Economic Growth and the Interests of Future (and Past and Present) Generations: A Comment on Tyler Cowen

Matthew D. Adler†

What principles vis-à-vis future generations should govern our policy choices? Tyler Cowen argues for a "Principle of Growth": "We should make political choices so as to maximize the rate of sustainable economic growth."[1] Economic growth means the growth of inclusive gross domestic product (GDP), not just marketed goods and services, but also leisure time, household production, and environmental amenities.[2]

Cowen, as I read him, suggests that the Principle of Growth is justified by a welfarist, consequentialist moral theory. The core of the argument is that increasing economic growth both increases average well-being in future generations, and increases the well-being of the worst off in future generations (so there's a distributive benefit as well).[3] He worries about the nonconsequentialist moral precepts of commonsense morality, and suggests that those might be handled through a rights constraint, incorporated in a Modified Principle of Growth.[4]

I find welfare consequentialism quite plausible, so for purposes of this Response I will ignore rights and other nonconsequentialist constructs. I want instead to focus on a worry about the Principle of Growth that arises *within* welfare consequentialism.

Welfare consequentialists say that the moral appropriateness of a choice is determined by its possible outcomes, and that the goodness of an outcome is determined solely by facts about individual welfare. To put this in the language of normative economics, welfare consequentialism says to maximize a social welfare function whose arguments are individual utilities, measuring each individual's well-being.[5]

---

† Leon Meltzer Professor of Law, University of Pennsylvania Law School. Thanks to Eric Posner and Chris Sanchirico for comments, and to David Weisbach for inviting me to participate in the University of Chicago Law School conference on discounting.

1 Tyler Cowen, *Caring about the Distant Future: Why It Matters and What It Means*, 74 U Chi L Rev 5, 16 (2007).

2 See id at 17.

3 See id at 18–23.

4 Id at 17 ("We should push for sustainable economic growth, but not at the expense of inviolable human rights.").

5 See, for example, Matthew D. Adler, *Against "Individual Risk": A Sympathetic Critique of Risk Assessment*, 153 U Pa L Rev 1121, 1183–87 (2005) (explaining welfare consequentialism);

Before I explain why welfare consequentialists should worry about Cowen's Principle of Growth, let me make a few preliminary points. First, it is important to see that Cowen is *not* making a Pareto-dominance argument in the style of Louis Kaplow and Dexter Samida and David Weisbach.[6] He is not claiming that a policy that increases the growth rate can be coupled with intergenerational transfers so that everyone in both current and future generations is better off. Rather, the claim is that, even if a policy that increases the growth rate makes some individuals in the current or near-term generations worse off, it is a better policy. To put the point another way, the Principle of Growth is offered as a way to choose between Pareto-noncomparable allocations of well-being to different generations.

Second, Cowen suggests that the welfare-consequentialist social planner should be time neutral. He should not give less weight to the well-being of future individuals per se. Time neutrality, indeed, might seem to be the hallmark of any recognizably *moral* theory about how we should act vis-à-vis the future (as opposed to a predictive theory, which postulates how the political process will in fact operate, or a theory of individual rational choice, which might allow an individual to discount his own, and certainly someone else's, well-being). In fact, I am not sure that neutrality regarding future interests is morally unchallengeable, for two reasons. One involves the problem of infinite futures and infinite streams of well-being, which Geoffrey Heal discusses.[7] A second involves "existence" or "nonidentity" problems. If some present policy choice brings into existence a future individual—who would not exist were a different choice to be selected—does her well-being count in evaluating the goodness of that choice?[8]

However, bracketing infinite welfare streams and existence problems, I think it *is* reasonably clear that welfare consequentialism should be time neutral. Imagine that the universe has finite spatial

---

Matthew D. Adler and Chris W. Sanchirico, *Inequality and Uncertainty: Theory and Legal Applications*, 155 U Pa L Rev 279, 291–96 (2006) (same).

6 See Dexter Samida and David A. Weisbach, *Paretian Intergenerational Discounting*, 74 U Chi L Rev 145, 155–60 (2007); Louis Kaplow, *Discounting Dollars, Discounting Lives: Intergenerational Distributive Justice and Efficiency*, 74 U Chi L Rev 79, 86–87 (2007). Of course, Samida and Weisbach and Kaplow advance a Pareto-dominance argument in favor of discounting, not in favor of a principle of growth. But the more general point of their articles—I take it—is that, in formulating principles for intergenerational policy, we should keep in mind the possibility that some policy option coupled with a scheme of intergenerational wealth transfers might be Pareto-superior to another policy. Cowen, however, is *not* presenting this sort of argument for the Principle of Growth. See Cowen, 74 U Chi L Rev at 22 (cited in note 1) ("Pure Pareto improvements are few and far between. So at some level of the analysis, through some method or another, we must assert that the benefits to one group of people outweigh the losses to another.").

