

# Assessing the Distribution of Impacts in Global Benefit-Cost Analysis

Lisa A. Robinson and James K. Hammitt  
with supplement by Matthew D. Adler

March 2018

Guidelines for Benefit-Cost Analysis Project  
Working Paper No. 3

Prepared for the Benefit-Cost Analysis Reference Case Guidance Project  
Funded by the Bill and Melinda Gates Foundation

Visit us on the web: <https://sites.sph.harvard.edu/bcaguidelines/>

## Preface and Acknowledgements

The Bill and Melinda Gates Foundation is supporting the development of guidelines for the economic evaluation of investments in health and development, particularly in low- and middle-income countries (“Benefit-Cost Analysis Reference Case: Principles, Methods, and Standards,” grant number OPP1160057). These guidelines will supplement the existing international Decision Support Initiative (iDSI) reference case, which provides general guidance on the overall framework for economic evaluation as well as specific guidance on the conduct of cost-effectiveness analysis.

This working paper is part of a series of methods papers and case studies being conducted to support the extension of the reference case to include benefit-cost analysis. The methods papers were reviewed by selected experts, posted online for public comment, discussed in a November 2017 workshop at Harvard University, then finalized. Although the methods papers will provide the basis for the benefit-cost analysis reference case guidance, the reference case may ultimately deviate from their recommendations in some cases.

This paper was prepared by Lisa A. Robinson and James K. Hammitt (Harvard T.H. Chan School of Public Health, Center for Health Decision Science and Center for Risk Analysis); Matthew Adler (Duke University) prepared the supplement on social welfare functions. Excellent research assistance was provided by Lucy O’Keeffe. For their careful review and many helpful comments, we thank David Greenberg, Neal Fann, Sam Harper, Julian Jamison, Paul Kelleher, Carol Levin, Ole Norheim, and members of the London School of Hygiene and Tropical Medicine Economic Evaluation group (Michelle Remme, Catherine Pitt, Sedona Sweeney, Tom Drake, Fiammetta Bozzani, and Zia Sadique). For many helpful discussions, we also thank the other members of our Leadership Team (Dean Jamison and David de Ferranti), our Advisory Group members (<https://sites.sph.harvard.edu/bcaguidelines/advisory-group/>), and our Gates Foundation Program Officer (David Wilson) as well as the participants in the November workshop.

More information on the project is available at <https://sites.sph.harvard.edu/bcaguidelines/>

## Executive Summary

There is widespread agreement that benefit-cost analyses should be supplemented with information on how the impacts are distributed across individuals with different characteristics, such as varying income levels, yet reviews of completed analyses suggest that such information is rarely provided. The goal of this paper is therefore relatively simple: to encourage analysts to provide information on the distribution of net benefits throughout the population in addition to assessing the overall impact of the policy.

Conventionally, benefit-cost analysis focuses on economic efficiency, summing the values of a policy's costs and benefits based on the preferences of those affected. Decision-makers and other stakeholders typically find this information useful but insufficient; they also want to know who is harmed, who is helped, and by how much.

Responding to this question requires first identifying the characteristics of individuals and impacts of most concern, which will vary depending on the policy and decision-making context. In addition to income level, individuals of concern may be defined by attributes such as health status, geographic location, and educational attainment. Similarly, impacts of concern may relate to changes in health, longevity, education, and environmental conditions, as well as income and other contributors to wellbeing. At minimum, decision-makers and other stakeholders are often interested in the distribution of changes in health, longevity, and disposable income across income groups.

The starting point for distributional analysis is the assessment of social benefits and costs; i.e., the analysis of economic efficiency. The types of benefits and costs addressed will depend on the policy. For example, for health-related policies, the benefits analysis typically provides data on the net changes in the risk of mortality as well as of nonfatal illnesses and injuries attributable to the policy -- expressed as numbers of statistical cases averted. The distributional analysis then describes how these cases are distributed across individuals grouped by the attribute(s) of concern, such as differing income levels or geographic locations. While the benefit-cost analysis typically values these benefits using population-average estimates, for distributional analysis these estimates should be adjusted to reflect the values held by individuals with differing characteristics to the extent possible.

Assessing the distribution of costs is often more difficult, requiring that the analyst estimate the effect on disposable income. When costs are borne directly by individuals and households, estimating this distribution may be relatively straightforward. When costs are borne initially by the government, industry, donors, or other organizations, assessing the effects on individuals involves additional steps. For government programs, the analyst first must estimate how the costs translate into changes in taxes or user fees or are otherwise financed, prior to determining how these impacts are distributed across individuals. For programs operated by nonprofit or for-profit organizations, the analyst must first estimate how these costs translate into changes in unit prices (which have both income and substitution effects on consumer expenditures), in wages paid to employees, and in returns to capital that accrue to owners. Costs paid by external donors (e.g., aid from foreign governments or foundations) may have no

or little direct, immediate impact on resource allocation within the country. However, the donor agency may be interested in estimating how these costs would be distributed if the policy were instead funded using in-country resources as well as its longer term economic impacts.

Once the distribution of costs and benefits across individuals is estimated, the costs can be subtracted from the benefits to determine the distribution of net benefits. The results should be reported in tables, charts, or graphics that indicate the costs, benefits, and net benefits that accrue to individuals at different points in the income or other distribution; e.g., in each quintile. This distribution also can be described using standard inequality metrics, such as the Gini coefficient, concentration index, or Atkinson index. Each of these metrics has advantages and limitations, but can provide useful summary information. Clear communication of the implications and associated uncertainties is essential.

Scholars have developed approaches to more fully integrate the analyses of distribution and efficiency. These include conducting the benefit-cost analysis with distributional weights that reflect estimates of societal preferences, or using a social-welfare function to represent preferences for both the level and distribution of wellbeing. While promising, more work is needed to develop consensus on the approach, including the functional form and parameter values needed for practical application.

In the near-term, all benefit-cost analyses should at minimum describe the distribution of net benefits as follows.

- 1) Identify the individuals and impacts of concern.** In consultation with decision-makers and other stakeholders, analysts should identify the characteristics of individuals and impacts of concern. At minimum, the distributional analysis should address the effects of the policy on the health, longevity, and income of members of different income groups, including the distribution of net benefits. Analysts should consider whether other impacts and other groupings also should be addressed given the policy and decision-making context.
- 2) Determine the level of detail and degree of quantification for the distributional analysis.** The effort devoted to the distributional analysis should be proportionate to its importance for decision-making. “Importance” may depend on the likely magnitude of the distributional impacts and concerns about associated inequities; it may also depend on the need to respond to questions likely to be raised by decision-makers and others. The extent to which the analysis is quantitative or qualitative, as well as its level of detail, should reflect the relative importance of these concerns and take into account the data, time, and resources available.

Screening analysis that relies on easily accessible data is often useful in providing preliminary information on the possible direction and magnitude of the effects, and in informing decisions about future work. Depending on the results, this screening may be followed by more detailed assessment that involves collecting additional data, refining the methods used, and expanding the scope of the analysis.

Both the benefit-cost analysis and the distributional assessment should be conducted concurrently with the policy development process, so that the preliminary results can be used to inform the options to be considered as well as future analytic steps.

- 3) Describe the distribution of benefits, costs, and net benefits across members of different population groups.** For each policy option assessed, analysts should report the results of the distributional analysis in text, tables, graphics, or other forms. These results should be reported as monetary values and in physical terms to the extent possible; e.g., as net benefits and as expected numbers of deaths, illnesses, and injuries averted. Measures of inequality, such as the Gini coefficient, may also be used; the advantages and limitations of the selected measure(s) should be discussed along with the results. The ultimate results of the distributional analysis should be reported in a clearly labeled, separate section of the documentation, that describes the data sources and methods used as well as the results.

All analyses should address the distribution of net benefits (benefits minus costs) to members of each population subgroup (e.g., each income quintile), to ensure that potentially counterbalancing impacts are not ignored. In addition, analysts should explore the extent to which there is heterogeneity in the net benefits that accrue across members of each group; e.g., because some members are more vulnerable to the health condition than others.

- 4) Address uncertainty:** Analysts should assess uncertainties in the estimates both qualitatively and quantitatively; e.g., by conducting sensitivity or probabilistic analysis, discussing the quality of the evidence, and indicating the implications for decision-making.

While the screening analysis under step 2 may lead to the conclusion that qualitative discussion will suffice, in most cases quantitative data on distribution will be more informative and desirable. Where data limitations or other issues limit the ability to quantify the effects, analysts should use “what if” or bounding analysis to explore the likely distribution of net benefits.

Over the long term, additional work is needed to provide examples of how to assess the distribution of the impacts of different types of policies, recommendations for the application of specific inequality metrics, and options for distributional weighting using social welfare functions.

- 1) Develop case studies.** Assessment of distribution may be complicated and difficult in many contexts. Case studies that illustrate best practices for conducting such assessments will substantially aid in encouraging the more routine completion of these analyses. These case studies could, for example:
- a. Explore the data available in different contexts and the options for qualitative and quantitative assessment.
  - b. Examine how to best communicate the analytic approach, the results, and the implications to diverse decision-makers and other stakeholders.
  - c. Investigate the use and usefulness of distributional information in decision-making, to identify needed analytic improvements.

