

DUKE ENVIRONMENTAL AND ENERGY ECONOMICS WORKING PAPER SERIES
organized by the
NICHOLAS INSTITUTE FOR ENVIRONMENTAL POLICY SOLUTIONS
and the
DUKE UNIVERSITY ENERGY INITIATIVE

A Better Calculus for Regulators: From Cost-Benefit Analysis to the Social Welfare Function

Matthew Adler*

Working Paper EE 17-01
March 2017

*Duke University School of Law

The Duke Environmental and Energy Economics Working Paper Series provides a forum for Duke faculty working in environmental, resource, and energy economics to disseminate their research. These working papers have not necessarily undergone peer review at the time of posting.



**A BETTER CALCULUS FOR REGULATORS:
FROM COST-BENEFIT ANALYSIS TO THE SOCIAL WELFARE FUNCTION**

Matthew D. Adler, Richard A. Horvitz Professor of Law and Professor of Economics,
Philosophy and Public Policy, Duke University. adler@law.duke.edu

Working Paper, March 2017

Introduction

I. The Social Welfare Function (SWF) Framework

A. *Ethics, Weak Consequentialism, Welfarism*

1. Weak Consequentialism
2. Welfarism

B. *The SWF Framework: Rules and Axioms*

C. *The Well-Being Function*

D. *Uncertainty*

II. The SWF Framework as a Guide to Regulatory Policy

A. *The SWF Framework as a Basis for Environmental, Health, and Safety Regulation*

B. *Cost-Benefit Analysis (CBA)*

C. *A Worked Example*

1. Uniform Risk Regulation and Cost Incidence
2. Other Policies
3. A Summary

III. Why the SWF Approach Improves on CBA: An Argument

A. *Defenses of CBA?*

1. Long Run Pareto
2. Potential Pareto
3. Rough Proxy for Overall Well-Being

B. *Objections to the SWF Approach*

1. Value Choices and Legal Legitimacy
2. The Tax System

C. *SWFs as a Regulatory Methodology: A Nuanced Approach*

Conclusion

Appendix

A BETTER CALCULUS FOR REGULATORS: FROM COST-BENEFIT ANALYSIS TO THE SOCIAL WELFARE FUNCTION

Introduction

Cost-benefit analysis (CBA) has become the dominant governmental methodology, in the U.S., for evaluating regulatory policy. Since 1981, a Presidential order has directed regulatory agencies in the executive branch to comply with a cost-benefit standard, where statutory discretion exists to do so, and to prepare an analytic document describing the costs and benefits of major rules.¹ These Presidential orders are enforced by OIRA, an oversight body within the Executive Office of the Presidency, which—along with policy offices within each regulatory agency—constitutes a significant bureaucratic structure implementing the construct of CBA.²

CBA is certainly an advance on what existed before 1981. A regulatory agency subject to the CBA order is instructed to characterize the effects of its proposed rules along multiple dimensions of human well-being—beyond the specific dimension (health, safety, environmental quality, etc.) that its organic statute highlights—and to make some attempt, using the metric of dollars, to commensurate effects along these multiple dimensions. This is progress.

But CBA is not the endpoint of good regulatory policy analysis. It can be bettered. In this Article, I show how. We can improve CBA-based policy assessment by using the “social welfare function” (SWF) as a tool to evaluate proposed regulations.

The SWF framework conceptualizes the status quo and each policy alternative as a pattern of *well-being* across the population of concern (or, given uncertainty, as a probability distribution across such patterns). CBA uses money as the metric for gauging policy effects on each individual, but money has “diminishing marginal utility”: a \$10,000 increase in the money holdings of a millionaire is not the same, in well-being terms, as a \$10,000 increase in the holdings of someone with average income: the millionaire gains *less* in welfare. The SWF approach corrects for the diminishing marginal utility of money by using an appropriately constructed measure $w(\cdot)$ of individual well-being as the indicator of how well each person is doing, and how much he or she stands to gain or lose from a given policy.

A related point is that the SWF framework is sensitive to distributional considerations, while CBA is not. The framework can be specified in various ways—both with respect to the specific steps taken in constructing the well-being measure $w(\cdot)$, and with respect to the rule adopted for ranking patterns of well-being. One such rule is “utilitarian”; the utilitarian SWF says that policy P is better than policy P^* if the sum total of individual well-being is greater with P . The utilitarian SWF is sensitive to the distribution of *income*. *Ceteris paribus*, the utilitarian SWF favors the transfer of a unit of income (dollars) from higher- to lower-income individuals,

¹ Exec. Order No. 12,866, 3 C.F.R. 638 (1994).

² For a summary of the system of regulatory review in the U.S., see ANDREA RENDA, LAW AND ECONOMICS IN THE RIA WORLD ch. 2 (2011).

since income has declining marginal utility. So-called “prioritarian” SWFs are yet more sensitive to the arrangement of income and other well-being attributes across the population. SWFs within the prioritarian family prefer, not merely to equalize the distribution of income, but indeed to equalize the distribution of well-being itself. *Ceteris paribus*, these SWFs favor the transfer of a unit of *well-being* from someone at a higher level of well-being, to someone at a lower level.

The SWF framework has never (as far I’m aware) been put to use by U.S. regulatory agencies. However, it has deep roots in the academic literature. It originates in theoretical welfare economics, in work by Abram Bergson and Paul Samuelson from the 1930s and 1940s and, somewhat later, by Amartya Sen in response to the Arrow impossibility theorem.³ It is the linchpin for the contemporary literature in economics on optimal taxation. James Mirrlees, in path-breaking scholarship from the 1970s, used an SWF to study the problem of optimizing the schedule of income tax rates—balancing the gains to overall well-being from the redistribution of a fixed “pie” of income, against the disincentive to income-generating activities that occurs with taxes.⁴ Since Mirrlees’ work (work that earned him the Noble Prize), the leading methodology used by economists to address normative questions regarding taxation—the specification of income-tax rates, the choice between different types of taxes, and so on—has been the SWF.⁵

But the range of application of the SWF approach goes well beyond tax policy. It can, in principle, be used to regiment the normative assessment of *any* type of tax *or* non-tax policy. Indeed, SWFs are widely employed in the field of climate economics.⁶ How to make tradeoffs between the material costs of reducing carbon emissions (costs that may be distributed in various ways among the richer and poorer members of both present and future generations, depending on the policy choice), and the benefits (both environmental and material) of slowing warming (these benefits, too, being distributed both within and across the generations in one or another manner),

³ See MATTHEW D. ADLER, WELL-BEING AND FAIR DISTRIBUTION: BEYOND COST-BENEFIT ANALYSIS 79-86 (2012) (describing the origins of the SWF concept).

⁴ James Mirrlees, *An Exploration in the Theory of Optimum Income Taxation*, 38 REVIEW OF ECONOMIC STUDIES 175 (1971).

⁵ For overviews of the literature on so-called optimal taxation, including the role of the SWF therein, see LOUIS KAPLOW, THE THEORY OF TAXATION AND PUBLIC ECONOMICS (2008); MATTI TUOMALA, OPTIMAL REDISTRIBUTIVE TAXATION (2016); MATTI TUOMALA, OPTIMAL INCOME TAX AND REDISTRIBUTION (1990); GARETH MYLES, PUBLIC ECONOMICS (1995); BERNARD SALANIE, THE ECONOMICS OF TAXATION (2003); Mikhail Golosov and Aleh Tsyvinski, *Policy Implications of Dynamic Public Finance*, 7 ANNUAL REVIEW OF ECONOMICS 147 (2015); N. Gregory Mankiw, Matthew Weinzierl, and Danny Yagan, *Optimal Taxation in Theory and Practice*, 23 JOURNAL OF ECONOMIC PERSPECTIVES 147 (2009); Peter Diamond and Emmanuel Saez, *The Case for a Progressive Tax: From Basic Research to Policy Recommendations*, 25 JOURNAL OF ECONOMIC PERSPECTIVES 165 (2011); ROBIN BOADWAY, FROM OPTIMAL TAX THEORY TO TAX POLICY (2012).

⁶ See generally W.J. Wouter Botzen and Jeroen C.J.M. van den Bergh, *Specifications of Social Welfare in Economic Studies of Climate Policy: Overview of Criteria and Related Policy Insights*, 58 ENVIRONMENTAL AND RESOURCE ECONOMICS 1 (2014). On prioritarian SWFs and climate economics, see in particular Matthew D. Adler and Nicholas Treich, *Prioritarianism and Climate Change*, 62 ENVIRONMENTAL AND RESOURCE ECONOMICS 279 (2015).

is a complex ethical problem that has been much explored—and, I suggest, much illuminated—with the SWF construct. Other areas of application of the SWF construct include refinements of GDP as a measure of social condition,⁷ health care policy,⁸ environmental, health and safety regulation,⁹ and inequality metrics (which have a conceptual link to SWFs).¹⁰ The “Green Book” (the official policy-assessment document in the U.K., applicable to regulations as well as other types of governmental interventions, such as infrastructure spending) generally instructs decisionmakers to conduct CBA with “distributive weights.”¹¹ CBA with distributive weights is a refinement of the standard unweighted technique, and can be used to approximate an SWF.¹²

Part I of this Article provides an overview of the SWF framework. Part II describes how that framework might be used as a guide to regulatory policy. The focus of Part II is environmental, health and safety regulation (a.k.a “risk regulation”): governmental interventions designed to reduce individual fatality and morbidity risks.¹³ Since environmental, health and safety regulation has been the leading substantive area of application of CBA to regulatory policy within the US government, it will also serve in this Article as the exemplar—with reference to which I’ll illustrate how the SWF methodology would function as a tool for regulators.

Part II includes a stylized case study: a regulatory intervention reduces fatality risks to various population groups (differentiated by age, income, and health), at some cost to individuals’ incomes. Whether the policy is justified depends upon the pattern of risk reduction and income loss among the various groups—and upon the parameters of the SWF used to evaluate these changes. Using the case study, I illustrate the utilitarian and prioritarian SWFs, and compare these SWFs both to each other and to (two variants of) CBA.

⁷ See, e.g., Paul Schreyer, *GDP*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY 21-46 (Matthew D. Adler and Marc Fleurbaey eds., 2016); Marc Fleurbaey, *Beyond GDP: The Quest for a Measure of Social Welfare*, 47 J. ECON. LIT. 1029 (2009); Koen Decanq and Erik Schokkaert, *Beyond GDP: Measuring Social Progress in Europe* (working paper, KU Leuven, April 2013); World Bank, *A Measured Approach to Ending Poverty and Boosting Shared Prosperity: Concepts, Data and the Twin Goals* (Policy Research Report, 2015).

⁸ See Adler, *supra* note 3, at 87 n. 56 (citing sources).

⁹ See, e.g., Matthew D. Adler, James K. Hammitt and Nicolas Treich, *The Social Value of Mortality Risk Reduction: VSL versus the Social Welfare Function Approach*, 35 JOURNAL OF HEALTH ECONOMICS 82 (2014); R. Baker et al., *Valuing Lives Equally: Defensible Premise or Unwarranted Promise?* 36 J. RISK AND UNCERTAINTY 125 (2008); Luc Bovens and Marc Fleurbaey, *Evaluating Life and Death Prospects*, 28 ECONOMICS AND PHILOSOPHY 217 (2012); T. Gajdos et al., *Shared Destinies and the Measurement of Social Risk Equity*, 176 ANNALS OF OPERATIONS RESEARCH 409 (2010).

¹⁰ See Adler, *supra* note 3, at 114 n. 101 (citing sources); Frank Cowell, *Inequality and Poverty Measures*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 82.

¹¹ HM TREASURY, THE GREEN BOOK: APPRAISAL AND EVALUATION IN CENTRAL GOVERNMENT 91-94 (2003).

¹² Matthew D. Adler, *Benefit-Cost Analysis and Distributional Weights: An Overview*, 10 REVIEW OF ENVIRONMENTAL ECONOMICS AND POLICY 264 (2016); Robin Boadway, *Cost-Benefit Analysis*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 126.

¹³ The proposal in this Part builds on an existing literature regarding the use of SWFs to evaluate risk regulation. See sources cited *supra* note 9.

Having set forth the SWF methodology, both in a general way (Part I) and with specific reference to regulatory policy (Part II), I turn in Part III to a head-to-head *normative* appraisal of that methodology as contrasted with CBA. Why believe that the SWF approach is indeed an improvement on CBA when it comes to environmental, health and safety regulation—or regulation more generally?

A word on timing. This Article was largely drafted in the months before November 8, 2016, when it seemed quite likely that the 45th U.S. President would keep in place the system of OIRA-based regulatory review built up over 35 years, and plausible that he or she might consider refinements to CBA designed to make it more sensitive to distributional considerations. The path of regulatory review under President Trump is, as of now, quite unclear. The Trump Administration may well retain CBA¹⁴; but it seems hard to imagine that it will have much interest in utilitarian or prioritarian SWFs.

Still, our 45th President will be succeeded by a 46th. And CBA is not confined to the U.S. It is, increasingly, used by governments abroad. Most of the analysis in this Article is applicable to any regime of regulatory review. (The one exception is the U.S.-centric discussion in Part III of legal legitimacy and divided government.) Scholars of the regulatory process shouldn't take the 2016 election as a signal to cease the work of contributing to our collective knowledge about good regulatory policy and policy-analysis. The knowledge will be there to tapped, now or in the future, by political leaders who understand its value.

I. The SWF Framework

This Part provides an overview of the SWF approach, highlighting its key features.¹⁵ The Part draws upon (and ruthlessly compresses) a large body of scholarship, concerning SWFs or related topics, in social choice theory, welfare economics, and normative ethics.

¹⁴ On January 30, 2017, President Trump issued an Executive Order on “Reducing Regulation and Controlling Regulatory Costs.” <https://www.whitehouse.gov/the-press-office/2017/01/30/presidential-executive-order-reducing-regulation-and-controlling>. The order leaves in place the existing system of CBA and OIRA oversight as per Executive Order 12,866: that order is not repealed. Rather, the new order adds to 12,866 a requirement that the costs of any new regulation be offset by the repeal of two existing regulations; and that agencies not impose overall regulatory costs beyond an annual regulatory budget, which for 2017 is set to \$0 (which overall number is presumably net of savings achieved by repealing prior rules). What exactly the new order means will become clearer over time.

¹⁵ For more comprehensive and rigorous discussions of the SWF approach, see Adler, *supra* note 3; CHARLES BLACKORBY, WALTER BOSSERT AND DAVID DONALDSON, POPULATION ISSUES IN SOCIAL CHOICE THEORY, WELFARE ECONOMICS, AND ETHICS, chs. 3-4 (2005); ROBIN BOADWAY AND NEIL BRUCE, WELFARE ECONOMICS ch. 5 (1984); Walter Bossert and John Weymark, *Utility in Social Choice*, in HANDBOOK OF UTILITY THEORY 1099 (S. Barbera et al. eds, 2004, vol. 2); Claude d’Aspremont and Louis Gevers, *Social Welfare Functionals and Interpersonal Comparability*, in HANDBOOK OF SOCIAL CHOICE AND WELFARE 459 (Kenneth J. Arrow et al. eds., vol. 1., 2002); Amartya Sen, *Social Choice Theory*, in HANDBOOK OF MATHEMATICAL ECONOMICS 1073 (Kenneth J. Arrow and Michael Intriligator eds., vol. 3, 1986); Phillippe Mongin and Claude d’Aspremont, *Utility Theory and Ethics*, in HANDBOOK OF UTILITY THEORY 371 (S. Barbera et al. eds, vol. 1, 1998); John Weymark, *Social Welfare Functions*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 126.

The presentation is quite abstract. Abstraction is—I believe—the best way to convey the essence of the SWF methodology. Still, some readers may be impatient for practical application and institutional context. These will come in Parts II and III.

The SWF methodology is a tool for *ethical* deliberation, more specifically ethical deliberation that is (weakly) consequentialist and welfarist. These matters are addressed in Section A. The methodology combines a measure $w(\cdot)$ of well-being, with some rule for ranking patterns of well-being. See Sections B and C. The utilitarian SWF and prioritarian family of SWFs emerge as especially plausible. The application of SWFs under uncertainty is the topic of Section D—with a specific focus on the three-way choice between utilitarianism, “ex ante” prioritarianism, and “ex post” prioritarianism.”

The presentation, although abstract, is relatively non-technical in eschewing the heavy use of mathematical formalism that one finds in most treatment of SWFs. A brief, technical summary of the methodology is provided in the Appendix.

A. *Ethics, Weak Consequentialism, Welfarism*

The SWF approach, as I conceptualize it, is a framework for ethical deliberation.¹⁶ It functions to guide a decisionmaker in determining what she ought to do, ethically. By “ethical” (or equivalently “moral”) deliberation I mean normative deliberation that is impartial (among everyone in some population of concern—at the limit, all persons and other beings with ethical standing) and that is not constrained by existing social practices, but rather is willing to criticize such practices as ethically misguided or suboptimal.

David Hume is famous for observing that “ought” can’t be deduced from “is,”¹⁷ and the reader would do well to intone Hume’s dictum. Some economists seem to think that identifying an SWF is a purely factual matter—a topic for social science.¹⁸ This is a double mistake. First, modern, pluralistic societies are characterized by ethical disagreement, and thus it is very implausible (as a factual matter) that everyone or mostly everyone within a given society espouses the very same SWF. Second, to give policy advice is to engage in normative discourse: it is to tell the decisionmaker that she *ought* to do something (choose a particular policy, reject a policy, seek more information, etc.). Whatever facts the economist might observe, the inference *from* those *to* some conclusion about what the decisionmaker “ought” to do cannot be grounded upon facts alone. Some of the premises grounding that inference must be normative.

¹⁶ For fuller discussion of this basic point, see Adler, *supra* note 3, ch. 1.

¹⁷ DAVID HUME, A TREATISE OF HUMAN NATURE, section 3.1.1 (L.A. Selby-Bigge ed., Oxford University Press, 1978).

¹⁸ Such a thought is suggested by the literature that seeks to infer a SWF from existing laws or policies. See, e.g., Peter J. Lambert et al., *Inequality Aversion and the Natural Rate of Subjective Inequality*, 87 JOURNAL OF PUBLIC ECONOMICS 1061 (2003); Benjamin B. Lockwood and Matthew C. Weinzierl, *Positive and Normative Judgments Implicit in U.S. Tax Policy, and the Costs of Unequal Growth and Recessions* (NBER Working Paper 21927) (January 2016); *id.* at 2 n.2 (citing sources).

This Article offers ethical advice, not in the form of an “ought” statement regarding a specific policy choice (“Enact the Clean Power Plan”...), but in the form of a very general and second-order “ought” statement—namely that, in the various policy-choice situations that may arise, regulatory decisionmakers ought to deliberate about which specific policies are ethically choice-worthy by employing the SWF framework. In reasoning to this second-order conclusion, I will repeatedly make ethical claims—sometimes wholeheartedly, sometimes in the more tentative form of describing some premise as “plausible.” The reader will have to decide for herself whether she accepts the claims: she will have to make an ethical “value judgment.”

A further question is institutional: Whatever the attractiveness of the SWF methodology as a framework for *ethical* deliberation, is it *legally* appropriate for regulatory officials to guide their choices using this methodology? This question implicates complicated issues about the intersection of law and ethics. I grapple with the question in Part III, as one part of the head-to-head ethical *and* legal comparison between the SWF approach and CBA. For now, I leave aside problems of legal authority and appropriateness, and focus on describing the approach; explaining where it sits in the landscape of ethical views; and summarizing the ethical case for specific types of SWFs (the utilitarian and prioritarian SWF).

Let’s turn, first, to a bird’s eye view of the ethical landscape. The SWF framework presupposes *weak consequentialism* together with *welfarism* concerning the goodness ranking of outcomes.

1. Weak Consequentialism

I’ll now introduce the concept of an “outcome.” Philosophers use the term “possible world” to mean a full description of a possible history of the world, from start to finish.¹⁹ Ethical decisionmakers with unbounded cognitive powers would be able to think about the “possible worlds” that might result from their decisions, and to deliberate about whether a given “world” would be better or worse than another. Human beings, of course, don’t have unbounded cognitive powers, and an explicit ranking of possible worlds cannot play a role in human ethical deliberation. Instead, what *is* feasible for humans is to think in terms of “outcomes.” An outcome is a cognitively tractable *model* of possible world. It is a stripped down—radically stripped down—description of a full possible world, highlighting some of the main features of ethical interest.”²⁰

“Consequentialists” endorse the following, very general methodology for ethical deliberation.²¹ The decisionmaker faced with a choice between options $\{a, b, c, \dots\}$ should have in mind *some* set of possible outcomes $\{x, y, z, \dots\}$. She should construct a “goodness ranking”

¹⁹ On possible worlds, see E.J. LOWE, A SURVEY OF METAPHYSICS ch. 5, 7 (2002).

²⁰ On outcomes as cognitively tractable models of whole possible worlds, see Adler, *supra* note 3, at 39-40, 241-58.

²¹ See *id.* at 22-32.

of that set of outcomes. And she should make her selection among the choices *in light of* the goodness ranking of the outcomes.

The goodness ranking of a set of outcomes has the following structure. Pick any two outcomes in the set, and either (1) the first is at least as good as the second, (2) the second is at least as good as the first; (3) both; or (4) neither, i.e., the outcomes are “incomparable.” Moreover, the “at least as good” relation is transitive: if x is at least as good as y , and y is at least as good as z , then x is at least as good as z .²² Finally the more intuitive notions of “better than” and “equally good as” can be defined in terms of the single relation, “at least as good as.”

A further simplifying assumption, often adopted in policy-analysis scholarship, is *completeness*. This rules out incomparability, option (4). Throughout the Article, my presentation of the SWF framework and its ethical foundations will assume completeness.

In short, for consequentialists, the “at least as good” relation between outcomes plays a foundational role in determining what “ought” to be done as an ethical matter.

We can now differentiate between two variants of consequentialism, “strong” and “weak.” Strong consequentialists believe that ethical requirements are *reducible* to the goodness ranking of outcomes. If the decisionmaker is selecting among a group of choices under conditions of certainty, then the ethically appropriate choice a^* is the choice with the best outcome. Under conditions of uncertainty, with each choice corresponding to a probability distribution across outcomes, there is some criterion for identifying the best choice a^+ in light of the goodness ranking of outcome plus this probability information; and the ethically appropriate choice is now a^+ .

Weak consequentialists believe that the goodness ranking of outcomes plays a foundational role in determining ethical requirements, but do not insist that ethical “oughts” are *wholly* determined by this ranking. In particular, weak consequentialists can affirm that ethics is a *combination* of the outcome ranking and non-consequentialist factors, such as rights or deontological constraints.²³

Much of the academic literature on the SWF is strongly consequentialist. However, the defense of the framework needs no more than weak consequentialism—and, to avoid controversy on this issue, nothing more than that will be presupposed here. The reader inclined to strong consequentialism, *if* persuaded that the SWF framework indeed helps us get a handle on the goodness ranking of outcomes, will see it as providing all-things-considered guidance concerning which policy choices are ethically recommended. The reader inclined to *weak*

²² The statement that x is better than y is equivalent to saying that x is at least as good as y , but y is not at least as good as x . The statement that x is equally good as y is equivalent to saying that x is at least as good as y , and y is at least as good as x .

²³ See Adler, *supra* note 3, at 24-29 (discussing deontological constraints and providing references to the philosophical literature). On non-consequentialist factors, see SHELLY KAGAN, *NORMATIVE ETHICS* (1998).

consequentialism will (if thus persuaded) see the framework as furnishing *prima facie* guidance with respect to “ethical” oughts—which guidance is then to be combined with rights, constraints, or whatever other non-consequential factors the reader finds appealing.

2. Welfarism

Welfarism says that the ethical goodness ranking of outcomes is reducible to the well-being of persons and, perhaps, other “welfare subjects” (sentient animals that are not persons but, still, can be characterized as having some level of welfare).²⁴ The case of non-person welfare subjects will be ignored here. Leaving aside this (important!) complication, welfarists say: whether outcome *x* is at least as good as outcome *y* is wholly determined by the well-being of the persons who exist in *x* and *y*. In particular: If each person’s well-being level in *x* is the same as her well-being level in *y*, the two outcomes are equally ethically good.

What is *well-being*?²⁵ Abstractly, we might say that well-being is *person-relative* goodness. To say that Claudia has more well-being in outcome *x* than *y* is just to say that outcome *x* is better *for her*—that *her* life goes better in the first outcome. Welfarism amounts to the (contested) normative position that moral goodness is determined by person-relative goodness. Given two outcomes *x* and *y*, we start with questions of the following sort—Is *x* at least as good as *y* from the perspective of person 1? Is *x* at least as good as *y* from the perspective of person 2? ... —and it is the answers to such questions that determine whether *x* is ethically at least as good as *y*, full stop.

A different issue—also an ethical, rather than purely factual issue—concerns the specific nature of well-being. Some ethicists believe that an individual’s welfare is a function of her pains, pleasures, and other hedonic or (more broadly) experiential states; some argue, instead, that welfare consists in preference-satisfaction; yet others adopt an “objective good” view of welfare. I’ll come back to this controversy below²⁶; welfarism as such is agnostic with respect to it.

