Journal of Benefit-Cost Analysis (2024), 1–28
doi:10.1017/bca.2024.22

CAMBRIDGE UNIVERSITY PRESS

ARTICLE

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

Zachary Liscow[1] and Cass R. Sunstein[2]

[1]Yale Law School, Yale University, New Haven, CT, USA
[2]Harvard Law School, Harvard University, Cambridge, MA, USA
Corresponding authors: Zachary Liscow and Cass R. Sunstein; Emails: zachary.liscow@yale.edu; csunstei@law.harvard.edu

Keywords: administrative law; benefit–cost analysis; efficiency; law and economics; welfare

JEL classifications: D61; H50; H54; H23; Q52; R58; D63; K23

Abstract
Both Republican and Democratic administrations make regulatory and funding decisions with close reference to benefit–cost analysis (BCA). With respect to regulation, there has been a great deal of academic 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. Office of Management and Budget (OMB) Circular A-94 sets out guidelines for the BCA required when people are applying to many federal discretionary grant programs. Through Circular A-94, OMB has long required applicants to demonstrate that the benefits of their projects would exceed the costs. But under Circular A-94 as it stood for many years, efficiency-based BCA could produce results that fail to maximize welfare and that are also highly inequitable. The 2023 revision of Circular A-94 focuses more directly on welfare and equity, which are now – not uncontroversially – being brought directly into policy. At the same time, the new Circular A-94 raises fresh questions about how best to promote welfare, and to consider equity, in practice. This article explains the economic foundations for promoting welfare through distributional weighting – and how the old BCA guidance fell short. It then offers recommendations on how to operationalize distributional weighting on the ground specifically for government spending programs – and for BCA more broadly.

1. 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 costs.

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 costs.
4. Poor community, hard-to-quantify benefits. A poor community applies 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 (FEMA, 2023), but also the Department of Transportation, the Department of Agriculture, the Army Corps of Engineers, and other agencies across U.S. government.[1] To allocate its limited resources, the government is likely to 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 has the best available information about specific numbers?

Since 2003, federal *regulation* has been assessed under the framework established by Office of Management and Budget (OMB) Circular A-4, Regulatory Analysis (Office of Management and Budget, 2023a), which has been subject to extended academic discussion (Hemel, 2022; Kniesner & Viscusi, 2023). Since at least 1972 (OMB, 1972), *funding decisions* for discretionary grants 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 (OMB, 1992), which has received exceedingly little academic discussion (Liscow, 2024).[2]

Originally, Circular A-94 covered regulation as well, before the issuance in 1996 of new guidance on regulations (OMB, 1996). Circular A-94 is much shorter than Circular A-4, but it has comparable importance, covering $40–50 billion of spending annually (Mitchell, 2023). It establishes the criteria that federal agencies use when deciding whether and when to give out discretionary grants. In April 2023, the OMB proposed significant revisions of both circulars – not without controversy (Dudley & Viscusi, 2023; Harberger *et al.*, 2023) – with all former presidents of the Society of Benefit-Cost Analysis opposing some of the proposed changes, including those on equity.[3] Nonetheless, important changes were adopted. (Disclosure: Liscow led the revision to Circular A-94 as Chief Economist at the OMB in 2022–2023; Sunstein was involved in its revision as Senior Counselor to the Secretary of Homeland Security and later as Special Government Employee.) The new Circular A-94

[1] 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.

[2] 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. See Liscow (2024) for a discussion of other changes that Circular A-94 made and how they differ from those in the regulatory context.

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

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

and the new Circular A-4 were finalized in November 2023.[4] Circular A-94, not Circular A-4 (Sunstein, 2024), is our focus here, though much of what we say will bear on the new Circular A-4 as well. In that way, we offer a lens onto broader questions about how and whether to consider welfare and fair distribution in benefit–cost analysis (BCA).

In its old version from 1992, Circular A-94 announced that its goal was "to promote efficient resource allocation" (OMB, 1992). It emphasized the importance of formal BCA, which it described as the recommended "technique to use in a formal economic analysis of government programs or projects"[5] (OMB, 1992). It called for "comprehensive estimates of the expected benefits and costs," with close reference to "[t]he principle of *willingness to pay*,"[6] noting that "[m]arket prices provide an invaluable starting point" (OMB, 1992). Importantly, Circular A-94 also recognized that there might be "significant distributional effects," and asked agencies to analyze those effects by reference to income class, geographical region, or demographic group (OMB, 1992). We shall return to this direction in due course.

Under the old 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. The community might well be out of luck.

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, not in terms of monetary measures but 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 old Circular A-94 might well have pointed in the wrong direction (and the new Circular A-94 shows awareness of that point).

The fundamental and broader problem is that monetized BCA might not promote welfare, if people who gain (say, $X) are poor and if people who lose (more than $X) are rich. The reason is that a given amount of money is worth more to people who are poor than to people who are rich. This is a point about welfare, which is our focus – hence our title.

Second, the community in case 2 is poor, and purely on non-welfarist 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

[4] See Issuance of Revised OMB Circular No. A-94, 88 Fed. Reg. 77615 (November 13, 2023a,b); Issuance of Revised OMB Circular No. A-4, 88 Fed. Reg. 77615 (November 13, 2023a,b).

[5] Circular A-94 also allows consideration of cost-effectiveness analysis, but pointedly describes it as "a less comprehensive technique" (OMB, 1992).

[6] It is worth noting that the old A-94 referred 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) (Kahneman *et al.*, 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 revised Circular A-94 is alert to these points and does not opt for willingness to pay as opposed to willingness to accept.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

public resources (and the new Circular-94 shows awareness of that point as well).[7] The fundamental and broader problem is that apart from welfare effects, public officials might believe that if funds are being given out, poor people have more of a claim on them than do wealthy people, a conventional view in political philosophy (for relevant discussion, see Adler & Holtug, 2019).

For three reasons, it is especially important to probe Circular A-94, both in its old and new versions, and the BCA of spending. First, spending programs have of course always mattered, and they may be more important now than ever. For example, the US faces multiple climate-related risks, including wildfire, extreme heat, drought, and flooding. Many federal programs provide significant resources to increase resilience against those risks (FEMA, 2023). In addition, there have been large expansions in recent years in this grant-based funding requiring BCA.[8] 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 well seem preposterous.

Spending is a propitious place to focus on considering equity in BCA, mostly because the stakes are so high but also because for technical reasons, the legal regime makes legal challenges less likely for spending than for regulation (Cecot & Viscusi, 2015; Revesz & Yi, 2022).[9] The distributive challenges are also particularly stark and compelling, with a 2022 *Politico* headline reading: "How FEMA helps white and rich Americans escape floods" (Frank, 2022b). Importantly, and as we will see, the tendency to direct funds toward richer places is actually baked into FEMA's longstanding methods for its prospective resilience funding (thus prompting recent efforts at reform). To say the least, it is important to get funding decisions right.[10]

Second, funding is relevantly different from regulation, both analytically and in practice. Suppose, for example, that the 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. If, by contrast, the government subsidizes energy-efficiency refrigerators, no such effect should be expected. To be sure, there will be *some* systemic effects; taxpayers have to pay for the subsidies, and consumers may receive only some of the benefits if refrigerator manufacturers receive some of the benefits. Still, while the effects may be mitigated, it remains very likely that the intended beneficiaries will in fact benefit, making the distributional impacts of spending at least somewhat more straightforward than for regulation.[11] 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 (or regulating aggressively or more leniently). And unlike for regulation, parties – often poorly resourced local governments, not well-versed in the requirements of BCA – apply directly, creating distinctive challenges with a high degree of analytical complexity.

[7] "Prioritarians," who give priority to those at the bottom of the welfare ladder, would be strongly inclined in this direction (Adler, 2000; Adler & Holtug, 2019).

[8] For example, consider the Infrastructure Investment and Jobs Act of 2021 (Bipartisan Infrastructure Law), Pub. L. No. 117-58, 135 Stat. 429 (2021).