- 2) Explore inequality metrics.** As noted earlier, inequality metrics such as the Gini coefficient can provide useful summary measures of the degree of equality of the policy outcome. More criteria-driven review of the potential applicability of these indices in global benefit-cost analysis, and of their normative foundations, is needed to develop recommendations for the indices that are most appropriate across different policy and decision-making contexts.
- 3) Investigate the use of social welfare functions.** Distributional analysis using social welfare functions is a useful supplement or complement to traditional benefit-cost analysis accompanied by descriptive information on distribution. More work is needed to further develop parameter values and to provide practical advice on the implementation of this approach in global benefit-cost analysis.

Estimating the distribution of net benefits allows decision-makers and others to weigh the extent to which benefits and costs are counterbalancing for each group as well as the overall distribution of net benefits across groups. The distributional analysis then makes the trade-offs between economic efficiency and distributional concerns explicit. Decision-makers may choose an economically-efficient policy that maximizes net benefits, perhaps addressing distributional concerns through other policies, or may choose a less efficient option to ameliorate distributional impacts or achieve other policy goals.

## Table of Contents

Preface and Acknowledgements .....	2
Executive Summary.....	3
1.0 Introduction .....	8
2.0 Basic Concepts .....	11
3.0 Methods for Describing Distribution .....	14
3.1 Estimating the Distribution of Benefits.....	15
3.2 Estimating the Distribution of Costs .....	16
3.3 Describing the Distribution of Net Benefits .....	18
4.0 Summary and Recommendations.....	23
3.1 Near-Term Recommendations.....	23
3.2 Long-Term Recommendations.....	25
Supplement: Social Welfare Functions .....	26
References .....	31

## 1.0 Introduction

There is widespread agreement that benefit-cost analyses should be supplemented with information on how the impacts are distributed across different members of the population. This agreement is reflected in the international Decision Support Initiative (iDSI) Reference Case (Wilkinson et al. 2016), which provides general guidance on the economic evaluation of policies in low- and middle-income countries, as well as in numerous guidance documents that address the conduct of benefit-cost analysis in other settings (see Robinson et al. 2017).<sup>1</sup> Yet reviews of completed analyses suggest that such information is rarely provided. Analyses labelled as “distributional” often focus on the allocation across sectors (e.g., industry, government, taxpayers) rather than across individuals; yet the primary concern of many involved in policy development is understanding how individuals with different attributes are affected.

For example, in Robinson et al. (2017), we review selected benefit-cost analyses conducted in low- and middle-income countries, and find that little information is provided on the distribution of impacts across individuals. Similarly, in Robinson, Hammitt, and Zeckhauser (2016) we find little such information in analyses of U.S. regulations, despite government-wide guidance requiring assessment. The reason for this inattention is unclear; it may reflect the difficulties associated with conducting these analyses, issues related to the decision-making context, or other challenges. Regardless, without more and better analyses of distribution, we lack the information needed to explore the potential usefulness and implications of the results.

The goal of this paper is therefore relatively simple: to encourage analysts to provide information on the distribution of net benefits in addition to assessing the overall impacts of the policy. It is one of a series of methods papers which will ultimately be used to expand the iDSI Reference Case to address the conduct of benefit-cost analysis. More information on this project is provided at: <https://sites.sph.harvard.edu/bcaguidelines/>.

In this introductory chapter, we briefly summarize the need for information on distribution. In the following chapters, we explore related concepts, discuss approaches for describing the distribution of impacts, and summarize our conclusions and recommendations. The supplement, authored by Matthew D. Adler, discusses the application of social welfare functions.

Conventionally, benefit-cost analysis focuses on the net impact of a policy on social welfare, using market prices or nonmarket valuation methods to estimate the preferences of the individuals affected for the harms and improvements they would likely experience. The values held by those affected are then summed to estimate the net benefits of the policy; i.e., its economic efficiency.

Decision-makers and other stakeholders typically find this information useful but insufficient; they also want to know who is harmed, who is helped, and by how much. Do the benefits primarily affect the

---

<sup>1</sup> In this and other papers prepared for this project, we use the term “policy” generically to refer to all types of programs and interventions that may be the focus of a benefit-cost analysis.

disadvantaged while the costs primarily affect the advantaged? Or vice-versa? Or do both the benefits and costs accrue to the same group? What is the relative magnitude of the impacts? Answering these questions provides information that can be useful in determining whether and how to tailor the policy to address distributional concerns as well as in choosing among the policy options.

To respond to these questions, analysts must first identify the types of individuals and impacts of concern. Individuals of concern may be defined by attributes such as income, health status, geographic location, educational attainment, and so forth.<sup>2</sup> Similarly, impacts of concern may relate to changes in income, health, longevity, education, environmental conditions, and other contributors to wellbeing.

In this paper, we focus on the changes in income and in health and longevity that accrue to members of different income groups. While this focus in part reflects the impacts and groups that are often of greatest concern, it is largely for ease of exposition. We recognize that analysts, decision-makers, and other stakeholders are likely to wish to address other groupings and other impacts within the context of particular analyses.

In our reviews of completed benefit-cost analyses, we have found that – if distribution is considered at all – the analysis is usually incomplete. Often, it considers the initial distribution across broad sectors (e.g., industry or government, or all beneficiaries) rather than across individuals with different attributes.<sup>3</sup> There is a tendency to focus on narrow measures, such as impacts only on those living below the poverty line rather than across all income groups, or on particular types of impacts, such as health improvements without consideration of costs, rather than the net effects of the policy.

Such narrow focus is problematic. First, any dividing line raises difficult questions about the rationale both for choosing that threshold and for ignoring impacts on those who are above it, perhaps only by a small amount. Arguments about defining the threshold can also divert attention and analytic effort away from more fundamental and important tasks related to estimating and evaluating the distribution.

Second, focusing on only a subset of the population (such as the poor) or only a subset of the impacts (such as changes in health risks) ignores the implications of the overall distribution for policy design and

---

<sup>2</sup> This question is closely related to the issue of standing (or perspective); i.e., determining whose preferences are counted in the benefit-cost analysis, as discussed in Robinson et al. (2017). However, in determining standing, the analyst is concerned with delineating the total population to be addressed, whereas in assessing distribution, the analyst is concerned with defining subgroups across that population.

<sup>3</sup> For example, when assessing major U.S. regulations, agencies are required to assess impacts on small businesses and on state, local, and tribal governments and the private sector under the Regulatory Flexibility Act and the Unfunded Mandates Act (U.S. Office of Management and Budget 2003, U.S. Department of Health and Human Services 2016). In analyses of U.S. social programs, impacts are at times assessed separately for two groups, program beneficiaries and other members of society (at times referred to as “taxpayers” or subdivided into non-program participants and the government), as well as for society as a whole (Long et al. 1981, Greenberg and Cebulla 2008, Greenberg et al. 2009). Such analyses do not report the distribution across individuals (e.g., at different income levels or with attributes other than program participation), which is the focus of this paper, but provide important information and a useful starting point for more disaggregate analysis.

decision-making. For example, a policy redesign that shifts costs from the wealthiest to middle-income groups may be less desirable than the opposite, even if in both cases the poor receive the majority of the benefits. An analysis that only considers impacts on the poor would not help inform this type of redesign decision.

Third, disregarding the net effects provides incomplete information on how the policy affects the wellbeing of different groups. For example, considering only health gains that accrue to the poor ignores the extent to which they also bear the costs of achieving these gains. Any distributional effect involves both “from” and “to” sides of the equation; who gains may be as important as who loses. Thus throughout this paper, we focus on approaches for estimating the full distribution of the net impacts -- both positive and negative; e.g., costs and benefits that accrue to all income quintiles rather than solely to the poor.

Estimating the distribution of net benefits allows decision-makers and others to weigh the extent to which benefits and costs are counterbalancing for each group as well as the overall distribution of net benefits across groups. The distributional analysis then makes the trade-offs between economic efficiency and distributional concerns explicit. Decision-makers may choose an economically-efficient policy that maximizes net benefits, perhaps addressing distributional concerns through other policies, or may choose a less efficient option to ameliorate distributional impacts or achieve other policy goals.

## 2.0 Basic Concepts

Benefit-cost analysis is based on two fundamental elements: the notion that each individual is the best (or the most legitimate) judge of how a change in policy or other circumstances affects his or her wellbeing, and a method to compare improvements for some people against harms (or forgone improvements) to others. The first element focuses attention on how a policy affects an overall measure of wellbeing, typically summarized by an individual's utility (evaluation of wellbeing).<sup>4</sup> The second element addresses the circumstances under which it is appropriate to adopt a policy that enhances the wellbeing of some individuals while at the same time diminishing the wellbeing of others.

The conventional normative basis for using benefit-cost analysis in decision-making begins with the Pareto Principle: a policy is desirable if it makes someone better off and no one worse off.<sup>5</sup> While attractive in theory, few policies meet this criterion. Almost any policy will harm at least some people; for example, by raising the prices they pay by more than the value of the benefits they receive. To address this limitation, the Kaldor-Hicks criterion was developed. This criterion suggests that a policy is desirable if it makes the "winners" (who gain from the policy) better off by a large enough amount that they could compensate the "losers" (who are harmed by the policy), and that it should be rejected if the losers could compensate the winners for not pursuing the policy -- assuming compensation is costless.<sup>6</sup> This compensation is hypothetical; the criterion does not demand that compensation be paid or even contemplated.