Nonwelfarists *deny* that the outcome ranking is solely a function of well-being. They say: *x* may be better than *y* even though each person (and every other welfare subject) is equally well off in the two outcomes.

Various arguments can be mounted in favor of non-welfarism; the most powerful concerns individual *responsibility*.²⁷ Imagine that Charlie and Danny are badly off in *x*, and equally so. Charlie is badly off because he squandered his resources and education; Danny has

²⁴ See Adler, *supra* note 3, at 32-56; Andrew Moore and Roger Crisp, *Welfarism in Moral Theory*, 74 AUSTRALASIAN JOURNAL OF PHILOSOPHY 598 (1996); L.W. SUMNER, WELFARE, HAPPINESS AND ETHICS ch. 7 (1996); T.M. SCANLON, WHAT WE OWE TO EACH OTHER ch. 3 (1998).

²⁵ See sources cited *infra* note 40.

²⁶ See *infra* text accompanying notes 40-43.

²⁷ See Adler, *supra* note 3, at 33-39; Richard J. Arneson, *Welfare Should be the Currency of Justice*, 30 CANADIAN JOURNAL OF PHILOSOPHY 497 (2000).

tried hard to make a good life for himself, but through sheer bad luck his efforts have failed. We are faced with the choice of improving Charlie’s welfare by a certain amount or Danny’s by the same amount—assume we can’t help both. Whom should we aid? There’s a strong case to be made that outcome y (in which we aid Charlie) is worse than outcome y^* (in which we aid Danny). Surely Danny has a stronger claim to our assistance than Charlie. Yet welfarists will count y and y^* as equally good.

I take this to be a cogent objection to welfarism. In particular, the SWF framework (as presented in this Article) is vulnerable to the criticism that it ignores responsibility facts. However, the framework *can* be refined to take account of such facts. Indeed, there is a growing literature in economics—the formal and empirical literature on “equality of opportunity”—that does so.²⁸ Moreover, CBA also ignores responsibility facts. Moving from CBA to the SWF framework would be progress. From SWFs to responsibility-adjusted SWFs would be a further advance—but the details of that second step will have to be left for another day.

B. *The SWF Framework: Rules and Axioms*

The SWF framework is a formal expression of welfarism.²⁹ The framework combines weak consequentialism *plus* welfarism about the goodness ranking of outcomes *plus* a numerical indicator of well-being. The numbers assigned by the indicator *represent* the well-being facts. Each outcome corresponds to a list (“vector”) of well-being numbers, and the comparative goodness of any two outcomes can be summarized by a rule for ordering these “vectors.” This greatly facilitates systematic thinking about the “at least as good” relation.

As explained above in Section A, welfarism says that the combination of well-being facts about individuals in the population of interest is what determines whether a given outcome x is better than, worse than, or equally good as some other outcome y . The SWF framework adds to the mix the formal device of well-being measurement. $w(\cdot)$ denotes an indicator or measure of well-being, which assigns a number to each person in a given outcome. $w_i(x)$ is the well-being number of individual i in outcome x , $w_j(y)$ the well-being number of individual j in outcome y , and so forth.

These numbers *mirror* the well-being facts. In particular, they *mirror* intra- and interpersonal facts about well-being levels. If Ariel is better off in x as compared to y , then $w_{Ariel}(x) > w_{Ariel}(y)$. If Ariel in x is better off than Banksy in y , then $w_{Ariel}(x) > w_{Banksy}(y)$. The

²⁸ See Adler, *supra* note 3, at 579-84; Francisco H.G. Ferreira and Vito Peragine, *Individual Responsibility and Equality of Opportunity*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 746; John E. Roemer and Alain Trannoy, *Equality of Opportunity*, in HANDBOOK OF INCOME DISTRIBUTION 217 (Anthony Atkinson and Francois Bourguignon eds., vol. 2, 2015); John E. Roemer, *On Several Approaches to Equality of Opportunity*, 28 ECONOMICS AND PHILOSOPHY 165 (2012); Xavier Ramos and Dirk Van de Gaer, *Approaches to Inequality of Opportunity: Principles, Measures and Evidence*, 30 JOURNAL OF ECONOMIC SURVEYS 855 (2016).

²⁹ See Adler, *supra* note 3, at 61-78 (describing the SWF approach as a combination of a well-being measure and a rule for ranking well-being vectors).

numbers also mirror intra- and interpersonal facts about well-being differences. If the difference between Ariel’s well-being in x and y is greater than the difference between Ariel’s well-being in z and zz , then $w_{Ariel}(x) - w_{Ariel}(y) > w_{Ariel}(z) - w_{Ariel}(zz)$. And similarly for interpersonal difference comparisons.

But how are interpersonal comparisons to be made? And how, specifically, should a well-being indicator $w(\cdot)$ be constructed? See Section C below. First, though, let’s see what a well-being measure can accomplish.

Recall that “outcomes” are partly described possible realities. Concretely, an outcome will describe some of the major welfare-relevant attributes of individuals in the population. Policy analysts using tools grounded in modern economics (including CBA, the SWF framework, and other tools) tend to focus on attributes such as: an individual’s income, longevity, health state, amount of leisure, and/or public goods such as environmental quality.³⁰ An outcome, thus, will be a complicated, multidimensional object—in which each individual (or each homogeneous group of individuals) is assigned some lifespan, some income amount (or, with greater detail, an intertemporal sequence of income amounts for each period the individual is alive), some health state (or, with greater detail, an intertemporal sequence of health states), etc.

The well-being measure $w(\cdot)$ allows us to compress this mass of information into a list (“vector”) of numbers, one for each person in the population. Assume that there are three people in the population, Immanuel, Javier, and Kate. In a given outcome $(x, y, \text{etc.})$, each of these three individuals has some possible bundle of income, health, longevity, etc. If the w -number of Immanuel, Javier and Kate in x are, respectively, 3, 10, and 15, then x becomes the vector of numbers (3, 10, 15). If their w -numbers in y are, respectively, 4, 7, and 17, then y becomes the vector (4, 7, 17).

The vector (3, 10, 15) is a numerical representation of the well-being pattern in x , and the vector (4, 7, 17) a numerical representation of the well-being pattern in y . We can now determine the comparative ethical goodness of the two outcomes *via* a rule for comparing their corresponding well-being vectors. This is what the SWF framework does.

A bit more abstractly, let $(w_1(x), \dots, w_N(x))$ be the vector of well-being numbers corresponding to x , and $(w_1(y), \dots, w_N(y))$ the vector of well-being numbers corresponding to y . Let E be some rule for ranking well-being vectors. Then the SWF approach says: outcome x is at least as good as outcome y just in case $(w_1(x), \dots, w_N(x))$ is ranked by E at least as good as $(w_1(y), \dots, w_N(y))$.

What should be the rule E ? This is an ethical question. In specifying the rule E , we are taking a stance on an *ethical* issue: namely, how the pattern of well-being facts in any two given

³⁰ See *id.* at 243-45 (citing sources).

outcomes determines the “at least as good” relation between them. Indeed, I chose “*E*” as the abbreviation for the rule that ranks well-being vectors so as to remind the reader that this rule is one moving part within a (weakly consequentialist) framework for ethical deliberation, and that the choice of rule is a matter for ethical judgment.

But how should we think about the appropriate form of the rule *E*? Scholars in the SWF tradition do so by positing *axioms*: proposed constraints on the rule.³¹ When an ethical deliberator accepts a given axiom, she thereby eliminates rules inconsistent with the axiom. This is a powerful technique for deliberating about the rule *E*, since, first, a small list of very plausible axioms dramatically shrinks the “space” of possible rules; and, second, it is often easier to think in terms of axioms rather than directly in terms of rules.

What are some especially plausible axioms?³² These include the following: (1) *Ordering*. We have posited that the goodness ranking of outcomes should be complete and transitive, and in order for this to occur *E* will need to be complete and transitive. (2) *Anonymity*. Ethical assessment is impartial. Each person’s interests are given equal weight. This is captured in the “Anonymity” axiom. If the well-being numbers in *x* are a rearrangement of the well-being numbers in *y*—the same pattern, differing only in the identities of the individuals at the various well-being levels—then *x* and *y* are equally good. For example, let the first entry in the vector be Immanuel’s well-being level in an outcome, the second Javier’s, and the third Kate’s (as in the example above). So if *x* corresponds to the vector (3, 10, 15), and *y* the vector (10, 3, 15), Anonymity requires that *x* and *y* be ranked by *E* as equally good. (3) *Pareto*. Assume that everyone is at least as well off in *x* as compared to *y*, and at least one person is better off. Then *x* is better than *y*.³³ For example, the vector (8, 12, 3) is ranked above (7, 12, 3).

Two further axioms, although less compelling, are quite attractive on pragmatic grounds. They markedly facilitate the ease of using the SWF format as a deliberative tool, for reasons discussed in the notes.³⁴ One is (4) *Separability*. Let’s say that an individual is “affected by the *x/y* pair” if she is *not* equally well off in the two outcomes. It matters to her well-being whether *x* or *y* occurs. Separability then stipulates that the ranking of *x* and *y* is determined by the pattern of well-being among those affected by the *x/y* pair. A second pragmatic axiom is (5) *Continuity*. If the well-being vector for *x* is ranked better than the well-being vector for *y*, this is also true for all vectors sufficiently close to *x*.

³¹ See sources cited *supra* note 15.

³² For a discussion and defense of the axioms discussed in the remainder of this Section (Ordering, Anonymity, Pareto, Separability, Continuity, Pigou-Dalton), see Adler, *supra* note 3, at 52-56, 70-71, 307-78.

³³ The observant reader may note that I am toggling between stating the axioms as constraints on the ranking of outcomes, and stating them as constraints on the rule *E* for ranking well-being vectors. The latter statement is, strictly, more correct, but to avoid clunky articulation I’m doing both.

³⁴ A complete ordering of well-being vectors that satisfies Continuity can be represented by a continuous function. See EFE A. OK, REAL ANALYSIS WITH ECONOMIC APPLICATIONS 242 (2007). If the ordering also satisfies Separability, it can be given an additive representation. See Matthew D. Adler, *Justice, Claims and Prioritarianism: Room for Desert?* 31-33 (working paper, October 24, 2016), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2859553.

Putting together the axioms of Ordering, Anonymity, Pareto, Separability, and Continuity, we essentially eliminate all SWFs except the utilitarian SWF and the “prioritarian” family of SWFs.³⁵ The utilitarian SWF uses a straight summative rule. It assigns a given outcome x the sum total $w_1(x) + w_2(x) + \dots + w_N(x)$. Outcomes are ranked according to these simple sums. For example, if x corresponds to the well-being vector (3, 10, 15) and y the vector (7, 7, 13), the utilitarian SWF assigns x the number $3 + 10 + 15 = 28$, and y the number $7 + 7 + 13 = 27$, and says that x is the better outcome.

In order to understand the “prioritarian” SWF, we need the idea of an “increasing, concave” function. This is a mathematical function which ascends, but at a decreasing rate—as in Figure 1 below. I’ll denote such a function as $g(\cdot)$. A prioritarian SWF, instead of using the straight summative rule of utilitarianism, sums well-being numbers “transformed” by a strictly increasing and concave function. For example, if x corresponds to the well-being vector (3, 10, 15), and y the vector (7, 7, 13), the prioritarian SWF will use *some* strictly increasing and concave $g(\cdot)$ function; will assign x the sum $g(3) + g(10) + g(15)$; will assign y the sum $g(7) + g(7) + g(13)$; and will compare the outcomes by comparing these sums.

Since there are *many* (an infinity) of g -functions, the prioritarian SWF is a whole family of SWFs. One illustrative g -function is the square root. In that case, x is assigned the number $\sqrt{3} + \sqrt{10} + \sqrt{15} = 8.767$, and y the number $\sqrt{7} + \sqrt{7} + \sqrt{13} = 8.897$. Note that this reverses the utilitarian ranking. I’ll often use the term “prioritarian SWF,” but the reader should keep in mind that this is always a shorthand for “prioritarian family of SWFs.” See the notes on the question of how to identify a specific $g(\cdot)$ -function.³⁶

Both the utilitarian and prioritarian SWFs are rooted in philosophical scholarship. The idea of utilitarianism has, of course, been a mainstay in philosophical writing about ethics since


³⁵ More precisely, the following can be shown. If (a) the set of well-being vectors corresponding to the outcomes under consideration satisfies a “richness” requirement (namely, that this set is the N -fold Cartesian product of a nondegenerate interval of real numbers), and (b) the rule E produces a complete and transitive ordering of this set satisfying the Anonymity, Separability, and Continuity axioms just stated, then: E can be represented as the sum of individual well-being transformed by a single $h(\cdot)$ function. $((w_1(x), \dots, w_N(x))$ is ranked by E at least as good as $(w_1(y), \dots, w_N(y))$ iff $\sum_{i=1}^N h(w_i(x)) \geq \sum_{i=1}^N h(w_i(y))$. Because E satisfies Pareto, $h(\cdot)$ must be strictly increasing. See Adler, *supra* note 3, at 357 n.83; Adler, *supra* note 34, at 31-33.

In theory, $h(\cdot)$ can be any strictly increasing function. However, as a substantive philosophical matter, the only plausible possibilities are utilitarianism (in which case $h(\cdot)$ is the identity function, i.e., $h(w_i) = w_i$), or prioritarianism, in which case $h(\cdot)$ is strictly concave. As discussed in the text immediately below, the prioritarian E satisfies the axioms of Ordering, Anonymity, Pareto, Separability, and Continuity *and* Pigou-Dalton. Pigou-Dalton forces $h(\cdot)$ to be strictly concave. See Adler, *supra* note 34, at 31-33.

³⁶ If the rule E satisfies a very plausible axiom of ratio-rescaling-invariance (on top of Ordering, Anonymity, Pareto, Separability, Continuity and Pigou-Dalton), then the $g(\cdot)$ function must take the Atkinson/power form. $g(w_i) = (1-\gamma)^{-1}w_i^{1-\gamma}$, with $\gamma > 0$, $\gamma \neq 1$, or $g(w_i) = \ln w_i$, which is the special form of the Atkinson $g(\cdot)$ in the case of $\gamma = 1$. The parameter γ determines the degree of priority for the worse off. As γ increases, the worse off take greater priority; mathematically, $g(\cdot)$ becomes increasingly concave. The level of γ can be specified via ethical thought experiments, such as “leaky transfer” experiments, or by polling citizens for their views. See Adler, *supra* note 3, at 378-99.

Bentham.³⁷ Prioritarianism is considerably more recent, deriving from work by Derek Parfit from the 1990s,³⁸ but this is now a much-discussed view among academic philosophers.³⁹

Figure 1



³⁷ See, e.g., GEOFFREY SCARRE, *UTILITARIANISM* (1996); AMARTYA SEN AND BERNARD WILLIAMS, EDs., *UTILITARIANISM AND BEYOND* (1982).

³⁸ Derek Parfit, *Equality or Priority?*, in *THE IDEAL OF EQUALITY* 81 (Matthew Clayton and Andrew Williams eds., 2000) (delivered as the Lindley Lecture at the University of Kansas in 1991).

³⁹ See Adler, *supra* note 3, at 360-67 (reviewing and citing literature on prioritarianism). More recent works on prioritarianism include the articles in the symposium in *Utilitas* (vol. 24, no. 3, Sept. 2012). The symposium in *Economics and Philosophy* (vol. 31, no. 2, July 2015) includes the important then-unpublished articles by Fleurbaey, Hausman, and Broome, cited in Adler, *supra* note 3, at 360 n.90.

Explanation: Figure 1 illustrates a strictly increasing and concave $g(\cdot)$ function. Note that the Pigou-Dalton principle is satisfied. A transfer of Δw from a higher level (w_2) to a lower one (w_1) increases the sum of $g(\cdot)$ values.

The intuitive idea of prioritarianism is that well-being changes affecting worse off individuals have greater weight. This is captured, axiomatically, in the *Pigou-Dalton* axiom. Pigou-Dalton says that a pure, non-rank-switching transfer of well-being from a better-off to a worse-off individual, affecting no one else, is an ethical improvement. Assume that, in outcome x , one individual (Lower) is worse off than a second (Higher). In outcome y , Lower's well-being has increased by a certain amount, and Higher's has decreased by the same amount; Lower in y is still no better off than Higher; and everyone else has the same well-being level in x as she does in y . Then Pigou-Dalton stipulates that y is better than x .

Pigou-Dalton is the key axiomatic *difference* between the utilitarian and prioritarian SWFs. The utilitarian SWF does not satisfy the axiom, while Figure 1 illustrates that the rule of summing $g(\cdot)$ -transformed well-being numbers necessarily does. The Pigou-Dalton axiom strikes many as a powerful reason to choose prioritarianism. However, when we turn to the context of uncertainty, it will emerge that prioritarianism faces certain dilemmas, and these may push us back to utilitarianism.

To sum up: the SWF framework serves to structure deliberation about the “at least as good” relation. The master formula is: x at least as good as y iff the well-being vector corresponding to x is ranked by rule E at least as good as the well-being vector corresponding to rule y . We can make progress in specifying E by positing axioms that, we judge, it should satisfy. The combination of five quite plausible axioms—Ordering, Anonymity, Pareto, Separability, and Continuity—leaves us with only two possibilities for E : the utilitarian rule or the prioritarian family of rules. Prioritarian rules satisfy the five axioms just mentioned *plus* a sixth, the Pigou-Dalton axiom; the utilitarian rule satisfies these five *but not* Pigou-Dalton.

C. *The Well-Being Function*

Having discussed the SWF format, we can now drill down on the well-being measure $w(\cdot)$.

One question concerns the nature of well-being. This is a topic much debated by philosophers, with three main categories of views emerging from the debates.⁴⁰ *Preference*

⁴⁰ For overviews of the philosophical literature on well-being, see Richard Arneson, *Human Flourishing versus Desire Satisfaction*, in HUMAN FLOURISHING 113 (Ellen F. Paul et al. eds., 1999); Richard Arneson, *Desire Formation and Human Good*, in PREFERENCES AND WELL-BEING 9 (Serena Olsaretti ed., 2006); BEN BRADLEY, WELL-BEING (2015); JAMES GRIFFIN, WELL-BEING: ITS MEANING, MEASUREMENT, AND MORAL IMPORTANCE, chs. 1-4 (1986); Scanlon, *supra* note 24, at 108-43; Sumner, *supra* note 24; GUY FLETCHER, ED., THE ROUTLEDGE HANDBOOK OF PHILOSOPHY OF WELL-BEING (2015).

accounts analyze well-being in terms of an individual's actual or idealized preferences.⁴¹ Preferentialists then engage in intramural disagreement about the specific type of preference that is the basis for well-being. *Experientialists* affirm the following: If someone has the very same mental states in two outcomes, then she is equally well off in the two. Experientialists then engage in intramural disagreement about the specific type or types of mental states (pains and pleasures, feelings of happiness, cognitions, perceptions, emotions, etc.) that are the basis for well-being.⁴²

Experientialist and preference views of well-being are distinct. An individual (actually or ideally) can have preferences over outcomes that are partly driven by features other than her own mental states: for example, her physical health, the quality of her relationships with friends and families, her state of knowledge, how much freedom she has, or her success in accomplishing career goals.

The third type of well-being account is objective.⁴³ An objective account sets forth a list of goods, which are (a) posited as enhancing an individual's well-being, independent of whether the individual (actually or ideally prefers them), and are such that (b) not all of the goods are purely mentalistic.⁴⁴

Each of these families of welfare views is compatible with the SWF format. On each given view, the well-being measure $w(\cdot)$ is such as to embody intra- and interpersonal comparisons of well-being levels and differences *according to the view*. For example, $w_{Sofia}(x) > w_{Sofia}(y)$ indicates that Sofia's well-being level in outcome x is greater than her well-being level in y . If a preference account of well-being is adopted, this means—more specifically—that Sofia has an (actual or idealized) preference for x over y . On an experientialist account, this means that Sofia's package of mental states in x is better than her package of mental states in y .

⁴¹ On preference views of well-being, see Adler, *supra* note 3, at 159-62; Chris Heathwood, *Desire-Fulfillment Theory*, in THE ROUTLEDGE HANDBOOK OF PHILOSOPHY OF WELL-BEING, *supra* note 40, at 135; Krister Bykvist, *Preference-Based Views of Well-Being*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 321.

⁴² The contrast between experientialist and preference-based views of well-being stated here is fleshed out in Matthew D. Adler, *Happiness Surveys and Public Policy: What's the Use?*, 62 DUKE L.J. 1509, 1521-38 (2013). For a general discussion of experientialist views of well-being, see Adler, *supra* note 3, at 162-65; Daniel Haybron, *Mental State Approaches to Well-Being*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 347 (2016); Alex Gregory, *Hedonism*, in THE ROUTLEDGE HANDBOOK OF PHILOSOPHY OF WELL-BEING, *supra* note 40.

⁴³ See Adler, *supra* note 3, at 165-70; Thomas Hurka, *Objective Goods*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 379; Guy Fletcher, *Objective List Theories*, in THE ROUTLEDGE HANDBOOK OF PHILOSOPHY OF WELL-BEING, *supra* note 40, at 148.

⁴⁴ Hybrid views, which mix experientialist, preference-based, and/or objective elements are also possible. See Christopher Woodward, *Hybrid Theories*, in THE ROUTLEDGE HANDBOOK OF PHILOSOPHY OF WELL-BEING, *supra* note 40, at 161.

(according to the criterion for ranking such bundles set forth by the account).⁴⁵ On an objective account, this means that Sofia has a better bundle of objective goods in x .

For the remainder of this Article, I generally adopt a preference view of well-being. I have elsewhere defended such a view—specifically, that well-being should be analyzed in terms of preferences that are “idealized” in the sense of being well-informed, rational, and self-interested.⁴⁶ Moreover, the preference view remains the dominant approach in welfare economics⁴⁷ (notwithstanding the rise of happiness economics)⁴⁸, and thus the bulk of extant scholarship on SWFs does equate well-being with preference-satisfaction. Such a view, finally, is the basis for CBA as standardly understood—and so adopting a preference-based rather than experientialist of objective account of well-being will facilitate the comparison of the SWF approach with CBA.

Let us now turn to the so-called “informational” content of the well-being measure $w(\cdot)$.⁴⁹ I have already stated repeatedly that $w(\cdot)$ represents both intra- and interpersonal well-being information, concerning both well-being levels and well-being differences. Why insist that $w(\cdot)$ contain information about well-being *differences* (not merely well-being levels), and indeed that it contain *interpersonal* information about levels and differences, not merely intrapersonal information?

Table 1 below illustrates why $w(\cdot)$ needs to provide information about well-being *differences*, not merely well-being levels. That table displays a group of outcomes $\{x, y\}$ and persons (Able and Barry), with well-being numbers assigned to the persons using two different well-being measures $w(\cdot)$ and $w^*(\cdot)$. As can be seen in the table, these two measures contain the very same information about well-being levels, but different information about well-being differences. If information about well-being levels were *sufficient* for the SWF format, the ranking of the outcomes should be *invariant* to whether $w(\cdot)$ or $w^*(\cdot)$ is used. But Table 1 illustrates that the ranking of $\{x, y\}$, by both the utilitarian and prioritarian SWFs, is different using the two well-being measures. This shows that both utilitarianism and prioritarianism require information about well-being differences—not merely well-being levels.

⁴⁵ John Bronsteen, Christopher Buccafusco and Jonathan Masur have proposed that CBA be replaced with “well-being analysis” (WBA), which is effectively a utilitarian SWF applied to a hedonic measure of well-being.

⁴⁶ See Adler, *supra* note 3, ch. 3.

⁴⁷ For a canonical statement of economic theory in which preferences are central, see ANDREU MAS-COLLEL, MICHAEL D. WHINSTON AND JERRY R. GREEN, *MICROECONOMIC THEORY* (1995).

⁴⁸ On happiness economics, see Carol Graham, *Subjective Well-Being in Economics*, in *THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY*, *supra* note 7, at 424.

⁴⁹ On this topic, see especially Bossert and Weymark, *supra* note 15.

Table 1

	Well-being numbers according to $w(\cdot)$		According to $w^*(\cdot)$	
	<u>Outcome x</u>	<u>y</u>	<u>Outcome x</u>	<u>y</u>
Able	25	16	3.5	3
Barry	4	9	1	2
Utilitarian SWF	29	25	4.5	5
Prioritarian SWF	7	7	2.87	3.15

(with $g(\cdot) = \sqrt{\cdot}$)

In a parallel fashion, Table 2 shows that *intrapersonal* well-being information is insufficient for both SWFs. In this case, $w(\cdot)$ and $w^+(\cdot)$ are identical in terms of their within-person comparisons of both levels and differences. However, the two measures make different interpersonal comparisons. Table 2 shows that neither the utilitarian SWF nor the prioritarian SWF is invariant to swapping $w(\cdot)$ for $w^+(\cdot)$.

