

WELL-BEING, INEQUALITY AND TIME: THE TIME-SLICE PROBLEM AND ITS POLICY IMPLICATIONS

Matthew D. Adler¹
July 2007

INTRODUCTION

- I. EQUITY-REGARDING WELFARISM
 - A. *Social Welfare Functions (SWFs)*
 - B. *Lifetime versus Sublifetime Views*
 - C. *Why Welfarism?*

- II. LIFETIME VERSUS SUBLIFETIME VIEWS: PHILOSOPHICAL ARGUMENTS
 - A. *The Basic Case for the Lifetime View*
 - B. *McKerlie's Response*
 - C. *Parfit's Response*

- III. EX ANTE VERSUS EX POST APPROACHES TO SOCIAL CHOICE UNDER UNCERTAINTY

- IV. IMPLICATIONS OF THE LIFETIME VIEW: MEASURING EQUALITY AND POVERTY
 - A. *Equality Measurement*
 - B. *Poverty Measurement*

- V. IMPLICATIONS OF THE LIFETIME VIEW: CHOOSING POLICIES
 - A. *Redistributive Taxes*
 - B. *Should Non-Tax Instruments be Used for Redistribution?*
 - C. *Distributively Sensitive Policy Analysis*

CONCLUSION

INTRODUCTION

Probably no topic has been more central to legal scholarship than inequality. But one aspect of the problem has been insufficiently discussed. This is the question of the *time-slice*. Should egalitarians be concerned with the inequality of individuals' well-being or resources over their lifetimes, or with inequality during what I shall "sublifetimes," such as annual or even momentary inequality?

This Article engages the time-slice problem, and does so from the perspective of welfarism. In other words, I assume that the "currency" for equality is human well-

¹ Leon Meltzer Professor, University of Pennsylvania Law School. Many thanks to workshop participants at the law schools of Arizona State, Boalt Hall, Chicago-Kent, Georgetown, Harvard, Illinois, and the University of Pennsylvania, for their very helpful comments.

being,² and then ask whether the legal system should aim to secure the equality of individual lifetime or sublifetime well-being. My conclusion is that the lifetime view is more attractive.

Part I of the Article summarizes the welfarist framework for thinking about problems of policy choice. Among other things, it explains my reasons for addressing the time-slice problem from a welfarist rather than extra-welfarist perspective. In short: the concept of “well-being” is an inclusive one, potentially encompassing a range of objective goods or aspects of human flourishing, and not merely preference satisfaction. Further, welfarists need not be utilitarians. The formalism of an “equity-regarding social welfare function” provides a tool for thinking about law and policy which is analytically rigorous *and* sensitive to distributional considerations.

Part II defends the lifetime view. The time-slice problem, while generally overlooked by legal scholars, has generated a small body of philosophical scholarship. Although the dominant view, here, *is* a lifetime view, a number of dissenters have advanced important challenges to this position, which I grapple with in Part II. Specifically, Derek Parfit has presented a deflationary picture of personal identity over time which, he suggests, supports a sublifetime view. And Dennis McKerlie has argued that -- quite apart from considerations of personal identity -- there are powerful intuitions supporting the sublifetime view. Part II responds to McKerlie’s and Parfit’s challenges to the lifetime view.

The thrust of Part II is that the optimal distributively-sensitive legal rules and institutions -- within the welfarist framework -- are those identified by the application of an equity-regarding social welfare function to *lifetime* rather than *sublifetime* utilities. But what, exactly, does that mean given conditions of uncertainty? If the policymaker is uncertain about the distribution of lifetime utilities, should she aim to equalize the distribution of *expected lifetime well-being*, or should she instead focus on the expected degree of inequality of *actual lifetime well-being*? Should she care about equality of expectations (the “ex ante” approach), or about expected inequality (the “ex post” approach). Drawing on my work with Chris Sanchirico, Part III clarifies the distinction between the two approaches and summarizes the case for the “ex post” approach.

The upshot of Parts II and III is a framework for thinking about legal rules and institutions which cares about equality, focuses on lifetime rather than sublifetime well-being, and is outfitted for choice under uncertainty. Parts IV and V survey the multiple, concrete implications of this account. In particular, I discuss what a lifetime view implies for the measurement of inequality; for the measurement of poverty; for the design of the tax-and-transfer system; for the contested question whether non-tax instruments, such as environmental regulations or expenditures on public goods, should ever be used for redistributive purposes; and (because the answer to this question is, plausibly, yes) for how policy-analytic tools, such as cost-benefit analysis, should be rendered sensitive to distributive considerations.

² See generally Richard Arneson, *Welfare Should be the Currency of Justice*, 30 CAN. J. PHIL. 497 (2000).

Whether or not the reader is ultimately moved by my arguments in favor of a lifetime view, this Article should at least persuade her that the choice between lifetime and sublifetime perspectives --in short, the time-slice problem -- is a vital one for law and policy. Consider the following examples, discussed (among others) in Parts IV and V.

-- A very large body of work within applied economics applies quantitative metrics of inequality, such as the so-called Gini coefficient, to the distribution of income. The bulk of this scholarship attends to the distribution of annual income, but there are scholars who have tried to measure the inequality of lifetime income, using longitudinal data, proxies for lifetime income (such as annual consumption), or simulation models. The upshot of this research is that the distribution of lifetime income is more equal, both in the U.S. and abroad, than the distribution of annual income. Further, time trends may look different. For example, the inequality of annual income has been increasing in the U.S., but the inequality of lifetime income has not.³

-- Health economists have followed the lead of the income-inequality scholars, and have started to measure the population distribution of health or longevity. Here, too, the lifetime and sublifetime perspectives about the extent or time trend of inequality can differ.⁴ For a simple example, imagine an outbreak of a fatal disease that tends to affect those who are already in chronically bad health. From a lifetime perspective, the outbreak increases inequality. Those affected by the disease are, generally, already badly off in terms of lifetime well-being. The disease lowers the lifetime well-being of this badly off group, by shortening their lifespans. But, from a sublifetime perspective, the effect of the disease on inequality is more ambiguous. The outbreak may actually reduce annual or momentary health inequality, going forward, because those who remain in the population tend to be in better health.⁵

-- Tax scholars have long sought to measure the “incidence” of taxes. A tax is “regressive” if its burden, as a fraction of the taxpayer’s income, tends to decrease as the taxpayer’s income increases. A tax is “progressive” if its burden tends to increase as the taxpayer’s income increases. Recent incidence scholarship shows that the time-slice problem matters critically for evaluating the degree of regressivity/progressivity of various taxes. Are we asking about the tax burden as a fraction of *annual* income or of *lifetime* income? Sales and excise taxes are, famously, regressive from a sublifetime perspective, but are much less so from a lifetime perspective. The personal income tax is progressive from a sublifetime perspective, but is less progressive from a lifetime perspective.

-- Poverty researchers have begun to examine “chronic” poverty. Traditional poverty metrics look to whether the individual falls below some periodic or momentary threshold with respect to some attainment critical to well-being - for

³ See *infra* Part IV.A.1.

⁴ See *infra* Part IV.A.3.

⁵ See *infra* Part IV.A.2.

example, whether the individual's annual income is below the poverty line, or whether he is badly nourished or lacks shelter at present. Chronic poverty metrics look to whether the individual repeatedly falls below the periodic threshold, or whether his long-term attainments fall below some long-term threshold. It turns out that the determinants of transient and chronic poverty, and the appropriate policy responses, are quite different. Chronic poverty results from low human capital and geographic or social isolation. Individuals who experience transient rather than chronic poverty generally do so because of transitory "shocks" from which they rebound, such as the loss of a wage-earner in the household, an economic downturn, or a natural disaster. Policy tools such as unemployment insurance or disaster relief, which mitigate transient but not chronic poverty, will tend to look quite attractive from the sublifetime perspective but less so if we care about lifetime well-being.⁶

- Louis Kaplow and Steven Shavell, among others, have argued against using non-tax instruments for redistributive purposes. The core of the analysis is this: Whenever the policymaker rejects a policy which has certain "efficiency" characteristics, on distributive grounds, it will be possible to couple the efficient policy with changes to the income tax system so that everyone is better off. But what, exactly, does "efficiency" mean? And how does uncertainty figure into the argument? If one adopts a lifetime perspective together with an "ex post" account of choice under uncertainty, one policy is Pareto-superior to another -- the second would be rejected by any social welfare function -- only if the first policy increases everyone's *lifetime well-being in every state of the world*. Call this EP-L Pareto superiority. An "efficient" policy, in turn, is one that can be transformed into an EP-L Pareto superior policy via the tax system. This is a demanding criterion which, it seems, few real-world policies will satisfy. Thus, the Kaplow and Shavell argument turns out to have limited scope.

-- The "suspect class" framework is a useful heuristic for evaluating distributional impacts. A policy is distributionally problematic, and merits closer scrutiny, if it discriminates against, or has a disparate impact on, racial minorities, women, and perhaps other groups. What groups precisely? The specification of suspect classes implicates the time-slice problem -- because a suspect class is one whose members tend to be badly off, and one then needs to ask whether that means "badly off" in lifetime or sublifetime terms. This implication is clearest in the case of *age*. From the sublifetime perspective, there is a strong case for thinking of the aged as a suspect class. Older individuals tend to have poorer health, lower incomes, fewer friends and pursuits, and so forth. From the lifetime perspective, that case evaporates or at least becomes much weaker. Absent secular changes in individuals' income, health, etc., the distribution of lifetime well-being among individuals who have reached a point of relatively advanced age (e.g., 85) should be no different from the distribution of lifetime well-being among individuals who are now younger but will live to that point. And older individuals will tend

⁶ See *infra* Part IV.B.

to have a *higher* level of lifetime well-being than younger individuals who will die prematurely, before ever reaching that age.

In short, a wide range of real-world questions that egalitarians should care about -- questions about how to measure inequality and how to design policies to redress it -- are enmeshed with the time-slice problem. It is surprising -- deeply puzzling, really -- that legal scholars have paid the problem so little attention.

I. EQUITY-REGARDING WELFARISM

This Part sets forth the welfarist account of equality that undergirds the remainder of the Article. Section A explains what a “social welfare” function (SWF) is, and distinguishes between three different kinds of equity-regarding SWFs: “prioritarian,” “comparativist,” and “sufficientist.” Section B explicates the distinction between *lifetime* and *sublifetime* views, within the SWF framework. Section C addresses the worry that welfarism and the SWF framework is an inadequate basis for thinking about problems of equality.

A. *Social Welfare Functions (SWFs)*

The SWF framework, which is little used by philosophers but is standard within theoretical welfare economics and certain areas of applied economics and legal scholarship, works as follows.⁷ Assume a fixed population of N individuals, each of whom lives for K periods.⁸ Each possible outcome is represented by a vector or list of utilities, representing the lifetime or sublifetime well-being of the individuals in that outcome. In the case of lifetime welfarism, each possible outcome will be represented as a vector of N utilities (1 for each lifetime well-being of each of the N individuals). In the case of sublifetime welfarism, each possible outcome will be represented as a vector of $N \times K$ utilities (1 for the well-being of each of the N individuals during each of the K stages in her life.) However, let us ignore the distinction between lifetime and sublifetime utilities for the moment. On both sorts of views, each outcome is representable as a vector, with an appropriate number of entries, representing individuals’ well-being. To put this formally, there is a utility function u that maps a given outcome O onto a vector $u(O)$, with N or $N \times K$ entries.

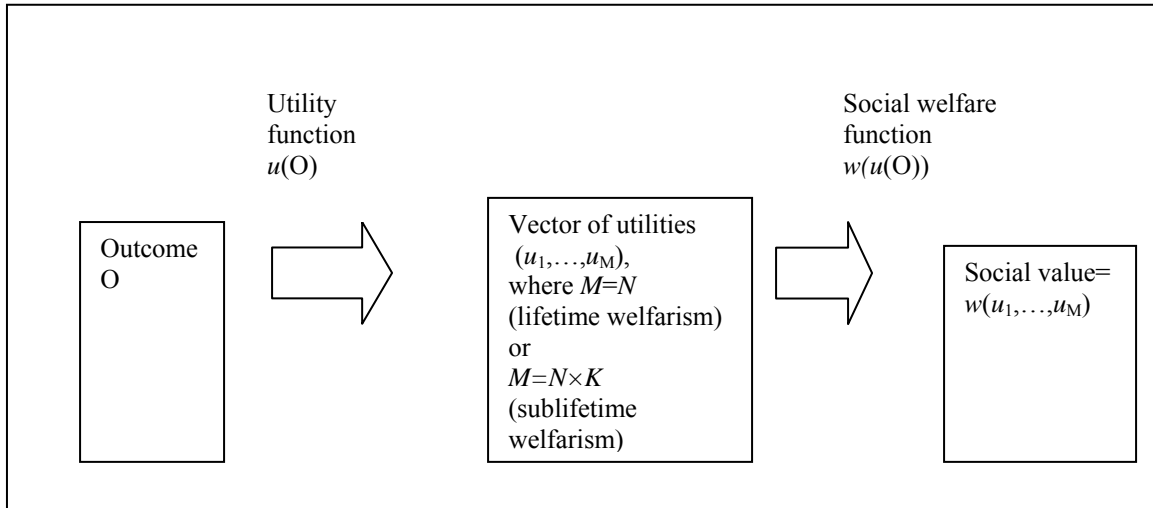
An SWF is a mathematical function w which takes the utility vector for each outcome, and maps it onto a single (scalar) number, representing the social value of that

⁷ See generally Matthew D. Adler & Chris William Sanchirico, *Inequality and Uncertainty: Theory and Legal Applications*, 155 U. PA. L. REV. 279, 291-304 (2006).

⁸The “fixed population” assumption means that the same individuals exist in all possible outcomes of the policy choices on hand. A full theory of social choice will indeed need to deal with variable-population cases, where the composition of the population may vary across outcomes. On the specific issue of the time slice, however, it is very hard to see why a shift from fixed- to variable population cases should shift the relevant time slice from sublifetime to lifetime utilities. In any event, and in order to simply the analysis, I will not discuss variable-population issues here. On such issues, see generally CHARLES BLACKORBY ET AL., *POPULATION ISSUES IN SOCIAL CHOICE THEORY, WELFARE ECONOMICS, AND ETHICS* (2005).

outcome. $w(u(O)) > w(u(O^*))$ if and only if O is a better outcome than O^* , where w is the SWF and $u(O)$ and $u(O^*)$ are the utility vectors representing outcomes O and O^* , respectively.

THE SWF FRAMEWORK



An SWF can be thought of as a mathematical representation of a welfarist moral view. A welfarist view offers moral verdicts of a certain sort: verdicts about how well off different individuals are in different outcomes; and verdicts about how outcomes compare with each other, which depend solely on the verdicts about individual well-being. (This is what makes the view welfarist: its ranking of outcomes and, then, choices, must be a product solely of well-being information.) The SWF framework formalizes all this with numbers.

A more precise statement of the correspondence between welfarist moral views and SWFs is this: A given welfarist moral view corresponds to a *set* of pairs of utility functions and social welfare function, $\{(u^1, w^1), (u^2, w^2), \dots (u^L, w^L)\}$, where each u^i maps a given outcome onto a vector of N or $N \times K$ utilities, and each w^i then maps utility vectors onto scalars representing the social value of outcomes. The reasons for thinking of a welfarist moral view as represented by a set of pairs of utility functions and SWFs, rather than a single such pair, have to do with the nonuniqueness of utility numbers and the possibility of incommensurability.⁹ These issues are orthogonal to the time-slice

⁹ Social choice theorists typically associate a single SWF with a family of utility functions, recognizing that the given SWF applied to the original utility function *or* admissible transformations thereof will produce the same ordering of outcomes. For example, the utilitarian SWF, applied to utility vectors produced from outcomes using a given u , will produce the same ordering of outcomes if u is multiplied by any positive constant and any constant is added. The sum of the square root of utilities, a kind of equity-regarding SWF, will produce the same ordering of outcomes if u is multiplied by a positive constant.

My suggestion that a given welfarist moral view might correspond to a set of SWFs (each paired with a corresponding utility function), rather than a single SWF, is a natural extension of the basic idea that utility functions and SWFs are mathematical *representations* of a moral view and that such representations

problem, however, and the arguments below in favor of a lifetime view will generally adopt the simplifying assumption that moral views correspond to a single utility function and social welfare function.

For a welfarist moral view to provide useful guidance to decisionmakers, it must do more than rank outcomes. It must also rank choices, in a given choice situation, as a function of the possible outcomes of each choice. This is the problem of outfitting an SWF for choice under uncertainty, which will be postponed until Part III.

What, exactly, are “outcomes”? In the case of unboundedly rational decisionmakers – decisionmakers who may not be omniscient but have unbounded reasoning abilities and mental storage capacities --- each outcome is what philosophers call a “possible world”: a fully specified possible history of the universe. No humans are unboundedly rational. SWFs with possible worlds as outcomes would be unusable by humans. Rather, SWFs must operate on *simplified* possible worlds: simplified pictures of possible realities. How to identify the appropriate simplification is a huge, unsolved problem for welfarism, and consequentialism more generally, which I will not attempt to address here. The problem is orthogonal to the time-slice problem and does not, I think, affect my argument for the lifetime view or assessment of its legal implications.

So the SWF framework maps outcomes – simplified possible worlds – onto vectors of utilities, and then uses an SWF to map these vectors onto numbers representing the ranking of the outcomes. Economists virtually always make the further assumption that SWFs satisfy the Pareto principle for outcomes, and most philosophers (even those who care about equality) find that principle congenial. The Pareto principle for outcomes says that, if at least one individual is better off in one outcome as compared to a second, and no individual is worse off, the first outcome must be ranked higher. For the remainder of the Article, the reader should understand (unless otherwise noted) that all SWFs being discussed satisfy the Pareto principle for outcomes.¹⁰

What, then, is the distinction (within the family of SWFs satisfying the Pareto principle for outcomes) between SWFs which are sensitive to distributional considerations -- for short, “equity-regarding” SWFs – and SWFs which are not?

Economists have a simple answer. An SWF is equity-regarding if it satisfies the so-called Pigou-Dalton axiom.¹¹ The Pigou-Dalton axiom says that a transfer of a unit of utility from a higher utility individual to a lower utility individual, holding constant total utility, must increase social value. Note that the utilitarian SWF, which simply sums


may not be unique. Further, that extension may be helpful in allowing the SWF framework to cope with incommensurability -- where a moral view provides a partial rather than complete ordering of outcomes. The view corresponds to $\{(u^1, w^1), (u^2, w^2), \dots (u^l, w^l)\}$; each pair (u^i, w^i) provides a complete ordering of outcomes; the partial ordering is the “intersection” of these complete orderings, i.e., one outcome is better than another pursuant to the view iff it is better according to every (u^i, w^i) pair.

¹⁰ See Adler & Sanchirico, *supra* note 7, at 293-94. All SWFs also satisfy an anonymity constraint, *id.*, which is even less controversial. Two vectors with the same utilities, but in a different order, must be ranked the same.

¹¹ See Adler & Sanchirico, *supra* note 7, at 298-300.

utilities, and is the hallmark of an SWF that is *not* equity-regarding, does not satisfy the Pigou-Dalton axiom.

THE PIGOU-DALTON PRINCIPLE



How does this simple answer relate to the contemporary philosophical literature, which has coalesced around three different conceptions of equality: a view commonly known as “prioritarian,” a view which I shall term “comparativist,” and a view commonly known as “sufficientist”?

“Prioritarians” do not essentially care about the pattern of well-being.¹² Rather, they give greater social value to increments to well-being that are realized by individuals at lower levels of well-being. They believe that individuals who are worse off in absolute terms have a stronger moral claim. There *is* a standard and intuitively compelling way to represent “prioritarianism” in the SWF framework: namely, a prioritarian SWF is one that sums an *increasing, strictly concave function* g of individual utilities. Assume we have a vector $u(O)$ with M entries (where $M=N$ in the case of lifetime welfarism, and $M=N \times K$ in the case of sublifetime welfarism). In other words, $u(O) = (u_1, u_2, \dots, u_M)$, where u_1 is the first utility number in the vector, u_2 the second, up to u_M . Then a prioritarian SWF w has the following form: $w(u(O)) = \sum_{j=1}^M g(u_j)$, where g is an increasing, strictly concave function.

An increasing, strictly concave g -function, such as the square root or the logarithm, has a positive first derivative and a negative second derivative.¹³ This means

¹² See *id.* at 300-01.

¹³ This is imprecise, because strictly concave functions need not be fully differentiable. A precise definition is that a strictly concave function g is such that the weighted average of the g -value of any two

(1) that an increase in some individual's utility, *ceteris paribus*, increases the value of the SWF, that is, $\sum_{j=1}^M g(u_j)$, which is just what the Pareto principle for outcomes says; and (2) that the increase in the value of the SWF produced by a small increase in an individual's utility (the first derivative of g) *decreases* as the level of that utility increases, which is just a mathematical restatement of the prioritarian idea that changes in an individual's well-being have greater moral importance the lower the level of the individual's well-being.