[9] In the funding context, issues of standing and reviewability may limit plaintiffs' ability to bring such challenges. Lujan v. Defenders of Wildlife, 504 U.S. 555 (1992); Allen v. Wright, 468 U.S. 737 (1984).

[10] More broadly, we are trying to give some substance to a very specific context within the broader debate on equity-informed legal rules (Kaplow & Shavell, 1994; Sanchirico, 2000; Liscow, 2014).

[11] For some complications, see Allcott and Greenstone (2017).

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

Third, spending provides an exceedingly useful lens on BCA more generally, including both its past and its future. For decades, Circular A-94 was admirably clear about what it was doing and why, which matters for both spending and regulation, since guidance for regulation grew out of it. The goal was maximizing efficiency.[12] The justification for doing so was in large part the assumption 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 for government practice, providing insights for the past, present, and future.

In Section 2, we describe the foundations of basing BCA on welfare, rather than efficiency, in the context of the history of Circular A-94. We then describe different kinds of objections to the old Circular A-94, which call for somewhat different reforms. For example, the potential disconnect between efficiency and welfare, on which we will be insisting here, will justify the use of distributional weights, designed to "upweight" benefits (as measured with the efficiency metric) 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 (Resources for the Future, 2023). In Section 3, we explore the new issues that agencies will need to confront as they undertake this work on the ground under the new Circular A-94, providing a series of recommendations meant to be applicable to spending programs and BCA generally.

2. Efficiency, Kaldor–Hicks, and Welfare

2.1. Efficiency in the old Circular A-94

The old Circular A-94 grounded 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.[13] By an "efficient" outcome, people sometimes mean one without available Pareto improvements, in which at least one person is made better off while no one is made worse off. But funding decisions do not involve Pareto improvements (too many potential parties are involved, and some will be made worse off), which means that for good reason, the old A-94 referred to *potential* Pareto improvements, not actual Pareto improvements. It said, with admirable and (to us) surprising clarity, that its goal of maximizing net benefits was "based on the premise that gainers could fully compensate the losers and still be better off" (OMB, 1992). This explicit foundation, forgotten in subsequent decades, provides an opportunity for assessing the neglected history.

The basic idea of efficiency in the old Circular A-94 was 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. 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

[12] As we discuss below, the term can be understood in different ways. We will generally understand the term to refer to Kaldor–Hicks efficiency (Coleman, 1980).

[13] This is of course a longstanding debate (Posner, 1979; Dworkin, 1980; Posner, 1985). Coleman (1980) offered an especially clear outline of alternative conceptions of efficiency.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

$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 be the replacement cost of the structures protected from flooding, since that is what the levee saves people from needing to pay.

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 the best of our knowledge has never been followed in the context of spending, the old A-94 even said that "[t]he presence or absence of such compensation should be indicated in the analysis," along with discussion of any "significant distributional effects" (OMB, 1992). 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 lose money, and we might at least ponder the possibility of compensating them (though regulators typically lack legal authority to do so). In the context of funding, the issue is much fuzzier, because taxpayers are unlikely to be reimbursed by recipients of their money.

The old Circular never said so explicitly, but it adopted the logic of Kaldor–Hicks efficiency (Coleman, 1980; Adler, 2000), which has long been subject to powerful and frequent academic objections (Adler, 2019), but which has also long been influential – with its influence continuing to the present day. A policy is Kaldor–Hicks efficient when it places resources in the hands of those who value them most, based on their willingness to pay for them. The criterion says that if the winners from the policy change can potentially compensate the losers and retain a surplus, then there has been a Kaldor–Hicks efficiency improvement (hence a Kaldor–Hicks improvement is sometimes described as a *potential* Pareto improvement). A policy is Kaldor–Hicks efficient when no more such improvements are available. This is when resources are in the hands of those who are willing to pay the most for them.

The problem is that, without the compensation, the change might be bad. Consider a case in which the losers, who are already miserable, become still more miserable, and in which those who are now flourishing flourish more. Suppose that the losers lose less in monetary terms than the winners gain in those terms. Is this a desirable change? Possibly not. One reason involves distribution. A Kaldor–Hicks improvement might increase distributional unfairness. Another reason involves welfare: Even if there is a net efficiency gain, welfare may not increase. If rich people gain $100 million and poor people lose $50 million, is there a gain in welfare? Not necessarily. A loss of $50 million to poor people might produce a decrease in welfare that it is higher than the increase in welfare that comes from a $100 million gain by rich people. To be sure, there is a welfare gain if the rich end up compensating the poor for their loss; but what if they do not? If rich people gain $100 million and taxpayers pay $50 million, if there is a gain in welfare? 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, 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 disproportionally benefit from a project, might a project with a benefit–cost ratio (BCR) a

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

little below 1, measured on efficiency grounds, become cost-justified (because potentially welfare-justified)? Suppose that communities apply for funding for a project that would cost $500 million and deliver $550 million in benefits. If the rich disproportionally benefit, would a project with a BCR a little above 1, measured on efficiency grounds, no longer be cost-justified (because potentially not welfare-justified)? None of this was ever specified. And since again to the best of our knowledge, the exhortation to do the distributional analysis (which would require more work – and sometimes not easy work[14]) 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 FEMA grant because the BCR is below 1: the community is not willing to pay for the cost of the protection. The rich communities, in contrast, do get funding.

2.2. *Welfare in the new Circular A-94*

The new Circular A-94 offers a different account of the goal of BCA.[15] The new Circular states that its goal "is to promote social welfare" (OMB, 2023c), reflecting the emphasis in economics on welfare or well-being[16] as the ultimate social goal (more so, that is, than efficiency) and the questionable empirical plausibility of that idea that winners will compensate losers (Calabresi, 1991; Adler & Posner, 2006; Adler, 2011; Fennell & McAdams, 2014; Liscow, 2018b). Efficiency is removed as the ultimate goal for BCA in the new version, though (as it notes) of course all else equal greater efficiency is still better.

The new Circular nevertheless continues to envision a rigorous quantitative analysis. It does not by any means abandon BCA. The new Circular seeks to mend it, not to end it. As with the old Circular, so with the new version: If a project would cost $50 million and produce benefits of $500 million (even calculated under the old Circular's methods) agencies should almost certainly fund it. The new version 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 new version, 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 (measured in efficiency terms) are upweighted and higher-income net benefits are downweighted. FEMA has been working on a pilot to incorporate this method (Resources for the Future, 2023).

The new Circular 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

[14] See the discussion of the challenges in measuring ultimate incidence below.

[15] In that sense, it is consistent with the general approach by Adler and Posner (2006).

[16] We use this term interchangeably with "welfare" or "social welfare," which we also use interchangeably, though "social welfare" can be taken to be more explicit about allowing for tradeoffs among the welfare of different individuals.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

negative of the elasticity of marginal utility (ε).[17] Note that this assumes a constant elasticity of marginal utility. Consistent with the new Circular A-4 (OMB, 2023a; OMB, 2023b), the guidance recommends a default elasticity of marginal utility of 1.4, based on estimates of within-person risk-aversion, hedonic studies on happiness, and other methods (OMB, 2023b), mostly included in a meta-analysis of 1,636 studies in this journal (Acland & Greenberg, 2023). We will shortly (under the "Declining marginal utility of income" subsection) explain more what it means to use these forms of evidence. For now note that 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.[18] 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.[19] This is problematic because not only are poorer communities less likely to get funding (which is arguably unfair on purely distributive grounds, welfare to one side) 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, and those who have most to gain, in welfare terms, will lose out.

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). Assume too that all benefits are incident on the homeowners; that is, the homeowners will enjoy those benefits.[20] Assume finally that utility can be modeled as the logarithm of income, as noted a conventional assumption in economics (partly 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.[21] For a logarithmic utility function, the weight is conveniently equal to one divided by income.

[17] The UK Green Book uses a broadly similar method (HM Treasury, 2022). There are some small international differences in the weights (the UK uses an elasticity of marginal utility of 1.3, whereas the US is using 1.4), which remain to be explained.