Benefit-cost analysis can be used to address whether a policy meets this criterion.<sup>7</sup> Under the standard (neoclassical) model, some economists argue that decisions on policies designed to improve health or achieve other social goals should be based solely on economic efficiency, to ensure that resources are invested in those activities that produce the greatest benefit. Supporters of this argument note that distributional goals can be achieved more comprehensively and effectively, at lower cost, by transferring money through the tax system or programs that provide supplementary income (e.g., Kaplow and Shavell 2006).<sup>8</sup> Money transfers can be targeted on the outcome and the population of concern; for example, by distributing income directly to the poor. Other types of policies often provide more

---

<sup>4</sup> This conception of wellbeing tends to be self-interested, taking little account of the interactions in wellbeing between individuals. The way in which benefit-cost analysis should incorporate altruism is subtle and not fully resolved; in part, it depends on whether the altruism is "pure" (the altruist cares only about other people's self-assessed wellbeing including both the costs and the benefits they accrue) or "paternalistic" (the altruist cares about some aspects of other people's wellbeing but not others, e.g., cares about their health but not about the pleasure or satisfaction they obtain from an unhealthful activity).

<sup>5</sup> See Brouwer et al. (2008), Adler (2012), Nyborg (2012), and Hammitt (2013) for more discussion of this and other normative framings.

<sup>6</sup> If more than one policy provides positive net benefits, the preferred choice is the one that yields the greatest net benefits.

<sup>7</sup> In reality, because not all outcomes can be feasibly quantified and valued and the results are uncertain, benefit-cost analysis provides related insights but does not necessarily resolve this question.

<sup>8</sup> Such transfers are not costless. They involve administrative costs and may affect behavior; for example, by discouraging employment.

heterogeneous benefits to more heterogeneous populations. This approach may be described as expanding the “social welfare pie” (the total quantity of goods and services available), potentially enabling everyone to consume more. The distribution and possible reallocation of the pie can, in principle, be evaluated and addressed separately.

However, constantly tweaking the tax and income support system to compensate for inequities introduced by other policies is clearly impossible. As a result, decision-makers and other stakeholders generally desire information on distribution that can be considered along with information on net benefits in deciding whether and how to implement a particular policy.<sup>9</sup>

Concerns about distribution generally reflect interest in both the equality and the equity of the outcomes. Equality is a relative concept that describes the distribution of a quantity (such as health or income) across individuals or groups. Equity involves judging the extent to which the distribution is fair or just. Typically, the guidance for benefit-cost analysis focuses on describing how impacts are distributed, leaving it up to the decision-maker to determine whether that distribution is equitable.

Two normative frameworks are frequently referenced in this context.<sup>10</sup> The first is utilitarianism, which focuses on measuring and maximizing utility rather than simply economic efficiency as represented by unadjusted monetary values. Benefit-cost analysis based on the sum of the unweighted costs and benefits does not take into account the likelihood that an incremental dollar received by a poor person yields a greater increase in wellbeing than the same amount received by a rich person; i.e., that the marginal utilities of income differ. Implementing a utilitarian approach requires assessing distribution by income level, then applying weights that reflect expected differences in marginal utility. This approach has been implemented in some contexts (see, for example, HM Treasury 2011), but such weights appear to be rarely used and more work is needed to determine the weights appropriate in different contexts.

Prioritarianism is similar but counts changes in the utility of individuals who are worse off as more important than comparable changes to individuals who are better off. It goes beyond utilitarianism by incorporating preferences for others’ wellbeing. Prioritarianism seems consistent with many people’s intuition about what might be just or fair. Implementing this framework involves a two-step process. First, the effects on individual utility must be estimated. Second, these utility measures must be transformed into measures of welfare. For this approach to be widely-used, more work is needed on how to define those who are worst off and how to weight alternative allocations.<sup>11</sup> This framework is discussed in more detail in the Adler supplement to this paper.

---

<sup>9</sup> Robinson, Hammitt, and Zeckhauser (2016) discuss additional reasons why such information may be useful.

<sup>10</sup> Examples of other frameworks include egalitarianism, in which case increased equality is valued no matter how it is achieved, and maximin, which ranks impacts based on the effects on the worst-off members of society. See, for example, Adler (2012), Eyal et al. (2013) and Adler and Fleurbaey (2016) for more discussion of these and other frameworks.

<sup>11</sup> The Prioritarianism in Practice research network is devoted to improving understanding of prioritarianism as a methodology for policy evaluation. See: <https://law.duke.edu/laweconomicsandpublicpolicy/conferences/pip/>.

To implement any normative framework, descriptive information on how costs, benefits, and net benefits are distributed is needed. A function that evaluates or weights impacts accruing to members of different population subgroups cannot be applied without this basic information.

In the discussion that follows, we focus on supplying this descriptive information for two reasons. First, such information is often missing from benefit-cost analyses; encouraging greater reporting is an important initial step in moving towards greater consideration of distribution. Second, development and application of weights requires substantial additional work, given the complexities of the issues. These include the lack of consensus on the appropriate normative framework and on how it can be best implemented in the context addressed by this project: i.e., health and development policies implemented in low- and middle-income countries. We return to this need for more work in the recommendations at the end of this paper.

### 3.0 Methods for Describing Distribution

The starting point for distributional analysis is the benefit-cost analysis; i.e., the analysis of economic efficiency. In the sections that follow, we first provide a general overview of the steps involved in estimating the distribution of the net benefits and discuss some options for describing the equality of their distribution. Our goal is to introduce the issues and options, indicating sources analysts can consult for more information. The detailed approach for such analyses will vary significantly depending on the decision-making context, the nature of the policy, the characteristics of its benefits and costs, the population groups of interest, and the data and other analytic resources available. As noted earlier, for ease of exposition we focus on the distribution of health and longevity and disposable income across individuals in different income groups, while recognizing that analysts, decision-makers, and other stakeholders often will be interested in additional impacts.<sup>12</sup>

In this discussion, we follow the categorization of benefits and costs introduced in our scoping report (Robinson et al. 2017) and elsewhere (e.g., U.S. Department of Health and Human Services 2016). Whether an impact is categorized as a “cost” or “benefit” is arbitrary and varies across analyses.<sup>13</sup> Here we define “costs” as the inputs or investments needed to implement and operate the policy – including real resource expenditures such as labor and materials regardless of whether these are initially incurred by government, firms, or individuals. Benefits are then the outputs or outcomes of the policy; i.e., the changes in welfare such as reduced risk of death, illness, or injury. Under this framework, counterbalancing effects should be assigned to the same category as the impact they offset. For example, “costs” might include expenditures on improved technology as well as any cost-savings that result from its use; “benefits” might include the reduction in disease incidence as well as any offsetting risks, such as adverse reactions to vaccines.

Transfer payments need not necessarily be included in benefit-cost analysis, but must be considered in distributional analysis. Transfers are monetary payments between persons or groups that do not affect the total resources available to society. The transfer itself is a benefit to recipients and a cost to payers, with zero net effect. For example, taxes and fees are usually transfer payments, and should be included in any distributional analysis.<sup>14</sup>

---

<sup>12</sup> We also focus on the initial consequences. Improved health is likely to affect future wealth (e.g., by allowing the individual to spend more time working), and increased wealth is likely to affect future health (e.g., by allowing the individual to live in a safer environment and pay for medical care). These subsequent consequences are discussed in more detail in Deaton (2013) and other sources and should also be considered by analysts. Economy-wide models, as discussed in the methods paper prepared by Strzepek, Amanya, and Neumann (2018) for this project, may also be useful in estimating the distribution of these broader impacts, particularly for policies with sizable effects.

<sup>13</sup> As long as the sign is correct (positive or negative), the categorization of an impact as a cost or a benefit will not affect the estimate of net benefits, but will affect the ratio of benefits to costs. If categorized inconsistently, benefit-cost ratios, total costs, and total benefits cannot be compared across analyses.

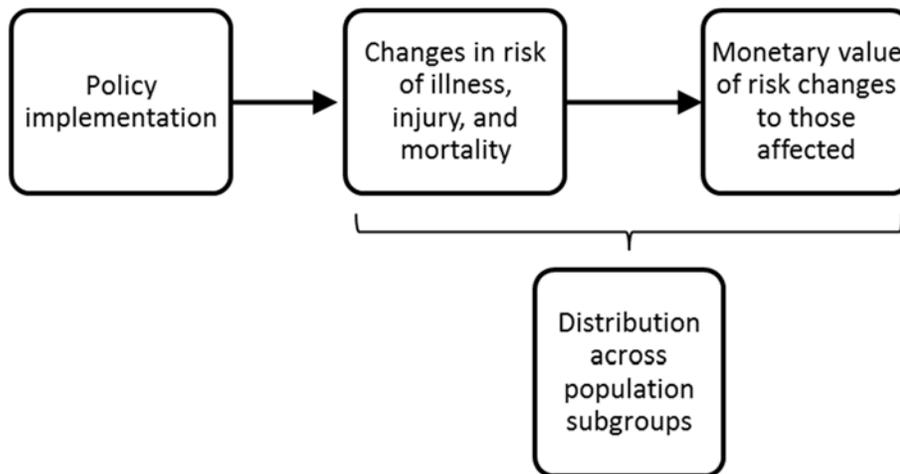
<sup>14</sup> While the transfers themselves often may be ignored in benefit-cost analysis, they may lead to behavioral changes that significantly affect resource allocation and the calculation of net benefits. Any such changes should

### 3.1 Estimating the Distribution of Benefits

In the case of health and longevity as well as other policy outcomes, there are several options for measuring the effects on individuals with different attributes, such as varying income levels. Analysts can count the number of statistical cases averted (by multiplying the average expected individual risk reduction by the number of people affected); use integrated nonmonetary measures to estimate the joint effect on health-related quality of life and longevity (such as quality-adjusted life years, QALYs, or disability-adjusted life years, DALYs); and use monetary measures that indicate the amount those affected would be willing to pay for the risk reductions. At minimum, distributional analysis of health-related policies should provide estimates of the number of averted cases of deaths, illnesses, or injuries across the members of groups of concern, and of the monetary value of these cases for each group to support the calculation of net benefits. (As noted earlier, in this paper we focus on health and longevity for simplicity; however, a similar process would be followed in the case of non-health benefits.)

