# Discounting Dollars, Discounting Lives:
# Intergenerational Distributive Justice and Efficiency

## Louis Kaplow†

*The view that intergenerational distributive justice and efficiency should be treated separately is familiar, yet controversial. This Article elaborates the often-implicit justifications for separate treatment and provides a more express statement of how and when such treatment is appropriate. Substantial attention is devoted to an approach that holds constant the intra- and intergenerational distribution of well-being, which proves to be a valuable analytical device even for intergenerational policies that are not distribution neutral. Also explored are possible interrelationships between intergenerational distributive justice and efficiency, the choice of interest rate for discounting dollars, and how the present approach relates to those that would employ direct social weights to dollars at different points in time.*

## I. INTRODUCTION

The thesis of this Article is that questions of intergenerational distributive justice and of intergenerational efficiency are substantially distinct in principle. A pure issue of intergenerational distributive justice would be presented in a simple, hypothetical world, in which the only decision with intergenerational consequences concerned the rate of use of a single resource: more for any one generation would entail less for another. The argument here is that, however that sort of allocation should ideally be made or is in fact determined, all further consideration of intergenerational dimensions of policies—whether involving the environment, infrastructure, research and development, education, or social security—dissolves almost entirely into matters of efficiency.[1] This claim does not entail that the remaining efficiency

[1]    The question of population size is distinct, but note that it may arise even in the simple world in which there is only one resource to allocate over time. Additionally, population size issues might be classified as intragenerational in nature. In any event, population size will not be considered further here.

questions are straightforward or uncontroversial. It does imply that they are essentially the same questions that arise when considering policies whose impact is confined to a single generation.

Part of this claim may be restated using the rubric of discounting. Whether lives—or, more broadly, utility—in different generations should be discounted is one way to define some positions in debates about intergenerational distributive justice.[2] From the welfarist perspective generally followed here—under which social welfare is taken to be solely a function of individuals' well-being—a discount rate (including a zero rate) may be employed in specifying the social welfare function that aggregates the well-being of individuals in different generations.[3] Different discount rates on utility in this setting correspond

---

[2]    Some favor the discounting of utility or lives to avoid imposing what are perceived as excessive savings obligations on earlier generations. One strand of such work builds on F.P. Ramsey, *A Mathematical Theory of Saving*, 38 Econ J 543 (1928), which presents a model where early generations' savings rates should be very high. See, for example, K.J. Arrow, et al, *Intertemporal Equity, Discounting, and Economic Efficiency*, in James P. Bruce, Hoesung Lee, and Erik F. Haites, eds, *Climate Change 1995: Economic and Social Dimensions of Climate Change* 125, 137 (Cambridge 1996); Kenneth J. Arrow, *Discounting, Morality, and Gaming*, in Paul R. Portney and John P. Weyant, eds, *Discounting and Intergenerational Equity* 13 (Resources for the Future 1999). Additionally, in some models a zero discount rate over an infinite horizon leads to conundrums, such as no consumption of exhaustible resources ever being permitted (when such resources are the only source of consumption). See, for example, Geoffrey M. Heal, *The Optimal Use of Exhaustible Resources*, in Allen V. Kneese and James L. Sweeney, eds, 3 *Handbook of Natural Resource and Energy Economics* 855, 868–78 (Elsevier 1993). In any event, a range of scholars (including Ramsey) have opposed any discounting of utility, the most commonly cited reason being that it violates equal treatment (although some would nevertheless impose limits on the sacrifice required of early generations if a pure no-discounting approach required great sacrifice). See, for example, John Broome, *Counting the Cost of Global Warming* 52–112 (White Horse 1992); Tyler Cowen, *Consequentialism Implies a Zero Rate of Intergenerational Discount*, in Peter Laslett and James S. Fishkin, eds, *Justice Between Age Groups and Generations* 162 (Yale 1992); Tyler Cowen and Derek Parfit, *Against the Social Discount Rate*, in Laslett and Fishkin, eds, *Justice Between Age Groups* 144; Derek Parfit, *Reasons and Persons* 480–86 (Oxford 1984); Richard L. Revesz, *Environmental Regulation, Cost-Benefit Analysis, and the Discounting of Human Lives*, 99 Colum L Rev 941 (1999); and Henry Sidgwick, *The Methods of Ethics* (Chicago 7th ed 1907). Interestingly, Rawls, although arguing from an original position, does not reach this conclusion, in part because he stipulates that only a single generation is taken to be behind the veil of ignorance (even though that generation is ignorant of which generation it is). See John Rawls, *A Theory of Justice* 284–98 (Harvard 1971). This approach has been criticized. See, for example, Brian Barry, *Theories of Justice* 189–203 (California 1989); B.M. Barry, *Justice Between Generations*, in P.M.S. Hacker and J. Raz, eds, *Law, Morality and Society: Essays in Honour of H.L.A. Hart* 268 (Oxford 1977). A revised view is presented in John Rawls, *Political Liberalism* 274 (Columbia 1993). For a libertarian approach that would impose more limited obligations on present generations, see David Gauthier, *Morals by Agreement* (Oxford 1986). For collections of views, see Laslett and Fishkin, eds, *Justice Between Age Groups,* and Portney and Weyant, eds, *Discounting and Intergenerational Equity*.

[3]    Although much of the analysis will be explicitly welfarist, many of the ideas are applicable to a variety of non-welfarist theories for the reasons suggested in the conclusion. For contrasting views on welfarism, see, for example, Amartya Sen and Bernard Williams, eds, *Utilitarianism and Beyond* (Cambridge 1982). For a defense, see, for example, R.M. Hare, *Moral Think-*

to different views of the appropriate social weight on different genera-
tions. By contrast, whether and how dollars in different generations
should be discounted is often viewed as a question of efficiency. An
implication of the present view, as will be seen, is that efficiency does
require discounting dollars—and, in particular, at the market rate of
return to capital—and that this is true independently of whether and
how lives should be discounted. Note that, since this conclusion does
not depend on the appropriate rate of discount on lives, it is possible a
priori that either of the two discount rates could be higher.

   The view that distribution and efficiency should be separated in
the intergenerational context is hardly novel.[4] Yet adherents largely
take its validity for granted while others reject it. The purpose of the
present investigation is to analyze more precisely the foundation for
the position.[5] This effort should generate greater acceptance and lead
to a better appreciation of the broad applications of the proposed
framework. In any event, a more explicit articulation of justifications
also serves to expose the underlying premises more sharply, which
should enhance understanding of pertinent qualifications as well as
provide a more focused target for those who might disagree.

   A number of perspectives on intergenerational efficiency are de-
veloped in Part II. First, it is suggested that standard methods of cost-
benefit analysis within generations combined with ordinary discount-
ing of returns to capital between generations validate the use of dollar
discounting for all intergenerational policy decisions. (How this en-
compassing conclusion can be fully reconciled with a wide range of
views on intergenerational distributive justice will be explained.) Sec-
ond, deviations from this approach are shown to be inefficient; specifi-

---

*ing: Its Levels, Method, and Point* (Oxford 1981), and Louis Kaplow and Steven Shavell, *Fairness Versus Welfare* (Harvard 2002).

   [4]    See, for example, Arrow, et al, *Intertemporal Equity* at 125 (cited in note 2); David F. Bradford, *On the Uses of Benefit-Cost Reasoning in Choosing Policy toward Global Climate Change*, in Portney and Weyant, eds, *Discounting and Intergenerational Equity* 37 (cited in note 2); John Broome, *Discounting the Future*, 23 Phil & Pub Aff 128 (1994); Charles D. Kolstad and Jeffrey A. Krautkraemer, *Natural Resource Use and the Environment*, in Kneese and Sweeney, eds, *Handbook of Natural Resource* 1219 (cited in note 2); Robert C. Lind, *A Primer on the Major Issues Relating to the Discount Rate for Evaluating National Energy Options*, in Robert C. Lind, et al, *Discounting for Time and Risk in Energy Policy* 21 (Resources for the Future 1982); Revesz, 99 Colum L Rev 941 (cited in note 2); W. Kip Viscusi, *Discounting Health Effects for Medical Decisions*, in Frank A. Sloan, ed, *Valuing Health Care: Costs, Benefits and Effectiveness of Pharmaceuticals and Other Medical Technologies* 125 (Cambridge 1995); Martin L. Weitzman, *"Just Keep Discounting, But . . . ,"* in Portney and Weyant, eds, *Discounting and Intergenerational Equity* 23 (cited in note 2).

   [5]    See also Dexter Samida and David A. Weisbach, *Paretian Intergenerational Discounting*, 74 U Chi L Rev 145 (2007); Cass R. Sunstein and Arden Rowell, *On Discounting Regulatory Benefits: Risk, Money, and Intergenerational Equity*, 74 U Chi L Rev 171 (2007); W. Kip Viscusi, *Rational Discounting for Regulatory Analysis*, 74 U Chi L Rev 209 (2007).

cally, policies in violation can be modified in the stipulated direction so as to make all generations better off. Third, it is explained that future generations permitted to make allocative choices for the present generation regarding actions with intergenerational effects would choose policies entirely on grounds of efficiency.

Each of the foregoing three perspectives is somewhat loose, including with regard to questions of *intra*generational distribution that may be affected by intergenerational decisions. The fourth argument in this Part treats all distributive concerns more directly. Specifically, it is described how more efficient policies may be implemented in a manner that is distribution neutral both within each generation and among all generations. When distribution is held constant, matters of (intra- and intergenerational) distributive justice are moot. Greater efficiency can be to the benefit of every individual in every generation. Extensions of and qualifications to this argument are also considered.

Possible relationships between intergenerational distributive justice and efficiency are explored in Part III. First, just as in the intragenerational setting, which policies are efficient may depend on the distribution of resources. Second, private intergenerational transfers and liquidity constraints may alter the effects of government intergenerational transfers. Third, the possibility that future generations may be systematically richer (or poorer) is examined. Although each consideration complicates the determination of optimal intergenerational policy, the conclusion from Part II remains intact.

Because of its centrality to the analysis of intergenerational efficiency, the market interest rate on capital is examined in Part IV. Factors considered are: uncertainty in government and private projects and about the interest rate, capital taxation, and endogeneity of the market interest rate. It is noted that, just as with many of the other complications considered in this Article, the issues are not distinctively intergenerational. That is, they would arise in a world with only one generation (the lives of which span some period of time), and their resolution for intergenerational purposes should be no different from when the problem is viewed within a single generation. One way to put this point is that a dollar invested, in whatever manner, over a period of time does not know or care whether it is an intragenerational dollar (part of someone's lifecycle savings) or an intergenerational dollar.

In Part V, the approach to discounting dollars of using a market interest rate (with appropriate adjustments) is contrasted with the often-advanced alternative of making explicit reference to social weightings of dollars in different time periods. This comparison serves to highlight the interrelationships among the arguments in earlier parts of the Article. Concluding remarks are offered in Part VI. The discussion returns to the core question of distribution versus efficiency in the intergen-

erational context and offers some comments concerning why demonstrations from the welfarist perspective should be seen as persuasive from a number of non-welfarist positions.

## II. Intergenerational Efficiency

### A.  Valuing Utility, Discounting Lives

Suppose that one wishes to compare the relative merits of policies that differentially affect individuals' well-being in the present and in the future. The claim in this section is that one can (1) convert future utility into future dollars using standard principles of valuation, then (2) transform future dollars into present dollars by discounting at the market rate of return on capital, and finally (3) convert present dollars into present utility using standard principles of valuation. Although this approach (and that in sections B and C) has been suggested previously,[6] it is useful to elaborate on its rationale.

For example, consider the construction from scratch of a new urban transit system involving an elaborate system of tunnels that will take a generation to complete. Costs, both financial and nonpecuniary (such as disruption), will be incurred in the present generation, whereas much of the benefits will accrue to future generations. Using the aforementioned technique, future benefits can be valued in dollars and discounted to present value to compare with the costs. Everything could be considered in present dollars (converting present nonpecuniary costs into dollars), as is often done, or everything could be evaluated in present utility. As long as a single common denominator is used and consistent conversion rates are employed, the result will be the same.[7]

The basis for this procedure is straightforward. Steps (1) and (3) each involve conversions between utility and dollars within a time period (generation).[8] Standard cost-benefit principles govern how these conversions are properly made. As long as the appropriate method is employed, these steps are justified. (Note that the methods

---

[6]    See note 4.