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

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

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

[21] 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 ($150 k) – log ($150 k – $21 k)] = 33. For Project B, this is equal to: 10 years * 100 households * [log ($50 k) – log ($50 k – $9 k)] = 86. So, the welfare gain is much larger for Project B.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

Table 1. *Example of distributional weighting*

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

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. The weight for Project B is (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.

The new version of Circular A-94 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 (OMB, 2023c). It is important to see that this is arguably what is already done for the value of a statistical life (VSL; Kniesner & Viscusi, 2023). Agencies use a number – now in the vicinity of $12 million (U.S. Department of Transportation, 2021) – 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 far less to reduce statistical risks than are wealthy people (or if poor people demand far less to face statistical risks than do wealthy people) (Sunstein, 2004). 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.[22]

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 might be given an average value.[23] There are downsides to this approach: in particular, it ignores variations in home prices, since sometimes lower-income communities have some or many higher-valued homes (which are especially valuable to protect), leading to a worse targeting of resources from a welfare perspective. But it is at least facially simpler to implement and seems to have a certain ethical appeal.

Return in this light to the stylized cases with which we began and consider them in light of the new version of Circular A-94. In cases 2 and 3, involving poor communities, it is

[22] In the context of regulation, the incidence question raises many complications; it is not clear that upweighting the monetary value of 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 (Hemel, 2022).

[23] As the new rules note, weights should not be applied to measures, like the value of a 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 the monetary value of a statistical life, though there are ethical and economic questions about doing so (Hemel, 2022). In any case, though, that weighting should consider the averaging that has already taken place.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

plausible that under the new version, 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 case 2 *poor community, wildfire risks*, and in case 3 *poor community, extreme heat*, at least, it is perfectly imaginable that the adjusted numbers would allow for funding.

The impact of the new Circular should be to increase welfare and to promote equity. Out of the fixed funds that agencies have, the old Circular wrongly in effect forbade money from being directed to lower-income areas in cases like 2 and 3.[24] It is nonetheless true that in terms of direct effects (i.e., ignoring possibly beneficial effects of greater equity on the broader economy), the new version 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 provide modest benefits (in welfare terms) in low-income communities but that cost a great deal. The new version does not authorize such a project, and as long as agencies do not go too far in considering equity, what it does authorize will increase welfare.

2.3. Four distinct justifications

This shift from efficiency to welfare is justified on four distinct but interconnected grounds. Each of them would justify a shift from the old 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 to increase welfare.

1. The declining marginal utility of income: The first reason is the declining marginal utility of income, a justification explicitly mentioned in the new Circular. Essentially, the argument goes, welfare increases more when $1,000 goes to poor people than when the same amount goes to rich people. If each of the 10 poorest people in the US is given $1,000, they would gain far more, in terms of welfare, than each of 10 richest people in the US would lose if they lost $1,000. 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 new BCA guidance builds on average behavior. The assessment in Circular A-94 comes from several sources, of which we discuss two main ones here. First is measuring risk aversion within individuals, by asking how much less do individuals value a dollar as they

[24] Technically, the old Circular A-94 did allow for "a policy [that] is intended to benefit a specified subgroup of the population," and thus would presumably be inefficient (OMB, 1992). We are unaware of this ever having been followed.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

get richer, based on how much they are willing to pay for things such as insurance, which effectively transfers money from times when they are richer to times when they are poorer (Chetty, 2006). The second is hedonics: asking how much happier do people say they are in surveys as they get richer (Deaton, 2008; Stevenson & Wolfers, 2008; Gandelman & Murillo, 2015; Killingsworth *et al.*, 2023). Averaging studies of these measures, along with others, yields the weights suggested in the guidance, though consideration of the same general factors could of course end up producing different specific weights.

Granted, some economists object to interpersonal comparisons of utility.[25] We do not believe that the objection is convincing, but consider a few responses. It is true that an especially greedy rich person, who is made much happier by every increment in consumption, will really enjoy an extra $1,000, and that an abstemious poor person, who cares little or not at all about consumption, could conceivably get little or no utility out of that same amount.[26] There is nothing about the economic situation of two individuals to reveal, as a matter of mathematics or logic, that one's utility from a given gain will higher than the other's. Behavior itself might tell us less than we need to know. But this does not defeat the usefulness of the idea of the declining marginal utility of income. On average, it makes sense to believe that a gain of a given amount of money to those who are poor produces a greater increase in welfare than an equivalent gain to the rich (consistently with the empirical analysis in Matsumori *et al.* (2021) and Killingsworth *et al.* (2023)). As a matter of government policy, it strikes us as reasonable and right – as a matter of administrability and ethics – to ignore potentially differing preferences at the individual level and instead focus on average utility (or welfare) at a given income level (Samuelson, 1947; Bergson, 1954). That itself is not totally straightforward, but rather must be approximated, beginning from the eminently reasonable judgment that on average $1,000 brings more social value to a starving person than to a multibillionaire.

Estimating the decline of marginal utility using measures like risk aversion and hedonics does not produce definitive numbers; the empirical and conceptual issues are challenging. But if one agrees that achieving the most welfare as possible in society is a desirable goal, it makes sense to use the best available evidence to achieve that goal. It is certainly better than doing something that does not even try to achieve that goal – and for which there are good reasons to think that it will not come close to achieving it, as the examples above illustrate.

Government is not a debating club. Money is going out the door. Interpersonal judgments are inherent in making those funding decisions. Officials should make those decisions with the help of the best available evidence about how to increase welfare. And we emphasize that the evidence will improve over time.

2. Correcting "rich-bias": The second reason is bias-correcting (Dworkin, 1980; Liscow, 2018b).[27] As one of us mathematically defined it elsewhere, a policy is "rich-biased" if, "as one gets richer, one tends to get more legal entitlements from efficient policies," such as getting more funding for protecting homes from discretionary grant programs (Liscow, 2018b). Efficient policies tend to be rich-biased because richer people tend to have a higher willingness to pay to receive a piece of infrastructure, which is what drives efficiency analysis. For example, one of the main benefits of transportation infrastructure is saving

[25] For various perspectives, see Interpersonal Comparisons of Well-Being (1991).

[26] This is the notion of the "utility monster" (Nozick, 1974).

[27] 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 (Giraldo & Liscow, 2024).

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

time. Richer people have higher earnings and so are willing to pay more for time savings; the Department of Transportation has a higher value of time for projects benefitting higher income-earners, which will tend to direct funds toward the rich (Liscow, 2018b).

Similarly, as the examples 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 cheaper houses. Under a purely efficiency-based framework, this tendency will direct spending to richer places. Relative to an egalitarian goal of having an equal shot at funding across the income distribution (or perhaps a prioritarian goal of spending to help the least well-off), the efficiency-based framework reflects a bias.

This reasoning on bias correction is mathematically tied to the declining marginal utility of income, but it is conceptually distinct (Liscow, 2018b). In particular, the declining marginal utility of income means that "[d]ollars are 'cheap' to the rich because they already have so many of them," leading them to spend more on things like housing and more generally be willing to pay a lot (in dollar terms) for amenities that the government provides – and thus for an efficiency-minded government to direct resources to them (Liscow, 2018b).

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 (Adler & Posner, 2006). But the logic that efficient BCA can maximize well-being depends upon some empirical foundations that turn out to be flawed.

3. Failure to compensate losers: Arguably the old guidance did not say what its *ultimate* goal was, since it was 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.)[28] This Paretian logic is appealing for those who emphasize the challenges of measuring social welfare. But as we have seen, the winners do not compensate the losers when funds are awarded. This is the third reason for the switch: Because compensation does not occur, eliminating the underlying Paretian logic, it makes sense to shift the focus from efficiency to welfare. The idea of the old Circular was that, with compensation to losers, 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 (Raskolnikov, 2021). For example, famously, after the US allowed China to 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 (Autor *et al.*, 2013). And after state supreme courts direct legislatures to spend more on poor school districts, those who lose out receive no compensation; the consequences stick (Liscow, 2018a). The broader point is that when some agency gives out

[28] In addition, as Kaplow and Shavell (1994) and Kaplow (2020) have argued on welfarist grounds, under certain assumptions, it is better to adopt the efficient non-tax policies (like BCA for spending) and redistribute through taxes than to redistribute through non-tax rules.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

grant money – for example, to reduce the risks associated with flooding – no one compensates the losers.[29]

It is not even 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 the best of our knowledge, there was never an assessment of compensation in the spending context. It is a good thing to be discarding this logic.

