

ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

EFFICIENCY VS. WELFARE IN BENEFIT-COST ANALYSIS:
THE CASE OF GOVERNMENT FUNDING

Zachary D. Liscow
Cass R. Sunstein

Discussion Paper No. 1108

11/2023

Harvard Law School
Cambridge, MA 02138

This paper can be downloaded without charge from:

The Harvard John M. Olin Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center

The Social Science Research Network Electronic Paper Collection:
<https://ssrn.com/abstract=4589563>

Efficiency vs. Welfare in Benefit-Cost Analysis: The Case of Government Funding

Zachary Liscow* and Cass R. Sunstein**

Abstract

In Republican and Democratic administrations, regulatory and funding decisions have both been made with close reference to benefit-cost analysis (BCA). With respect to regulation, there has been a great deal of scholarly discussion of BCA and its limits, but almost no attention has been paid to the role of BCA in government funding. That is a serious gap, not least in connection with climate-related risks, such as wildfire, drought, extreme heat, and flooding. In OMB Circular A-94, the Office of Management and Budget has long required applicants for federal funding to demonstrate that the benefits of their projects would exceed the costs. Under Circular A-94, efficiency-based BCA can produce results that fail to maximize welfare and that are also highly inequitable. The 2023 draft revision of Circular A-94, focused on welfare and equity, reflects an effort to incorporate new academic thinking over the past three decades, which is now—not uncontroversially—being brought directly into policy. At the same time, the new draft Circular A-94 raises fresh questions about how best to promote welfare, and to consider equity, in practice. Pressing issues involve the use of distributional weights in funding decisions and also the use of averages across populations, which might be seen as a form of distributional weighting. More broadly, the trajectory of this benefit-cost guidance, which predates the guidance for regulation and originally covered regulation, helps uncover the logic under which BCA has been operating and deeper challenges and tensions within BCA, in the past and going forward.

* Professor of Law, Yale Law School. Disclosure: Liscow led the revision to Circular A-94 during his time as Chief Economist at the Office of Management and Budget in 2022-2023.

** Robert Walmsley University Professor, Harvard University. Disclosure: Since early 2021, Sunstein has served as Senior Counselor to the Secretary of Homeland Security (now as Special Government Employee), and in that capacity, he was involved in various discussions of Circular A-94 and its revision. Both Liscow and Sunstein wrote extensively on benefit-cost analysis before government service, and this essay, written solely in their academic capacities, uses only publicly available information. Nothing here represents an official position in any way. Thanks to Raj Bhargava, Krister Rasmussen, Marisa Sylvester, and Victoria Yu for excellent research assistance. Thanks too to Nathan Hendren and participants at the Tax Policy and the Economy conference for helpful comments.

Table of Contents

I. Introduction: Funding Projects with Winners and Losers	2
II. Efficiency, Kaldor-Hicks, and Welfare	7
A. Efficiency in the Current Circular A-94	7
B. Welfare in the Draft Revised Circular A-94.....	9
C. Four Distinct Justifications	12
III. From Theory to Practice: Pursuing Welfare with Benefit-Cost Analysis	17
A. Direct Incorporation of Equity into BCA	17
B. Optional Use of Distributional Weights	18
C. What Distributional Weights to Use	19
D. Measuring Income of Affected Communities	19
E. Incidence of Spending	20
F. Uncertainty and Equity	21
G. Net Benefits vs. Benefit-Cost Ratio	21
H. Administrative Burden.....	23

I. Introduction: Funding Projects with Winners and Losers

By what criteria should the national government fund projects?

For orientation, consider four stylized cases:

1. *Wealthy community, flood risks.* A wealthy community applies to the federal government for funding for a levee designed to reduce the risks associated with flooding. The community seeks \$40 million. Because of high property values in the community, the expected benefits of the project exceed the costs.
2. *Poor community, wildfire risks.* A poor community applies to the federal government for funding for a project designed to reduce the risks associated with wildfire. The community seeks \$40 million. Because of low property values in the community, the expected benefits of the project are significantly lower than the cost.
3. *Poor community, extreme heat.* A poor community applies to the federal government for funding for a project designed to reduce the mortality and morbidity risks associated with extreme heat. The community seeks \$40 million. The expected benefits of the program are somewhat lower than the cost.

4. *Poor community, uncertainty.* A poor community applies for to the federal government for funding for a project designed to reduce the risks associated with flooding. The community seeks \$40 million. The community does not know whether the benefits of the project would exceed the costs. It lacks information about benefits.

Stylized though they are, these problems are broadly reflective of actual ones faced by government agencies, perhaps above all the Federal Emergency Management Agency,¹ but also the Department of Transportation, the Department of Agriculture, the Army Corps of Engineers, and other agencies across government.² To allocate its limited resources, the government might well want to know about both costs and benefits before it agrees to fund projects. But what it should *do* with that knowledge is not entirely obvious. In each of the four cases given above, is it so clear how government should proceed, if it had an assortment of specific numbers?

Since 2003, federal *regulation* has been assessed under the framework established by OMB Circular A-4, *Regulatory Analysis*,³ which has been subject to extended academic discussion.⁴ Since at least 1972,⁵ *funding decisions* have been made under the framework established by OMB Circular A-94, now called *Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs*,⁶ which has received exceedingly little academic discussion.⁷ Originally, Circular A-94 also covered regulation as well, before the issuance in 1996 of new guidance on regulations.⁸ Circular A-94 is much shorter than Circular A-4, but it has comparable importance, covering \$40-50 billion of spending annually.⁹ It establishes the criteria that federal

¹ See, e.g., *Building Resilient Infrastructure and Communities*, FEMA, <https://www.fema.gov/grants/mitigation/building-resilient-infrastructure-communities> (last updated Sept. 23, 2023).

² There are international analogues, as in decisions faced by the United States Agency for International Development, the World Bank, and the Asian Development Bank. We do not explore those decisions, but the discussion obviously has implications for them, not least in the context of efforts to fund projects designed to increase resilience against climate-related risks.

³ See OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, CIRCULAR A-4, REGULATORY ANALYSIS (2003).

⁴ See generally Daniel Hemel, *Regulation and Redistribution With Lives in the Balance*, 89 U. CHI. L. REV. 649 (2022); Thomas J. Kniesner & W. Kip Viscusi, *Promoting Equity Through Equitable Risk Tradeoffs* (IZA Inst. of Labor Econ. Discussion Paper, No. 15771), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4294396 (last updated Apr. 14, 2023).

⁵ See OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, CIRCULAR A-94, DISCOUNT RATES TO BE USED IN EVALUATING TIME-DISTRIBUTED COSTS AND BENEFITS (1972). The original version of Circular A-94 dates to 1969.

⁶ See OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, CIRCULAR A-94, GUIDELINES AND DISCOUNT RATES FOR BENEFIT-COST ANALYSIS OF FEDERAL PROGRAMS (1993).

⁷ A Google Scholar search of “Circular A-94” reveals just a few academic citations before the recent revisions, and those tended to be on discount rates.

⁸ OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, ECONOMIC ANALYSIS OF FEDERAL REGULATIONS UNDER EXECUTIVE ORDER 12866 (1996). And, on its face, the version current through 2023 still appears to cover regulation, though it does not in practice.

⁹ David Mitchell, *Proposed Update to Federal Cost-Benefit Analysis Guidelines Correctly Focuses on Accounting for Inequality in Regulations*, WASHINGTON CTR. FOR EQUITABLE GROWTH (May 23, 2023),

agencies use when deciding whether and when to give out grants. The Office of Management and Budget has proposed significant revisions of both circulars—not without controversy,¹⁰ including all former presidents of the Society of Benefit-Cost Analysis opposing some of the proposed changes, including those on equity.¹¹ Circular A-94 is our focus here, and we shall explore what we call the “draft revised” circular. Some of what we say will bear on the proposed revision of Circular A-4 as well.

In its current version from 1993, Circular A-94 announces that its goal “is to promote efficient resource allocation.”¹² It emphasizes the importance of formal benefit-cost analysis, which it describes as the recommended “technique to use in a formal economic analysis of government programs or projects.”¹³ It calls for “comprehensive estimates of the expected benefits and costs,” with close reference to “[t]he principle of *willingness to pay*,”¹⁴ noting that “[m]arket prices provide an invaluable starting point.”¹⁵ Importantly, Circular A-94 also recognizes that there might be “significant distributional effects,” and asks agencies to analyze those effects by reference to income class, geographical region, or demographic group.¹⁶ We shall return to this direction in due course.