[7]    Suppose that, in present dollars, costs are *C* and benefits *B*; the project passes a cost-benefit test if $B > C$. If present dollars convert to present utility at the rate *a*, then in terms of present utility costs are *aC* and benefits are *aB*, so the test is whether $aB > aC$, which is the same test.

[8]    The text speaks of the conversion of utility into dollars, whereas it is often common to speak of the conversion of lives into dollars, under the rubric of the value of a statistical life (VSL). With regard to some environmental issues and other matters (for example, medical research), future statistical lives are indeed an important part of what is at stake. Furthermore, regarding distributive justice, it is more common to refer to discounting lives than utility. Nevertheless, influencing risks to life are but one way that policies may influence utility, and the methods of determining VSL are utility-based in any event, so it is more useful to speak in terms of the more encompassing notion of utility.

are qualitatively the same in each period, even though the conversion ratios could differ for various reasons, as will be noted below.)

Step (2) is a pure transformation of dollars between time periods. Whatever discount rate is appropriate for such purposes—which will be the same as the discount rate that is used to assess an ordinary capital project whose cost is entirely in present dollars and whose benefit is entirely in future dollars—should be employed here. Discounting dollars is justified because the market discount rate is a price that signifies resource costs just like any other price, such as the price for steel, concrete, or other inputs for the project.

It follows, therefore, that since each step in the proposed method for the assessment of intergenerational projects is independently valid, the entire procedure is valid. Notice furthermore that each of the steps is justified independently of any special consideration of the intergenerational setting. Steps (1) and (3) are explicitly intragenerational because each operates within a single time period. Step (2) is intertemporal, but it need not be intergenerational. That is, the same reasoning applies to a project undertaken by individuals when they are young that is for their own benefit when they are old. (Recall the reference in the introduction to lifecycle savings, such as for retirement.)

What about distributive concerns? Steps (1) and (3) involve intragenerational distribution, familiar from the debate about whether distributive weights should be employed in cost-benefit analysis.[9] There is a strong presumption against using distributive weights because any redistribution is best undertaken (or offset) through the income tax,[10] a view that will be elaborated more explicitly in sections D and E. At this point, however, it is sufficient to note that, if one does accept the appropriateness of distributive weighting, then those weights should be employed in steps (1) and (3) and the procedure

---

[9]    See, for example, Burton A. Weisbrod, *Income Redistribution Effects and Benefit-Cost Analysis*, in Samuel B. Chase, Jr., ed, *Problems in Public Expenditure Analysis* 177 (Brookings 1968); and the survey by Jean Drèze and Nicholas Stern, *The Theory of Cost-Benefit Analysis*, in Alan J. Auerbach and Martin Feldstein, eds, 2 *Handbook of Public Economics* 909 (North-Holland 1987). Arrow and coauthors state that, in addition to aggregate economic benefits and costs, "equity within and across generations" may sometimes be important and that "a good analysis will also identify important distributional consequences"; nevertheless, they also state that "environmental, health, and safety regulations are neither effective nor efficient tools for achieving redistributional goals." Kenneth J. Arrow, et al, *Is There a Role for Benefit-Cost Analysis in Environmental, Health, and Safety Regulation?*, Science 221, 222 (April 12, 1996).

[10]   This argument is first developed in Aanund Hylland and Richard Zeckhauser, *Distributional Objectives Should Affect Taxes but Not Program Choice or Design*, 81 Scandinavian J Econ 264 (1979). It is applied explicitly to legal rules in Steven Shavell, *A Note on Efficiency vs. Distributional Equity in Legal Rulemaking: Should Distributional Equity Matter Given Optimal Income Taxation?*, 71 Am Econ Rev 414 (1981), and Louis Kaplow and Steven Shavell, *Why the Legal System Is Less Efficient Than the Income Tax in Redistributing Income*, 23 J Legal Stud 667 (1994). See also the sources cited in note 12.

would be valid. Keep in mind that this distributive weighting is only concerned with intragenerational distributive justice.

Step (2), by contrast, does potentially raise questions about intergenerational distribution. Specifically, if costs and benefits are not allocated across generations in the same proportions, then the discounting in step (2) might be questioned as ignoring distribution by equating present and market-discounted future dollars. The efficiency point still stands, however, if one compares projects with a similar intergenerational distributive incidence or in other ways holds intergenerational distribution constant, as will be done explicitly in the sections that follow. (Furthermore, Parts III and IV will discuss how both intragenerational valuations and the market discount rate already reflect the intergenerational distribution that prevails.)

## B.   Efficient Tradeoffs

Policies in violation of the aforementioned cost-benefit approach, which combines intragenerational dollar valuation of utility and intergenerational discounting of dollars, can be modified to reduce the extent of deviation in a manner that benefits all pertinent generations. Suppose, for example, that the proposed transit project is to be undertaken on a smaller scale than indicated by the three-step cost-benefit test. (The analysis for a project undertaken to excess would be the same, *mutatis mutandis*, and thus will not be considered separately.)

It is useful to state more precisely what this description implies. Suppose that the cost of an incremental increase in the scale of the project to the present generation, converted into present dollars, is $C_p$, and that the benefit of this modification to the future generation, in future dollars, is $B_f$. Furthermore, let $R$ denote the return on capital between the relevant time periods; that is, one dollar invested in an ordinary (marginal) capital investment in the present yields $1 + R$ dollars in the future. Therefore, the statement that the project expansion would pass section A's cost-benefit test means that $C_p < B_f/(1 + R)$, or, equivalently, $C_p(1 + R) < B_f$. (In the former inequality, $B_f/(1 + R)$ is the future benefit, in future dollars, discounted to present dollars; in the latter inequality, $C_p(1 + R)$ is the present cost, in present dollars, converted to future dollars.)

Consider now a reform that expands the transit project in the manner just described and simultaneously contracts an ordinary capital investment that benefits the future generation by an equivalent amount, $C_p$. For the latter, one could imagine anything done on behalf of the future generation; as just one example, the government could issue debt in the amount of $C_p$, with all accumulated interest and the original principal to be paid off by the future generation. The future generation will thus have to sacrifice an additional $C_p(1 + R)$ to cover

the forgone investment or debt. However, it was stipulated that its benefit from the project, $B_f$, is greater than $C_p(1 + R)$, so the future generation is better off. The present generation breaks even: it pays $C_p$ for the project expansion but saves $C_p$ from the other financial adjustment. Hence, the future generation is better off and the present generation is no worse off.

Obviously, one could also share the net benefit between the two generations. For example, the latter adjustment (forgone investment or increased present generation debt) could be raised from $C_p$ to $C_p + \frac{1}{2}[B_f/(1 + R) - C_p]$. The term in brackets is the net efficiency gain from the project expansion, evaluated in present dollars. With this adjustment, each generation is better off by half of that amount. Hence, as asserted at the outset, if there is any deviation from the three-step test outlined in section A, it is possible to move in the direction indicated by that procedure in a manner that benefits all affected generations. In this case, the intergenerational distribution is indeed held constant, providing a clear justification for relying purely on an efficiency test.[11]

## C. A Future Generation's Preferred Choices for a Present Generation's Allocations

Suppose that some overall allocation between the present generation and a future generation has been selected. (The analysis to follow does not depend on whether this choice accords with any particular theory of intergenerational distributive justice.) One can imagine that whatever amount is to be invested on behalf of the future generation is placed in a segregated pool, a trust of sorts, and that a guardian who will faithfully carry out the will of the future generation has complete authority regarding how the pool is to be invested. The claim in this section is that the guardian, seeking to maximize the overall well-being of the future generation, will be guided solely by considerations of intergenerational efficiency, consistent with the conclusion of the preceding two sections.

Possible investments for this pool include ordinary private capital investments (for example, machines), improvements to infrastructure, research and development, education of the young who will become the adults of the future generation, environmental protection, resource conservation, building government institutions (including enhanced future security through defense or foreign aid), and fostering

---

11   As noted in section A, there are also intragenerational distributive issues, but they may be addressed either separately or through using distributive weights in valuing the costs and benefits within each generation.

institutions and activities that will favorably influence social norms. In addition, there would be pure financial claims, such as the extent to which social security is prefunded (versus pay-as-you-go) and the magnitude of government debt more broadly. The guardian is authorized to expend more resources on any such investments, but each incremental expenditure must be offset by reductions in other investments: the composition of the investment portfolio, so to speak, may be freely adjusted, but its total size is fixed.

It is apparent that the guardian, in seeking to maximize the benefit to the future generation, will be guided solely by considerations of efficiency. Each investment will be pursued to the point at which its marginal return just equals the market interest rate on capital (investment in which, in various ways, is one of the choices open to the guardian). Indeed, the problem is conceptually no different from one involving purely intragenerational allocation within the time period of the future generation. Expenditures that must be made in different periods have different prices—the interest rate, it is well understood, simply constitutes the market price for one particular input. For example, if it is cheaper to clean up an environmental problem sooner rather than later, that is what the guardian would choose, by allocating some of the pool to environmental protection. However, if it is cheaper to clean up later, that can be done by allocating resources instead, say, to some market investment, wherein the future proceeds would be expended subsequently on cleanup. In choosing between the two alternatives, of course, the costs would be compared using the market interest rate. (Either the present cost would be multiplied by $1 + R$ to put it in common terms with the future outlay, or the future outlay would be discounted—divided by $1 + R$—to put it in common terms with the present cost.)

What of distributive concerns? The intergenerational distribution is held constant under this construct and hence may be ignored. Likewise, the present generation's intragenerational distribution is unaffected. Regarding the intragenerational distribution within the future generation, the same remarks offered in section A apply: that too might be held constant—for example, through an income tax adjustment, as will be elaborated in section D—or, if thought appropriate, distributive weights would be employed; that is, they would be employed regarding the costs and benefits both of projects funded in the present generation from the pool and of those funded in the future generation, whether financed by proceeds from investment of the pool or out of then-current production.

The present perspective implies that it is incomplete and potentially misleading to suggest that the present generation does (or does not) have an obligation to a future generation to do one specific thing or another, such as cleaning up the environment, conserving nonre-

newable resources, or avoiding accumulation of a large debt. As a matter of intergenerational distributive justice, there may be an obligation that the pool in question be of a certain size. How it should be allocated, however, is a matter of efficiency. Thus understood, it would nevertheless be meaningful to state that some particular obligation exists. The meaning would be that, given the amount of resources that the current generation is supposed to make available for a future generation, it would be efficient to spend certain amounts on particular projects. Hence, the statement of obligation would entail a judgment about efficiency, itself conditional on a judgment about the intergenerational distribution (on which, see Part III.A).

