out:

“Save the dog from a broken leg, damn the uncertainty, the potential variance of outcomes from the invasion decision is high in any case.”

But that’s not my answer. I’m willing to accept there is something to Lenman’s basic point, at least as outlined in the example at hand. My reply is this: “Stop the nuclear bomber in Manhattan, about the dog’s leg I couldn’t say. Maybe Lenman is right and this D-Day case is up for grabs.” We are then left with the view that consequentialism is strongest when we pursue values which are high in absolute importance. You can debate where to draw the line between the bomb in Manhattan and the dog’s leg, but once a distinction is granted between cases which differ in terms of size of costs, we have something to work with.

The use of a dog’s broken leg, as the relevant cost, is designed to be murky. To be sure we may find merit in the notion of animal welfare. But do we really know how to weight the welfare of dogs against the welfare of humans, or against other species? We start off being uncertain about the value of a healthy leg for the dog and then the epistemic critique intensifies this uncertainty. We wonder something like: “given the possible variance of outcomes at stake, could the dog’s leg really be so important as to sway the issue”? This entire example, by focusing on the dog’s leg, a relatively small and also potentially ambiguous value, gives the epistemic critique the appearance of more power than it merits.



“Let’s round it off to 2 and go to lunch.”

To put those questions into stark relief, consider another moral dilemma, again involving an invasion decision. Instead of preventing the broken leg of a dog, assume that one beach invasion will kill one hundred innocent civilians. The other beach will involve no comparable danger. By assumption, these civilian lives have nothing to do with the final outcome of the war, but one hundred innocent lives carry greater moral value than did the dog.

Obviously we should choose the landing that has no chance of killing the civilians, given that we have no other reason for favoring one beach over the other. Consequentialism gives us a straightforward answer, even though our actions will set off unforeseeable long-run effects of massive long-run importance. Once the upfront benefit from one course of action is sufficiently high, the epistemic critique has less force, though we remain uncertain as to whether we will choose the correct beach for defeating Hitler. We remain uncertain about the long-run “remixing” effects of our choice. Still, we must pursue large benefits when we can—like the one hundred lives—at least provided there is no good reason not to.

Now let's look at some additional intermediate cases and see what happens. We'll find further support for the notion that a modified form of consequentialism, focusing on large benefits and costs, does fine when faced with the epistemic critique.

[1] Lenman (2000) appears to favor "ethical theories for which the focus is on the character of agents and the qualities of their wills, for theories that are broadly Kantian or Aristotelian in spirit."

. . .



Aesthetics, Goodness, & **MORAL** REASONING

Aesthetics, Goodness, and Moral Reasoning

The epistemic critique may be drawing on a different moral principle altogether, and this is in principle which pops up frequently in pluralistic approaches.

Let us consider what I call *The Principle of Roughness*: “Some of our choice options will differ in complex ways. We might nonetheless, ex ante, make a reasoned judgment that they are roughly equal in value, and that we should be roughly indifferent across the two options. After making a small improvement to one of these choices, we still might be roughly indifferent to which option is better.” Philosophers often write of the closely related ideas of incommensurability and incomparability.



We often resort to The Principle of Roughness in **aesthetics**. Assume we are trying to judge whether Beethoven or Mozart is the better composer. We might judge the two composers as being roughly equal, or judge that neither composer can be elevated over the other. Assume then that we discover one new work by Beethoven, a lovely two-minute bagatelle for piano. We are not now obliged to assert that Beethoven is the better composer. Our original judgment of equality was sufficiently “rough” that it can survive this new discovery. In contrast, a very exact comparison of equality, such as that of weight or length, could be upset by a small change at the appropriate margin of measurement. For this reason, The Principle of Roughness seems especially likely to apply to aesthetic comparisons.[1]



We also find The Principle of Roughness in some of our judgments of **goodness**. We might, for instance, be choosing between a new health care program, and a new poverty reduction program, each with significant and complex benefits for different groups of people. We might judge that the two policies are roughly equal in value. We might then discover that one of the policies was slightly better than we previously had thought, perhaps because it cost a thousand dollars less than expected, due to an initial calculation error. Yet we would not then have to declare that the now-cheaper policy suddenly is clearly better than the other policy. It is still likely, by construction of the example, that the two policies are roughly equal in value. Of course at some cost differential, this judgment of rough equality no longer will hold. We might say the programs are roughly equal in value if one turns out to be a dollar cheaper, but not if one turns out to be \$100 billion cheaper.

I would not defend The Principle of Roughness when we have very exact information and measurements and a well-defined, single dimensional standard of what counts as good. But how often is that the case? Our judgments, as we find them in the real world, very often seem to be rough, if only for practical reasons such as lack of good data and the incomplete nature of moral reasoning.