Any SWF which takes the form of summing an increasing, strictly concave function of individual utilities will satisfy the Pigou-Dalton axiom. Note, further, that any SWF which takes the form of summing an increasing, strictly concave function of individual utilities will satisfy an "indifference" axiom. This says that the ranking of two utility vectors should be unaffected by "indifferent" utilities: those that are the same in the two vectors. "Indifference" is a precise and, once more, intuitively attractive way to express the notion that prioritarians do not essentially care about the pattern of welfare levels.

THE INDIFFERENCE AXIOM

Consider the initial pair of utility vectors (a,b,x,y,c,d) , (a^*,b^*,x,y,c^*,d^*) , with the same utilities in the 3rd and 4th positions. Change those "indifferent" utilities to produce a new pair of utility vectors (a,b,x^+,y^+,c,d) , $(a^*,b^*,x^+,y^+,c^*,d^*)$. The "indifference" axiom says that $w(a,b,x,y,c,d) > w(a^*,b^*,x,y,c^*,d^*)$ if and only if $w(a,b,x^+,y^+,c,d) > w(a^*,b^*,x^+,y^+,c^*,d^*)$

For example, if the social value of $(2,10,100,100,3,5)$ is greater than that of $(4,4,100,100,4,4)$, then the social value of $(2,10,4,4,3,5)$ must be greater than that of $(4,4,4,4,4,4)$.

In the case of prioritarianism, then, the translation into SWF language is pretty straightforward. Things get a bit trickier when we turn to "comparativism." "Comparativists" do care essentially about patterns. To put this formally, a "comparativist" SWF does *not* satisfy the indifference axiom. But we need to say more,

numbers is less than the g value of their weighted average, for all weights. *See, e.g.,* RANGARAJAN K. SUNDARAN, A FIRST COURSE IN OPTIMIZATION THEORY 173-85 (1996).

because there are lots of SWFs which fail indifference but we would surely not want to describe as equity-regarding. The simple answer is that a comparativist SWF is one that satisfies the Pigou-Dalton axiom (and thus is equity-regarding), but fails indifference (and in that sense essentially cares about patterns).¹⁴ This is rough, because there are comparativists who think that the Pigou-Dalton axiom is too demanding.¹⁵ Once we depart from prioritarianism and start to care about patterns, why insist that every Pigou-Dalton transfer must be an improvement in the pattern? A somewhat better definition is that a comparativist SWF fails indifference and counts a total-utility-preserving move to perfect equality as an improvement. This may be too weak, but seems closer to the truth,¹⁶ and shall be the definition I use here.¹⁷

COMPARATIVIST SWFs

An SWF that fails Indifference but satisfies the Pigou-Dalton principle: the sum of rank-weighted utilities. If (a,b,c,d) are utilities in ascending order, then $w=4a+3b+2c+d$. The Pigou-Dalton principle is satisfied: $w(1,3,7,12)=39$, and if we transfer 2 units from the fourth to the third entry, producing $(1,3,9,10)$, $w(1,3,9,10)=41$. However, the Indifference Axiom is not satisfied. While $w(5,10,12,16)<w(6,10,12,13)$, $w(5,1,2,16)>w(6,1,2,13)$.

An SWF that fails Indifference and Pigou-Dalton but prefers perfect equalization: An SWF that sums utilities, but adds some premium k to the sum if the utilities are perfectly equal. Note how this fails Pigou-Dalton. $w(1,3,7,12)=23$, and $w(1,3,9,10)=23$. Note also that it fails Indifference. For example, if k is 6, then while $w(2,10,100,100)>w(4,4,100,100)$, $w(2,10,4,4)<w(4,4,4,4)$.

“Sufficientists,” like prioritarrians, deny that the pattern of well-being is itself of moral significance.¹⁸ Rather, a concern for “equity” arises in virtue of the fact that individuals who are worse off in absolute terms have a stronger moral claim. But, unlike prioritarrians, they believe that this claim evaporates once the individuals reach a level of “sufficient” or “adequate” well-being. In other words, “sufficientists” incorporate a

¹⁴ See *id.* at 300-02.

¹⁵ See *id.* at 303; LARRY S. TEMKIN, *INEQUALITY* 83-84 (1993).

¹⁶ It is plausible that an SWF which generally ranks outcomes in light of total utility, but gives a bonus to outcomes where all the utilities are exactly equal, should be counted as equity-regarding.

¹⁷ Note that, if an SWF counts *every* Pigou-Dalton transfer as an improvement (the stronger definition), then it counts a move to total equality as an improvement (the weaker definition), because a move to equality can be achieved by a series of Pigou-Dalton transfers. See Adler & Sanchirico, *supra* note 7, at 299 n.40.

¹⁸ See Roger Crisp, *Equality, Priority, and Compassion*, 113 *ETHICS* 745 (2003); Harry Frankfurt, *Equality as a Moral Ideal*, 98 *ETHICS* 21 (1987).

morally critical threshold level of well-being into their theory. Those below the threshold have stronger moral claims to welfare than those above. Those above the threshold do not have stronger claims than those yet further above. To put this in SWF terms, “sufficientists” (like prioritarrians) accept the indifference axiom but accept the Pigou-Dalton axiom only with respect to transfers across the threshold, and not with respect to transfers between individuals above it. For example, imagine that the utility threshold is 20. Then a sufficientist SWF w will say that $w(1,5,25,35)=w(1,5,27,33)$ -- despite the transfer of 2 units from the fourth to the third entry -- but that $w(1,5,25,35)<w(1,7,23,35)$.

The following table summarizes the features of the three families of equity-regarding SWFs, as just explicated.

SWFs: A SUMMARY

-- *The Utilitarian SWF* (simple sum of utilities)

-- *Equity Regarding SWFs*

-- Prioritarian SWFs: $w(u_1, u_2, \dots, u_M) = \sum_{j=1}^M g(u_j)$, where g is increasing and strictly concave

Prioritarian SWFs satisfy the Pigou-Dalton principle and the Indifference axiom.

-- Comparativist SWFs: These do *not* satisfy the Indifference axiom. However, they *do* either (1) satisfy the Pigou-Dalton principle, or if not (2) at least count perfect equalization of utilities, holding total utility constant, as an improvement.

-- Sufficientist SWFs: These incorporate a well-being threshold, and satisfy the Pigou-Dalton principle with respect to transfers across the threshold, but not with respect to transfers above it. They do satisfy the Indifference axiom.

A final point of clarification. Prioritarian, comparativist, and sufficientist SWFs are equity-regarding, in the sense that they care about the distribution of well-being (more precisely, in satisfying the Pigou-Dalton principle, generally or across a threshold, or at least preferring perfect equalization). But they are also sensitive to considerations other than equity. If we add an increment to the largest utility in the vector, then that cannot be an improvement with respect to equality, but still (given the Pareto principle for outcomes) must be counted as a social improvement. To put the point another way, prioritarian, comparativist, and sufficientist SWFs represent all-things-considered moral judgments, integrating *both* considerations of distribution *and* considerations of Pareto efficiency or overall well-being. No one who accepts the Pareto principle for outcomes believes that our all-things-considered moral judgments are solely a matter of equity.

B. *Lifetime versus Sublifetime Views*

The difference between lifetime and sublifetime equity-regarding views can now be stated. I will assume that each outcome can be described as a series of temporally locatable facts – facts that occur in arbitrarily small periods. (Because the facts of relevance to welfare either consist of, or supervene upon, physical or mental facts

somewhere in the world, this assumption seems plausible.) There are N individuals in the population. The lifetime equity-regarding view represents a given outcome as a vector of N utilities, representing the lifetime well-being of each individual in that outcome, where each utility is a function of *all the facts about that outcome*; and applies some equity-regarding SWF (prioritarian, comparativist, or sufficientist) to that N -entry vector. The sublifetime equity-regarding view cares about the distribution of individual well-being during some temporal fraction of a lifetime. It divides the outcome into K stages (these could be decades, years, months, or moments); represents the well-being of each individual in a given stage with a stage utility number *which is a function of all the facts in that stage*; and applies some equity-regarding SWF (prioritarian, comparativist, sufficientist) to the $N \times K$ entry vector that results.

To put this a bit more formally, if F^O are all the facts about outcome O , then $F^O = (F_1^O, F_2^O, \dots, F_K^O)$, where F_1^O are all the facts about outcome O in period 1, F_2^O are the facts about outcome O in period 2, up to F_K^O , the facts about outcome O in period K . In the case of lifetime welfarism, $u(O)$, the utility vector for outcome O , has N entries (u_1, u_2, \dots, u_N) , where u_i is the lifetime utility of individual i . Each lifetime utility number u_i is a function of -- potentially influenced by -- the facts about the outcome in *all* periods, that is by the whole array $(F_1^O, F_2^O, \dots, F_K^O)$. In the case of sublifetime welfarism, $u(O)$, the utility vector for outcome O has $N \times K$ entries $(u_{1,1}, u_{2,1}, \dots, u_{N,1}; u_{1,2}, u_{2,2}, \dots, u_{N,2}; \dots; u_{1,K}, u_{2,K}, \dots, u_{N,K})$, where $u_{1,1}$ is the utility of individual 1 during period 1, $u_{1,2}$ the utility of individual 1 during period 2, and so forth, up to $u_{N,K}$, the utility of the last individual N during the last period K . Each $u_{i,k}$ is a function of the contemporaneous facts, the facts in period k , that is F_k^O . So sublifetime utilities $u_{1,1}, u_{2,1}$, up to $u_{N,1}$ are a function of F_1^O ; sublifetime utilities $u_{1,2}, u_{2,2}$, up to $u_{N,2}$ are a function of F_2^O ; and so forth.

Note some critical points about this conceptualization of the lifetime/sublifetime distinction. First, I assume that a given individual's sublifetime well-being during some period is a function solely of the facts during that period. The proponents of sublifetime accounts seem to adopt this premise implicitly,¹⁹ and it is intuitively plausible. How well Sam's life is going during a year depends on what happens during that year, and is not dependent on what happens to him later on, or what happened to him earlier.

Second, I do *not* assume that a given individual's lifetime or sublifetime well-being depends just on facts about his own mental states or physical body. As discussed in a moment, the SWF framework is meant to mesh with a wide variety of conceptions of

¹⁹ As further discussed below, in Part II, McKerlie presents cases in which we (allegedly) prefer to equalize sublifetime well-being, such as the noble-peasant case. These are cases in which the difference in sublifetime well-being during some period supervenes on facts in that period -- for example, facts about who is a noble and who a peasant at a particular point in time. He and Parfit also discuss our preference to relieve temporary suffering -- the idea, here, is that a person is particularly badly off at a given time because she is suffering at that time.

well-being. An individual's well-being during a period, or over his lifetime, may depend on what other individuals are doing or thinking (for example, on how well his children are doing, or what his social status is), or other facts about the world outside his mind and body.

Finally, I do *not* assume that an individual's lifetime utility is *additively separable* into the sum of some function of sublifetime utilities (either the straight sum of sublifetime utilities, or the sum of some transformation thereof).²⁰ Although economists sometimes assume additive separability, it is an oversimplification – and, for our purposes, a problematic one. Additive separability rules out the possibility of cross-temporal interaction effects, between welfare-relevant facts in different periods, in producing lifetime well-being. For example, it rules out the possibility of an improvement effect -- namely that a life where facts with respect to some aspect of well-being get better over time is better just by virtue of this improvement. There is strong experimental support for an improvement effect, and for other failures of additive separability, and there is no reason to think that the best theory of lifetime well-being satisfies additive separability.²¹

²⁰ Formally, additive separability says that u_i , the lifetime utility of individual i , can be decomposed as follows: $u_i = v(u_{i,1}) + v(u_{i,2}) + \dots + v(u_{i,K})$, where v is some function and $u_{i,k}$ is the sublifetime utility of individual i during period k .

²¹The issues, here, involve what might be called an “independence” condition with respect to the relationship between temporal facts and lifetime well-being. (This is structurally very similar to the “indifference” principle discussed earlier, although that principle involved the connection between utilities and social welfare, while this one involves the connection between sublifetime facts in one individual's life and his lifetime well-being.) The independence condition says the following. Consider two arrays of facts in periods 1 through K , with identical facts in one or more periods. Produce two new arrays of facts, by taking the periodic facts that are the same in both arrays and substituting any other pair of identical periodic facts. Then the ordering of the arrays must be preserved. If the original two arrays produced equal lifetime utility for individual i , then the two new arrays must do so. Similarly, if one of the two arrays produced greater lifetime utility for individual i , then the new array created from this one must produce greater lifetime utility than the new array created from the other array. To illustrate, if (A,B,W,X,C) and (A*,B*,W,X,C*) are two arrays of facts (in a 5-period case), then the temporal facts are identical in periods 3 and 4 across the arrays. We produce two new arrays by substituting new facts: Y for W and Z for X. So we now have two new arrays (A,B,Y,Z,C) and (A*,B*,Y,Z,C*). If (A,B,W,X,C) produces a higher lifetime utility for i than (A*,B*,W,X,C*), then (A,B,Y,Z,C) must produce a higher lifetime utility than (A*,B*,Y,Z,C*). And if (A,B,W,X,C) produces the same lifetime utility as (A*,B*,W,X,C*), then (A,B,Y,Z,C) must produce the same lifetime utility as (A*,B*,Y,Z,C*).

It is easy to see that additive separability implies that the “independence” condition is satisfied. Remember that each sublifetime utility for individual i (that is, each $u_{i,k}$, where k ranges from 1 to K) is *just* a function of the facts in period k . On the independence condition and additive separability, see, e.g., Han Bleichrodt & Amiram Gafni, *Time Preference, The Discounted Utility Model, and Health*, 15 J. HEALTH ECON. 51-53, 55; Shane Frederick et al., *Time Discounting and Time Preference: A Critical Review*, in TIME AND DECISION 13, 21 (George Loewenstein et al., 2003)

The “independence” condition, in turn, precludes an improvement effect. For example, if we are inclined to say that steadily improving sequences of mental states such as (unhappy, neutral, happy) are better for well-being as such, we will violate the “independence” condition. We will say that (unhappy, neutral, happy) is better than (unhappy, happy, neutral), but that (happy, neutral, happy) and (happy, happy, neutral) are equally good, which violates “independence.” For philosophical argument in favor of an improvement effect, see J. David Velleman, *Well-Being and Time*, in THE METAPHYSICS OF DEATH 329, 331 (John Martin Fischer ed., 1993). This argument is buttressed by empirical work. “One of the most

Indeed, I do not assume that lifetime utility is a function of sublifetime utility *at all*. Someone's lifetime well-being in an outcome depends on everything that happens in that outcome, in turn parsable into everything that happens during each period; but it is a further, and contestable, assumption that the rich information about temporal facts in each period can be compressed into a single utility number, with lifetime well-being in turn a function of the utility number in each period. There is no reason to saddle the lifetime view with that assumption either.

C. *Why Welfarism?*

The SWF framework is, I believe, an attractive framework for thinking about policy problems. As already mentioned, the framework is agnostic about the nature of well-being. It can be meshed with objectivist accounts, preferentialist accounts, or mental state accounts (the three leading families of views about well-being), or indeed with any other account.²² For any moral view, whatever its account of well-being, an individual's lifetime or sublifetime utility represents her lifetime or sublifetime attainment with respect to *that*.

Further, as should now be abundantly clear, SWFs are potentially sensitive to equity considerations. Utilitarianism is simply one *variant* of welfarism. Finally, the use of a formal, mathematical vocabulary helps sharpen important distinctions relevant to discussions of equality – for example, what prioritarianism means, what cross-temporal interaction effects are, or (as we shall see in Part III) what the different conceptions of choice under uncertainty involve.

Still, the reader might object to the framework. She might object, first, that the precision mathematical representation imports is purchased at a cost – the cost of ignoring *incommensurability*. The framework tries to represent a moral view's verdict

robust findings in research about assessment of experiences is the clear preference for improvement over time. ... Preference for improvement has been demonstrated in many domains, including monetary payments, experiences such as vacations, queuing events, pain, discomfort, medical outcomes and treatments, gambling, and academic performance.” Dan Ariely & Ziv Carmon, *Summary Assessments of Experiences: The Whole is Different from the Sum of its Parts*, in TIME AND DECISION, *supra*, at 323, 327 (citations omitted). See also Frederick, *supra*, at 28-29; George Loewenstein & Drazen Prelec, *Preferences for Sequences of Outcomes*, 100 PSYCH. REV. 91 92-93 (1993). Other empirical findings inconsistent with the “independence” condition include a preference to spread rather than cluster enjoyable preferences: “maximal endurable time” preferences with respect to health states; and the “peak-end” rule for retrospective evaluations of experiential sequences. See Frederick, *supra*, at 29-30; Aki Tsuchiya & Paul Dolan, *The QALY Model and Individual Preferences for Health States and Health Profiles over Time: A Systematic Review of the Literature*, 25 MED. DECISION MAKING 460, 461 (2005); Ariely & Carmon, *supra*, at 326-27

To be sure, the question how the temporal components of well-being produce lifetime well-being is normative, not empirical. The welfarist is not obliged to incorporate all of the observed violations of the “independence” condition into her account of lifetime well-being. (For example, she might see “peak-end” evaluations as a kind of irrationality.) But I suggest that many plausible accounts of lifetime well-being will violate the “independence” condition in *some* way.

²² For a discussion of these various accounts, see MATTHEW D. ADLER & ERIC A. POSNER, NEW FOUNDATIONS OF COST-BENEFIT ANALYSIS 28-35 (2006).

with respect to the ordering of outcomes by an SWF. But, in fact, there are plausible moral views which incompletely order outcomes – which include pairs of outcomes such that neither is better than the other, nor are they equally good – and no such ordering can be represented by a single SWF operating on a single utility vector for each outcome.²³ I have already suggested, however, that the problem of incommensurability can be handled by conceptualizing a given welfarist moral view as corresponding to a set of pairs of utility functions and SWFs, rather than a single pair.²⁴

A second and quite different objection is that the SWF framework ignores considerations of choice and responsibility. In ranking outcomes, we should be sensitive to information about the causal origins of individuals' well-being attainments, and not simply to those attainments. Did individuals end up where they did because of "brute luck," or because of the foreseeable result of their own choices? Dworkin, Sen, and other prominent philosophers have challenged welfarist accounts of equality on these grounds.²⁵ The difficulty is that no one has yet provided a general framework for integrating considerations of overall well-being and the distribution of well-being (which the SWF framework accomplishes) with information about individual choice and responsibility. This is an active area of research, where attractive solutions may well emerge, but for now the question of how a *responsibility-sensitive* welfarist would analyze the time-slice problem is not one I can address.

II. LIFETIME VERSUS SUBLIFETIME VIEWS: PHILOSOPHICAL ARGUMENTS

Having explicated the SWF framework and the conceptual difference between lifetime and sublifetime views, I now turn to substantive philosophical analysis. Section A presents the basic case for lifetime welfarism. Sections B and C address two important responses to that case: one, by Dennis McKerlie, which adopts a standard picture of personal identity over time; and a second, by Derek Parfit, which trades on a deflationary view of personal identity.

A. *The Basic Case for the Lifetime View*

A number of prominent philosophers of equality have adopted a lifetime view. This includes Thomas Nagel, who writes: "the subject of an egalitarian principle is not the distribution of particular rewards to individuals at some time, but the prospective quality of their lives as a whole, from birth to death"²⁶ It also includes John Rawls and Ronald Dworkin, who believe that distributive justice requires a fair allocation of lifetime shares of "primary goods" or "resources."²⁷ The philosophers Norman Daniels and Dan Brock, who have theorized about the distribution of health care, have also

²³ An SWF assigns a single real number to each utility vector. Every real number is greater than, less than, or precisely equal to every other real number.

²⁴ See *supra* note 9 and accompanying text.

²⁵ See Arneson, *supra* note 2.

²⁶ THOMAS NAGEL, EQUALITY AND PARTIALITY 69 (1991). See also THOMAS NAGEL, MORTAL QUESTIONS 120, 124 n.16 (1979).

²⁷ See JOHN RAWLS, A THEORY OF JUSTICE 78, 178 (1971); Ronald Dworkin, *What Is Equality? Part 2: Equality of Resources*, 10 PHILOSOPHY & PUB. AFF. 283, 304-05 (1981).

argued for the lifetime perspective, as has Alan Williams, a leading health economist.²⁸
To quote Brock:

People's lives extend continuously over time, and our moral concern should be for the lives they lead, not simply for how good their lives are at a particular point in time, whether now or sometime in the past or the future. A concern for the worse off should reflect this concern for people's lives as a whole and so not ignore the duration of people's poor health. Moreover, in other areas besides health we often take what a person has had or suffered in the past to be relevant in distributing scarce benefits now; if one child has had little opportunity for travel in the past in comparison with her well-traveled sibling, fairness supports giving a travel opportunity to her now that can only go to one of them²⁹

The basic argument in favor of a lifetime view is straightforward. I will articulate this argument in terms of well-being, because (again) the framework for this Article is welfarist. Egalitarian considerations enter morality because it is important to treat separate persons fairly.³⁰ It is unfair that some individuals end up with more well-being than others, at least where the worse-off individuals are not responsible for their plight. Another way to put the same idea is this: individuals are moral equals; as equals, individuals have equal claims to flourish, to have good lives; and therefore a distribution of well-being that departs from equality is *prima facie* problematic. Such a distribution may be justified, all-things-considered – for example, a Pareto-superior move away from perfect equality will be justified, all-things-considered – but it has the troubling feature that some individuals' claims to well-being are more fully satisfied than others'.