Hicks similarly made the argument, called the "classical creed," that if society made "all alterations" meeting the Kaldor–Hicks criterion, then (even without compensation) "there would be a strong probability that almost all (individuals) … would be better off after the lapse of a sufficient length of time" (Hicks, 1941). Insofar as the criterion is understood in terms of monetary measures, Hicks was too optimistic. His speculation seems most unlikely given the "bias" toward funding the rich with efficient programs. This is even more so given recent theoretical work showing that initial inequality can lead to increasingly "snowballing" inequality over time. Efficient rules create "vicious cycles" in which a program is biased toward the rich; they get even richer; their willingness to pay goes up; and then in the future, the program is even more biased (Giraldo and Liscow, 2024).

4. Failure to redistribute sufficiently through the tax code: 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. This argument is distinct from the third argument because, while the third argument is about compensatory transfers for *each project*, this argument is about the *baseline level* of redistribution. Both arguments are ultimately driven by the benefits of cash redistribution. This fourth argument is not explicitly stated in the old guidance, but it has long been prominent, dating back in public finance economics to Richard Musgrave's 1959 textbook (Musgrave, 1959)[30] that separated the redistributive from the efficiency-achieving branches of government. It is fair to say that the argument is implicit in the old guidance.

The essential argument is that it maximizes welfare to redistribute through taxation, rather than spending or regulation. If the goal is to redistribute income from rich to poor, it is best to require the rich to transfer money to the poor, rather than (for example) to impose an occupational safety regulation that gives more safety to mostly poor people and that increases the cost of goods that are purchased mostly by rich people. Cash (via taxation) lets the recipients choose how to spend, allowing them to spend it on the things most valuable

[29] Of course, sometimes there are policy changes that offset distributional impacts. For example, Baicker and Gordon (2006) and Boylan and Mocan (2014) both find that, after court orders benefitting lower-income people, there are subsequent reductions in welfare expenditures. Our claim is that such compensatory changes are often, even typically, not present.

[30] After this, luminaries such as James Tobin and Richard Posner discussed this argument.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

to them. Everyone values a dollar at a dollar, while higher-income people typically value things, such as goods 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 (in monetary terms) and provides them to those who value them less (also in monetary terms). (And to be clear, "cash" means money itself, not things funded by cash.)

Consider the numerical levee example above. When member of the poorer community get 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. One more time: Cash 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 maximize social welfare. The strategy is for spending (like levee 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 – unless there are actually compensatory taxes and transfers for each spending project – it requires an *optimal tax code* that redistributes the socially optimal amount. Otherwise, other policies – like BCA – should seek to move society closer to that optimal distribution. As noted, those compensatory taxes and transfers are unlikely: It is not plausible that the poor community of 100 families will receive the $9.5 million in new cash transfers. FEMA cannot transfer it. And there is little evidence that other institutions compensate for distributional impacts project-by-project (and, of course, if it did, then FEMA could alter its benefit-cost procedures). So that leaves the system of taxes and transfers to promote distributive fairness. Is it plausible to think that now or in the future the tax code will do all the necessary work to achieve an optimal distribution?

As Liscow (2022) argues, there is good reason to think that we should not expect to have an optimal tax code, in the sense of maximizing social welfare on its own. Notably, there is strong public support for greater economic equality, especially through providing necessities like health care (Liscow, 2022) 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 taxation, which can appear (i) confiscatory with respect to the rich's income and (ii) undeserved with respect to transfers toward the poor who do not work for it (Sheffrin, 2013; Liscow, 2022). For example, the public is much more comfortable transferring an account that can be spent only on necessities than an equivalent amount of cash that can be spent on anything (Liscow & Pershing, 2022). 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 suggest (Hendren, 2020; Liscow, 2022). The public simply does not treat public policy as if there is just "one pie" from which to redistribute most efficiently. Instead the public compartmentalizes their views into different policy-by-policy silos. And expecting taxation to pick up all the redistributive slack would mean asking about an area of policy where the public is *particularly averse* to redistribution, setting up a failure to achieve a just distribution of resources (Liscow, 2022).

There is just no reason to think that our political system operates in a way that treats different areas of policy fungibly, finding the one that redistributes most efficiently. So, in

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

the end, 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 new Circular is to authorize agencies to achieve that goal, and thus to allow funding in imaginable variations on the second and third 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 could be a sensible approach. For example, FEMA could use average valuations of things like housing – though with the caveat discussed above that this leads to some inefficiency because of heterogeneity in housing within communities of a given income. A way to correct for the bias while taking heterogeneity into account would be to adjust for the change in willingness to pay for a given asset as income increases – which would just be a specific income weight. And, indeed, the new Circular reflects this reasoning. It notes that "[w]eights can also be based on the tendency of market-based methods to value things that low-income people use at low amounts" (OMB, 2023c).

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 at least certain kinds of complete compensation, there should be no weighting. For example, in our example above, if as a result of choosing the project in the rich community, the government transfers $9.1 million from the rich to the poor, then both groups are better off than under the weighted scheme. The rich get a levee that they value at $10.5 million minus a tax of $9.1 million, and the poor get $9.1 million, which is more than they would have valued the levee ($9 million). 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 new Circular allows agencies to do precisely that. If there were intermediate amounts of compensation, there could be intermediate weights.

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 new Circular, would increase welfare while also promoting distributional goals in a way that should be broadly agreeable. And, as we discuss in the next part, the system of distributional weighting adjusts itself as the tax code (and other policies) redistribute.

Overall, then, the 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.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

3. From theory to practice: Pursuing welfare with benefit–cost analysis

Turning from theory to practice, we now attempt to provide guidance for agencies on the ground, as they seek to maximize welfare and consider equity in their spending programs and benefit-cost analyses more generally, now that they are pursuing those goals in earnest. Importantly, some of these issues are specific to spending, and the analysis would be different in the regulatory context (Sunstein, 2024). Experimenting with implementing income weighting now is not a huge challenge; agencies must choose weights (likely the ones suggested by OMB), know the income of affected communities (which is easily available in the Census), and apply the weights to values that are not already implicitly weighted (recall that the VSL is implicitly weighted). But perfecting the methods will take more conceptual and practical work, both by academics and by those in government – not necessarily much more work when implemented (which is largely plugging values into existing formulas), but more up-front development work.

3.1. 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 the old Circular A-94 told agencies to analyze it. The federal government just never did so for spending (to the best of our knowledge) – and did not do so very often for regulation.[31] 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 (Dudley & Viscusi, 2023; Harberger *et al.*, 2023). On this view, CBA should be undertaken without regard to distribution, but – before they make final spending decisions – 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 (Frank, 2022a; Frank, 2022b). The OMB Director has even apologized for the current guidance in front of Congress for not giving a fair shake to lower-income communities (Alex Padilla, U.S. Senator for California: Press Releases, 2023). 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 both its old and revised forms, sets BCA up with the goal of being *quantitative* (OMB, 1992, 2023c).[32] So, when a feature that matters, such as distribution, can be quantified, it stands to reason that it is most helpful to quantify that feature. Doing so helps make better decisions by being more rigorous.

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

[32] The 2023 Circular A-94 notes: "Quantifying costs and benefits to the extent possible is expected, even if assigning monetary values is not feasible" (OMB, 2023c). The 1992 Circular A-94 defines "benefit–cost analysis" as a "systematic quantitative method of assessing the desirability of government projects or policies when it is important to take a long view of future effects and a broad view of possible side-effects" (OMB, 1992).