Under the current Circular A-94, some of the cases given above are relatively straightforward. In Case 1, the wealthy community can receive federal funds. In Cases 2 and 3, the poor community almost certainly cannot. In Case 4, the question is why, exactly, the community does not know the costs and benefits. If the reason is that it has not made reasonable efforts to find out, it almost certainly cannot receive federal funds. If the reason is that numbers cannot be found, the question is whether some kind of additional analysis would be helpful.

<https://equitablegrowth.org/proposed-update-to-federal-cost-benefit-analysis-guidelines-correctly-focuses-on-accounting-for-inequality-in-regulations/>.

¹⁰ See, e.g., Susan Dudley & W. Kip Viscusi, *Biden’s OMB Politicizes Cost-Benefit Analysis*, WALL ST. J. (AUG. 28, 2023, 5:44 PM), <https://www.wsj.com/articles/bidens-omb-politicizes-cost-benefit-analysis-regulation-social-justice-2534e819>; Susan E. Dudley et al., *Letter to OIRA Administrator on Circular A4*, GEO. WASH. U. REGUL. STUD. CTR. (AUG. 28, 2023), <https://regulatorystudies.columbian.gwu.edu/letter-oira-administrator-circular-a4>.

¹¹ See, e.g., Public Comment on Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs, 88 Fed. Reg. 20913 (2023) (OMB requesting comments in April 2023 for proposed changes to Circular A-94).

¹² See CIRCULAR A-94, *supra* note 6, at 2.

¹³ *Id.* at 4. Circular A-94 also allows consideration of cost-effectiveness analysis, but pointedly describes it as “a less comprehensive technique.” *Id.*

¹⁴ *Id.* at 6-7. It is worth noting that the current A-94 refers to willingness to pay, not willingness to accept. Willingness to pay is dependent on ability to pay, which leads to differences between the two measures (as does the endowment effect). See Daniel Kahneman, Jack L. Knetsch & Richard H. Thaler, *Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias*, 5 J. ECON. PERSPECTIVES. 193 (1991). Poor people might be willing to pay very little for goods that would greatly improve their lives, and wealthy people might be willing to pay a great deal for goods that would only modestly improve their lives. The draft revised Circular A-94 is alert to these points and does not opt for willingness to pay as opposed to willingness to accept.

¹⁵ *Id.* at 7.

¹⁶ *Id.* at 12.

But many questions might be asked about these results. Let us focus in particular on the contrast between Case 1 and Case 2. It is agreed that the monetized net benefits of the project in Case 1 are far higher than the monetized net benefits of the project in Case 2. But should that be the end of the matter? Consider two points. First, the benefits to the community in Case 2 might be much higher than the benefits to the community in Case 1, purely in terms of welfare. If you are rich and lose \$200,000, you might not lose much; if you are poor and lose \$200,000, you lose everything. If welfare is what matters, then the current Circular A-94 might well point in the wrong direction. Second, the community in Case 2 is poor, and purely on distributive grounds, that community might have a stronger claim to public resources than the community in Case 1. Even if we do not see an improvement in welfare from a grant in Case 2, we might go forward on the ground that those who receive the grant have a special claim on public resources. “Prioritarians,” who give priority to those at the bottom of the welfare ladder, would be strongly inclined in this direction.¹⁷

For three reasons, it is especially important to probe Circular A-94, both in its current and draft revised versions, and the benefit-cost analysis of spending. First, spending programs have of course always mattered, and they may be more important now than ever. For example, the United States faces multiple climate-related risks, including wildfire, extreme heat, drought, and flooding. Many federal programs provide significant resources to increase resilience against those risks.¹⁸ In addition, there have been large expansions in recent years in this grant-based funding requiring BCA.¹⁹ But who should receive the relevant funding? Poor communities are often especially vulnerable to climate-related risks. Should they receive no funding, or less funding, because of low property values? That might seem preposterous. The whole area is a propitious place to focus on considering equity in BCA, mostly because the stakes are so high but also because the legal regime makes legal challenges less likely for spending than for regulation.²⁰ The distributive challenges are also particularly stark and compelling, with a recent *Politico* headline reading: “How FEMA helps white and rich Americans escape floods.”²¹ Importantly, and as we will see, the tendency to direct funds toward richer places is actually

¹⁷ See generally Matthew D. Adler, *Theory of Prioritarianism*, in *PRIORITARIANISM IN PRACTICE* 37 (Matthew D. Adler & Ole F. Norheim eds., 2022) (outlining the theory of prioritarianism as a branch of welfare consequentialism); Matthew D. Adler & Nils Holtug, *Prioritarianism: A Response to Critics*, 18 *POL., PHIL. & ECON.* 101 (2019) (defending prioritarianism against objections).

¹⁸ See, e.g., *Building Resilient Infrastructure and Communities*, *supra* note 1.

¹⁹ See, e.g., Infrastructure Investment and Jobs Act of 2021 (Bipartisan Infrastructure Law), Pub. L. No. 117-58, 135 Stat. 429.

²⁰ On legal challenges to regulation, see Caroline Cecot & W. Kip Viscusi, *Judicial Review of Agency Benefit-Cost Analysis*, 22 *GEO. MASON L. REV.* 575 (2015); Richard L. Revesz & Samantha P. Yi, *Distributional Consequences and Regulatory Analysis*, 52 *Env't. L.* 53 (2022). In the funding context, issues of standing and reviewability may limit plaintiffs' ability to bring such challenges. See, e.g., *Lujan v. Defenders of Wildlife*, 504 U.S. 555 (1992); *Allen v. Wright*, 468 U.S. 737 (1984).

²¹ Thomas Frank, *How FEMA Helps White and Rich People Escape Floods*, *POLITICO* (May 27, 2022, 4:30 AM), <https://www.politico.com/news/2022/05/27/unfair-fema-climate-program-floods-00032080>.

baked into FEMA’s longstanding methods (thus prompting recent efforts at reform). To say the least, it is important to get funding decisions right.²²

Second, funding is relevantly different from regulation, both analytically and in practice. Suppose, for example, that government requires refrigerators to be more energy-efficient. If so, consumers are likely to have to pay more for refrigerators, at least upfront, and it is possible that the regulation could end up hurting poor people.²³ If, by contrast, the government subsidizes energy-efficiency refrigerators, no such effect should be expected. This absence of price effects makes the distributional impacts of spending more straightforward than for regulation. Furthermore, the fixed budgets of agencies make the tradeoffs different: it is a matter of directly shifting financial resources between parties, rather than regulating or not. And unlike for regulation, parties—often poorly-resourced local governments, not well-versed in the requirements of cost-benefit analysis—apply directly, creating distinctive challenges with a degree of analytical complexity.

Third, spending provides a useful lens on benefit-cost analysis more generally, including both its past and its future. Circular A-94 has been admirably clear about what it is doing and why, which matters for both spending and regulation, since regulation grew out of it. It is maximizing efficiency and doing so because it assumes that the winners would (or at least could) compensate the losers. Unearthing this foundation helps us understand the logic of past BCA and where it could go in the future. Showing what BCA was really trying to do, and on what grounds, reveals tensions, challenges, and questions in past and future practice, providing insights for the past, present, and future. We place particular emphasis on different kinds of objections to the current Circular A-94, which call for somewhat different reforms. For example, the frequent disconnect between efficiency and welfare will justify the use of distributional weights, designed to “upweight” benefit to those without much money and to “downweight” benefits to those with a great deal of it. This is particularly important as FEMA works on piloting weighting, which would make it the first time the method has been used by the federal government.²⁴ We explore the new issues that agencies will soon need to confront as they make this work on the ground, with implications for benefit-cost analysis generally.

²² More broadly, we are trying to give some substance to a very specific context within the broader debate on equity-informed legal rules. See Louis Kaplow & Steven Shavell, *Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income*, 23 J. LEGAL STUD. 667 (1994); Chris Sanchirico, *Taxes Versus Legal Rules as Instruments for Equity: A More Equitable View*, 29 J. LEGAL STUD. 797, 805–06 (2000); Zachary Liscow, *Reducing Inequality on the Cheap: When Legal Rule Design Should Incorporate Equity as Well as Efficiency*, 123 YALE L. J. 2134 (2014).

²³ See Hemel, *supra* note 4.