Table 2

	Well-being numbers according to $w(\cdot)$			According to $w^+(\cdot)$		
	<u>x</u>	<u>y</u>	<u>z</u>	<u>x</u>	<u>y</u>	<u>z</u>
Able	49	64	36	4.9	6.4	3.6
Barry	16	9	25	16	9	25
Utilitarian SWF	65	73	61	20.9	15.4	28.6
Prioritarian SWF	11	11	11	6.21	5.53	6.90

(with $g(\cdot) = \sqrt{\cdot}$)

We now face a serious puzzle. On a preference view of well-being, intrapersonal comparisons of well-being levels are reasonably straightforward. But *interpersonal* comparisons of well-being levels, and comparisons of well-being differences (intra- or interpersonal), are more puzzling. Sofia, for example, (actually or ideally) prefers x to y to z to w , while Tim prefers y to x to w to z . We now ask, “Is Sofia in x better off, worse off, or equally well off as Tim in w ?” How does the data about Tim’s and Sofia’s preference rankings enable us to answer this question (concerning their comparative well-being levels)? Similarly, how does this data enable us to answer a question about well-being differences, be it intrapersonal (“Is the difference between Sofia’s well-being in x and y greater than, less than, or equal to the difference in her well-being between z and w ?”) or interpersonal (“Is the difference between Tim’s well-

being in y and x greater, less than, or equal to the difference between Sofia's well-being in z and w ?)?)?

Indeed, there is a long intellectual history in economics of skepticism about difference comparisons⁵⁰ and interpersonal comparisons⁵¹; and some contemporary economists, outside the SWF school, continue to voice such skepticism.

A plausible response to this skepticism—pioneered in theoretical work by John Harsanyi⁵²—is to look to *risk* preferences. Assume first (this should not be controversial) that a given individual's preferences over the outcomes depends upon her *attributes* in the outcomes—her income, health, leisure, public goods, etc. Sofia prefers x to y if she prefers her attribute bundle in x to her attribute bundle in y . Assume, next, that the individual has a ranking, not only of attribute bundles, but of lotteries over such bundles, and that this ranking is represented by a so-called von Neumann/Morgenstern (vNM) utility function $u(\cdot)$.⁵³ She prefers one lottery to a second if the first has greater expected utility (as calculated using the vNM function $u(\cdot)$).

How might we use vNM utility functions to generate inter- and intrapersonal comparisons of well-being levels and differences? The easiest case is where individuals are assumed to have common preferences over attribute bundles and bundle lotteries.⁵⁴ In the case of common preferences over attribute bundles and bundle lotteries, it is straightforward to use the common vNM utility function to generate intra- and interpersonal comparisons of levels and differences, as illustrated in Table 3. The reason this is successful is because of a well-known technical feature of vNM utility functions: they are unique up to linear transformations, i.e., are “cardinal” (thus difference comparisons) and not merely “ordinal.”

⁵⁰ See Peter H. Farquhar and L. Robin Keller, *Preference Intensity Measurement*, 19 ANNALS OF OPERATION RESEARCH 205 (1989) (describing skepticism about difference comparisons).

⁵¹ See Adler, *supra* note 3, at 185-87.

⁵² JOHN HARSANYI, RATIONAL BEHAVIOR AND BARGAINING EQUILIBRIUM IN GAMES AND SOCIAL SITUATIONS, ch. 4 (1977). I discuss Harsanyi's account of well-being measurement in Adler, *supra* note 3, at 192-202; Matthew D. Adler, *Extended Preferences and Interpersonal Comparisons: A New Account*, 30 ECONOMICS AND PHILOSOPHY 123 (2014); and Matthew D. Adler, *Extended Preferences*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 480-84 (hereafter *Extended Preferences*).

⁵³ On vNM utility functions, see ITZHAK GILBOA, THEORY OF DECISION UNDER UNCERTAINTY (2009); DAVID KREPS, NOTES ON THE THEORY OF CHOICE (1988).

⁵⁴ Note that this does not imply common individual rankings of *outcomes* (which would be absurd). For example: Sheila and Tom have identical rankings of health-income bundles, and lotteries over such bundles, but Sheila prefers x to y since Sheila's health-income bundle in x is better (according to the common bundle ranking) than Sheila's health-income bundle in y ; while Tom prefers y to x since Tom's health-income bundle in y is better.

Table 3

Attributes		vNM utility				
		original function		linear transformation		
	<u>Outcome x</u>	<u>Outcome y</u>	<u>x</u>	<u>y</u>	<u>x</u>	<u>y</u>
Elijah	\$90K, 50yrs	\$50K, 60 yrs	570.4	649.2	7.0	14.9
Mia	\$30K, 80 yrs	\$70K, 70 yrs	824.7	780.9	32.5	28.1

	Elijah- x	Elijah- y	Mia- x	Mia- y
Elijah- x	0	-78.8	-254.3	-210.6
Elijah- y	78.8	0	-175.5	-131.8
Mia- x	254.3	175.5	0	43.8
Mia- y	210.6	131.8	-43.8	0

	Elijah- x	Elijah- y	Mia- x	Mia- y
Elijah- x	0	-7.9	-25.4	-21.1
Elijah- y	7.9	0	-17.6	-13.2
Mia- x	25.4	17.6	0	4.4
Mia- y	21.1	13.2	-4.4	0

Explanation: In each of two outcomes (x and y), each of two individuals (Elijah and Mia) has an attribute bundle consisting of an annual income and a lifespan. A vNM utility number is assigned to that bundle using the formula $u(\text{income, lifespan}) = \ln(\text{income}) \times \text{lifespan}$. Consider now a linear transformation $u^+(\cdot)$ of that original function, defined as follows: $u^+(\text{income, lifespan}) = (1/10)u(\text{income, lifespan}) - 50$. The utility numbers according to the two functions are displayed in the top right.

There are four “histories” (a combination of an individual and an outcome): Eijah- x , Elijah- y , Mia- x , and Mia- y . It can be seen by inspection that the utility levels of these 4 histories, as per vNM function $u(\cdot)$, have the same order as the utility levels of these four histories as per vNM function $u^+(\cdot)$.

The bottom tables show the utility differences between all 16 ordered pairs of histories. The left table does this for function $u(\cdot)$ and the right table for function $u^+(\cdot)$, with all values rounded to one decimal place. It can be seen that the ordering of these 16 differences is unchanged. Indeed, it’s clear mathematically (and visible by comparing the two tables) that the difference between a pair according to $u^+(\cdot)$ will be one-tenth the difference according to $u(\cdot)$, so that the ordering of the 16 differences must be the same for the two functions.

The premise that all individuals in the population have common preferences over attribute bundles and lotteries, and thus a common, cardinal, VNM utility function, is fairly standard in the SWF literature. This premise is admittedly unrealistic—empirical research into preferences calls it into question—but might be defended as just one instance of the use of simplified models that is pervasive in both positive and normative economics. Still, it would be good to have a methodology for making well-being comparisons that is robust to heterogeneity in individual rankings of attribute bundles and bundle lotteries.

Such a methodology is on hand. In very brief compass, the idea is this. Assume a set \mathbf{A} of attribute bundles $\{a, a^*, a^{**}, \dots\}$. A preference, formally, is a ranking of \mathbf{A} and lotteries over \mathbf{A} . Let's allow for heterogeneous preferences. The symbol R denotes one such preference; R^+ another; and so forth. Each such preference is represented by a vNM utility function. We achieve a common scaling of these utility functions by identifying two bundles, “high” and “low,” and scaling all the functions so that they have the same value at the high bundle, and the same value at the low bundle. The well-being number of a given person, in a given outcome, is a function of *two* things: her attribute bundle (be it a, a^*, a^{**}, \dots) and her preference (be it R, R^+, R^{++}, \dots). Specifically, it is the utility number assigned to that bundle by the rescaled vNM utility function corresponding to that preference.⁵⁵

In conclusion, it is worth noting that difference comparisons and interpersonal comparisons (of levels or differences) are *not* at odds with ordinary talk about well-being. Such comparisons are made all the time. For example, we might say that Philly stands to benefit more from some health intervention than Rupert (an interpersonal difference comparison); or that Amy, who is unemployed, has low income, and is depressed, is much worse off than Bob, who has a stable job, high income, and good mental health (an interpersonal level comparison). Indeed (as a matter of ordinary discourse) the refusal to make any such comparisons will seem pedantic or (in a clear case, like Bob and Amy) outrageous. The challenge is to provide a rigorous foundation for the kinds of well-being judgments that are already in play as a matter of ordinary ethical talk—and the Harsanyi-esque proposal to do so by looking to risk preferences and vNM utility functions is one plausible approach.⁵⁶

D. *Uncertainty*

The SWF framework, as discussed to this point, has two key components: a well-being measure $w(\cdot)$ that translates each outcome into a vector of well-being numbers⁵⁷; and a rule E for ranking these well-being vectors. Under conditions of certainty, these two components are sufficient to provide choice-guidance to a decisionmaker. But with uncertainty added to the picture, they are not. In the uncertainty case, each choice corresponds to a probability

⁵⁵ See Adler, *Extended Preferences*, *supra* note 52. This “extended preferences” approach to interpersonal comparisons is a further refinement of the account presented in Adler, *supra* note 3, ch. 3. As explained in *Extended Preferences*, high-low scaling is only one version of the approach. For an empirical illustration, see Koen Decancq and Dirk Neumann, *Does the Choice of Well-Being Measure Matter Empirically?*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 588.

⁵⁶ A measure of well-being that reflects intra- and interpersonal comparisons of well-being levels and differences is sufficient for utilitarianism, but not prioritarianism. See Adler, *supra* note 3, at 341-42. However, information about intra- and interpersonal comparisons of well-being levels, differences and *ratios* is sufficient for the prioritarian SWF of the Atkinson form described *supra* note 36. We construct ratio comparison by identifying a zero point. See Adler, *supra* note 3, at 216-21, 379-82; Adler and Treich, *supra* note 6, at 289-93. For an empirical illustration of the significance of the zero point within prioritarianism, see Matthew Adler et al., *Priority for the Worse Off and the Social Cost of Carbon* (CESifo working paper no. 6032) (August 2016).

⁵⁷ The discussion of uncertainty that follows does not presuppose a preference-based account of well-being. I assume only that $w(\cdot)$ measures well-being according to whichever view is adopted, be it preference-based, experientialist, or objective.

distribution across outcomes.⁵⁸ With $w(\cdot)$ in hand, each choice translates into a probability distribution across well-being vectors—and an “expected well-being” number for each person in the population. We now need some procedure for ranking choices in light of this probabilistic well-being information.

Three such procedures suggest themselves: utilitarianism-under-uncertainty (UUU), “ex post prioritarianism” (EPP) and “ex ante prioritarianism” (EAP). Each of the three assigns a numerical value to any given choice; the set of choices is then ranked correspondingly. The three procedures differ in how numbers are attached to choices.⁵⁹ See Table 4. Mathematical formulas for the three approaches are provided in the Appendix.

If we’ve adopted the utilitarian rule for ranking outcomes, the straightforward procedure under uncertainty is to rank choices in light of *the expected sum of individual well-being*. This is UUU. Each of the possible outcomes of a given choice is assigned a utilitarian value by summing up the individual well-being numbers in that outcome. This outcome value is then probabilistically discounted, i.e., multiplied by the probability of the outcome. The UUU value of the choice is, in turn, the sum total of the probabilistically discounted utilitarian values of its possible outcomes.

The second procedure, EPP, is structurally similar to UUU, but differs in using the prioritarian rather than utilitarian formula for assigning a value to outcomes. EPP orders choices in light of *the expected sum of transformed individual well-being*. Each of the possible outcomes of a given choice is assigned a *prioritarian* value by summing up the $g(\cdot)$ values in that outcome—the individual well-being numbers, each transformed by the prioritarian function $g(\cdot)$. This outcome value is then probabilistically discounted, i.e., multiplied by the probability of the outcome. The EPP value of the choice is, in turn, the sum total of the probabilistically discounted *prioritarian* values of its possible outcomes.

The third procedure, EAP, is built upon the vector of expected well-being numbers associated with any given choice. This information about individuals’ well-being expectations is

⁵⁸ A more general treatment would allow for the absence of precise probabilities (sometimes termed “ambiguity”). See Adler, *supra* note 3, at 542-47. Space constraints preclude a discussion of this topic here.

⁵⁹ For overviews of the literature on applying an SWF under uncertainty, see Adler, *supra* note 3, at 477-530; Philippe Mongin and Marcus Pivato, *Social Evaluation under Risk and Uncertainty*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 711. On the closely related issue of Harsanyi’s so-called aggregation theorem, see John Weymark, *A Reconsideration of the Harsanyi-Sen Debate on Utilitarianism*, in INTERPERSONAL COMPARISONS OF WELL-BEING (Jon Elster and John Roemer eds., 1991); Blackorby, Bossert and Donaldson, *supra* note 15, at 222-33; Mongin and d’Aspremont, *supra* note 15, at 425-44. Note that Adler, *supra*, uses the term “EU prioritarianism” rather than “ex post prioritarianism” to refer to EPP.

Two other approaches closely related to UUU and EPP apply a non-linear and strictly increasing function to the sum of well-being or transformed well-being, and maximize the expected value of this function. See Adler, *supra* note 3, at 493-94, 518-24; Adler and Treich, *supra* note 6, at 301-02; Adler, Hammitt, and Treich, *supra* note 9; *infra* Appendix. Such approaches are worth consideration, but on balance—I believe—are less attractive than UUU and EPP, and are not further discussed here. Their key flaw is a failure to respect an axiom of ex ante separability: if some individuals are sure to be unaffected by a choice, then their levels of well-being shouldn’t matter.

adjusted in a prioritarian manner. Specifically, EAP takes each expected well-being number; transforms that number using the prioritarian transformation function $g(\cdot)$; and sums up. This sum is the EAP value of the choice. In short, EAP orders choices in light of the *sum of transformed expected individual well-being*.

Table 4

	<u>Choice A</u>		<u>Choice B</u>			
	State s ($p=.5$)	<u>State s^*</u> ($p=.5$)	Expected well-being	State s ($p=.5$)	<u>State s^*</u> ($p=.5$)	Expected well-being
Juan	100	9	54.5	36	64	50
Sarah	25	81	53	36	64	50
$\sum w_i$	125	90		72	128	
$\sum \sqrt{w_i}$	15	12		12	16	

	<u>Choice A</u>	<u>Choice B</u>	<u>Preferred Choice</u>
UUU value (expected $\sum w_i$)	107.5	100	A
EPP value (expected $\sum \sqrt{w_i}$)	13.5	14	B
EAP value ($\sum \sqrt{\text{exp. } w_i}$)	14.7	14.1	A

Much ink has been spilt on the choice between UUU, EPP, and EAP (as well as other approaches to SWFs under uncertainty that some take to be plausible).⁶⁰ Without belaboring that discussion, it seems that there are compelling reasons to *reject* EAP.⁶¹ One is the axiom of stochastic dominance. EAP might recommend choice *a* over choice *b* even though *every possible outcome of choice a is worse than every possible outcome of choice b*. It seems deeply problematic for a consequentialist to endorse a framework for ethical deliberation that violates the stochastic-dominance axiom. Two other arguments against EAP rest upon the so-called “sure thing” axiom and an axiom of time-consistency.

⁶⁰ See Adler, *supra* note 3, at 497 n.40 (citing sources). Other contributions include Michael Otsuka and Alex Voorhoeve, *Why it Matters that Some are Worse off than Others: An Argument against the Priority View*, 37 PHILOSOPHY AND PUBLIC AFFAIRS 171 (2009); Luc Bovens and Marc Fleurbaey, *Evaluating Life or Death Prospects*, 28 ECONOMICS AND PHILOSOPHY 217 (2012); Hilary Greaves, *Antiprioritarianism*, 27 UTILITAS 1 (2015); Luc Bovens, *Concerns for the Poorly Off in Ordering Risky Prospects*, 31 ECONOMICS AND PHILOSOPHY 397 (2015); David McCarthy, *The Priority View*, ECONOMICS AND PHILOSOPHY (forthcoming); David McCarthy, *Distributive Equality*, 124 MIND 1045 (2015).

⁶¹ See Adler, *supra* note 3, at 506-18; Matthew D. Adler and Chris Sanchirico, *Inequality and Uncertainty: Theory and Legal Applications*, 155 U. PA. L. REV. 279, 344-47 (2006).

This leaves the choice between UUU and EPP. This is a genuinely hard choice. Not surprisingly, one key difference concerns the Pigou-Dalton principle. EPP, but not UUU, satisfies an axiom of *Equity Dominance*: If two choices a and b are such that every possible⁶² outcome of a has the same total well-being as every possible outcome of b , but more equally distributed, then a is the better choice. “More equal” is understood, here, in terms of the Pigou-Dalton principle: well-being is more equally distributed in outcome x than y if x can be reached from y by a series of pure well-being transfers from someone better off to someone worse off.⁶³

Those who find the Pigou-Dalton principle an attractive axiom with respect to the goodness ranking of outcomes will naturally gravitate towards Equity Dominance and thus EPP. On the other hand, UUU but not EPP satisfies the *Ex Ante Pareto Principle*: If each person’s expected well-being with choice a is at least as large as her expected well-being with choice b , and at least one person has greater expected well-being with a , a is a better choice.⁶⁴

Economists may find it astonishing that a methodology for policy assessment would reject the ex ante Pareto principle. Isn’t the Pareto principle a core commitment of welfarism? But we should distinguish between the Pareto principle as a constraint on the ranking of outcomes, which has intense normative appeal—How can it be worse to make the world better for some, and worse for none?—and the ex ante principle. Well-being expectations incorporate imperfect information: an individual might in fact be better off with choice b than a but, lacking that information, have greater *expected utility* with a . Relatedly, choice a might be such that the ethical decisionmaker (i) knows for certain someone will be worse off with a rather than b , and (ii) cannot be sure that anyone will be better off with a , and yet the ex ante Pareto principle favors a over b .⁶⁵ Why should she feel compelled to choose a in this sort of case?

Further, the ex ante Pareto principle has logical implications that some may find unacceptable, in particular a conflict with the principle of Equity Dominance. If the ethical deliberator finds Equity Dominance to be the more compelling of the two axioms, then she

⁶² Strictly, possible and with non-zero probability.

⁶³ If two outcomes x and y are such that they have the same total well-being, but x can be reached from y by a series of Pigou-Dalton transfers, then the utilitarian SWF will be indifferent between the two, while the prioritarian SWF will prefer x (because that SWF satisfies the Pigou-Dalton principle and transitivity). The following is thus immediate from the formulas for EPP and UUU (as given in the Appendix): if policies P and P^* are such that, in every state of nature with non-zero probability, the outcome of P can be reached from that of P^* by a series of Pigou-Dalton transfers, then: P is preferred to P^* by EPP, but UUU is indifferent. Call the axiom that requires a preference for P “Statewise Equity Dominance.” Statewise Equity Dominance in turn implies Equity Dominance as stated in the text.

⁶⁴ Because UUU is equivalent to the sum of individual expected well-being (see *infra* Appendix), it is immediate that UUU satisfies ex ante Pareto. On the inconsistency between EPP and ex ante Pareto, see Adler, *supra* note 3, at 497-506.

⁶⁵ For an illustration, assume that there are three equally probable states of nature (see *infra* Appendix) and two individuals affected by the choice, Jim and June. Choice a is such that, in one state, Jim gets a well-being level of 15 and June 1; in a second state, it’s June who gets 15 and Jim 1; in the third state, both receive 1. Choice b is such that both individuals are sure to receive 3. With choice a , whatever the state of nature, one of the two will receive less than 3, and there’s a 1/3 probability that both will.

should conclude that she must tolerate violations of ex ante Pareto. Table 5 illustrates the conflict between the ex ante Pareto principle and Equity Dominance.

Some readers will be unpersuaded by these arguments. Others will embrace Equity Dominance and reject ex ante Pareto. Again, the choice between UUU and EPP is a hard one; and I will keep both in play for the remainder of this Article in my development of the SWF framework as a regulatory tool.

Table 5

	<u>Choice A</u>			<u>Choice B</u>		
	<u>State s</u> (p=.5)	<u>State s*</u> (p=.5)	Expected well-being	<u>State s</u> (p=.5)	<u>State s*</u> (p=.5)	Expected well-being
Able	90	10	50	$50-\epsilon$	$50-\epsilon$	$50-\epsilon$
Baker	10	90	50	$50-\epsilon$	$50-\epsilon$	$50-\epsilon$

Explanation: By Pigou-Dalton, the well-being vectors (90, 10) and (10, 90) are worse than (50, 50). Thus Equity Dominance would require choice B over A if the well-being vector in each state for B were (50, 50). If the ranking of actions, further, satisfies a continuity requirement, then B as stated here must be better than A for ϵ sufficiently small. But the ex ante Pareto principle prefers Choice A.

II. The SWF Framework as a Guide to Regulatory Policy

The SWF framework, as set forth in Part I, has the widest scope of application. It can be used to provide guidance for any kind of choice, and in particular for any kind of decision regarding governmental policy. Indeed, as mentioned in the Introduction, SWFs are the basis for scholarly work by economists in numerous policy areas.

The specific focus of this Article is governmental regulation, with environmental, health and safety regulation (a.k.a. “risk regulation”) taken as the central case—since this is where CBA has been most significant in U.S. governmental practice. Section A explains how SWFs can inform risk regulation. Section B reviews the structure of CBA. Section C uses a stylized case study to illustrate SWFs, CBA, and the contrast between the two. The aim of these sections is expository, not argumentative. The SWF framework and CBA *both* provide guidance in balancing the money costs of risk regulation (lost income for members of the population of interest) as against the benefits (improved health states and longer life expectancies); but the two methodologies *differ* in how benefits and costs are quantified and commensurated. Once the reader is clear about these differences, she’ll be in a position to adjudicate the normative considerations pro and con the two methodologies—to assess whether the SWF methodology is indeed a better regulatory calculus than CBA, as I’ll argue in Part III below.

A. *The SWF Framework as a Basis for Environmental, Health and Safety Regulation*

Going forward, I assume that the measure $w(\cdot)$ of well-being is based upon vNM utilities; that the rule E is either utilitarian or prioritarian; and that the procedure under uncertainty (corresponding to these two rules) is either UUU or EPP.

One important aspect of the well-being measure, not yet stressed but now relevant, is that it quantifies an individual's *lifetime well-being*. The number $w_{Rhonda}(x)$ is a numerical indicator of the overall quality of Rhonda's life in x , from birth to death, as determined by its length and by what occurs to Rhonda in each of its time slices. Lifetime well-being is the appropriate input into the SWF, since personal identity normally continues over a whole human lifetime⁶⁶—and indeed most of the literature on SWFs (explicitly or implicitly) uses a lifetime-well-being indicator.⁶⁷

But what, now, are the features that should be incorporated in outcomes so that each individual can be assigned a lifetime well-being level? Recall that an outcome is a *model* of a possible world—a stripped-down description thereof, setting forth *some* of the characteristics of ethical interest. But what *are* the characteristics that should be part of this stripped-down description, in the specific context of environmental, health and safety regulation?

As mentioned earlier, an individual's well-being is determined by her *attributes*. We can think of an outcome as assigning bundles of attributes to individuals, attributes of two kinds: *welfare* attributes, namely the individual attributes that directly determine lifetime well-being; and *further* attributes, namely additional types of individual attributes that are useful to incorporate in the outcome-description, for reasons to be discussed momentarily.

Let's start with *welfare attributes*. One such determinant, surely, is individual longevity. Each outcome should include information about the length-of-life of each person in the population of interest. We can't possibly assess someone's lifetime well-being without knowing her lifespan! Further, lifespan information is of key relevance for risk regulation. A fatality cuts short an individual's well-being by truncating her lifespan; fatality *risks*, such as environmental toxins, dangerous products, and so forth, are *risks* to someone's well-being in virtue of raising the chance of a shorter lifespan and, thus, a lower level of lifetime welfare.

Economic modelling almost invariably takes an individual's *consumption* (the value at market prices of the goods and services she uses) or, as a proxy for consumption, her *income*, as

⁶⁶ Adler, *supra* note 3, ch. 6.

⁶⁷ See *id.* at 245, 441. For a review of the structure of lifetime utility and the application of SWFs to lifetime utilities, see Gregory Ponthiere, *Lifetime Well-Being*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 871. See also the literature on optimal taxation in a dynamic perspective, reviewed in Golosov and Tsyvinski, *supra* note 5. A recent illustrative example is Jean-Baptiste Michau, *Optimal Redistribution: A Life-Cycle Perspective*, 111 J. PUB. ECON. 1 (2014).

a key welfare-relevant attribute.⁶⁸ Thus an outcome should at least describe, for each person in the population of interest, not only the lifespan of that person, but also her income/consumption during each period alive.

Income information, along with information about mortality rates, is most readily available on an *annual* basis. So the simplest periodization of outcomes is in terms of *years*.

Thus far, we have that an outcome should at least describe the lifespan of each person in the population of interest and her income/consumption during each year alive. Further, since environmental, health and safety risks constitute threats to individual *health* as well as lifespan, it would be useful if outcomes contained some information about individuals' health states across their lifetime (most tractably, in the form of a "QALY"-like measure of health quality in each year)⁶⁹.