D.   Intra- and Intergenerational Distribution Neutrality

The perspectives offered in sections A through C varied in the explicitness with which they addressed intra- and intergenerational distribution. It is useful to state more precisely a notion of distribution neutrality that makes statements about the separation of efficiency and distribution more rigorous, particularly from an intergenerational perspective. In addition, it is desirable to have a framework in which one can identify the role of the income tax, transfers, and other fiscal instruments pertinent to distributional matters. Accordingly, this section explains how the formal analysis that has been previously developed in the intragenerational context may be interpreted and extended to encompass intergenerational distribution and efficiency.[12]

To begin, it is helpful to introduce some notation. Let $U_i(y - T_i(y,z), z)$ denote the utility of an individual in generation $i$ who has consumption of $y - T_i(y,z)$ (income $y$ net of taxes and transfers $T_i(y,z)$, on which more in a moment) when the prevailing regime is $z$. We can think of the simple case largely employed above in which there are two generations, in which case $i$ can equal 1 or 2, or a se-

---

[12]   Important precursors are cited in note 10. More formal treatments and further elaboration appear in Louis Kaplow, *The Optimal Supply of Public Goods and the Distortionary Cost of Taxation*, 49 Natl Tax J 513 (1996); Louis Kaplow, *On the (Ir)Relevance of Distribution and Labor Supply Distortion to Government Policy*, 18 J Econ Persp 159 (Fall 2004); Louis Kaplow, *On the Undesirability of Commodity Taxation Even When Income Taxation Is Not Optimal*, 90 J Pub Econ 1235 (2006); Louis Kaplow, *Public Goods and the Distribution of Income*, 50 Eur Econ Rev 1627 (2006); Louis Kaplow, *Optimal Control of Externalities in the Presence of Income Taxation* (National Bureau of Economic Research Working Paper Series No 12339, 2006), online at http://www.nber.org/papers/w12339 (visited Jan 22, 2007); Louis Kaplow, *The Theory of Taxation and Public Economics* (Princeton forthcoming). Much of the work can be understood as generalizations of and extensions to Atkinson and Stiglitz's seminal paper on optimal income and commodity taxation (albeit using a different approach). See A.B. Atkinson and J.E. Stiglitz, *The Design of Tax Structure: Direct Versus Indirect Taxation*, 6 J Pub Econ 55 (1976). The discussion in the text presents only a heuristic argument, but one that is highly suggestive of how these previous analyses can readily be interpreted or slightly extended to make the present statements rigorous.

quence of discrete generations (there might be $n$ generations, or they might extend into the indefinite future).[13] Individuals' income levels $y$ may differ, which makes intragenerational as well as intergenerational distributive questions interesting.[14] For the moment, however, incomes will be taken as given. (Section E.1 will make income endogenous, a function of individuals' choices of labor effort, which introduces the issue of distortionary income taxation.) The variable $z$ is taken to represent a complete depiction of the regime that prevails in all generations.

As noted, $T_i(y,z)$ denotes the income tax and transfer scheme applicable in each generation. Note first that this is, indeed, taken to be a tax *and transfer* scheme, allowing, for example, net payments to lower-income individuals (such as through welfare programs) as well as transfers across generations (that is, it is not required that the government's budget balances in each generation, but only that there is long-run fiscal balance). Further, in making each generation's tax schedule a function of the regime $z$, it is contemplated that different tax schedules may be employed as part of the shift to another regime.

Let us now consider any reform from an initial regime, $z°$, to some other regime, $z*$. As a general matter, the utility of individuals in various generations will change. These changes, in turn, may differ both across generations (such as when the present generation makes an investment that benefits a future generation) and within generations (those at different income levels may gain or lose by different amounts).

Next, consider a specific (although hypothetical) adjustment to the income tax schedules in each generation, from $T_i(y,z°)$ to $T_i'(y,z*)$, such that the utility of every individual is kept constant. That is, whatever is the effect of the regime shift from $z°$ to $z*$ on the utility, say, of an individual with income $y$ in generation $i$, adjust the tax schedule for generation $i$ at income level $y$ in a precisely offsetting manner. If the individual in question experiences a gain that, when converted to dollars in that individual's generation, is \$100, the individual's taxes would be raised by \$100; if he or she loses by \$50, the individual's taxes would be reduced by \$50.[15] (Note that these adjustments are likely, in general, to depend not only on an individual's generation but

---

[13]   For some analytical purposes, it is useful to let $i$ be a continuous variable. The present formulation allows that interpretation as well (although some of the statements to follow may in that case fail to hold on sets of measure zero, which does not materially alter the conclusions).

[14]   In formal analysis that employs a continuous distribution of income, one would have, for example, a density function $f_i(y)$ and cumulative distribution function $F_i(y)$ for income in each generation $i$ (with the population of each generation normalized to one). This addition allows one to take into account how many individuals there are at each income level, which is important because costs and benefits of projects (regimes) in general vary across the income distribution.

[15]   Formally, the tax adjustments may have income effects; however, it is straightforward how the tax adjustments need to be defined so as to keep utility constant.

also on an individual's income).[16] Such adjustments to each generation *i*'s tax schedule are to be made at every level of income *y*.

Combining the regime shift from $z°$ to $z*$ with the income tax schedule adjustment from $T_i(y,z°)$ to $T_i'(y,z*)$, we can now analyze both the distributive and efficiency consequences of the combined reform. First, by definition (of the income tax schedule adjustment in particular), each individual, regardless of generation, has precisely the same utility before the reform and after. Accordingly, the resulting distribution of well-being both within each generation and between any two generations is identical to what it was previously. That is, the current formulation indeed defines a reform that is intra- and inter-generationally distribution neutral.

What about the efficiency of the shift in regimes? To determine this, consider whether the hypothesized tax adjustment, from $T_i(y,z°)$ to $T_i'(y,z*)$, generates a budget surplus or deficit. How much net revenue is raised or lost? The additional tax paid, or reduction received, by each individual just equals (by construction) the individual's gain or loss from the change, measured in dollars. If we sum these dollar gains and losses—and discount them to a common time period using the market rate of interest—we have a surplus if and only if the conventional cost-benefit test (as modified for intergenerational projects, following section A) is satisfied.[17] That is, if benefits exceed costs, there is a surplus by construction, and if costs exceed benefits, there is a deficit. Hence, efficient reforms generate surpluses, and inefficient reforms create deficits.

If the reform is efficient, therefore, it is possible to further adjust each generation's income tax schedule, now from $T_i'(y,z*)$ to $T_i(y,z*)$, by rebating surplus pro rata to each individual in each generation until the point at which the surplus is exhausted.[18] Now, the first step of the reform—the regime change plus the initial, hypothetical tax adjustment—kept every individual's utility constant, and this second step, a

---

[16]   Concretely, consider the case in which benefits or costs are additively separable in utility; that is, the effect of the reform on the *utility* of individuals in a given generation is the same. To convert those constant utility benefits into dollars, one must divide by the marginal utility of consumption. If the utility of consumption is determined, say, by the natural log of consumption, then the marginal utility of consumption is the inverse of (one over) consumption; hence, an individual with twice the income has one half the marginal utility and thus twice the dollar benefit associated with a given utility increase. Accordingly, such richer individuals would have their taxes raised twice as much to offset the benefit of the reform. As will be elaborated in Part III.C, similar implications arise between generations if some generations are richer than others.

[17]   Any costs or benefits directly to the government fisc may be added to this total or may be imputed to the individuals who would otherwise have their tax obligations adjusted to produce fiscal balance.

[18]   It is contemplated that the rebates be equal in present value, although nominal equality or various other allocations would also generate a Pareto improvement.

pro rata rebate, increases everyone's utility. Hence, every individual in every generation is made better off. That is, when efficient reforms are undertaken in a distribution-neutral fashion, it follows that there is an intra- and intergenerational Pareto improvement.

Similar analysis applies to inefficient reforms, which produce deficits. Obviously, if the deficit is financed by a pro rata tax increase, every individual in every generation is made worse off. Consider, however, a reform in the opposite direction, for example, shrinking rather than expanding the contemplated project. If expansion failed the cost-benefit test, generating a deficit, then a contraction will generate a surplus—that is, will be efficient and permit a Pareto improvement, as just described. The exception is when neither expansion nor contraction creates a surplus, which typically holds when (at the margin) costs just equal benefits, the standard criterion for a project that is of optimal size.[19]

The foregoing analysis shows more precisely the manner in which regimes may be changed in a manner that is distribution neutral—both intra- or intergenerationally—in which case all policies (other than purely redistributive ones, see section E.3) are properly assessed purely in terms of efficiency. This procedure can be seen as underlying the more informal statements in sections A though C that purported to justify a separation of efficiency and distribution.[20]

## E.   Extensions and Qualifications

Despite the apparent generality of the arguments in sections A through C and the more explicit framework developed in section D, it is appropriate to consider both further extensions and possible qualifications. This section will confine attention to three of the most im-

---

[19]   Qualifications pertaining to nonmarginal changes and multiple optima in the presence of nonconvexities are ignored, as the adjustments necessary in light of such complications are straightforward.

[20]   The analysis in the present section is a good deal more abstract and general than that in the preceding sections. For example, projects may have any pattern of effects across any number of generations. The discussions in sections A through C proceeded as if projects have all of their costs in one generation and all of their benefits in a single future generation. However, many projects—including environmental protection, infrastructure, and research and development—undoubtedly involve benefits to many generations, including the present one. In the framework in this section, this sort of complication can be seen to be irrelevant. The distribution-neutral tax adjustment in each generation reflects both the costs and benefits attributable to that generation from any sort of reform whatsoever. With a bit more elaboration, one can also incorporate this consideration into the previous analyses. Indeed, the analysis in section A did not require that benefits or costs were confined to a single generation. For sections B and C, one can consider further adjustments, such as to the level of debt carried between the generations, to provide appropriate offsets. (For example, if a present project designed to benefit a future generation has a positive benefit in the present, the cost can be reduced accordingly. One can also imagine that a Pigouvian subsidy accompanies the project, which, in section C's metaphor of a pool of funds, injects additional revenue into the pool.)

portant, involving labor supply and distortionary taxation, heterogeneity among individuals (aside from differences in generation and income), and the applicability of the analysis to reforms that are not distribution neutral. As will be seen, these considerations—along with numerous others that might have been examined—are not distinctive to the intergenerational context.[21] Indeed, for the most part, they are presented by purely intragenerational policy decisions and are not substantially altered, if at all, by adding an intergenerational dimension. Additional considerations are taken up in Parts III and IV.

1.  Labor supply and distortionary taxation.

The analysis in section D took individuals' levels of (pre-tax) income as given. A more complete formulation of the problem must recognize that individuals' incomes are determined by their labor effort, which in turn is influenced by income taxation (and potentially by projects). Specifically, taxing labor income distorts labor supply downward, so that greater redistribution is accompanied by greater distortion, the familiar equity-efficiency tradeoff. This distortion also applies to intergenerational redistribution.[22]

The problem of labor supply distortion is of particular relevance in the present context because of the possibility that second-best considerations interfere with the separation of distribution and efficiency, the latter now interpreted to concern the efficiency of intergenerational investments (as distinct from efficiency in redistributive taxation per se). Especially since the 1970s, a great deal of attention has been devoted to this question with regard to the financing of public goods and the fiscal consequences of environmental policies, both of which are important in the intergenerational context. Issues addressed include the appropriateness of distributive weights (since although one can redistribute more or less through the income tax, this method is distortionary) and the concern that the financing of government allocative policies may directly supplement pre-existing distortions or indirectly affect their magnitude.[23]

---

[21]   For a catalog and some analysis of many of these factors, see Kaplow, *The Theory of Taxation and Public Economics* at § 6.C (cited in note 12).

[22]   In addition, income taxes and other types of taxation may distort capital—that is, the intertemporal allocation of resources. See Part IV.B.

[23]   On public goods, see, in addition to the sources on distributive weights cited in note 9, work that builds on A.C. Pigou's classic, *A Study in Public Finance* 53 (Macmillan 1928). This includes A.B. Atkinson and N.H. Stern, *Pigou, Taxation and Public Goods*, 41 Rev Econ Stud 119 (1974); Peter A. Diamond and James A. Mirrlees, *Optimal Taxation and Public Production II: Tax Rules*, 61 Am Econ Rev 261 (1971); J.E. Stiglitz and P. Dasgupta, *Differential Taxation, Public Goods, and Economic Efficiency*, 38 Rev Econ Stud 151 (1971); and the commentary in Charles L. Ballard and Don Fullerton, *Distortionary Taxes and the Provision of Public Goods*, 6 J Econ

It turns out, however, that in a basic, benchmark case, labor supply is unaffected when a distribution-neutral income tax adjustment of the sort described in section D is employed. It is useful to review the argument, which extends as well to the intergenerational context.[24] First, recall that when the initial (hypothetical) distribution-neutral income tax adjustment is employed in conjunction with the change from regime $z^o$ to regime $z^*$, each individual in each generation has the same utility. More precisely, the income tax adjustment was constructed so that, in each generation, those earning any level of income had the same utility after the reform, including the income tax adjustment, as they had before.