²⁴ *Modernizing Regulatory Review: Exploring OMB’s Updated Benefit-Cost Guidance*, RES. FOR THE FUTURE (Apr. 11, 2023) <https://www.rff.org/events/rff-live/modernizing-regulatory-review-exploring-ombs-updated-benefit-cost-guidance/> (announcing the pilot).

II. Efficiency, Kaldor-Hicks, and Welfare

A. Efficiency in the Current Circular A-94

The current Circular A-94 grounds BCA in the goal of efficiency, but it is not entirely clear what that goal was understood to entail, or why the government should pursue it.²⁵ On a standard view, efficiency refers to Pareto improvements²⁶: at least one person is made better off while no one is made worse off.²⁷ But funding decisions do not involve Pareto improvements,²⁸ which means that for good reason, the current A-94 refers to *potential* Pareto improvements, not actual Pareto improvements. It says, with admirable and (to us) surprising clarity, that its goal of maximizing net benefits “is based on the premise that gainers could fully compensate the losers and still be better off.”²⁹ This explicit foundation, forgotten in subsequent decades, provides an opportunity for assessing the neglected history.

The basic idea must be that if the gainers gain (say) \$500 million, and the losers lose (say) \$200 million, the project is efficient in the sense that the gainers could transfer \$200 million to the losers, leaving a \$300 million surplus and no one worse off. In the context of funding decisions, such transfers do not happen in practice. To anchor the discussion in reality, a project designed to reduce the risk of flooding might cost \$200 million but deliver \$500 million in benefits. To understand what those figures mean, we would want to know much more, of course. Typically the \$200 million in cost would represent the financial cost of the project (and hence the amount of the federal grant), though it could also refer to state and local costs, and perhaps to costs to private parties (as might occur if, for example, a street was being relocated). The benefits would typically include the elimination or reduction in the likelihood of having to incur the replacement cost of the structures protected from flooding.

On the stipulated numbers, the benefits exceed the costs, but the recipient of funds is most unlikely to return \$200 million to FEMA or the U.S. Treasury. In fact, the very idea seems absurd. In an exhortation that to our knowledge has never been followed in the context of spending, the current A-94 even says that “[t]he presence or absence of such compensation should be indicated in the analysis,” along with discussion of any “significant distributional effects.”³⁰ With respect to compensation, we are not sure what the authors of the circular had in mind. Who, exactly, would be compensated? For what would they be compensated? The case of regulation is clearer; perhaps consumers or workers are being forced to pay more, or to

²⁵ This is of course a longstanding debate. See RICHARD A. POSNER, *UTILITARIANISM, ECONOMICS, AND LEGAL THEORY* (1979); Richard A. Posner, *Wealth Maximization Revisited*, 2 NOTRE DAME J. L., ETHICS & PUB. POL’Y 85 (1985); Ronald Dworkin, *Is Wealth A Value?*, 9 J. LEGAL STUD. 191 (1980); Jules Coleman, *Efficiency, Utility, and Wealth Maximization*, 8 HOFSTRA L. REV. 512 (1980). Coleman offered an especially clear outline of alternative conceptions of efficiency.

²⁶ See Guido Calabresi, *The Pointlessness of Pareto: Carrying Coase Further*, 100 YALE L.J. 1211 (1991).

²⁷ *Id.* at 1215.

²⁸ *Id.*

²⁹ CIRCULAR A-94, *supra* note 6.

³⁰ *Id.* at 12.

lose money, and we might at least ponder the possibility of compensating them (though regulators will typically lack authority to do so). In the context of funding, the issue is much fuzzier.

The Circular never says so explicitly, but it adopts the logic of Kaldor-Hicks efficiency.³¹ This was standard economic thinking in the 1990s³² and notwithstanding powerful and frequent academic objections,³³ it continues to have influence today. If we were speaking in terms of well-being, we might understand the Kaldor-Hicks criterion to be a form of applied utilitarianism: If the winners gain more than the losers lose, in terms of well-being, there is by hypothesis an overall increase in well-being. Still, we might not embrace the change: If those who are now miserable become more miserable, and those who are now flourishing flourish more, it is not clear that we should favor the change even if there is a net welfare gain. And if we are speaking in terms of money, things are even less clear. If rich people gain \$100 million and poor people lose \$50 million, is there a gain in well-being? Yes, if the rich compensate the poor; but what if they do not? If rich people gain \$100 million and taxpayers pay \$50 million, is there a gain in well-being? Yes, if the rich return (say) \$60 million to taxpayers; but what if they do not?

The examples are sufficient to show that in the context of funding decisions,³⁴ which of course do involve money, the Kaldor-Hicks approach raises immediate questions: Who are the losers, exactly? What follows from the fact that the winners are not going to compensate the losers? What role would the analysis of “significant distributional effects” then play? Suppose that communities apply for funding for a project that would cost \$500 million and deliver \$450 million in benefits. If the poor disproportionately benefit from a project, would a project with a benefit-cost ratio a little below 1, measured on efficiency grounds, become cost-justified? Suppose that communities apply for funding for a project that would cost \$500 million and deliver \$550 million in benefits. If the rich disproportionately benefit, would a project with a benefit-cost ratio a little above 1, measured on efficiency grounds, no longer be cost-justified? None of this was ever specified. And since again to our knowledge the exhortation to do the distributional analysis (which would be more work – and sometimes not easy work³⁵) was never followed, we do not know what would have happened had the analysis been done.

As we have noted, the outcome of most of the stylized examples sketched earlier is clear under past practice. The poor communities with low home values are unlikely to get a

³¹ For an outline, see Matthew Adler, *Beyond Efficiency and Procedure: A Welfarist Theory of Regulation*, 28 FLA. ST. U. L. REV. 241 (2000). See also Coleman, *supra* note 25.

³² See Diane Coyle et al., *Is It Time to Reboot Welfare Economics? Overview*, 44 FISCAL STUD. 109 (2023). For a normative defense, see Posner, *supra* note 25.

³³ See generally MATTHEW ADLER, *MEASURING SOCIAL WELFARE* (2019).

³⁴ And elsewhere. See *id.*

³⁵ Analysis of the incidence of regulatory costs and benefits may be challenging. Who, for example, benefits from clean air regulation, and who loses? Who benefits from road safety regulation, and who loses? For various accounts, see Hemel, *supra* note 4; Matthew E. Kahn, *The Beneficiaries of Clean Air Act Regulation*, 24 REGUL. 34 (2001).

FEMA grant because the benefit-cost ratio is below 1: the community is not willing to pay for the cost of the protection. The rich communities, in contrast, do get funding.

B. Welfare in the Draft Revised Circular A-94

The draft revised Circular A-94 offers a new account of the goal of BCA.³⁶ The revision states that its goal “is to promote social welfare,”³⁷ reflecting an evolution in thinking about CBA over the past three decades, as scholars have emphasized the stronger foundations of well-being³⁸ as the ultimate goal and the questionable empirical plausibility of that idea that winners will compensate losers.³⁹ Efficiency is essentially entirely removed as the ultimate goal for benefit-cost analysis in the draft revision, though (as it notes) of course all else equal greater efficiency is still better.

The draft Circular nevertheless continues to envision a rigorous quantitative analysis. It does not by any means abandon cost-benefit analysis. It seeks to mend it, not to end it. As with the current Circular, so with the draft revision: If a project would cost \$50 million and produce benefits of \$500 million (even calculated under the current Circular’s methods) agencies should fund it. The draft revision does not suggest that, if a project would have benefits of \$50 million and costs of \$500 million (calculated under the old methods), agencies should go forward. In fact it would be very hard, under the draft revision, to proceed with such numbers. How, then, should agencies analyze the stylized examples with which we began?

The Circular proposes as an option, though not a requirement, that agencies like FEMA use distributional weighting, which changes the measurement of costs and benefits. Distributional weighting works by measuring the net benefits of a project across the income distribution. And then, according to a pre-specified formula, lower-income net benefits are upweighted and higher-income net benefits are downweighted. FEMA has been working on a pilot to incorporate this method.⁴⁰

The draft revision proposes a specific formula, $w_i = \left(\frac{\bar{y}_i}{y_{med}}\right)^{-\varepsilon}$, in which the weight w on subgroup i is the average income y of group i divided by the median income, raised to the negative of the elasticity of marginal utility (ε).⁴¹ Note that this assumes a constant elasticity of

³⁶ In that sense, it is consistent with the general approach in MATTHEW ADLER & ERIC POSNER, *NEW FOUNDATIONS OF COST-BENEFIT ANALYSIS* (2006).