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

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 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 potentially arbitrary.[33]

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. Including distributional weights as part of that metric accomplishes that goal. The remarkable dearth of cases in which distribution was considered in spending decisions suggests that the previous regime did not.

3.2. Optional use of distributional weights under OMB guidance

Another notable feature of the 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 the best of 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 might well 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" – as the U.S. government becomes accultured to the methods.

At the same time, it is wise for OMB to issue guidance on the subject of weighting for several reasons. First, OMB has unusual technical expertise. There are economies of scale in reading the literature and understanding these issues. OMB is well-positioned to do this for the federal government, especially as agencies are often under-staffed and focused on implementation. Second, it is valuable for the purposes of promoting (if not necessarily dictating) uniformity across agencies as appropriate, which helps to avoid inconsistency and potential unfairness. It remains true that practices could legitimately change across administrations. There are normative issues here, and different administrations might care differently about distributive issues.

[33] 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.)

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

3.3. 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.[34]

To be sure, Congress could dictate the weights. But if it did so, it would have to engage the relevant technical issues, and it seems unlikely that it would have an incentive to do that, or even to specify some direction on whether or how to use distributional weights. Existing statutes governing discretionary grant programs subject to BCA for spending do not speak to the underlying questions. They certainly do not require "efficient" spending, in the sense of spending more in rich areas than in poor areas where efficiency so suggests.[35] In the absence of guidance from Congress, the executive branch has significant latitude.[36]

It is clear 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 or $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 (NBER, 2023). In-kind transfers such as 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. This is not only because estimated costs 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 (Finkelstein *et al.*, 2019).

Ignoring this adjustment would lead to excessive redistribution, with equity gains not justified by efficiency losses, taking as fixed the welfare goal implied by a given set of distributional weights. 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 BCA 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.

[34] A separate question is whether to weight by aspects other than income. Income has the advantage of being relatively well-measured and doing a decent job of capturing many aspects of disadvantage. But it surely does not capture everything. So other aspects could be considered. We do note, however, that race-based weighting would likely face serious constitutional challenges. *See* Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll., 600 U.S. 181 (2023).

[35] It is notable that the OMB director was pushed by members of Congress to be more equitable – and that Politico finds it a scandal that FEMA money is spent so disproportionately in rich places (Alex Padilla, U.S. Senator for California: Press Releases, 2023; Frank, 2023b).

[36] Of course, executive branch policymakers may be wary of political objections that could undermine political support for the programs – and will want to be concerned with political legitimacy more broadly. But there is just very little reason to think that equity weighting lacks political legitimacy. As we have said, weights of one on an efficiency metric that is tilted toward funding the rich is still a weighting scheme.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

3.4. Measuring income of affected communities

Measuring the income of affected parties might seem relatively straightforward, but it requires making some conceptual choices. Importantly, communities, rather than individuals, are typically affected by grant funding, which eases the measurement burden (as well as making 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 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 (U.S. Census Bureau, 2022)? We would suggest using Census tracts, which are used for Opportunity Zone tax credit designations, for example (IRS, 2022). 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 BCA. Agencies have long needed to know what the geographic scope of impacts, in order to know the magnitude of benefits. The new rules merely suggest the possibility of considering the relative wealth and poverty of whoever is within that scope.

3.5. Incidence of spending

As the old Circular A-94 noted, 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 at the pre-regulation price; they have to *pay* for the cars as they are priced post-regulation. 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. A reasonable (though imperfect) assumption is 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 – and higher 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 in an amount equivalent to a portion of the value of the levee. It is not totally clear ex ante what the

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

impact of ignoring incidence would be. But our sense is that it will tend to be the case that, through price changes, some rich people benefit in poor places, and some poor people benefit in rich places (e.g., through employment), in ways that ignoring economic incidence would cause to be missed. So it is plausible that ignoring economic incidence would lead to excessive redistribution, with equity gains not justified by efficiency losses.

There are rules of thumb that agencies can follow; for example, the more elastic the housing supply, the more renters will capture the benefits.[37] In any case, because distribution has not been incorporated much if at all into BCA 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 less well – in extreme cases, agencies might spend funds with much less in the way of benefit because they neither promote efficiency nor actually reach those in need.[38]

3.6. Difficult-to-quantify impacts

The old Circular A-94 largely assumed that the impacts are known. But what if important aspects are difficult to quantify, as case 4 in the introduction suggested? For example, the disruption to livelihoods from natural disasters, especially for the poor, may be hard to quantify. The new Circular A-94 continues to put a premium on quantification, but it notes that "[s]ome important benefits and costs may be either difficult to monetize or difficult to quantify" (OMB, 2023c). In such cases, it remains "important that these be included the analysis" (OMB, 2023c). In some cases, the non-monetized benefits or costs might be significant, and "a threshold or break-even analysis may be considered for inclusion in the benefit–cost analysis" (OMB, 2023c). This general point might matter for equity, because poor communities might find it especially challenging to produce a comprehensive BCA for grant applications.

In such cases, the new 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" (OMB, 2023c). Where data are sparse, breakeven analysis may be the best that agencies can do (Sunstein, 2014). It should be a useful tool for grant applications, and it might make a significant difference for poor communities (Sunstein, 2014; Revesz & Yi, 2022; Sunstein, 2024).

By arguing that distributional weights should be used, we do not mean to suggest that BCA can quantify everything. It would be absurd to suggest that things that are difficult to quantify should not be considered. Quantification – including of distributional impacts – is valuable, but incorporating everything is beyond our current capabilities.[39]

[37] 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.

[38] Another aspect of incidence worth further exploration is the extent to which grants to localities displace state and local spending on the things that the federal government is funding. It might be a reasonable guess that displacing local funding may have incidence not too far off from where there is no displacement. Displacing state funds is a harder question.

[39] And, indeed, some questions of value remain outside the basic skillset of technocrats, suggesting that direct systematic citizen input may be useful. For example, the National Institute for Health and Care Excellence in the UK has done this (Liscow & Markovits, 2022). Liscow and Markovits (2022) explore how "deliberative polling" could be incorporated into rulemaking in the case of policy designed to correct bounded rationality.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

3.7. Net benefits vs. benefit–cost ratio

For regulation, it makes sense to focus on net benefits, not benefit-cost ratios (BCRs). 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 regulation with benefits of $500 million and costs of $250 million, with a BCR of 2 to 1 (OMB, 2003). The analysis is not quite the same for individual spending projects, for which the ratio is most relevant. Total net benefits are always what ultimately matter, and the net benefits rule must indeed be applied to the total package of projects, but there are wrinkles in the context of spending for arriving at that total package. With a fixed budget of (say) $250 million, choosing individual projects with the highest BCRs will maximize net benefits and thus welfare. Spending all $250 million on one project with a BCR of 2 to 1 would not be as good as spending funds on smaller projects, each with higher BCRs 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. Choosing Projects B and C ultimately produces higher net benefits ($150 million) than choosing Project A ($100 million). So it is better to choose individual projects (which is what agencies do) by their BCR, not their net benefits.

Under current government practice for programs subject to BCA, there are two relevant differences between regulation and spending. The first is that BCA applies to each individual project 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.[40] 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.[41] In contrast, for spending, agencies' need to choose individual projects subject to a budget constraint drives the need to focus on BCR.

[40] We are supposing that no regulatory budget is in place. The Trump Administration did impose a kind of regulatory budget of various kinds, with a rule of "one in, two out" and with a mandated annual regulatory cost of no more than $0. Exec. Order 13,771, 82 Fed. Reg. 9339 (January 30, 2017). (We bracket some complications and qualifications.) If the constraint is 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. It is also true, of course, that statutes may impose restrictions on the imposition of costs, as through feasibility constraints or requirements of BCA.