Consider the implication of this feature for any individual's choice of labor supply. Different choices of labor supply are associated with different levels of (before-tax) income. However, under the present construct, each level of income maps to precisely the same level of utility both before and after the reform (implicitly taking labor supply to be constant). Hence, whatever level of income (determined by a particular level of labor supply) maximized utility initially will continue to do so.[25] Therefore, no one's labor supply will be affected. Accordingly, the conclusions that were valid previously—when effects on labor supply were ignored—will continue to hold.

As stated at the outset, this analysis holds for a benchmark case. One assumption of that case is that individuals within a given generation have the same utility function; that is, there is no preference heterogeneity. This assumption will be considered further in subsection 2. Another assumption is that the mix of private consumption, public goods, and externalities (including environmental amenities) that contributes to a given level of utility does not influence the relative preference for labor versus leisure. To illustrate how this condition may

---

Persp 117 (Summer 1992). On environmental policies, see the survey in A. Lans Bovenberg and Lawrence H. Goulder, *Environmental Taxation and Regulation*, in Alan J. Auerbach and Martin Feldstein, eds, 3 *Handbook of Public Economics* 1471 (Elsevier 2002), and the collection of articles in Lawrence H. Goulder, ed, *Environmental Policy Making in Economies with Prior Tax Distortions* (Edward Elgar 2002).

[24]   This point is central in most of the references cited in note 12. A number of those sources explicitly consider public goods and externalities, as well as different commodities that individuals may choose to consume. Each of these may be interpreted as vectors, with elements corresponding to the levels or effects of public goods, externalities, and commodities in different time periods (generations). Likewise, as presented in section D, one can allow a separate income tax schedule for each period. Making these interpretations and adjustments, the formal arguments presented in prior work would hold true. The requirement in those proofs that individuals have identical subutility functions of public goods, externalities, and commodities—in a setting in which labor is weakly separable—would, with multiple generations, only require this commonality within each generation, which would, for example, allow the effects of an externality or of a public good on utility to vary over time.

[25]   Again, for a more formal statement, see the sources cited in note 12.

not hold, consider environmental improvements that make recrea-
tional venues more attractive. These enhancements may in turn in-
crease the relative desirability of leisure, inducing individuals to work
somewhat less. Given that labor supply is already distorted downward
due to redistributive income taxation, this result is socially costly, mak-
ing it optimal to undertake somewhat less of such environmental im-
provements. However, other environmental projects may make travel
to and time at or near one's workplace more inviting. The effect in this
case would be to induce additional labor effort, a social benefit in light
of the preexisting distortion, so it would be optimal to expand such
projects beyond the level that would otherwise be efficient.[26]

Observe that, in either case, the adjustment does not depend on
the effects arising in a different generation from the project expendi-
tures. That is, the requisite adjustment is one that arises in the purely
intragenerational setting. Additionally, one might view this sort of con-
sideration as going more to efficiency, in that it affects the distortion
associated with the income tax, and, in addition, the nature of the ad-
justment (notably, its direction) does not depend upon the extent of
redistribution undertaken (although the magnitude of the adjustment
generally does).

### 2. Heterogeneity.

The analysis in section D assumed that individuals' utility func-
tions were of the form $U_i(y - T_i(y,z),z)$.[27] Through the subscript $i$ on
the utility function, it was explicitly allowed that individuals in differ-
ent generations may have different utility functions; put another way,
a given regime $z$—including attributes such as levels of environmental
protection and investment in infrastructure—may affect individuals in
different generations differently, as one would expect. Furthermore,
within each generation, utility was taken to depend on income, $y$.
What this formulation implicitly rules out, however, is heterogeneous
preferences within generations and at given levels of income. For ex-
ample, all individuals in a specified future generation who have in-
come of \$100,000 were taken to obtain the same utility under regime $z$.

More realistically, individuals will vary, say, in how much they
value environmental protection. Nature lovers benefit more from con-

---

[26] This idea is originally developed by W.J. Corlett and D.C. Hague, *Complementarity and
the Excess Burden of Taxation*, 21 Rev Econ Stud 21 (1953), in the context of differential com-
modity taxation.

[27] If notation had been introduced for the extension in subsection 1, the utility function
would have had a third argument because utility was taken to depend on both after-tax income
and labor supply, in addition to the regime. This consideration is largely independent of the
present one and thus will be ignored here.

serving natural habitats than do couch potatoes. Accordingly, the claim that adopting more efficient policies is necessarily to everyone's benefit requires qualification. The income tax adjustment, say, for a program of enhanced habitat conservation, will, for future beneficiary generations, reflect only average benefits (measured in dollars) at each income level. Accordingly, efficient policies implemented in a distribution-neutral fashion benefit each income group on average, but do not necessarily help each and every individual.[28] Future couch potatoes who earn $100,000 thus may pay more on account of the tax adjustment than the dollar value of the benefits they receive; likewise, nature lovers will gain even more (they benefit from their pro rata share of the efficiency gain and also because their benefits from conservation exceed average benefits for their income group).[29]

This qualification should have little impact on the analysis. Note that it is unrelated to intergenerational distributive justice. Indeed, it does not even relate to intragenerational distributive justice, as ordinarily construed, for the tax adjustment does achieve distribution neutrality between income groups both within and across generations. The unequal treatment is within income groups that are within the same generation.

Inequality of this sort does tend to be suboptimal.[30] Furthermore, different regimes may increase or reduce inequality of this sort. For example, if nature lovers are already better off (perhaps a side benefit of their taste for nature is better health), then additional conservation increases inequality in individuals' well-being. On the other hand, if the environmental improvements are especially valuable to individuals with unusual sensitivity to pollution, then further protection reduces inequality. For these reasons, some projects should be carried out to a greater or lesser extent than indicated by standard efficiency tests that omit consideration of such heterogeneity. It should be emphasized, however, that this is not a distributive adjustment of a con-

---

[28] Yew-Kwang Ng, *Quasi-Pareto Social Improvements*, 74 Am Econ Rev 1033 (1984), refers to this situation as involving a quasi-Pareto improvement.

[29] Even this residual heterogeneity could be removed if individuals differed in observable ways that would allow more fine-tuned tax adjustments. For example, programs that disproportionally benefit or harm individuals living in certain regions could be accompanied by tax adjustments that vary geographically. Alternatively, user charges could be employed; these would reduce heterogeneity but would introduce distortions in behavior (unless the efficient user charge, equal to full marginal costs, including congestion costs, happened to equal benefits, which would be true for the marginal user but not inframarginal ones).

[30] Such inequity may reduce social welfare if favored individuals have lower marginal utilities of income or if the social welfare function itself is strictly concave in individuals' utilities. This sort of inequity is sometimes viewed distinctly, under the rubric of horizontal inequity, but it is a mistake to analyze such inequity as a separate fount of social welfare rather than as indirectly related to welfare, ordinarily construed. See Louis Kaplow, *Horizontal Equity: Measures in Search of a Principle*, 42 Natl Tax J 139 (1989).

ventional sort and, in particular, has nothing to do with the fact that a policy has intergenerational effects.

### 3.    Reforms not implemented in a distribution-neutral manner.

The preceding analysis, particularly the explicit framework employed in section D, assumes that reforms will be financed in a manner that is overall distribution neutral, combining the effects of the policy change and accompanying tax adjustment. This subsection considers regime changes that are not distribution neutral. (Although this Article is addressed to intergenerational issues, the analysis is largely the same, at least conceptually, if the lack of distribution neutrality is within one or more generations or across generations, so for the most part these dimensions will not be distinguished here.)

First, the distribution-neutral framework helps to illuminate non-distribution-neutral reforms. To see this, observe that any such reform can be decomposed into two components: (1) a distribution-neutral reform, consisting of the underlying policy change and the (hypothetical) distribution-neutral tax adjustment specified in section D, and (2) a purely redistributive reform, one that substitutes the actual method of finance for the distribution-neutral one.[31] This decomposition is helpful because it is clear how each of the components should be analyzed. Component (1) presents an efficiency question, as outlined previously. Component (2) presents a pure question of distributive justice.

It is worth elaborating on the second component. One virtue of isolating the distributive element is that the analysis of it is generic. That is, if an aggregate policy transfers a given amount of resources from some future generation to the present, the same distributive analysis is applicable regardless of whether the source of the intergenerational redistribution is from environmental degradation, failing to maintain infrastructure, or running up a larger national debt. To be sure, the efficiency analysis of such policies may differ, but taking that as given, the distributive analysis is the same. Recognition of this point should clarify many policy debates and facilitate useful academic specialization (environmental specialists can focus on the environment, civil engineers on infrastructure, and others on distribution).

Next, consider how the second component should affect the overall analysis of a reform—assuming, that is, that the change in redistribution is indeed specially linked to the particular policy change. Suppose, for example, that the reform increases redistribution (say, it

---

[31]    This approach is advanced in Kaplow, *The Theory of Taxation and Public Economics* (cited in note 12); Kaplow, 18 J Econ Persp 159 (cited in note 12), and Kaplow, 49 Natl Tax J 513 (cited in note 12).

moves resources from a richer to a poorer generation). This increase would be desirable if the extent of redistribution was too little, undesirable if already too much, and indifferent if just right. Redistribution, after all, is something that should be done optimally, not maximally, because of the tradeoff with distortion. At the redistributive optimum, the benefits of further redistribution just equal the costs of additional distortion. Hence, when at the optimum, small changes are a matter of indifference, and when near the optimum, they matter only slightly. Even when away from the optimum, there is still a partial offset. For example, if redistribution were inadequate, increasing it would be desirable, but the redistributive benefit would nevertheless be accompanied by some increase in distortion.

Finally, although this Article is not primarily concerned with political considerations, it is worth some reflection on the plausibility of distribution neutrality. A conjecture is that, if one had to predict a priori the most likely long-run distributive impact of a policy change, distribution neutrality would be the best guess. The explanation (aside from insufficient reason to make any particular contrary assumption) is that the extent of redistribution that arises in a society typically reflects some balance of political forces. If that balance remains the same, so too would the approximate extent of redistribution. Furthermore, if that balance should change, one would expect the extent of redistribution to adjust accordingly, without regard to whether one or another particular substantive policy is implemented.[32] To repeat, the foregoing is merely a conjecture of what may tend to be true roughly, on average, and in the long run, not a precise description of any particular political reality.[33]

It should also be noted that the conceptual separation represented by the two-part decomposition has a counterpart in terms of real policymaking. Income taxes can be modified in myriad ways, and frequently are. Moreover, when one considers the panoply of policies that may be adjusted to influence the intergenerational distribution of income, it seems clear that a policymaking body could, if it wished, adopt either of the two components independently. That is, if it liked some

---

[32]  For example, if the weight on future generations reflects the extent of altruism toward descendants, one can imagine that, ordinarily, this weighting does not change very much in the near term and, furthermore, that if the weighting does shift, the change may be acted upon in any number of ways (as mundane as reducing or raising national debt, at the cost or to the benefit of currently living generations).

[33]  Competing conjectures are most likely to reflect particular pending proposals. However, the momentary packaging of a proposal with significant intergenerational effects may be a poor guide to its actual, long-run distributive consequences. Pre-enactment packaging often changes, implemented reforms are subsequently amended, and other components of the system may ultimately adjust in reaction to a redistributive shift that is contrary to the balance of political forces.

substantive policy but had no intention of changing distribution, it could implement that policy in a distribution-neutral fashion. And if it wished to change the extent of redistribution but found to be inefficient a substantive policy that, in isolation (without an offsetting tax adjustment), had the desired distributive effect, it could reject the policy and achieve the desired shift in redistribution directly.[34] An additional, optimistic note is that if more attention were devoted to formulating distribution-neutral policy packages, policymaking might be improved. After all, efficient distribution-neutral policies benefit everyone (or at least each income group in each generation) and thus may be likely to garner support, and inefficient distribution-neutral policies are to everyone's detriment and thus should be difficult to implement. Of course, politics does not always work this way, in part because some costs and benefits are harder to perceive—and for other reasons.