This process is illustrated in Figure 3.1 and discussed below. This process should be iterative. Preliminary results from screening analysis, that uses easily accessible data and a reasonable range of assumptions, can be used to inform the policy development process in addition to informing decisions about future research and analytic steps.

**Figure 3.1. Distribution of Health Benefits**



The starting point for estimating the distribution of health effect incidence (as well as the distribution of other outcomes) is the benefit analysis, which provides estimates of the change in the risk of mortality as well as nonfatal illnesses and injuries attributable to the policy -- expressed as the number of statistical cases averted. The challenge is then to identify how these cases are allocated across members of the groups of concern. One question is whether the research design and available data allow the analyst to predict the effectiveness of the intervention across individuals with different attributes with

---

be included in the benefit-cost analysis. In addition, the benefit-cost analysis should include the resource costs of implementing the transfers, such as administrative costs and deadweight losses.

reasonable certainty. In the absence of such data, analysts will need to make assumptions to estimate this distribution and discuss the implications of associated uncertainties. The characteristics of the policy often aid in estimating this distribution. For example, if a vaccination program reduces the risks of tuberculosis, the vaccine is administered throughout the population, and the distribution of tuberculosis across different income or other groups is known, one default might be to assume the vaccine is equally effective across all members of the population. The risk assessment and disease modeling that supports the benefit-cost analysis will often provide related information. The underlying research is likely to also summarize or reference available data on populations that may be particularly sensitive or vulnerable to the risks as well as particularly responsive to the policy options under consideration.<sup>15</sup>

Once the likely number of averted statistical cases is estimated, the next step is to estimate the monetary value of these risk reductions. These values are discussed in separate methods papers that address the valuation of mortality risk reductions (Robinson, Hammitt, and O’Keeffe 2018) and nonfatal health risk reductions (Robinson and Hammitt 2018), and hence are not described in detail here. However, benefit-cost analysis typically uses population-average values. For distributional analysis, ideally analysts should adjust the values to reflect the preferences of individuals with differing characteristics. These values reflect the affected individuals’ willingness to trade-off spending on risk reductions for spending on other things that money could buy, and are limited by the total resources available to that individual. As a result, individual willingness to pay for these risk reductions is expected to decrease as income decreases, as well as vary depending on other attributes of the individual and the risk. Spending to accrue particular risk reductions means less money is available for other, potentially more essential, goods and services as well as for other risk-reducing activities.

Due to in part to widespread misunderstanding of how affordability affects valuation within the benefit-cost analysis context, analysts are often reluctant to use different values for changes in risks that accrue to different segments of the population.<sup>16</sup> These estimates do not reflect the value that the government or the analyst places on averting deaths, illnesses, injuries, or other disabilities; they instead reflect the value the affected individual places on the outcomes. Applying a population-average value to the risk reductions experienced by those in different income groups likely overstates the values held by poorer individuals and understates the values held by wealthier individuals. Relying on these averages obscures the value of the benefits that accrue to each group. It also can lead to misleading conclusions about the distribution of benefits as well as about the extent to which the benefits that accrue to each group are worth more or less than the costs each group incurs.

### **3.2 Estimating the Distribution of Costs**

In the case of costs (and off-setting savings), analysts are typically interested in the monetary expenditures needed to implement the policy and the ultimate effect on the disposable income of the

---

<sup>15</sup> For example, see Fann et al. (2011).

<sup>16</sup>For further investigation of related issues from an ethical perspective, see for example Fleurbaey and Schokkaert (2009).

groups of concern.<sup>17</sup> If costs are borne directly by individuals, the main challenge is determining how the costs are distributed across those who belong to different groups -- identified, for example, by income quintile.<sup>18</sup> If costs are borne initially by the government, industry, donors, or other organizations, assessing the effects on individuals requires additional steps. For government programs, the analyst first needs to estimate how the costs translate into changes in taxes or user fees or are otherwise financed, then estimate the incidence of these taxes or fees. For programs operated by nonprofit or for-profit organizations, the analyst must determine how costs are allocated among owners, workers, and consumers. This allocation will be affected by how the costs translate into changes in unit prices (which have both income and substitution effects on consumer expenditures), in wages paid to employees, and in returns to capital that accrue to owners. Costs paid by external donors (e.g., aid from foreign governments or foundations) raise other issues. In the short-term, these costs may have little or no impact on those affected by the policy. However, the donor agency may be interested in estimating how these costs would be distributed if the policy were instead funded using in-country resources.

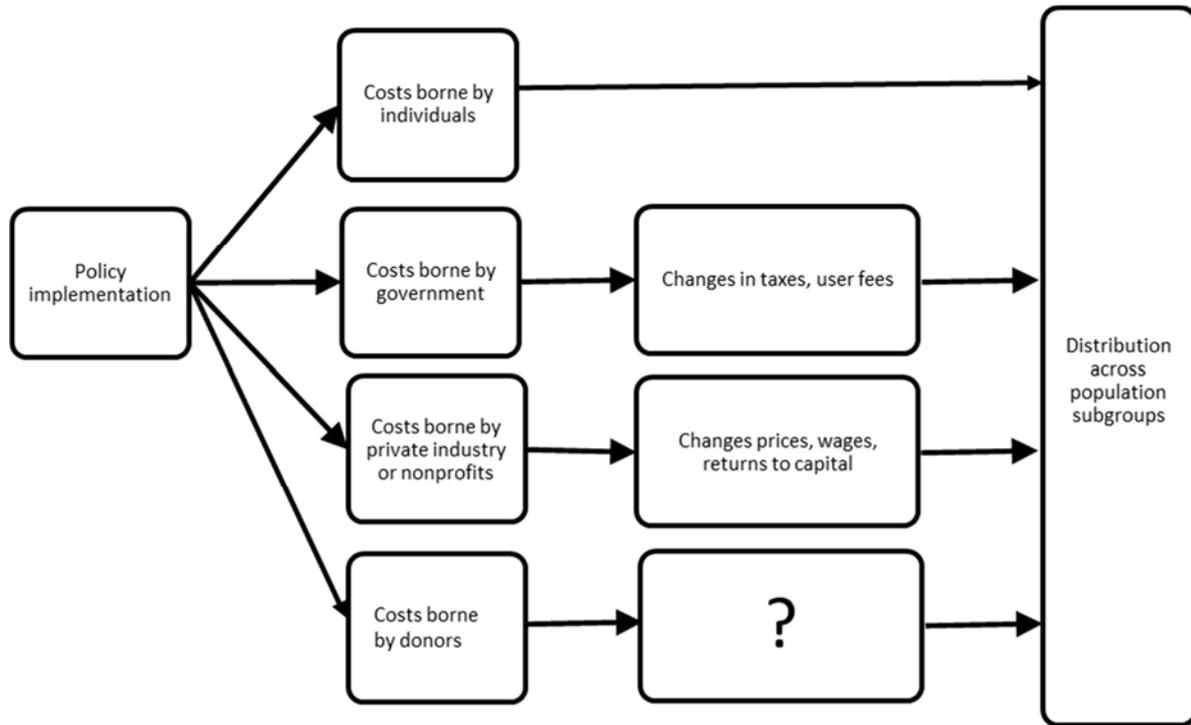
Figure 3.2 illustrates this process, which should be iterative. The initial results may have important implications for policy design as well as for decisions regarding how to best refine the analysis. As discussed in Strzepak et al. (2018), economy-wide modelling is often useful for estimating these effects, particularly for policies with relatively large impacts.

---

<sup>17</sup> There are many ways to define and measure income. For example, it can be defined at a particular point in time or over the individual's lifetime, and may include or exclude various types of payments (e.g., government subsidies as well as work-related earnings) and of wealth (e.g., investment income and real property). In general, more comprehensive measures are preferable to less comprehensive measures where feasible, so as to accurately reflect the resources available over time to the individuals affected by the policy.

<sup>18</sup> Consumer behavior will also affect the distribution of these costs. For example, if the price of a food is increased, some may substitute an alternative food. This substitution may affect both the costs and the benefits incurred, and such behavioral responses may vary across population groups.

**Figure 3.2. Distribution of Costs**



The estimation of these effects is more complicated than can be conceivably covered by this paper; analysts should consult economics and benefit-cost analysis texts for more information.<sup>19</sup> In many cases, detailed assessment may not be feasible and analysts may use “what if” or bounding analysis to explore the possible consequences. Such analysis uses the available data to explore the effects under different scenarios. For example, given what is known about the likely distribution of benefits, what would be the distribution of net benefits if all of the costs were allocated across members of the highest income group? Of the lowest income group? Or equally across all income groups?