On this picture of the moral force of equality, well-being is a quantifiable characteristic of each person in each outcome. Call it “property W”. Individuals have equal claims to property W; and the unequal distribution of property W is, *prima facie*, morally problematic.

Let us adopt, now, a traditional premise about personal identity: namely, that each human is a single person, from birth to death. The upshot is that *individuals* or *persons* endure over time. A single person has multiple temporal stages. What, then, is property W? We cannot say that property W is the individual's sublifetime well-being, full stop – because (on the assumption of multiple temporal stages) there are multiple sublifetime well-being properties associated with each person. We *might* adopt the view

²⁸ See Dan W. Brock, *Priority to the Worst Off in Health-Care Resource Prioritization*, in MEDICINE AND SOCIAL JUSTICE 362, 370 (Rosamond Rhodes et al. eds., 2002); NORMAN DANIELS, JUSTICE AND JUSTIFICATION 257-83 (1996); NORMAN DANIELS, AM I MY PARENTS' KEEPER (1988); NORMAN DANIELS, JUST HEALTH CARE 86-113 (1985); Alan Williams, *Intergenerational Equity: An Exploration of the “Fair Innings” Argument*, 6 HEALTH ECON. 117 (1997).

²⁹ Brock, *supra* note 28, at 370.

³⁰ Louis Kaplow and Steven Shavell have emphasized the distinction between “fairness” and “welfarism,” so it may seem oxymoronic for the equity-regarding *welfarist* to appeal to fairness considerations. See LOUIS KAPLOW & STEVEN SHAVELL, FAIRNESS VERSUS WELFARE (2002). But a view which understands well-being itself as the “currency” of distributive justice, and explains the weight of distributional considerations by pointing out the unfairness of giving some individuals more than others, is fully welfarist in the standard sense and indeed in Kaplow and Shavell's. Such a view satisfies the supervenience requirement definitional of welfarism: no outcome is ranked better than another unless there is some well-being difference between the outcomes.

that property W is each individual's sublifetime well-being during some stipulated stage (say, the last year of her life, or the first year, or the middle year). But picking one stage over another as the stipulated stage seems arbitrary. How the individual does during the stipulated stage is no more a property of that individual, and no more a well-being property, than how she does during any other stage.

Further, for any stipulated stage, we can imagine lowering some person P's well-being during that stage, with a compensating increase in a different stage, so that her lifetime well-being remains the same. It would be odd to think that this shift, without more, could change our assessment of the fairness of the distribution of well-being as between P and other individuals. The extent to which P's claim to well-being is satisfied, and to which other individuals' claims to well-being are satisfied, seem to be unchanged by the shift. P has been (seemingly) compensated for the loss in well-being during the stipulated stage by an increase in well-being at another time. But, if property W is well-being during the stipulated stage, this temporal shift in P's well-being – without more – *could* affect our assessment of the fairness of distribution of well-being as between P and other individuals.

Here's another way to make this last point. If individuals endure over time, then – as a matter of honoring individuals' claims to well-being – it should always or at least often be possible to compensate an individual for a loss of well-being during one period with a gain during another. Call this “the possibility of intrapersonal, intertemporal compensation.” (PIIC). But, in turn, PIIC will undercut the identification of property W with an individual's sublifetime well-being during any stipulated stage. With PIIC in hand, we can always construct cases in which a change in individuals' well-being during the stipulated stage, with compensating changes later or earlier in their lives, produces no true change in how equally well-being is distributed across persons, but *would* be recorded as changing that distribution by a sublifetime view focused on the stipulated stage. As Nagel explains:

[The possibility of intrapersonal compensation] implies ... that *if* an egalitarian theory is accepted, it should apply only across lives rather than within them. It is a reason for taking individual human lives, rather than individual experiences, as the units over which any distributive principle should operate.³¹

The upshot is that property W is lifetime well-being. Given the persistence of individual identity over time, it seems hard to reject PIIC; and PIIC, in turn, means that the well-being characteristic whose fair distribution equality requires cannot be well-being during any stipulated sublifetime stage.

Thus the basic case for a lifetime account. We shall now consider two arguments for a sublifetime view: one, principally advanced by Dennis McKerlie, which does not make revisionary claims about personal identity over time; and another, advanced by Derek Parfit, which does.

³¹ NAGEL, MORTAL QUESTIONS, *supra* note 26, at 120.

B. McKerlie's Response

In a series of articles, Dennis McKerlie has argued for the moral relevance of sublifetime well-being *without* challenging the standard view that each human is one and the same person from birth until death.³² In the early articles, McKerlie adopts what might be seen as a kind of sublifetime comparativism. In the more recent work, McKerlie has shifted to a sublifetime prioritarianism. I shall consider both of these arguments, as well as a third issue – which McKerlie doesn't raise – concerning the consistency of the lifetime view and sufficientism.

1. Intuitions about Simultaneous Sublifetime Inequalities: A Sublifetime Comparativism?

In the early work, McKerlie appeals to our intuitions about synchronizing certain sublifetime sources of well-being across persons which, allegedly, show that we prefer to equalize sublifetime well-being even when lifetime well-being is held constant.³³ He provides the following sort of example.

[Imagine a society] that contains great inequality, with happier lives attached to certain social positions. But at a fixed time people change places and switch from a superior position to an inferior one or vice versa. One example would be a feudal society in which peasants or nobles exchange roles every ten years. The result is that people's lives as wholes are equally happy. Nevertheless during a given time period the society contains great inequality If equality between complete lives were all that mattered, an egalitarian could not object to it. But I think that many egalitarians would find it objectionable.³⁴

To give a modern example, as between two societies in which individuals sometimes have periodic incomes of \$20,000, sometimes periodic incomes of \$100,000, we prefer (ceteris paribus) the society in which everyone has the lower income at the same time.

In other words, McKerlie suggests that we intuitively prefer outcome II to outcome I in the following sort of case.

³² See Dennis McKerlie, *Egalitarianism and the Difference Between Interpersonal and Intrapersonal Judgments*, in *EGALITARIANISM* (Nils Holtug & Kasper Lippert-Rasmussen ed., 2007); Dennis McKerlie, *Justice Between the Young and the Old*, 30 *PHILOSOPHY & PUB. AFF.* 152 (2002); Dennis McKerlie, *Dimensions of Equality*, 13 *UTILITAS* 263 (2001); Dennis McKerlie, *Equality Between Age-Groups*, 21 *PHILOSOPHY & PUB. AFF.* 275 (1992); Dennis McKerlie, *Equality and Time*, 99 *ETHICS* 475 (1989)..

³³ My analysis in this subsection is indebted to work by Klemens Kappel and R.I. Sikora, who provide a variety of criticisms of McKerlie's sublifetime comparativism. See Klemens Kappel, *Equality, Priority, and Time*, 9 *UTILITAS* 203 (1997); R.I. Sikora, *Six Viewpoints for Assessing Egalitarian Distributions Schemes*, 99 *ETHICS* 492 (1989).

³⁴ McKerlie, *Equality and Time*, *supra* note 32, at 479. Larry Temkin offers a similar example. See TEMKIN, *supra* note 15, at 235-38.

	Outcome I			
	<i>Time Period</i>			
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>Person A</i>	Peasant	Noble	Peasant	Noble
<i>Person B</i>	Noble	Peasant	Noble	Peasant

	Outcome II			
	<i>Time Period</i>			
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>Person A</i>	Peasant	Noble	Peasant	Noble
<i>Person B</i>	Peasant	Noble	Peasant	Noble

OR

	Outcome I			
	<i>Time Period</i>			
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>Person A</i>	\$20,000	\$100,000	\$20,000	\$100,000
<i>Person B</i>	\$100,000	\$20,000	\$100,000	\$20,000

	Outcome II			
	<i>Time Period</i>			
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>Person A</i>	\$20,000	\$100,000	\$20,000	\$100,000
<i>Person B</i>	\$20,000	\$100,000	\$20,000	\$100,000

Let us grant McKerlie the intuition. The question is how one explains it. One *could* explain the intuition by assuming that being a peasant or having an income of \$20,000 has a low context-invariant sublifetime utility (say, 2) and that being a noble or having an income of \$100,000 has a high context-invariant sublifetime utility (say, 10). This is the picture McKerlie seems to be tendering. In other words, McKerlie seems to be suggesting that the sublifetime and lifetime well-being attainments of the individuals in outcome I are accurately represented by the utility numbers in I*, and that the sublifetime and lifetime well-being attainments of the individuals in outcome II are accurately represented by the utility numbers in II*.

I* (Possible sublifetime and lifetime utilities corresponding to Outcome I)						
	<i>Time Period</i>				<i>Lifetime Utility</i>	
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>		
<i>Person A</i>	2	10	2	10	X	
<i>Person B</i>	10	2	10	2	X	

II* (Possible sublifetime and lifetime utilities Corresponding to Outcome II)						
	<i>Time Period</i>				<i>Lifetime Utility</i>	
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>		
<i>Person A</i>	2	10	2	10	X	
<i>Person B</i>	2	10	2	10	X	

Clearly, the lifetime view cannot explain a preference for II* over I*. Any SWF, applied to lifetime utilities, will rank the outcomes the same, because the vectors of lifetime utilities are identical. Note also that a sublifetime prioritarian view – a sum of an increasing, strictly concave function g of sublifetime utilities – cannot explain a preference for II* over I*. Whatever the g -function, the sum of that function applied to the sublifetime utilities is the same in I* and II*. However, there *are* equity-regarding SWFs, applied to sublifetime utilities, which would favor II* over I*. In particular, a comparativist view that focused on the pattern of sublifetime utility within each time period could explain a preference for II* over I*. Here’s a simple example: an SWF which aggregates sublifetime utilities, but gives an outcome a premium for every time period in which everyone’s sublifetime utility is exactly equal.

But why assume that outcomes I and II are accurately represented by I* and II*? The lifetime welfarist will see outcome I as an instance in which individual inequality in some time period with respect to some source of well-being – social status in the noble/peasant case, income in the \$100,000/\$20,000 case— itself lowers lifetime well-being. In other words, it is worse for individual P to be a peasant during some portion of his life *if other members of the society are nobles at the same time* than to be a peasant during some portion of his life *when everyone else in the society is also a peasant*. Being a peasant with nobles in the vicinity reduces P’s lifetime well-being because it is demoralizing, or frustrating, or lowers self-respect. The impact on lifetime well-being of the status of peasant during some portion of a life is not context invariant. Similarly, it is worse for P to have a life in which he has an income of \$20,000 during some period *when some members of the society have higher incomes during that period* than to have that very same income during a period when everyone else has the same income.

In short, the lifetime welfarist can see outcomes I and II as accurately represented by I+ and II+.

I+ (Possible lifetime utilities
corresponding to Outcome I -- noble/peasant version)

	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
Person A	Peasant	Noble	Peasant	Noble	X
Person B	Noble	Peasant	Noble	Peasant	X

II+ (Possible lifetime utilities
corresponding to Outcome II -- noble/peasant version)

	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
Person A	Peasant	Noble	Peasant	Noble	Y>X
Person B	Peasant	Noble	Peasant	Noble	Y>X

OR

I+ (Possible lifetime utilities
corresponding to Outcome I -- \$100,000/\$20,000 version)

	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
Person A	\$20,000	\$100,000	\$20,000	\$100,000	X
Person B	\$100,000	\$20,000	\$100,000	\$20,000	X

II+ (Possible lifetime utilities
corresponding to Outcome II -- \$100,000/\$20,000 version)

	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
Person A	\$20,000	\$100,000	\$20,000	\$100,000	Y>X
Person B	\$20,000	\$100,000	\$20,000	\$100,000	Y>X

Any SWF that respects the Pareto principle for outcomes and is applied to lifetime utilities, including any equity-regarding SWF, will prefer II+ to I+.³⁵

McKerlie has given us a hypothetical case – a case in which we improve an outcome by synchronizing individuals’ attainments with respect to some source of well-being (here, being a peasant versus being a noble, or having \$20,000 versus having \$100,000) – and argues that the social improvement occurs by equalizing sublifetime well-being, without changing the individuals’ lifetime well-being. I suggest, instead, that

³⁵ For a simple way to generate a lifetime utility number Y in II+ that is greater than X in I+, assign a sublifetime utility of 2 to being a peasant with other peasants around (or having \$20,000 where everyone else has \$20,000 at the same time), 10 to being a noble with other nobles around, 1 to being a peasant with nobles around, and 10.5 to being a noble with peasants around; and then make lifetime utility the sum of these context-variant sublifetime utilities.

improvement occurs because the synchronization itself changes the individuals' lifetime well-being.

What favors my interpretation? To begin, we have the basic case for lifetime welfarism, presented earlier. In addition, and crucially, in cases where the synchronization clearly does *not* affect lifetime well-being, our intuition in favor of synchronization is much weaker. Imagine a society in which everyone works 5 days a week, spends one weekend day doing chores, and one weekend day having fun. In the first outcome, half the population does chores on Saturdays and has fun on Sundays; the other half has the fun on Saturdays. In a second outcome, weekend activities are synchronized so that everyone does chores on Saturdays. Or (to use a longer time slice), in some society everyone has a minor depressive episode, lasting several years, at one point during his or her life. In one outcome, these episodes occur at different times; in a second outcome, they are synchronized. The sublifetime comparativist who sees the noble-peasant case or the \$100,000/\$20,000 income case as characterized by I* and II* will also presumably prefer to synchronize weekend chores and fun, or to synchronize depressive episodes, in the cases just described. But the intuition in favor of doing so (I suggest) is much weaker in these cases.

Indeed, as McKerlie has come to acknowledge in more recent work, the problem with a sublifetime comparativism that prefers to reduce simultaneous sublifetime inequality is that the moral role of simultaneity is suspect.³⁶ Why should we be bothered by the fact that person A has a lower sublifetime well-being during some period as compared to person B's sublifetime well-being during the same period, but not that person A has a lower sublifetime well-being during some period as compared to person B's sublifetime well-being during some other period? Consider, again, outcomes I* and II*. In each outcome, there are 16 pairings of the form $(u_{1,k}; u_{2,l})$ with $u_{1,k}$ the sublifetime utility of the first person, person A, during period k , and $u_{2,l}$ the sublifetime utility of person B during period l , k and l each ranging from one to four. Note further that, in each outcome, there are four such pairings with the numbers (2,2); eight pairings (10,2) or (2,10); and 4 pairings (10,10). And, of course, in each outcome the two individuals' lifetime utilities are the same. So why *should* we morally prefer II* to I*? Such a preference has no theoretical warrant, given a standard picture of personal identity; it is not needed to explain our intuitions in cases such as the noble-peasant or \$100,000/\$20,000 cases, as I have argued; and it is inconsistent with intuitions in other cases (such as weekend chores or depressive episode).

2. Sublifetime Prioritarianism

McKerlie now advocates a sublifetime prioritarianism. He suggest, to begin, that sublifetime prioritarianism actually provides the best account of individual "prudence" or

³⁶ For recent work in which McKerlie has moved away from sublifetime comparativism and towards sublifetime prioritarianism, see *Egalitarianism and the Difference between Interpersonal and Intrapersonal Judgments*, *supra* note 32; *Dimensions of Equality*, *supra* note 32, at 273-88; *Justice Between the Young and the Old*, *supra* note 32, at 163-74.

self-interest.³⁷ On this account, a prudent individual -- one who maximizes her lifetime well-being -- will give greater weight to changes in her sublifetime well-being that occur when she is at a low level of sublifetime well-being, as compared to changes in her sublifetime well-being that occur when she is at a higher level of sublifetime well-being. In short, an individual's lifetime well-being in an outcome is the sum of an increasing, strictly concave function of her sublifetime well-being, not a simple sum of her sublifetime well-being.

McKerlie's Account of Prudence

	<i>Time Period</i>				<i>Lifetime Utility (g increasing and strictly concave)</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
Person A	$u_{1,1}$	$u_{1,2}$	$u_{1,3}$	$u_{1,4}$	$g(u_{1,1}) + g(u_{1,2}) + g(u_{1,3}) + g(u_{1,4})$
Person B	$u_{2,1}$	$u_{2,2}$	$u_{2,3}$	$u_{2,4}$	$g(u_{2,1}) + g(u_{2,2}) + g(u_{2,3}) + g(u_{2,4})$

Why does this matter? Note that, even if McKerlie's account of prudence is correct, there is still a difference between an equity-regarding view that evaluates outcomes by applying a given prioritarian SWF to lifetime utilities, and an equity-regarding view that applies that same SWF to sublifetime utilities. The account of prudence does not show that the two approaches converge.³⁸

³⁷ See *Dimensions of Equality*, supra note 32, at 281-88; *Justice Between the Young and the Old*, supra note 32, at 168-74.

³⁸ The sublifetime prioritarian would evaluate outcomes by applying the increasing and strictly concave g -function to sublifetime utilities, and summing. The lifetime prioritarian would evaluate outcomes by using the g -function once to determine each individual's lifetime utility as a function of her sublifetime utilities, and then by calculating the sum of the g -function applied to these lifetime utilities. For example, imagine that the lifetime and sublifetime prioritarian both use the square root as the g -function, and that each individual's lifetime utility is the sum of the square root of sublifetime utility. Then the lifetime prioritarian will say that outcome II below is worse than outcome I, while the sublifetime prioritarian will say that they are equally good.

	Outcome I (showing sublifetime and lifetime utilities)				
	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	4	9	4	9	10
<i>Person B</i>	9	4	9	4	10

Sum of square root of sublifetime utilities: 20

Sum of square root of lifetime utilities: 6.32

	Outcome II (showing sublifetime and lifetime utilities)				
	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	4	25	4	9	12
<i>Person B</i>	9	4	9	0	8

What it would mean, more narrowly, is that sublifetime prioritarianism would be the correct approach to choosing between outcomes *if* individuals were identical.³⁹ In that sense (perhaps?!) McKerlie's account of prudence might be seen as offering a kind of analogical support for sublifetime prioritarianism in the more general and realistic case where individuals are not identical. Further, on McKerlie's account of prudence, sublifetime prioritarianism would never conflict with the Pareto principle applied to lifetime utilities (even where individuals are heterogeneous).⁴⁰ The divergence between

Sum of square root of sublifetime utilities: 20
 Sum of square root of lifetime utilities: 6.29 (which is less than 6.32)

³⁹ In other words, in comparing outcomes like I and II below, a lifetime prioritarian view that determines each individual's lifetime utility as the sum of a *g*-function of sublifetime utilities (*g* increasing and strictly concave), and then calculates social value as the sum of that *g*-function of lifetime utilities, will reach the same result as a sublifetime prioritarianism that simply applies the *g*-function to all the sublifetime utilities and sums. This is illustrated by using the square root as the *g*-function, but is true for any increasing, strictly concave *g*-function.

Outcome I (showing sublifetime and lifetime utilities)

	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	4	9	4	9	10
<i>Person B</i>	4	9	4	9	10

Sum of square root of sublifetime utilities: 20
 Sum of square root of lifetime utilities: 6.32

Outcome II (showing sublifetime and lifetime utilities)

	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	25	36	0	0	11
<i>Person B</i>	25	36	0	0	11

Sum of square root of sublifetime utilities: 22
 Sum of square root of lifetime utilities: 6.63

In such cases of identical individuals, if in one outcome the sum of the *g*-function of sublifetime utilities is greater, than each individual's lifetime utility must be greater as well, and (because the *g* function is increasing, i.e., the lifetime prioritarian satisfies the Pareto principle for outcomes) the lifetime prioritarian must see that outcome as better.

⁴⁰ Imagine that one outcome is a Pareto-improvement in lifetime utilities over a second. In other words, either every individual is better off in the first outcome, or some individuals are better off and the others are equally well-off. In either event, it would follow -- since each individual's lifetime utility is the sum of the *g*-function of her sublifetime utilities -- that the sum of the *g*-function of sublifetime utilities is greater in the first outcome. But then the sublifetime prioritarian must say that the first outcome is better. For example, consider the following case, more once using the square root as the *g*-function for the sake of illustration.

Outcome I (showing sublifetime and lifetime utilities)

	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	4	9	4	9	10
<i>Person B</i>	9	4	9	4	10

sublifetime and lifetime approaches would be limited to pairs of outcomes I and II where some individuals have greater lifetime well-being in one outcome and others have greater lifetime well-being in the other. Thus, McKerlie’s account of prudence – if true – would soften the conflict between sublifetime and lifetime approaches in the case of heterogeneous individuals, and thus might be seen as bolstering any preexisting argument for the sublifetime approach.

But McKerlie’s is almost surely *not* the correct account of prudence. If there are interactions between an individual’s sublifetime well-being in different stages in producing her lifetime well-being, such as improvement effects, additive separability fails and lifetime well-being cannot be represented as a function of sublifetime well-being – whether a simple sum of sublifetime utility, or the sum of an increasing, strictly concave function of sublifetime utility.⁴¹

A second and stronger argument for sublifetime prioritarianism, advanced by McKerlie as well as some other philosophers, appeals to our intuitions about pain and other forms of suffering. The claim is that, intuitively, someone who is currently in great pain has a stronger claim on our assistance, as opposed to someone who is currently in a less painful state, even if the first person’s lifetime well-being is greater. To simplify the analysis, here, I will ignore non-hedonic sources of well-being and assume that lifetime and sublifetime well-being is just a function of negative affects (pains) and positive affects (pleasures).