Let's call this combination of information—longevity, income/consumption in each year alive, and health status each year alive—an LHI history. "LHI" stands for "longevity, health income." The welfare attributes of a given person, in a given outcome, should consist in at least⁷⁰ (1) her LHI history in that outcome, and (2) information about her preferences, sufficient to assign her a vNM utility function. With this combination of data about any given person, we have sufficient information to assign that person a lifetime well-being level, as per formula (1) immediately below.

Let $T_i(x)$ denote the lifespan of individual i in outcome x . Let $c_i^t(x)$ be the income/consumption of individual i in outcome x during her t -th year of life. And let $h_i^t(x)$ be her health quality in x during the t -th year of life. Finally, let $u_i(\cdot)$ be the rescaled vNM utility function of individual i . Then the well-being of individual i in outcome x is as follows.

$$(1) w_i(x) = u_i(c_i^1(x), \dots, c_i^{T_i(x)}(x); h_i^1(x), \dots, h_i^{T_i(x)}(x)).$$

We can simplify this formula in various ways. First, economists typically assume that lifetime utility is the sum of period (annual) utility.⁷¹ This greatly eases the calculation of lifetime utility. Second, if preference heterogeneity is ignored, the individual-specific subscript

⁶⁸ See Adler, *supra* note 3, at 243-45.

⁶⁹ On QALY measures of health, see Jose-Luis Pinto-Prades, Carmen Herrero, and Jose Maria Abellan, *QALY-Based Cost-Effectiveness Analysis*, in THE OXFORD HANDBOOK OF WELL-BEING AND PUBLIC POLICY, *supra* note 7, at 170-75. To be clear, what is being advocated in this Article is not cost-effectiveness analysis with health-quality-adjusted longevity as the "effectiveness" metric (the standard use of QALYs for policy assessment, see *id.*). Rather, a QALY-type measure of someone's health quality in each year, along with income in each year, and longevity, would determine that person's lifetime utility for purposes of an SWF. On health-income interactions in producing utility, see sources cited in James Hammitt, Kevin Haninger, and Nicolas Treich, *Effects of Health and Longevity on Financial Risk Tolerance*, 34 GENEVA RISK AND INSURANCE REVIEW 117, 122-23 (2009).

⁷⁰ Other attributes might also be added, for example individual labor hours. Leisure/labor attributes are a mainstay of the use of SWFs to determine optimal tax policies. See sources cited *supra* note 5.

⁷¹ Adler, *supra* note 3, at 284-97; Ponthiere, *supra* note 67, at 876-80.

from the utility function in formula (1) can be dropped. We can reframe that formula (or the time-additive version thereof) in terms of a single utility function representing the common preferences.

To repeat, an individual's *welfare* attributes in an outcome (for purposes of risk regulation) consist of her LHI history in that outcome plus her preferences (or, more simply, her LHI history, if common preferences are assumed). What about her *further* attributes? These are either (a) attributes that have some causal role in producing LHI histories (for example, demographic characteristics such as gender, educational attainment, occupation, place of birth or residence; biological or psychological attributes; environmental conditions; and so forth); or (b) attributes that function to identify an individual (her name, date of birth, SSN, DNA signature, etc.).

Let's now shift from *outcomes* to *choices*. A "choice," in the context at hand, is some regulatory intervention open to the decisionmaker. This includes the baseline or "status quo" choice of doing nothing. Each choice (as discussed in Part I) is modelled as a probability distribution across outcomes. We assign a value to each choice using either the UUU or EPP formula.

But now a bit of magic will be performed.⁷² A choice is not merely a probability distribution across outcomes. Since each outcome is a combination of LHI histories for each individual in the population, each choice *a* is *also* a probability distribution ("lottery") across LHI histories for each such individual.

Further, since UUU and EPP are equivalent to (respectively) the sum of individual expected well-being or individual expected transformed well-being, these procedures have a very useful invariance feature. Let two choices *a* and *a** be such that each choice gives rise to the very same lottery over LHI histories for person 1; each choice gives rise to the very same lottery over LHI histories for person 2; and so forth for everyone in the population. Then UUU assigns the very same value to *a* and *a**; and EPP assigns the very same value to *a* and *a**.

In short, rather than evaluating policy choices by thinking of them as lotteries over whole outcomes, we can instead do so by thinking of them as lotteries over LHI histories for the individuals of ethical concern.

And now to conclude the magic trick: Although the SWF format is meant to inform a process of ethical deliberation, whereby the decisionmaker takes on a stance of impartial concern regarding each and every person in some population of interest, the SWF methodology (with UUU and EPP) doesn't actually *require* the explicit use of identifying attributes such as name,

⁷² See *infra* Appendix.

SSN, or DNA signature. CBA analysts certainly don't refer to specific individuals by name or SSN in the course of their cost-benefit analyses; and the SWF analyst doesn't need to either.

Why not? Let's define a *cohort* as a group of individuals who (1) share some cluster of defining attributes (welfare attributes and/or further attributes) that are *exogenous* to the regulatory choices under consideration,⁷³ and (2) have, or at least are treated for modeling purposes as having, the same preferences. And let a *cohort structure* be a collection of cohorts that *partitions* the population of interest, i.e., each person falls into one and only one cohort. Finally, assume that we have at hand some causal model whereby regulatory choices are mapped onto lotteries over LHI histories for each cohort, as a function of the choices and the welfare/further attributes of cohort members. Then we have sufficient information to assign each regulatory choice a UUU or EPP value. We don't need a list of names, SSNs, or other unique identifiers that pick out each person in the population.

This isn't to say that no such identifying characteristics exist. They surely do, and in many other contexts governments rely upon them—for example, in implementing tax and benefits programs. It's just to say that the SWF format (at least UUU and EPP) can circumvent the explicit use of individuals' identifying attributes.

So what, then, *is* a plausible cohort structure for purposes of environmental, health and safety regulation? It would be foolhardy for me to offer a definitive answer to this question. It depends upon the available information and causal models whereby clusters of attributes defining a candidate cohort might be mapped onto lotteries over LHI histories. Still, a few tentative suggestions might be hazarded.

Imagine a regulatory decision, as between the status quo and various policy interventions, undertaken at a particular point in time (during year Y^*). What might be the cohort structure for purposes of this decision?

First, *year of birth* (or, equivalently, age at the beginning of year Y^*) should be one of the defining attributes of a cohort. That is to say, a cohort should be identical to, or a subset of, what demographers call a "birth cohort." Why? Each birth cohort has its own probabilities of surviving to various lifespans—its own "survival curve."⁷⁴ Moreover, even if the survival curves at birth of different birth cohorts are identical, the probability of surviving to a given age of those born in year Y^* itself is different from the probability of surviving to that age of those who were born in $Y^* - 1$ and have survived to year Y^* ; which in turn is different from the probability of

⁷³ The idea of a cohort structure is that members of each cohort can be assigned a probability distribution over LHI histories, for every policy under consideration. These probabilities are meant to guide the decisionmaker's choice, and thus to be independent of which policy she picks. Therefore the attributes that define each cohort need to be exogenous; they can't vary depending upon the policy.

⁷⁴ Cohorts and survival curves are central concepts in the field of demography. See SAMUEL H. PRESTON, PATRICK HEUVELINE, AND MICHEL GUILLOT, *DEMOGRAPHY* (2001).

surviving to that age of those who were born in year $Y^* - 2$ and have survived to year Y^* ; and so forth. Thus information about year of birth will usefully *differentiate* cohorts, for purposes of assessing lifespan probabilities. We can form a more accurate assessment of someone's current probability of living any given lifespan if we know her current age.

For similar reasons, *prior-year income* and some marker of *prior-year health quality* should likely be among the defining attributes of a cohort.⁷⁵ We can form a more accurate assessment of someone's current probability of having a particular income this year or in any future year if we know her prior-year income, and analogously for health quality.

Thus a cohort should, plausibly, be defined as all individuals with a common (a) year of birth, (b) prior-year income or income tranche; (c) prior-year health quality or health quality tranche; and (d) some cluster of additional attributes (demographic, physical, environmental, etc. attributes, or perhaps income or health attributes further in the past than the prior year).

Each regulatory choice $\{s, a, b, c, d \dots\}$ defines a lottery over LHI histories for cohort members. To flesh this out: we have the baseline choice of doing nothing, denoted as “s” for status quo, and various regulatory interventions. Each such intervention can be seen as a *perturbation* to each cohort's baseline lottery over LHI histories.

The cohort in the baseline is characterized by a survival curve, given that it has survived to year Y^* : a probability of surviving to the end of year Y^* , of surviving to the end of year $Y^* + 1$, of surviving to the end of year $Y^* + 2$, and so forth. The cohort also faces some income-generating process (as determined by prior-year income and other cohort attributes), which defines a probability distribution over income amounts in Y^* for cohort members; a probability distribution over income amounts in year Y^*+1 for those who survive to the beginning of year Y^*+1 ; a probability distribution over income amounts in year Y^*+2 for those who survive to the beginning of that year; and so forth. Finally, the cohort faces some health-generating process (as determined by prior-year *health* and other attributes) defining a probability distribution over health levels in Y^* ; a probability distribution over health levels in $Y^* + 1$ for those who survive to the beginning of that year; etc.

A regulatory policy *a* will reduce risks: specifically, by improving the cohort's survival curve. It may also improve the probability distribution over health quality states in Y^* and future years. But the policy will also have costs, which take the form of perturbing the income-generating process going forward (most simply, reducing income in Y^* or subsequent years by some specific amount). By virtue of the combination of these fatality-risk, health, and income

⁷⁵ A yet fuller and more complicated cohort structure would individuate cohorts by income and health quality for multiple years prior to the current year—at the limit, back to birth.

changes, a now corresponds to a different lottery over LHI histories for cohort members than the baseline option s .⁷⁶

B. *Cost-Benefit Analysis (CBA)*

There is now a voluminous literature on CBA, both in general and as applied to environmental, health, and safety regulation,⁷⁷ as well as many years of governmental practice. A lengthy description of CBA is, thus, not needed here. What follows is a capsule summary—just sufficient to review the key features of the technique, which are well documented and elaborated in the cited sources.

CBA in governmental practice differs from the textbook version.⁷⁸ I'll henceforth refer to “textbook CBA” and “governmental practice” CBA. CBA without a qualifier means textbook CBA.

CBA evaluates policy choices by summing individuals' *monetary equivalents*. An individual's monetary equivalent for a policy, broadly speaking, is the change to her money

⁷⁶ The conceptualization of regulatory choices here, as perturbations to lotteries over LHI histories for various cohorts, which are then assessed by a SWF, builds from many literatures: (1) demography, with its focus on the evolution, over a lifetime, of the period (annual) survival probabilities for individuals in a given cohort, see SAMUEL H. PRESTON, PATRICK HEUVELINE, AND MICHEL GUILLOT, *DEMOGRAPHY* (2001); (2) the economics literature on the lifetime evolution of income and consumption, see, e.g., Orazio Attanasio and Guglielmo Weber, *Consumption and Savings: Models of Intertemporal Allocation and their Implications for Public Policy*, 48 J. ECON. LIT. 693 (2010); Mark Aguiar and Erik Hurst, *Deconstructing Life-Cycle Expenditure*, 121 J. POL. ECON. 437 (2013); Ponthiere, *supra* note 67, at 880-83; (3) the parallel literature on the evolution of individual health over a lifetime, see, e.g., Michael Kuhn et al., *Optimal Choice of Health and Retirement in a Life-Cycle Model*, 158 J. ECON. THEORY 186 (2015), Ponthiere, *supra* note 67, at 888-91; (4) the massive literature on the relation between income and health/longevity, see, e.g., Raj Chetty et al., *The Association between Income and Life Expectancy in the United States, 2001-2014*, 315 JAMA 1750 (2016); David Cutler Angus Deaton, and Adriana Lleras-Muney, *The Determinants of Mortality*, 20 J. ECON. PERSP. 97 (2006); sources cited in Hans van Kippersluis et al., *Socioeconomic Differences in Health over the Life Cycle in an Egalitarian Country* 5-7 (Tinbergen Institute Discussion Paper, January 2009); (5) scholarship on the structure of lifetime utility, see Ponthiere, *supra* note 67, at 876-80; Adler, *supra* note 3, at 245, 418-22; (6) inequality metrics applied to lifetime income, longevity, or life expectancy, see, e.g., Jeremy Arkes, *Trends in Long-Run Versus Cross-Section Earnings Inequality in the 1970s and 1980s*, 44 REV. INCOME AND WEALTH 199 (1998); Mark Huggett et al., *Sources of Lifetime Inequality*, 101 AMERICAN ECON. REV. 2923 (2011); Sam Pelzman, *Mortality Inequality*, 23 J. ECON. PERSP. 175 (2009); Ryan D. Edwards and Shripad Tujapurkar, *Inequality in Life Spans and a New Perspective on Mortality Convergence across Industrialized Countries*, 31 POP. AND DEVELOPMENT REV. 645 (2005)); and finally (7) scholarship that applies an SWF to lifetime utilities, see sources cited *supra* note 67. Richard Cookson and co-authors (including myself) have advocated using LHI histories for health policy. See Richard Cookson et al., *Years of Good Life Based on Income and Health* (CHE Working Paper 132, University of York, July 2016). For an earlier version of the approach set forth in this Article, see Matthew D. Adler, *Risk Equity: A New Proposal*, 32 HARV. ENV'T'L. L. REV. 1 (2008).

⁷⁷ See, e.g., MATTHEW D. ADLER AND ERIC POSNER, *NEW FOUNDATIONS OF COST-BENEFIT ANALYSIS* (2006); A. MYRICK FREEMAN III, *THE MEASUREMENT OF ENVIRONMENTAL AND RESOURCE VALUES* (2d ed. 2003); PER-OLOV JOHANSSON AND BENGT KRISTROM, *COST-BENEFIT ANALYSIS FOR PROJECT APPRAISAL* (2016); RICHARD E. JUST, DARRELL L. HUETH, AND ANDREW SCHMITZ, *THE WELFARE ECONOMICS OF PUBLIC POLICY: A PRACTICAL APPROACH TO PROJECT AND POLICY EVALUATION* (2004); ROBIN W. BOADWAY AND NEIL BRUCE, *WELFARE ECONOMICS* (1984); Boadway, *Cost-Benefit Analysis*, *supra* note 12.

⁷⁸ See *infra* text accompanying notes 81-82.

holdings that produces a change to her well-being (as defined by her preferences) which is equal in magnitude to the policy's impact on her.⁷⁹

More precisely, “monetary equivalents” can be specified as so-called “equivalent variations” or “compensating variations.” Equivalent variations are formally better behaved than compensating variations,⁸⁰ and so my summary will be pitched in terms of equivalent variations (henceforth “EVs”).

Consider a set of regulatory options $\{s, a, b, c, \dots\}$, with s as above the status quo option of inaction. And consider some individual in the population, say, Chloe. Chloe's EV for policy a is the change to Chloe's current income in the status quo such that Chloe is indifferent between the status quo with her income altered by that amount, and policy a . Note that if Chloe prefers policy a to the status quo, then her EV for policy a is positive; her status quo current income would need to be *increased* for her to reach a point of indifference as between the status quo and a . If Chloe prefers the status quo to policy a then her EV for policy a is negative; her status quo current income would need to be decreased for her to reach a point of indifference between the status quo and a . Finally, if Chloe is already indifferent between the status quo and a , her EV is zero.

In discussing the SWF approach, I conceptualized each regulatory choice as a probability distribution across outcomes. CBA conceptualizes choices the very same way. Status quo option s gives rise to one probability distribution across outcomes; some other option a gives rise to a different distribution. Chloe's EV for option a is such that, if we were to “tweak” the status quo choice s as follows—we increase Chloe's current income in every outcome by the EV, and hold constant the probabilities of outcomes given s —Chloe would now be indifferent between the status quo and choice a .

Chloe's preferences (as standard for the preference view of well-being) are summarized in a vNM utility function $u_{\text{Chloe}}(\cdot)$. So we can say that Chloe's EV for option a is such that, if her status quo current income were increased by this amount, her *expected utility* (the expected value of $u_{\text{Chloe}}(\cdot)$) for the status quo s and for option a would be the same.

All this rehearses the concept of monetary equivalent (EV). CBA's rule for ranking choices is simple. Each choice is assigned a value equaling the sum of EVs, for all individuals in the population of interest; the choices are ranked in the order of these values.

We can now turn to the specific context of environmental, health and safety regulation. In the status quo, Chloe has a particular current-year survival probability (her probability of surviving to the end of the year); she faces some probability distribution over current-year income; and, finally, she has some probability distribution over current-year health states. A

⁷⁹ See *infra* Appendix for a formal definition of money equivalents.

⁸⁰ See Adler, *supra* note 3, at 95-98; Adler, *supra* note 12, on-line supplementary materials, at <http://reep.oxfordjournals.org/content/10/2/264/suppl/DC1>.

particular regulatory option a changes Chloe's current-year survival probability, probability distribution over current income amounts, and/or probability distribution over current health states. Chloe's EV for a is the change to Chloe's current-year income in s that just suffices to compensate her for the differences (in current income, survival probabilities, and health) between s and a .

This generalizes to the case in which option a changes Chloe's survival probabilities, income, or health in *later* years. We can *still* ask about the change to Chloe's *current* income that counterbalances the variety of differences (in either current or subsequent income, survival risk and health) as between s and a .

Actually, CBA can be readily expressed in terms of the concept of LHI histories that I used earlier to set forth the SWF framework. (This will be useful, below, in comparing and contrasting the two methodologies). In the status quo, Chloe faces one lottery over LHI histories; with regulatory option a , she faces a different lottery. Chloe's EV for a is such that, if her current-year income in every history in the status quo lottery were increased by that amount, she would now be indifferent between the two lotteries. Her expected utility for the s lottery and the a lottery would be the same.

Also, the same move from individuals to cohorts, described above with reference to the SWF methodology, also occurs (implicitly) in CBA practice. Monetary equivalents are not literally assigned to individuals on an individual-by-individual basis. Nor need they be. Assume Chloe is part of some cohort of individuals with the same preferences and the same probabilities over LHI histories, for each of the regulatory choices under consideration $\{s, a, b, c, \dots\}$. Then everyone in the cohort will have the same EV for option a , option b , option c , etc.

A basic feature of textbook CBA is that valuations are *heterogeneous*. Consider a policy that improves the survival probability for members of cohort C by a certain amount (relative to the status quo); that improves the survival probability for members of cohort C^* by the same amount; and that has no other effect on these individuals. Even though the reductions in fatality risk are the same, the EV for each member of the first cohort may well be different from the EV for each member of the second. Why? The cohorts might have different preferences. Moreover, even assuming common preferences, a given person's EV for a risk reduction is a function not only of the risk reduction, but also the individual's baseline income, survival probabilities, and health state. The cohorts may well differ in these respects.

I distinguished above between "textbook" and "governmental practice" CBA, and now can explain the difference. Governmental practice CBA, in the area of risk regulation, generally works as follows. First, policy impacts are *decomposed* with respect to the three dimensions of cost, fatality risk, and health. That is, a given policy a will be assigned an overall monetary amount equaling the sum of (1) the total "costs" of a (its total effect on individual incomes); (2) its total monetized risk-reduction benefit; and (3) its total monetized health-improvement

benefit.⁸¹ Second, in monetizing the risk-reduction and health-improvement benefits from a policy, governmental-practice CBA generally uses *population-average* valuations rather than valuations specific to the individuals being benefitted.⁸²

For example, assume that a policy (among other effects) improves the survival probabilities of members of various cohorts 1 through M . Cohort 1 members receive a reduction in fatality risk of Δp_1 each, cohort 2 members a reduction of Δp_2 each, and so forth. There are N_1 individuals in cohort 1, N_2 in cohort 2, and so forth. Note that the total risk reduction experienced by the entire population is $N_1 \times \Delta p_1 + \dots + N_M \times \Delta p_M$. Governmental-practice CBA will assign the policy a monetized risk-reduction benefit by taking this total reduction ($N_1 \times \Delta p_1 + \dots + N_M \times \Delta p_M$) and multiplying through by the population-average willingness to pay per unit of risk reduction (that is, the population average “value of statistical life”). By contrast, the risk-reduction benefit of the policy as per textbook CBA is equal to the total risk reduction of the first cohort multiplied by that group’s willingness to pay per unit of risk reduction, plus the total reduction of the second cohort multiplied by that group’s willingness to pay per unit of risk reduction, and so forth. In general, these two sums are *not* the same.

C. *A Worked Example*

This section uses a simulation model to illustrate the differences between the SWF format, in either its utilitarian or prioritarian form, and CBA. Specifically, I consider UUU; EPP with a moderate degree of priority for the worse off (denoted EPP^{mod}); EPP with a higher degree of priority for the worse off (EPP^{high}); textbook CBA; and governmental-practice CBA.

The population is divided into cohorts defined by age and income. In the baseline, a given age-income cohort faces a survival curve—which states the probability of surviving the current year, the probability of surviving the year after that, and so forth. These survival curves are calibrated to the official U.S. government data about survival probabilities for the entire U.S. population, and then adjusted to take account of the effect of income on survival curves: higher-income groups have more favorable survival curves and, thus, higher life expectancies.

In order to simplify the simulation and facilitate exposition, I assume that the individuals affected by the policy choices at hand fall into five age groups: those currently aged 20, 30, 40, 50 and 60. I assume five different income strata, with income amounts corresponding to the 10th, 30th, 50th, 70th, and 90th percentiles of the U.S. distribution of (after tax) income. This yields twenty-five cohorts, formed by combining each of the five age groups with the five levels of annual income. I assume an equal number of individuals in each of the twenty-five cohorts.

⁸¹ See EPA, *Economy-Wide Modelling: Benefits of Air Quality Improvements White Paper* (Sept. 22 2015), at p. 6-14 (describing EPA’s “effect-by-effect” approach to CBA).

⁸² See Lisa A. Robinson, *How US Government Agencies Value Mortality Risk Reductions*, 1 REVIEW OF ENVIRONMENTAL ECONOMICS AND POLICY 283 (2007).

For short, I refer to the five levels of income as “10th,” “30th,” “50th,” “70th” and “90th” percentile, although to be clear these are estimated percentiles in the actual U.S. population; in the simulated population used for the case, these are the *only* five income levels. I also refer to the levels as “bottom,” “lower,” “middle,” “higher,” and “top” income, respectively.

I assume common preferences, and ignore health differences between the cohorts other than differences in fatality risks. Thus, a LHI history takes the form of lifespan and an income amount each year while alive. Lifetime utility is assumed to be the sum of annual utility, which in turn is the logarithm of income. Note that members of each given cohort face a common probability distribution over LHI histories, arising from the common cohort survival curve and common income amounts earned in each year of survival if alive.

See the Appendix for a more detailed description of the construction of the model.

1. Uniform Risk Reduction and Cost Incidence

Consider, now, a regulatory policy that produces a population-wide reduction in current mortality risk of 1-in-100,000, and (2) “pays” for this benefit by reducing each person’s current income by some uniform $\$D$ (the same amount for each person, regardless of income). For short, let’s refer to this policy as “Uniform Risk Reduction and Uniform Cost Incidence.” We can then calculate, for each methodology, the *break-even* amount such that the policy is seen as beneficial if $\$D$ is less than this break-even amount.

These break-even (or “maximum”) amounts are displayed in Table 6 below.

Table 6: Uniform Risk Reduction and Uniform Cost Incidence

UUU	EPP ^{mod}	EPP ^{high}	Textbook CBA	Govt. CBA
\$48	\$59	\$78	\$91	\$91

Consider, first, the difference between the UUU maximum (\$48) and the CBA maximum (\$91), for both textbook and governmental practice CBA. While CBA and utilitarianism are sometimes conflated, this example underscores that the methodologies are not identical and can, in practice, produce quite different recommendations.

To understand why UUU and CBA diverge (dramatically) in this case, it’s helpful to think of the two methodologies as using distinct *units of ethical value*. In the case of UUU, the unit of ethical value is expected well-being: a policy is assigned a number equaling the sum of

individual expected well-being.⁸³ For CBA, by contrast, the unit of ethical value is the monetary equivalent (EV): a policy is assigned a number equaling the sum of EVs.

Table 7A displays how UUU assigns ethical value to risk reductions incurred by the various cohorts. The numbers in Table 7A are normalized so that 1 means the value of a 1-in-100,000 risk reduction for a 60-year-old in the bottom income group. Note that the 60-year-old/middle income cohort is assigned a value of 1.6. This means that a 1-in-100,000 risk reduction for a member of that group produces an increase in UUU value (the sum of individual expected well-being) that is 1.6 times the increase in UUU value produced by the same risk reduction for a member of the 60 year old/bottom income cohort.

Table 7B does the same exercise, but now for CBA. Again, the numbers in Table 7B are normalized so that 1 means the value of a 1-in-100,000 risk reduction for 60 year olds in the bottom income group. Note that the 60-year-old/middle income cohort is assigned a value of 3.9. This means that a 1-in-100,000 risk reduction for members of that group produces an increase in CBA value (the sum of individual EVs) that is 3.9 times the increase in CBA value produced by the same risk reduction for members of the 60 year old/bottom income cohort.