### 3.3 Describing the Distribution of Net Benefits

Once costs and benefits are estimated for members of each group of concern, they can be combined to determine the distribution of net benefits. In addition to reporting the results for each group, analysts should explore the extent to which there is heterogeneity within the groups. For example, within an income quintile, some may be more vulnerable than others to a health hazard and may accrue a disproportionate share of the net benefits compared to others within the group.

At minimum, the results should be reported as a table, chart, or graphic that indicates the costs, benefits, and net benefits that accrue to individuals at different points in the distribution; e.g., to income quintiles, as illustrated by Table 3.1 below. For simplicity, we assume that policy only affects mortality; we provide a more detailed example in Robinson and Hammitt (2017).

<sup>19</sup> See, for example, Boardman et al. (2011).

**Table 3.1. Distribution of Net Benefits** (stylized example; numbers provided only for illustration)

<b>Income Range</b>	<b>Deaths Averted</b>	<b>Benefits</b> (value of deaths averted)	<b>Costs</b>	<b>Net benefits</b> (benefits minus costs)
\$0 – \$500	10	\$600,000	\$100,000	\$500,000
\$500-\$1,000	5	\$310,000	\$50,000	\$260,000
etc.	etc.	etc.	etc.	etc.
<b>Total</b>				

The types of data displayed in this table are an essential starting point for any distributional analysis. However, while such a table describes how net benefits are calculated and distributed, it does not provide a summary measure of the equality of the impacts.

Standard inequality indices can be used to summarize the distribution, and can be useful for comparing distributions. Several such measures are available, each of which has advantages and limitations. We briefly introduce some commonly-used indices and provide references for further information on their application and interpretation.

Note that selecting among these indices requires considering many factors that we do not discuss in detail, including practical concerns such whether the variables are ordinal or cardinal, and normative concerns such as whether those involved are interested in absolute or relative measures of inequality.<sup>20</sup> In addition, in some cases decision-makers may be interested solely in the distribution of the incremental policy impacts, in other cases they may also want information on the extent to which this distribution exacerbates or ameliorates existing (baseline) inequalities.

These and other concerns are discussed in the cited references and elsewhere. Given the complexities of the issues, analysts lacking experience with the use of these measures may wish to consult with relevant experts to determine how to best proceed in a particular policy context. In some cases, it may be desirable to report and compare the results using multiple indices. In other cases, the types of descriptive information provided in Table 3.1 may suffice and no index may be needed.

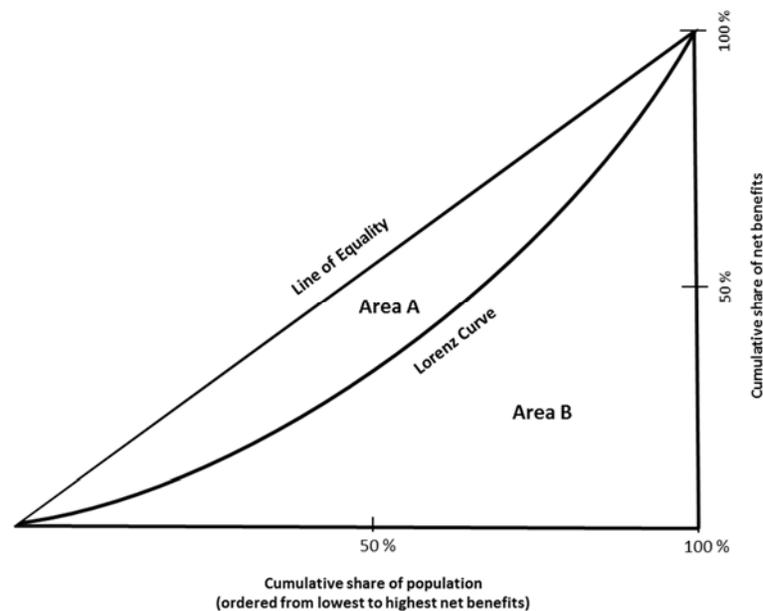
One approach, as illustrated in Figure 3.3, involves first using a Lorenz curve to show the degree of inequality that exists in the distribution of a variable, such as the fraction of net benefits that accrue to each fraction of the population, represented as a cumulative distribution. The Gini coefficient is a numerical measure of the degree of inequality between the variables, typically applied using Lorenz curve data. It is calculated by dividing Area A in Figure 3.3 by the sum of Area A and Area B (because the sum of area A and area B is one-half, the Gini coefficient is also equal to twice area A). It measures the

---

<sup>20</sup> See Harper et al. (2010) for discussion of these normative concerns in the health context.

departure from a uniform distribution and ranges from a value of zero to one, where zero represents perfect equality and one represents maximum inequality. For example, a value of zero would result if each one percent of the population received one percent of the net benefits; a value of one would result if one individual received all the net benefits.

**Figure 3.3. Lorenz Curve and Gini Coefficient**



While the Gini coefficient is perhaps the most commonly used measure of equality, it is not without limitations.<sup>21</sup> For example, a transfer between individuals will have differing effects on the Gini coefficient depending on their ranking within the distribution and the shape of the distribution. In addition, it cannot be easily decomposed by subgroup and only considers a single dimension of inequality (net benefits in the example).

There are many other such measures that can be used to summarize the degree of inequality.<sup>22</sup> Some of the most frequently applied include the concentration index (which is based on a concentration curve, a variant of the Lorenz curve illustrated above) and the Atkinson index. A concentration curve can be used to display the share of an outcome (such as net benefits) experienced by the cumulative share of the population ranked by a measure of socio-economic status such as income, and compare it to the line

---

<sup>21</sup> One important issue in evaluating these measures is whether they satisfy the Pigou-Dalton principle, which is attractive when evaluating social welfare. This principle states that a transfer of utility from the well-off to the less well-off is desirable (assuming the transfer is small enough that it does not reverse the order of wellbeing). See Adler (2015) for more discussion.

<sup>22</sup> For example, Harper et al. (2013) identify and evaluate 20 measures with a particular focus on health, building on work by Levy et al. (2006) and others.

of equality. A concentration index can then be used to indicate the equality of the outcomes; this index is defined as twice the area between the concentration curve and the line of equality. The index is zero if there is no inequality, with values ranging from -1.0 to +1.0. A value of -1.0 indicates the outcome is entirely borne by the individual with the lowest socio-economic status; a value of +1.0 indicates that the entire outcome accrues to the individual with the highest status; more generally negative values indicate that a greater share of the outcome accrues to the worse off (see O'Donnell et al. 2008 for more discussion).

A measure with more desirable normative properties is the Atkinson Index. For example, it disaggregated by subgroup. This index incorporates a measure of inequality aversion that can range from zero to infinity, where zero reflects indifference between the individuals who receive benefits and high values reflect a preference for transfers to those who are worse off. The analyst can calculate the index using several values to determine how sensitive it is to different weights (see Norheim 2013, Adler 2015 for more discussion).

The Kolm-Pollak index is similar to the Atkinson index but less frequently used. While the Atkinson index depends on proportional differences (e.g., one individual receives twice as much benefit as another), the Kolm-Pollak index depends on absolute differences (e.g., one individual receives ten units more benefit than another). The choice between these measures depends on whether proportional or absolute differences are more important. The Kolm-Pollak index can more easily be adapted to incorporate "bad" outcomes, such as pollution exposure (see Maguire and Sheriff 2011 for more discussion).

These indices provide additional information on the equality of the distribution of the effects throughout the population, supplementing the presentation of raw data as in Table 3.1. Because they are mathematically complex, these indices can be challenging to communicate to a nontechnical audience. Clear discussion of the calculations and the implications of the results is needed.

Providing better information about the distribution does not solve the decision-making problem, however. These metrics do not provide a guide to determining whether the distributional effects are severe enough to warrant selection of a policy that is less efficient but provides a more desirable distribution. Decision-makers still need to decide how to measure and weigh the desirability of the distributional effects.

Distribution and efficiency can be more fully integrated by conducting weighted benefit-cost analysis, in which the net benefits accruing to different groups are multiplied by distributional weights that reflect estimates of society's preferences for distribution (see, for example, Adler 2016a). Alternatively, policies can be evaluated using a social-welfare function to represent preferences for both the level and distribution of wellbeing. Such functions typically have two key components: one that describes individual wellbeing (e.g., a utility function that allows comparison of wellbeing between people), and a second that describes the weighting or ranking of individual wellbeing (e.g., societal preferences for equality of wellbeing).

Applying these more integrative approaches requires first agreeing on the normative framework (or frameworks) to be presented, then agreeing on how to implement the framework in terms of the mathematical formulation and parameter values. While uncertainty related to the appropriate framing could be represented by presenting the analytic results using multiple approaches, implementing these frameworks requires numerous complex decisions. Thus more guidance is needed to facilitate implementation if such functions are to be routinely applied. We provide an example of such an approach in the supplement to this paper.

## 4.0 Summary and Recommendations

While there is widespread agreement that benefit-cost analyses should be supplemented with information on how the impacts are distributed across individuals with different attributes (such as varying incomes), reviews of completed analyses suggest that such information is rarely provided. The reasons for this gap are unclear: perhaps analysts see this information as unimportant or unnecessary, lack the time and resources to conduct the assessment, need more methodological guidance, or are worried about what they might find.