[41] There can be minor exceptions, for example, in the case of project "clumpiness," in which one big high-BCR project takes up too many resources to allow for a more sizable, but somewhat lower-BCR project from being funded. Imagine that an agency has $100 to spend and has three projects available to it: Project A has a BCR of 3 (and cost of $60), Project B has a BCR of 2.5 (and cost of $100), and Project C has a BCR of 1 (and cost of $40). The agency should choose Project B, with net benefits of $150. It should not do Project A (net benefits of $120) and C (net benefits of $0), even though Project A's BCR is higher than Project B's. This shows how maximizing net benefits is still the ultimate goal for any pool of funds, even as – when agencies consider large numbers of projects – the BCR is typically (but not always) the best ranking rule in practice.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

One more time: 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 BCRs.

Finally, some might wonder whether certain costs should go in the numerator or the denominator.[42] For example, a project might benefit most people, but reduce the livelihoods of some people. Do the reductions of livelihoods go into the numerator or the denominator? There is actually a simple rule to follow: The financial cost borne by the government goes into the denominator, and everything else goes into the numerator.[43] Again, this is a difference from regulation, which raises some challenges in deciding whether to put certain quantities in the numerator or denominator.[44] In contrast, there is a single question that we want to ask for spending: for a $1 net outlay of government funds, what are the net benefits? The answer can then be compared on equal footing to the corresponding answer for other projects, producing the most beneficial use of government funding.[45]

3.8. Addressing taxes in the benefit–cost ratio

Spending decisions can affect tax revenue, raising the question of how that tax revenue should be incorporated into a weighted BCR. 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. This extra piece of information raises the information burden, but ideally it would be incorporated, especially since taxes can be substantial. Second, changes in tax revenue that result from spending should be included

[42] It may not be immediately obvious why it matters whether something (such as tax revenue resulting from a project) is included in the numerator or subtracted from the denominator. To see how it matters, consider a choice between two projects: Project A has $220 of nontaxed benefits and $100 of costs to FEMA. Project B has only $200 of benefits (e.g., increased earnings) taxed at 20% and $100 of costs to FEMA. 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.

[43] This is the "marginal value of public funds" approach, in which tax revenue is subtracted from the benefits (numerator) and the costs (denominator) (Hendren & Sprung-Keyser, 2020; Kaplow, 2020).

[44] This arbitrariness has led some to be skeptical of the benefit–cost ratio. For example, Harvey Rosen writes, "As a basis for comparing admissible projects, however, the benefit–cost ratio is virtually useless," partly on these grounds (Rosen & Gayer, 2014). As noted though, this critique does not apply to spending, where a clear decision rule is available: the question is how much benefit the government gets from expending (all-told) per dollar.

[45] Subtracting tax revenue from the denominator does implicitly treat the government as if it is fungibly directing resources to places with value similar to that of the project that produces the tax revenue, since – by subtracting tax revenue from the denominator – the approach treats tax revenue the same way as the grant funding. This strikes us as a plausible baseline assumption, even though of course tax revenue that results from FEMA grant programs is not directed to FEMA. If one has a strong view that government revenue would be not used very effectively, this baseline assumption could be problematic. For example, considering the two projects in footnote 42, if the tax revenue were be given lump sum back to everyone in the population (which is equivalent to keeping those funds in the numerator with a weight of 1), then Project A would be more desirable because of that project's higher net benefits. More generally, if those funds will be spent somewhere with known benefits (or offset distortionary taxes), then those benefits (per dollar) are the weights on the tax revenue in the numerator. However, we cannot think of real-world cases in which we would know how marginal government revenue would be spent, leading us to believe that subtraction from the numerator makes sense, so that agencies can ask how much benefit arises from $1 of net government outlays.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

in the denominator, for the reasons just discussed: the goal of the BCR is see how much net benefit arises per dollar of net government expenditures.

There are, in turn, conceptually two ways that government spending can affect government revenues. First, often benefits are taxed directly. 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. Second, there can be a more subtle behavioral responses to spending. For example, if the government funds an expansion of LaGuardia Airport, people might earn more money knowing that they can benefit from easier access to the flights that they would want to take if they had a higher-paying job. Or, if $1 spent on low-income housing is likely to have more funding protecting it than $1 spent on high-income housing, then people might earn a little less in anticipation of spending less on housing. With distributional weighting, it will tend to be the case that the types of things that benefit lower-income people are funded, relative to an "efficient" baseline, creating an incentive to earn less, thus lowering tax revenue. Ideally, both of these types of effects on tax revenue would be incorporated into the benefit-cost formula.

The first way that government spending affects government revenue is easier to handle. If one knows the average income of those affected, it is easy to calculate the taxes, as noted earlier, with the TAXSIM program online.[46]

The second is trickier, since we do not know the behavioral impact on earnings that results from expanding LaGuardia or from, say, valuing all housing equally for the purpose of FEMA resilience funds. There is a resemblance between the impact of distributional weights and taxation: if one gets richer, he must pay more in taxes, in the same way that, if one gets richer and moves into a bigger home, the federal government is less likely to spend as much protecting a dollar of home value. The behavioral impact of taxation is extremely well-studied. However, even those impacts are still hotly debated. And we are unaware of any estimates of the behavioral impact of redistributive elements in grant funding. There are good reasons to think that people would be more responsive to taxation than to distributional weights, especially since income taxes are very salient,[47] whereas distributional weights – and more generally how government spending impacts how much people wish to earn – is not salient and may be hard for individuals to even mentally process.[48]

The direction of the effect of this second refinement – accounting for revenue impacts from behavioral changes – is clear: redistribution through BCA will tend to reduce revenue, lowering the benefit/cost ratio of projects using distributional weights. But the magnitude of the effect is not clear at all. We urge more research in this area. But if all other aspects of BCA are done perfectly, the absence of this adjustment will tend to move spending toward being excessively redistributive for the purpose of maximizing welfare with a given set of distributional weights.

Finally, note again the connection to the fourth justification for considering equity in BCA (that taxes do not redistribute enough). As more redistribution happens through taxation (or other policies), the distortionary cost of additional redistribution through

[46] Using average income to calculate taxes is imperfect because of the progressivity of the tax code, but it strikes us as close enough for practical purposes.

[47] People are considerably less responsive even to less salient taxes, such as sales taxes, which are not included in the price of goods (Chetty *et al.*, 2009).

[48] Daniel Hemel (2023) suggests directly treating distributing weights like taxation. We admire his discussion, but are cautious about recommending that approach because, as noted in-text, we do not know empirically how grant funding affects income-earning via behavioral responses.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

spending should increase, at least in principle. As a result, the negative revenue impacts of redistribution through spending will increase as well. In fact, they will increase up to the point at which their impact on the BCR exactly negates the impact of the distributional weights: that is the point at which the background level of inequality reflects what the distributional weights suggest it should be. And then there is no need to redistribute through BCA anymore. As we argued earlier, though, there is good reason to think that the US is far from that point now. But it is good to know that the method appropriately adjusts on its own as inequality changes.

3.9. Administrative burden

For regulations and spending, there is an important difference in who bears the administrative burden of conducting BCA. 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 more[49]; President Biden changed that number to $200 million.[50]

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 BCA itself.* It is worth asking: Does BCA survive BCA?

This question suggests another dimension of equity in BCA: ensuring administrative burdens, or "sludge" (Sunstein, 2021), 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 U.S. Government Accountability Office (2021). 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 communities from applying for funds because of the additional requirements, especially among the most disadvantaged. This suggests once more that the analysis should be done by the agency itself, which presumably can also benefit from familiarity with the relevant issues and economies of scale.

4. Conclusion

If people receive a federal subsidy, they are highly likely to be better off. At the same time, the old 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 old Circular did not consist in its recognition of the importance of BCA. The problem was that it made economic efficiency, rather than welfare, foundational. The old Circular failed to note

[49] Exec. Order 12,866, 58 Fed. Reg. 190 (October 4, 1993).
[50] Exec. Order of April 6, 2023, Modernizing Regulatory Review (April 6, 2023).

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

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 BCA, the new Circular A-94 directs agencies to focus on welfare, not efficiency. Among other things, and for the first time, it explicitly 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 BCA will promote rather than frustrate its fundamental goal, which is to increase welfare.