Table 7A: UUU Risk Reduction Values

	<i>Income percentile 10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	2.8	3.5	4.0	4.5	5.7
<i>30</i>	2.3	2.9	3.4	3.8	4.8
<i>40</i>	1.8	2.4	2.7	3.1	3.9
<i>50</i>	1.4	1.8	2.1	2.4	3.1
<i>60</i>	1.0	1.3	1.6	1.8	2.3

Table 7B: CBA Risk Reduction Values

	<i>Income percentile 10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	2.8	6.2	9.9	15.6	44.6
<i>30</i>	2.3	5.2	8.3	13.1	37.8
<i>40</i>	1.8	4.2	6.7	10.7	31.0
<i>50</i>	1.4	3.2	5.2	8.4	24.5
<i>60</i>	1.0	2.3	3.9	6.2	18.4

What leaps out from the comparison of Tables 7A and 7B is first, that UUU is biased towards those with higher incomes in valuing risk reduction; but, second, that CBA is substantially *more* biased than UUU. In Table 7A, as we moved rightwards (up the income ladder) within each row, the UUU value of a risk reduction increases; the same occurs for CBA values in Table 7B; but the rate of increase in the second table is considerably greater than in the first. For example, consider 40-year-olds: UUU says that a 1-in-100,000 risk reduction for a 40 year old/bottom income individual yields 1.8 times the ethical value of that same risk reduction

⁸³ Again, UUU can be expressed as the expected sum of individual well-being or, equivalently, as the sum of individuals' expected well-being. See *infra* Appendix. Thus expected well-being, summed across individuals, is the unit of ethical value for UUU.

for a 60-year-old/bottom income person; moving up the income ladder among 40 year olds, the numbers are 2.4, 2.7, 3.1, and 3.9. for lower, middle, higher, and top-income individuals, respectively,. By contrast, CBA says that a 1-in-100,000 risk reduction for a 40 year old/bottom income individual yields 1.8 times the ethical value of that reduction for a 60 year old/bottom income individual; moving up the income ladder among 40 year olds, the numbers are 4.2, 6.7, 10.7, and 31.0.

UUU is biased towards the rich in valuing risk reduction because the well-being of each year alive increases with income. Thus, increasing the probability that a higher-income person of a given age will live an incremental year creates a bigger increase in expected individual well-being than increasing the probability that a lower-income person will live that incremental year. But CBA is *more* biased towards the rich than UUU because individual monetary equivalents are inflated by the declining marginal utility of money. UUU translates a risk reduction for a given individual into a change in her expected well-being. CBA translates a risk reduction for a given individual into her EV. The EV, in turn, can be shown to be equal to the change in the individual's expected well-being, thence translated into money at a conversion rate determined by the individual's marginal utility of income.⁸⁴ Because the marginal utility of income is decreasing, the very same change in expected well-being produces a larger EV for a richer individual. This is what explains the different patterns across the rows in Table 7A as compared to 7B—why CBA is *more* biased to the rich than UUU.

A second difference between UUU and CBA concerns the way in which units of ethical value are assigned to reductions in individual income. While Table 7A displays the comparative ethical value (as per UUU) of risk reductions conferred upon different cohorts, Table 8A displays the comparative ethical value (as per UUU) of increases in present income (by \$1) conferred upon different cohorts. These numbers are normalized so as to be consistent with Table 7A. That is, the number 1 indicates the change in expected well-being that occurs with a 1-in-100,000 reduction in the risk of a member of the 60-year-old/bottom income cohort. Thus, in Table 8A, the numbers are *multiples* of this change. For example, in Table 8A, those in the 50-year-old/middle income cohort have the number 0.05. This means that an increase in present income by \$1 for a 50-year-old/middle income individual produces a change in individual expected well-being that is 0.05 the magnitude of the change that occurs by reducing the mortality risk of a 60-year-old/bottom income individual by 1-in-100,000.

⁸⁴ See *infra* Appendix, showing how EV for a risk reduction is approximately equal to VSL times the risk reduction—with VSL equaling the change in expected well-being per unit of risk reduction *divided by* the change in expected well-being per unit of consumption.

Table 8A: UUU Values of Increased Income

	<i>Income percentile</i>				
	<i>10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	.12	.07	.05	.03	.02
<i>30</i>	.12	.07	.05	.03	.02
<i>40</i>	.12	.07	.05	.03	.02
<i>50</i>	.12	.07	.05	.03	.02
<i>60</i>	.12	.07	.05	.03	.02

Table 8B: CBA Values of Increased Income

	<i>Income percentile</i>				
	<i>10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	.12	.12	.12	.12	.12
<i>30</i>	.12	.12	.12	.12	.12
<i>40</i>	.12	.12	.12	.12	.12
<i>50</i>	.12	.12	.12	.12	.12
<i>60</i>	.12	.12	.12	.12	.12

Table 8B provides analogous information, but for CBA. The numbers indicate the changes in CBA ethical value (the sum of EVs) that occur with increases in current income (by \$1) for the members of the various cohorts. Table 8B is synched to Table 7B, so that the number 1 indicates the change in the sum of EVs that occurs with a 1-in-100,000 risk reduction for a 60-year-old/bottom income person.

Note that the numbers in Table 8B do not change! For CBA purposes, a “dollar is a dollar is a dollar.” Anyone’s EV for increasing her income by \$1 is, well, \$1.⁸⁵ Thus every number in Table 8B is 0.12, which simply equals the ratio between (a) the EV for increasing the current income of the member of any cohort by \$1, i.e., \$1; and (b) the EV of a 60-year-old/bottom income individual for reducing her risk by 1-in-100,000, which is \$8.33.

By contrast, the numbers in Table 8A vary among cohorts—specifically, they *decrease* as we move rightwards in each row, moving up the income ladder for each age group. The difference between the row-wise decreasing pattern for UUU in Table 8A, as compared to the constant pattern for CBA in Table 8B, is—once more—the result of the declining marginal utility of money. A change in a richer individual’s income produces a *smaller* change in expected individual well-being than a change in a poorer individual’s income.

Let’s return to where we began, with the policy of Uniform Risk Reduction and Uniform Cost Incidence. The numbers in Tables 7A, 7B, 8A and 8B not only give general insight into how UUU and CBA assign comparative ethical value to risk reduction and income reduction incurred by individuals in various cohorts. They also, more specifically, can be used to derive the UUU and CBA maximum dollar reductions (\$48 versus \$91) and to understand why these differ. The ethical value, according to UUU, of a 1-in-100,000 risk reduction for the members of every cohort is just the sum of the values in Table 7A (scaled up to account for the size of the population). The ethical value of a \$D reduction in the current incomes of the members of every

⁸⁵ This is true given the approximating formula for EVs that has been used for all the calculations in this Section, but isn’t true if EVs are exactly measured. However, the sensitivity of CBA to cost incidence, even with exact EV measurement, is approximately nil if all policies are relatively close to the status quo. See *infra* Appendix for the approximating formula and a more detailed discussion of CBA’s insensitivity to cost incidence.

cohort is just $-\$D$ times the sum of the values in Table 8A (again scaled up for population size). We can (approximately)⁸⁶ identify the break-even point as that value $\$D$ which, multiplied by the sum of the values in Table 8A, just equals the sum of the values in Table 7A. This is \$48. The parallel exercise for CBA (using now Tables 7B and 8B) yields \$91.⁸⁷

Let's now turn to the EPP case. While UUU sums individuals' expected well-being values, and CBA sums their monetary equivalents (EVs), EPP sums individuals' expected transformed well-being values. Thus the unit of ethical value, for EPP, is expected transformed well-being. With this in mind, we can produce analogues of Tables 7A, 7B, 8A and 8B for EPP. I do so here for EPP^{high}; the tables for EPP^{mod} are in the Appendix.

Table 7C shows how EPP^{high} assigns ethical value to risk reduction for the various cohorts. We'll again use 1 to denote the ethical value (here, increase in expected transformed well-being) produced by a 1-in-100,000 risk reduction for a 60-year-old/bottom income individual. Thus, for example, the number 1.4 in Table 7C for the 50-year-old/middle income individual means that a 1-in-100,000 risk reduction for such an individual produces an increase in the sum of individual expected transformed well-being that is 1.4 times the increase in that sum produced by the very same reduction for a 60-year-old/bottom income person.

Table 8C shows how EPP^{high} assigns ethical value to income increases for the various cohorts. This table is harmonized with Table 7C by using 1 to mean the same thing in both tables (the change in ethical value for a 1-in-100,000 risk reduction for a 60 year old/bottom income individual), just as we did above for the corresponding UUU and CBA tables.

⁸⁶ We are using here the total-differential approximation from calculus—as applied here to mean that the change in UUU value for a pattern of risk reductions and income changes in a population is approximately the sum across the individuals of the UUU change per unit of risk reduction for each individual multiplied by the number of units of risk change for her, plus the UUU change per unit of income change for each individual multiplied by the units of income change for her. (The total differential approximation to a continuously differentiable real-valued function is discussed in any good calculus textbook. See, e.g., C.H. EDWARDS, JR., *ADVANCED CALCULUS OF SEVERAL VARIABLES* 63-76 (1994).) The use of the total-differential approximation to describe how SWFs value risk reduction is undertaken in Adler, Hammitt, and Treich, *supra* note 9. The “social value of risk reduction,” as described there, is the derivative of the SWF with respect to survival probability.

⁸⁷ The reader might observe that the CBA values for risk reduction in each cohort, relative to UUU, are adjusted for the marginal utility of income for that cohort; and that the CBA values for income change in each cohort, relative to UUU, are adjusted in the same way. So why don't these perfectly cancel out, yielding the very same $\$D$ for the policy of Uniform Risk Reduction and Cost Incidence? However, since the adjustments are cohort-specific, there's no reason mathematically to think they will cancel out.

Table 7C: EPP^{high} Risk Reduction Values Table 8C: EPP^{high} Values of Increased Income

	<i>Income percentile</i> <i>10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	9.1	7.6	7.0	6.4	5.4
<i>30</i>	5.0	4.2	3.9	3.6	3.0
<i>40</i>	2.9	2.5	2.3	2.1	1.8
<i>50</i>	1.7	1.5	1.4	1.3	1.1
<i>60</i>	1.0	0.9	0.8	0.8	0.7

	<i>Income percentile</i> <i>10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	0.13	0.05	0.03	0.02	0.004
<i>30</i>	0.12	0.04	0.02	0.01	0.004
<i>40</i>	0.12	0.04	0.02	0.01	0.004
<i>50</i>	0.11	0.04	0.02	0.01	0.004
<i>60</i>	0.10	0.04	0.02	0.01	0.003

How does EPP^{high} value risk reduction, as compared to UUU and CBA? By contrast with UUU and CBA, EPP^{high} is biased toward the *poor*. Moving rightwards in each row, up the income ladder for each age, the EPP^{high} numbers *decrease* rather than increasing as they did for UUU and CBA. For example, among 40 year olds: a 1-in-100,000 risk reduction for a bottom-income individual yields 2.9 units of ethical value, while the very same reduction for a low, middle, higher, or top-income individual yields, respectively, 2.5, 2.3, 2.1 and 1.8 units of ethical value. This pattern is *one* aspect of the way in which EPP gives priority to the worse off. Why? Consider two individuals of the same age, one with an income of \$30,000, the other an income of \$50,000. The first individual is worse off in lifetime terms. Adding an incremental year of \$30,000 produces a smaller increase in lifetime well-being than adding an incremental year of \$50,000; but the first increase befalls an individual with lower lifetime well-being, and so is up-weighted by EPP.

Also note that, in valuing risk reduction, EPP^{high} has a more pronounced bias toward the *young* than UUU or CBA. This can be seen by moving *down the column* within each income group. For example, among middle-income individuals, the ethical value of a 1-in-100,000 reduction in risk is 7.0 units for 20 year-olds; 3.9 for 30 year olds; 2.3 for 40 year olds; 1.4 for 50 year olds; and 0.8 for 60 year olds. The parallel UUU and CBA values also decrease column-wise (moving up the age ladder), but not nearly as steeply. Again, the distinctive pattern of EPP is tied to its priority for the worse off. Younger individuals have a greater chance of a shorter lifespan; a 20 year old has some chance of only living to 25, while a 30 year old will have a lifespan of at least 30 for sure. Among any given income group, an individual with a shorter lifespan is worse off in lifetime terms than an individual with a longer lifespan. Adding an incremental year of life produces the very same benefit for either individual, but EPP gives greater ethical weight to the benefit in the first case because the individual is worse off.

Let's turn, finally, to the valuation of income reduction by EPP^{high} (Table 8C). While the ethical value of income reduction is uniform for CBA ("a dollar is a dollar is a dollar") and decreases row-wise for UUU, Table 8C shows that it decreases row-wise yet more steeply for EPP^{high}. This sharp decrease is a compound effect of the decreasing marginal utility of money and the prioritarian up-weighting of well-being gains to the worse off.

The EPP^{high} optimal value for $\$D$ in the case of Uniform Risk Reduction and Cost Incidence ($\$78$) can be arrived at by combining Tables 7C and 8C—just as the CBA and UUU optimal values were determined by combining the corresponding tables.⁸⁸

2. Other Policies

We have thus far explicated how the various methodologies (CBA, UUU, and EPP) value risk reduction and income reduction; and used these patterns to explain one type of policy choice (Uniform Risk Reduction and Uniform Cost Incidence). We can gain further insight into the methodologies by varying the policy choice.

First, consider a variation in the cost incidence of the policy. Uniform Cost Incidence assumes that each individual, regardless of age or income, pays the same amount for the population-wide risk reduction. (This would occur, for example, if the cost of the reduction were passed through to consumers via an increase in the price of goods or services consumed equally by all cohorts.) By contrast, we can imagine patterns of cost incidence that are more favorable to those at lower income. Richer individuals might bear a greater share of the costs of the population-wide risk reduction, notably (a) because those costs are passed through to consumers via an increase in the price of goods or services that are consumed more heavily by those with higher incomes and/or (b) because those costs are spread through the income tax system, which imposes a greater absolute burden on those with higher income.

Specifically, consider two variations in cost incidence. One such variation is “Proportional Cost Incidence,” whereby everyone pays the same fraction of her income; the other variation is “Highly Progressive Cost Incidence,” whereby all costs are borne by top-income individuals. Table 9 below displays the largest per-capita cost $\$D$ that the methodologies are willing to incur for a Uniform Risk Reduction combined with these various incidence assumptions. (In this table and below, per-capita costs are calculated for the entire population even if the expense of reduction is wholly borne by a specific group.)

Table 9: Uniform Risk Reduction with Alternative Cost Incidence Assumptions

	UUU	EPP^{mod}	EPP^{high}	Textbook CBA	Govt practice CBA
Uniform Cost Incidence	\$48	\$59	\$78	\$91	\$91
Proportional Cost Incidence	\$77	\$103	\$159	\$91	\$91
Highly Progressive Cost Incidence	\$184	\$360	\$804	\$91	\$91

⁸⁸ Here, as above, this is an approximation. See *supra* note 86.

Note that the break-even amounts for UUU increase as costs are shifted to those in a higher income group—first, in a proportional manner, and even more so in a highly progressive manner. UUU is willing to pay more in aggregate for a given risk reduction if those aggregate costs fall more heavily on the rich. Break-even values increase from \$48 to \$77 to \$184. This is a consequence of the declining marginal utility of income (exactly what we saw in Table 8A, showing how the UUU value of income increments decreases row-wise). The effect is intensified by EPP. This is a compound result of the declining marginal utility of income *and* priority for those at lower income (again, exactly what we saw in Table 8C). Break even values for EPP^{mod} increase more sharply (from \$59 to \$103 to \$360), and yet more sharply for EPP^{high}.

By contrast, CBA is insensitive to cost incidence. The break-even cost for CBA is \$91 regardless of how much of the costs are shifted to the rich. This is just because “a dollar is a dollar,” as shown by the constant pattern in Table 8B. Assume that $\$D$ is the break-even value with Uniform Cost Incidence. A policy that confers a uniform risk reduction for all cohorts, and imposes a $\$D$ cost on all, has a net zero sum of EVs. Now imagine a new policy, which holds constant the risk reduction, but shifts costs, so that some cohorts pay less than $\$D$ per member, while others pay more than $\$D$ —holding constant total costs. The cohorts that now pay less than $\$D$ will have larger EVs for the new policy, while the cohorts that pay more than $\$D$ will have smaller EVs, but these changes will balance out, so that the sum of EVs for the new policy will still be zero. And if total costs for the new policy are larger than for the initial policy, the sum of EVs will be negative regardless of the incidence pattern.⁸⁹

Second, consider a variation in the beneficiaries of the policy. Imagine, specifically, that the policy addresses some risk that is associated with low income. Let’s call this Income-Targeted Risk Reduction. Assume, specifically, that the policy reduces the risk of bottom-income individuals by 5-in-100,000. (Thus the total risk reduction is the same as in the Uniform Case, where a five-fold greater population incurred a reduction of 1/5th the size.) Table 10 shows break-even values for this policy, with a variety of Cost Incidence assumptions—uniform, proportional, highly progressive, and now the additional possibility that the poor themselves wholly bear the costs of the policy (“Beneficiary Incidence”).

⁸⁹ See *infra* Appendix for a qualification to the statement that CBA is neutral to cost incidence.

Table 10: Income-Targeted Risk Reduction

	UUU	EPP ^{mod}	EPP ^{high}	Textbook CBA	Govt practice CBA
Uniform Cost Incidence	\$32	\$55	\$98	\$15	\$91
Proportional Cost Incidence	\$51	\$98	\$201	\$15	\$91
Highly Progressive Cost Incidence	\$122	\$337	\$1008	\$15	\$91
Beneficiary Incidence	\$15	\$22	\$34	\$15	\$91

This table shows, for the first time in this section, the divergence between governmental practice and textbook CBA. As explained earlier, governmental practice CBA ignores variation among cohorts in EVs for risk reduction. Thus the break-even value for governmental practice CBA is the same as in the Uniform Risk Reduction case—even though bottom-income individuals have substantially lower EVs for a given risk reduction than those with higher incomes. By contrast, textbook CBA substantially lowers the break-even as compared to governmental practice CBA (\$15 as compared to \$91) because the average EV for the bottom-income individuals for a 5-in-100,000 risk reduction is much lower than the average EV of all the cohorts for that reduction.⁹⁰

Textbook CBA converges with UUU if the costs are borne by the beneficiaries. But the break-even costs diverge once we allow for other patterns of incidence—showing again that UUU is sensitive to cost incidence while CBA is not. EPP tolerates somewhat higher break-even costs than UUU and CBA in the case of Beneficiary Incidence. The breakeven values become much greater as costs are shifted to higher-income groups. As compared to UUU, EPP up-weights the risk-reduction benefits to lower income groups (Table 7A versus Table 7C), and downweights the cost of income reduction for higher-income groups (8A versus 8C).

Third, consider a different kind of variation in the benefits of the policy. Now, the beneficiaries are individuals in poorer health. Specifically, I assume that 1/5th of the population (in every income-age group) consist of a “fragile” group, with a less favorable survival curve.

⁹⁰ The average EV of the bottom-income cohorts for a 5-in-100,000 risk reduction is \$77. This number, multiplied by the fraction of bottom-income individuals in the whole population (1/5), yields the \$15 in Table 10. (Recall that the values in Table 10 are per capita costs among the whole population). The population average EV for a 5-in-100,000 risk reduction is \$457. If we use this number rather than \$77 to value each bottom-income individual’s risk reduction, and then once more multiply by the fraction of bottom-income individuals in the whole population, we arrive at \$91.

Thus there are fifty cohorts rather than twenty-five. The policy reduces the risk of the fragile individuals by 5-in-100,000—for example, by abating some toxin that specifically affects these vulnerable individuals, or by improving some therapy that they use. Call this policy “Health-Targeted Risk Reduction.” Table 11 below shows CBA, UUU, and EPP break-even values for a Health-Targeted Risk Reduction and varying cost-incidence assumptions—Uniform Incidence, Beneficiary Incidence, or Incidence on the Healthy.

Table 11: Risk Reduction for “Fragile” Individuals

	UUU	EPP ^{mod}	EPP ^{high}	Textbook CBA	Govt practice CBA
Uniform Cost Incidence	\$40	\$50	\$69	\$77	\$88
Beneficiary Cost Incidence	\$41	\$47	\$59	\$77	\$88
Incidence on Healthy	\$40	\$51	\$72	\$77	\$88

Note that the break-even UUU value for Health-Targeted Risk Reduction and Uniform Cost Incidence (\$40) is lower than the corresponding value in Table 6 for Uniform Risk Reduction and Uniform Cost Incidence (\$48). This is fairly intuitive: as between fragile and non-fragile individuals of the same age and income, a one-time reduction in mortality risk produces *less* increase in expected well-being for the first group (who are likelier to die in all future periods). Further, a change in current income has roughly the same UUU benefit for fragile and non-fragile individuals of the same age and income.⁹¹

However, perhaps more surprisingly, break-even EPP values (\$50 and \$69 in Table 11) for Health-Targeted Risk Reduction and Uniform Cost Incidence are *also* lower than the corresponding values in Table 6 (\$59 and 78). This occurs, in part, because EPP assigns a lower ethical value to a unit risk reduction for a fragile individual as opposed to a non-fragile individual of the same age and income.⁹² This is illustrated in Table 12 immediately below.

⁹¹ These are only slightly different, since the fragile individual has a slightly greater chance of dying without consuming the income.

⁹² EPP also assigns a higher disvalue to costs on fragile individuals. A reduction in income produces roughly the same change in expected well-being for fragile and non-fragile alike, but in the former case this loss is suffered by individuals who tend to be worse off in lifetime terms. This effect, together with the effect displayed in Table 12, is what accounts for the lower EPP breakeven values in Table 11 as compared to Table 6.

Table 12: EPP^{high} Nonfragile/Fragile Risk Reduction Values

	<i>Income percentile 10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	9.1/ 8.6	7.6/ 7.3	7.0/ 6.7	6.4/ 6.2	5.4/ 5.2
<i>30</i>	5.0/ 4.5	4.2 / 3.9	3.9/ 3.6	3.6/ 3.3	3.0/ 2.8
<i>40</i>	2.9/ 2.6	2.5 / 2.2	2.3/ 2.1	2.1/ 1.9	1.8/ 1.7
<i>50</i>	1.7/ 1.4	1.5/ 1.3	1.4/ 1.2	1.3/ 1.1	1.1/ 1.0
<i>60</i>	1.0/ 0.8	0.9/ 0.7	0.8/ 0.7	0.8/ 0.6	0.7/ 0.6

Explanation: The entry to the left of the “/” symbol in each cell is the EPP^{high} value of a 1-in-100,000 risk reduction for a non-“fragile” individual in the age-income cohort, while the entry to the right of the “/” symbol is the value for a “fragile” individual of the same age and income. These are normalized so that 1 denotes the EPP value of a 1-in-100,000 risk reduction for a nonfragile individual in the age 60/bottom income cohort.

The mechanism at work in Table 12 is subtle. To see it in a simpler way, imagine that individuals of a given age (say, 40) and income (say, \$30,000) face a risk of death in the next year. The cohort of such individuals has two subcohorts. A “fragile” subcohort has a probability p of dying at the age of 60, if they survive the current year, and a $(1-p)$ chance of dying at 80. A non-fragile subcohort has a smaller probability ($p^* < p$) of dying at 60, and a correspondingly greater chance $(1-p^*)$ of living until 80. Is there greater EPP value in eliminating the risk of death in the next year for the fragile or the non-fragile subcohort?

The answer is the latter. Eliminating the fragile cohort’s risk yields a lottery with probability p of a lower level of transformed lifetime well-being (the transformed lifetime well-being associated with a 60-year life at income \$30,000), and a probability $(1 - p)$ of a higher level of transformed lifetime well-being. The EPP value in *that* lottery is less than the value in the same lottery with p replaced by p^* .

3. A Summary

This worked example illustrates important features of the SWF approach.

First, the SWF framework, like CBA, provides a tractable methodology that can be used to answer concrete policy issues. For example, like CBA, it can tell us the maximum justified cost to be expended to prevent a fatality risk.

Second, the framework, because it is built around quantitative measures of well-being and social value, offers insight into policy questions that would be difficult for policy analysts equipped with the concepts of “utilitarianism” and “prioritarianism” but lacking the framework’s formalism and numbers. For example, does priority for the worse off mean less or greater effort to produce a given risk reduction for “fragile” individuals whose baseline mortality risk in every

year is greater? Utilitarianism already gives extra weight to reducing risk among the young as opposed to the old. How much does this weight increase if we shift to prioritarianism? More generally, the quantification of the comparative ethical value, across cohorts, for risk reduction and income change (as per Tables 7A/B/C and 8A/B/C) is, I believe, very useful information for (weak or strong) welfarists, whether their commitments are utilitarian or prioritarian.