³⁷ OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, *CIRCULAR A-94 DRAFT FOR PUBLIC REVIEW, GUIDELINES AND DISCOUNT RATES FOR BENEFIT-COST ANALYSIS OF FEDERAL PROGRAMS 3* (2023).

³⁸ We use this term interchangeably with “welfare.”

³⁹ See Adler & Posner, *supra* note 36; Calabresi, *supra* note 26; MATTHEW ADLER, *WELL-BEING AND FAIR DISTRIBUTION* (2011); Liscow, *Is Efficiency Biased?* 85 U. CHI. L. REV. 1649 (2018); *FAIRNESS IN LAW AND ECONOMICS* (Lee Anne Fennell & Richard H. McAdams eds., 2014).

⁴⁰ Resources for the Future, *supra* note 24 (announcing the pilot).

⁴¹ The UK Green Book uses a similar method. HM TREASURY, *THE GREEN BOOK* (2022).

marginal utility. Consistent with proposed revisions to Circular A-4,⁴² the guidance recommends a default elasticity of marginal utility of 1.4, based on hedonic studies on happiness, estimates of within-person risk-aversion, and so forth.⁴³ For reference, logarithmic utility (commonly used for analytical convenience), in which a doubling of income results in a halving of the welfare weight, would have an elasticity of marginal utility of 1. So the default estimate is considerably more redistributive than that reference point.

A stylized example shows how weighting would work in practice. Suppose that FEMA has \$10 million to spend on a levee. It can spend the funds on one of two projects, detailed in Table 1. In Project A, FEMA could spend it saving 50 owner-occupied homes worth \$210,000 from certain destruction in a higher-income community (for total conventionally-measured benefits of \$10.5 million). Or, in Project B, it could save 100 \$90,000 homes in a lower-income community (for total conventionally-measured benefits of \$9 million). Suppose that these are the only benefits.⁴⁴ As BCA is conventionally practiced, Project A will win. In the lower-income community, the benefits do not justify the costs. And it is easy to see how richer communities will tend to be more likely to get funding: they usually have more valuable homes to protect.⁴⁵ This is problematic because not only are poorer communities less likely to get funding but also the poorer community values a dollar of housing value saved more than the richer community does; it is more valuable to them in welfare terms. Those most in need are least likely to get funding.

Consider how distributional weighting would change the analysis. Suppose that annual household income is \$150,000 in the richer community (Project A) and \$50,000 in the poorer community (Project B). And assume that all benefits are incident on the homeowners.⁴⁶ And assume that utility can be modeled as the logarithm of income, as noted a conventional assumption in economics (partially for its ease of analysis). If that is the case, then the project in the lower-income community has vastly greater welfare benefits, with welfare gains of 86 utils for the lower-income community and only 33 utils for the higher-income community, even as the higher-income project is funded under conventional analysis.⁴⁷ For a logarithmic utility function, the weight is conveniently equal to one divided by income. To make the weights easier to handle, the weights are divided by the weight for median income (\$75,000). So the weight for Project A is $(1/\$150,000)/(1/\$75,000)=1/2$. And the weight for Project B is

⁴² See OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, PREAMBLE: PROPOSED OMB CIRCULAR NO. A-4, "REGULATORY ANALYSIS" (2023); OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, CIRCULAR A-4 DRAFT FOR PUBLIC REVIEW (Apr. 6, 2023).

⁴³ *Id.* at 73 n.136.

⁴⁴ For example, we are putting aside the important disruptions to people's lives from having their homes destroyed.

⁴⁵ Technically, FEMA uses the cost of replacement rather than the market value in its calculations.

⁴⁶ So, assume that insurance rates adjust based on reduced harm, so that insurers do not benefit, for example.

⁴⁷ Assume that the cost of replacing the housing is born in equal amounts over 10 years. The welfare gains are the welfare with the project minus the welfare without the project, which are equal to: the number of years * the number of households * [log (full income) - log (income after annual housing losses)]. For project A, this welfare gain is equal to: 10 years * 50 households * [log (\$150k) - log (\$150k - \$21k)] = 33. For project B, this is equal to: 10 years * 100 households * [log (\$50k) - log (\$50k - \$9k)] = 86. So, the welfare gain is much larger for Project B.

$(1/\$50,000)/(1/\$75,000)=1.5$. Recalculating the benefits with these weights yields benefits of \$5.25 million (= \$10.5 million * ½) for Project A and benefits of \$13.5 million (= \$9 million * 1.5) for Project B. Now the low-income community receives the funding. Under this analysis, the lower-income community has a fairer shot at getting funds: twice as many homes are saved, and those who most need the funds get them.

Table 1: Example of Distributional Weighting

	Project A: higher-income community	Project B: lower-income community
Conventionally-measured benefits	50 houses * \$210,000 = \$10.5m	100 houses * \$90,000 = \$9m
Household income	\$150,000/year	\$50,000/year
Weight	$(1/\$150k) / (1/\$75k) = 1/2$	$(1/\$50k) / (1/\$75k) = 1.5$
Reweighted benefits	$\$10.5m * 1/2 = \$5.25m$	$\$9m * 1.5 = \$13.5m$

The draft revision also offers as an option that agencies can “income-average,” or use an average value for a given benefit or cost across the whole population.⁴⁸ It is important to see that this is arguably what is already done for the value of statistical life (VSL).⁴⁹ Agencies use a number – now in the vicinity of \$12 million⁵⁰ – regardless of whether poor people or rich people are facing the statistical mortality risk. Agencies do not say (for example) that the VSL is \$2 million for poor people and \$20 million for wealthy people, even if (as seems likely) poor people are willing to pay or accept far less to reduce statistical risks than are wealthy people.⁵¹ In this sense, a poor person’s life is “upweighted,” at least if it is measured in terms of willingness to pay to reduce statistical risks or willingness to accept to face such risks.⁵² Something similar could be done for things like home values. Instead of treating poor people’s homes as worth less than those of wealthy people, homes, as such, might be given an average value.⁵³ There are downsides to this approach: in particular, it ignores variations in home prices,

⁴⁸ CIRCULAR A-94 DRAFT FOR PUBLIC REVIEW, *supra* note 37.

⁴⁹ Kniesner & Viscusi, *supra* note 4.

⁵⁰ See *Departmental Guidance on Valuation of a Statistical Life in Economic Analysis*, U.S. DEP’T OF TRANSP. (Mar. 23, 2021), <https://www.transportation.gov/office-policy/transportation-policy/revised-departmental-guidance-on-valuation-of-a-statistical-life-in-economic-analysis>.

⁵¹ See Cass R. Sunstein, *Valuing Life: A Plea for Disaggregation*, 54 Duke L.J. 384 (2004).

⁵² In the context of regulation, the incidence question raises many complications; it is not clear that upweighting poor people’s lives (with respect to safety) helps poor people, if it makes them pay for levels of safety for which they are not willing to pay. See Hemel, *supra* note 4.

⁵³ As the draft rules note, weights should not be applied to measures, like the value of statistical life, that have already been income averaged in ways that ignore that previous averaging. In welfare terms, a life is worth the same whether it is a rich or poor person. Of course, there could still be weighting of values like VSL, though there

since sometimes lower-income communities have higher-valued homes (which are especially valuable to protect) and vice versa, leading to a worse targeting of resources from a welfare perspective. But it is at least facially simpler to implement and may have a certain ethical appeal.

Return in this light to the stylized cases with which we began and consider them in light of the draft revision. In cases (2) and (3), involving poor communities, it is plausible that under the draft revision, funding would be available, even if the monetized benefits are lower than the monetized costs under the old efficiency metric. Everything would depend, of course, on how the analysis works out. But in *poor community, wildfire risks*, and in *poor community, extreme heat*, at least, it is perfectly imaginable that the adjusted numbers would allow for funding.

The impact of the draft revision should be simultaneously to increase welfare and to promote equity. Out of the fixed funds that agencies have, the current circular wrongly forbids money from being directed to lower-income areas in cases like (2) and (3). It is nonetheless true that in terms of direct effects (i.e., ignoring possibly beneficial effects of greater equity on the broader economy), the draft revision does harm efficiency as standardly defined. In the levee example with distributional weights, funds will go toward protecting homes that are worth less as judged by willingness to pay to protect them: \$10 million are spent protecting homes worth \$9 million. But the result is to increase welfare while promoting equity. To be sure, it would be possible to act in a way that decreases welfare but increases equity – as, for example, by funding projects that cost a great deal and provide benefits that, while modest, are concentrated in low-income communities. But the draft revision does not authorize that, and as long as agencies do not go too far in considering equity, what it does authorize will increase welfare.