The Principle of Roughness, to the extent it applies, implies we should discriminate against relatively small benefits and losses. For this reason The Principle of Roughness may be operating in some instances of the epistemic critique. The future changes at stake—all of the remainder of human history being up for grabs—seem so large that relatively small changes in upfront

benefits and costs, such as the dog's leg, do not matter much and do not move the comparison out of the category of the blurry.




In most applications of The Principle of Roughness (e.g., Mozart vs. Beethoven), small changes (e.g., discovery of an extra sonata) are swamped by high absolute totals (of achievement) in the first place. In the D-Day example, the small change—the dog's leg—is swamped by the high variance in our estimates of consequences. In other words, the epistemic critique extends one version of the Principle of Roughness to comparisons involving uncertainty.

The Principle of Roughness, however, does not refute consequentialism or even downgrade the importance of consequences. Instead the Principle, like the arguments presented above, reflects the importance of pursuing *large*

benefits. The rougher the initial comparison, the more that large additional benefits would be required to make the new comparison a clear one.

To borrow a metaphor, anything we try to do today is “floating in a sea of long-run radical uncertainty.” Only big, important goals will, in reflective equilibrium, stand above the ever-present froth and allow the comparison to be anything more than a rough one. When small goals are at stake, our moral intuitions become confused—properly or not—and as a result we downgrade the importance of those small goals. If there is any victim of the epistemic critique, it is focusing on small benefits and costs, but not consequentialism more generally. If we bundle appropriately and “think big” and pursue Crusonia plants, our moral intuitions will rise above the froth of long-run variance.

Critics of consequentialism would like to establish something like the following: “We find it hard to predict consequences. Therefore consequences do not matter very much, relative to other factors, such as deontology or virtue ethics. We should abandon consequentialist morality.”



But epistemic considerations have yet to produce a strong argument for this view. The arguments instead support the downgrading the importance of minor consequences and the upgrading the importance of major consequences. We should not count small events for nothing, but epistemic issues, and The Principle of Roughness, lower their importance, that is relative to highly significant benefits and costs. In other words, yet another part of moral theory is directing our attention toward the pursuit of Crusonia plants.

[1] On these ideas, see Chang (1997, 2002). On the related concepts of “vagueness,” see the work of Timothy Williamson, for instance “Vagueness in Reality,” (2003).

. . .



HOW TO BE **MOST LIKELY CORRECT**

How to Be Most Likely Correct


What are the practical implications of these arguments? The arguments above have (at least) two practical implications for what we should believe, how we should believe, and how we should act. I will consider agnosticism and individual rights in turn.



1. How to be a good agnostic

We should be skeptical of political ideologies that suggest they know how to get everything right in terms of means-ends relationships. After all, why should we be so sure that our favored ideology will in fact bring good consequences? Given the radical uncertainty of the more distant future, we do not usually have a good idea how to achieve our preferred goals over longer time horizons. Our attachment to particular means therefore should be highly tentative, highly uncertain, and radically contingent.

Our specific policy recommendations, though we believe them to be the best available, will stand only a slight chance of being correct. They ought to stand the highest chance of being correct, of all available views, but this chance will not be very high in absolute terms. We should think of the details of our political views as analogous to betting on a slightly crooked roulette wheel, designed to land on the number seven more than a proportionate amount of the time. We should bet on the slightly favored outcome, namely the number seven, and by doing so we improve our prospects. But most of the time we are likely to predict the wrong number, as we will be betting on seven and some other number will come up.



Our political stances and policy recommendations should be accordingly tolerant. Imagine a world where your chance of being right is two percent, and your chance of being wrong is ninety-eight percent. Each opposing view, however, is right with a chance of only one and a half percent, slightly less than the chance for your view being right. Furthermore, if any one of these other people is right, and you are wrong, your view will have grave negative consequences, such as bringing the premature end of civilization. That is, your view has grave negative consequences with probability 0.98. In this scenario, how intellectually arrogant should you be about the details of your beliefs? Above and beyond your basic loyalty to achieving good and significant ends, how firm should your dogmatism be about means-ends relationships?

We should not pat ourselves on the back and feel that we are on the correct side of an issue. We should choose the course that is most likely to be correct, but at the end of the day we are more likely wrong than right. Our particular views, in politics and elsewhere, should be no more certain than our assessments of how to play that roulette wheel. With this attitude, political posturing loses much of its fun and indeed it ought to be viewed as disreputable or perhaps even as a sign of our own personal delusions.

2. Why the case for rights is compelling and which rights are the important ones?

The epistemic critique also helps us understand why we should respect individual rights rather than overturning them in favor of better consequences. They also help us outline the limits of those individual rights.