Third, the SWF approach gives *different* answers to policy questions than CBA. This is true both of UUU and of EPP. UUU and EPP breakeven values for a wide range of hypothetical policies (varying in the beneficiary and in cost incidence) are *substantially* divergent from CBA values. Further, CBA values both risk reduction and income change quite differently than both UUU and EPP. CBA is significantly more biased to the rich than UUU, let alone EPP. CBA is neutral to cost incidence, while UUU and EPP give greater weight to income changes experienced by those at lower income. In each case, the divergence between the SWF approach and CBA occurs because of CBA's failure to correct for the declining marginal utility of income.

Finally, the choice between UUU and EPP is also a consequential one. UUU differs from prioritarianism with even a moderate degree of priority for the worse off (EPP^{mod}), let alone EPP^{high}, in assigning comparative (across-cohort) ethical value to risk and income reductions. EPP and UUU do not show a great divergence in break-even cost values for policies with a uniform pattern of incidence, but EPP becomes willing to impose much higher costs than UUU if the costs are more heavily borne by those at higher incomes.

III. **Why the SWF Approach Improves on CBA: An Argument**

Thus far, the Article has set forth the SWF framework—both in general, and as applied to the domain of regulatory choice—and has explained how it differs from CBA. The presentational strategy is: first explication, then argument. But with the exposition concluded, we can now ask: why believe that the SWF approach is an improvement on CBA?

My argument will proceed as follows. Section A will consider existing defenses of CBA: the long-run Pareto defense, the potential-Pareto defense, and the argument that CBA is a rough proxy for overall well-being. The first defense is a failure; the second two, if well understood, actually show that CBA can be improved by incorporating the SWF framework into regulatory decisionmaking.

Section B considers and rebuts two important objections to SWF-based regulation: first, that it puts regulators in the position of making contestable and legally problematic “value choices”; and second, that it goes against the appropriate institutional structure of government, since the tax system alone is the appropriate channel for redistribution.

I don't claim that the SWF framework should wholly replace CBA in governmental policy analysis. Rather, I favor the *nuanced* use of an SWF, with an SWF (UUU or EPP) as the prime assessment tool, and CBA providing useful information about possible alterations to the tax code that might be coupled with a regulatory intervention so as to increase the value of the SWF. This nuanced approach is described in Section C.

A. *Defenses of CBA?*

CBA has been debated over many years. We'll now consider the most prominent defenses. What's striking about these is that they work within welfarism: appealing either to the Pareto principle, an axiom accepted in some form by virtually all welfarists, or to overall well-being, which is one specification of welfarism. But, of course, welfarism is also the foundation for the SWF framework. It's not as if CBA is grounded in ethical commitments wholly distinct from those that underpin SWFs—in rights versus well-being, or deontology versus consequentialism. Rather, they grow from the same ethical soil. Anyone who likes *that* bit of ethical terrain should see the SWF approach as bettering CBA.

1. Long Run Pareto

One recurrent defense of CBA says this: everyone is better off, in the long run, if governmental agencies employ CBA to choose policies.⁹³ The argument seeks to provide a practice-level defense of CBA. The argument is not that CBA is justified on a one-off basis, but rather that an ongoing governmental practice of CBA is justified, specifically by virtue of Pareto superiority.

Let's distinguish, preliminarily, between "actual-outcome Pareto" and "ex ante Pareto." One policy or practice *P* is "actual-outcome" Pareto superior to a second policy or practice *P** if, in the actual state of nature⁹⁴, the outcome resulting from *P* is Pareto superior to the outcome resulting from *P**. One policy or practice *P* is "ex ante" Pareto superior to a second policy or practice *P** if each person's expected well-being with *P* is at least as large as her expected well-being with *P**, with at least one person's expected well-being strictly larger with *P*.

When it comes to risk regulation, it's highly implausible to believe that the practice of CBA is actual-outcome Pareto superior relative to plausible alternative practices. We're talking about regulatory interventions that are going to save some lives. Some of the individuals who die prematurely because of government's failure to adopt regulations vetoed by CBA under a

⁹³ See Adler, *supra* note 3, at 566-67 (citing sources). A recent example is Richard O. Zerbo and Tyler A. Scott, *(Almost) Everybody Wins: A True Pareto Justification for Practical Welfare Economics and Benefit-Cost Analysis* (working paper, July 14, 2014), at file:///C:/Users/mda18/Downloads/SSRN-id2466101.pdf.

⁹⁴ The standard format for conceptualizing choice under uncertainty, pioneered by Leonard Savage, stipulates that there are possible "states of nature" assigned probabilities. See Adler, *supra* note 3, at 484-85; *infra* Appendix. A given action or policy associates an outcome with each "state of nature." One can think of each "state of nature" as a possible combination of prior facts at the time of choice, together with causal regularities. What the upshot (outcome) of a given choice will be depend upon prior facts and causal regularities.

superiority.¹⁰¹ (6) *CV/EV divergence*. CBA is defined in terms of EVs: hypothetical changes to income in the status quo. But the transfers from ex ante winners to losers that would be required to render one policy Pareto superior to another would occur in the policy outcomes, not in the status quo. It is possible that policy *a* has a positive sum of EVs as compared to policy *b*, and yet (even leaving aside hidden attributes, incentive effects, administrative costs, behavioral effects, and price effects) policy *a* is not in fact potentially ex ante Pareto superior relative to policy *b*.

It is beyond the scope of the Article to discuss all of these in detail. Let me instead elaborate on the first, which I believe is quite relevant to risk regulation and also perhaps not sufficiently widely appreciated. It's useful, here, to reintroduce the concept of a cohort structure that was used earlier. In considering whether policy *a* is ex ante Pareto superior to *b*, economists and policy analysts don't literally go down the population name by name, inquiring into the expected well-being of each distinct person from Aaron Abbott to Zuleika Zux. Rather, the population is partitioned into groups of individuals defined by generic attributes: attributes that serve to predict how policies causally impact individuals' well-being. Now, it might well be the case that some of the attributes that we use to group individuals into these cohorts, for purposes of predicting well-being impacts, are not readily taxable. That is, a given cohort might be defined as individuals who have the very same attributes with respect to *M* types of attribute, A_1, A_2, A_3, \dots through A_M . We might believe that individuals with a specific combination of these types of attributes ($A_1^*, A_2^*, \dots, A_M^*$) can expect to gain from policy *a* relative to policy *b*, while individuals with a different combination of these types of attributes ($A_1^+, A_2^+, \dots, A_M^+$) can expect to lose. In arranging compensation from the first group to the second, we'd want to ascertain whether a given individual (filing her income tax return or her claim for some sort of government benefit) has level * or level + for the various attribute types. But some of these attributes might not be ascertainable to the taxing authorities at reasonable expense; or, alternatively, social norms of tax legitimacy might prevent the authorities from relying on them.¹⁰²

In the Appendix, I provide a simulated case based on the model from Part II.C. above, to show how hidden/non-taxable attributes can interfere with transforming a policy that passes a CBA test (relative to some alternative) into a genuinely ex ante Pareto superior option (relative to that alternative).

¹⁰¹ See Robin Boadway, *The Welfare Foundations of Cost-Benefit Analysis*, 84 *ECON J.* 926 (1974); Boadway and Bruce, *supra* note 15, ch. 9; Charles Blackorby and David Donaldson, *A Review Article: The Case against the Use of the Sum of Compensating Variations in Cost-Benefit Analysis*, 23 *CANADIAN J. ECON.* 471 (1990).

¹⁰² In the case of risk regulation, where interventions most directly affect individual well-being by changing lifespan, health, and income, I suggested that cohorts be defined (at least) as groups with common preferences, of the same age, and with common prior-year income and health quality. It is difficult to identify individuals' preferences or health quality. Age can be readily identified on a tax form, but there may be norms against strong differentiation of taxes by age. See Mankiw, Weinzierl and Yagan, *supra* note 5, at 163 (discussing theoretical case for use of "tags" in the tax code, and noting that the restricted use of age tags by a few countries is "a relatively mild version of tagging relative to what the theory would suggest").

Third: Even if option *a* is indeed potentially ex ante Pareto superior relative to option *b*, this fact does not as such justify the choice of *a*. Assume that ex ante Pareto superiority is endorsed as an ethical principle, with justificatory force. Then, if *a* is potentially ex ante Pareto superior relative to *b*, it follows that *a* coupled with appropriate compensatory transfers from ex ante winners to ex ante losers *would be justified* over *b*. Yet it is an obvious fallacy to leap from the fact that an option *would be justified* if modified in certain ways (here, if coupled with certain compensations), to the conclusion that the option *is* justified even without that modification.¹⁰³

But if option *a* is indeed potentially ex ante Pareto superior relative to option *b*, why wouldn't *a* be converted into a policy that is *actually* ex ante Pareto superior over *b*? In considering this question, we need to focus on the preferences of those who control the tax system (specifically, legislators). If *a* is enacted, compensation for the ex ante losers (relative to *b*) will not be forthcoming if the legislature (or a blocking coalition within the legislature) prefers policy *a* without compensations to policy *a* with them.¹⁰⁴

We will return to these issues below, in discussing the “tax system” challenge to the use of SWFs as a regulatory tool. In considering that challenge, we will need to carefully keep in mind, first, that CBA is not logically equivalent to potential ex ante Pareto superiority; and, second, that the potential for ex ante Pareto superiority may not, in fact, be realized (the scheme of compensations required to make policy *a* actually ex ante Pareto superior to *b* may not be adopted). My ultimate conclusion will be this: the reader who endorses the principle of ex ante Pareto superiority should see the nuanced use of UUU (UUU as the prime regulatory assessment methodology, with CBA functioning as a signal of potential Pareto improvements) to be an improvement over straight CBA.

3 Rough Proxy for Overall Well-Being

The third defense of CBA says that the methodology is a workable proxy for overall well-being.¹⁰⁵ This defense circumvents the contingency of the nexus from CBA to ex ante Pareto superiority. It says: if *a* is ranked by CBA over *b*, then we have reasonable evidence that *a* increases overall well-being relative to *b*. If so, we have a normative basis for choosing *a* even if it's not Pareto superior (directly or after compensations) relative to *b*.

However, CBA is certainly not a perfect proxy for overall well-being. The UUU version of the SWF approach is our best technology for measuring overall well-being, and we saw above with the worked example in Part II.C. that UUU and CBA can diverge substantially. CBA is guaranteed to be a perfect proxy for overall well-being only under special conditions, where income and other determinants of well-being are symmetrically distributed as between the winners and losers of any given pair of policies.

¹⁰³ This fallacy has been stressed in a long tradition of critical thinking about CBA, starting with Sen. Amartya Sen, *The Welfare Basis of Real Income Comparisons: A Survey*, 17 J. ECON. LIT. 1 (1979).

¹⁰⁴ See *infra* text accompanying note 124.

¹⁰⁵ This is the defense of CBA put forth by Adler and Posner, *supra* note 77.

If the SWF format were, for some reason, not usable by regulators as a decision procedure—if it were some kind of abstruse measuring device, suitable for the laboratories of welfare economies but not useful outside the academic campus—the roughness of CBA in proxying overall well-being might be acceptable. But there’s nothing in the format that precludes regulators from using it as a decision procedure. If the reader is gripped by the normative force of overall well-being, she should see UUU as an improvement over CBA.

Alternatively, the reader might not be gripped by the normative force of overall well-being. Utilitarianism, albeit a longstanding and storied ethical view, is contested. Prioritarians challenge utilitarianism, as we saw above. They embrace the Pigou-Dalton principle (the key axiomatic difference between utilitarianism and prioritarianism). EPP, not UUU, is the version of the SWF format that operationalizes a normative commitment to prioritarianism. Now it is possible, I suppose, that CBA under certain conditions is a rough proxy for prioritarianism. But CBA is obviously not a perfect proxy for prioritarianism—as Part II.C. shows. If the reader is gripped by the normative force of the Pigou-Dalton principle, she should see EPP as an improvement over CBA.

B. *Objections to the SWF Approach*

I here consider two important objections to the use of the SWF approach by regulators.

1. Value Choices and Legal Legitimacy

The SWF approach, as I see it, is a framework for ethical deliberation. It functions to guide a decisionmaker in determining what she ought to do, ethically.

We can of course debate the ethical merits of the SWF approach. Does it really give *good* ethical guidance? The framework works within a particular, contested, tradition of ethical thinking: it is (weakly) consequentialist and welfarist. Non-consequentialists and non-welfarists will challenge it. Moreover, the operational specification of the format that I am here recommending—UUU or EPP—can be controverted even by welfare consequentialists. Utilitarians will challenge EPP, and prioritarians UUU. Some welfarists will reject the combination of ethical axioms common to EPP and UUU,¹⁰⁶ and reject both.

Here and (even more so) elsewhere, I have tried to defend the credentials of the SWF approach *qua* framework for ethical deliberation. But the proposal of this Article is not merely that private citizens, or academics, or philanthropists, employ the methodology. The proposal, specifically, is that government *regulators* do so. And this proposal prompts a distinct kind of objection: not ethical, but *legal*. To what extent is it legally permissible for ethical

¹⁰⁶ Recall that those axioms are Ordering, Anonymity, Pareto, Separability, and Continuity. The axiomatic characterization of SWFs is generally discussed in the sources cited *supra* note 15. Adler, *supra* note 3, ch. 5, specifically considers the plausibility of dropping Separability or Continuity. It is also possible to adopt a utilitarian or prioritarian ranking of outcomes but endorse a procedure for choice under uncertainty different from both UUU and EPP (such as EAP or some other). See Adler, *supra* note 3, ch. 7.

considerations to be brought to bear upon government's regulatory decisions? Which officials in the regulatory process have legal authority to select among contending ethical approaches (all controversial), and to employ the selected approach as one basis for enacting legally binding regulations?

Let's begin with statutes. It is undoubtedly legally permissible for ethical considerations to enter into law at the level of legislation. Academic experts can rely on SWFs in giving advice to congressional committees about drafting regulatory statutes, just as they can do¹⁰⁷ in advising congressional committees about the drafting of the tax code. Congress might take the advice: either in shaping the specific terms of a statute, or in issuing a general mandate to one or more agencies to employ some SWF. It has at various times been proposed that Congress put in place a "super-mandate" directing agencies to use CBA. These proposals failed, but not because of doubts about the legality of this mandate. Nor would there be reason to doubt the legality of a UUU or EPP statutory supermandate, if such were proposed.

Absent an SWF supermandate, though, why believe that regulators are legally permitted to inject ethical considerations and, specifically, SWFs into their decisions?

It is a touchstone of analytical jurisprudence that ethical considerations do play a systematic role in a legal system, not merely at the level of the drafting of statutes, but at the level of statutory application, by executive and judicial officials.¹⁰⁸ It occurs when legislators explicitly delegate rulemaking authority to regulatory bodies, thereby inviting regulators to make policy choices, or when statutes employ value-laden terms such as "reasonable" or "appropriate." "Reasonable" is naturally read to invite a consideration of the balance of "reasons," i.e., normative considerations. Ethics also plays a role in specifying the content of law when statutes are framed in (seemingly) non-value-laden terms, which however turn out to have some degree of "open texture"; ethical reasoning fills the gap.

To be sure, for a regulator to have legal authority to inject ethical considerations *of some sort* into her decisions is not, necessarily, legal authority to rely upon the SWF approach. We'd need to look carefully at the various statutes governing environmental, health and safety agencies. E.O. 12,866 currently instructs agencies to use a cost-benefit test in drafting regulations "where permissible by statute."¹⁰⁹ EO 12,866 thereby recognizes that the legal authority of regulatory bodies to inject CBA into their decisions varies from statute to statute; and the same is true of the SWF approach.

¹⁰⁷ For examples of the use of optimal tax theory as a source of practical guidance on tax policy, see, e.g., Diamond and Saez, *supra* note 5; Mankiw, Weizsacker and Yagan, *supra* note 5.

¹⁰⁸ This is a point of agreement for our greatest jurisprudents, Hart and Dworkin, although they disagree as to whether ethical considerations are introduced specifically when officials choose to do so within their zones of legal discretion (Hart) or rather have a yet more pervasive role in determining the content of law (Dworkin). H.L.A. HART, *THE CONCEPT OF LAW* (2d ed. 1994); RONALD DWORKIN, *LAW'S EMPIRE* (1986).

¹⁰⁹ Exec. Order No. 12,866, 3 C.F.R. 638 (1994).

Still, in considering the legal authority of regulators to use SWFs, we should (legally) give them due deference. This is what *Chevron*¹¹⁰ says. *Chevron* generally means that it is legally permissible for agencies to incorporate ethical considerations of some specific sort into their decisions unless doing so is clearly inconsistent with the text, structure, and legislative history of the governing statute. Or at least no court should step in to nullify such incorporation absent clear conflict.

Moreover, in considering (with due *Chevron* deference) whether a specific statute authorizes an agency to use SWFs, we should recognize that the same kind of language that licenses CBA also can be read to license an SWF. The courts have blocked CBA under statutes that focus agency attention on safety, as opposed to the economic (or other non-safety) costs of regulation. The paradigmatic example is *Whitman v. American Trucking*,¹¹¹ which barred CBA under section 109 of the Clean Air Act, instructing the EPA to set air quality standards at a level “requisite to protect the public health” with an “adequate margin of safety.” Such a provision surely also bars an agency from basing its choice on an SWF. The SWF format, like CBA, takes account of economic impacts—in the former case, because economic impacts change individuals’ interpersonally comparable well-being numbers; in the latter case, because they affect individuals’ EVs. A safety-focused statute that precludes an agency from attending to economic impacts in the second way surely also precludes it from doing so in the first.

Conversely, statutes sometime say specifically that agencies should “consider” costs, or use other language that mandates such consideration.¹¹² Could a statutory cost-consideration mandate be fulfilled by an agency’s use of an SWF rather than CBA? Plausibly the answer is yes—certainly if that answer is filtered through a *Chevron* lens. The two formats surely differ in *how* they consider costs; but absent explicit statute language (or legislative history) demanding that costs be assessed via the specific methodology of summing EVs, the SWF approach would also seem to be statutorily permissible.

Finally, as in *Entergy v. Riverkeeper*,¹¹³ courts may, depending on the details of statutory language, structure, and history, construe text that seems to emphasize safety or feasibility as permitting or even requiring cost consideration.¹¹⁴ In *Entergy*, the EPA was permitted to use CBA in determining how stringently to regulate electric utilities’ cooling water intake structures pursuant to a provision that required the structures to “reflect the best technology available for minimizing adverse environmental impact.” What justifies the *Entergy* decision is the ethical premise that costs ought to be considered in shaping regulatory policies, with such premise then motivating a canon of statutory interpretation to the effect that Congress’ intent to preclude the

¹¹⁰ *Chevron v. Natural Resources Defense Council*, 467 U.S. 837 (1984).

¹¹¹ 531 U.S. 457 (2001).

¹¹² See Cass Sunstein, *Cost-Benefit Default Principles*, 99 MICH. L. REV. 1651 (2001). The paradigm of other such language would be the term “reasonable.”

¹¹³ 556 U.S. 208.

¹¹⁴ Sunstein, *supra* note 112.

consideration of costs needs to be clear.¹¹⁵ Since the ethical case for an SWF of some sort is at least as strong as that for CBA, an agency's choice of the SWF format rather than CBA when faced with a statute similar to the EPA's in *Entergy* should also be statutorily permissible.

We can now focus on the “who” question. Assume that an agency has regulatory authority with respect to certain environmental, health or safety risks. If the agency's governing statute clearly requires the use of the SWF framework, then the legal issue is straightforward. The basic principle of legislative supremacy means that officials within the agency are legally obligated to employ the framework in selecting among regulations. Executive branch officials (including the President) lack discretion to countermand the use of the framework, and courts (even filtering statutory interpretation through the lens of *Chevron*) will enforce the SWF requirement.

But this case of clear statutory mandate is (currently) entirely hypothetical. What about the (as of now) more realistic case in which an agency's governing statute permits but does not clearly require the SWF approach? Who within the agency or elsewhere in government legally gets to choose between the SWF approach and other statutorily permissible assessment frameworks (such as CBA)? And, specifically, to choose between different versions of the SWF approach, in particular UUU versus EPP?

In the case of independent agencies, the answer would seem to be: agency heads.¹¹⁶ In the case of executive agencies, the answer would seem to be: the President. A standard view of Presidential power is that the President has the legal authority to direct the rulemaking choices of executive agencies, up to the limits of statutory requirement. Acting on this view, Presidents since 1981 have directed executive agencies to use CBA as their general assessment framework—as currently required by EO 12,866.¹¹⁷ Whatever ambiguity one might glean in that Executive Order's use of the language of “distributive impacts,” it is pretty clear that EO 12,866 is a pro-CBA, not pro-SWF document. And this reading is reinforced by Circular A-4¹¹⁸, the authoritative OIRA guidance document concerning EO 12,866. For an executive agency to shift from CBA to the SWF approach, if the statute permits but does not clearly require the latter, would be legally permissible only if EO 12,866 is amended.

¹¹⁵ For a related view, regarding the arbitrariness of not considering costs, see Cass Sunstein, *Cost-Benefit Analysis and Arbitrariness Review* (Harvard Law School Discussion paper No. 870, June 2016), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2752068

¹¹⁶ The standard view is that Congress can, within certain constitutional limits, create independent agencies whose heads are insulated from at-will Presidential removal and at-will Presidential policy control. *Humphrey's Executor v. United States*, 295 U.S. 602 (1935); *Morrison v. Olson*, 487 U.S. 654 (1988). A revisionary view holds that it would be unconstitutional to insulate any agency from at-will Presidential policy control. See Geoffrey Miller, *Independent Agencies*, 1986 S. CT. REV. 41. On this latter view, the policy-assessment framework used by an independent agency (CBA, the SWF framework, etc.) would also be ultimately a matter of Presidential choice. Note that the current CBA mandate does not apply to independent agencies. See Exec. Order No. 12,866, sec. 3(b), 3 C.F.R. 638 (1994). See also Exec. Order No. 13,579, 3 C.F.R. 256 (2012).

¹¹⁷ Exec. Order No. 12,866, 3 C.F.R. 638 (1994).

¹¹⁸ OMB, Circular A-4 (Sept. 17, 2003),

https://www.whitehouse.gov/sites/default/files/omb/assets/regulatory_matters_pdf/a-4.pdf.

To sum up: it is indeed crucial to differentiate between the ethical attractiveness of the SWF approach, and its legal permissibility. A regulatory official might come to believe (for example, after reading the academic literature on SWFs) that using an SWF to assess regulation is ethically preferable to using CBA. But it hardly follows that the official is legally authorized to undertake this shift. In this section, I have identified three prime channels for giving a legal imprimatur to the SWF approach: statutory language; Presidential order; and a decision to do so by the heads of an independent agency. Scholars (like myself) who ethically favor the approach and want to see it used in the U.S. government should direct their efforts to these channels.

2. The Tax System

A longstanding scholarly tradition in welfare economics argues that distributional considerations should be the concern of the tax system; other institutions should focus on efficiency.¹¹⁹ This line of thought has been most fully developed by Louis Kaplow.¹²⁰ Kaplow argues that non-tax bodies, such as regulatory bodies, should employ CBA to evaluate their choices. Kaplow's argument, in a nutshell, is this. Assume that the regulatory body adopts some procedure other than CBA. In a particular case, the non-CBA procedure favors option *b*, while CBA favors option *a*. Assume, further, that the income tax system is not an ideal "lump-sum" system, but (more realistically) a Mirrleesian income tax system, such that an individual's ability, wage rate, and labor-leisure choices are unobservable. Notwithstanding this hidden information, Kaplow shows the following, given appropriate assumptions (for short, the "benchmark assumptions") regarding the "separability" structure of individual preferences as well as other relevant factors: there is a potential change to the income tax schedule which, if coupled with *a*, makes everyone better off than *b*.¹²¹ Thus, Kaplow concludes, anyone whose fundamental ethical criterion respects the Pareto principle should want regulatory and other non-tax bodies to employ CBA in assessing their choices.

¹¹⁹ This line of thought goes back, at least, to Richard Musgrave's proposal that allocation and distribution functions be separated within government. See *THE THEORY OF PUBLIC FINANCE: A STUDY IN POLITICAL ECONOMY* (1959). Another important precursor to Kaplow is Aanund Hylland and Richard Zeckhauser, *Distributional Objectives Should Affect Taxes but not Program Choice or Design*, 81 *SCANDINAVIAN J. ECON.* 264 (1979).

¹²⁰ See sources cited *supra* note 98. Kaplow's analysis in these sources focuses on public good provision. He and Steven Shavell have famously developed a parallel analysis for judicial doctrine. For a discussion of this work and the subsequent academic literature concerning it, see Lee Anne Fennell and Richard H. McAdams, *The Distributive Deficit in Law and Economics*, 100 *MINN. L. REV.* 1051 (2016). David Weisbach's critique of CBA with distributional weights builds upon the Kaplow analysis. *Distributionally Weighted Cost-Benefit Analysis: Welfare Economics Meets Organizational Design*, 7 *J. LEGAL ANALYSIS* 151 (2015).