One additional aspect of the political dimension is that some policies with intergenerational consequences impose costs and benefits that cross national boundaries. (This is true of purely intragenerational policies as well.) At this point in time, there does not exist an apparatus to provide for smooth distributively offsetting adjustments. There is no international tax and transfer authority and no multinational mechanism that can readily carry out international transfers of various sorts. Some side payments are made, and progress in reaching international agreement often proceeds bilaterally or regionally. In any event, international distributive effects may significantly influence the second, distributive component of the decomposition.

---

[34]   It is sometimes suggested that a society wishing to benefit distant future generations, say, through reducing emissions that contribute to global warming, may not find it feasible to benefit those generations through alternatives such as ordinary capital investments because interim generations may consume the proceeds. See, for example, Arrow, et al, *Intertemporal Equity* at 132 (cited in note 2); Robert C. Lind, *Analysis for Intergenerational Decisionmaking*, in Portney and Weyant, eds, *Discounting and Intergenerational Equity* 173 (cited in note 2). The concern is that subsequent generations will be less generous toward the future, even as that future becomes closer in time. Specific reasons for this supposition are not offered, but, assuming it to be true, it is hardly clear that choosing one rather than another present policy to benefit the distant future avoids this problem. Just as proceeds from capital investments can be consumed rather than reinvested, so public debt can be run up, conserved resources can be utilized, and even many environmental protections can readily be reversed or allowed to lapse. (In the case of the latter, for example, if the interim generation does not care so much about the future, it can expend less to further restrict emissions or to control new sources). Note also that it is not necessary for a relatively selfish interim generation to benefit itself at the expense of the future through reversing the precise policy by which a previous generation meant to benefit the distant future. The interim generation may be expected to choose whichever mix of policies would provide it the greatest benefit at a given cost to the future. If restrictions in the present generation are not in fact the most efficient way to generate future benefits, there will be less overall for the interim generation to divide between itself and the future. Accordingly, for a given weighting of the two generations, the result of the interim generation's actions will most plausibly make the distant future generations even worse off than if the present generation simply adopted the most efficient policies going forward.

## III. Implications of Intergenerational Distributive Justice for Intergenerational Efficiency

The central claim of this Article is that, in principle, intergenerational efficiency should be pursued independently of how questions of intergenerational distributive justice are answered. This Part explores some interdependencies between efficiency and distribution, none of which, it turns out, disturbs the original claim.

### A.  Efficient Policies May Depend on Distribution

It is familiar that, in a setting with only a single generation, which policies are efficient depends on the distribution of income. Even in the first-best world of the two fundamental theorems of welfare economics, the second theorem holds that for any distribution (specifically, for any Pareto optimal allocation), there exists a set of prices that can sustain the allocation as a competitive equilibrium. In general, different allocations give rise to different relative prices. Prices depend on distribution because the relative demand for different goods depends on income, and because preferences may vary among individuals, as noted in Part II.E.2.

Efficient government policies likewise depend on prices. The cost of providing a public good or cleaning the environment depends on input prices. Moreover, individuals' valuations of public projects depend on their incomes.[35] To take one example, valuations of at least some environmental amenities are generally believed to rise disproportionately with income; if so, a more unequal distribution of income would be associated with a greater marginal social valuation for environmental improvements.[36]

This basic principle is equally applicable in the intergenerational context.[37] In addition to the dependence of project costs and benefits on relative prices (including prices in between generations), preference heterogeneity obviously matters since the costs and benefits of projects with intergenerational effects are borne by different generations. If a project's benefits rise at all with income (even less than proportionately), a distribution of income more favorable to a future generation will tend to make projects that benefit the future more

---

[35]  See, for example, note 16.

[36]  Suppose that individuals' marginal valuations are proportional to the square of income. If each of two individuals has income of $y$, the sum of their marginal valuations is (proportional to) $2y^2$. If instead their incomes are $y - a$ and $y + a$, the sum of their marginal valuations is (proportional to) $2y^2 + 2a^2$, which is increasing in $a$ and thus in the extent of inequality.

[37]  One price of particular relevance in the intergenerational setting (but also relevant intragenerationally) is the interest rate, which itself may depend on the intergenerational distribution of income. See Part IV.C.

valuable. One way of viewing the matter is to note that the pool of resources available for a future generation (see Part II.C) is larger, so in general more will tend to be allocated for most uses, just as a single individual who receives a larger income will typically buy more of many goods. Furthermore, because relative valuations tend to change with income, the efficient relative allocation of resources across projects will depend on the intergenerational distribution. Taking a concrete example, if a future generation is richer, it will place a higher value on clean air and water, and the most efficient way of producing environmental improvements for that generation will involve some further expenditure toward that end in the prior generation.

For the most part, the analysis in Part II took the intergenerational distribution as given and asked what intergenerational policies were efficient. The valuations of costs and benefits to present and future generations were in all cases based on the rates of conversion between utility and dollars within each generation. This was explicit in Part II.A; Part II.B employed the same procedure as did II.A; as just noted, the use of the pool metaphor in Part II.C entailed the use of the future generation's level of income in evaluating its benefits; and the tax adjustments for each generation in Part II.D were set equal to individuals' dollar valuations of costs and benefits and therefore depended on the income levels of individuals in the pertinent generation. The present discussion merely highlights the point that those valuations depend on the overall level of income in each generation (as well as the distribution of income within the generation) and thus may be different if the distribution is different. But as long as those income-dependent valuations are used, the relevant impact of distribution on efficiency is taken into account.[38]

Put another way, arguments that we should, for example, expend more to protect the environment for the future because we should care more about future generations are in one important sense correct, despite the proposed separation of intergenerational distributive justice and efficiency. If the future matters more, resources in general should be allocated more in favor of those who will live in the future. The optimal allocation of those additional resources will, as noted, plausibly include a greater allocation to environmental protection. This result reflects an ordinary income effect. The same logic suggests that more resources should be allocated to other intergenerational

---

[38]   A further subtlety is that changing one project may affect the efficiency of another project, both because of utility interdependencies and because funding more or less of one project may change the marginal utility of income and thus the valuation of another project. These points affect the outcomes of proper efficiency analysis but not the merits of basing policy choices on efficiency.

investments, such as infrastructure and research and development, whether for medical research or improving the entertainment value of toys, for most such things will be valued more in the future if future generations are to receive a greater allocation of resources. It may well be that, as incomes rise, relatively more would be allocated to the environment than to other goods, but to the extent that this is true it will automatically be reflected in an efficiency analysis that employs dollar valuations from the generations that bear the costs and benefits of various possible projects.

### B.   Private Intergenerational Transfers, Liquidity Constraints, and Intergenerational Redistribution

Although this Article is only indirectly about intergenerational redistribution, it is useful to explore certain of its properties that may have efficiency implications (aside from the standard point that redistribution per se tends to be distortionary).[39] A central consideration is the feasibility of intergenerational redistribution in light of offsetting private behavior.

The notion of Ricardian equivalence, advanced by Robert Barro, holds in its purest form that changes in government intergenerational transfers will be matched by equal and opposite changes in private intergenerational transfers.[40] Suppose that members of the present generation (parents) are altruistically linked to corresponding members of the subsequent generation (children). In simple models of altruism, the parents' aggregate utility is taken to be a function of two components: utility from their own consumption and their children's utility. Such altruistic parents will allocate resources between themselves and their children to equate the marginal contributions to their aggregate utility from each component. Suppose this involves a net transfer of $X$ to their children. Now if for some reason the parents discovered that they had one dollar less and that their children had one dollar more to begin with, the parents would instead choose a

---

[39]   Note that if only intergenerational distributive justice (if it could be defined independently of intragenerational distributive justice) were a concern, redistribution could be implemented in a nondistortionary fashion. If the only relevant feature determining the quantity and direction of taxes and transfers was the generation to which one belonged, then payments would be solely a function of an observable and immutable characteristic, in which case no behavior would be distorted. However, when inter- and intragenerational distributive justice are combined, it is likely that a change in the extent of intergenerational redistribution would optimally involve some changes in marginal tax (inclusive of transfer) rates within generations, which would influence the extent of distortion.

[40]   See Robert J. Barro, *Are Government Bonds Net Wealth?*, 82 J Polit Econ 1095 (1974), building on the model of altruism in Gary S. Becker, *A Theory of Social Interactions*, 82 J Polit Econ 1063 (1974).

transfer of $\$(X - 1)$, producing the same net result as they had (optimally) obtained previously.

Government intergenerational transfers have this feature. Imagine that the government in some fashion transfers an additional per capita dollar from the current generation to the next. For example, it might raise current taxes to retire some of the government's debt, so less needs to be repaid by the next generation. In response, this theory implies that parents would transfer one dollar less, leaving each generation at the same level of income it was at previously. Accordingly, the government's attempt at intergenerational redistribution would have failed. Similar logic applies if the government were to raise taxes to fund infrastructure or environmental protection designed to benefit the next generation. In this case, however, the offset would involve the future generation receiving fewer dollars but more infrastructure or a cleaner environment than it received previously, which raises a question of efficiency.[41]

If intergenerational redistribution were impossible, intergenerational distributive justice would be moot, but it would remain true that intergenerational efficiency should be pursued. Indeed, the constructs employed in Part II to hold the intergenerational distribution of income constant would be unnecessary if this result would be produced automatically by adjustments to private transfers.

It is generally thought, however, that although there may be some offsets, Ricardian equivalence does not hold. Not everyone has children, destroying complete linkage. Not all parents are altruistic (in the formal sense); that is, they may be motivated to make transfers to children, but not according to the utility function described above. And not all altruistic parents will operate in the relevant range; for example, if their children are sufficiently well off relative to themselves, they may make no transfers despite their altruism, so they cannot reduce transfers to offset the government's increase in intergenerational redistribution.

A further point of interest is that changes in intergenerational redistribution can have efficiency implications when private transfers respond. For example, suppose that some parents provide their marginal transfers in the form of human capital, by devoting time and resources to their children's education. For them to reduce transfers may require reducing this sort of investment even if its return exceeds the market interest rate. This result could be avoided if children could essentially borrow the funds from their parents to finance efficient

---

[41] There is also a subtle feedback on the extent of net transfers, for the incremental benefit through greater infrastructure or environmental protection may not reduce the future generation's marginal utility of income to the same extent as did the dollars the parents would otherwise have transferred.

investments in their own human capital, but such contracts are not generally feasible.

The primary relevance of the foregoing discussion concerns the character of proper efficiency analysis of intergenerational policies, not whether efficiency analysis should be conducted. Notably, the return to human capital investments may exceed the ordinary market return in equilibrium due to liquidity constraints. Accordingly, allocating more of the pool of resources available for a future generation to education and the like rather than to ordinary capital investments and other projects may be efficient. Likewise, this reallocation may be facilitated by providing loans to younger generations to avoid the detrimental impact of liquidity constraints on investment in human capital (which allows the members of the next generation themselves to choose to allocate more resources in this fashion). Of course, most governments pursue such policies, although there are many other motivations that may help to explain this activity.[42] In addition, to the extent that such policies do not fully remedy the problem, the optimal intergenerational distribution may be affected, for the marginal welfare gain from redistributing dollars to generations containing liquidity-constrained children may accordingly be greater.

Finally, it is worth observing that the present issue can be viewed in purely intragenerational terms. To the extent that some individuals are liquidity constrained, the marginal social value of a dollar to them—at the times of their life during which they are liquidity constrained—is greater. This is true whether liquidity constraints prevent investments in human capital, in entrepreneurship, or simply in consumption smoothing (borrowing from future, higher-earning years to raise consumption in earlier, lower-earning years). Accordingly, greater intragenerational redistribution than otherwise may turn out to be optimal (assuming that lower-income individuals are more likely to be liquidity constrained); changing the timing of tax obligations over the lifecycle may be advantageous; and simple projects such as spending more on a park may be less attractive because some who would have to pay more in taxes have a higher marginal utility of income.[43]

---

[42] For example, education may be thought to have social spillovers, so that the social return exceeds the private return. Additionally, some (or most) parents may be thought to allocate too little to their children's futures due to misinformation, myopia, or self interest, making in-kind provision by the government optimal.