Consider first a preference for outcome X over outcome Y in the case below. Affects can be strongly, moderately, or mildly negative; neutral; or mildly, moderately, or strongly positive. (In the table, “mod.” is short for “moderately”.)

Sum of square root of sublifetime utilities: 20
 Sum of square root of lifetime utilities: 6.32

Outcome II (showing sublifetime and lifetime utilities)

	<i>Time Period</i>				<i>Lifetime Utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	4	25	4	9	12
<i>Person B</i>	9	4	9	4	10

Sum of square root of sublifetime utilities: 22
 Sum of square root of lifetime utilities: 6.63

⁴¹ On additive separability, see *supra* note 21 and accompanying text.

	Status quo			
	<i>Time Period</i>			
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>Person A</i>	mod. positive	mod. positive	strong negative	mod. positive
<i>Person B</i>	mild negative	mild negative	mild negative	mild negative

	Outcome X (relieving A's pain in period 3)			
	<i>Time Period</i>			
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>Person A</i>	mod. positive	mod. positive	mod. negative	mod. positive
<i>Person B</i>	mild negative	mild negative	mild negative	mild negative

	Outcome Y (relieving B's pain in period 3)			
	<i>Time Period</i>			
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>
<i>Person A</i>	mod. positive	mod. positive	strong negative	mod. positive
<i>Person B</i>	mild negative	mild negative	neutral	mild negative

In the status quo, person A's affects are moderately positive in all periods except period 3, where he suffers great pain, while person B's affects are mildly negative in all the periods. In outcome X, we give A the pain relief in period 3, thereby moving him from a strongly negative state in that period to a moderately negative state. In outcome Y, we give B the pain relief in period 3, thereby moving him from a mildly negative state to a neutral state.

Grant the intuition that we prefer outcome X to outcome Y in this case. This sort of intuition about pain relief is one that the lifetime welfarist can accommodate. It is possible that the lifetime well-being A would experience, as a result of a strongly negative affect in one period, is worse than the lifetime well-being of all the other lives in the example. The strongly negative affect in that period may "ruin" that whole life, either because of non-additivities, or even within the context of an additively separable model of lifetime well-being, so that the choice between X and Y is as follows.

	Outcome X (relieving A's pain in period 3)				
	<i>Time period</i>				<i>Lifetime utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	mod. positive	mod. positive	mod. negative	mod. positive	56
<i>Person B</i>	mild negative	mild negative	mild negative	mild negative	48

	Outcome Y (relieving B's pain in period 3)				
	<i>Time period</i>				<i>Lifetime utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	mod. positive	mod. positive	strong negative	mod. positive	45
<i>Person B</i>	mild negative	mild negative	neutral	mild negative	49

Any prioritarian SWF, applied to lifetime utilities, will pick outcome X over outcome Y.⁴²

But there are other examples involving pain relief that a lifetime view has more difficulty handling. Klemens Kappel provides the following example:

In a casualty department in a hospital two persons, A and B, are in pain. Unfortunately, only one of them can be given pain relieving treatment since there is a shortage of the particular painkiller needed. A's pain is much stronger than B's pain. However, B did in the past suffer from pains similar to those he has now, and these pains went on for a considerably longer time and were not treated. ... Suppose, finally, that whoever we treat, the benefit will be the same because A's stronger pain will be only partly relieved, while B's pain will be fully relieved.⁴³

Now, the outcomes are X* and Y*. B is in strongly negative affects in all periods except period 3; A is in a strongly negative affect in period 3, otherwise in a moderate positive state; and we have to choose between giving the pain relief to A in period 3 (outcome X*) or to B in period 3 (outcome Y*).

	Status quo*			
	<i>Time Period</i>			
	1	2	3	4
<i>Person A</i>	mod. positive	mod. positive	strong negative	mod. positive
<i>Person B</i>	strong negative	strong negative	mild negative	strong negative

	Outcome X* (relieving A's pain in period 3)			
	<i>Time Period</i>			
	1	2	3	4
<i>Person A</i>	mod. positive	mod. positive	mod. negative	mod. positive
<i>Person B</i>	strong negative	strong negative	mild negative	strong negative

	Outcome Y* (relieving B's pain in period 3)			
	<i>Time Period</i>			
	1	2	3	4
<i>Person A</i>	mod. positive	mod. positive	strong negative	mod. positive
<i>Person B</i>	strong negative	strong negative	neutral	strong negative

In this kind of case, by contrast with the first pain case, it seems very clear that A's lifetime well-being will be higher than B's even if we give B the pain relief and not

⁴² To get these lifetime utilities within the context of an additively separable model, assume that lifetime utility is the simple sum of sublifetime utility; that the sublifetime utility of a strongly negative state is 0; and that the sublifetime utilities of moderately negative, mildly negative, neutral, mildly positive, moderately positive, and strongly positive states are, respectively, 11, 12, 13, 14, 15 and 16. Given anonymity, the lifetime prioritarian must count the second outcome as equal to one in which person A gets 49 and person B gets 45, which in turn is Pareto inferior in lifetime utilities to the first outcome.

⁴³ Kappel, *supra* note 33, at 211. See also DEREK PARFIT, REASONS AND PERSONS 341, 344 (1984). Derek Parfit, *Comments*, 96 ETHICS 832, 869-70 (1986).

A. Therefore, an intuition that we should choose outcome X *(relieving A's sublifetime pain) is tougher for the lifetime welfarist to handle.

But maybe not impossible. *If* giving A the pain relief produces a greater change in lifetime well-being than giving B the pain relief, a lifetime prioritarian could justify the choice of X* over Y*. In other words, the facts in X* and Y* might correspond to the utilities in X+ and Y+. Assume that strongly, moderately, and mildly negative affects have a sublifetime utility of 0, 3, 5 respectively; a neutral affect has a sublifetime utility of 7; mildly, moderately, and strongly positive affects have a sublifetime utility of 9, 11, and 13 respectively and that (absent an improvement affect) lifetime utility is the simple sum of sublifetime utility.

X+ (possible sublifetime and lifetime utilities corresponding to outcome X*)

	<i>Time period</i>				<i>Lifetime utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	11	11	3	11	36
<i>Person B</i>	0	0	5	0	5

Y+ (possible sublifetime and lifetime utilities corresponding to outcome Y*)

	<i>Time period</i>				<i>Lifetime utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	11	11	0	11	33
<i>Person B</i>	0	0	7	0	7

Observe that a utilitarian SWF, applied to lifetime utilities, *will* choose outcome X+ over outcome Y+ (41 is greater than 40) and a prioritarian SWF which is sufficiently close to utilitarianism and which is applied to lifetime utilities will do so as well.

Kappel, however, has tried to word the case so that the pain relief produces the very same change in sublifetime and lifetime well-being whomever we treat. He intends that X* and Y* be represented by the sublifetime and lifetime utilities in X++ and Y++. (I have changed the sublifetime utility of a strongly negative affect from 0 to 1, so that the change in sublifetime and lifetime utility is 2 whether we give the pain relief to A or B.)

X++ (possible sublifetime and lifetime utilities corresponding to outcome X*)

	<i>Time period</i>				<i>Lifetime utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	11	11	3	11	36
<i>Person B</i>	1	1	5	1	8

Y++ (possible sublifetime and lifetime utilities corresponding to outcome Y*)

	<i>Time period</i>				<i>Lifetime utility</i>
	<i>1</i>	<i>2</i>	<i>3</i>	<i>4</i>	
<i>Person A</i>	11	11	1	11	34
<i>Person B</i>	1	1	7	1	10

Choosing X++ over Y++ would seem impossible for a lifetime equity-regarding welfarist to justify. Any lifetime prioritarian will choose Y++⁴⁴; a lifetime sufficientist will choose Y++ or be indifferent; and although in principle a lifetime comparativist view might fall within my definition of equity-regarding and choose X++, that just shows the overbreadth of that definition.

So has the lifetime welfarist been cornered? Maybe not. If X* and Y* are accurately represented by the utility numbers in X++ and Y++, we may no longer have the intuition that we should choose X*.

Even if we *do* have the intuition that we should choose X* over Y* when represented by the utility numbers in X++ and Y++, that is not necessarily strong evidence in favor of sublifetime welfarism, as against lifetime welfarism. The cases that are deployed in the literature to show a preference for aiding someone at a low sublifetime level of well-being during some period, even when she has a higher level of lifetime well-being, all involve pain and suffering. It is difficult to construct such cases where the well-being differentials involve the absence of a “good,” rather than an affirmative “bad” (suffering resulting from physical pain, malnourishment, homelessness, etc.) Imagine that B has had a life without enjoyment, and will likely lack enjoyment in the future, but this month is enjoying himself moderately; A has had an enjoyable life in the past, and will likely enjoy himself in the future, but this month is bored. We can increase B’s enjoyment by some amount, or A’s enjoyment by the same amount. The sublifetime prioritarian will say that we should increase A’s enjoyment – but *that* seems counterintuitive, or at least is less strongly intuitive than giving A the pain relief in the X* /Y* case.

A joint intuition (1) in favor of X* over Y* even when represented by the utility numbers in X++ and Y++, but (2) to give the enjoyment to B rather than A in the enjoyment case, is not evidence in favor of sublifetime prioritarianism. Rather it points

⁴⁴ This is because Y++ is produced from X++ by effecting a Pigou-Dalton transfer in lifetime utilities, namely transferring 2 units of lifetime utility from A, at level 36, to B, at level 8.

beyond welfarism entirely, to a nonwelfarist duty to alleviate current pain and suffering even when our best welfarist account (lifetime equity-regarding welfarism), would require a different course of action. Seeing the pain cases this way *could* lead us to reject welfarism. Alternatively, however, in the familiar process of “reflective equilibrium,” we might stick by the theory of welfarism – for all the other reasons that warrant it – and see the intuition as a real but recalcitrant one, inconsistent with the moral view that we take, on balance, to be best justified.

3. Sublifetime Sufficiency

We have considered challenges to lifetime welfarism in favor of, first, sublifetime comparativism and, then, sublifetime prioritarianism. For symmetry, we should consider a third challenge: that *sufficiency* is best construed in sublifetime rather than lifetime terms.

The worry is that lifetime sufficiency seems implausible. The essence of sufficiency is the notion that there is some threshold of well-being, above which claims for redistribution become qualitatively weaker. What would that threshold be, in terms of lifetime well-being? The point is not just that it would be hard to justify a particular lifetime threshold, rather than taking it as a primitive grounded in intuitions. (After all, it might be hard for prioritarians of any stripe to justify a particular curvature for their SWF.) The point is rather that (I hazard to guess) most of us lack any intuitions at all about what a “sufficient” lifetime well-being would be. Instead, intuitive thresholds tend to be momentary or sublifetime: being well fed now, being sufficiently clothed now, having shelter now, having enough income for this year, and so forth.

However, some sufficientists have entertained the possibility of lifetime sufficiency. Roger Crisp, a leading proponent of sufficiency, does so.

[I claimed] that the compassion principle provides the best account of when, and why, we should give priority to the worse off. Compassion for any being gives out, I suggested, when that being has enough welfare, and it may be based on assessment of a being’s life as a whole or on the individual’s position at the time of assessment.⁴⁵

Equity-regarding welfarists, in short, are not compelled to choose between (1) sufficiency, and the (2) the basic case for lifetime welfarism, given the continuity of personal identity over time. Lifetime sufficiency incorporates both (1) and (2). Alternatively, the equity-regarding welfarist could respond to the counterintuitiveness of lifetime sufficiency by abandoning (1), preserving equity-regard plus a whole-life time-slice in the form of lifetime comparativism or prioritarianism.

To be sure, another response is to abandon (2) but not (1): to adopt sublifetime sufficiency. Sublifetime sufficiency is, admittedly, intuitively more appealing than lifetime sufficiency. Momentary or periodic well-being poverty lines seem less odd than a lifetime well-being poverty line. On the other hand, even a momentary well-being

⁴⁵ Crisp, *supra* note 18, at 763.

poverty line is hard to specify-- a point borne out by the literature on poverty measurement, where the problem of identifying the threshold of "utility" poverty is seen as a difficult, unresolved problem. Sublifetime sufficientism, I suggest, is not a moral view which is sufficiently powerful and unproblematic to warrant us in rejecting the basic case for lifetime welfarism.

C. *Derek Parfit: Equity-Regarding Welfarism plus Non-standard personal identity*

What if we revise the standard account of personal identity? The standard view says that a single human being is, determinately, a single person (in normal cases) from birth to death and that personhood and personal identity are morally significant properties. One possible revision to this view says that different stages of a human life are, determinately, different persons. A single human is one person from the age of zero to 20, a different person from the age of 20 to 40, and so forth. But this view, on a little reflection, seems unappealing. There is normally no abrupt change in psychological abilities, or break in psychological connectedness, at any point in life (at least after the development of basic abilities in early childhood), that would constitute the determinate ending of the existence of one person and the beginning of another.

Derek Parfit, in his seminal book *Reasons and Persons*, takes a very different tack.⁴⁶ Parfit's views are subtle and complicated, and I can only crudely summarize them here. A human, at one point in time, may be more or less psychologically connected with himself at a prior point in time -- or for that matter with a different human. A human-at-a-time H is psychologically connected with a human-at-a-time H* if there are links of memory, belief, desire, intention and so forth between the two. For example, H* and H have the same belief. Or H* remembers something that H experiences. Note that psychological connectedness is intransitive: if H* is connected with H+, and H+ with H++, it doesn't follow that H++ is connected with H*. A different relation is continuity, namely the existence of overlapping chains of connectedness, which is transitive. If H* is connected to H+, and H+ to H++, then H*, H+, and H++ are all continuous with each other. Parfit then suggests that personal identity consists in psychological continuity and/or connectedness with the right cause -- paradigmatically having the same brain.

One implication of Parfit's account is that personal identity can be indeterminate. It might be the case that two particular persons are neither determinately identical to each other, nor determinately not identical. This is striking, because we normally think of the identity relation -- at least with respect to objects such as persons -- as being determinate.

⁴⁶ See PARFIT, *supra* note 43, at 199-347. Parfit also discusses personal identity in *Comments*, *supra* note 43; and Derek Parfit, *The Unimportance of Identity*, in PERSONAL IDENTITY 292 (Raymond Martin & John Barresi eds., 2003). For an excellent recent treatment of personal identity that discusses Parfit's views, see JEFF MCMAHAN, THE ETHICS OF KILLING 3-94 (2002). For a discussion focused on the implications of Parfit's account of personal identity for morality, see Bart Schultz, *Persons, Selves, and Utilitarianism*, 96 ETHICS 721 (1986).

Importantly for our purposes, however, Parfit does *not* argue that personal identity is indeterminate within the confines of a normal human life. His examples of indeterminacy involve esoteric cases where a spectrum of brain operations produce fewer and fewer psychological connections, or less and less preservation of brain matter, between consecutive stages of a human life. As for the case of a normal human life, Parfit seems to say that personal identity over time is determinate, or at least exists to a very high degree. He writes:

In ordinary cases, questions about our identity have answers. In such cases, there is a fact about personal identity, and Reductionism [Parfit's account] is one view about what kind of fact this is. On this view, personal identity just consists in physical and/or psychological continuity. ...

In the problem cases [such as the spectrum cases], things are different.⁴⁷

How can Parfit say that ordinary human lives *aren't* "problem cases" for personal identity? Imagine that it is a high degree of psychological continuity, not connectedness, that makes for a high degree of personal identity. In other words, two temporal stages of a human life are parts of the same person, to a high degree, if the human's brain is the same and there is a chain of temporal stages linking the two -- such that each link in the chain is psychologically connected to the next, to a high degree. On this version of Parfit's account, the absence of strong direct psychological connections between childhood and old age would be consistent with a high degree of psychological continuity (and hence personal identity) from childhood to old age.⁴⁸

Parfit *does* challenge the whole lifetime account of equality, even for normal humans -- he suggests that distributive justice may end up taking the form of a sublifetime prioritarianism⁴⁹ -- but his challenge does *not* involve the claim that personal identity is normally indeterminate. Rather, the challenge rests on a different aspect of Parfit's account of personal identity: its reductive or deflationary character. Once we see that personal identity, determinate or not, is reducible to facts about psychological and physical connectedness and/or continuity, personal identity will become less *significant*. "Personal identity is not what matters."⁵⁰ It will become less significant as a matter of individual rationality: whether some experience happens to me or some other person will be less significant in shaping my rational desires. And personal identity will become less significant as a matter of morality. "On the Reductionist View, we believe that the fact of personal identity over time is less deep, or involves less. We may therefore claim that this fact has less moral importance."⁵¹

For example, imagine two cases: one in which some person suffers a loss at one time, and another person experiences an equal gain at a later time; a second in which one person suffers a loss at one time, and that same person experiences an equal gain at a later

⁴⁷ Parfit, *The Unimportance of Identity*, *supra* note 46, at 303.

⁴⁸ See McMAHAN, *supra* note 46, at 39-41.

⁴⁹ See PARFIT, REASONS AND PERSONS, *supra* note 43, at 329-47; Parfit, *Comments*, *supra* note 43, at 837-43, 869-72.

⁵⁰ PARFIT, REASONS AND PERSONS, *supra* note 43, at 217.

⁵¹ *Id.* at 337-38.

time. Standardly, egalitarian theorists see these cases as sharply distinct -- because in the second case, it is the same person who suffers the effects, while in the first case it is two different persons. Parfit suggests that, once we deflate personal identity to its physical and psychological basis, our inclination to see the cases as qualitatively morally distinct will weaken.

There are a number of reactions to this suggestion. First, it is not clear that reducing some moral concept to its physical or psychological basis weakens the hold of that concept. If, for example, we reduce the concept of "life" to its physical basis, does that mean that morality must care less about killing?⁵² More generally, most moral theorists now accept that moral facts supervene on physical and psychological facts -- more crisply, on physical facts, since psychological facts themselves supervene on the physical. So there is a sense in which all moral concepts -- well-being, intention, harm, injury, choice, action -- are reducible to the physical and psychological.

This fact might mean that morality itself is less significant; but it is hard to see how it would favor one moral view over another. In other words, Parfit's deflationary view of personal identity, if simply one part of a general deflation of moral properties to psychological and physical properties, might be plausibly thought to lack implications *within* morality. In particular, it might be thought not to favor sublifetime over lifetime egalitarianism.

Second, assume that Parfit has succeeded in especially deflating the concept of a person. We continue to care about doing the morally right thing, and (if we are welfarists), to think of that in terms of human well-being -- even though we recognize that, ultimately, well-being is just a matter of the arrangement of molecules -- but are less impressed by the distinction between what occurs intrapersonally and interpersonally. The question, then, is why we would shift from lifetime to sublifetime egalitarianism. We might just shift to utilitarianism. The thought behind lifetime egalitarianism is that it may be unfair to burden one person for the sake of another -- an unfairness the utilitarian doesn't recognize -- but that effects within a normal life do not implicate this principle, because they are effects on the same person. Having excised or deflated the concept of "person," will we retain our concern for fair distribution but revise it to say that it may be unfair to burden one moment in a human life for the sake of another? Perhaps we might respond, instead, by saying that fair distribution -- across persons or temporal stages -- is not important at all.

Indeed, Parfit acknowledges as much. He suggests that his account of personal identity might plausibly change the scope of distributive principles (from lifetimes to sublifetimes), but acknowledges that it also might change their *weight*.

[Distributive] principles are often held to be founded on the separateness, or non-identity, of different persons. This fact is less deep on the Reductionist View, since personal identity is less deep. ... [I]t is [therefore] more plausible to give less weight to distributive principles. If we ...

⁵² See Mark Johnston, *Human Concerns without Superlative Selves*, in *PERSONAL IDENTITY*, *supra* note 46 at 260; Schultz, *supra* note 46, at 731.

come to believe that the unity of a life involves no more than the various relations between the experiences in this life, it becomes more plausible to be more concerned about the quality of experiences, and less concerned about whose experiences they are. This gives some support to the Utilitarian view.⁵³

In sum, there are a range of plausible reactions that the lifetime equity-regarding welfarist might have, once he comes to accept an account of personal identity, such as Parfit's, that allows for identity to be indeterminate in some cases, and that reduces it to physical and psychological facts. One reaction is to insist that similar deflations (and indeterminacies) can be created for all normative concepts, and to stick by his lifetime equity-regarding view. Another is to shift to utilitarianism. Yet a third is to shift to sublifetime equity-regarding welfarism. This third possibility *is* reasonable, but hardly compelled.

III. EX ANTE VERSUS EX POST APPROACHES TO SOCIAL CHOICE UNDER UNCERTAINTY

Part I grappled with the time-slice problem. This Part raises a cross cutting problem, namely uncertainty. Chris Sanchirico and I have engaged that problem at length in prior work.⁵⁴ We distinguish between "ex post" and "ex ante" approaches to welfarist social choice under uncertainty, and ultimately defend the "ex post" approach. I will not repeat our analysis but will quickly summarize the high points, specifically focusing on the case of lifetime welfarism. The "ex post" account of choice under uncertainty, together with lifetime welfarism, undergirds the array of practical prescriptions that are surveyed in Parts IV and V below.