7 See Geoffrey Heal, *Discounting: A Review of the Basic Economics*, 74 U Chi L Rev 59, 61–62 (2007).

8 See Derek Parfit, *Reasons and Persons* 351–79 (Oxford 1984).

extent, will definitely end, and contains a finite number of individuals whose identities are fixed relative to the choice at hand. In that case, how could the welfare consequentialist justify giving less weight to the interests of future individuals just by virtue of their temporal position? Note that a number of the standard arguments for discounting pressed by economists do *not* invoke a pure time preference—a preference for present over future well-being. Rather, these arguments point to opportunity costs; the proposition that future individuals will be wealthier (hence have a lower marginal utility of money, so that future dollars should be discounted); or uncertainty.[9]

A third, and final, preliminary point. In his article, Cowen raises the possible concern that increased GDP (even GDP defined in an inclusive sense) might not mean increased well-being. What if having a larger GDP does not make people happier?[10] I am not that worried about this issue. Well-being—as I have argued elsewhere—involves the satisfaction of self-interested, fully informed preferences.[11] The lists of objective goods proposed by objectivists about well-being, such as Martha Nussbaum's list—life; bodily health; bodily integrity; senses, imagination, and thought; emotions; practical reason; affiliation; other species; play; control over one's environment[12]—are, in turn, plausibly understood as estimates of what self-interested individuals with full information would prefer.[13] Happiness is one component of well-being, but not the sole component.[14] Whatever the connection between GDP and happiness, it seems very plausible that there is a strong empirical connection between a society's overall GDP and its overall well-being

---

9 See Matthew D. Adler and Eric A. Posner, *New Foundations of Cost-Benefit Analysis* 175–76 (Harvard 2006).

10 See Cowen, 74 U Chi L Rev at 23–27 (cited in note 1). Cowen ultimately concludes that "wealth and happiness commove in the longer run." Id at 24. But what matters to the welfarist is the connection between wealth (GDP) and well-being, not wealth and happiness. My point, in this paragraph, is that greater GDP plausibly correlates with greater overall well-being even if it doesn't correlate (or correlate as well) with greater overall happiness.

11 See Adler and Posner, *New Foundations of Cost-Benefit Analysis* at 38–39 (cited in note 9). Strictly, Posner and I argue that well-being consists in the satisfaction of self-interested, idealized preferences, without taking a position in the debate between full-information accounts of idealization and other accounts. I, in fact, find the full-information account particularly plausible. In any event, my claim, here, that well-being encompasses more than happiness, and that GDP correlates well with overall well-being regardless of its correlation with overall happiness, does not depend on the specifics of preference idealization.

12 See Martha C. Nussbaum, *Women and Human Development: The Capabilities Approach* 78–80 (Cambridge 2000).

13 See, for example, Adler and Posner, *New Foundations of Cost-Benefit Analysis* at 31–32, 51, 74–75 (cited in note 9); Matthew D. Adler, *Beyond Efficiency and Procedure: A Welfarist Theory of Regulation*, 28 Fla St U L Rev 241, 297–99 (2000).

14 See Adler and Posner, *New Foundations of Cost-Benefit Analysis* at 28–31, 50 (cited in note 9).

(or, equivalently, between its per capita GDP and the average well-being of its members).[15]

So why do I worry about the justifiability of Cowen's Principle of Growth within welfare consequentialism? The problem is this: the contemporary philosophical literature points to four plausible variants of welfare consequentialism, and it's far from clear that all would underwrite the Principle of Growth. At a minimum, Cowen hasn't done the analytic work to demonstrate that they all would.

The four variants are:

1. *Utilitarianism.* Utilitarians, of course, seek to maximize overall well-being.