[43] See Glenn R. Hubbard and Kenneth L. Judd, *Liquidity Constraints, Fiscal Policy, and Consumption*, 1 Brookings Papers on Econ Activity 1 (1986) (with comments by Robert E. Hall and Lawrence Summers); Glenn R. Hubbard and Kenneth L. Judd, *Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax*, 77 Am Econ Rev 630 (1987); A. Mitchell Polinsky, *Imperfect Capital Markets, Intertemporal Redistribution, and Progressive Taxation*, in Harold M. Hochman and George E. Peterson, eds, *Redistribution*

## C.  Relevance of Greater (or Lesser) Wealth of Future Generations

It is often hypothesized that future generations will be wealthier on account of continuing technological advances, in which case intergenerational redistribution should be favorable to present generations.[44] This justifies less saving (running up greater debt) and less effort to further research and development, enhance infrastructure, protect the environment, and maintain resource stocks. The socially optimal mix among these policies is dictated on pure efficiency grounds, as developed in Part II. Indeed, the same criterion is appropriate under the opposite assumption, that future generations will be worse off, which may arise due to overpopulation, pollution, and the greater lethality of future risks of war and terrorism.[45]

These familiar ideas, however, may not capture the entire phenomenon. Consider, for example, the prospect that enhanced knowledge will improve the well-being of future generations. Furthermore, suppose in particular that the gain will accrue due to more effective use of given resources. One way to think of this is that a dollar will buy more effective consumption. (Consider how the value per dollar of expenditures on personal computers has grown in the past few decades.) In that case, there are two competing effects. First, more should be allocated to generations with better technology because each dollar goes further. Second (the more familiar point), less should be allocated because, since those generations have higher effective wealth for a given allocation, their marginal utility is less. Which factor is greater is an empirical question; plausibly, the second is greater, although it must be kept in mind that a partial, possibly substantial offset is provided by the first.[46]

*Through Public Choice* 229 (Columbia 1974); Kaplow, *The Theory of Taxation and Public Economics* (cited in note 12).

44    See also Part V. The extent depends on the concavity of the social welfare function. For example, under a utilitarian formulation, richer individuals have a lower marginal utility of income and, to that extent, should be taxed more in order to finance transfers to poorer individuals. Further preferences for equality, favored for example by Rawls, *A Theory of Justice* (cited in note 2) and Amartya Sen, *On Economic Inequality* (Oxford enlarged ed 1997), would go further than this (assuming that incentive considerations cause utilitarian redistribution to stop short of egalitarianism).

45    An important distinction arises, however, if adverse outcomes reduce population size or the likelihood that life will continue to exist. In that case, many accept that there should be a probability discount on the future—not reflecting any difference in the value of life or utility but instead reflecting the reduced likelihood that the value will be realized. See, for example, Parfit, *Reasons and Persons* at 482 (cited in note 2).

46    A variant of this problem is modeled in Kaplow, *The Theory of Taxation and Public Economics* at § 12.A.2 (cited in note 12), and Louis Kaplow, *Optimal Distribution and the Family*, 98 Scandinavian J Econ 75 (1996), for the case of economies of scale in family consumption. Suppose, for example, that individuals in the better-off generation have utility of income given by $u(\alpha y)$ rather than by $u(y)$, where $\alpha > 1$. Then, the marginal utility of income is given by $\alpha u'(\alpha y)$

It has also been suggested that future generations' greater wealth implies that they will have higher marginal valuations, for example, higher values of statistical lives (VSL), which will offset what would seem to be the lower valuations on future lives that result from discounting dollars.[47] For example, if the VSL is $5 million in the present and in the future, then the discounting of dollars would imply that a future life is only worth $1/(1 + R)$ as much as a present life. If future generations are richer, so that their VSL is higher, say, by a factor of $(1 + R)$, then discounting dollars implies equivalent weights on lives. (Note that this logic also implies that, if future generations were poorer, their lives would count less for two reasons—their lower VSLs and the discounting of dollars—although it should be kept in mind that the discount rate may differ in this scenario.)

Whether and the extent to which future generations have higher VSLs depends on their relative wealth, which in turn is a function of the intergenerational distribution that is chosen.[48] To take a simple and highly stylized example, suppose that there is a fixed pot of resources in the present that is to be divided between two individuals, one living presently and one in the future, and that the intergenerational interest rate is $R$. Consider an equal division of the pot. The half that is saved for the future individual will, when the future arrives, be worth $1 + R$ times as much (in that period's dollars) as the half was worth to the present individual (in that period's dollars). Now, assume further that, when resources are larger by a given factor, then VSL is larger by precisely the same factor. In that case, an equal division would indeed produce VSLs that were equal in discounted present value. Moreover, it can be shown that, under these assumptions about the relationship between individuals' utility functions (which determine VSLs), the postulated equal division of the initial pot is the optimal distribution from a utilitarian and other standard welfarist perspectives.[49]

---

rather than by $u'(y)$. In the case in which utility is given by the natural log of income, marginal utility is the inverse of (effective) income, so the first expression equals $\alpha(1/\alpha y)$ and the second is $1/y$; obviously, these values are equal, which is to say that the two effects are precisely offsetting. If utility were more concave (which is equivalent to stating that the coefficient of relative risk aversion exceeds one, as is commonly supposed), then the first marginal utility would be less than the second. A more concave social welfare function would have similar implications.

[47]    See, for example, Peter Singer, *One World: The Ethics of Globalization* 25–26 (Yale 2d ed 2004).

[48]    Furthermore, as explored in Part IV.C, the intergenerational distribution also may influence $R$.

[49]    Just as in note 46, the pertinent relationship is that utility is a function of the natural log of resources. To translate the analysis there to the present setting, let $\alpha = 1 + R$. Then if the same present-period income is allocated to each individual, call this amount $y$, the previous expressions apply, in which case effective marginal utility is equalized. Also by analogy to the preceding case, greater concavity in the utility function would imply that effective marginal utility is equalized by a division somewhat more favorable to the present generation.

This latter discussion indicates that, not surprisingly, valuations of projects (whether involving different risks to life or otherwise) will depend, as section A states, on the distribution of income. As before, however, whatever that intergenerational distribution is or should be, if it is taken as given, projects should be chosen on grounds of efficiency.

## IV. MARKET INTEREST RATE

Following the prescription of Part II that future project costs and benefits, measured in dollars, be discounted to the present, it is necessary for the assessment of intergenerational efficiency to determine the proper market interest rate. This choice is not straightforward, in part because there is not a single such rate that is appropriate in all circumstances.[50] Or, to put the point differently, one can choose a single benchmark rate, but for some projects it may be necessary to make certain adjustments. There are also additional complications, some of which are also considered in this Part.[51] As with much of the analysis in this Article, however, most of the refinements are not unique to the intergenerational context. That is, even if the discounting were only over a single year or if a long-term investment had purely intragenerational effects (consider young workers setting aside funds for retirement or possible nursing home care), an interest rate would have to be chosen and most of the factors considered here would have the same implications as in the intergenerational context. Indeed, as noted previously, dollars invested, say, in a machine that will produce goods in the future (which in turn will be sold for dollars) neither know nor care (if they were capable) whether the proceeds are to be distributed to individuals of an earlier generation (perhaps at retirement) or a later one.

### A. Uncertainty

The market interest rate—the private return to capital—is taken to be the relevant interest rate for purposes of discounting dollar costs

---

[50]   Actually, if there were a single interest rate relevant for all purposes, the implication of choosing the wrong rate would be entirely distributive. For example, choosing too high a rate would undervalue everything occurring in the future to the same extent. That is, placing too high of a discount rate on dollars is equivalent to excessively discounting utility (relative to whatever is the appropriate normative rate).

[51]   The list here does not attempt to be exhaustive. Among others, macroeconomic and general equilibrium effects of public projects are not considered; nor are (positive or negative) externalities to consumption growth. Most pertinent issues have been addressed in a substantial literature. See, for example, Arrow, et al, *Intertemporal Equity* (cited in note 2); Coleman Bazelon and Kent Smetters, *Discounting Inside the Washington D.C. Beltway*, 13 J Econ Persp 213 (Fall 1999); Lind, et al, *Discounting for Time* (cited in note 4); William D. Nordhaus, *Managing the Global Commons: The Economics of Climate Change* ch 6 (MIT 1994).

and benefits from public projects because market investments are viable alternatives. As the analysis in Part II.B explains, if a public project is attractive when discounting at the market rate, it can be increased and market investment reduced in an amount sufficient to finance the project, producing a net gain. Likewise, if the public project is unattractive at this rate, reducing its scope will free resources that can be invested in the private capital market for a better return.

The market interest rate, as is well known, contains a premium for systematic (market-wide) risk.[52] It can be thought of as a riskless rate (representing the pure time value of money) and an additional component to compensate for riskiness. This feature raises two issues.

First, not all projects have the same risk profile. In the capital market, investments that are negatively correlated to the market as a whole are more attractive in terms of their impact on the owner's risk and thus pay a lower return, whereas those that vary with and to a greater extent than the market must pay a higher return. Likewise, in determining the appropriate discount rate for public projects, it is appropriate to employ an interest rate that reflects the relationship between the project's risks and those of the economy as a whole.[53]

Second, it is sometimes suggested that there should be no risk adjustment for public projects—they should be discounted at the risk-free rate of return—because the government is able to pool risks widely.[54] It is well understood, however, that this view is based on assumptions that do not generally match reality. The government can, at best, achieve diversification across the economy as a whole. Accordingly, it should ignore idiosyncratic project risks, which are thereby eliminated. Of course, with diversified ownership (directly, or through insurance or other arrangements), idiosyncratic risks in the private market do not ordinarily command a risk premium either. Regarding economy-wide risks, however, the government should in principle employ a discount rate reflecting systematic risk to the extent that the project in question involves such risk, as discussed in the preceding paragraph. A possible exception concerns risks that span generations, where the government may have some advantage relative to incom-

---

[52] Nominal interest rates also include a component to offset the effects of inflation. The discussion in the text should be interpreted as focusing on real interest rates.

[53] See, for example, Arrow, et al, *Intertemporal Equity* (cited in note 2); Martin J. Bailey and Michael C. Jensen, *Risk and the Discount Rate for Public Investment*, in Michael C. Jensen, ed, *Studies in the Theory of Capital Markets* 269 (Praeger 1972); Lind, *A Primer on the Major Issues* (cited in note 4); Agnar Sandmo, *Discount Rates for Public Investment Under Uncertainty*, 13 Intl Econ Rev 287 (1972); Robert Wilson, *Risk Measurement of Public Projects*, in Lind, et al, *Discounting for Time* 205 (cited in note 4).

[54] See, for example, Kenneth J. Arrow and Robert C. Lind, *Uncertainty and the Evaluation of Public Investment Decisions*, 60 Am Econ Rev 364 (1970).

plete private markets, which may influence the appropriate discount rate for projects spanning into the distant future.[55]

In light of the foregoing points, the government could proceed in a number of ways.[56] It could employ as its standard discount rate the market rate, which includes systematic risk, with adjustments to the extent that the government's project should carry a higher or lower risk premium. Or it could use as a benchmark the riskless rate, with adjustments (more often upward, and to a greater degree) from that starting point. If it were the case that the average government project had a profile similar to that of the average private project, it is probably more sensible to use the market rate, inclusive of the standard risk premium, as the default. Another approach would be to employ different starting points for different agencies or different types of projects. In particular, certain projects—like strategic reserves, protection against large-scale natural disasters, and provisions for national defense—might ordinarily tend to offset economy-wide risk; for them, a benchmark below the riskless rate might be sensible.

An additional type of uncertainty concerns interest rates themselves. Especially when discounting over long periods of time (beyond which even long-term financial instruments do not provide much of a prediction of future interest rates), interest rate uncertainty can have a substantial impact on how discounting should be done. The main reason is that it is inappropriate simply to employ the mean predicted interest rate. When compounding over long periods, discounting at rates that may be high or may be low, taking the probability-weighted average of the outcome, does not produce results equivalent to discounting at the average rate. Instead, uncertainty implies a substantially lower rate of discount in the long run.[57]

---

[55]  See, for example, Robert J. Shiller, *Social Security and Institutions for Intergenerational, Intragenerational, and International Risk-Sharing*, 50 Carnegie-Rochester Conference Series Pub Pol 165 (1999).