The old 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, including specification of distributional weights, measurement questions, and administrative burdens. We have offered some preliminary thoughts on how best to approach these challenges. For example, we identified several ideal, but challenging, refinements – measuring incidence, shrinking the weights for the receipt of taxes and in-kind transfers, and considering the distortion to earnings from redistributing through spending – that, if not implemented, would probably leave BCA biased toward excessive redistribution. There may be offsetting errors in the other direction – including, very plausibly, some hard-to-measure effects such as disruption to livelihood from natural disasters, which may particularly harm the poor. Effects of this kind may or may not compensate for the biases the other way.

In our view, there is little doubt that the changes to Circular A-94 are moving and will move national policy with respect to funding in the direction of greater welfare and more distributive justice. Agencies are already beginning to implement the new methods. They have adequate analytical capabilities, and they appear to believe, reasonably enough, that they can best learn about them through experimenting with real-world use – especially if there is vigorous *ex post* evaluation of impacts on equity and efficiency. But implementing the changes optimally will require over time a combination of good theory, good practice, continued academic work, experimentation, and gradual iterative development.

More generally, the case of BCA for spending shows the importance of assessing distributive questions in various institutional contexts, rather than in the abstract, since the right approach can differ in different contexts in surprising ways. And it illustrates a general point at the foundation of economic analysis of policy: if one considers (1) the particular institutional context and its relevant political economy, (2) the governing legal regime, and (3) the core economics of who gets what, things are often going to seem far less clean than in textbooks. And if we are to make progress in increasing social welfare and achieving a fair society and 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 independent grounds of distributive justice.

Acknowledgments. Thanks to Raj Bhargava, Krister Rasmussen, Marisa Sylvester, and Victoria Yu for excellent research assistance. We are grateful to Matthew Adler, Susan Dudley, Dan Hemel, Nathan Hendren, Christine Jolls, Louis Kaplow, Paul Kelleher, Arik Levinson, and Eric Posner, and participants at the Tax Policy and the Economy conference, Yale Executive Action Workshop, and Yale Law School faculty workshop for valuable comments. We are also grateful to the anonymous reviewers.

Competing interest. As noted, both Liscow and Sunstein were involved in government in the revision of Circular A-94. Both 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.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

References

Acland, Daniel J., and David H. Greenberg. 2023. "Distributional Weighting and Welfare/Equity Tradeoffs: A New Approach." *Journal of Benefit-Cost Analysis*, *14*(1): 68–92.

Adler, Matthew D. 2000. "Beyond Efficiency and Procedure: A Welfarist Theory of Regulation." *Florida State University Law Review*, *28*: 241–338.

Adler, Matthew D. 2011. *Well-Being and Fair Distribution*. Oxford: Oxford University Press.

Adler, Matthew D. 2019. *Measuring Social Welfare*. Oxford: Oxford University Press.

Adler, Matthew D., and Eric Posner. 2006. *New Foundations of Cost-Benefit Analysis*. Cambridge, MA: Harvard University Press.

Adler, Matthew D., and Nils Holtug. 2019. "Prioritarianism: A Response to Critics." *Politics, Philosophy & Economics*, *18*(2): 101–144.

Alex Padilla, U.S. Senator for California: Press Releases. 2023. "Padilla Presses Administration to Prioritize Pajaro River Flood Projects and Protect Low-Income Communities." https://www.padilla.senate.gov/newsroom/press-releases/padilla-presses-administration-to-prioritize-pajaro-river-flood-projects-and-protect-low-income-communities/.

Allcott, Hunt, and Michael Greenstone. 2017. "Measuring the Welfare Effects of Residential Energy Efficiency Programs." National Bureau of Economic Research Working Paper No. 23386. https://www.nber.org/system/files/working_papers/w23386/w23386.pdf.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, *103*(6): 2121–2168.

Baicker, Katherine, and Nora Gordon. 2006. "The Effect of State Education Finance Reform on Total Local Resources." *Journal of Public Economics 90*: 1519–1535.

Bergson, Abram. 1954. "On the Concept of Social Welfare." *Quarterly Journal of Economics*, *68*(2): 233–252.

Boylan, Richard T., and Naci Mocan. 2014. "Intended and Unintended Consequences of Prison Reform." *Journal of Law, Economics, & Organization 30*(3): 558–586.

Breyer, Stephen. 1984. *Regulation and its Reform*. Cambridge, MA: Harvard University Press.

Calabresi, Guido. 1991. "The Pointlessness of Pareto: Carrying Coase Further." *Yale Law Journal*, *100*: 1211–1237.

Cecot, Caroline, and W. Kip Viscusi, 2015. "Judicial Review of Agency Benefit-Cost Analysis." *The George Mason Law Review*, *22*: 575–617.

Chetty, Raj. 2006. "A New Method of Estimating Risk Aversion." *The American Economic Review*, *96*(5): 1821–1834.

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *The American Economic Review*, *99*(4): 1145–1177.

Coleman, Jules. 1980. "Efficiency, Utility, and Wealth Maximization." *Hofstra Law Review*, *8*(3): 509–551.

Deaton, Angus. 2008. "Income, Health, and Well-Being around the World: Evidence from the Gallup World Poll." *The Journal of Economic Perspectives*, *22*(2): 53–72.

Dudley, Susan, and W. Kip Viscusi. 2023. "Biden's OMB Politicizes Cost-Benefit Analysis." *Wall Street Journal*. https://www.wsj.com/articles/bidens-omb-politicizes-cost-benefit-analysis-regulation-social-justice-2534e819.

Dworkin, Ronald. 1980. "Is Wealth a Value?" *The Journal of Legal Studies*, *9*(2): 191–226.

Exec. Order 12,866, 58 Fed. Reg. 190. October 4, 1993.

Exec. Order 13,771, Reducing Regulation and Controlling Regulatory Costs. January 30, 2017.

Exec. Order of April 6, 2023, Modernizing Regulatory Review. April 6, 2023.

Federal Emergency Management Agency (FEMA). 2023. "Building Resilient Infrastructure and Communities." https://www.fema.gov/grants/mitigation/building-resilient-infrastructure-communities.

Fennell, Lee Anne, and Richard H. McAdams. 2014. *Fairness in Law and Economics*. Cheltenham: Edward Elgar Publishing.

Finkelstein, Amy, Nathaniel Henren, and Erzo F.P. Luttmer. 2019. "The Value of Medicaid, Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy*, *127*(6): 2836–2874.

Frank, Thomas. 2022a. "FEMA Flood Program Could Violate Civil Rights Law." *Politico*. https://www.politico.com/news/2022/06/16/fema-flood-program-civil-rights-00037261.

Frank, Thomas. 2022b. "How FEMA Helps White and Rich People Escape Floods." *Politico*. https://www.politico.com/news/2022/05/27/unfair-fema-climate-program-floods-00032080.

Gandelman, Néstor, and Rubén Hernández Murillo. 2015. "Risk Aversion at the Country Level." *Review - Federal Reserve Bank of St. Louis*, *97*(1): 53–66.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

Giraldo Paez, Daniel, and Zachary Liscow. 2024. "Inequality Snowballing." *International Review of Law and Economics*, *77*: 106180. https://doi.org/10.1016/j.irle.2023.106180.

Harberger, Arnold C., Richard Zerbe, Glenn Jenkins, Joseph Cordes, David Weimer, Lisa Robinson, W. Kip Viscusi, Susan Dudley, Lynn Karoly, Donald Kenkel, Clark Nardinelli, Craig Thornton, Dale Whittington, Thomas Kniesner, and William Hoyt. 2023. "Letter to OIRA Administrator on Circular A4." *The George Washington University Regulatory Studies.* https://regulatorystudies.columbian.gwu.edu/letter-oira-administrator-circular-a4.

Hemel, Daniel. 2022. "Regulation and Redistribution with Lives in the Balance." *University of Chicago Law Review*, *89*: 649–734.