¹²¹ What exactly are the benchmark assumptions? The ones mentioned by Kaplow concern labor separability, unobserved preference heterogeneity, the administrative cost of the tax system, and behavioral effects. See Kaplow, *supra* note 5, at 137-44. I have pointed to others, concerning hidden attributes (in addition to individuals' hidden "ability," wage rate, and preferences), price effects, and CV/EV divergence. See *supra* text accompanying notes 97-101.

By "benchmark assumptions," I mean *whatever* assumptions need to be made with respect to these seven issues such that there is a potential change to the income tax schedule which, if coupled with *a*, makes everyone better off than *b*.

A notable feature of the Kaplow argument is that the argument does *not* presuppose CBA itself to be a bedrock ethical norm. Indeed, Kaplow himself has expressed a foundational commitment to the SWF framework, not CBA.¹²² Rather, the structure of the Kaplow argument is to start with a Pareto-respecting ethical view (such as a Pareto-respecting SWF) and to argue from those foundational commitments to the *institutional* conclusion that non-tax bodies should use CBA.

Kaplow's analysis does not focus on the case of policy choice under uncertainty—where the costs and benefits of intervention are probabilistic. He generally simplifies matters by assuming that the effects of a policy on given person are known for certain—if not by government, then by the individual herself. However, once we take into account uncertainty, there is no reason to believe that CBA together with changes to the tax system tracks Pareto superiority in the actual-outcome sense. CBA at most tracks ex ante Pareto superiority. So I'll assume that Kaplow's analysis generalizes to the case of policy choice under uncertainty by showing this: with appropriate benchmark assumptions, there is a change to the Mirrleesian income tax schedule which renders a CBA-approved policy ex ante Pareto superior to any alternative.

Is the Kaplow argument persuasive? It will clarify matters to assess the argument from the perspective of the President. We saw above that the President has a major role in determining which assessment tools are used by regulatory agencies. The President is also, of course, a central player in shaping the U.S. tax system. The tax code is put in place by statute, and the President must agree to the enactment or amendment of a statute (absent a veto override). So we can evaluate the Kaplow argument as follows: would it be sensible for this key actor with respect to both regulatory and tax institutions to agree (as per the Kaplow argument) to impose a CBA mandate on regulators?

Let's consider three cases: (1) Durably Unified Government with a Shared Commitment to UUU; (2) a Presidential commitment to UUU, but Government is not Durably Unified; or (3) a Presidential commitment to EPP. In all three cases (I'll suggest), there are reasons for the President *not* to want a CBA mandate. The best alternative to a CBA mandate is for regulatory bodies to use UUU or EPP in a “nuanced” fashion—an approach that will be described in section C below, after I rehearse the considerations against a CBA mandate.¹²³

¹²² See Kaplow, *supra* note 5, ch. 3, 13-15. See also LOUIS KAPLOW AND STEVEN SHAVELL, FAIRNESS VERSUS WELFARE (2002).

¹²³ The analysis that follows provides a response, I believe, to David Weisbach's important critique of distributionally weighted CBA and defense of the use of traditional CBA by non-tax bodies. Weisbach, *supra* note 120. Weisbach's key addition to the Kaplow argument is institutional. Weisbach stresses the advantages to social welfare maximization from institutional specialization.

Weisbach's arguments are most plausible on the assumption that all key actors durably share a commitment to UUU—case (1) below. Even in this case, I believe there are good institutional arguments for non-tax bodies engaging in nuanced UUU-based analysis (see below, Part III.C) rather than CBA. Weisbach stresses that non-tax bodies are set up to address market failures, and that CBA is the proper tool for doing so. But market failures

Case (1): Durably Unified Government with a Shared Commitment to UUU.

Assume that the President endorses UUU. A majority of Congress also does so. Moreover, the President (she) expects that her successors for the foreseeable future will also endorse UUU, as will future Congresses. The legislature and the Presidency is unified by this shared commitment to UUU; and the commitment is a durable one, extending into the foreseeable future.

In deciding whether to favor a CBA mandate, the President considers the following scenario. A regulator at some point in the future is choosing between option *a*, which is preferred by CBA, and option *b*, which is preferred by UUU. A CBA mandate would require the regulator to choose *a*. Does the President, indeed, want the regulator in such a scenario to choose *a*? If the benchmark assumptions hold true, there is a change to the income tax system (let's denote this as " Δ "), which is such that the hybrid policy of adopting *a* and then enacting the tax change Δ is ex ante Pareto superior to *b*. Call this hybrid policy *a-plus- Δ* .

Moreover, if the benchmark assumptions hold true, the President predicts that Congress at the time of the regulatory choice would in fact enact the tax change Δ , and the contemporaneous President would sign Δ into law, were the regulator to adopt *a*. By enacting *a*, the regulator prompts a statutory response that produces the hybrid policy *a-plus- Δ* —or so the President predicts. The President predicts this because she believes that Congress and the President in the future (like the President herself now) will favor UUU. A future Congress and President who embrace UUU can be predicted to prefer *a-plus- Δ* to *b* and *b* to *a*, and thus by transitivity will prefer *a-plus- Δ* to *a*.

However, the benchmark assumptions may not hold true. If so, there is no Δ such that *a-plus- Δ* is ex ante Pareto superior to *b*. Thus the regulatory choice of *a* rather than *b* might not be followed by Congressional tax changes that yield a result which is preferred by UUU to *b*.

(departures from Pareto optimality arising from externalities, information asymmetries, etc.) will be disfavored by UUU, too, since it respects ex ante Pareto superiority. Weisbach also emphasizes the lack of expert knowledge in non-tax bodies about appropriate departures from traditional CBA; but such knowledge is endogenous. (Non-tax bodies would develop models and information to implement UUU, models and information these bodies currently lack, if under a binding legal instruction to do UUU-based analysis—just as agencies developed expertise in CBA that was lacking prior to the issuance of a CBA mandate by President Reagan in 1981.) Finally, Weisbach describes the advantages in coordinating tasks among governmental bodies. Distribution can be hived off to the income tax system. However, these gains from coordination must be balanced against the losses in utilitarian social welfare that arise when the benchmark assumptions fail and CBA and UUU diverge—when the policy preferred by CBA, even after accompanying income tax changes, is worse than the policy recommended by UUU.

Further, once we allow for ethical disagreement among political actors (case (2) below), a gap emerges between optimal institutional design for the case of a shared SWF and the directives that a given actor (for example, the President) rationally endorses in her actual political environment. Finally (case (3) below), the Weisbach and Kaplow premise—that a well-functioning tax system *could* (at least if the benchmark assumptions hold) fully correct losses to social welfare that occur when an agency uses CBA—does not apply to EPP, since EPP can violate ex ante Pareto superiority. (Weisbach might respond that EPP should be flat rejected just because of such violation—but it might be responded that SWFs that satisfy ex ante Pareto superiority should be rejected because they violate Equity Dominance.)

In short, even in the case of Durably Unified Government with a Shared Commitment to UUU, the possible failure of the benchmark assumptions is a reason why the President might not favor a CBA mandate.

Case (2): A Presidential Commitment to UUU, but Government is not Durably Unified

The case of a durably unified governmental commitment to UUU, or any other SWF, is highly unrealistic. A given SWF is the formalization of an ethical view (in the case of UUU, utilitarianism). But citizens in a modern pluralistic society such as the U.S. espouse a diversity of conflicting views, as do the governmental actors they elect. The party system may constrain ethical diversity among Congress to some extent, but certainly is not strong enough to assure that, at any given point in time, there is always a majority in Congress with a single, shared, ethical framework. The separation-of-powers structure of the Constitution allows for intratemporal ethical conflict between the President and Congress. And, of course, the occurrence of Presidential and Congressional elections at regular intervals works to undermine the intertemporal persistence of whatever framework a given President or Congress may endorse.

Kaplow's analysis, and other work on the "tax system" argument for CBA, does not grapple in a sustained way with the diversity of SWFs among government actors.¹²⁴ This may just be a matter of modeling simplicity. In any event, once we consider such diversity, the argument for a CBA mandate weakens.

Why? Consider the case in which a utilitarian President favors UUU, but this view is not durably shared (either because a majority of the current Congress is *not* utilitarian; or because, while the current Congress is utilitarian, the President has no reason to believe that future Congresses or Presidents will also be).

This President considers a scenario in which a regulator at some point in the future is choosing between option *a*, which is preferred by CBA, and option *b*, which is preferred by UUU. A CBA mandate would require *a*. Does the President, indeed, want the regulator in such a scenario to choose *a*?

The President, now, has *two* reasons to answer "no." First, as above, the benchmark assumptions may fail. But assume they hold true. Even so, the President may not want the regulator to choose *a*. *a*-plus- Δ is *ex ante* Pareto superior to *b*. However, since the officials who control the tax code at the time of the enactment of *a* (that is, the contemporaneous Congress and President) may not be utilitarian, the President may not be confident that the regulator's choice of *a* would trigger the enactment of the tax change Δ .

Why not? For *a* to lead to the hybrid policy *a*-plus- Δ , it must be the case that a majority of the contemporaneous Congress plus the contemporaneous President prefer *a*-plus- Δ to *a*. But

¹²⁴ Neither Weisbach, *supra* note 120, nor Hylland and Zeckhauser, *supra* note 119, does so.


if they do not endorse UUU, they may not have this preference. See Figure 2 for a graphical illustration. As a concrete example, imagine that a majority of the contemporaneous Congress are libertarians who are committed to defending property rights, not utilitarians. They are opposed to limiting property rights by regulatory imposition, unless the property owner is compensated; and they oppose redistributive taxation. Assume that b is a regulation that limits property rights to protect some risk-bearers; a is no regulation. The regulation passes a utilitarian test, but not CBA, and thus it is not possible to fully compensate the property owners. UUU prefers b to a , while the contemporaneous Congress prefers a to b and a to a -plus- Δ .¹²⁵

The critical point is that the UUU President doesn't want the enactment of a alone. She wants a followed by Δ , or at least by some other change to the tax system which yields a hybrid policy that is better than b by utilitarian lights. However, if government is not durably unified in its UUU commitment, the President may fear that the regulator's choice of a would produce simply a . This is a second reason (quite apart from the failure of the benchmark assumptions) for her not to favor a CBA mandate.

Case (3): A Presidential Commitment to EPP

The EPP case is straightforward. EPP can violate the ex ante Pareto principle (by virtue of satisfying the principle of Equity Dominance, which UUU rejects). The EPP President imagines a future regulatory choice in which b is better than a according to EPP, while a is better than b according to CBA. Even if the benchmark assumptions hold true, so that there exists an a -plus- Δ which is ex ante Pareto superior to b , it does not follow that a -plus- Δ is better than b according to EPP. Although there might be some change to the tax system such that a with that change is better than b , CBA does not automatically indicate such a change, even with the benchmark assumptions.

¹²⁵ This analysis is in keeping with the work of Fennell and McAdams, *supra* note 120, who generally stress the possibility of political impediments to tax-based redistribution.



Explanation: The x -axis shows the expected well-being of one cohort, while the y -axis shows the expected well-being of a second. Assume that no other cohorts are affected by the policies. Each solid line is a UUU social indifference curves (points ranked equally good by UUU). Note that a -plus- Δ is ex ante Pareto superior relative to b , and is indeed ranked by UUU above b , which in turn is better than a . However, if the legislators who control the tax system do not endorse UUU, they may not rank a -plus- Δ above a .

C. *SWFs as a Regulatory Methodology: A Nuanced Approach*

The “tax system” defense of a CBA mandate does not succeed—or so I have argued. Still, that defense contains a critical insight for the appropriate design of SWF-based regulatory analysis. The tax system isn’t static. A regulatory intervention might cause, or be deliberately coordinated with, an amendment to the tax statutes. Policy analysts should take account of the possibility of hybrid policies (combining regulatory interventions and tax-code changes). “Nuanced” SWF-based regulatory assessment is structured to do so.

Consider, first, regulatory analysis implementing UUU. There is a set of regulatory policies $\{s, a, b, \dots\}$, and some option b (“best”) is ranked highest by UUU. How should the analyst proceed? If her analytical resources were unlimited, she would carefully consider every policy other than b and, for each, a wide range of tax changes—asking whether that alternative together with that change is ranked above b by UUU. But of course this approach is, in practice,

infeasible. Rather, the analyst needs some threshold test for determining when to devote analytical resources to considering hybrid policies.

What should that threshold test be? A natural answer is that the test should incorporate CBA. For each policy other than b , the analyst should ask: Does CBA prefer that policy to b ? If a given policy (a) is preferred by CBA to b , there is then good reason to search for a tax system change Δ^* such that $a + \Delta^*$ is preferred by UUU to b . Why is this so? If the benchmark assumptions hold true, then there exists a tax system change Δ^* such that the policy plus Δ^* is ex ante Pareto superior to b —and thus preferred by UUU. Even if the benchmark assumptions are false, it is worth considering the magnitude of the factors that defeat ex ante Pareto superiority. If these are not too “large,” there may be a hybrid policy involving a that is preferred by UUU to b even though no such hybrid policy is ex ante Pareto superior to b . (That $a + \Delta^*$ is ex ante Pareto superior to b is a *sufficient* condition for it be preferred by UUU, not a necessary condition.)

To be sure, CBA need not be the *only* component of the threshold test used by the UUU analyst to search for hybrid policies. One can envision a disjunctive, multi-component approach, whereby the analyst (1) asks whether a given alternative a is preferred to b *either* by CBA or by some other threshold screen, and if so (2) searches for hybrid policies of the form a plus tax change Δ^* that are ranked above b by UUU. Yet it’s hard to know what that additional screen would be. Further, there would be minimal “start-up” costs in using CBA as a screen for hybrid policies (since regulatory bodies now have decades of experience with CBA), but significant “start-up” costs for additional screens.

Assume now that a given alternative a is indeed preferred to b by the screening test (be it CBA alone or a multi-component threshold test). What next? Considerations of political feasibility should presumably inform the search for hybrid policies. For example, if a blocking coalition in the Congress is opposed to higher income tax rates, then a close analysis of hybrid policies that include a rate hike seems a waste of time. At the very least, even if the analyst ignores political feasibility in *searching* for hybrid policies, she surely should not do so in her final recommendation. Assume that a passes the threshold test and that, after a search by the analyst for hybrid policies, it emerges that $a + \Delta^*$ beats b according to UUU. It would be folly for the analyst to leap immediately to a *recommendation* of policy a . The actual chance of Δ^* passing Congress might be low.

Individual regulatory agencies are, of course, less well situated than the Presidency to make judgments of political feasibility. Thus the process of searching for hybrid policies could be undertaken by agencies in conjunction with some entity within the Executive Office of the President—or perhaps wholly shunted off to such an entity. It’s easy to see how this process could be one part of a UUU-based system of regulatory review that is overseen and coordinated by OIRA (the linchpin of the current oversight mechanism). A regulatory agency prepares an analytic document. The document outlines the policy alternatives and describes their impacts on

individuals' fatality risks, income, health, and other attributes. The policy alternatives are then evaluated, in two ways. First, they are ranked using UUU. Then, the document uses CBA (or some multicomponent screen) to compare every other alternative to the one ranked tops by UUU ("b"). Having identified *b* and the successfully screened alternatives, the agency sends the document to OIRA. OIRA checks the analysis and (in consultation with the White House) determines whether there is a feasible change to the tax system which, in conjunction with some successfully screened alternative, is better in light of UUU than *b*. If so, it instructs the agency to adopt that alternative; if not, it approves the adoption of *b*.

Let's now swap SWFs. What is the appropriate structure of regulatory analysis with EPP rather than UUU as the SWF?

The importance of nuanced assessment doesn't diminish. Even though *b* is ranked best by EPP, there could well be some change to the tax system which, if conjoined with some other alternative, yields a hybrid policy preferred by EPP to *b*. Nor does the shift to EPP obviate the need for judgments of political feasibility in searching for hybrid policies. What does change, perhaps, is the nature of the screening test that precedes a search. Should CBA be part of this test? EPP, recall, can violate ex ante Pareto superiority. Why, then, does the fact that *a* beats *b* in a CBA test provide a useful signal to search for a feasible alternative Δ^* such that *a* plus Δ^* is preferred by EPP to *b*?

As the prioritarian transformation function $g(\cdot)$ becomes less concave, EPP approaches UUU. Thus, at least for moderate degrees of concavity, CBA does have some screening value (since it does for UUU). Further, as mentioned, CBA is a low-cost screen, given accumulated agency expertise in this methodology. More research is needed in the usefulness of CBA as a screen for EPP. It would be helpful to characterize the conditions under which EPP violates ex ante Pareto superiority, and how these map onto various possible distributions of attributes across cohorts and various possible regulatory interventions.

Conclusion

The goals I've set for this Article are ambitious: to provide an accessible but accurate rendering of the SWF framework for policy evaluation; to illustrate, in detail, how both utilitarian (UUU) and "prioritarian" (EPP) SWFs can be used to evaluate environmental, health and safety regulation; to contrast the framework with CBA; and to offer a reasonably comprehensive *defense* of the approach, both showing why it improves upon CBA, and rebutting objections that sound either in legal legitimacy or the supposed monopoly of the tax system (excluding regulatory bodies) over distributional considerations. Whether the goals have been met or missed is for the reader to decide.

In prior work, I defended CBA as a rough proxy for overall well-being. My analysis here is fully consistent with that prior work—since what I do now is to construct in detail a *precise* rather than rough indicator of overall well-being (UUU), and also an alternative approach (EPP) that is both welfarist and equity regarding. The roughness of CBA is that it quantifies well-being impacts with a money metric. This metric is distorted by the diminishing marginal utility of money. Both UUU and EPP eliminate this distortion, as shown in detail in Part II.C. above.

Many policy-assessment frameworks are *worse* than CBA.¹²⁶ Anyone who cares about good regulation should be vigilant that the Trump Presidency (or any other Presidency) doesn't shift from CBA in *this* direction. CBA was a large step forward on the path of rational deliberation about when and how government should regulate. But it's not the end of that path. We can do better, and this Article shows—or at least tries to show—how.

¹²⁶ See Adler and Posner, *supra* note 77. A framework focused simply on reducing compliance costs, without taking account of the many benefits of regulation, would also be much worse than CBA.

Appendix

A. *Social Welfare Functions: A Brief Technical Presentation*¹²⁷

1. Structure and Functional Forms

Assume a set of outcomes $\{x, y, \dots\}$; a fixed population of persons $\{1, \dots, N\}$; and a well-being measure $w(\cdot)$, which maps a given outcome x onto a vector of well-being numbers $(w_1(x), \dots, w_N(x))$. Then a social welfare function (SWF) is some rule E that ranks well-being vectors. Outcomes are then ranked accordingly:

x at least as good as y iff $(w_1(x), \dots, w_N(x))$ ranked by E at least as good as $(w_1(y), \dots, w_N(y))$

The utilitarian SWF is the sum of well-being numbers. The ranking of outcomes by the utilitarian SWF is as follows:

x at least as good as y iff $\sum_{i=1}^N w_i(x) \geq \sum_{i=1}^N w_i(y)$

The prioritarian SWF uses some strictly increasing and strictly concave function $g(\cdot)$ and sums g -transformed well-being numbers:

x at least as good as y iff $\sum_{i=1}^N g(w_i(x)) \geq \sum_{i=1}^N g(w_i(y))$

An attractive class of prioritarian SWFs is the Atkinson family. These are the only prioritarian SWFs that satisfy an axiom of ratio rescaling invariance. Atkinson SWFs are parameterized by an inequality-aversion parameter γ that can be any positive number.

$g(w_i(x)) = (1 - \gamma)^{-1} (w_i(x))^{1-\gamma}$ or $\ln w_i(x)$ in the special case $\gamma = 1$.

2. Well-Being Measure

In each outcome, a given individual is characterized by a bundle of “welfare attributes”: properties that determine her well-being. Let \mathbf{A} be the set of possible welfare-attribute bundles. Assume individuals have identical ordinal and risk preferences with respect to \mathbf{A} : a common ranking of \mathbf{A} and of lotteries over \mathbf{A} . Then this ranking, if well-behaved in the sense of satisfying the axioms of vNM (von Neumann/Morgenstern) theory, can be represented by a vNM utility function $u(\cdot)$. It is well-known that vNM functions are unique up to a positive affine transformation: some other function $u^*(\cdot)$ also represents the ordinal and risk preferences iff there are constants $d > 0$, e , such that $u^*(a) = du(a) + e$ for all a . Thus $u(\cdot)$ captures information

¹²⁷ The material in Section A.1 is covered in more detail in Adler, *supra* note 3, ch. 5, as well as the other sources cited *supra* note 15. The material in Section A.2 is covered in Adler, *Extended Preferences*, *supra* note 52. On the choice of a zero point to arrive at ratio measurement of well-being, see sources cited *supra* note 56. The material in Section A.3 is covered in Adler, *supra* note 3, ch. 7, as well as the other sources cited *supra* note 59.

about well-being levels and differences. If a, a^+, a^{++}, a^{+++} are any four bundles: $u(a) \geq u(a^+)$ iff $u^*(a) \geq u^*(a^+)$; and $u(a) - u(a^+) \geq u(a^{++}) - u(a^{+++})$ iff $u^*(a) - u^*(a^+) \geq u^*(a^{++}) - u^*(a^{+++})$. Further, if a zero bundle is chosen, $u(\cdot)$ can be rescaled and will be unique up to a positive ratio transformation.

$u(\cdot)$ can now be used as the well-being measure. Let a_i^x be the attribute bundle of individual i in x . Now define $w(\cdot)$ as follows: $w_i(x) = u(a_i^x)$.

Using vNM functions as the basis for the well-being measure generalizes to the case of heterogeneous preferences. Let R_i be the preference structure of individual i . $u^{R_i}(\cdot)$ is a vNM utility function representing these preferences. $s(R_i) > 0$ and $t(R_i)$ are scaling factors. Then $w_i(x) = s(R_i)u^{R_i}(a_i^x) + t(R_i)$. The idea here is to equate the well-being of individual i in outcome x with the utility number assigned to her attributes there by her (appropriately scaled) utility function $u^{R_i}(\cdot)$. This approach yields well-being numbers unique up to a positive affine transformation and (after specification of a zero point) unique up to a positive ratio transformation.

For simplicity, I'll henceforth drop the scaling factors from the notation and use $u^{R_i}(\cdot)$ to mean the vNM function representing preferences R_i after application of $s(R_i)$ and $t(R_i)$ and, if necessary, specification of a zero point.

3. Uncertainty

I'll present the SWF framework under uncertainty using the Savage representation of uncertainty, which is very standard in economics. Assume a finite set \mathbf{Z} of "states of nature," with each state z having probability $\pi(z)$. A given choice (whether the status quo policy s of inaction, or the choice of some alternative policy P ¹²⁸) is a mapping from states to outcomes. Let $x^{s,z}$ be the outcome of the status quo option s in state z , and $x^{P,z}$ the outcome of policy P in state z .

UUU is the most straightforward way to apply a utilitarian SWF under uncertainty. Policies are ranked according to the expected sum of individual well-being. That is, the UUU rule is as follows:

$$\text{Policy } P \text{ is at least as good as policy } P^* \text{ iff } \sum_{z \in \mathbf{Z}} \pi(z) \sum_{i=1}^N w_i(x^{P,z}) \geq \sum_{z \in \mathbf{Z}} \pi(z) \sum_{i=1}^N w_i(x^{P^*,z}) .$$

Turning to prioritarianism, we first have EPP (ex post prioritarianism), which ranks policies according to the expected sum of individual well-being transformed by the strictly increasing and concave $g(\cdot)$ function.

¹²⁸ In the main text, I use "a" and "b" to refer to policies, but here use "P" and "P*" because "a" is now being used to abbreviate attribute bundles.

Policy P is at least as good as policy P^* iff $\sum_{z \in \mathbf{Z}} \pi(z) \sum_{i=1}^N g(w_i(x^{P,z})) \geq \sum_{z \in \mathbf{Z}} \pi(z) \sum_{i=1}^N g(w_i(x^{P^*,z}))$

Second, we have EAP (ex ante prioritarianism), the sum of transformed expected individual well-being. Here, the $g(\cdot)$ function is applied to individuals' well-being expectations rather than to realized well-being

Policy P is at least as good as policy P^* iff $\sum_{i=1}^N g\left(\sum_{z \in \mathbf{Z}} \pi(z) w_i(x^{P,z})\right) \geq \sum_{i=1}^N g\left(\sum_{z \in \mathbf{Z}} \pi(z) w_i(x^{P^*,z})\right)$

Note, finally, that UUU, the expected sum of individual well-being, is mathematically equivalent to the sum of individual expected well-being.

$\sum_{z \in \mathbf{Z}} \pi(z) \sum_{i=1}^N w_i(x^{P,z}) = \sum_{i=1}^N \sum_{z \in \mathbf{Z}} \pi(z) w_i(x^{P,z})$. EPP, the expected sum of individual transformed well-being, is mathematically equivalent to the sum of expected transformed well-being.