Let us consider, for instance, the right of an innocent baby not to be murdered. Let's say you believe in such a right, as I do, and then you are presented with a counterexample where killing that innocent baby will, in the short run, raise national income by \$5 billion. Normally, economists would value a life at much less than \$5 billion, typically in the neighborhood of about \$5 million, which is a big difference. Yet in this instance it is wrong to set up the comparison of "baby's life vs. \$5 billion" and then have to choose. The correct comparison is "baby's life vs. a froth of massive uncertainty with a gain of \$5 billion tossed in as one element of that froth." When it is phrased that way, it is easier to side with preventing the murder of the baby. There is even a good chance—albeit a less than fifty percent chance—that stopping the murder of the baby will be good for gdp too.



**"Well, last year I kicked gambling ...
the odds are 3 to 1 the New Year will be a good one."**

In other words, very often rights do not conflict with consequences in the simple ways set up by philosophic thought experiments. And so now we can see a shift in how we think about radical uncertainty concerning consequences. Rather than letting it paralyze us, think of the uncertainty as giving us the freedom to act morally, without fear of the certain knowledge that we are engaging in consequentialist destruction.

We also can see this radical uncertainty as supporting a new “enchantment” of human life and choice, one where most or all of our actions will have consequences we cannot possibly predict. On average these consequences will be positive, just as average economic growth is positive, but we will always be wondering which future consequences we have set in motion. We will wonder about our strange and almost magical powers in this regard. For all the confusion we should feel about the marginal product of an individual act, this is also an empowering notion and it relates to the idea that all fruitful societies are based on some notion of faith. In this case we can hold on to our faith in doing the right thing, and indeed doing the right thing for its own sake, without being brutally beaten back by the knowledge that we must be bringing a consequentialist disaster upon ourselves.

Now let’s consider another crazy philosophic thought experiment, namely that either we murder a specified baby or aliens from Alpha Centauri destroy the entire earth. Here we have more room to be what is called a “rights consequentialist.” Murdering the one baby is wrong, but if we don’t do it, by construction of the example, even more babies will be murdered, many millions more in fact. On one side of the equation, we have “murder one baby.” On the other side of the equation we have “all babies get murdered, everything on earth is lost, and, since human life on earth ends, no more Crusonia plants, no significant froth of uncertainty to follow.”


It is possible to see why we probably would opt to kill the baby in such a case. For one thing, the cost of not murdering the baby is now much higher. For another—and this is significant—if the entire world ends there is no residual uncertainty about what will happen next. Arguably we should pursue the better consequences and there is no remaining “froth of uncertainty” to justify sticking with the rights of the individual baby. And in that sense the notion of rights postulated is not strictly absolute against all possible external consequences, as might be dreamt up by philosophers. Still, most or all of the hypothetical examples where rights should be violated are not very relevant to the real world choices we have had to face so far.



To be sure, we do not know exactly where this comparison should end and at what point the case for murdering the innocent baby looks strong. What if the choice is murdering one baby or the aliens destroy only the nation of Nepal? The number of lives affected by the alien threat could be made larger or smaller until the rights theorist cries the proverbial “Uncle!” We don’t know exactly where to draw that line and you can think of this as one of the weaknesses of the rights theory I am putting forward. Still, we have a rights theory where rights are indeed absolute, at least provided the examples we consider match some very brute and basic facts about the real world, for instance the existence of Crusonia plants and the froth of consequentialist uncertainty.

Using this understanding of the epistemic problem, we also can see some further micro-foundations for why “lifeboat ethics” might differ from our more usual and more practical ethical recommendations. I define lifeboat ethics as the ethics which should hold as the end of the world—or the end of

some sufficiently segmented part of the world—approaches. People in a (not-to-be-rescued) lifeboat cannot look forward to great improvements in their future welfare or much economic growth. The sharks are circling and they expect their supplies of food and water to run out. By construction of the example, the lifeboat is not connected to the broader froth of uncertainty in the world at large.



So what does that mean? In lifeboat settings, the benefits at stake typically will be small precisely because lifeboats, even the relatively large ones, are small. Rights therefore acquire greater force in relative terms. No matter what you do, you can't produce large social benefits in lifeboat examples, and so there is a stronger case for simply doing the right thing. Don't toss that weak guy overboard or cook his flesh, it is the wrong thing to do and there is only so much gain to be had from it. Making an omelette may require the breaking of some eggs, but here the omelette is so small and also not very tasty, so I say leave those eggs intact. Once again, the case for human rights is stronger than it appears at first.

. . .



SUSTAINABLE GROWTH, FRAGILITY, & *Our Future*

Conclusion: Where have we landed?

Sustainable Growth, Fragility, and Our Future

Some of my arguments have concerned the theme of distance. I have claimed that the inhabitants of the future are less distant from us, in moral terms, than many other views would indicate. Therefore we should take more


seriously the implications of our choices for that future.