C. Four Distinct Justifications

This change from efficiency to welfare is justified on four distinct but partially interconnected grounds. Each of them would justify a shift from the current Circular A-94 guidance toward an approach that promotes welfare better, values equity more, and spends more in lower-income communities. The first two are normative goals: attending to the declining marginal utility of income and correcting the bias of efficiency-based analysis to direct resources toward the well-off. The third and fourth reasons are empirical: the premises underlying efficiency-based analysis do not hold up if the ultimate goal is welfare.

The first reason is the declining marginal utility of income, a justification explicitly mentioned in the proposed new circular. Essentially, the argument goes, welfare increases more when \$100 goes to poor people than when the same amount goes to rich people. If each of the ten poorest people in the United States are given \$100, they would gain far more, in

are ethical and economic questions about doing so. See Hemel, *supra* note 4. In any case, though, that weighting should consider the averaging that has already taken place.

terms of welfare, than each of ten richest people in the United States would lose if they lost \$100. Looking *within* an individual, this is why people would want home insurance: the \$1 million to rebuild a home is particularly valuable to you when you have been made poorer by losing a \$1 million asset when your house burns down. Distributional weighting shifts grant funding to those with lower incomes and thus a higher marginal utility, thereby increasing welfare. Income averaging in effect does the same, though not necessarily in a way that precisely corresponds to what the marginal utility of income would imply should happen.

The second reason is bias-correcting.⁵⁴ This is mathematically tied to the first reason, but conceptually distinct.⁵⁵ As the draft revision notes, “[w]eights can also be based on the tendency of market-based methods to value things that low-income people use at low amounts.”⁵⁶ As the example earlier noted, low-income people cannot pay as much for housing, so they are going to tend to be willing to pay less for housing and live in smaller houses. Under a purely efficiency-based framework, this tendency will direct spending to richer places. Relative to an egalitarian goal of spending equally across the income distribution as a matter of outcome (or perhaps a prioritarian goal of giving spending to help the least well-off, or perhaps a goal of equality in treatment of individuals as a matter of *process*), the efficiency-based framework presents a bias. Income averaging directly corrects for this bias. Distributional weighting does so as well, though not necessarily to the precise degree that would correct for this bias.

The third and fourth reasons for the switch toward *directly* accounting for welfare in BCA arise from the possibility that versions of well-being have actually been the ultimate goal all along, with efficient BCA being an intermediate instrumental goal to maximize overall well-being.⁵⁷ But the logic that efficient BCA can maximize well-being depends upon some empirical foundations that turn out to be flawed.

Arguably the current guidance did not say what its *ultimate* goal was, since it is premised on the idea that policy winners could (would?) compensate the losers. On one view, the logic is Paretian, in which case it is unobjectionable, at least as an improvement from the status quo if not necessarily as an optimal solution for welfarists. (The reason is that redistribution from rich to poor might be even better, on welfarist grounds, than a Pareto improvement.) This Paretian logic is appealing for those who emphasize the challenges of measuring social welfare.⁵⁸ But as we have seen, the winners do not compensate the losers when funds are awarded. Because compensation does not occur, eliminating the underlying Paretian logic, it makes sense (and this is the third reason for the switch) to shift the focus from efficiency to welfare. The idea of the current Circular was that, with compensation to losers,

⁵⁴ See Liscow, *supra* note 39; Dworkin, *supra* note 25. Furthermore, the bias can grow over time as resources are directed to the rich, driving up their willingness to pay, thus driving more resources to them. William Giraldo & Zachary Liscow, *Inequality Snowballing*, <https://ssrn.com/abstract=3327460> (mimeo).

⁵⁵ *Id.*

⁵⁶ CIRCULAR A-94 DRAFT FOR PUBLIC REVIEW, *supra* note 37 at 16.

⁵⁷ See Adler & Posner, *supra* note 36.

⁵⁸ For example, the challenges of measuring cardinal welfare and the incomparability of utility among individuals.

agencies could maximize the size of the pie without compromising distributional goals. From our perspective, this never made any legal sense. How *could* the recipients of FEMA grants compensate the losers, even aggregating across income groups? FEMA simply does not have that authority.

At a higher level of abstraction, perhaps Congress could compensate the losers. But how realistic is that? If a funded project results in a relocation of a highway, will Congress pay those who lose from the relocation? There is scant evidence that this ever happens. For example, famously, after the US let China accede to the WTO, predictably (given prevailing theories of trade) causing significant harm to working-class Americans and benefitting better-off Americans, there was little compensation to the harmed workers.⁵⁹ And after state supreme courts direct legislatures to spend more on poor school districts, those who lose out receive no compensation; the consequences stick.⁶⁰ The broader point is that when some agency gives out grant money – for example, to reduce the risks associated with flooding -- no one compensates the losers.

More deeply, it is not clear who the “losers” are in the context of spending. In the regulatory context, a business may have higher costs (for example, as a result of pollution abatement); if so, the business is the loser. (It is true that a business may pass along costs to customers and workers, but the starting point for finding the loser is at least clear.) For spending, are the losers those parties who have to pay something, or who lose some good or service, as a result of the funded project? Or are the losers those who do not receive spending? Or are taxpayers as a whole the losers? And, if so, how do we attribute impacts across the income distribution? Should individuals receive a pro rata share? In proportion to their tax burden? Actually, given the Paretian logic, the issue presumably is not so much dividing the population into income bins but rather assessing individuals *individually*. Little wonder then that, to our knowledge, there was never an assessment of compensation in the spending context. It is a good thing to be discarding this logic.

Fourth, a related but distinct empirical point involves the now-familiar argument that the best way to redistribute—even from the perspective of social welfare—is through taxes-and-transfers, not through regulations or spending in ways other than transferring cash. That argument is not explicitly stated in the current guidance, but it has long been prominent, dating back in public finance economics to Richard Musgrave’s 1959 textbook⁶¹ that separated the redistributive from the efficiency-achieving branches of government. It is fair to say that the argument is implicit in the current guidance. In recent decades, the argument has become associated with its most articulate expositors, Louis Kaplow and Steve Shavell.⁶²

⁵⁹ See David H. Autor, David Dorn & Gordon H. Hanson, *The China Syndrome: Local Labor Market Effects of Import Competition in the United States*, 103 AM. ECON. REV. 2121 (2013).

⁶⁰ Zachary Liscow, *Are Court Orders Sticky? Evidence on Distributional Impacts from School Finance Litigation*, 15 J. OF EMPIRICAL LEGAL STUD. 4 (2018).

⁶¹ RICHARD A. MUSGRAVE, *THE THEORY OF PUBLIC FINANCE: A STUDY IN PUBLIC ECONOMY* (1959). After this, luminaries such as James Tobin and Richard Posner discussed this argument.

⁶² See Kaplow & Shavell *supra* note 22.

The essential argument is that it maximizes welfare to redistribute through taxation, rather than spending or regulation, because everyone values a dollar at a dollar, while higher-income people typically value things, such as good and services, more than lower-income people do (in monetary terms). The result is that any kind of non-cash redistribution is often inefficient because it in effect denies things to those who value them more and provides them to those who value them less. (And, to be clear, “cash” means money itself, not things funded by cash.) Consider the numerical levee example above. When the poorer community gets a levee funded by cash, they have received a levee valued at less (\$9 million) than the value of the cash spent on it (\$10 million). In fact, they would rather have received \$9.5 million of cash than the more expensive levee. This is the nature of deviating from the efficient rule. Cash (via taxation) lets the recipients choose how to spend, allowing them to spend it on the things most valuable to them.

By this logic, perhaps the goal has been well-being all along, but if so, efficiency was just an instrumental goal to achieve maximize social welfare. The strategy is for spending like levies funds to go disproportionately to the rich, make them better off by protecting their expensive homes, and then tax them more to transfer more cash to the poor. In other words, grow the size of the economic pie so that more can be redistributed through taxation.