Given certainty, each choice available to the decisionmaker can be represented as an N -entry vector of lifetime utilities: the particular set of lifetime utilities that would result were she to make that choice. Given uncertainty, each choice can be represented as a set of utility vectors, or matrix: one N -entry vector for each possible prior state of the world. The decisionmaker, we assume, can attach probabilities to these different possible states. With those probabilities in hand, plus her social welfare function, the decisionmaker might calculate a vector of expected utilities for each choice, and then apply the SWF to that. This is the "ex ante" approach. Alternatively, the decisionmaker might apply the SWF to each of the possible utility vectors, discount the SWF value of each vector by the probability of that vector, and sum. This is the "ex post" approach.

Crucially, the "ex ante" and "ex post" methods for applying an SWF under uncertainty can diverge. To see that possibility, and to get a better sense of how the approaches work, consider the following, extremely simple example, involving two persons (Jim and June), two possible choices A and B, two equiprobable states, and a prioritarian social welfare function, w , equaling the sum of the square root of utilities. The vectors of utilities in each state, for each choice, represent Jim's and June's lifetime utilities in that state.

⁵³PARFIT, REASONS AND PERSONS, *supra* note 43, at 346.

⁵⁴ See Adler & Sanchirico, *supra* note 7.

Choice A

	<u>State 1</u> p = .5	<u>State 2</u> p = .5	
Jim	4	9	Jim's expected utility = 6.5
June	0	4	June's expected utility = 2
	$w(4,0) = \sqrt{4+\sqrt{0}} = 2$	$w(9,4) = \sqrt{9 + \sqrt{4}} = 5$	

w applied ex ante to choice A: $\sqrt{6.5} + \sqrt{2} = 3.96$

w applied ex post to choice A: $2 \times .5 + 5 \times .5 = 3.5$

Choice B

	<u>State 1</u> p = .5	<u>State 2</u> p = .5	
Jim	3.5	3.5	Jim's expected utility = 3.5
June	3.5	3.5	June's expected utility = 3.5
	$w(3.5,3.5) = \sqrt{3.5+\sqrt{3.5}} = 3.74$	$w(3.5,3.5) = 3.74$	

w applied ex ante to choice B: $\sqrt{3.5} + \sqrt{3.5} = 3.74$

w applied ex post to choice B: $3.74 \times .5 + 3.74 \times .5 = 3.74$

Consider, first, the “ex ante” approach. Using that approach, we *first* calculate the vector of expected utilities for each choice, and *then* apply the SWF to that. Here, the vector of expected utilities associated with choice A is (6.5, 2) and the vector associated with choice B is (3.5, 3.5). So A is assigned a value of 3.96 ($\sqrt{6.5} + \sqrt{2}$), while B is assigned a value of 3.74 ($\sqrt{3.5} + \sqrt{3.5}$).

Consider, next, the “ex post” approach. Now, we *first* apply our SWF in each possible state, and *then* take probabilities into account, discounting each SWF value by its probability and adding. In the case of choice A, the SWF value in state 1 is $\sqrt{4 + \sqrt{0}} = 2$. The SWF value in state 2 is $\sqrt{9 + \sqrt{4}} = 5$. The expected w value yielded by A is $.5 \times 2 + .5 \times 5 = 3.5$. Turning to choice B, it emerges that the ex post approach assigns that a value of 3.74.

So what has happened is that the “ex ante” approach yields the judgment that choice A is better than choice B ($3.96 > 3.74$), while the “ex post” approach yields the judgment that choice B is better than choice A ($3.74 > 3.5$). The “ex ante” approach thinks of a given choice as a package of individual lotteries over outcomes, one for each individual in the population, and aggregates the values of these lotteries to arrive at an overall valuation of the choice. The “ex post” approach thinks of a given choice as a package of possible outcomes (each yielding a particular distribution of utilities), and values each outcome, then calculates an expected outcome value, to arrive at an overall valuation of the choice. As the simple example shows, these two approaches can diverge.

Does the example generalize? Indeed it does -- quite a bit! It turns out that the “ex ante” and “ex post” approaches can diverge for *any* SWF other than the utilitarian. For any possible choice situation, the utilitarian SWF will rank choices the same way regardless of whether we apply it “ex post” or “ex ante.” It makes no difference whether utilitarians conceptualize choices as vectors of expected utilities, or probability distributions over possible vectors of actual utilities. But, when we shift from utilitarianism to equity-regarding SWFs, the two approaches are genuinely different. For any equity-regarding SWF (comparativist, prioritarian, sufficientist), or indeed any nonutilitarian SWF, there will be some choice situations where the two approaches yield different recommendations.

In other words, lifetime equity-regarding welfarists must decide whether they care about the distribution of individuals’ expected lifetime well-being, or about the expected distribution of individuals’ lifetime well-being. These are just different approaches.

So which approach to choice under uncertainty should equity-regarding welfarists adopt? The ex ante approach certainly has its proponents. Some welfare economists who have written specifically about choice under uncertainty favor it. And, more generally, much scholarship about law and policy tends to implicitly adopt it -- to think about different policies as different package of expected utilities for the different individuals in the population.

Sanchirico and I contend, however, that the “ex post” approach is the better specification of welfarism -- and in this have some support in the existing literature. The core of the argument is the “sure thing” principle: what many take to be a basic premise of rationality, and in any event seems to be the hallmark of consequentialism. The sure-thing principle says that, if a choice involves a probability distribution over outcomes, and a revised choice is produced by swapping some of the outcomes for outcomes that the decisionmaker regards as equally good (without changing their probabilities), the decisionmaker should be indifferent between the original and revised choices. It is easy to see that the application of an equity-regarding SWF using the ex post approach always satisfies the sure thing principle, while the ex ante application of an equity-regarding SWF will sometimes violate that principle. A second important argument is that the “ex ante” approach is time-inconsistent.

In sum: A substantial case can be made that lifetime welfarism should be outfitted for conditions of uncertainty by adopting the “ex post” approach, which tells the decisionmaker to maximize the expectation of her SWF (be it comparativist, prioritarian, or sufficientist). Let us now see what lifetime welfarism, thus outfitted, implies for law and policy.

IV. IMPLICATIONS OF THE LIFETIME VIEW: MEASURING EQUALITY AND POVERTY

This Part discusses what lifetime equity-regarding welfarism, together with an ex post account of choice under uncertainty, implies for the *measurement* of equality and poverty. The next Part turns to policy choice.

A. *Equality Measurement*

A substantial body of empirical work in economics seeks to measure inequality. Much of this work focuses on the equality of incomes. A close related body of work consists of “incidence analysis”: quantifying whether taxes or other policy measures are “regressive” or “progressive,” depending on how their burdens are borne across the income distribution. More recently, health economists have begun to employ the equality metrics developed in the equality-of-income literature to study the distribution of health or longevity.

The choice between lifetime and sublifetime welfarism has direct and significant implications for all these projects.

1. Income Equality

Income-inequality studies -- which estimate the Gini coefficient, the coefficient of variation, or some other quantitative measure of the degree of inequality of some distribution of incomes -- are informational inputs for egalitarian welfarists.⁵⁵ Income or some transformation of income (such as the logarithm of income) is a measurable proxy for well-being. At least some inequality metrics correspond to particular SWFs -- a point stressed by Anthony Atkinson in seminal work.⁵⁶ To be sure, inequality measurement alone should not guide policy choice. A policy might produce a more equal distribution of well-being, but with a lower total amount. Any sensible SWF should be sensitive to both considerations, not merely the degree of inequality. But the well-established enterprise of measuring income inequality can readily be seen, by the welfarist, as furnishing important data for policymakers.

Most work in this area has focused on quantifying the inequality of annual incomes. But there are many economists who believe that the inequality of lifetime

⁵⁵ For overviews of this large scholarly field, see, e.g., HANDBOOK OF INCOME INEQUALITY MEASUREMENT (Jacques Silber ed., 1999); F.A. Cowell, *Measurement of Inequality*, in 1 HANDBOOK OF INCOME DISTRIBUTION 87 (2000); HILDE BOJER, DISTRIBUTIONAL JUSTICE: THEORY AND MEASUREMENT 65-133 (2003).

⁵⁶ See A.B. Atkinson, *On the Measurement of Inequality*, 2 J. ECON. THEORY 244 (1970).

income has greater normative relevance -- a position buttressed, of course, by the arguments of this Article -- and some have sought to measure it.⁵⁷

Take lifetime income (or some transformation of lifetime income) as a proxy for lifetime welfare. Why, in turn, isn't the distribution of annual income a good proxy for the distribution of lifetime income? There are (at least) two fundamental mechanisms that drive a wedge between the equality of annual versus lifetime income.⁵⁸ One is the age distribution of earnings (and thus income). Individuals earn more as they age and become more educated and skilled, up to a point; then, with retirement, earnings decline. To see how this alone can produce a wedge between the equality of lifetime and annual income, imagine that the age profile of earnings is completely fixed, and no one dies prematurely. Everyone in the population, then, has the same lifetime earnings; but because the population is composed of different age cohorts, earnings in any year are unequally distributed.

The second mechanism is income mobility within age cohorts. Imagine that everyone in the population is part of the same age cohort -- they were all born in the same year -- and that the cohort's mean earnings do not change at all with age. However, there is variation around the mean. In any given year, some in the age cohort earn above the mean, others below it. Further, there is earnings mobility -- at the extreme, whether someone is above the mean one year does not influence whether she is above it the following year. Then, it is easy to see, lifetime income could be much more equally distributed than annual income.

The precise relation between the equality of lifetime and annual income will depend, not just on income dynamics in the population, but also on the equality metric used. Still, economists have made some progress in formalizing the point that annual studies cannot generally be taken as good proxies for lifetime studies. Anthony Shorrocks has shown that, for a large class of inequality measures, inequality of lifetime income can never be greater than the average of inequality of annual income (and may be smaller, at the limit zero).⁵⁹

Is it, in fact, feasible to measure the distribution of lifetime income? How can empirical work of this kind be undertaken? Various approaches have been used. The most straightforward is to employ data from a longitudinal (or "panel") dataset, where a sample of the population is selected and then tracked over multiple years. For example, the Swedish Level of Living Survey interviewed a large sample of Swedes over many

⁵⁷ On the time dimension in income-inequality measurement, see generally A.B. Atkinson & F. Bourguignon, *Introduction: Income Distribution and Economics*, in HANDBOOK OF INCOME DISTRIBUTION, *supra* note 55, at 38-40; John Creedy, *Lifetime versus Annual Income Distribution*, in HANDBOOK OF INCOME INEQUALITY MEASUREMENT, *supra* note 55, at 513; Morton Paglin, *The Measurement and Trend of Inequality: A Basic Revision*, in 65 AM. ECON. REV. 598 (1975).

⁵⁸ See Daniel Millimet et al., *Bounding Lifetime Income Using a Cross Section of Data*, 49 REV. INCOME & WEALTH 205 (2003).

⁵⁹ See Anthony Shorrocks, *Income Inequality and Income Mobility*, 19 J. ECON. THEORY 376, 379-81 & n.7 (1978). The class of inequality measures includes those derived from the Atkinson family of SWFs; the square of the coefficient of variation; and the Theil coefficient.

years, and contains information about each individual's income for a period of 38 years. Anders Bjorklund calculated the Gini coefficient of the present value of individuals' total income over this 38-year period for different age cohorts, as well as the Gini coefficient of annual incomes in various years. He found that the long-term Ginis were 35 to 40 percent lower than the annual Ginis.⁶⁰ Similarly, Jeremy Arkes examined data from the Panel Study of Income Dynamics, which has conducted annual interviews with the same families since the 1960s. Arkes looked at the income of male heads of households, comparing the Gini coefficient of the distribution of 5-year earnings for different periods with the average of the Gini coefficient of annual earnings.⁶¹ In every comparison, 5-year earnings were more equally distributed than annual earnings.

A second approach to estimating the distribution of lifetime income involves simulation analysis.⁶² A third approach is to measure the distribution of some proxy for lifetime income. Many economists take an individual's annual consumption -- meaning her annual expenditures on goods and services -- to be a better proxy for lifetime income than annual income. The basic insight is that rational individuals with access to intertemporal financial markets will "smooth" consumption over a lifetime, saving some income in high-income periods for expenditure in later periods, and borrowing or spending from savings to finance consumption in low-income periods. A substantial body of empirical work has therefore examined the inequality of annual *consumption*, as opposed to annual income, as a way to estimate the inequality of lifetime income.⁶³

Does the choice between annual and lifetime income as the temporal unit for studying income equality turn out to be a significant choice? Do empirical studies that rely upon lifetime income as the temporal unit regularly reach substantially different conclusions than studies that rely upon annual income? The answer, pretty clearly, is yes. To begin, the choice clearly affects estimates of the degree of inequality. As already mentioned, both Arkes and Bjorklund, who used longitudinal studies and the Gini coefficient as the measure of inequality, found that lifetime or long-term income inequality was lower than annual income inequality. More generally, virtually every work that has applied some inequality metric to lifetime or long-term versus annual income has reached the same result -- not surprisingly, given Shorrocks' proof. Lifetime inequality is always lower. We hardly live in an egalitarian world; but the world is revealed to be less marred by inequality, once one takes the lifetime perspective, than it appears to be from the annual perspective.

⁶⁰ Anders Bjorklund, *A Comparison Between Actual Distributions of Annual and Lifetime Income: Sweden 1951-89*, 39 REV. INCOME & WEALTH 377 (1993).

⁶¹ Jeremy Arkes, *Trends in Long-Run Versus Cross-Section Earnings Inequality in the 1970s and 1980s*, 44 REV. INCOME & WEALTH 199 (1998).

⁶² See, e.g., ANN HARDING, *LIFETIME INCOME DISTRIBUTION AND REDISTRIBUTION* (1993); Christopher J. Flinn, *Labor Market Structure and Inequality: A Comparison of Italy and the United States*, 69 REV. ECON. STUD. 611 (2002).

⁶³ See DANIEL T. SLESNICK, *CONSUMPTION AND SOCIAL WELFARE* (2001); Dirk Krueger & Fabrizio Perri, *Does Income Inequality Lead to Consumption Inequality? Evidence and Theory*, 73 REV. ECON. STUD. 163 (2006); *id.* at 163-64 (citing sources).

The choice of annual versus lifetime income may lead to different conclusions, not just about the amount of inequality in some population, but also about the change in equality over time, or about the causes of inequality, or about the comparative degree of inequality in different populations. Daniel Slesnick's scholarship on consumption provides an excellent illustration of the possible link between the choice of unit and conclusions about equality dynamics. Many believe that the U.S. has become a less egalitarian society in recent years, and this view is borne out by annual-income studies. For example, the Gini coefficient of annual income decreases from the 1950s through the late 1960s, but then changes direction and increases pretty steadily from 1970 through the late 1990s. (This is the so-called "U-turn in income inequality"⁶⁴). Slesnick challenges this pessimistic picture with evidence from the Consumer Expenditure Survey, a cross-sectional study of household consumption. Slesnick's conclusion is as follows:

[T]he widely reported U-turn in inequality in the United States is an artifact of the inappropriate use of family income as a measure of welfare. When well-being is defined to be a function of per equivalent consumption, inequality either decreased over the sample period or remained essentially unchanged depending on the choice of equivalence scales. This conclusion arises because the distribution of consumption differs sharply from the distribution of income, which, in turn, is due to the distribution of savings and dissavings across the population. . . .

[This conclusion] has important policy implications. The conventional wisdom is that rising earnings inequality has been offset by income transfers until 1973. The increasing level of income inequality since 1973 has led some analysts to the conclusion that the government must redouble its efforts to bring about greater equality in the distribution of well-being. My results suggests that there is no U-turn to explain and no increase in inequality to overcome.⁶⁵

Much empirical work on inequality tries to illuminate the causes of inequality in some population by decomposing overall inequality into inequality between certain groups plus inequality within the groups. Here, too, the choice of temporal perspective can be significant. Dirk Krueger and Fabrizio Perri, like Slesnick, examine trends in the inequality of annual income and consumption in the U.S. They reach conclusions similar to Slesnick's, namely that "despite the surge in income inequality in the U.S., consumption inequality has increased only moderately," and also find the following: "[I]ncome inequality has increased substantially, both between and within groups of households with the same characteristics (such as education, sex, and race), but even though between-group consumption inequality has tracked between-group income inequality quite closely, within group *consumption* inequality has increased much less than within-group income inequality."⁶⁶ Finally, for a nice illustration of the effect of temporal perspective on comparisons of inequality between populations, consider Christopher Flinn's use of a simulation model to compare earnings inequality in the U.S. with Italy, an economy where workers tend to change employers less often and have less mobile earnings. Flinn finds that the inequality of hourly wages is three times greater in the U.S. than in Italy, but that the inequality of lifetime earnings in the two countries is the same.⁶⁷

⁶⁴ See SLESNICK, *supra* note 63, at 2-3.

⁶⁵ *Id.* at 154.

⁶⁶ Krueger & Perri, *supra* note 63, at 164.

⁶⁷ See Flinn, *supra* note 62.

2. Incidence Studies

Economists working in the area of taxation often employ a distinctive approach to characterizing inequality: incidence analysis. Incidence analysis predicts the distribution of individual “tax burdens” for a given tax; calculates the individual tax burdens as a proportion of individual income; and characterizes the tax as “regressive,” “proportional,” or “progressive” if this fractional burden decreases, stays flat, or increases with individual income.

Classically (as in the famous early studies by Joseph Pechman),⁶⁸ the income measure employed in incidence analysis is *annual* income.⁶⁹ The conventional wisdom that the personal income tax and the corporate income tax are progressive, but that their progressive impact is largely offset by the regressive impact of payroll taxes, sales taxes, and excise taxes, derives in substantial part from Pechman’s studies.⁷⁰ Subsequent incidence studies sometimes challenged these conclusions, but they almost invariably followed Pechman in using the tax burden as a fraction of annual, not lifetime, income as the basis for characterizing a tax’s regressivity or progressivity.

In a 1989 article, James Poterba challenged the annual-income perspective on tax incidence. In line with scholarship in the inequality measurement literature, Poterba took the view that, “[i]f households base their spending plans on their expected lifetime income, then consumption provides a more accurate measure of lifetime resources than does annual income.”⁷¹ He examined the Consumer Expenditure Survey, which provides data about both household income and household consumption, and assumed (as had Pechman) that the burden of an excise tax is borne by the consumer. Poterba then considered the regressivity of excise taxes on gasoline, alcohol, and tobacco. The typical expenditure on each of these items as a percentage of annual income declines steeply. From the annual perspective, then, such taxes appear quite regressive. But the expenditure on each of these items as a percentage of total annual expenditure declines much less steeply. From the lifetime-income perspective (with annual consumption taken as a proxy for lifetime income), the gasoline, alcohol and tobacco excise taxes appear substantially less regressive.

Since Poterba’s work, some incidence scholarship has taken a lifetime perspective by actually estimating lifetime income, rather than taking consumption as a proxy. For example, Erik Caspersen and Gilbert Metcalf used the Panel Study of Income Dynamics,

⁶⁸ JOSEPH A. PECHMAN & BENJAMIN A. OKNER, WHO BEARS THE TAX BURDEN? (1974); JOSEPH A. PECHMAN, WHO PAID THE TAXES, 1966-85? (1985).

⁶⁹ On this point, and the choice between annual and lifetime approaches to incidence analysis, see DON FULLERTON & DIANE LIM ROGERS, WHO BEARS THE LIFETIME TAX BURDEN 1-21 (1993); Thomas A. Barthold, *How Should We Measure Distribution?*, 46 NAT’L TAX JOURNAL (); Don Fullerton & Diane Lim Rogers, *Lifetime versus Annual Perspectives on Tax Incidence*, 44 NAT’L TAX JOURNAL (); James M. Poterba, *Lifetime Incidence and the Distributional Burden of Excise Taxes*, 79 AM. ECON. REV. 325, 325 (1989).

⁷⁰ See FULLERTON & ROGERS, *supra* note 69, at 12.

⁷¹ Poterba, *supra* note 69, at 325.

a longitudinal study, to predict lifetime income as a function of various individual characteristics. That function was then applied to the Consumer Expenditure Survey, enabling Caspersen and Metcalf to calculate the burden of the value added tax (VAT) both as a fraction of annual income and as a fraction of estimated lifetime income.⁷² The most ambitious work in this vein has been undertaken by Don Fullerton and Diane Lim Rogers in their 1993 book, *Who Bears the Lifetime Tax Burden?* They use the Panel Study of Income Dynamics to estimate wages as a function of age and other characteristics, and incorporate these lifetime wage profiles in a general equilibrium model.

[The model] encompasses all major U.S. taxes, many industries, ...and consumers. It is ... a life-cycle model in which each individual receives a particular inheritance, a set of tax rules, a wage profile, and a transfer profile. Each then plans an entire lifetime of labor supply, savings, goods demands, and bequests. We also look at each industry's use of labor, capital, and intermediate units. We can then simulate the effects of a tax change on each economic decision through time. . .