2. *Prioritarianism*. The idea here is to give greater weight to changes in well-being that affect individuals who are worse off. The effect on the social calculus of a change in some individual's well-being is a function, not just of the size of the change, but also of the individual's level of well-being. More technically, prioritarians do not sum utilities (as do utilitarians) but sum an increasing concave function of each individual's utility. (The sum of the square root of each individual's utility would be an example of a prioritarian social welfare function.) An equivalent definition is that prioritarians maximize a social welfare function which is Paretian and equity-regarding in the Pigou-Dalton sense and separable in individual utilities. The limiting point of prioritarianism is the leximin principle (defined on welfare, not primary goods).[16]

3. *Comparativism*. This maximizes a Paretian social welfare function which is equity-regarding in the Pigou-Dalton sense but *not* separable in individual utilities.[17] The idea here (nontechnically) is that the comparativist is interested in the pattern of welfare levels while the prioritarian is not. Consider a policy that, with certainty, increases person P's welfare from fifty to fifty-eight, and de-

---

15 I say "empirical connection" because, even if GDP is defined inclusively to take account of leisure time, household production, and environmental amenities, it is *conceptually* possible for GDP and overall well-being to diverge. Nussbaum's list, for example, includes items that would not be subsumed by an inclusive GDP measure.

16 On prioritarianism, see Adler and Sanchirico, *Inequality and Uncertainty* at 296–304 (cited in note 5); Adler and Posner, *New Foundations of Cost-Benefit Analysis* at 56–58 (cited in note 9). Leximin ranks two outcomes by comparing the welfare levels of the individuals who are worst off in each outcome and picking the outcome where the worst-off individual's welfare is higher—or, if the worst-off individuals are equally well off, comparing the second-to-worst-off individuals, then (if these are equal), the third-to-worst-off, and so on.

17 See Adler and Sanchirico, *Inequality and Uncertainty* at 296–304 (cited in note 5).

creases Q's from thirty to twenty-seven, and affects no one else. For the utilitarian and the prioritarian, this is all the information we need to evaluate the policy. By contrast, comparativists need more information. Comparativists evaluate policies by considering not merely welfare changes and the welfare levels of affected individuals (those whose welfares change), but also the welfare levels of unaffected individuals. For example, if the social welfare function being maximized is the sum of rank-weighted utility—one kind of comparativist function—the policy will increase social welfare if there is one other individual in the population and she is worse off than P and Q, but not if her well-being level is in between P's and Q's.[18]

4. *Sufficientism*. The idea here is to give greater weight, or perhaps absolute priority, to welfare changes that affect individuals below some well-being threshold—call it the poverty line, or the line of basic functioning.[19]

Utilitarianism corresponds to a specific social welfare function: the unweighted sum of utilities. By contrast, prioritarianism, comparativism, and sufficientism each correspond to a distinct family of social welfare functions[20]—each to a different, generic departure from the simple utilitarian formula of maximizing overall well-being.

It is pretty straightforward to see that prioritarian, comparativist, and sufficientist social welfare functions need not endorse a policy that maximizes economic growth. Consider a very simple example, which involves a finite, fixed population—meaning that the total number of individuals across time is the same finite amount regardless of which policy is picked—and no problems of *intra*generational equity or uncertainty. There are two generations, with F individuals in the first generation and S in the second. $S > F$; the size of the living population

---

[18] Imagine that the third individual, M, has utility level m in both cases. If M is worse off than both P and Q (in both outcomes), the rank-weighted sum of utilities in the status quo equals $(3 \times m) + (2 \times 30) + (1 \times 50)$, and the rank-weighed sum of utilities with the policy equals $(3 \times m) + (2 \times 27) + (1 \times 58)$. In this case, the policy is better. By contrast, if M is in between P and Q (in both outcomes), then the rank-weighted sum of utilities in the status quo equals $(3 \times 30) + (2 \times m) + (1 \times 50)$, and the rank-weighted sum of utilities with the policy is $(3 \times 27) + (2 \times m) + (1 \times 58)$. In this case, the status quo is better.

[19] For a defense of sufficientism, see generally Roger Crisp, *Equality, Priority, and Compassion*, 113 Ethics 745 (2003).

[20] More precisely, each corresponds to a family of social welfare *orderings*. An ordering is a ranking of outcomes. Some orderings cannot be represented by functions (where a function is a mapping from each outcome to a number that represents the place of the outcome in the ordering). For example, a leximin ordering of an uncountably infinite set of outcomes cannot be represented by a function. However, since the term "social welfare function" is more familiar than "social welfare ordering," I use the former term in the text.

is growing over time.[21] The status quo is zero economic growth.[22] Everyone in the first generation is at utility level ten, and everyone in the second generation is at level ten. There's a policy that produces economic growth—at some cost. (Because the Principle of Growth is meant to adjudicate between Pareto-noncomparable outcomes, it will, inter alia, approve some policies that purchase growth at some cost to the current generation.) Let us assume that we can trade a small decrease in per capita GDP and average well-being in the first generation for a substantially larger increase in per capita GDP and average well-being in the second generation. In particular, let us assume that the growth policy reduces the well-being of each first-generation individual to nine, and increases the well-being of each second-generation individual to fifteen.