That greater concern for the future induces us to rethink a variety of moral questions, including the importance of economic growth, the best kinds of redistribution, and also it makes the stability of social systems a higher priority, among other considerations. The previous chapter on agnosticism and radical uncertainty made an additional argument for being willing to think big. If our values are to rise in importance above the froth of long-run uncertainty about the effects of our actions, we must look to relatively large and important values. That again returns us to a sphere where Crusonia plants reign supreme and sustainable economic growth is all-important.



Here is my short, three-point summary of where these arguments have brought us:

First, believing in the overriding importance of sustained economic growth is more than philosophically tenable and it may be philosophically imperative. We should pursue large rather than small benefits and we should have a deep concern for the more distant future, rather than discounting it exponentially. Our working standard for evaluating choices should be to increase sustainable economic growth, because those choices overcome aggregation problems and they are decisively good. That provides us with a broad quantitative proxy for the long-run development of human civilization, and it constitutes one means of finding and promoting co-moving plural values.



S econd, there is plenty of room for our morality, including our political morality, to be strict and based in the notion of rules and rights. We should subject ourselves to the constraint of respecting human rights, noting that only semi-absolute human rights will be strong enough to place any constraint on pursuing the benefits of a higher rate of sustainable economic growth.

At the end of that tunnel we have not “The Best Ethical Theory,” as a philosopher might wish to derive, but rather some good decision rules to live by and also some standards for how we might imagine a much brighter future.

To be sure, I still have made no attempt to derive the existence of rights. Still, once we are thinking naturally in terms of big, packaged changes, a belief in rights fits in quite naturally. We have some rules for what to do—maximize sustainable growth—and other rules—rights—which place some constraints on those choices. In other words, the lower-order rules stand within some higher-order rules, namely respecting the rights. Across the entire map we should stick to our chosen priorities, and our chosen rules, rigidly. Rather than rights and consequentialist considerations representing warring or contrasting approaches to philosophy, the door is open for consilience and compatibilism to reign, and that is all within a rules-nested approach to thinking about the both the rights and the practical, consequentialist side of the equation.

We need not defend such rules-based perspectives on the grounds that they are a highly practical “noble lie.” It is nice to see the practical benefits of rules

recognized, but the “noble lie” approach is too cynical. It is assuming that rules are a philosophical weak sister to begin with, when they are not. So rather than seeing belief in strict rules as a noble lie, it just might be part of a very important noble truth.



That said, on the purely practical side it is unlikely that democratic real world decision-makers will think too big. The nature of politics is more likely to produce too much small ball rather than too little. Policymakers often make decisions on a day-to-day, case-by-case basis, simply hoping to survive the

next election cycle. I don't mean that as a cynical criticism of human nature, rather it is an implication of political competition with relatively short electoral cycles, relative to the time horizons over which policy matters. (The U.S. House of Representatives has a two-year voting cycles for policies which may have effects over twenty or thirty years, or in the case of environmental policies longer yet.) That is one reason why our politics makes as many mistakes as it does. So from a purely pragmatic or consequentialist point of view, there is an additional argument in terms of imagining our political choices in terms of broader bundles of choices and rules, at least provided we get the rules right.

Third, we should be very cautious in our attitudes about specific policies. Even if we succeed in taking true aim at what we think are the best courses of action, the chance that we are right on the specifics—even if as high as possible—still is not very high. It's like trying to guess at the origin of the universe. The best you can do is to pick what you think is right with 1.05 percent, rather than siding with what you think is right at 1.03 percent. Most likely you're wrong—even if some others are likely to be even more wrong—and thus your attitudes should be correspondingly modest in the epistemic sense.

For some more concrete recommendations, I'll suggest the following:

- a. Policy should be more forward-looking and more concerned about the more distant future.*
- b. Governments should place a much higher priority on investment than is currently the case, whether that concerns the private sector or the public sector. Relative to what we should be doing, we are currently living in an investment drought.*
- c. Policy should be more concerned with economic growth, properly specified, and policy discussion should pay less heed to other values. And yes, your favorite value gets downgraded too. No exceptions, except of course for the semi-absolute human rights.*



d. We should be more concerned with the fragility of our civilization.

e. The possibility of historical pessimism stands as a challenge to this entire approach, because in that case the future is dim no matter what and there may not be a more distant future to resolve the aggregation dilemmas involved in making decisions which affect so many diverse human beings.

f. At the margin we should be more charitable but we are not obliged to give away all of our wealth. We do have obligations to work hard, save, invest, and fulfill our human potential, and we should take these obligations very seriously.

g. We can embrace much of common sense morality, while knowing it is not inconsistent with a deeper ethical theory. Common sense morality also can be reconciled with many of the normative recommendations which fall out of a more impersonal and consequentialist framework.

i. When it comes to most “small” policies, affecting the present and the near-present only, we should be agnostic because we cannot overcome aggregation problems to render a defensible judgment. The main exceptions here are the small number of policies which benefit virtually everybody.

. . .