[W]e evaluate the effects of each U.S. tax by comparing its estimated burdens with those of a proportional tax. ... In our lifetime framework, a progressive tax is one in which the lifetime tax burden as a fraction of lifetime income rises as lifetime income rises, and a regressive tax is one in which the lifetime tax burden as a fraction of lifetime income falls as lifetime income [rises].⁷³

As with the companion literature on income inequality, discussed above, it is the clear that the choice of temporal unit makes a difference to incidence analysis. Every study which has compared lifetime and annual perspectives concludes that at least some taxes look different from the two perspectives, sometimes quite substantially. For example, Poterba, as stated, finds that excise taxes on gasoline, alcohol and tobacco are less regressive in a lifetime perspective. Fullerton and Rogers find that property taxes are progressive in an annual perspective but have a U-shaped burden in a lifetime perspective; that the corporate income tax shifts from being progressive to being regressive; and that the progressivity of the income tax decreases. Caspersen and Metcalf find that the VAT is substantially less regressive in the lifetime perspective. Davies and co-authors find that the personal income tax is less progressive.⁷⁴

3. Health Inequality

The vast majority of scholarly work that quantifies the degree of inequality in the population distribution of some item, with an inequality metric such as the Gini coefficient or the coefficient of variation, concerns the distribution of income (annual or lifetime). But the approach is applicable to many other items, and the welfarist would surely support its export to other areas. Income is hardly the sole constituent or resource for well-being, nor the sole one which is measurable or unequally distributed.

⁷² See Erik Caspersen & Gilbert Metcalf, *Is A Value Added Tax Regressive? Annual Versus Lifetime Incidence Measures*, 47 NAT'L TAX JOURNAL 731 (1994).

⁷³ FULLERTON & ROGERS, *supra* note 69, at 4-5.

⁷⁴ See James Davies et al., *Some Calculations of Lifetime Tax Incidence*, 74 AM. ECON. REV. 633 (1984).

C.J.L. Murray, a prominent health economist, has indeed argued that inequality metrics should be employed to quantify the population distribution of *health*.⁷⁵ And there is a small, but growing, body of work in this vein.

What does the lifetime perspective imply for measurement of population health inequality? To begin, it implies that the focus of empirical work should be the distribution of lifetime rather than sublifetime health status -- be it momentary health status or health status during some longer fraction of a lifetime.

An analysis by Angus Deaton and Christina Paxon, in the course of a broader study on the correlations between age, income, and health, illustrates the momentary approach. The Panel Study on Income Dynamics contains, inter alia, information on "self-reported health status." (SRHS) Respondents are asked to rate their health on a scale from 1 to 5. SRHS values are a kind of rough proxy for momentary QALYs.⁷⁶ Deaton and Paxon calculate the variance in SRHS within age cohorts.⁷⁷ Although they do not do the calculations, it would be straightforward to use the Panel Study on Income Dynamics data to estimate the population distribution of SRHS and to measure the inequality of that distribution.

From the lifetime perspective, however, such an approach to health-inequality measurement is problematic. For example, imagine that lifespans are increasing in a given country but that older individuals have more variance in SRHS, momentary QALYs, or some other such measure of momentary health. Then the distribution of lifetime QALYs or some other measure of lifetime health may be becoming more equal, even though the distribution of momentary health is not. Or, imagine that a fatal disease increases mortality in a country, and that those who survive are healthier and have less variance in momentary health status than the pre-disease population. Then, if we use a momentary inequality metric, it appears that the disease has reduced health inequality.

⁷⁵ See C.J.L. Murray et al., *Health Inequalities and Social Group Differences: What Should We Measure?*, 77 BULL. WORLD HEALTH ORG. 537 (1999). Most empirical work on health inequalities focuses, instead, on health differences between socioeconomic, racial, or gender groups. The population approach advocated by Murray would seem to be more directly connected to the SWF framework, and thus is my focus here. But measures of social-group differences might also be informative for welfarist social planners - a point I discuss below with reference to anti-discrimination and disparate impact tests. See *infra* Part V.C.1. On the social-group approach to health inequality measurement, see Paula Braveman, *Health Disparities and Health Equity: Concepts and Measurement*, 27 ANN. REV. PUB. HEALTH 167 (2006); Murray, *supra*; Johan P. Mackenbach & Anton E. Kunst, *Measuring the Magnitude of Socio-Economic Inequalities in Health: An Overview of Available Measures Illustrated with Two Examples from Europe*, 44 SOC. SCI. MED. 757 (1997); Adam Wagstaff & Eddy van Doorslaer, *Overall versus Socioeconomic Health Inequality: A Measurement Framework and Two Empirical Illustrations*, 13 HEALTH ECON. 297 (2004).

⁷⁶ QALYs are a widely used metric in health research and, increasingly, government decisionmaking. The value of health states is measured on a 0-1 scale, using various survey techniques -- with 1 representing perfect health and 0 a health state no better than death -- and the value of a temporally extended health profile is the sum of the value of each component state multiplied by its duration. See generally Matthew D. Adler, *QALYs and Policy Evaluation: A New Perspective*, 6 YALE J. HEALTH POLICY, L. & ETHICS 1 (2006).

⁷⁷ See Angus S. Deaton & Christina H. Paxon, *Aging and Inequality in Income and Health*, 88 AM. ECON. REV. 248, 250 (1998).

What that calculation ignores is the inequality in lifetime health between those who die from the disease and those who survive it.

Similar criticisms would apply to inequality measures of the distribution of nonmomentary but sublifetime (rather than lifetime) health -- for example, the distribution of annual QALYs in a population over the last year.

There are, in fact, health-inequality researchers who have adopted the lifetime perspective. A number of scholars have examined the distribution of longevity. For example, Julian Grand looked at age-at-death statistics from different countries in (roughly) the same year. These statistics show the age at death of everyone who has died in the country in that year. He calculated the Gini coefficient and other inequality measures of the age-at-death distribution, for the different countries.⁷⁸ Ryan D. Edwards and Shripad Tuljapurkar compared variance of age-at-death in different countries and at different times.⁷⁹ Similar work has been undertaken by Jacques Silber and by Douglas Hicks.⁸⁰

All this scholarship employs a crude measure of lifetime health - namely, lifespan. Lifetime QALYs would be a better measure of lifetime health. Indeed, in an important paper, Murray and co-authors, E.E. Gakidou and J. Frenk, have proposed using the distribution of *health expectancy* as the basis for studying health inequality; and one of the WHO's annual *Health Reports* adopts this approach.⁸¹ An individual's health expectancy is his expected lifetime health, given his risk of dying at each age, his age-conditional risk of different health states, and the value of those states on a zero-one scale. In other words, an individual's health expectancy is just his expected lifetime QALYs.

The Gakidou, Murray and Frenk proposal therefore combines a lifetime perspective with the "ex ante" rather than "ex post" approach to social choice under uncertainty. Part III has criticized the "ex ante" view, and the criticisms are relevant, not merely to the Gakidou/Murray/Frenk proposal, but to the enterprise of inequality-measurement more generally -- as I shall now explain.

4. Inequality of Individual Expectations Versus Expected Inequality of Attainments

⁷⁸ See Julian Le Grand, *An International Comparison of Distributions of Ages-At-Death*, in HEALTH INEQUALITIES IN EUROPEAN COUNTRIES 75 (John Fox ed., 1989); Julian Le Grand, *Inequalities in Health: Some International Comparisons*, 31 EUR. ECON. REV. 182 (1987). See also Raymond Illsley & Julian Le Grand, *The Measurement of Inequality in Health*, in HEALTH AND ECON. 12 (Alan Williams ed., 1987).

⁷⁹ Ryan D. Edwards & Shripad Tuljapurkar, *Inequality in Life Spans and a New Perspective on Mortality Convergence Across Industrialized Countries*, 31 POP. & DEVEL. REV. 645 (2005)

⁸⁰ See Douglas A. Hicks, *The Inequality-Adjusted Human Development Index: A Constructive Proposal*, 25 WORLD DEVELOPMENT 1283 (1997); Jacques Silber, *Health and Inequality: Some Applications of Uncertainty Theory*, 16 SOC. SCI. MED. 1663 (1982).

⁸¹ See E.E. Gakidou et al., *Defining and Measuring Health Inequality: An Approach Based on the Distribution of Health Expectancy*, 78 BULL. WORLD HEALTH ORG. 42 (2000); THE WORLD HEALTH ORGANIZATION, THE WORLD HEALTH REPORT 2000, at 146 (2000).

Given some inequality measure *M*, and some proxy for lifetime well-being *X*, plus uncertainty about individual attainments with respect to *X*, we can measure the *inequality of individual expectations* or rather *the expected inequality of attainments*. We can determine the population distribution of individual expectations with respect to *X*, and apply *M* to that; or we can determine the expected value of *M*, by applying *M* to each possible population distribution of individual attainments, discounting by the probability of that distribution, and summing. The inequality-of-expectations approach maps onto an “ex ante” account of social choice under uncertainty; the expected-inequality-of-attainments approach maps onto an “ex post” account of social choice under uncertainty.

Gakidou, Murray and Frenk argue for the inequality-of-expectations approach, with QALYs as the proxy. They propose to measure the inequality of expected lifetime QALYs, not the expected inequality of lifetime QALYs. Why?

One observation they make is that individuals with the same health expectancies “may have very different healthy lifespans because of chance.”⁸² Imagine, at the extreme, a birth cohort in which each baby has the same risk of dying at every age, and the same age-conditional chance of incurring any health state, regardless of the babies’ social classes, ethnicities, genders, or other observable characteristics. In this sort of case, expected lifetime QALYs are perfectly equal within the cohort at the time of birth, but the lifetime QALYs actually attained by cohort members will not be equal, “because of chance.” Some members of the cohort will die young; some will not but will suffer burdensome diseases; some will live long, healthy lives. But it should hardly be an ambition of egalitarian welfarists to ignore differences in well-being that are simply a matter of chance.

Gakidou, Murray and Frenk also note that “equality of healthy lifespan could only be realized if risks of incidence and remission of non-fatal health outcomes, and risks of mortality, were either zero or one for the entire population.”⁸³ Perfect equality of realized lifetime QALYs may be impossible -- but so what? We can still measure the expected degree of inequality in lifetime QALYs; and policymakers can still take steps to reduce that. (Parenthetically, perfect equality of expected lifetime QALYs seems just as unrealistic as perfect equality of realized lifetime QALYs.)

A further worry about focusing on the distribution of realized, rather than expected, lifetime QALYs is this. How can we employ that approach to measure *future* health inequality? As described above, scholars such as Le Grand, Edwards and Tuljapurkar, Silber, and Hicks have examined age-at-death statistics to determine the distribution of realized lifespan in a past population; but, given uncertainty about the future, what would it mean to measure the distribution of realized lifespan, or realized lifetime QALYs, in a future population? The answer, here, is that the expected-inequality-of-attainments framework is robust to uncertainty, for example, uncertainty about what

⁸² Gakidou et al., *supra* note 81, at 43.

⁸³ Gakidou et al., *supra* note 81, at 44.

the health of future cohorts will be. Indeed, integrating a concern for equality with uncertainty is the very point of *both* this framework as well as the competing equality-of-expectations approach. Neither framework assumes certainty, or is merely retrospective. We can use current information to estimate the probability of different possible population distributions of longevity or lifetime QALYs; measure the degree of inequality in each such distribution, for example by using the Gini coefficient; and then calculate the expected degree of inequality across possible distributions, for example, the expected Gini coefficient of realized longevity or lifetime QALYs.

The choice between inequality-of-individual-expectations and expected-inequality-of-attainments has been most clearly highlighted by Gakidou, Murray and Frenk, who are health researchers, but scholars working on income inequality must confront the same question. In estimating (say) the Gini coefficient of the inequality of lifetime income in some population, should researchers estimate the Gini coefficient of expected lifetime income, or the expected Gini coefficient of lifetime income? Note that the question arises most obviously with respect to the prediction of future income equality; but it also lurks under the surface of retrospective work, given measurement error. Similarly, lifetime tax-incidence analysis might focus on the expected lifetime tax burden as a fraction of expected lifetime income; or on the probability distribution of possible incidence scenarios, in each of which individuals at different lifetime income levels incur some tax burden.

B. *Poverty Measurement*

Poverty measurement, like the companion field of inequality measurement, is a large area of research within empirical economics.⁸⁴ Development economists, in particular, have devoted considerable effort to quantifying poverty, but work in the field also extends to the United States and other developed countries. Much existing research focuses on income poverty: measuring the extent to which a population is impoverished in virtue of its low incomes. The standard methodology starts by defining a poverty line: either an absolute poverty line, often meaning the amount of income required to purchase sufficient food to avoid malnutrition; or a relative poverty line, meaning some low percentile of the population's income distribution (for example, the 20th percentile).⁸⁵ The poverty index is then some function of the number of individuals below the poverty line, their incomes, and the size of the whole population. Two simple indices are the headcount ratio (the fraction of the population which is impoverished); and the poverty gap (the average shortfall between the income of each impoverished person and the poverty line). More sophisticated indices are also available. Measures within the class of so-called Foster, Greer, and Thorbecke indices, which take the form

⁸⁴ See generally BOJER, *supra* note 55, at 118-22; POVERTY AND INEQUALITY (David B. Grusky & Ravi Kanbur eds., 2006); Bhaskar Dutta, *Inequality, Poverty and Welfare*, in 1 HANDBOOK OF SOCIAL CHOICE AND WELFARE 597, 619-27 (2002); Caterina Ruggeri Laderchi et al., *Does It Matter that We Do Not Agree on the Definition of Poverty? A Comparison of Four Approaches*, 31 OXFORD DEV. STUD. 243 (2003).

⁸⁵ See sources cited *supra* note 84; Martin Ravallion, *Poverty Lines in Theory and Practice* (The World Bank, LSMS Working Paper No. 133, 1998).

$1/N \sum [(Z-X_i)/Z]^\alpha$, $\alpha \geq 0$, with X_i the individual's income and Z the poverty line, are most commonly used.⁸⁶

What are the implications, here, of lifetime welfarism? As suggested earlier, the lifetime view *might* be seen to undermine the variant of equity-regarding welfarism -- sufficientism -- which most directly supports the practice of poverty measurement.⁸⁷ Lifetime sufficientism would be defined around a single threshold level of lifetime well-being that would make a categorical difference with respect to redistributive claims. Identifying that lifetime level is not an easy matter. Consider the criterion of being well-nourished -- the typical basis in practice for setting absolute poverty lines. Someone who is badly nourished at all moments throughout a lifetime is impoverished with respect to lifetime well-being -- but this sets the poverty threshold extremely low, and makes sufficientism collapse to utilitarianism for the vast majority of the population. At the other extreme, the fact that an individual is hungry once in her life can hardly mean that she falls in the class of lifetime "poor" individuals who have stronger redistributive claims on the rest of the population. And any cutoff percentage Q which is greater than 0 and less than 1, which purports to define the percentage of moments that an individual must be malnourished to be lifetime malnourished, may well seem arbitrary -- arbitrary in a way that the basic notion of being malnourished is not.

In short, the lifetime view may radically change the theoretical basis for the field of poverty measurement. This is not to say that the field would disappear. Poverty measures can provide important information for utilitarians, prioritarists, and comparativists, not just for sufficientists. For example, an income-poverty line might track qualitative changes in the welfare productivity of money. The marginal utility of income might decrease dramatically at the point at which income suffices to purchase basic nourishment. But abandoning sufficientism would mean that poverty measurement would need to locate some other philosophical justification within welfarism, and in particular that the poverty line could not be set with reference to some morally basic cut-off level of well-being. Giving up sufficientism means giving up that notion.

Still, it is *possible* to marry sufficientism with the lifetime perspective on well-being. As mentioned earlier, Roger Crisp, a leading proponent of sufficientism, entertains just this possibility.⁸⁸ Poverty measurement rooted in lifetime sufficientism would naturally be focused on the distribution of lifetime income in some population (or the distribution of some proxy for lifetime income) and a poverty line defined in terms of lifetime income (or some proxy).

This is not a utopian idea. Poverty researchers, like inequality researchers, generally employ short-term measures of income; but, as in the inequality field, there actually is a reasonably substantial body of scholarship that attempts to take a longer-term view. The rubric for this work is the study of "chronic" rather than "transient"

⁸⁶ See Udo Ebert & Patrick Moyes, *A Simple Axiomatization of the Foster, Greer and Thorbecke Poverty Orderings*, 4 J. Pub. Econ. Theory 455, 456 (2002).

⁸⁷ See *supra* Part II.B.3.

⁸⁸ See *id.*

poverty.⁸⁹ Much of this research employs longitudinal data, with chronic poverty defined either in terms of multi-period income (or consumption), or in terms of the frequency with the individual is below the periodic income (or consumption) poverty line. For example, Jyotsana Jalan and Martin Ravallion examined panel data over six years from China, which provided information about a household's consumption and other characteristics.⁹⁰ They defined a household's *chronic poverty* to be the squared poverty gap between its *mean* consumption over the six years and the annual poverty line. Since annual consumption was normalized by the annual poverty line, this is the same as taking the squared poverty gap between a household's normalized six-year consumption and a six-year poverty line equaling six times the annual poverty line. Jalan and Ravallion also calculated "transient" poverty, defined as the difference between the household's average squared poverty gap and its chronic poverty. Raghav Gaiha and Anil Deolalikar used panel data from India over 9 years, and determined whether the household was below the annual income poverty line for zero, one, or more years. Some households, the most "persistently" poor, had annual incomes below that line for all nine years.⁹¹

A different approach, again paralleling inequality-measurement techniques, is to employ a proxy for long-term poverty. For example, the depth of poverty in one period might be taken as a crude measure of whether the household's long-term income is above or below some long-term line.⁹² Or annual consumption might be used. Slesnick examined not only consumption inequality in the U.S., but consumption poverty as well⁹³; and Krishna Pendakur examined consumption poverty in Canada.⁹⁴

If individuals' annual incomes were completely immobile, annual income would be no worse an indicator of lifetime well-being than lifetime or long-term income, and poverty studies could use a long- or short-term income time slice interchangeably. The immobility premise is, of course, false, and studies of chronic poverty therefore reach distinctive conclusions: about the extent of poverty, its causes and correlates, and policy responses. For example, Jalan and Ravallion find that the amount of poverty would be reduced by half if transient poverty were excluded; that one of the four Chinese provinces studied had a much higher proportion of transient poverty than the others; and that, when a household's chronic and transient poverty were regressed on household characteristics, the determinants of the two were quite different. Slesnick, in the U.S. context, finds that the percentage of the population suffering consumption poverty is substantially lower than the percentage that are counted as impoverished using traditional annual-income

⁸⁹ For overviews, see David Hulme & Andrew Shepherd, *Conceptualizing Chronic Poverty*, 31 *WORLD DEV.* 403 (2003); Andrew McKay & David Lawson, *Assessing the Extent and Nature of Chronic Poverty in Low Income Countries: Issues and Evidence*, 31 *WORLD DEV.* 425 (2003).

⁹⁰ See Jyotsana Jalan & Martin Ravallion, *Is Transient Poverty Different? Evidence for Rural China*, 36 *J. DEVELOPMENT STUD.* 82 (2000).

⁹¹ Raghav Gaiha & Anil B. Deolalikar, *Persistent, Expected and Innate Poverty: Estimates for Semi-Arid Rural South India, 1975-1984*, 17 *CAMBRIDGE J. ECON.* 409 (1993).

⁹² See McKay & Lawson, *supra* note 89, at 434.

⁹³ See SLESNICK, *supra* note 63, at 156-89.

⁹⁴ See Krishna Pendakur, *Consumption Poverty in Canada, 1969 to 1998*, 27 *CANADIAN PUB. POLICY* 125 (2001).

measures, and that consumption poverty (by contrast with the traditional measure) has decreased over time.

More generally, the picture from multiple studies suggests that:

[C]hronic poverty ... is strongly associated with disadvantages that are difficult to reverse quickly such as a lack of assets (physical and human ...), being trapped in low-productivity activities ..., a high dependency ratio[], and location in remote or otherwise disadvantaged areas. Transient poverty by contrast is more typically associated with cases where households have very little ability to insure themselves against fluctuations due to either external factors such as prices, climate or job availability, or household level shocks such as serious illness or death.⁹⁵

As for policy implications, the literature on chronic poverty suggests – very crudely – that government policies such as unemployment insurance, good health care coverage, worker retraining, or disaster relief, will be applauded by the sublifetime sufficientist, but will be viewed more skeptically by the lifetime sufficientist. These policies may help many whose short-term well-being is episodically low, without doing much for the badly educated, socially isolated, and low-asset portion of the population whose long-term well-being is low.

I have, until now, bracketed issues of uncertainty -- but those must inevitably be confronted by poverty researchers, as in the companion field of inequality measurement. Given some proxy for lifetime well-being X (such as lifetime income), some poverty line Z , and some poverty measure M , we might calculate M as a function of Z and individual expectations with respect to X ; or, alternatively, we might calculate the expectation of M , that is, the expected degree to which individuals' actual attainments with respect to X fall short of Z . The first approach corresponds to the “ex ante” account of social choice under uncertainty; the second, to the “ex post” approach. A number of researchers have taken the first approach. For example, Gaiha and Deolalikar inter alia employ a regression model to determine the household's expected income, and calculate poverty as a function of household expected income.⁹⁶ But it would be better -- if the “ex post” account is correct -- to try to estimate the probability of different possible distributions of long-term income in the population of interest, and then to calculate the expected degree of poverty as a probability-weighted average of the poverty measure applied to each of these possible income distributions.⁹⁷

⁹⁵ McKay & Lawson, *supra* note 89, at 436-37.