Regardless, decision-makers and other stakeholders often express interest in this distribution. Estimating the distribution of net benefits allows them to weigh the extent to which benefits and costs are counterbalancing for members of each group as well as the overall distribution of net benefits across groups. It makes the trade-offs between economic efficiency and distributional concerns explicit. Decision-makers may choose an economically-efficient policy that maximizes net benefits, perhaps addressing distributional concerns through other policies, or may choose a less efficient option to ameliorate distributional impacts or achieve other policy goals.

### 3.1 Near-Term Recommendations

In the near-term, all benefit-cost analyses should include information on the distribution of net benefits. At minimum, analysts should report the effects of the policy on the health and longevity as well as income of individuals across different income groups. Analysts should also consider whether other impacts and other groups should be considered given the policy and decision-making context. The extent to which the analysis is quantitative or qualitative, as well as its level of detail, should be proportional to the importance of these concerns. The distributional analyses should be reported in a clearly labeled, separate section of the benefit-cost analysis, that discusses the available evidence and related uncertainties as well as the implications.

Analysts should proceed as follows.

- 1) Identify the individuals and impacts of concern.** In consultation with decision-makers and other stakeholders, analysts should identify the characteristics of individuals and impacts of concern. At minimum, the distributional analysis should address the effects of the policy on the health, longevity, and income of members of different income groups, including the distribution of net benefits. Analysts should consider whether other impacts and other grouping also should be addressed given the policy and decision-making context.
- 2) Determine the level of detail and degree of quantification for the distributional analysis.** The effort devoted to the distributional analysis should be proportionate to its importance for decision-making. “Importance” may depend on the likely magnitude of the distributional impacts and concerns about associated inequities; it may also depend on the need to respond to questions likely to be raised by decision-makers and others. The extent to which the analysis is quantitative or

qualitative, as well as its level of detail, should reflect the relative importance of these concerns and take into account the data, time, and resources available.

Screening analysis that relies on easily accessible data is often useful in providing preliminary information on the possible direction and magnitude of the effects, and in informing decisions about future work. Depending on the results, this screening may be followed by more detailed assessment that involves collecting additional data, refining the methods used, and expanding the scope of the analysis.

Both the benefit-cost analysis and the distributional assessment should be conducted concurrently with the policy development process, so that the preliminary results can be used to inform the options to be considered as well as future analytic steps.

- 3) Describe the distribution of benefits, costs, and net benefits across members of different population groups.** For each policy option assessed, analysts should report the results of the distributional analysis in text, tables, graphics, or other forms. These results should be reported as monetary values and in physical terms to the extent possible; e.g., as net benefits and as expected numbers of deaths, illnesses, and injuries averted. Measures of inequality, such as the Gini coefficient, may also be used; the advantages and limitations of the selected measure(s) should be discussed along with the results. The ultimate results of the distributional analysis should be reported in a clearly labeled, separate section of the documentation, that describes the data sources and methods used as well as the results.

All analyses should address the distribution of net benefits (benefits minus costs) to members of each population subgroup (e.g., each income quintile), to ensure that potentially counterbalancing impacts are not ignored. In addition, analysts should explore the extent to which there is heterogeneity in the net benefits that accrue across members of each group; e.g., because some members are more vulnerable to the health condition than others.

- 4) Address uncertainty:** Analysts should assess uncertainties in the estimates both qualitatively and quantitatively; e.g., by conducting sensitivity or probabilistic analysis, discussing the quality of the evidence, and indicating the implications for decision-making.

While the screening analysis under step 2 may lead to the conclusion that qualitative discussion will suffice, in most cases quantitative data on distribution will be more informative and desirable. Where data limitations or other issues limit the ability to quantify the effects, analysts should use “what if” or bounding analysis to explore the likely distribution of net benefits.

### 3.2 Long-Term Recommendations

Over the long term, more work is needed to provide examples of how to assess the distribution of the impacts of different types of policies, recommendations on the application of specific inequality metrics, and options for distributional weighting using social welfare functions.

- 1) **Develop case studies.** Assessment of distribution may be complicated and difficult in many contexts. Case studies that illustrate best practices for such assessments are likely to substantially aid in encouraging the more routine completion of these analyses. These case studies could, for example:
  - a. Explore the data available in different contexts and the options for qualitative and quantitative assessment.
  - b. Examine how to best communicate the analytic approach, the results, and the implications to diverse decision-makers and other stakeholders.
  - c. Investigate the use and usefulness of distributional information in decision-making, to identify needed analytic improvements.
  
- 1) **Explore inequality metrics.** As noted earlier, inequality metrics such as the Gini coefficient can provide useful summary measures of the degree of equality of the policy outcome. More criteria-driven review of the potential applicability of these indices in global benefit-cost analysis, and of their normative foundations, is needed to develop recommendations for the indices that are most appropriate across different policy and decision-making contexts.
  
- 2) **Investigate the use of social welfare functions.** Distributional analysis using social welfare functions is a useful supplement or complement to traditional benefit-cost analysis accompanied by descriptive information on distribution. More work is needed to further develop parameter values and to provide practical advice on the implementation of this approach in global benefit-cost analysis.

In conclusion, this paper provides recommendations for describing the distribution of net benefits that can be feasibly implemented based on the research now available. It also identifies areas in need of more research both to demonstrate the implementation of these recommendations and to further explore measures of inequality and application of social welfare functions. The results of such analysis allow decision-makers and others to weigh distributional concerns along with other policy impacts and determine how to best address these concerns within a particular context.

## Supplement: Social Welfare Functions

By Matthew D. Adler

The social welfare function (SWF) is a framework for policy assessment that is very well developed in the academic literature (Adler 2012, Weymark 2016.) The underlying theory has been fully elaborated by welfare economists, and the framework is widely used by academic economists in various specific policy-focused literatures, such as the “optimal tax” literature and scholarship concerning climate policy.

Stated most abstractly, the SWF framework has two components: (1) a wellbeing measure, and (2) a rule for ranking lists of wellbeing numbers. An “outcome” is a possible consequence of policy choice. The wellbeing measure converts each outcome into a list of wellbeing numbers, one for each person in the population of interest. An individual’s wellbeing number in a given outcome depends upon her bundle of attributes in that outcome (income, health, longevity, leisure, education, etc.).

These wellbeing numbers are interpersonally comparable. If Jim in one outcome is assigned a wellbeing number of 100 (for example), and Sara a wellbeing number of 230, then this indicates that Jim is worse off than Sara. Further, most SWF scholars assume (in line with the standard view in economics) that an individual’s wellbeing depends upon her preferences. Developing a wellbeing measure that is both interpersonally comparable and based upon individual preferences poses a challenge for the SWF framework—but the literature has demonstrated how this challenge can be overcome. If individuals in the population of interest have common preferences, then the wellbeing measure is just a so-called utility function representing the common preferences.<sup>23</sup> With heterogeneous preferences, the number of a person’s attribute bundle in a given outcome is its utility according to a utility function representing that particular person’s preferences (Adler 2016b).<sup>24</sup>

Consider now the second component of the SWF framework: the rule for ranking lists of wellbeing numbers. There are many different possibilities, but two are especially plausible and have been dominant in the literature. The utilitarian SWF simply adds up wellbeing numbers. One outcome is ranked better than a second if it has a greater sum total of wellbeing. A prioritarian SWF sums up “transformed” wellbeing numbers. Each number is plugged into a “transformation” function that has the concave shape as indicated in Figure 1 below.<sup>25</sup> One outcome is ranked better than a second if it has a greater sum total of transformed well-being.