What will our different social welfare functions say about the policy?

The utilitarian will approve the policy. Total well-being in the status quo is $10F + 10S$. Total well-being with the growth policy is $9F + 15S$. Because $S > F$, the policy increases overall well-being.

The prioritarian may not approve the growth policy. To begin, regardless of the size of F and S, a leximin social welfare function will certainly not approve the growth policy. In the status quo, the worst-off person is at level ten. With growth, the worst-off person (namely, everyone in the first generation) is at level nine. What about less extreme versions of prioritarianism—those that sum an increasing concave function of individual utilities and therefore (by contrast with leximin) allow a sufficiently small loss to the welfare of the worst-off individual to be compensated by a sufficiently large gain for individuals who are better off? Formally, $\mathrm{SW}(O) = \sum \mathrm{W}(U_i)$, where SW is the social welfare function, U_i is the utility of individual i, and W is an increasing concave function. The social value of the status quo equals $(F \times \mathrm{W}(10)) + (S \times \mathrm{W}(10))$. The social value of the policy equals $(F \times \mathrm{W}(9)) + (S \times \mathrm{W}(15))$. Which is larger depends on the shape of W.[23]

---

21 Talk of "population change" is ambiguous. In the example at hand, the population is temporally variable (the number of individuals in each generation is not the same) but *modally* fixed. The number of individuals at each time, and thus throughout time, is the same in each possible world—whether the policy or status quo is chosen. It is modal, not temporal, variation in population that leads to puzzles for welfarists—the existence/nonidentity puzzles.

22 The growth metric used in this example is the growth of average well-being from generation to generation. In principle, as mentioned in note 15, average well-being and GDP per capita might diverge, but I'm ignoring that possibility here.

23 Consider the family of functions $\mathrm{W}(U_i) = -1 \times (U_i)^N$, where $N < 0$. These functions have positive first derivatives and negative second derivatives, so are increasing and concave. If N is sufficiently close to zero, the growth policy is better than the status quo. But if N is sufficiently far from zero, the status quo is better than the growth policy. Consider, for concreteness, the case where there are twice the number of individuals in the second generation as in the first: $S = 2F$.

The comparativist may not approve the growth policy. To be sure, the growth policy increases total well-being, which grows from $10F + 10S$ to $9F + 15S$. On the other hand, the policy increases the inequality in welfare levels in the intertemporal population—that is the entire population of $F + S$ individuals taken as a group. In the status quo, there is no gap at all in welfare levels. Everyone in both generations is at level ten. With the policy, the gap becomes six. The better-off individuals (namely everyone in the second generation) are at level fifteen; the worse-off individuals (namely everyone in the first generation) are at level nine. The distribution of utilities with the growth policy, (9, 9, 9, . . . , 15, 15, 15, . . .), is certainly a less equal distribution than the distribution of utilities in the status quo, (10, 10, 10, . . . , 10, 10, 10, . . .), and the comparativist might judge that it is worse, all things considered, notwithstanding the increase in total well-being.

Finally, the sufficientist may not approve the growth policy. Let us imagine that the threshold, for purposes of the sufficientist theory, is just at level ten. Then the growth policy increases overall well-being but also increases the number of individuals below the threshold, from zero to F. Depending on how the sufficientist trades off below-threshold and above-threshold utility, she may or may not approve the growth policy. Alternatively, imagine that the sufficientist threshold is five. In this case, the sufficientist will approve the growth policy, because it increases overall well-being and the number of below-threshold individuals is zero both with growth and in the status quo.[24]

The lesson of this example, I suggest, is that the justifiability of Cowen's Principle of Growth remains open to question within welfare consequentialism, and more generally that the rejection of nonconsequentialism and nonwelfarism hardly settles questions of social policy vis-à-vis future generations. Even with welfare consequentialism in hand, we still need to do the hard, philosophical work of figuring out which specific social welfare function policymakers should use. And, if the specification of a social welfare function is seen as an irreducibly subjective matter, appropriate for legislatures rather than scholars, we will need to refer the issue to the political process and await its answer before recommending intergenerational policies. Except in the limit-

---

In this case, if for example $N = -5$, the growth policy is better. But if, for example, $N = -12$, the status quo is better.