⁹⁶ See Gaiha & Delolalikar, *supra* note 91.

⁹⁷ There is a general point to be noted about the deviation between “ex ante” and “ex post” approaches to implementing an SWF or to measuring inequality or poverty. Let $V_1 \dots V_n$ be random variables representing individual attainments (either utilities or income levels). F is a function of these (such as a social welfare function, an inequality measure, or a poverty measure). Then, if E is the expectation operator, $E(F(V_1 \dots V_n))$ is only guaranteed to be equal to $F(E(V_1), E(V_2) \dots, E(V_n))$ if F is a linear function of the V_i . See Adler and Sanchirico, *supra* note 7. Since this is not true for equity-regarding SWFs, inequality measures, or poverty measures, the “ex ante” and “ex post” approaches will deviate in all three cases.

V. IMPLICATIONS OF THE LIFETIME VIEW: CHOOSING POLICIES

Here, I discuss the implications of lifetime equity-regarding welfarism for policy choice. Rather than discussing the implications in particular policy domains (education, environmental law, health care, natural hazards, unemployment insurance, etc.), which would consume many articles, I approach the issue in a more general way. First, I look at the implications for redistributive taxation. The credentials of the tax system, as a redistributive mechanism, are unimpeachable. Then, I examine the implications for the well-known view, pressed by Louis Kaplow and Steven Shavell and others, that the tax system is the *sole* legal mechanism appropriately employed for redistribution, regardless of the SWF. Finally (because there are plausible grounds to dispute this view), I examine the implications for the policy-analytic tools we might use to evaluate the equity effects of non-tax policies: in particular, “suspect class” analysis and the development of integrated policy-analytic tools that will be sensitive to both overall well-being and equity.

A. *Redistributive Taxes*

Debates about the appropriate tax structure, for example the long-running debate about consumption versus income taxation, are commonplace both in various scholarly fields and in the political process. Equity-regarding welfarism suggests that the first-best technique for choosing among tax regimes is by specifying a SWF, or range of SWFs, and comparing possible regimes in light of this function(s). The traditional criteria for evaluating tax policy, such as “horizontal equity” or “vertical equity” or “ability to pay,” are second-best, proxy criteria, which are justifiably employed in some contexts given the computational demands of the first-best technique.⁹⁸ The choice between sublifetime and lifetime perspectives has direct relevance, both to the proxy criteria and to first-best policy evaluation by application of a SWF.

Start with the proxy criteria. Horizontal equity means equal tax burdens for similarly situated individuals. Vertical equity means different tax burdens for differently situated individuals. A standard way to define an individual’s “situation,” for purposes of these criteria, is in terms of her income. This makes sense, for welfarists, given the role of income as an important determinant of well-being. Horizontal equity then requires that individuals with equal incomes bear equal tax burdens, and vertical equity that individuals with greater incomes pay more tax (at least in absolute and perhaps fractional terms) -- since the utilitarian and, a fortiori, the equity-regarding welfarist will favor some transfers of income (bracketing disincentive effects) from rich to poor. Adopting a lifetime view of well-being, in turn, motivates a definition of horizontal and vertical equity in terms of lifetime income rather than sublifetime well-being. As Fullerton and Rogers explain, “[T]he lifetime perspective provides a useful yardstick to help evaluate any ... tax system. For ‘horizontal equity,’ two individuals with similar lifetime incomes

⁹⁸ For critiques of these traditional criteria as first-best criteria, see LOUIS KAPLOW, *THE THEORY OF TAXATION AND PUBLIC ECONOMICS* ch. 15 (forthcoming 2007); Thomas D. Griffith, *Should “Tax Norms” Be Abandoned: Rethinking Tax Policy Analysis and the Taxation of Personal Injury Recoveries*, 1993 WIS. L. REV. 1115.

should pay similar total lifetime taxes. In addition, for ‘vertical equity,’ higher lifetime incomes should be associated with higher lifetime tax burdens.⁹⁹”

The temporal framing of horizontal and vertical equity norms clearly affects whether particular regimes are seen as more or less “equitable” in light of these norms. Consider the most obvious example: William Vickrey’s well-known proposal for cumulative income taxation.¹⁰⁰

A cumulative tax system consists of a final tax schedule that specifies total tax payments and of procedures for collecting tax payments along the way . . . Vickrey proposed that the final tax schedule be derived from the existing annual tax schedule by calculating for any total income the tax that would have been paid under the annual tax schedule if the total income had been earned in equal amounts in each of the years of the averaging period. Total income was to be calculated as the sum of annual taxable income over the averaging period plus an imputation of accumulated interest on tax payments made through the averaging period.¹⁰¹

Making the averaging period an entire lifetime is what horizontal equity, framed in terms of equal taxation for individuals with equal lifetime incomes, would require.¹⁰² Under the current income tax system in the U.S., which uses an annual period for measuring income and increasing rather than constant marginal rates, two individuals with identical lifetime incomes can readily pay different lifetime taxes depending on the temporal pattern of income. A very simple example¹⁰³: Imagine that the marginal tax rates are 10% for the first \$30,000 of income and 30% for income thereafter. A and B live for two years. A has a fluctuating income and earns \$10,000 in the first year and \$90,000 in the next, while B has a constant income and earns \$50,000 in both years. Then A will owe \$1,000 in taxes in the first year and \$21,000 in the second, for a total of \$22,000, while B will owe \$9,000 in taxes in each year, for a total of \$18,000. A therefore owes more in total taxes than B, even though both have identical lifetime incomes. From an annual perspective, horizontal equity framed in income terms is simply inapplicable to A and B - - who have different incomes in both years -- but from a lifetime perspective the differential treatment of A and B appears arbitrary. As Vickrey explained, “no taxpayer should bear a heavier or lighter burden merely because certain items of his income happen to be earned or realized in one year or another . . .”¹⁰⁴ Vickrey’s system ensures that individuals with equal total lifetime incomes will pay an equal present value of taxes.

Specifying the horizontal-income-equity norm in terms of lifetime rather than annual income would also argue for adopting Vickrey-like averaging mechanisms for any

⁹⁹ FULLERTON & ROGERS, *supra* note 69, at 19.

¹⁰⁰ See William Vickrey, *Averaging of Income for Income-Tax Purposes*, 47 J. POL. ECON. 379 (1939); WILLIAM VICKREY, AGENDA FOR PROGRESSIVE TAXATION 164-97 (1947); William Vickrey, *Cumulative Averaging After Thirty Years*, in MODERN FISCAL ISSUES 117 (Richard M. Bird & John G. Head eds., 1972).

¹⁰¹ Jeffrey B. Liebman, *Should Taxes Be Based on Lifetime Income? Vickrey Taxation Revisited* (John F. Kennedy School of Government, December 2003), at 6.

¹⁰² On the horizontal-equity argument for Vickrey’s and similar proposals, see Kirk J. Stark & Lee Anne Fennell, *Taxation over Time* 22-27 (UCLA Law and Economics Research Paper No. 05-24, Sept. 6, 2005), at 22-27; Daniel Shaviro, *Permanent Income and the Annual Income Tax*, at 9-10.

¹⁰³ This example is taken from Fennell and Stark, *supra* note 102, at 23-24.

¹⁰⁴ Vickrey, *Averaging of Income*, *supra* note 100, at 381.

income-based tax or transfer program. Lily Batchelder argues for multi-year averaging in calculating the Earned Income Tax Credit.¹⁰⁵ Jeffrey Liebman notes that workers with identical lifetime earnings can pay different amounts of Social Security Tax, which is a constant percentage but has a cap.¹⁰⁶ A pure welfare program, designed just to alleviate poverty -- understood in lifetime rather than annual terms -- might well key annual payments to cumulative income, not annual income.¹⁰⁷

A lifetime-income specification of traditional horizontal- or vertical-equity tax norms not only strengthens the case for Vickrey-style income-averaging structures, which in actual practice have been rarely adopted.¹⁰⁸ That specification may also change how commonplace tax structures are evaluated in light of horizontal or vertical equity norms. Indeed, the main function of lifetime-incidence studies of tax structures is to inform such evaluation. As Fullerton and Rogers explain:

A misconception here is that the lifetime perspective mandates lifetime accounting to calculate current tax liability. Not so. Tax collections can still be based on annual accounts. But the lifetime perspective provides a useful yardstick to help evaluate any such tax system.¹⁰⁹

Consider the choice between the VAT and the existing income tax. The VAT is regressive in terms of annual income, while the existing income tax is progressive. The VAT is substantially less regressive in terms of lifetime income, while the existing income tax is less progressive. So the “vertical equity” argument for the existing income tax, over the VAT, weakens when we take the position that individuals are differently situated for purposes of “vertical equity” when they have different lifetime (not annual) incomes.

To be sure, for egalitarian welfarists, the proxy criteria are just that -- rough indices that economize on evaluation costs but can also produce errors. Individuals with equal lifetime incomes will not necessarily have equal lifetime utilities. In particular -- a point stressed by the scholarly literature on Vickrey’s proposal¹¹⁰ -- individuals with equal lifetime income may end up with different consumption sequences (pace classic economic wisdom) and therefore different levels of lifetime well-being. One response is to rework the notions of horizontal and vertical equity, defining them in terms of some function of consumption rather than income. Another is to move to first-best analysis. In fact, tax-policy scholarship is one area where policy proposals are regularly evaluated through the application of a SWF. The practice of doing so goes back to Mirlees’ work,

¹⁰⁵ See Lily L. Batchelder, *Taxing the Poor: Income Averaging Reconsidered*, 40 HARV. J. LEGIS. 395 (2003).

¹⁰⁶ Liebman, *supra* note 101, at 4

¹⁰⁷ Shaviro, *supra* note 102, 15-16. For helpful general discussions of income averaging, see *id.*; Stark & Fennell, *supra* note 102.

¹⁰⁸ A different set of esoteric structures, which may also be better than the straight income tax at equalizing lifetime income, are age-related taxes. See Stark & Fennell, *supra* note 102.

¹⁰⁹ FULLERTON & ROGERS, *supra* note 69, at 19.

¹¹⁰ See Stark & Fennell, *supra* note 102; Shaviro, *supra* note 102; Liebman, *supra* note 101.

on optimal taxation.¹¹¹ In this context, the lifetime perspective counsels -- of course -- that the SWF be defined as a function of lifetime utilities.

Some optimal-tax scholarship does exactly this. For example, Jeffrey Liebman studied a version of Vickrey's income-averaging proposal.¹¹² He assumes a population of three types of individuals, with steady, declining, or rising earnings over time. The individuals, in one scenario, smooth their consumption; in the other, they are "present focused" and consume the earnings when realized. Given a utilitarian SWF, Vickrey taxation is always better than an annual income tax in the smoothing scenario; whether it is in the "present-focused" scenario depends on the proportion of the three types of individuals. Liebman, then, compares the two tax structures using an equity-regarding SWF that depends on the lifetime utilities of the three types of individuals. The specific results depend on the parameters of the SWF but -- as it turns out -- tend to be similar to the utilitarian case. These conclusions about the effect of Vickrey averaging may or may not hold up with further analysis; my point, here, is to endorse Liebman's lifetime-focused methodology for optimal-tax modeling.

B. *Should Non-Tax Instruments be Used for Redistribution?*

Kaplow and Shavell have published a series of influential articles arguing that the income tax system rather than non-tax legal rules, such as tort doctrines, should be used to redistribute income.¹¹³ Their work is closely related to a body of scholarship in applied economics which argues against the use of "distributive weights" in cost-benefit analysis.¹¹⁴

An important analytic step in this work, certainly in Kaplow and Shavell's, involves Pareto-superiority. Imagine that a government decisionmaker is considering two non-tax policy options, P_{Equity} and $P_{\text{K-H}}$. The two options could, for example, be different sets of legal doctrines, or different regulations, or different expenditures on public goods. The decisionmaker has in hand an equity-regarding SWF, which favors P_{Equity} . But $P_{\text{K-H}}$ is Kaldor-Hicks efficient relative to P_{Equity} : those who benefit from $P_{\text{K-H}}$ could compensate those who are made worse off, via costless lump-sum transfers, making everyone better off than with P_{Equity} .

What is the argument for choosing $P_{\text{K-H}}$? One line of argument is not persuasive: namely that $P_{\text{K-H}}$, coupled with perfect redistributive taxation, is a Pareto improvement relative to P_{Equity} . By "perfect redistributive taxation," I mean just a system of costless lump-sum transfers -- one that has zero administrative costs and has no effect on

¹¹¹ See James A. Mirrlees, *An Exploration in the Theory of Optimum Income Taxation*, 38 REV. ECON. STUD. 175 (1971).

¹¹² See Liebman, *supra* note 101.

¹¹³ See in particular Louis Kaplow & Steven Shavell, *Why the Legal System is Less Efficient than the Income Tax in Redistributing Income*, 23 J. LEGAL STUD. 667 (1994); 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).

¹¹⁴ See Olof Johansson-Stenman, *Distributional Weights in Cost-Benefit Analysis--Should We Forget About Them?*, 81 LAND ECON. 337, 338-39 (2005) (citing sources).

individuals' welfare except to increase or decrease their holdings by the amount taxed or granted. The problem with this argument is that the actual tax system is imperfect, given administrative costs and hidden information.

Kaplow's and Shavell's contribution was to show that, in a class of cases, P_{K-H} , coupled with a kind of imperfect taxation, will be Pareto-superior to P_{Equity} . They focus on the income tax -- which is imperfect because it creates a disincentive to work. They show that, under certain conditions, where a legal rule is Kaldor-Hicks inefficient, the adoption of an efficient alternative together with matching changes in income tax rates will leave work incentives unchanged and raise a surplus -- which can be distributed in lump-sum form, making everyone better off than with the inefficient rule.

For their result to go through, Kaplow and Shavell have to assume that income, labor effort, and other characteristics enter into individual utility functions in a particular way. The debate, within legal scholarship, about their argument has largely focused on these assumptions.¹¹⁵ The question, in effect, is about the conditions under which a Kaldor-Hicks efficient policy can be converted into a Pareto-superior policy.

I want to come at the issue from a different direction. The lifetime perspective clarifies the *kind* of Pareto-superiority that arguments such as Kaplow and Shavell's need to establish.

The crucial point to see is that the four different variants of equity-regarding welfarism, defined by different answers to the time-slice problem and different views of how social welfare functions should be applied under conditions of uncertainty, generate four corresponding -- and *distinct* -- senses of Pareto-superiority.

Different Variants of Pareto Superiority Corresponding to Different Variants of Equity-Regarding Welfarism

<i>What is the Correct Approach to Social Choice Under Uncertainty?</i>			
		<u>Ex Ante</u>	<u>Ex Post</u>
<i>What is the Correct Time Slice?</i>	<u>Sublifetime</u>	EA-S (Ex Ante Sublifetime) Pareto Superiority	EP-S (Ex Post Sublifetime) Pareto Superiority
	<u>Lifetime</u>	EA-L (Ex Ante Lifetime) Pareto Superiority	EP-L (Ex Post Lifetime) Pareto Superiority

¹¹⁵ See Chris William Sanchirico, *Deconstructing the New Efficiency Rationale*, 86 CORNELL L. REV. 1003 (2001); Ronen Avraham et al., *Revisiting the Roles of Legal Rules and Tax Rules in Income Redistribution: A Response to Kaplow & Shavell*, 89 IOWA L. REV. 1125 (2004). See also Johansson-Stenman, *supra* note 114.

In particular, if one adopts the view argued for here -- the combination of lifetime welfarism and an “ex post” approach to social choice under uncertainty -- the matching sort of Pareto superiority is *ex post, lifetime Pareto superiority*. For short, let’s call this EP-L Pareto superiority. A policy P* is EP-L Pareto superior to alternative P if, in every state of the world, at least some individuals would have greater lifetime well-being if P* were chosen than they would if P were chosen, and no individual would have lower lifetime well-being if P* were chosen.

Along the temporal dimension, EP-L Pareto superiority is a relaxed notion of Pareto superiority. EP-sublifetime Pareto-superiority is more stringent. To see this point, consider any lifetime view that ascribes sublifetime as well as lifetime utilities.¹¹⁶ Let us bracket uncertainty for the moment, and imagine that we have two, certain outcomes O and O*, and various individuals. In some periods (e.g., some years), some individuals are worse off in O* than in O. However, no one has lower lifetime well-being in O*, and some have higher lifetime well-being. Sublifetime welfarists who accept the basic intuition of the Pareto principle can deny that O* is better than O. Paretian sublifetime welfarists would say that O* is Pareto superior to O only if, in every period, no individual is worse off in O* and some are better off. But lifetime welfarists surely cannot construe the Pareto principle so stringently. Lifetime welfarists think of well-being as an attribute of a whole human life. Paretians find it hard to see how an outcome that makes everyone better off could be a worse outcome. Lifetime welfarists might deny the Pareto principle; what is very hard to see is why they would accept it but insist on Pareto superiority in each period taken individually. In short, lifetime welfarists must count O* as a Pareto improvement over O -- in accordance with the EP-L construal of Pareto superiority.

Genuine Pareto improvements are therefore easier to secure for lifetime welfarists than for sublifetime welfarists. If one outcome produces greater sublifetime well-being for some, and no less for anyone, in every period, then it surely is Pareto superior in terms of lifetime well-being; but an outcome can be Pareto superior in terms of lifetime well-being without being Pareto superior in each period, as in the example just discussed.

Along the dimension of uncertainty, by contrast, EP-L Pareto superiority is a stringent notion of Pareto superiority. EP-L Pareto superiority demands that policy P* be a Pareto-improvement over P in every possible state of the world. I am relying, here, on Savage’s classic framework for formalizing choice under uncertainty. The decisionmaker expresses her uncertainty as a probability assignment to a set of possible states (where a “state” might be understood as a completely prior history of the world plus a set of possible causal laws). Each possible action is a mapping which assigns an outcome to each possible state. EP-L Pareto superiority requires that, for each possible state, the outcome that would occur if P* were chosen be a Pareto improvement over the outcome that would occur if P were chosen. In each state, some individuals must have greater lifetime well-being in the outcome that results from P*, and none may have lower lifetime well-being.

¹¹⁶ As mentioned earlier, not all lifetime views do this.

Why EP-L Pareto superiority? Why not ex ante-lifetime Pareto superiority (EA-L dominance), which says that P* is Pareto-superior to P if some have greater *expected lifetime well-being* with P* and none have lower expected lifetime well-being? If P* is EP-L Pareto-superior to P, then *any* SWF which respects the Pareto principle in its ranking of outcomes, and is applied in accordance with the correct account of choice under uncertainty -- the ex post account-- will prefer P* to P. By contrast, EA-L Pareto superiority doesn't guarantee that. Indeed, in our work on choice under uncertainty, Sanchirico and I highlight a stronger and counterintuitive result: For any continuous SWF which respects the Pareto principle in its ranking of outcomes and the Pigou-Dalton principle, there will be at least one pair of policies P+ and P++ such that P+ is EA-L Pareto superior to P++ but P++ is actually preferred by the SWF, if applied in accordance with the ex post approach, to P+.¹¹⁷

Return, now, to the question of using tax versus non-tax instruments to engage in redistribution. Those who argue that policymakers in some non-tax domain should ignore distributive considerations, and assert -- as do Kaplow and Shavell -- that this claim holds true regardless of the form of the SWF, must show that efficient policies in the domain can be converted, via the stipulated tax instrument, into EP-L Pareto-superior policies. Unless they can make *that* showing, the assertion that any Paretian equity-regarding SWF will prefer the efficient policy together with compensating tax changes will be untrue.

For a dramatic illustration of the point, consider risk regulation. Imagine that everyone in a reasonably large group of individuals is at risk of dying from a pollutant, a workplace hazard, a dangerous product, or some other risk source. A policy P prevents some premature deaths in the group, at some cost in consumption both to group members and nonmembers. Our equity-regarding SWF prefers P to a less interventionist policy P*. However, because the choice of P* simply increases the *risk* of death for group members, it is quite possible that P* is ex ante Kaldor-Hicks efficient relative to P. P* reduces each group member's expected lifetime well-being by some amount (by increasing the fatality risks to which group members are subject); P* increases the expected well-being of non-members by some amount (by increasing their consumption); and it is quite possible that there is a costless lump-sum redistribution of money, from group nonmembers to group members, such that P* coupled with this redistribution increases everyone's *expected well-being* as compared to P.