That argument is not without its logic, but note that it does require an *optimal tax code* that redistributes the socially optimal amount. Few people think that we have such a tax code. It is not plausible that the poor community of 100 families will receive the \$9.5 million of new cash transfers. FEMA cannot transfer it. And the tax system does not seem posed to do so (and, of course, if it did, then FEMA could alter its benefit-cost procedures). There is good reason to think that we should not *ever* expect to have an optimal tax code, in the sense of maximizing social welfare on its own.⁶³ Notably, there is strong public support for greater economic equality, especially through providing necessities like health care⁶⁴ and, we suspect, through having an equal shot – whether rich or poor - at FEMA funding to protect from disasters (which is considered “redistribution” to relative to economists’ efficient baseline). At the same time, there is public aversion to redistribution through specifically taxation, which can appear confiscatory regarding the rich’s “earned” income and undeserved regarding transfers toward the poor who do not work for it.⁶⁵ For example, the public is much more comfortable transferring an amount that can be spent only on necessities than an equivalent amount of cash that can be spent on anything.⁶⁶ So it makes sense that the tax code that we see implies social views that are vastly less redistributive than individuals’ risk aversion or hedonic studies

⁶³ Zachary Liscow, *Redistribution for Realists*, 107 IOWA L. REV. 495 (2022); STEVE SHEFFRIN, TAX FAIRNESS AND FOLK JUSTICE.

⁶⁴ *Id.*

⁶⁵ *Id.*

⁶⁶ Zachary Liscow and Abigail Pershing, *Why Is So Much Redistribution In-Kind and Not in Cash? Evidence from a Survey Experiment*, 75 NAT’L TAX J. 75 313 (2022).

suggest.⁶⁷ That redistribution needs to take place elsewhere, such as via FEMA, where it is worth noting that Congress has never said that the spending should be “efficient” at the expense of the poor. This failure to redistribute sufficiently through taxation provides a final reason to consider distribution through government spending.

All four of these reasons directionally support the shift toward direct incorporation of distribution in BCA. But they vary in their implications for the precise way or amount to incorporate distributional concerns to maximize well-being.

First: If the goal is to address the declining marginal utility of income, then simply using distributional weights that correspond to those marginal utilities maximizes welfare. This is what the distributional weights formula does. The minimal goal of the draft revised circular is to authorize agencies to achieve that goal, and thus to allow funding in imaginable variations on the second and fourth scenarios in our opening list.

Second: If the goal is to correct the bias of spending more in richer places or equivalently to treat individuals the same way (in the sense of valuing their preferences equally), then using average values would be the best tailored approach. The draft revised circular allows agencies to correct that bias. For example, FEMA could use the average housing values in the country in its analysis. How distribution-focused income-averaging is relative to distributional weights, of course, depends upon the distributional weights; the higher the weights being considered, the more likely they will do more to help the poor than income-averaging.

Third: If the goal is to address the failure to compensate for distributional impacts, then complicated questions of political economy (typically little discussed in government guidance documents) arise: How much compensation is there? And compensation by whom? The more compensation, the more efficiency-based the analysis should be. On one extreme, with complete compensation, a weighting scheme could still make sense, on one of the first two grounds: redistributing to those with a higher marginal utility of income or equal treatment. But, at least on marginal utility grounds, it would make sense to redistribute less with compensation than if there were no compensation. On the other extreme, with no compensation, agencies would just go back to the underlying welfare goal, such as accounting for the declining marginal utility of income. The revised circular allows agencies to do precisely that.

Finally: If the goal is to compensate for the insufficiency of redistribution through the tax code, then the question immediately arises of *by how much* the tax code fails to redistribute. The more the tax code is redistributing, the less is needed through non-tax means. It would be challenging and perhaps impossible, of course, for agencies to try to resolve such questions in the context of spending programs; in any case, they lack the authority to decide on the appropriate amount of redistribution. Distributional weights, of the sort suggested by the

⁶⁷ See Nathaniel Hendren, *Measuring Economic Efficiency Using Inverse-Optimum Weights*, 187 J. OF PUB. ECON (2022); *Redistribution for Realists*, *supra* note 63.

draft revised circular, would increase welfare while also promoting distributional goals in a way that should be broadly agreeable.

Overall then the proposed changes to Circular A-94 give agencies the opportunity and options to achieve the welfare-maximizing outcome, to the extent consistent with law, depending upon what their ultimate underlying thinking is.

III. From Theory to Practice: Pursuing Welfare with Benefit-Cost Analysis

Turning from theory to practice, we emphasize that several questions need to be worked out for agencies to maximize welfare and consider equity in their spending programs and benefit-cost analyses more generally, now that agencies are beginning to pursue the goal in earnest. Importantly, some of these are specific to spending, and the analysis would be different in the regulatory context.

A. Direct Incorporation of Equity into BCA

A remarkable feature of BCA of spending is that, for decades, it was *always supposed* to be considering distribution, given that Circular A-94 told agencies to analyze it. The federal government just never to our knowledge did so for spending – and did not do so often for regulation.⁶⁸ Part of this reticence could have resulted from the fact that it was unclear how to do so.

A commonly made argument is that agencies should consider distribution, but that CBA itself should not.⁶⁹ On this view, CBA should be undertaken without regard to distribution, but agencies should be entitled to consider distribution separately or after they have engaged in CBA. That view is not implausible, and cannot be rejected in the abstract, but two arguments militate against it. First, decades of practice suggest that this approach simply does not work, perhaps especially in the domain of spending. Distribution has made little or no appearance in spending decisions, which is part of why FEMA has been harshly criticized for “systematic favoritism toward wealthy and white people” in its spending.⁷⁰ The OMB Director has apologized for the current guidance in front of Congress for not giving a fair shake to lower-income communities.⁷¹ Second, and more fundamentally, it is *useful* to have distribution be explicitly part of the decision criteria through weights (or alternatively income averaging). Circular A-94, in its present and proposed forms, sets BCA up with the goal of being

⁶⁸ For regulatory impact analysis, distribution was supposed to be considered under EO 13563 and EO 12866, but not under EO 12291.

⁶⁹ See, e.g., Dudley & Viscusi, *supra* note 10; Susan E. Dudley et al., *supra* note 10.

⁷⁰ Frank, *supra* note 21. See also Thomas Frank, *FEMA Flood Program Could Violate Civil Rights Law*, POLITICO (June 16, 2022), <https://www.politico.com/news/2022/06/16/fema-flood-program-civil-rights-00037261>.

⁷¹ See Press Release, Alex Padilla, Padilla Presses Administration to Prioritize Pajaro River Flood Projects and Protect Low-Income Communities (Mar. 15, 2023), <https://www.padilla.senate.gov/newsroom/press-releases/padilla-presses-administration-to-prioritize-pajaro-river-flood-projects-and-protect-low-income-communities/>.

quantitative.⁷² So, when a feature that matters, such as distribution, can be quantified, it stands to reason that it is most helpful to quantify the factor. Doing so helps make better decisions by being more rigorous.

The case of FEMA funding shows the value of direct incorporation of equity in BCA. The way FEMA funding has historically worked is that, to be considered at all, the benefit-cost ratio (BCR) needs to be above 1. Applicants could not receive funding if the BCR was below 1, meaning that low-income communities with low housing values could not even be considered in some cases, even if the need for their protection was desperate. Presumably it would be possible to allow a project with an efficiency-based BCR of any value to apply – and just let FEMA decide based on distributional considerations. But that would be problematic for two reasons. First, it is costly to apply for grants, and it is helpful to know beforehand about one’s likelihood of getting funding. Second, to take distribution into account, FEMA would still need some criteria to weigh distributional concerns – and being explicit and systematic is better government practice than being vague, ad hoc, and arbitrary.⁷³

All this suggests that, institutionally, it is helpful to be clearer in guidance about how to consider distribution, including a specific method for doing so. Of course, time will tell what the ultimate impact is. But BCA tends toward a single metric to maximize welfare. Distributional weights as part of that metric accomplish that goal. The remarkable dearth of cases in which distribution was considered in spending decisions suggests that the previous regime did not.

B. Optional Use of Distributional Weights

Another notable feature of the proposed revisions to Circular A-94 is that the use of distributional weights is explicitly optional rather than mandatory. In our view, this is also wise as a matter of institutional design. To our knowledge, distributional weights have never been used in any form by the federal government. Allowing agencies to use them, by making that an option rather than a requirement, seems prudent. One reason is that there could be legal constraints on their use in some cases. Another reason is that issues of feasibility and administrability matter, and agencies are likely to be in the best position to consider those issues. This is a promising area for institutional learning – for “learning by doing.”