[56]  An additional complication is that, with more than two periods, a single discount rate does not ordinarily provide an appropriate risk adjustment unless the relationship between risk and return follows a particular time pattern (specifically, the ratio of the certainty equivalent to the expected value must be falling at a certain rate). See, for example, Lind, *A Primer on the Major Issues* at 63–67 (cited in note 4); Joseph E. Stiglitz, *The Rate of Discount for Benefit-Cost Analysis and the Theory of the Second Best*, in Lind, et al, *Discounting for Time* 151, 197 (cited in note 4). It is appropriate more generally to convert each period's returns into a certainty equivalent and then discount these (at the riskless rate). For any particular project, one can always use this method to determine a single discount rate that produces the same result, but the proper rate is project specific and can only be determined as a byproduct of determining the proper result in the first instance.

[57]  See Richard Newell and William Pizer, *Discounting the Distant Future: How Much Do Uncertain Rates Increase Valuations?*, 46 J Envir Econ & Mgmt 52 (2003); Martin L. Weitzman, *Why the Far-Distant Future Should Be Discounted at Its Lowest Possible Rate*, 36 J Envir Econ & Mgmt 201 (1998). Taking a simple example, if $1.00 ten years in the future is discounted to the

Another factor is that uncertainty regarding the well-being of fu-
ture generations may bear on the optimal distribution between gen-
erations. Taking the simple case in which income of a subsequent gen-
eration is uncertain, it follows that transfer of a certain dollar will have
a higher expected marginal payoff than otherwise. The principle is the
same as that involved in determining optimal precautionary savings
for a single individual facing earnings uncertainty.[58] Analogous reason-
ing may apply to uncertainty regarding public projects where effects
are nonlinear (such as is the case with the marginal utility of income).[59]

## B.   Taxation of the Market Return to Capital

It is common for income and other taxes to impose some burden
on the market return to capital.[60] This causes a distortion because the
full, pre-tax return to capital (which reflects the marginal product of
capital) is not realized by investors, who receive only the lower, after-tax
return. This presents the question of whether the full, pre-tax return, the
lower, after-tax return realized by investors, or some weighted average
of the two should be employed in evaluating government projects.[61]

There is no simple answer. On one hand, the full, pre-tax return is
the social return, which seems to make it appropriate to discount at that
rate. Moreover, the government's use of the same rate as the private

---

present at a 5% rate, the value is approximately $0.61. If it is discounted at 10% the value is
approximately $0.39 and at 0% the value is $1.00, the latter two averaging to approximately
$0.69, which is greater than $0.61, the result of discounting at the average rate of 5%. (Discount-
ing at about 3.8% would produce the result of approximately $0.69.) The intuition is that higher
rates have an ever-decreasing relative impact on the result as the time period increases, so low
rates have a greater effect. Note that the phenomenon illustrated here would be much greater if
the time period were substantially longer.

[58]   For standard utility functions, reducing income by a given amount raises marginal utility
by more than increasing income by the same amount reduces marginal utility; hence, increased
uncertainty increases the average, or expected, marginal utility of income.

[59]   For example, if the amount of future environmental degradation is uncertain and the
harm of degradation rises nonlinearly, then more present protection is optimal than if the ex-
pected (average) degradation were the same but certain.

[60]   For a pure accrual income tax imposed at the rate $t$, an annual before-tax interest rate $r$
yields the after-tax return $(1 - t)r$, which can reduce the long-term return substantially when this
reduced rate is compounded. In addition, many countries also tax capital held in the corporate
form. However, because most income taxes are not pure accrual taxes, sometimes providing
reduced rates for dividends and capital gains, taxing on a realization basis (with possible forgive-
ness at death), and exempting some types of investments and certain investors (creating oppor-
tunities for arbitrage, especially given the tax deductibility of interest), the effective tax rate on
capital tends to be much lower. See, for example, Roger Gordon, et al, *Toward a Consumption
Tax, and Beyond*, 94 Am Econ Assoc Papers and Proceedings 161 (2004). (Since relief is pro-
vided unevenly, however, much interasset distortion may nevertheless result.)

[61]   See, for example, Kenneth J. Arrow, *The Rate of Discount of Public Investments with
Imperfect Capital Markets*, in Lind, et al, *Discounting for Time* 115 (cited in note 4); Arrow, et al,
*Intertemporal Equity* at 135–36, 138 (cited in note 2); Lind, *A Primer on the Major Issues* (cited in
note 4); and Stiglitz, *The Rate of Discount* (cited in note 56).

sector avoids biasing government-private decisions.[62] On the other hand, if the government does undertake a project that looks favorable when benefits are discounted at the lower, after-tax rate but not at the higher, before-tax rate, individuals would tend to view such decisions positively, since they discount benefits at the after-tax rate.[63]

A commonly proposed solution is to discount at a weighted average rate. A higher weight is to be applied to full, pre-tax return to the extent that the social investment project is expected to displace private investment (since then it is the full return on private investment that is forgone by society). Correspondingly, to the extent that consumption is what is displaced, greater weight would be accorded to the investors' after-tax return, the rate at which individuals discount consumption.

Although the literature does not usually pursue the matter in depth, it seems that the most plausible conjecture is that investment (savings) would be displaced.[64] Suppose that the government financed the project by raising present taxes and that individuals reduced present consumption to cover this obligation.[65] The proceeds are assumed to finance a public project that provides future benefits, which may be viewed as augmenting future consumption.[66] If individuals had previously allocated their consumption across periods (chosen savings rates) in a privately optimal fashion (given the interest rate they face), they would now find that they would have too much future consump-

---

[62] Indeed, an argument for using the full, pre-tax return is that government projects with a return below that rate are dominated by the government using the same funds to invest in marginal private projects, which by definition pay the full, pre-tax return. (Such investment could be undertaken directly, or through taking an equity stake, which might be held passively.) It has further been recognized that, taken to its limit, the government's ability to compete for private projects without having to pay taxes on its returns could lead to full government ownership of the economy, a result generally taken to be inefficient. See, for example, Lind, *A Primer on the Major Issues* at 32 (cited in note 4).

[63] This description applies to investments with current costs and future benefits. The logic is analogous for the opposite pattern, such as when greater use of natural resources or degradation of the environment produces immediate benefits but long-term costs.

[64] Lind, *A Primer on the Major Issues* at 45 (cited in note 4), briefly notes that, in a model in which individuals are simply stipulated to save a given percentage of their resources, a public project providing future benefits might be perceived as part of such savings, although he does not adopt this view in the remainder of his discussion.

[65] Alternatively, suppose that the project is financed by public debt, which is to be paid off (inclusive of interest) at the time the benefits are conferred. In that case, some of individuals' savings would have to be allocated to holding the debt, displacing private investment. By reasoning analogous to that in the text, individuals would find the resulting allocation of consumption over time optimal, so they would not be motivated to increase their savings.

[66] Whether public projects are a substitute for consumption is an empirical question that depends on how the projects enter individuals' utility functions. Public projects that do not have the effect of substituting for consumption may well have a different effect on the marginal utility of consumption and thus on savings than that hypothesized in the text. This phenomenon, however, could also arise with a project producing present benefits.

tion and too little in the present. To restore the prior balance, they would need to reduce present savings by the amount of the present tax.[67]

The first-best solution involves no market distortion of investors' decisions. If implemented, there would be no need for the government to choose between before- and after-tax interest rates. If not, there seems to be a good case for choosing the before-tax rate. In any event, the question is essentially the same if the problem is purely intragenerational, involving a generation's investments (in the private or public sector) for its own future benefit.

## C.  Endogeneity

Until now, the market interest rate was taken to be fixed. Its level, however, may be influenced by intergenerational policies. For example, if the current generation consumes more, saving less privately and running up a larger public debt, interest rates will tend to be higher.[68] Indeed, this point is merely a special case of the phenomenon discussed in Part III.A that prices depend on the distribution of income since, as previously noted, the market interest rate is one such price.

This interdependency, however, has no impact on the analysis thus far. As described previously, the fact that the efficient outcome may depend on the distribution does not disrupt the argument that the choice among projects, once a distribution is chosen, should be made on grounds of efficiency. Moreover, as emphasized in much of Part II, the present analysis ordinarily takes the intergenerational distribution to be held constant. To the extent this is so, then the market interest rate is indeed constant. (And, even if it were not, few projects

---

[67]  Note further that the break-even investment for the government would be one that paid the full, pre-tax rate of return, not one that paid the lower, after-tax return. Although only an amount reflecting the after-tax return would have to be bestowed on individuals to equal the return that they would realize on equivalent savings, the remaining amount (the difference between the before- and after-tax return) would be required to cover the forgone tax proceeds. (The difference between the before- and after-tax return is simply the tax levy on the return, and this amount no longer flows into the government fisc on account of individuals' reduction in savings.) In addition, the same logic applies to debt finance, considered in note 65. Public debt would have to be offered at the full, pre-tax rate of return; otherwise, individuals would not choose to hold it. (The debt could be offered at the after-tax interest rate if such interest were tax exempt; however, what the government would save on interest payments would be fully offset by what the government would lose from tax on the interest.) Therefore, the public project would need to earn that rate of return to be sufficient to finance repayment of the debt (with interest).

[68]  In a small, open economy, world capital markets fix the interest rate. For large economies like the United States, domestic activities may influence the interest rate; likewise, if there is an intergenerational shift worldwide, interests rates will change. Additionally, in practice domestic interest rates are not entirely determined by world markets. See Martin Feldstein and Charles Horioka, *Domestic Saving and International Capital Flows*, 90 Econ J 314 (1980); Maurice Obstfeld, *International Capital Mobility in the 1990s*, in Peter B. Kenen, ed, *Understanding Interdependence: The Macroeconomics of the Open Economy* 201 (Princeton 1995).

would be of such a scope to have a substantial impact on the interest rate, one sufficient to influence significantly which projects should be undertaken.[69])

A subtle complication is that the chosen mix of projects also could influence interest rates. For example, spending more on infrastructure or on certain types of research and development might increase the market return to private investment, bidding up interest rates. Contrariwise, devoting more of the resources dedicated to benefit future generations to environmental preservation might reduce the market return. In any case, it remains true that the effects of shifting a few projects (short of radical, worldwide shifts) are likely to be very small, and whatever they may be, projects should be assessed on efficiency grounds in terms of whatever interest rate actually will prevail.[70]

## V. MARKET INTEREST RATE VERSUS EXPLICIT SOCIAL WEIGHTINGS

Many economists ground their view toward the discounting of dollars by asking explicitly how society should weight dollars in different time periods. Under this approach, the social discount rate on dollars is the sum of two components. The first is a pure rate of time preference, indicating how units of utility should be compared across periods. (For those who believe utility should not be discounted, this component would be set equal to zero.) The second involves the conversion of dollars to utility in different time periods.[71] It is ordinarily expressed as the product of rate of growth of per capita consumption and the rate at which marginal utility declines with consumption (expressed as a positive term). The intuition behind the latter component relates to the analysis in Part III.C: the richer (or, if the growth rate is negative, poorer) are future generations relative to the present, and the faster is the decline in marginal utility as the standard of living increases, the less valuable is a future dollar compared to a present dollar, from a social point of view.[72]

---

[69] A very small change in interest rates would typically have a very small impact on the optimal extent to which a project should be carried out.

[70] In this respect too, the interest rate can be viewed as just another price. For example, if more is spent on roads, the demand for concrete would be higher, which may cause its price to rise, which will have a feedback effect on what mix of projects is efficient.

[71] See, for example, Arrow, et al, *Intertemporal Equity* at 130 (cited in note 2); Geoffrey M. Heal, *Intertemporal Welfare Economics and the Environment*, in Karl-Göran Mäler and Jeffrey R. Vincent, eds, 3 *Handbook of Environmental Economics* 1105 (Elsevier 2005) (which also includes a third component reflecting how the marginal utility of consumption is affected by the stock of environmental assets); Geoffrey Heal, *Discounting: A Review of the Basic Economics*, 74 U Chi L Rev 59 (2007).