Imagine that our income tax system, despite its imperfections, is good enough to actually accomplish ex ante Paretian redistribution. P*, together with income tax adjustments, is EA-L Pareto-superior to P. It doesn't follow that P* is EP-L Pareto-superior to P. After all, *some group members will die prematurely as a result of P**. All

¹¹⁷ See Adler & Sanchirico, *supra* note 7, at 331-34. This result may not hold for all equity-regarding SWFs in the broader sense employed in this Article, but that does not affect my analysis of the Kaplow and Shavell claim. EA-L Pareto superiority cannot be used as a litmus test for the goodness of a policy, which one can use while remaining agnostic about the specific form of the SWF, because there are some (indeed, a large and plausible subset) which will reject some EA-L Pareto superior policies.

group members receive tax transfers such that the expected increase in each member's utility produced by the transfer exceeds the expected loss in his utility associated with his higher fatality risk. But, in any state of the world in which some group member dies earlier as a result of P*, he loses the entire utility difference between full life and premature death, and gains only the transfer. In that state, he might well be worse off as a result of P*-plus-transfer. Indeed, if the deaths at issue are instantaneous -- so that the transfer occurs after his death -- he will surely be worse off than if policy P had been implemented. So it is not, in fact, guaranteed that every equity-regarding SWF, if applied in an ex post manner, will prefer P*-plus-transfer to P. In some states, individuals are worse off (those who die prematurely) and inequality may be widened (between them and everyone else).

It might be objected that a sufficiently well-functioning tax system *could* convert an efficient policy that imposes an increased risk of death on some group into an EP-L Pareto-superior policy. But such a system would be fantastical: it would have to sort within the group, identifying those who actually will die in a given state, and targeting payments to them in that state (payments that must occur sufficiently long before they die to compensate for premature death).

To conclude, I should stress that this Section does *not* show that non-tax instruments should in fact be used for redistribution. It may be true that, with a particular SWF or family of SWFs in hand, we can show why policy choices in some domain should be made with reference to the Kaldor-Hicks criterion, leaving distributive considerations to the tax system. But Kaplow and Shavell offer a different argument against non-tax redistribution, one that can apply to all equity-regarding SWFs (more precisely, all that accept the Pareto principle for outcomes) because the argument invokes Pareto superiority. The thrust of this Section is to clarify the requisite structure of *this* sort of argument.

C. *Distributively Sensitive Policy Analysis*

1. Anti-Discrimination Norms, Disparate Impact Tests, and Protected Classes

Public law in the U.S. employs anti-discrimination norms and, to a lesser extent, disparate impact tests to constrain governmental policy choices. An anti-discrimination norm covers policies that facially refer to, or are designed to injure, certain suspect groups, subjecting such policies to a heightened burden of justification. The Equal Protection Clause protects racial minorities and women via an anti-discrimination norm. A disparate impact test demands heightened justification for nondiscriminatory laws that disproportionately harm members of a suspect class, as opposed to nonmembers. Under current doctrine, the Equal Protection Clause does not proscribe disparate impact. However, various nonconstitutional sources of law ground disparate impact tests applicable to governmental actors -- for example, Title VI, or Executive Order 12898 --

and measuring disparate impact with respect to racial minorities and low-income groups is a key focus of the literature on “environmental justice.”¹¹⁸

Discrimination norms and disparate impacts test are -- for the equity-regarding welfarist -- the non-tax analogues of notions of horizontal and vertical equity. They are proxy tests. “Suspect” or “protected” classes are groups of individuals who tend to be badly off with respect to well-being; policies that widen the welfare gap between them and the rest of society, either through explicit discrimination or disparate impact, are likely to be rejected by any sensible SWF absent a sufficient countervailing gain to overall welfare.

What are the implications of lifetime welfarism for anti-discrimination and disparate impact tests? To begin, it should be noted that the shift from a sublifetime to a lifetime framework does *not* substantially change the case for recognizing race and gender as suspect classifications. In a society where law or social norms stigmatize blacks as second-class citizens, or one where such norms have evaporated but existed in the recent past, the status of racial minority correlates with a lower level of *both* sublifetime *and* lifetime well-being. A similar account can be given with respect to gender.

Lifetime welfarism does affect the meaning of low-income status, for purposes of environmental-justice measures or other tests that recognize low-income groups as “suspect.” Bracketing measurement costs, the best specification of the low-income class looks to lifetime, not temporal income. Indeed, in the empirical literature that tries to quantify disparate impact along an income gradient, some of the newer work focuses on lifetime income.¹¹⁹

Let us turn, now, to the status of age as a protected class -- for it is here that lifetime and sublifetime welfarism diverge most dramatically. Within sublifetime welfarism, the case for viewing age as a protected class is strong. Age correlates with lower levels of momentary well-being, once age is sufficiently advanced. Certainly it is plausible that (say) 80-year olds tend to have a lower level of momentary well-being than 35-year olds: because of health problems, lower incomes, and social isolation. So redistributing health care or income from the 80-year to the 35-year old is, from a momentary perspective, a transfer from badly off to better off.

By contrast, from the lifetime perspective, age should not be a suspect class. The average lifetime well-being of 80-year-olds is *greater* than the average lifetime well-being of 35-year olds -- since the 80-year-olds will end up living longer, on average, than the 35-year-olds.

¹¹⁸ See Matthew D. Adler, *Risk, Death and Harm: The Normative Foundations of Risk Regulation*, 87 MINN. L. REV. 1293, 1423-26 (2003).

¹¹⁹ See Tom Van Ourti, *Socio-economic Inequality in Ill-Health Amongst the Elderly: Should One Use Current or Permanent Income?* 22 J. HEALTH ECON. 219 (2003).

So what is the relevant suspect class? Youth? Will the lifetime welfarist view policies that discriminate against older individuals as a kind of affirmative action for the young? Some lifetime welfarists have actually defended this position.¹²⁰ To see why the position is incorrect, or at least much too crude, it is important to remember that lifetime welfarism should be married with an ex post, rather than ex ante, approach to social choice under uncertainty.

Consider a policy that imposes a tax on 80-year-olds to benefit 35-year-olds. Some of the 35-year-olds will live to be 80; some will die before 80, but live to experience many years of adulthood; a small fraction will die only a few years after reaching 35. Call these subgroups of the 35-year-olds subgroups A, B, and C. The policy (if kept in place permanently) simply shifts resources within the lives of the individuals in subgroup A, and may well harm rather than benefit them by doing so. Many individuals who live to advanced ages have a lifetime health and consumption profile that could be improved by shifting resources to later in life; the policy, in its effect on subgroup A, works in the opposite direction. Subgroup B is benefited by the policy (bracketing incentive effects), because they reap its benefits but will never live to be subject to the tax. Yet the individuals in subgroup B will have reasonably long lives. Shifting resources from the 80-year-olds to them may not be counted -- depending on the shape of the SWF -- as a large improvement in social welfare. It is a redistribution from the rich (in terms of lifetime well-being) to the merely prosperous.

In transferring resources from the 80-year-olds to subgroup C, the policy does work a clear redistributive benefit. But subgroup C may be only a small fraction of the total group of 35-year olds. The redistributive benefit of shrinking the gap between their lifetime well-being, and that of the 80-year olds, may not justify the overall costs of the policy -- administrative costs, reduced lifetime well-being of current 80 year olds, and reduced lifetime well-being of the 35-year olds who will live to be 80.

So, from the lifetime perspective, age is not a good candidate for a suspect class, but neither is youth -- at least in a society like the U.S. where many of the young will end up becoming fairly old. Within lifetime welfarism, the characteristic that generates redistributive claims is lifetime well-being; the lower the individual's level of lifetime well-being, the stronger her redistributive claim. Since *neither* the attribute of being old, *nor* the attribute of being young, is strongly correlated with having a low level of lifetime well-being, neither the old nor the young should be a "suspect" class for purposes of proxy tests that roughly implement lifetime welfarism.

A better candidate for "suspect" status -- odd as this may sound -- is *lifespan*. Longevity is not, of course, the sole determinant of lifetime well-being. Individuals with equal lifespans can surely end up at different lifetime welfare levels, if the experiential, health, or social components of their lives are different. Longevity is simply one dimension of well-being. But longevity has the special property of being, not just a dimension of well-being, but a precondition as well. Dead individuals don't have social lives, pleasant experiences, meaningful careers, and so on. This strengthens the case for

¹²⁰ See Klemens Kappel & Peter Sandoe, *QALYs, Age and Fairness*, 6 *BIOETHICS* 297 (1992).

thinking of the attribute of *having a short lifespan* as “suspect.” However much happiness and love is jammed into a life that ends at age five, that unfortunate child will end up much worse off in terms of lifetime flourishing than most others in her society. Those at the bottom of the lifespan distribution are a kind of impoverished group, analogous in this respect to those at the bottom of the (lifetime) income distribution, or those who bear the lifetime burden of being socially stigmatized as second-class citizens.

Should legal doctrine actually recognize the short-lifespanded as “suspect” for purposes of equality tests? That proposal is less outrageous than it may seem at first blush. Explicit discrimination against the short-lifespanded does not, of course, occur. But policies can have a disparate impact on them -- most obviously, policies that increase juvenile mortality. (Consider a policy to terminate funding for some health care measure targeted at a serious, possibly fatal, disease among infants and juveniles. Those who would be harmed by this policy are, disproportionately, short-lifespanded, as compared to those who would be benefited.).

The insight that short lifespan constitutes a kind of lifetime well-being poverty can also be integrated into policy choice in other ways than through a “disparate impact” test. For example, agencies undertaking cost-benefit analysis might use a greater monetary valuation of life for children’s lives. Or, reducing juvenile mortality might be generally taken as a high-priority goal of health care and development policy. To a substantial extent, policymakers do already appreciate the distributive significance of reducing juvenile mortality; but there is no consensus basis for doing so within the scholarly literature on equality and poverty. As David Hulme and Andrew Sheppard explain:

A key conceptual challenge for the study of chronic poverty is how to treat those who die preventable deaths. These need to be incorporated in our analysis as they experience the most acute form of deprivation.... This must be included or a household whose young children die may be reported as escaping poverty when a similar household that manages to raise all its children ... is seen as ‘failing’ to escape from poverty as its income/consumption levels per capita are lower because of its larger size and higher dependency ratio. . . . Concepts from the health sciences, such as disability adjusted life years [a kind of QALY], may provide a basis for starting to think through this theoretical frontier.¹²¹

Lifetime welfarism provides a firm normative basis for the cutting-edge poverty measures that Hulme and Sheppard propose, and for policy tests that give extra weight to lifesaving among the young -- be these in the form of proxies and presumptions, or in the form of more a systematic scheme of distributive weights, a topic to which we now turn.

2. Distributive Weights and other Integrated Policy-Analytic Tools

A substantial body of scholarship discusses the possibility of developing policy-evaluation tools that would be sensitive to *both* considerations of overall well-being or efficiency (unlike inequality or poverty measures) *and* distributive considerations (unlike

¹²¹See Hulme & Sheppard, *supra* note 89. See also Ravi Kanbur & Diganta Mukherjee, *Premature Mortality and Poverty Measurement* (ISER Working Paper 2003-6).

cost-benefit analysis or QALY-based cost-effectiveness analysis), and would incorporate “distributive weights” or some such device to systematically structure tradeoffs between overall well-being and equity.¹²² Administrative agencies in the U.S., by contrast with agencies in some other countries, do not currently employ distributionally sensitive policy-evaluation tools of this sort. They should consider doing so -- at least once tools that have sufficiently low decision costs and are sufficiently transparent to agency overseers are developed.¹²³

The questions whether well-being should be measured on a lifetime and sublifetime basis, and whether an ex post or ex ante perspective should be adopted, are absolutely fundamental to designing a distributionally sensitive procedure for policy analysis. These questions have surfaced, most clearly, in the scholarship authored by health economists that proposes incorporating equity weights in QALY-based analysis. Erik Nord proposes that equity weights, here, should take the form of “severity” weights.¹²⁴ The basic idea is that a change in an individual’s health should be weighted to account for his current health condition. Providing a .1-QALY improvement to an individual currently at QALY level .2 should count for more than providing the same improvement to an individual currently at level .4. More generally, we can estimate an individual’s expected QALYs in each future year, in the status quo and given some policy intervention. The expected improvement in the individual’s QALYs in each year, provided by the policy, should be weighted to take account of the level of expected QALYs that the individual would otherwise receive, in the status quo.

[Imagine that] a person’s utilities in the next five years in case of non-intervention are expected to be 0.7, .7, .6, .5, and 0.0 (dead), and the utilities would be .8 in all 5 years (and then dead) in case of intervention The severity approach implies the application of severity weights . . . to each of these annual utility gains.¹²⁵

From the perspective of lifetime, ex post, welfarism, Nord’s proposal is doubly problematic: first, for making the distributive weight applied to a health change a function of the annual rather than lifetime health status of the affected individual; and, second, for applying the weights to *expected* changes in individuals’ *expected* annual QALYs in each future year, thus implicitly adopting an “ex ante” account of distributional sensitivity under uncertainty.

The health economist Alan Williams disagrees with the sublifetime focus of Nord’s proposal. Williams’ suggestion, the so-called “fair innings” proposal, is that fairness concerns the distribution of lifetime QALYs. “A powerful part of the rhetoric

¹²² The literature on equity-weighted QALYs is large. For an overview, see F. Sassi et al., *Equity and the Economic Evaluation of Healthcare*, 5 HEALTH TECH. ASSESSMENT, no. 3 (2001). On distributionally weighted cost-benefit analysis, see Johansson-Stenman, *supra* note 114; Ian W.H. Parry et al., *The Incidence of Pollution Control Policies* (NBER Working Paper 11438), at 26-31.

¹²³ On decision costs and transparency as desiderata for administrative policy analysis tools, see generally ADLER & POSNER, *supra* note 22, at 80-88, 122-23.

¹²⁴ For a recent discussion by Nord, see Erik Nord, *Concerns for the Worse Off: Fair Innings versus Severity*, 60 SOC. SCI. MED. 257 (2005).

¹²⁵ *Id.* at 259.

about equity in health care employs the notion of equality applied to *people's whole lifetime experiences* and not just to their current situation.”¹²⁶ Williams proposes using survey research to determine how much population life expectancy citizens are willing to sacrifice for the sake of equalizing life expectancy. This research would generate weighting factors, which would be applied to expected QALY changes depending on the expected lifetime QALYs of those affected. From the perspective of ex post lifetime welfarism, Williams' proposal for equity-weighted QALY analysis is a significant improvement over Nord's, but it is still marred by its ex ante focus. “The fair innings weight applied to health gains of each individual is defined as a function of his/her *present expected lifetime QALYs*”¹²⁷

What would a better approach be? First, at least ideally, distributively sensitive policy analysis should consist in the application of a social welfare function. It should consist in “SWF” analysis. This proposal is less utopian than it may seem. The SWF is the best, systematic intellectual framework currently available for considering both the distributional and overall-welfare impacts of policies. As already mentioned, the scholarly literature on optimal taxation pervasively employs the SWF framework¹²⁸; much of the non-tax literature on distributive weighting draws on that framework as well¹²⁹; and there is some scholarly work that actually considers specific non-tax policy proposals using SWFs.¹³⁰ A crucial issue to be overcome, in rendering SWF analysis a feasible and sufficiently transparent tool for government bodies, is to specify the form of the SWF. Should it be prioritarian or comparativist? What should its parameters be? The so-called Atkinsonian family of SWFs is perhaps most widely used by scholars, and is a prioritarian SWF (with the degree of priority to welfare gains for worse-off individuals captured by an inequality aversion parameter ϵ).¹³¹ The rank-weighted SWF, corresponding to the popular Gini coefficient, is a kind of comparativist SWF.¹³² Thus a first stab at SWF analysis would evaluate policies using the Atkinsonian SWF (with a range of values of ϵ) and the rank-weighted SWF.

Second, distributively sensitive policy analysis should consider a variety of sources of well-being (up to the limits of feasibility, of course). QALY analysis, equity-weighted or not, is problematic for these purposes because it ignores income differences.¹³³ A better approach would be to take account of longevity, health *and* income. Each individual's lifetime utility would be a function of his health status and his

¹²⁶ Williams, *supra* note 28, at 120.

¹²⁷ Aki Tsuchiya, *QALYs and Ageism: Philosophical Theories and Age Weighting*, 9 HEALTH ECON. 57, 62 (2000) (emphasis added).

¹²⁸ See [to be filled in]

¹²⁹ See, e.g., Han Bleichrodt et al., *Equity Weights in the Allocation of Health Care: The Rank-Dependent QALY Model*, 23 J. HEALTH ECON. 157 (2004); Paul Dolan, *The Measurement of Individual Utility and Social Welfare*, 17 J. HEALTH ECON. 39 (1998); Lars Lindholm et al., *An Epidemiological Approach Towards Measuring the Trade-off Between Equity and Efficiency in Health Policy*, 35 HEALTH POLICY 205 (1996); Lars Peter Osterdal, *Axioms for Health Care Resource Allocation*, 24 J. HEALTH ECON. 679 (2005); Adam Wagstaff, *QALYs and the Equity-Efficiency Trade-Off*, 10 J. HEALTH ECON. 21 (1991);

¹³⁰ See Parry et al., *supra* note 122, at 27-28 (citing sources).

¹³¹ See, e.g., BOJER, *supra* note 55, at 109-110.

¹³² See Adler & Sanchirico, *supra* note 7.

¹³³ See Adler, *supra* note 76.

income in each period. In effect, this suggestion synthesizes the QALY literature on distributive sensitivity, which focuses on health differences; and the optimal tax and cost-benefit literatures, which focus on income and consumption differences. To be sure, in deriving an individual's lifetime *utility* from his health, income and longevity, we need a utility function. The existing literature often posits a utility function based on certain axiomatic properties (for example, declining marginal utility of income); a better approach, one I have discussed elsewhere, would be to use survey research about well-being to calibrate the utility function.¹³⁴

Finally -- and this is where lifetime welfarism plus the ex post perspective comes into play -- SWF analysis should take the form of *probabilistic population profile analysis*. A population profile is a set of individual life histories, consisting of a determinate income, health, and longevity trajectory for each individual. The status quo consists in a probability distribution over population profiles. A policy intervention perturbs this probability distribution, changing the population profile in some states by changing some individual(s)' life histories. A utility function maps each life history onto a utility. For a given SWF, be it the Atkinsonian SWF, the rank-weighted SWF, or some other, we calculate the expectation of the SWF in the status quo, and its expectation with the policy.

I discuss probabilistic population profile analysis at much greater length in other work.¹³⁵ My aim, here, is simply to sketch the approach and to suggest that it flows naturally from the general position argued for in this Article. If we take seriously the thought that distributively sensitive policy analysis should consist in SWF analysis, then -- given lifetime, ex post welfarism -- we should take seriously the suggestion that probabilistic population profile analysis is the appropriate template for the SWF approach.

It is possible, I suppose, that a different template is optimal. Perhaps, for example, probabilistic population profile analysis is too computationally demanding. Magnus Johannesson argues, along these general lines, against an ex post approach to equity-weighted QALYs.

From the *ex post* perspective, the equity weights have to be applied to all possible health profiles before they are multiplied by the probabilities of the different profiles If, for instance, an individual has a 50% chance of living 20 years in full health and a 50% chance of living 40 years in full health, the equity weights have to be applied to the 20 and 40 years, respectively, before the equity-weighted QALYs of each health profile are multiplied by .5. This implies that it would be much more impractical to take into account *ex post* equity than *ex ante* equity.¹³⁶

Yet current policy-analysis techniques often involve extensive, probabilistic, calculations. Consider toxic risk assessment, which relies upon complicated probabilistic models of the different steps by which a toxin is disseminated to the population, and uses

¹³⁴ See Matthew D. Adler, *Welfare Polls: A Synthesis*, 81 N.Y.U. L. REV. 1875 (2006).

¹³⁵ See Matthew D. Adler, *Risk Equity: A New Proposal*, HARV. ENVTL. L. REV. (forthcoming 2008).

¹³⁶ Magnus Johannesson, *Should We Aggregate Relative or Absolute Changes in QALYs*, 10 HEALTH ECON. 573, 574 (2001).

computationally demanding techniques such as Monte Carlo analysis to synthesize the different models. The biggest hurdle to SWF analysis by agencies would seem to be the normative task of specifying a particular SWF and individual utility function, not the modeling task of estimating the possible distribution of the determinants of well-being and applying the SWF “ex post” rather than “ex ante.” In any event, probabilistic population profile analysis deserves serious attention. Other templates for SWF analysis should be seen as second-best.

CONCLUSION

Should equality be evaluated from a lifetime or sublifetime perspective? I have engaged this question within the framework of welfarism -- taking equality of *well-being* as the goal of egalitarianism -- and have argued for the lifetime view, given both standard and revisionary accounts of personal identity over time. I have also surveyed the multiple implications of the lifetime view for law and policy: for the measurement of equality, the measurement of poverty, the design of redistributive taxes, the question whether non-tax instruments should be used for redistribution, and the nature of “suspect classes” and distributively sensitive policy tools.

The reader may not share my view that welfarism and the “SWF” framework provide insight into problems of social choice. She might prefer to ask: given some non-well-being currency for equality (for example, “resources” or “capabilities”), should we seek to equalize lifetime or sublifetime holdings of *that* currency? Or, she might be sympathetic to welfarism but be unpersuaded by the case in favor of a lifetime view (which, like much normative argument, is holistic rather than deductive -- there is nothing logically flawed or flatly incoherent about the sublifetime view). But I hope the reader does agree that the time-slice *question* is a vital one, not just for philosophers and theoretical economists, but for legal scholars, applied economists, public health scholars, and policymakers too. We cannot think seriously about how to measure and alleviate inequality without engaging that question.