⁷² CIRCULAR A-94 DRAFT FOR PUBLIC REVIEW, *supra* note 37 at 4-5 (“Quantifying costs and benefits to the extent possible is expected, even if assigning monetary values is not feasible.”).

⁷³ Agencies could, of course, conduct two versions of analyses: one with distributional weights of 1 for everyone and another with distributional weights different from 1. If there is an interest in public transparency in seeing the traditional purely efficiency-based numbers, that could be useful in the regulatory context. And nothing in the new rules precludes it. However, for spending decisions, unlike for regulation, it is not even the case that transparency would be served by running two sets of number because projects’ benefit-cost analyses are typically not released publicly. (If those numbers were released – which does seem useful to us for transparency purposes – then having the two sets of numbers could be helpful.)

C. What Distributional Weights to Use

The use of distributional weights immediately begs the question of what weights to use. We do not have a firm recommendation. As noted, the guidance offers a default based on hedonic studies of happiness and individuals' risk aversion. Ultimately, though, the choice is ineluctably normative – and also unavoidable. Having weights of 1 – the case under the standard efficiency analysis – is also a normative choice.

It is clear, though, that the weights should be based on *post-tax-and-transfer* income. If a community has average pre-tax household income of \$30,000, it is in a very different situation if – after taxes and transfers – it has an average post-tax household income of \$20,000 versus \$40,000. The weights should reflect that. Cash (or near-tax) transfers like food stamps or the Earned Income Tax credit (or equivalently income taxes) are relatively easy to measure. For example, taxes and transfers through the tax code can be estimated with great accuracy with the Natural Bureau of Economic Research's TAXSIM model.⁷⁴ In-kind transfers like healthcare are harder to measure. The fact in the median state Medicaid spending was \$8,436 per beneficiary does not mean that that amount should be added to the income of very low-income communities not only because this estimate can vary greatly across age but also because – to measure income in financial terms – it is not the spending that matters, but rather how much families value the spending *in financial terms*, which is typically considerably less than actual spending.⁷⁵ Estimating these things precisely initially is not essential. But implementers of distributional weighting should over time develop methods to account for these taxes and transfers to ensure the best targeting of scarce funds.

Finally, note the connection to the fourth justification above for considering equity in BCA: that taxes do not redistribute enough and, the more they redistribute, the less benefit-cost analysis should do so. This will happen naturally by using after-tax income weights: if transfers to the poor go up, their income weights go down.

D. Measuring Income of Affected Communities

Measuring the income of affected parties might seem relatively straightforward, but it requires making some conceptually difficult choices. Importantly, communities, rather than individuals, are typically affected by grant funding, which eases the measurement burden (as well as makes it less likely that weighting seems too extreme, because averaging across the individuals in communities will tend to yield less divergent weights than if individuals' incomes were used). Recent average income for communities is readily available from the Census. If the main goal is distributional weighting based on the average income of the community receiving the funds, it should be relatively simple to work from the Census. To the extent that agencies engage in income-averaging, assuming the same value for all homes—to adjust for the bias

⁷⁴ TAXSIM, NAT'L BUREAU OF ECON. RSCH, <https://www.nber.org/research/data/taxsim> (last accessed Sept. 30, 2023).

⁷⁵ Amy Finkelstein et al., *The Value of Medicaid, Interpreting Results from the Oregon Health Insurance Experiment*, 127 J. OF POLITICAL ECON. (2019).

toward funding to protect the homes of the well-off—would actually ease the measurement burden, since agencies would be required only to know the number of homes, rather than the number of homes and the value of each home.

Of course, questions remain. For example, assuming that weighting is done on the basis of the affected community rather than the particular affected individuals, at what level should the “community” be measured? For example, is the relevant geography a Census block (the smallest Census subdivision), a Census tract (about 4,000 inhabitants), a county, or something else?⁷⁶ We would suggest using Census tracts, which are used for Opportunity Zone tax credit designations, for example.⁷⁷ They are small enough to capture localized areas of a certain level of income without being so small that spillovers across space are ignored.

Of course, those spillovers could vary across projects, and agencies might want to consider exceptions when the geographic scope of a project is plausibly substantially larger or smaller than a Census tract. Note, though, that the issue of measuring the scope of impacts is nothing new to benefit-cost analysis. To measure benefits, agencies have long needed to know what scope of geography benefits. The proposed rules merely suggest the possibility of considering the relative poverty of whoever is within that scope.

E. Incidence of Spending

As Circular A-94 notes, distributional analysis should take into account the ultimate incidence. The incidence of spending is typically clearer than the incidence of regulation. If regulators require safer cars, consumers do not just receive safer cars; they have to *pay* for the cars. In contrast, consumers do not have to pay for the spending that they receive, making the incidence clearer.

Nevertheless, even for spending, incidence is not completely straightforward – and incidence matters with distributional weights in a way that it does not without them, since who gets what matters. Take the example of building a levee to protect homes. Suppose first that the homes are owner-occupied. There it is a reasonable (though imperfect) assumption that homeowners capture the benefit of the levee by having higher home values, at least if their home insurance rates reflect the reduced risk of flooding. But about a third of households – especially among lower-income families – do not own their homes. And here it is not reasonable to assume that the occupants of the homes will capture the benefits. In particular, they will likely pay higher rents, and their landlords will instead benefit. There is no guarantee that the owners will capture all the benefits; maybe rents will just go up a portion of the value of the levee.

⁷⁶ See *Glossary*, U.S. CENSUS BUREAU (Apr. 11, 2022), <https://www.census.gov/programs-surveys/geography/about/glossary.html>.

⁷⁷ *Opportunity Zones*, IRS, <https://www.irs.gov/credits-deductions/businesses/opportunity-zones> (last updated Nov. 18, 2022).

There are rules of thumb that agencies can follow; for example, the more elastic the housing supply, the more renters will capture the benefits.⁷⁸ In any case, because distribution has not been incorporated much if at all into benefit-cost analysis so far, the tools for measuring incidence in this context are not very well developed. It is important to develop them now. Otherwise, spending will be targeted worse – in some cases, spending funds with little benefit because they neither promote much efficiency or actually reach those in need.

F. Uncertainty and Equity

The current Circular A-94 largely assumes that the impacts are known. But what if there is uncertainty, as Case 4 in the introduction suggested? The draft revised Circular A-94 continues to put a premium on quantification, but it notes that “[s]ome benefits and costs may be either difficult to monetize or difficult to quantify.”⁷⁹ In such cases, it remains “important that these be included the analysis.”⁸⁰ In some cases, the nonmonetized benefits or costs might be important, and “a threshold or break-even analysis may be considered for inclusion in the benefit-cost analysis.”⁸¹ This general point might matter for equity, because poor communities might find it especially challenging to produce a comprehensive cost-benefit analysis for grant applications.

In such cases, the draft revised Circular A-94 states that an agency “asks what magnitude unmonetized benefits and costs would need to have for the project or program to yield positive discounted net benefits, or for one alternative to overtake another in terms of discounted net benefits.”⁸² The current Circular A-94 also recognizes the problem of uncertainty and offers guidance for how to handle it.⁸³ Where data are sparse, breakeven analysis may be the best that agencies can do.⁸⁴ It should be a useful tool for grant applications, and it might make a significant difference for poor communities.⁸⁵

G. Net Benefits vs. Benefit-Cost Ratio

For regulation, it makes sense to use net benefits. A regulation with benefits of \$10 and costs of \$1, has a ratio of 10 to 1, but it is not nearly as good as a program with benefits of \$500 million, and costs of \$250 million with a ratio of 2 to 1.⁸⁶ The analysis must be different for individual spending projects, for which the ratio is most relevant. With a fixed budget of (say)

⁷⁸ One can think about it this way: if the housing supply is inelastic, then there are fewer places for renters to move to, and the owners of the buildings can capture more of the benefits.

⁷⁹ See CIRCULAR A-94 DRAFT FOR PUBLIC REVIEW, *supra* note 37 at 7.

⁸⁰ *Id.*

⁸¹ *Id.*

⁸² *Id.* at 7-8.

⁸³ CIRCULAR A-94, *supra* note 6, at 10.

⁸⁴ See Cass R. Sunstein, *The Limits of Quantification*, 102 CAL. L. REV. 1369 (2014).

⁸⁵ See *id.*; Richard L. Revesz and Samantha Yi, *Distributional Consequences and Regulatory Analysis*, 52 Env. Law 53 (2022).

⁸⁶ See CIRCULAR A-4, *supra* note 3, for a clear discussion.