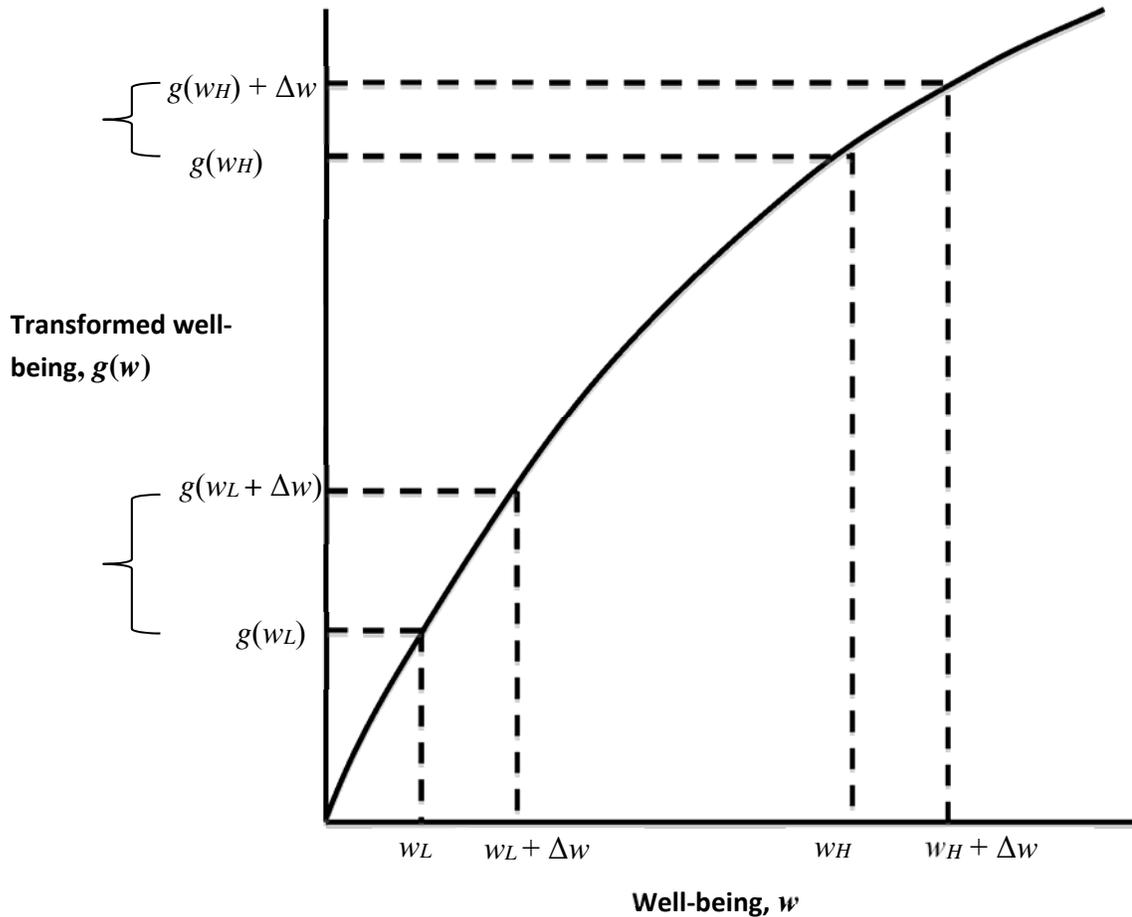
---

<sup>23</sup> Formally, a preference is understood as a ranking of attribute bundles and lotteries over bundles. Individuals have common preferences if they have the same such ranking.

<sup>24</sup> A different approach relies upon so-called “equivalent incomes.” See Fleurbaey (2016).

<sup>25</sup> More precisely: the prioritarian transformation function is strictly increasing and strictly concave.

Figure S-1



*Explanation:* Figure 1 displays a concave transformation function. Consider two individuals, one at lower well-being level  $w_L$ , the second at higher well-being level  $w_H$ . As illustrated by the dashed lines, increasing the worse-off individual's well-being by  $\Delta w$  produces a bigger increase in her transformed well-being than increasing the better-off individual's well-being by the same amount. This is the "priority" for the worse off.

Because the prioritarian transformation function is concave, a prioritarian SWF has the effect of giving greater weight ("priority") to wellbeing changes affecting worse-off individuals (those at lower wellbeing levels). This is illustrated in Figure S-1. Prioritarianism is an entire *family* of SWFs, each with its own transformation function. The more concave this function, the greater the priority for the worse off. At

the limit, the prioritarian SWF approaches the “leximin” rule, which gives absolutely priority to the worse off.

A properly constructed wellbeing measure will reflect the “diminishing marginal utility” of income. Each incremental unit of income produces a smaller increment in wellbeing. If Alex, Betty and Charles earn annual incomes of, respectively \$20,000, \$50,000 and \$100,000, then adding \$100 to Charles’ income has a smaller effect on wellbeing than adding \$100 to Betty’s, which in turn has a smaller wellbeing impact than adding \$100 to Alex’s.

Because both the utilitarian SWF and prioritarian SWFs incorporate a wellbeing measure reflecting the diminishing marginal utility of income, both are automatically sensitive to the incidence of a policy’s costs and benefits. In this regard, both differ from traditional benefit-cost analysis.<sup>26</sup> Assume that groups *A* and *B* are the same size. Policy *A* produces a benefit to group *A* members equivalent to  $\$D$  in additional income for each, i.e., each group *A* member is willing to pay  $\$D$  for the policy. Policy *B* produces a benefit to group *B* members also equivalent to  $\$D$  in additional income for each, i.e., each group *B* member is willing to pay  $\$D$  for the policy. Then traditional benefit-cost analysis will be neutral between the policies.

By contrast, the utilitarian and prioritarian SWFs will typically not be neutral. In particular, imagine that group *B* members have lower incomes than group *A* members and are otherwise identical. Then both types of SWF will prefer policy *B*.

While utilitarian and prioritarian SWFs are *both* sensitive to distribution, they differ in how they assess patterns of incidence. If Alex earns \$20,000 and Betty earns \$50,000, the utilitarian SWF will prefer adding \$100 to Alex’s income over adding \$100 to Betty’s. It will be indifferent between adding some amount  $\$Y$  to Alex’s income and \$100 to Betty’s, with  $\$Y$  less than \$100. A prioritarian SWF will be indifferent between adding some amount  $\$Y^*$  to Alex’s income and \$100 to Betty’s, where  $\$Y^*$  in turn is less than  $\$Y$ . As the prioritarian transformation function become increasingly concave,  $\$Y^*$  becomes smaller and smaller.

Critics of the SWF framework sometimes object that it involves contestable “value judgments.” This is certainly true. Value judgments do come into play at various junctures—in deciding how to construct a well-being measure; in choosing between the utilitarian SWF, prioritarian SWFs, and other possibilities; and in selecting a prioritarian transformation function (if that approach is followed). To the extent that the SWF framework is actually used to select governmental policies, these value judgments ought to be made by democratically accountable policymakers.

---

<sup>26</sup> The term “traditional benefit-cost analysis” is used in this supplement to mean the methodology that quantifies a policy’s impact on each individual by using a monetary equivalent—calculated as the individual’s willingness to pay/willingness to accept for the policy—and that scores policies by summing these monetary equivalents without distributional weights.

Traditional benefit-cost analysis, too, involves value judgments. The choice of an assessment framework that ignores the declining marginal utility of income is a (highly significant) value choice. So, too, are various choices made in fleshing out the details of benefit-cost analysis—for example, the use of population-average valuations. It is a mischaracterization to suggest that benefit-cost analysis is somehow “value neutral.”

A different criticism is that the SWF framework ignores potential Pareto improvements. Assume that traditional benefit-cost analysis prefers one policy to a second, while the SWF framework (be it utilitarian or prioritarian) prefers the second policy. Then the first policy is potentially Pareto superior (“Kaldor-Hicks efficient”) relative to the second. The first policy, together with an appropriate scheme of compensations from winners (those better off with the first policy) to losers (those worse off with the first) would make everyone better off than if the second policy were chosen.

Here, it is important to stress that utilitarian and prioritarian SWFs do rank outcomes consistently with the Pareto principle. If some individuals have higher wellbeing levels in outcome  $x$  than  $y$ , and no one has a lower wellbeing level, these SWFs will prefer  $x$  to  $y$ . With this in mind, consider more closely the case in the previous paragraph. If the first policy would *actually* be combined with a change to the tax system compensating the losers, then the SWF approach—now taking into account both the policy *and* the actual tax change predicted to accompany it—would favor the first policy. (The SWF approach prefers the second policy to the first policy without the tax change, but it prefers the first policy in combination with the tax change to the second policy.) If the scheme of compensations from winners to losers is *merely* hypothetical—if the policy analyst predicts that no such change to the tax system would actually be enacted to accompany the first policy—then the SWF framework will not pick the first policy. Why is it relevant that the first policy is *potentially* Pareto superior to the second, if that potential will not in fact be realized?

Let’s turn, specifically, to consider how the SWF framework can be used to evaluate health policy (Adler 2017). Various refinements to the general methodology come into play. (1) Each individual’s attribute bundle is a *lifetime* bundle. An individual’s wellbeing is her *lifetime* wellbeing, as determined by this whole-lifetime collection of attributes. Mortality-reduction policies, in particular, have the benefit of augmenting longevity and thereby increasing lifetime wellbeing. (2) Ideally, the attributes in a given bundle will include longevity, health across the lifespan, and income across the lifespan, since these are the attributes most directly affected by health policy. Let’s use the term “history” to mean a possible lifetime combination of longevity, health, and income. (3) Assuming common preferences, there is a common lifetime utility function. The wellbeing number of a history is its lifetime utility according to this common function.<sup>27</sup> (4) Uncertainty needs to be explicitly addressed. A typical policy will reduce various

---

<sup>27</sup> Let  $u(\cdot)$  denote the common lifetime utility function, and  $h$  a history. Often, economists assume that lifetime utility functions are additive. If  $h_t$  are the attributes in history  $h$  during period  $t$  of the individual’s life,  $u(h) = \sum_{t=1}^T v(h_t)$ , with  $T$  maximum length of life and  $v(\cdot)$  a sublifetime utility function. This approach generalizes to allow for heterogeneous preferences and thus heterogeneous lifetime utility functions, each temporally additive.

individuals' *risk* of dying, rather than saving identified persons. (5) Pulling this all together: each person in the status quo baseline faces a probability distribution over histories. A given policy produces its own probability distribution over histories. For example, if the policy reduces mortality risk and morbidity, at some cost, the individual will have a greater probability of histories with longer longevity or better health, and a greater probability of histories with lower income (the cost of the longevity and health gains). (6) Various modelling choices can be used to simplify the exercise. For example, the analyst might ignore income mobility and assume that each individual has the same income for each year alive. (7) One type of simplification is to assess probabilities on a group basis. The population can be divided into groups defined by current income, age, demographic characteristics, etc. It can then be assumed that each group member faces the same baseline probability distribution over histories, and the same policy distribution over histories associated with each policy under consideration.