[72] Consider also the caveats noted in Part III.C.

This approach to discounting dollars makes no reference to the market rate of interest, yet is generally regarded to be at least roughly consistent with using it. The reasoning makes reference to the optimization process. In perfect markets (including a lack of distortions such as from taxation), the socially optimal mix of investments and distribution of consumption across generations leads these two measures to be equal. As the efficiency analysis in Part II indicates, any investment mix that does not equate returns of all projects (private or public) at the margin is inefficient; moves in the direction of efficiency allow all generations to be made better off.

Moreover, at the optimal distribution, the market rate of return and social rate of tradeoff of dollars across time will be equated. To see this point, suppose that it were not so. For example, suppose as before that one dollar invested by the present generation will yield the market return of $R$, leaving $1 + R$ dollars for a future generation. And suppose further that one dollar in the present generation is socially valued as equivalent to $1 + S$ dollars in that future generation. If $R > S$, then an additional dollar of investment should be made, as it will yield $1 + R$ to the future generation, which by assumption exceeds $1 + S$ and thus is deemed more socially valuable than the one dollar given up by the present generation. Note that, as more is invested, two changes will occur. $R$ will fall because, as more investment is made, the marginal return to investment falls.[73] In addition, $S$ will rise because, as consumption in the future increases relative to that in the present, the marginal social value of consumption in the future falls and that in the present rises, so it will take more than the original $1 + S$ dollars in the future to be equivalent to one dollar in the present. As investment is increased, eventually $R$ and $S$ will be equated. (Similar analysis applies to the case in which $R < S$ initially.)

The foregoing discussion also implies, however, that the criterion of basing project (social investment) decisions on the market interest rate will differ from that of basing decisions on explicit social preferences when not at the distributive optimum. Which principle is correct in that case? The answer can best be understood by reference to the analysis in Part II.D, in which distribution was held constant, and that in Part II.E.3, in which non-distribution-neutral reforms were analyzed explicitly.

If in fact a project will be implemented in a distribution-neutral fashion (taking tax and other adjustments into account), then use of

---

[73] This effect recalls the point in Part IV.C that the market interest rate is endogenous because it depends on the intergenerational distribution. See also the qualification in note 68 on the effect of world capital markets on the interest rate.

the market interest rate is correct, as indicated previously. Actually, upon reflection, the explicit social preference approach would, if properly interpreted (adjusted), reach the same result. After all, if distribution is held constant, then the project does not involve relative changes in levels of dollars in the present and future, so the social weightings would be inapplicable. More precisely, if the reform is efficient and the benefits are shared across generations, there are only gains to discount, so the outcome will be favorable; likewise, an inefficient project under distribution neutrality would produce losses to all generations, so a socially weighted sum would be unambiguously negative.

A possible danger arises when a project on its face is not distribution neutral (many are not), leading it to be analyzed with the socially weighted approach, even though there will be other adjustments to the system, such as to taxes and debt in order to finance the project, that produce a distributive offset, the effects of which are not explicitly incorporated in the analysis.[74] If, as Part II.E.3 suggests, there is some tendency for whatever intergenerational distributive political equilibrium exists to prevail, then deviating from the market approach may well be incorrect. Furthermore, as previously discussed, if there is not distribution-neutral implementation, the second step of the proposed decomposition—isolating the redistributive feature—generally raises a distribution-distortion tradeoff, the primary distortion being of labor supply (as discussed in Part II.E.1). Placing social weights on dollar transfers while failing to analyze the tradeoff explicitly is tantamount to counting distribution while ignoring distortion, which also would produce misleading prescriptions.[75]

---

[74] The problem is analogous to the use of distributive weights in intragenerational cost-benefit analysis when the costs and benefits that are weighted are those of the project (viewed narrowly), ignoring how the project is to be financed.

[75] As Part II.E.1 notes, over the past few decades a substantial "second-best" literature in economics focusing on distortionary effects has emerged. There appears to be no analogue to this development regarding intergenerational applications of cost-benefit analysis. Fortunately, as Parts II.D and II.E indicate, a single unifying framework applies readily to both contexts; specifically, the approach developed in more recent work applicable (implicitly) to the problems of intragenerational public goods and the regulation of externalities, regarding both distribution and distortion, extends readily to the intergenerational version of these problems.

A further, related point is that much prior economic literature, for example, most of that discussed in Arrow, et al, *Intertemporal Equity* (cited in note 2), builds on variations of what is referred to as the Ramsey model of taxation in which there is a representative individual (in the present setting, for each generation), which rules out (intragenerational) income differences, and a standard labor income tax is disallowed (despite its prevalence in developed economies). As discussed by Atkinson and Stiglitz, 6 J Pub Econ 55 (cited in note 12), and elaborated in Kaplow, *The Theory of Taxation and Public Economics* (cited in note 12), and Kaplow, 90 J Pub Econ 1235 (cited in note 12), the principal results associated with such models are replaced by qualitatively different ones along the lines presented in Parts II.D and II.E once one shifts to a framework with income differences and redistributive income taxation.

The greatest potential shortcoming of using explicit social weightings is precisely the core idea identified throughout Part II: the method results in the inefficient selection of projects, which may come at the expense of every generation. Specifically, projects may be deemed socially desirable under this approach even when they are inefficient, as long as they implicitly produce a sufficient intergenerational distributive gain. Whenever that is the case, however, it would be better to shelve such projects and instead implement alternatives (including, notably, direct forms of intergenerational redistribution, such as through debt policy) that produce the same distributive gain without the direct efficiency cost from the project per se. Likewise, some projects that fail under the explicit weightings approach will be efficient and thus could be implemented with appropriate distributive offsets so as to benefit all generations. Both of these cases are obscured when different types of effects, some distinctive to the project and some not, are commingled under a unitary, all-things-considered assessment.

Use of explicit social weightings fails to separate questions of intergenerational efficiency and distribution. As just suggested, this failure can generate mistaken conclusions. Also of great importance, analysis becomes more complicated and is made more obscure. To properly use social weightings assumes that, for any project, one can identify the long-run distributive effect (that is, how the intergenerational distributive political equilibrium is likely to be altered, if at all), decide on the proper social weightings, and analyze the tradeoff with distortion—all in addition to the often formidable task of determining the direct costs and benefits of the project itself. As suggested previously, there are considerable advantages from specialization if these two sets of inquiries are separated. Furthermore, it seems plausible that greater clarity can be achieved, both among academic and other professional analysts and when the analysts attempt to communicate the results of their work to policymakers.

## VI. CONCLUSION

The core of this Article elaborates the thesis that intergenerational distributive justice and efficiency are largely distinct. This argument was developed from a number of complementary perspectives, many of which employed constructs that held the intergenerational distribution of resources constant. For the most part, it does not matter whether that given distribution constitutes a normative ideal in one or another sense, an equilibrium that results from extant political forces, or any other, even arbitrary, allocation.

Also examined were a variety of additional factors, including distortionary effects of redistribution, heterogeneity of preferences, feedbacks from distribution to efficiency, and complications relating to

the market interest rate at which dollars are discounted. These considerations do not materially alter the central conclusion. Nor do most of them raise questions that are distinctive to the intergenerational setting. Taken together, the analysis suggests that accepted principles of policy analysis developed in the intragenerational context carry over. Indeed, some complexities are purely intragenerational; that is, they influence analysis within each of the generations affected by policies but do not create any distinctive challenges along the intergenerational dimension.

Although no particular normative view was adopted concerning intergenerational distributive justice, much of the analysis employed a welfarist framework that, as the introduction notes, assesses policies exclusively by reference to their effects on individuals' well-being. A few words are in order about other normative perspectives.

First, because much of the analysis held distribution constant, one might suppose that a variety of views on distributive justice would be consistent with the present argument. This suggestion is not necessarily correct, however, since the very meaning of constant distribution was taken here to be denominated in terms of individuals' well-being.[76] If instead individuals' positions were assessed by some other metric, such as Rawls's primary goods or Sen's capabilities and functionings, the analysis could differ.[77]

Second, the analysis demonstrated that efficient policies with intergenerational consequences could be implemented in a fashion that makes everyone better off, that is, constitutes a Pareto improvement.[78] Hence, the adoption of other distributive notions, notably, nonwelfarist ones, entails acceptance of the social desirability of making everyone worse off if indeed the notions' prescriptions materially differ.[79] This stance seems difficult to defend and, moreover, inconsistent

---

[76]   See especially the more precise formulation of distribution neutrality in Part II.D.

[77]   See Rawls, *A Theory of Justice* 90–95 (cited in note 2); John Rawls, *Social Unity and Primary Goods*, in Sen and Williams, eds, *Utilitarianism and Beyond* 159 (cited in note 3); Amartya Sen, *Commodities and Capabilities* (North-Holland 1985); Amartya Sen, *Well-Being, Agency and Freedom: The Dewey Lectures 1984*, 82 J Phil 169 (1985). For a critique, including in particular the point that these approaches conflict with the Pareto principle and plausible notions of freedom and autonomy, see Louis Kaplow, *Primary Goods, Capabilities . . . or Well-Being?* (unpublished manuscript 2006).

[78]   The primary caveat, concerning heterogeneity within income groups within the same generation (see Part II.E.2), seems tangential to most imaginable notions of intergenerational distributive justice since the differences are purely intragenerational.

[79]   This is an application of the more general point that nonwelfarist principles violate the Pareto principle. See Kaplow and Shavell, *Fairness Versus Welfare* (cited in note 3); Louis Kaplow and Steven Shavell, *Any Non-Welfarist Method of Policy Assessment Violates the Pareto Principle*, 109 J Polit Econ 281 (2001); Louis Kaplow and Steven Shavell, *The Conflict Between Notions of Fairness and the Pareto Principle*, 1 Am Law & Econ Rev 63 (1999).

with typical appeals concerning intergenerational distributive justice. Regarding the latter, appeals of just savings, sustainability, leaving as good for those who come later, and so forth explicitly contemplate that the requisite sacrifices by the present generation are to benefit some or all future generations.[80] If, however, those generations to whom obligations are owed are hurt rather than helped, the positions become difficult to comprehend. An important possibility is that seemingly nonwelfarist moral intuitions about what is due to future generations are implicitly surrogates for directing sufficient concern toward those generations' welfare.

The most obvious exceptions to this point are those views that, for example, endow the natural environment with independent normative significance. Such objectives require a different analysis, although one that should possess many similarities with that presented here. For example, however nature and well-being are to be traded off, one could better advance both objectives by adopting more efficient policies. That is, both preservation and well-being should be advanced in the most efficient fashion possible, which requires employing a constant rate of tradeoff between the two across policies. Where the present analysis took the intergenerational distribution of well-being as given, one could also hold constant natural protection, asking what mix of policies was most efficient in the manner described herein.[81]

As a final, related note, observe that if a welfarist notion of intergenerational distributive justice is adopted, there is nothing a priori special about environmental protection, conservation of resources, or any other particular manner by which a present generation may affect future generations. All that matters is how one or another regime will affect subsequent individuals' well-being. It may be that due to the increasing marginal harm caused by certain forms of environmental degradation, the existence of irreversibilities, and certain forms of uncertainty, efficiency requires that a great deal of investment be made to preserve the environment. The extent to which this is true and, just as important, which investments are required should be determined by considerations of efficiency, once a principle governing intergenerational distributive justice is adopted. Such a distributive principle, it should be emphasized, would (at least in a welfarist framework) be

---

[80]  This point is equally applicable to principles designed to limit the sacrifices required by the present generation.

[81]  To even define such a view, it would be necessary to reduce relevant features of the natural environment to a common denominator—equivalently, to state how different traits should be traded off—unless one literally required that every aspect of nature remain unchanged (which seems infeasible and is difficult to define due to the need to specify a benchmark and the fact that evolution is part of nature).

the same even if it so happened that the course of the natural environment were entirely unaffected by human activity—for example, if the only decision to be made was how to allocate a single resource between consumption by the present generation and investment for the benefit of future generations.