\$250 million, choosing individual projects with the highest benefit-cost ratios will maximize welfare. Spending all \$250 million on the one project with a benefit-cost ratio of two would not be as good as spending funds on smaller projects, each with higher benefit-cost ratios but lower net benefits.

For example, suppose that there are three options for projects: A, B, and C. And suppose that FEMA has \$250 million to spend. For project A, benefits are \$200 million and costs are \$100 million. This results in a BCR of 2 and net benefits of \$100 million. For project B, benefits are \$125 million, and costs are \$50 million. This results in a BCR of 2.5 and net benefits of \$75 million. Project C has the same costs and benefits as project B. If FEMA ranks projects by net benefits, it will choose project A. If FEMA ranks projects by BCR, it will choose projects B and C. But choosing projects B and C ultimately produces higher net benefits (\$150 million) than doing project A (\$100 million). So it is better to choose individual projects (which is what agencies do) by their BCR, not their net benefits.

There are two crucial differences between regulation and spending. The first is that BCA applies to each individual projects for spending, whereas with regulation, there is typically BCA for one rule affecting several or many actors. The second is that spending is typically subject to strict statutory budget constraints, whereas agencies promulgating regulations are bound by no such tight constraints.⁸⁷ Of course, agencies promulgating regulations have staffing constraints, but that is a constraint on the inputs into *producing* regulations, not on the social resources expended as a result of the regulation. Given those constraints on producing regulations, each one should have the highest possible net benefits.⁸⁸ In contrast, for spending, agencies' need to choose individual projects subject to a budget constraint drives the need to focus on BCR.

⁸⁷ We are supposing that no regulatory budget is in place. The Trump Administration did impose a kind of regulatory budget in terms of the number of regulations. See Exec. Order 13,771, Reducing Regulation and Controlling Regulatory Costs (Jan. 30, 2017). Even under the Trump rule, since the constraint was the number of regulations rather than the social resources expended as a result of the regulation, regulating agencies should still maximize net benefits for each regulation allowed to them.

⁸⁸ Often benefits are taxed. For example, if a levee enhances the productivity of a farm (by making sure that its crops are not destroyed), the government will tax some of that. This raises the question of how that tax revenue should be incorporated into a weighted benefit-cost ratio. We have two recommendations. First, tax revenue should *not* be weighted as if it is received by the person who initially earns the money; the government gets it, so it should be treated as such. Of course, this extra piece of information raises the information burden, but ideally it would be incorporated, especially since taxes can be quite substantial.

Second, how to incorporate tax revenue into the BCR depends on how one thinks the revenue will be used. In the “marginal value of public funds” approach, tax revenue is subtracted from the benefits (numerator) and the costs (denominator). Nathaniel Hendren & Ben Sprung-Keyser, *A Unified Welfare Analysis of Government Policies*, 135 QUARTERLY J. ECON. 1209 (2020); Louis Kaplow, *A Unified Perspective on Efficiency, Redistribution, and Public Policy*, 73 NAT'L TAX J. 429 (2020). This is the ideal approach if a unified government fungibly directs resources to their most beneficial uses, since—by subtracting tax revenue from the denominator—the approach treats tax revenue the same way as the grant funding. If one believes instead that the tax revenue will not be used similarly to that, it should be included in the numerator, where the weighting likewise depends on assumptions about how the funds are used. For example, if the funds will be given lump sum back to everyone in the population, their weight is 1. If they will be spent somewhere with known benefits (or offset distortionary taxes), then those benefits are the weights. The approach matters. Consider a choice between two projects: Project A has \$220 of nontaxed benefits

To be sure, it remains true that the net benefits are ultimately what matter. But perhaps counterintuitively, the way to maximize net benefits for a given amount of spending is to choose the projects with the highest benefit-cost ratios.

H. Administrative Burden

There is also an important difference between who bears the administrative burden of conducting cost-benefit analysis for regulations and spending. In the regulatory context, there is typically one big regulation, and the rulemaking agency develops it, with interagency support. Notably, presidents have long required a full-scale regulatory impact analysis only for regulations with an annual economic impact of \$100 million or more⁸⁹; President Biden changed that number to \$200 million.⁹⁰

In the spending context, by contrast, large numbers of local communities apply for grants that are often small, sometimes as small as a hundred thousand dollars or less. (These grants often add up to a great deal, but may be individually small.) The fact that communities themselves must apply raises the question of the administrative burden *of the benefit-cost analysis itself*. It is worth asking: Does benefit-cost analysis survive benefit-cost analysis?

This question suggests another dimension of equity in benefit-cost analysis: ensuring administrative burdens, or “sludge,” that are not prohibitive, especially to communities least equipped to engage in BCA. And it suggests a significant benefit in a process in which the agency does as much of the BCA itself, rather than making communities use expensive consultants, potentially costing tens of thousands of dollars in some cases according to the GAO.⁹¹ The potentially complicated nature of the application process also suggests that it is valuable, when incorporating distributional consequences, not to make the process substantially more complicated for applicants. For example, one can imagine requiring communities to conduct their own distributional analyses to allow distributional weighting – and then for that to discourage applying for funds because of the additional requirements, especially among disadvantaged communities. This again suggests that as much as possible of the analysis is done by the agency itself, which presumably also can benefit from familiarity with the relevant issues and economies of scale.

and \$100 of costs. Project B has only \$200 of benefits (e.g., increased earnings) taxed at 20% and \$100 of costs. FEMA has \$100 to spend. Which project should FEMA choose? Project A has net benefits of \$120 and a BCR of 2.2. Project B has net benefits of \$100. Its BCR depends on how the tax revenue of \$40 is counted: if included in the denominator and numerator (with weight of one), the BCR is 2 (= \$200 / \$100); if excluded from both, the BCR is 2.7 (= (\$200 - \$40)/(\$100 - \$40)). So, how one treats the tax revenue determines which project has a higher BCR.

⁸⁹ See Exec. Order 12,866, 58 Fed. Reg. 190 (Oct. 4, 1993).

⁹⁰ See Exec. Order of April 6, 2023, Modernizing Regulatory Review (Apr. 6, 2023).

⁹¹ The Government Accountability Office white paper reported from interviews with local officials who had hired a contractor to perform BCAs for FEMA-funded mitigation activities. See U.S. GOV'T. ACCOUNTABILITY OFF., DISASTER RESILIENCE: FEMA SHOULD TAKE ADDITIONAL STEPS TO STREAMLINE HAZARD MITIGATION GRANTS AND ASSESS PROGRAM EFFECTS 18 (2021).

V. Conclusion

If people receive a federal subsidy, they are highly likely to be better off. At the same time, the current Circular A-94 rightly recognizes that if a project would deliver \$10 million in benefits at a cost of \$500 million, it is most unlikely to be a good idea. The problem with the current circular does not consist in its recognition of the importance of benefit-cost analysis; it rests in its emphasis on economic efficiency and its failure to note the possibility that even if the monetized costs of a proposed project are higher than the monetized benefits, it might increase social welfare. We have offered several examples of cases in which that might be so.

Responding to standard objections to efficiency-based cost-benefit analysis, the draft revised Circular A-94 directs agencies to focus on welfare, not efficiency. Among other things, and for the first time, it authorizes agencies to use distributional weights. The most straightforward justification of the proposed approach is that the use of such weights will increase the likelihood that cost-benefit analysis will promote rather than frustrate its fundamental goal, which is to increase welfare.

The current efficiency-based formula of Circular A-94 faces normative and empirical challenges. At the same time, the new welfare-based methods run into their own set of challenges, such as specification of distributional weights, measurement questions, and administrative burdens. We have offered some preliminary thoughts on how best to approach these challenges.

In our view, there is no doubt the proposed changes will move national policy, with respect to funding, in the direction of greater welfare and more distributive justice. But making that actually happen will require a combination of good theory, good practice, continued academic development, experimentation, and gradual iterative development.

More generally, the case of benefit-cost analysis for spending shows the importance of assessing distributive questions in institutional context, rather than in the abstract, since the right approach can differ in surprising ways. And it illustrates a general point at the foundation of economic analysis of policy: if one considers the particular institutional context and its relevant political economy, the governing legal regime, and the core economics of who gets what and what it costs, things are often going to seem less clean than in textbooks. And if we are to make progress on achieving a fair economy, we will often need to consider these all at once—and take significant steps to help out the neediest, both on welfarist grounds and on broader grounds of distributive justice.