The SWF (be it utilitarian or prioritarian) will then assign the baseline and each policy a score, depending upon the size of the affected groups, and what each group's distribution over histories is. In the case of the utilitarian SWF, this occurs as follows: each group member has an expected lifetime well-being with the baseline, and a (different) expected lifetime well-being with a given policy. The score assigned to a policy, or to the baseline, is just the sum total of expected well-being across the groups. The best policy is the one with the highest score.

This can be expressed formally as follows.

$$s^{Util}(P) = \sum_{g=1}^G N_g \sum_h \pi_g^P(h) w(h)$$

$s^{Util}(P)$  is the utilitarian score assigned to policy  $P$ . There are  $G$  groups in total in the population.  $N_g$  is the number of individuals in group  $g$ .  $h$  is a history, and  $w(h)$  is the lifetime well-being number of a history.  $\pi_g^P(h)$  is the probability of history  $h$  for group  $g$  given policy  $P$ . The formula for the prioritarian SWF is the same, except that it includes the transformation function.<sup>28</sup>

---

<sup>28</sup> This can be done in an "ex post" manner, yielding "ex post prioritarianism" (EPP), or an "ex ante" manner, yielding "ex ante prioritarianism" (EAP). With  $f(\cdot)$  the transformation function (I use the symbol  $f(\cdot)$  rather than  $g(\cdot)$  to avoid confusion with the "g" designating groups), the formulas are as follows.

$s^{EPP}(P) = \sum_{g=1}^G N_g \sum_h \pi_g^P(h) f(w(h))$ .  $s^{EAP}(P) = \sum_{g=1}^G N_g f\left(\sum_h \pi_g^P(h) w(h)\right)$ . On prioritarianism under uncertainty, see generally Adler (2012, ch. 7).

## References

(Includes links for publications that are freely available online.)

Adler, M.D. 2012. *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis*. New York: Oxford University Press.

Adler, M.D. 2015. Equity by the Numbers: Measuring Poverty, Inequality, and Injustice. *Alabama Law Review* 66(3): 551-607.

Adler, M.D. 2016a. Benefit–Cost Analysis and Distributional Weights: An Overview. *Review of Environmental Economics and Policy* 10(2): 264–285.

Adler, M.D. 2016b. Extended Preferences. *The Oxford Handbook of Well-Being and Public Policy* (M.D. Adler and M. Fleurbaey, eds.). New York: Oxford University Press.

Adler, M.D. 2017. A Better Calculus for Regulators: From Cost-Benefit Analysis to the Social Welfare Function. Working paper. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2923829](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2923829)

Adler, M.D. and M. Fleurbaey (eds). 2016. *The Oxford Handbook of Well-Being and Public Policy*. New York: Oxford University Press.

Boardman, A.E., D.H. Greenberg, A.R. Vining, D.L. Weimer. 2011. *Cost-Benefit Analysis: Concepts and Practice*, 4th edition. Upper Saddle River, N.J.: Pearson.

Brouwer, W.B.F. et al. 2008. Welfarism vs. Extra-Welfarism. *Journal of Health Economics* 27: 325-338.

Deaton, A. 2013. What Does the Empirical Evidence Tell Us About the Injustice of Health Inequalities? *Inequalities in Health: Concepts, Measures, and Ethics* (N. Eyal et al., eds.). New York: Oxford University Press.

Eyal, N. et al. (eds.). 2013. *Inequalities in Health: Concepts, Measures, and Ethics*. New York: Oxford University Press.

Fann, N., H.A. Roman, C.M. Fulcher, M.A. Gentile, B.J. Hubbell, K. Wesson, and J.I. Levy. 2011. Maximizing Health Benefits and Minimizing Inequality: Incorporating Local-Scale Data in the Design and Evaluation of Air Quality Policies. *Risk Analysis* 31(6): 908-922.

Fleurbaey, M. 2016. Equivalent Income. *The Oxford Handbook of Well-Being and Public Policy*, (M.D. Adler and M. Fleurbaey, eds.). New York: Oxford University Press.

Fleurbaey, M. and E. Schokkaert. 2009. Unfair Inequalities in Health and Health Care. *Journal of Health Economics* 28(1): 73–90.

Greenberg, D. and A. Cebulla. 2008. The Cost-Effectiveness of Welfare-to-Work Programs: A Meta-Analysis. *Public Budgeting and Finance*. Summer: 112-145.

Greenberg, D., V. Deitch, and G. Hamilton. 2009. Welfare-to-Work Program Benefits and Costs: A Synthesis of Research. Prepared for MDRC. <https://www.mdrc.org/publication/welfare-work-program-benefits-and-costs>

Hammitt, J.K. 2013. Positive vs. Normative Justifications for Benefit-Cost Analysis: Implications for Interpretation and Policy. *Review of Environmental Economics and Policy* 7(2): 199-218.

Harper S., N.B. King, S.C. Meersman, M.E. Reichman, N. Breen and J. Lynch. 2010. Implicit Value Judgements in the Measurement of Health Inequalities. *Milbank Quarterly*. 88(1): 4-29.

Harper, S., E. Ruder, H.A. Roman, A. Geggel, O. Nweke, D. Devon Payne-Sturges, and J.I. Levy. 2013. Using Inequality Measures to Incorporate Environmental Justice into Regulatory Analyses. *International Journal of Environmental Research and Public Health* 10: 4039-4059. <http://www.mdpi.com/1660-4601/10/9/4039>

H.M. Treasury. 2011. The Green Book: Appraisal and Evaluation in Central Government. <https://www.gov.uk/government/publications/the-green-book-appraisal-and-evaluation-in-central-government>

Kaplow, L. and S. Shavell. 2006. *Fairness versus Welfare*. Cambridge: Harvard University Press.

Levy, J.I., S.M. Chemerynski, and J.L. Tuchmann. 2006. Incorporating Concepts of Inequality and Inequity into Health Benefits Analysis. *International Journal for Equity in Health* 5:2. <https://equityhealthj.biomedcentral.com/articles/10.1186/1475-9276-5-2>

Long, D.A., C.D. Mallar and C.V.D. Thornton. 1981. Evaluating the Benefits and Costs of the Job Corps. *Journal of Policy Analysis and Management* 1(1): 55-76.

Maguire, K. and G. Sheriff. 2011. Comparing Distributions of Environmental Outcomes for Regulatory Environmental Justice Analysis. *International Journal of Environmental Research and Public Health* 8(5): 1707-1726.

Norheim, O.F. 2013. Atkinson's Index Applied to Health: Can Measures of Economic Inequality Help Us Understand Tradeoffs in Health Care Priority Spending? *Inequalities in Health: Concepts, Measures, and Ethics*. (N. Eyal et al., eds.). New York: Oxford University Press.

- Nyborg, N. 2012. *The Ethics and Politics of Environmental Cost-Benefit Analysis*. New York: Routledge.
- O'Donnell, O. et al. 2008. *Analyzing Health Equity Using Household Survey Data: A Guide to Techniques and Their Implementation*. Washington, DC: The World Bank.  
<https://openknowledge.worldbank.org/handle/10986/6896>
- Robinson, L.A. and J.K. Hammitt. 2017. Benefit-Cost Analysis in Global Health. Forthcoming in: *Global Health Priority-Setting: Beyond Cost-Effectiveness* (O.Norheim et al., eds.). Oxford, UK and New York, NY: Oxford University Press. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2952014](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2952014)
- Robinson, L.A., J.K. Hammitt, and R. Zeckhauser. 2016. Attention to Distribution in U.S. Regulatory Analysis. *Review of Environmental Economics and Policy* 10(2): 308–328.
- Robinson, L.A. et al. 2017. *Benefit-Cost Analysis in Global Health and Development: Current Practices and Opportunities for Improvement (Scoping Report)*.  
<https://sites.sph.harvard.edu/bcaguidelines/scoping/>
- Robinson, L.A., J.K. Hammitt, and L. O’Keeffe. 2018. Valuing Mortality Risk Reductions in Global Benefit-Cost Analysis. <https://sites.sph.harvard.edu/bcaguidelines/methods-and-cases/>
- Robinson, L.A. and J.K. Hammitt. 2018. Valuing Nonfatal Health Risk Reductions in Global Benefit-Cost Analysis. <https://sites.sph.harvard.edu/bcaguidelines/methods-and-cases/>
- Strzepek, K.M., C. Amany, and J. E. Neumann. 2018. Assessing Economy-wide Effects of Health and Environmental Interventions in Support of Benefit-Cost Analysis.  
<https://sites.sph.harvard.edu/bcaguidelines/methods-and-cases/>
- U.S. Department of Health and Human Services. 2016. *Guidelines for Regulatory Impact Analysis*.  
<https://aspe.hhs.gov/pdf-report/guidelines-regulatory-impact-analysis>
- U.S. Office of Management and Budget. 2003. *Circular A-4: Regulatory Analysis*.  
[http://www.whitehouse.gov/omb/inforeg\\_regpol\\_agency\\_review/](http://www.whitehouse.gov/omb/inforeg_regpol_agency_review/)
- Weymark, J. 2016. Social Welfare Functions. *The Oxford Handbook of Well-Being and Public Policy* (M.D. Adler and M. Fleurbaey, eds.). New York: Oxford University Press.
- Wilkinson, T. et al. 2016. The International Decision Support Initiative Reference Case for Economic Evaluation: An Aid to Thought. *Value in Health* 19(8): 921–928.