[24] The standard version of sufficientism stipulates that above-threshold individuals do not have egalitarian claims vis-à-vis each other. See Crisp, 113 Ethics at 755 (cited in note 19) ("[W]hen people reach a certain level, even if they are worse off than others, benefiting them does not, in itself, matter more."). So, in the limiting case where the number of below-threshold individuals is zero regardless of the policy chosen, sufficientism reduces to utilitarianism.

ing case where a policy is genuinely Pareto superior[25] and debates between utilitarians, prioritarians, comparativists, and sufficientists can be ignored, welfarist intergenerational policymaking depends on the shape of the social welfare function.

This claim does not depend on the unrealistic simplicity of the example I have been discussing. The example was chosen to illustrate the claim in an accessible and direct way. But it is hard to see why it would not be the case that, in more complicated and realistic cases, the optimal intertemporal policy choice, given welfare consequentialism, will depend on the shape of the social welfare function.

A second point illustrated by the example is that the welfare consequentialist (at least bracketing nonidentity problems and problems of infinite populations) should apply her social welfare function to the world's total intertemporal population, consisting of everyone in the first generation plus everyone in the second generation all the way through to everyone in the last generation—without reference to the identities of the individuals, other nonwelfare information, or their position in time. That is, after all, just what welfarism plus time-neutrality means—or at least how it is most naturally expressed. If there are F individuals in the first generation, S in the second generation and T in the third and last generation, and policymakers at any point in time are considering policies, then each possible policy corresponds to a vector of utilities with $F + S + T$ entries, one for each individual who exists at some time, whether in the past, the present, or the future.[26] Policies are chosen by applying the preferred social welfare function—utilitarian, prioritarian, comparativist, or sufficientist—to these utility vectors. And the social welfare function must be *anonymous*, meaning that reorderings of the same set of numbers must be ranked the same.[27] That is, (1, 1, 5, 5, 7, 7) must be ranked the same as (5, 1, 7, 5, 7, 1), and so forth. The crucial point about this formalism is that the *only* information about a given individual that influences the

---

25 Meaning that some individuals in some generations are better off and no one is worse off. Although it is possible to marry welfarism and non-Paretianism, the arguments for the Pareto principle are very strong and I therefore assume that any plausible social welfare function will approve a Pareto-superior outcome. See Adler and Sanchirico, *Inequality and Uncertainty* at 293–94 (cited in note 5).

26 This ignores uncertainty, and assumes that each policy corresponds to a particular outcome—a particular vector of utilities with $F + S + T$ entries. Given uncertainty, each policy corresponds to a set of possible outcomes—to a set of possible utility vectors, each with $F + S + T$ entries. See id at 304–09.

27 On "anonymous" or, equivalently, "symmetric" social welfare functions, see id at 294.

social welfare function is his well-being. His name, other nonwelfare information, *and* his position in time are all washed out.[28]

This welfarist and temporally neutral way of conceptualizing policy choice undermines Cowen's suggestion that nonutilitarians should approve the Principle of Growth because greater growth makes poor individuals in the future better off than they would be with less growth.[29] Even if that is true, the issue for the nonutilitarian, given time neutrality, is the distribution of welfare levels throughout time, not merely the distribution in the future. For example, if the policymaker employs a leximin function, she should lump all the generations together, identify the individual who is worst off in that combined population given each policy, and choose the policy that maximizes that individual's welfare level. If the policymaker's view is comparativist, then she should lump all the generations together and consider the pattern of utilities in that combined population associated with each policy. Increased economic growth can improve the position of badly-off individuals in future generations, but increased growth can also reduce the position of badly-off individuals right now; it can increase the gap between our welfare and the welfare of future generations; and it may in principle increase, rather than decrease, the extent to which people fall below "sufficientist" thresholds, depending on where those are set.

---

[28] It might be objected that nonutilitarians would care about the social position of individuals. Individuals in the same society have stronger redistributive claims on each other than on individuals in a different society—or so it might be thought. But that view amounts to a kind of nonwelfarism, and in any event doesn't argue for taking each generation as a unit (since members of different generations can be part of the same society, and members of the same generation can be part of different societies), or for other sorts of departures from time-neutrality.

[29] See Cowen, 74 U Chi L Rev at 21 (cited in note 1).