Hemel, Daniel. 2023. "Redistributive Regulations and Deadweight Loss." *The Journal of Benefit-Cost Analysis*, *14* (3): 407–436.

Hendren, Nathaniel. 2020. "Measuring Economic Efficiency Using Inverse-Optimum Weights." *Journal of Public Economics*, *187*: 1–12.

Hendren, Nathaniel, and Ben Sprung-Keyser. 2020. "A Unified Welfare Analysis of Government Policies." *The Quarterly Journal of Economics*, *135*(3): 1209–1318.

Hicks, John R. 1941. "The Rehabilitation of Consumers' Surplus." *The Review of Economic Studies*, *8*(2): 108–116.

HM Treasury. 2022. *The Green Book*. London: HM Treasury.

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

Internal Revenue Service. 2022. "Opportunity Zones." https://www.irs.gov/credits-deductions/businesses/opportunity-zones.

Interpersonal Comparisons of Well-Being (Jon Elster & John Roemer, eds.). 1991.

Issuance of Revised OMB Circular No. A-4, 88 Fed. Reg. 77615. November 13, 2023a.

Issuance of Revised OMB Circular No. A-4, 88 Fed. Reg. 77615. November 13, 2023b.

Issuance of Revised OMB Circular No. A-94, 88 Fed. Reg. 77615. November 13, 2023a.

Issuance of Revised OMB Circular No. A-94, 88 Fed. Reg. 77615. November 13, 2023b.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1991 "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives*, *5*(1): 193–206.

Kaplow, Louis. 2020. "A Unified Perspective on Efficiency, Redistribution, and Public Policy." *National Tax Journal*, *73*(2), 429–472.

Kaplow, Louis, and Steven Shavell. 1994. "Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income." *The Journal of Legal Studies*, *23*(2): 667–681.

Killingsworth, Matthew, Daniel Kahneman, and Barbara Mellers. 2023. "Income and Emotional Well-Being: A Conflict Resolved." *Proceedings of the National Academy of Sciences*, *120*(10): 1–6.

Kniesner, Thomas J., and W. Kip Viscusi. 2023. "Promoting Equity through Equitable Risk Tradeoffs." IZA Institute of Labor Economics Discussion Paper, No. 15771. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4294396.

Liscow, Zachary. 2014. "Reducing Inequality on the Cheap: When Legal Rule Design Should Incorporate Equity as Well as Efficiency." *Yale Law Journal*, *123*: 2478–2510.

Liscow, Zachary. 2018a. "Are Court Orders Sticky? Evidence on Distributional Impacts from School Finance Litigation." *Journal of Empirical Legal Studies*, *15*: 4–40.

Liscow, Zachary. 2018b. "Is Efficiency Biased?" *University of Chicago Law Review*, *85*: 1649–1718.

Liscow, Zachary. 2022. "Redistribution for Realists." *Iowa Law Review*, *107*(2): 495–561.

Liscow, Zachary. 2024. "The Benefit-Cost Analysis of Government Spending." Mimeo.

Liscow, Zachary, and Daniel Markovits. 2022. "Democratizing Behavioral Economics." *Yale Journal on Regulation*, *39*: 1274–1342.

Liscow, Zachary, and Abigail Pershing. 2022. "Why is so Much Redistribution In-Kind and Not in Cash? Evidence from a Survey Experiment." *National Tax Journal*, *75*(2): 313–354.

Matsumori, Kaosu, Kazuki Iijima, Yukihito Yomogida, and Kenji Matsumoto. 2021. "Interpersonal Comparison of Utility by Measuring Neural Activity." https://www.biorxiv.org/content/10.1101/2021.06.04.447048v1.full.

Mitchell, David. 2023. "Proposed Update to Federal Cost-Benefit Analysis Guidelines Correctly Focuses on Accounting for Inequality in Regulations." Washington Center for Equitable Growth. https://equitablegrowth.org/proposed-update-to-federal-cost-benefit-analysis-guidelines-correctly-focuses-on-accounting-for-inequality-in-regulations/.

Musgrave, Richard A. 1959. *The Theory of Public Finance: A Study in Public Economy*. New York: McGraw Hill.

National Bureau of Economic Research. 2023. "TAXSIM." https://www.nber.org/research/data/taxsim.

https://doi.org/10.1017/bca.2024.22 Published online by Cambridge University Press

Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.
Office of Management and Budget. 1972. "Circular A-94, Discount Rates to be Used in Evaluating Time-Distributed Costs and Benefits."
Office of Management and Budget. 1992. "Circular A-94, Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs."
Office of Management and Budget. 1996. "Economic Analysis of Federal Regulations Under Executive Order 12866."
Office of Management and Budget. 2023a. "Circular A-4, Regulatory Analysis."
Office of Management and Budget. 2023b. "Circular No. A-4: Explanation and Response to Public Input."
Office of Management and Budget. 2023c. "Circular A-94: Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs."
Posner, Richard A. 1979. "Utilitarianism, Economics, and Legal Theory." *The Journal of Legal Studies*, *8*(1): 103–140.
Posner, Richard A. 1985. "Wealth Maximization Revisited." *Notre Dame Journal of Law, Ethics and Public Policy*, *2*: 85–105.
Public Comment on Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs, 88 Fed. Reg. 20913. 2023a.
Public Comment on Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs, 88 Fed. Reg. 20913. 2023b.
Raskolnikov, Alex. 2021. "Distributional Arguments in Reverse." *Minnesota Law Review*, *105*: 1583–1666.
Resources for the Future. 2023. "Modernizing Regulatory Review: Exploring OMB's Updated Benefit-Cost Guidance." https://www.rff.org/events/rff-live/modernizing-regulatory-review-exploring-ombs-updated-benefit-cost-guidance/.
Revesz, Richard L., and Samantha P. Yi. 2022. "Distributional Consequences and Regulatory Analysis." *Environmental Law Review*, *52*(1): 53–98.
Rosen, Harvey S., and Ted Gayer. 2014. *Public Finance, Global Edition*, 10th ed. Maidenhead: McGraw-Hill Interamericana.
Samuelson, Paul A. 1947. *Foundations of Economic Analysis*. Cambridge, MA: Harvard University Press.
Sanchirico, Chris William. 2000. "Taxes Versus Legal Rules as Instruments for Equity: A More Equitable View." *The Journal of Legal Studies*, *29*: 797–820.
Sheffrin, Steve. 2013. *Tax Fairness and Folk Justice*. Cambridge, MA: Harvard University Press.
Stevenson, Betsey, and Justin Wolfers. 2008. "Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox." National Bureau of Economic Research Working Paper No. 14282.
Sunstein, Cass R. 2004. "Valuing Life: A Plea for Disaggregation." *Duke Law Journal*, *54*: 385–445.
Sunstein, Cass R. 2014. "The Limits of Quantification." *California Law Review*, *102*: 1369–1422.
Sunstein, Cass R. 2021. *Sludge: What Stops Us from Getting Things Done and What to Do About it*. Cambridge, MA: MIT Press.
Sunstein, Cass R. 2024. "The Economic Constitution of the United States." *Journal of Economic Perspectives*, *38*: 25–42.
U.S. Census Bureau. 2022. "Glossary." https://www.census.gov/programs-surveys/geography/about/glossary.html.
U.S. Department of Transportation. 2021. "Departmental Guidance on Valuation of a Statistical Life in Economic Analysis." https://www.transportation.gov/office-policy/transportation-policy/revised-departmental-guidance-on-valuation-of-a-statistical-life-in-economic-analysis.
U.S. Government Accountability Office. 2021. "Disaster Resilience: FEMA Should Take Additional Steps to Streamline Hazard Mitigation Grants and Assess Program Effects." https://www.gao.gov/assets/gao-21-140.pdf.

Cite this article: Liscow, Zachary, and Cass R. Sunstein. 2024. "Efficiency vs. Welfare in Benefit–Cost Analysis: The Case of Government Funding." *Journal of Benefit-Cost Analysis*, doi:10.1017/bca.2024.22