$\sum_{z \in \mathbf{Z}} \pi(z) \sum_{i=1}^N g(w_i(x^{P,z})) = \sum_{i=1}^N \sum_{z \in \mathbf{Z}} \pi(z) g(w_i(x^{P,z}))$ However, EPP and EAP are *not* equivalent. The expected sum of transformed well-being, or sum of expected transformed well-being, is not equivalent to the sum of transformed expected well-being.

Variations on UUU and EPP apply a non-linear and strictly increasing function to the sum of well-being or transformed well-being: $\sum_{z \in \mathbf{Z}} \pi(z) F\left(\sum_{i=1}^N w_i(x^{P,z})\right)$ or

$$\sum_{z \in \mathbf{Z}} \pi(z) F\left(\sum_{i=1}^N g(w_i(x^{P,z}))\right), F(\cdot) \text{ non-linear and strictly increasing.}$$

4. LHI Lotteries and Cohorts

A policy can be described as a list of lotteries over welfare-attribute bundles, one such lottery for each person in the population. Let $a_i^{P,z}$ be the attribute bundle of individual i in the outcome that results from policy P given state z (i.e., in outcome $x^{P,z}$). Now let $\rho_i^P(a^*)$ denote the probability that individual i receives bundle a^* , given policy P . This probability can be defined as follows: $\rho_i^P(a^*) = \sum_{z \in \mathbf{Z}: a_i^{P,z} = a^*} \pi(z)$.

UUU, EPP and EAP are such that the list of individual lotteries over attribute bundles created by a given policy P is *sufficient* information to assign a value to P . Note that the formulas for UUU, EPP, and EAP can each be restated as summing, across individuals, a value

that depends upon the lottery over attribute bundles she faces. UUU: $\sum_{i=1}^N \sum_{a \in \mathbf{A}} \rho_i^P(a) u^{R_i}(a)$. EPP:

$$\sum_{i=1}^N \sum_{a \in \mathbf{A}} \rho_i^P(a) g(u^{R_i}(a)) \text{ . EAP: } \sum_{i=1}^N g\left(\sum_{a \in \mathbf{A}} \rho_i^P(a) u^{R_i}(a)\right) \text{ .}^{129}$$

In the case of environmental, health and safety regulation, each attribute bundle is an “LHI” history—describing (at least) a possible longevity, income/consumption each year alive, and health status each year alive.

As for cohorts: let c be a bundle of exogenous (welfare and/or further) attributes such that everyone with bundle c has the same policy-specific lottery over a -bundles and the same tastes, for every policy P . Denote these, respectively, as $p_c^P(a)$ and R_c . Assume, further, that everyone in the population has one such bundle in a set \mathbf{C} of these exogenous bundles, with $n(c)$ the number of individuals with c . Then the formulas for UUU, EPP and EAP can be restated in terms of cohort probabilities and utilities. UUU: $\sum_{c \in \mathbf{C}} n(c) \sum_{a \in \mathbf{A}} \rho_c^P(a) u^{R_c}(a)$. EPP:

$$\sum_{c \in \mathbf{C}} n(c) \sum_{a \in \mathbf{A}} \rho_c^P(a) g(u^{R_c}(a)) \text{ . EAP: } \sum_{c \in \mathbf{C}} n(c) g\left(\sum_{a \in \mathbf{A}} \rho_c^P(a) u^{R_c}(a)\right)$$

B. Cost-Benefit Analysis

As mentioned earlier, CBA is specifically defined here as the sum of equivalent variations.¹³⁰ An individual’s equivalent variation for a given policy is the change to her status quo consumption that makes her equally as well off as the policy. In the case of multiperiod outcomes, it is the change to her status quo consumption during some specified time period (typically, the current period).

Attribute bundles should therefore describe each individual’s consumption (during the specified time period), denoted “ c ,” as well as other attributes, denoted “ b .” That is, bundle $a = (c, b)$. As above, $u^{R_i}(\cdot)$ is a vNM utility function representing preferences R_i .¹³¹

¹²⁹ By contrast, the variations on UUU and EPP that take the form $\sum_{z \in \mathbf{Z}} \pi(z) F\left(\sum_{i=1}^N w_i(x^{P,z})\right)$ or

$\sum_{z \in \mathbf{Z}} \pi(z) F\left(\sum_{i=1}^N g(w_i(x^{P,z}))\right)$ cannot be thus expressed. For these SWFs, it is not sufficient to know which attribute-

bundle lottery each individual faces; we also need to know how the individual lotteries are correlated.

¹³⁰ See *supra* text accompanying notes 79-81.

¹³¹ However, because CBA does not use interpersonal well-being information, $u^{R_i}(\cdot)$ can be any one of the family of vNM functions representing R_i , related by positive affine transformations. There is no need to pick scaling factors, since the determination of equivalent variations is independent of which $u^{R_i}(\cdot)$ is chosen.

Consider first the case of certainty. Pick one outcome s as the status quo. For any outcome x , individual i 's equivalent variation for x , denoted $EV_i(x)$, is defined as follows: $u^{Ri}(c_i^s + EV_i(x), b_i^s) = u^{Ri}(c_i^x, b_i^x)$.¹³² CBA now uses the following rule to rank outcomes:

$$x \text{ at least as good as } y \text{ iff } \sum_{i=1}^N EV_i(x) \geq \sum_{i=1}^N EV_i(y) .$$

We use the Savage framework, as above, to represent uncertainty. Each policy (P)—including a status quo policy (s)—is a mapping from states of nature to outcomes (and therewith attribute bundles for each individual). For a given policy P , $EV_i(P)$ is the change to status quo consumption in every state that make i 's expected well-being equal to her expected well-being with the policy. That is, $EV_i(P)$ is such that:

$$\sum_{z \in Z} \pi(z) u^{Ri}(c_i^{s,z} + EV_i(P), b_i^{s,z}) = \sum_{z \in Z} \pi(z) u^{Ri}(c_i^{P,z}, b_i^{P,z}) .$$

The CBA rule becomes: P at least as good as P^* iff $\sum_{i=1}^N EV_i(P) \geq \sum_{i=1}^N EV_i(P^*)$

C. *Simulation Model*

In the main analysis, there are 25 cohorts, defined by a combination of current age (meaning integer age, i.e., years of life lived at last birthday) and income. It's assumed that there are an equal number of individuals in each cohort. The age groups are 20, 30, 40, 50, and 60. The income amounts for the 10th, 30th, 50th, 70th, and 90th percentile incomes of the U.S. population were, respectively, \$15,373, \$27,232, \$38,149, \$54,024, and \$122,542. These were derived from the most recent CBO report on the distribution of household income, by taking the average after-tax household income amounts for each quintile, and dividing by a normalization factor to move from household to individual income.¹³³ Income mobility was ignored. For each given individual i , it is assumed that she earns an income amount y_i in each year of her life, if she survives to the end of that year; y_i is one of the five amounts just stated.

Each individual is assumed to face a “survival curve,” consisting of survival probabilities for each year of life. I here use the term “survival probability” for a given year of life to mean the individual's conditional probability of surviving to the end of the year, given that she is alive at the beginning; “mortality risk” is one minus survival probability.

¹³² To the extent EVs are modelled as depending upon market prices, these can be incorporated in the b attributes.

¹³³ Congressional Budget Office, *The Distribution of Household Income and Federal Taxes 2013* (June 2016). The average amounts for the five quintiles are \$24,500, \$43,400, \$60,800, \$86,100, and \$195,300. The normalization factor was the square root of 2.54, since the average household size for 2013 was 2.54. The CBO report suggests the square root of household size as a normalization factor. See id. at 27-28.

The status quo survival curve for individuals of the five current ages at 50th percentile income was based on the most recent U.S. life table.¹³⁴ Annual mortality rates for the other income groups were adjusted up or down by a multiplicative factor. These factors were, respectively, 1.5, 1.2, .9, and .7 for the 10th, 30th, 70th, and 90th income percentiles. The factors were loosely based upon recent work by Chetty and co-authors, looking at life expectancy for different income groups.¹³⁵ The upshot is the following life expectancies at birth for individuals of the different income groups (starting with the 10th): 74, 77, 79, 81 and 83 expected years of life.

It was assumed that individuals have common preferences; that lifetime utility is the sum of annual utility; and that annual utility is logarithmic in consumption. An individual engages in consumption during a given year of life only if she survives to the end of that year. Let c_i^l denote the consumption of individual i during year-of-life l , if she survives to the end of that year. c_i^l is equal to the individual's annual income (y^i), plus or minus adjustments (either changes in consumption caused by a policy or hypothetical changes to consumption considered to calculate equivalent variations for purposes of CBA).

In order to allow comparisons of lives of different lengths, a zero point of annual consumption needs to be chosen.¹³⁶ That was set at \$1000, based roughly on the World Bank

¹³⁴ U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, *National Vital Statistics Report* (vol. 65, no. 8, Nov. 28, 2016), Table 1.

¹³⁵ Chetty et al., *supra* note 76. Specifically, the multiplicative factor for a given income percentile was chosen so that, in the model, the ratio between life expectancy at age 40 for individuals in that percentile, and life expectancy at age 40 for 50th percentile individuals, was approximately equal to the ratio as estimated by Chetty. See *id.* at 1754, Figure 2.

¹³⁶ Assume that we have a family of eligible period (annual) utility functions, unique up to a positive affine transformation. All functions in the family equally well represent levels and differences of period utility. Moreover assume, as here, that lifetime utility is the sum of period utility numbers.

Now, let $u(\cdot)$ and $u^*(\cdot)$ both be eligible period utility functions—so that $u^*(\cdot) = ru(\cdot) + s$, r positive. Then

the two lifetime utility formulas $\sum_{l=1}^L u(c_l^i)$ and $\sum_{l=1}^L u^*(c_l^i)$ need not produce the same ranking of a set of life

histories, if this includes lives of different lengths. Thus the choice between $u(\cdot)$ and $u^*(\cdot)$ is not a matter of indifference for UUU (applied to lifetime utilities), let alone for EAP or EPP.

Within the family of eligible period utility functions, how shall we choose the “right” one for purposes of our calculation of lifetime utility? Note that the formula $\sum_{l=1}^L u(c_l^i)$ only sums over periods during which the

individual is alive; being dead during a period is assigned a period utility of 0. Thus, the right period utility function $u^+(\cdot)$ should have the following feature. Let c^{zero} be a consumption level so low as to be no better than death. Then $u^+(c^{\text{zero}}) = 0$.

$u^+(\cdot)$ can be specified by picking any one of the eligible period utility functions $u(\cdot)$, and then defining $u^+(\cdot)$ as follows. $u^+(c) = u(c) - u(c^{\text{zero}})$. By construction, then, $u^+(c^{\text{zero}}) = 0$. In the formula in the text above, $u(\cdot) = \ln(\cdot)$, and $c^{\text{zero}} = \$1000$.

Note that this approach automatically assigns a zero level of lifetime utility to any life in which the individual has c^{zero} in every period alive, or to a life of zero length. Note also that $u^+(\cdot)$ and $ru^+(\cdot)$, r positive, produce lifetime utility numbers that are ratio rescalings (by r) of each other. Thus by specifying $u^+(\cdot)$ in the manner just discussed, we have at the same time specified a lifetime utility function that is unique up to a ratio

level of extreme poverty.¹³⁷ Pulling all this together: If individual i dies at age L last birthday, with $L > 0$, her lifetime well-being is: $\sum_{l=1}^L (\ln(c_i^l) - \ln(1000))$. In the case of an individual who dies at age 0 last birthday, her lifetime well-being was set at 0.

In the model, a policy consists in a one-time change in mortality risk, specifically a change in morality risk during the current year, and a one-time change to consumption (again, a change to current consumption). The “current year” is, of course, the 21st year of life for individuals currently aged 20, the 31st year of life for individuals currently aged 30, etc. Because consumption occurs only if the individual survives the year, a “change to current consumption” means a conditional change: a change to the amount that the individual will consume in a given year if she survives to its end.

The formulas for UUU, EPP^{mod} , and EPP^{high} were used with lifetime well-being values as their inputs. The $g(\cdot)$ functions for EPP^{mod} and EPP^{high} were in the Atkinson family with γ set at, respectively, 1 and 2.

For purposes of CBA, an individual’s $EV_i(P)$ for a given policy P is the change to her current consumption that leaves her equally well off (in terms of lifetime well-being) as she is with the policy—given the combined change to her current mortality risk and current consumption produced by the policy. Equivalent variations were estimated using the following formula. Let Δc_i^P denote the policy change to the individual’s current consumption, and $\Delta \mu_i^P$ denote the change to her current survival probability. Let VSL_i denote the marginal rate of substitution between current survival probability and consumption, i.e., the individual’s “value of statistical life.” Then $EV_i(P) \approx \Delta c_i^P + (\Delta \mu_i^P) VSL_i$. For short, call this the “EV approximation formula.”¹³⁸ This type of formula is very standard in the literature on CBA.

The reasoning behind the EV approximation formula is as follows. Let U_i represent the status quo expected lifetime utility of individual i , currently of age l , as a function of her current consumption c_i^l , her current survival probability μ_i^l , and other attributes b_i ¹³⁹. Assume that a policy P changes c_i^l by Δc_i^P , and changes μ_i^l by $\Delta \mu_i^P$. Then, by the total differential approximation,

transformation—and thus adequate for UUU and EAP and EPP (using the Atkinson SWF) applied to lifetime utilities.

¹³⁷ That is, for purposes of the approach sketched in the footnote immediately above, $c^{\text{zero}} = \$1000$. The World Bank extreme poverty level is \$1.90/day. World Bank, *A Global Count of the Extreme Poor in 2012* (Policy Research Working Paper No. 7432, Oct. 2015). In the model, setting $c^{\text{zero}} = \$1000$ also yields a population average VSL of approximately \$9 million, which is a plausible level given current estimates of average VSL. Lisa A. Robinson and James K. Hammitt, *Valuing Reductions in Fatal Illness Risks: Implications of Recent Research*, 25 HEALTH ECON. 1039 (2016).

¹³⁸ See Adler, Hammitt and Treich, *supra* note 9.

¹³⁹ In our model, these other attributes are mortality risks and consumption amounts in other periods.

$$U_i(c_i^l + \Delta c_i^P, \mu_i^l + \Delta \mu_i^P, b_i) - U_i(c_i^l, \mu_i^l, b_i) \approx \frac{\partial U_i}{\partial c_i^l} \Delta c_i^P + \frac{\partial U_i}{\partial \mu_i^l} \Delta \mu_i^P$$

$EV_i(P)$ should be such that $U(c_i^l + EV_i(P), \mu_i^l, b_i) = U_i(c_i^l + \Delta c_i^P, \mu_i^l + \Delta \mu_i^P, b_i)$. We thus have that:

$$EV_i(P) \approx \Delta c_i^P + \frac{\partial U_i / \partial \mu_i^l}{\partial U_i / \partial c_i^l} \Delta \mu_i^P, \text{ with } VSL_i = \frac{\partial U_i / \partial \mu_i^l}{\partial U_i / \partial c_i^l}.$$

The formula provides a good approximation for EVs if policy consumption and survival probabilities are sufficiently small variations from the status quo. To illustrate the accuracy of this approximation, consider Table 13 immediately below. This compares, for 20 year olds, exact EVs (calculated to one decimal place) with the approximated EVs for a policy consisting in a 1-in-100,000 increase in current survival probability and a \$100 reduction in current consumption.

Table 13

20 year olds	10 th percentile <u>income</u>	<u>30th</u>	<u>50th</u>	<u>70th</u>	<u>90th</u>
Exact EV	-\$77.1	-\$48.7	-\$17.9	\$29.9	\$271.3
Approx.	-\$77.0	-\$48.6	-\$17.8	\$30.0	\$271.0

An immediate implication of the approximating formula is that CBA is neutral to shifts in current consumption between individuals. Assume that policy P^* is identical to policy P , except that individual i 's consumption increases by Δc , while individual j 's consumption decreases by the same amount. Then (using the approximating formula) $EV_i(P^*) - EV_i(P) = \Delta c$, while $EV_j(P^*) - EV_j(P) = -\Delta c$. Thus the sum of approximated EVs is the same for the two policies.

If EVs are calculated exactly, CBA need not be neutral to shifts in current consumption between individuals. Because any given individual i may have different non-consumption attributes (in the model here, mortality risks) with P and/or P^* , as compared to the status quo, it need *not* be exactly true that $EV_i(P^*) - EV_i(P) = \Delta c$. However, if all policies considered involve relatively small changes in non-consumption attributes relative to the status quo, the approximating formula *is* quite accurate and CBA will be nearly neutral to shifts in current consumption. Consider Table 14 below.

Table 14

	<i>Income percentile</i>				
	<i>10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	\$100.2	\$100.2	\$100.2	\$100.2	\$100.3
<i>30</i>	\$100.1	\$100.2	\$100.2	\$100.2	\$100.3
<i>40</i>	\$100.1	\$100.1	\$100.2	\$100.2	\$100.3
<i>50</i>	\$100.1	\$100.1	\$100.1	\$100.1	\$100.2
<i>60</i>	\$100.0	\$100.1	\$100.1	\$100.1	\$100.1

Explanation: Table 14 shows the difference (to one decimal point) between exact EVs for a policy consisting of a 1-in-100,000 risk reduction and a \$100 benefit, and exact EVs for a 1-in-100,000 risk reduction.

In the main analysis, CBA values were assigned to hypothetical policies using EVs calculated with the approximating formulas. Specifically, “textbook CBA” applied that formula with cohort-specific VSL values, while “government CBA” did so with population-average VSL value.

In Tables 6, 9 and 10, breakeven per capita cost amounts were calculated by finding the per capita cost such that a policy with the stipulated mortality-risk reduction, with that per capita cost, and with the stipulated cost incidence has the same UUU, EPP^{mod} , EPP^{high} , or CBA value (respectively) as the status quo.

Table 7A was calculated as follows. For each cohort, the difference in UUU value between a policy consisting in a 1/100,000 reduction in that cohort’s current mortality risk, and the status quo, was calculated. This cohort-specific change in UUU value was then divided by the corresponding such change for the 60-year-old/bottom income cohort. In table 8A, the difference in UUU value between a policy consisting in a \$1 increase in that cohort’s consumption, and the status quo, was calculated. This cohort-specific change in UUU value was then divided by the change in UUU value for a 1/100,000 risk reduction for the 60-year-old/bottom income cohort.

Tables 7B and 8B were calculated the same way using textbook CBA (with the approximating formula), while 7C and 8B used EPP^{high} . The analogous tables for $EPP^{moderate}$ is below.

Table 7D: EPP^{mod} Risk Reduction Values Table 8D: EPP^{mod} Values of Increased Income

	<i>Income percentile 10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	4.7	4.8	4.9	5.0	5.1
<i>30</i>	3.3	3.4	3.5	3.5	3.6
<i>40</i>	2.3	2.4	2.5	2.5	2.6
<i>50</i>	1.5	1.6	1.7	1.8	1.9
<i>60</i>	1.0	1.1	1.1	1.2	1.3

	<i>Income percentile 10th</i>	<i>30th</i>	<i>50th</i>	<i>70th</i>	<i>90th</i>
<i>Age = 20</i>	0.12	0.06	0.03	0.02	0.01
<i>30</i>	0.12	0.05	0.03	0.02	0.01
<i>40</i>	0.12	0.05	0.03	0.02	0.01
<i>50</i>	0.11	0.05	0.03	0.02	0.01
<i>60</i>	0.11	0.05	0.03	0.02	0.01

The values in these tables can be used to approximate the change in UUU and EPP values produced by a given policy, as against the status quo. Multiply each cohort’s change in mortality risk or income by, respectively, the value in the table. Sum up across cohorts and then multiply the sum by the change in UUU or EPP value (as the case may be) for a 1/100,000 mortality risk change to a 60-year-old/bottom income individual. The approximation is not precise; essentially we are again using the total differential approximation (now for UUU and EPP).¹⁴⁰ The policy change in CBA value (the sum of EVs) as calculated using these tables is just the change in CBA as per the approximating formula for EVs—not the change with EVs calculated exactly.

In an alternative analysis, it is assumed that each income-age cohort has a subset of “fragile” individuals. These are individuals with the cohort-specific mortality rates scaled up by a factor of 2. (This assumption produces life expectancies at birth of fragile individuals in the 10th, 30, 50th, 70th, and 90th income percentiles of, respectively, 65, 68, 71, 72, and 75 years.)

D. *Hidden Attributes and Potential Pareto Superiority: An Illustration*

The following case uses the simulation model above to illustrate how hidden/nontaxable attributes can produce a divergence between CBA and potential ex ante Pareto superiority. Although option *a* passes a CBA test relative to option *b*, there is no feasible scheme of taxes and transfers that, coupled with *a*, renders it ex ante Pareto superior to *b*.

Assume that the population is divided into age/income cohorts, with the specific mortality rates and income amounts described above. Further (as in the alternative analysis above), assume that each cohort has two subcohorts, one of fragile individuals, the other of non-fragile individuals. Fragile individuals have the characteristic described above: a higher mortality rate as compared to non-fragile (healthy) individuals of the same age and income. To make the illustration simple, assume an equal number of fragile and non-fragile individuals in each of the 50 subcohorts (hence an equal number in the whole population).

¹⁴⁰ See *supra* note 86.

Imagine a policy that reduces the current mortality risk of fragile individuals by 1 in 100,000. The average EV for this risk reduction, among fragile individuals, is approximately \$77. Thus, if the policy produces a one-time per capita consumption cost of less than \$77/2 among the population as a whole, it passes a CBA test relative to the status quo. However, there are some policies with less than \$77/2 per capita costs that will fail a utilitarian test (UUU) relative to the status quo.¹⁴¹

In particular, imagine that the policy produces a 0.07% one-time reduction in each person’s income. Hence the cost incidence of the policy is as follows (rounding to the nearest dollar)—\$11 for 10th percentile income individuals, \$19 for 30th percentile individuals, \$27 for 50th percentile individuals, \$38 for 70th, and \$86 for 90th. Per capita costs for the population as a whole are approximately \$36 < \$77/2, and the policy passes a CBA test. However, UUU prefers the status quo to this policy.

Table 15 shows EVs for the policy. As can be seen, most of the fragile subcohorts are net beneficiaries from the policy, while (of course) all of the non-fragile subcohorts are cost bearers. The sum of EVs across the population is positive.

Table 15

	10 th percentile income		30 th		50 th		70 th		90 th	
	Nonfragile	Fragile	Nonfrag.	Frag.	Nonfrag.	Frag.	Nonfrag.	Frag.	Nonfrag.	Frag.
20 year olds	-\$11	\$9	-\$19	\$25	-\$27	\$45	-\$38	\$75	-\$86	\$241
30	-\$11	\$5	-\$19	\$17	-\$27	\$32	-\$38	\$56	-\$86	\$186
40	-\$11	\$1	-\$19	\$9	-\$27	\$19	-\$38	\$36	-\$86	\$131
50	-\$11	-\$2	-\$19	\$2	-\$27	\$8	-\$38	\$18	-\$86	\$80
60	-\$11	-\$5	-\$19	-\$4	-\$27	-\$2	-\$38	\$2	-\$86	\$34

Assume, now, that the “fragility” attribute is hidden. Government and, specifically, the officials who administer the tax and transfer laws are not able to observe whether individuals have this characteristic. (Imagine, for example, that “fragility” is caused by a genetic defect that

¹⁴¹ This is because the fragile actually have a lower marginal utility of income in the simulation model, as compared to non-fragile individuals of the same age and income. Thus a policy that shifts income from nonfragile to fragile will be neutral in CBA terms but a loss in overall well-being; and some range of policies that are positive in CBA terms will also be a loss in overall well-being.

the individuals themselves may not be aware of.) Thus an individual's status as fragile/non-fragile is not a characteristic that the tax and transfer laws can take account of. Taxes and transfers can't be differentiated on the basis of this attribute. Call this the non-differentiation constraint.

With fragility hidden, is there a feasible scheme of taxes and transfers that, coupled with the policy, renders it ex ante Pareto superior to the status quo? Even assuming that the administrative costs of taxation are zero, and that tax can be levied on each age/income cohort without deleterious incentive effects (labor supply), the answer seems to be "no." Because non-fragile individuals are worse off with the policy, they must receive positive transfers. (Specifically, non-fragile individuals at 10th, 30th, 50th, 70th, and 90th percentile income must receive a transfer of at least \$11, \$19, \$27, \$38 and \$86, respectively.) But the non-differentiation constraint means that the amount paid or received under the tax/transfer scheme by the fragile group in each age-income cohort must be the same as the amount paid or received by the non-fragile group. Thus *fragile* individuals in the 10th, 30th, 50th, 70th, and 90th percentile must *also* receive positive transfers of at least \$11, \$19, \$27, \$38 and \$86, respectively. So the only tax/transfer scheme, consistent with the non-differentiation constraint, that compensates the non-fragile individuals for their losses from the policy must make a positive payment to everyone (fragile and non-fragile alike) in each of the age/income cohorts. A tax/transfer scheme that pays positive amounts to everyone in the population is not feasible, unless the money to fund the transfers somehow magically appears from outside the population.