

A LIMITED DEFENSE OF EFFICIENCY AGAINST CHARGES OF INCOHERENCY AND BIAS

BY JONATHAN H. CHOI*

Abstract: Scholars have long debated the appropriate balance between efficiency and redistribution. But recently, a wave of critics has argued not only that efficiency is less important, but that efficiency analysis itself is fundamentally flawed. Some say that efficiency is incoherent because there is no neutral baseline from which to judge inefficiency. Others say that efficiency is biased toward those best able to pay (generally, the rich). This essay contends that efficiency is not meaningfully incoherent or biased. The most widely discussed forms of efficiency do not require any particular baseline, and even those that do require a baseline can still serve as useful approximations of more theoretically sound but computationally demanding measures. Moreover, arguments of bias do not account for the source of funds in public projects, produce unintuitive results, and draw an arbitrary cutoff between bias and non-bias that elides important distributional details. Ultimately, the tradeoff between efficiency and redistribution remains the most useful frame for policy debate.

KEY WORDS: efficiency, Pareto, Kaldor-Hicks, wealth maximization, deadweight loss, law and economics

Efficiency pervades modern legal analysis—it is a lens through which we view legislation, regulation, court rulings, and more. Scholars have long debated the exact borders of efficiency’s proper domain, including whether it should take priority over redistribution. But in recent years, a new scholarly movement has begun to attack the concept of efficiency itself. Scholars writing in this vein argue that efficiency analysis is biased toward the rich or entirely incoherent.

This essay mounts a limited defense of efficiency as a useful and meaningful concept within legal theory. It contends that efficiency analysis is not incoherent because it produces unambiguous, determinate results. Contrary to recent claims that efficiency requires the selection of an arbitrary baseline, the most popular forms of efficiency (Pareto efficiency and Kaldor-Hicks efficiency) do not require any particular baseline in order to determine which states of the world are efficient. Moreover, forms of efficiency that do require a baseline (wealth maximization and deadweight loss analysis) may still be useful in practice as proxies for more theoretically sound but computationally difficult standards (like the maximization of social welfare).

* Law School, University of Minnesota, jonchoi@umn.edu. Competing Interests: The author declares none. Thanks to Neil Buchanan, Erick Sam, David Schmidt, the editors of the journal *Social Policy and Philosophy*, and the other contributors to this volume for helpful comments.

The essay further contends that efficiency is not meaningfully biased. Critiques that public spending favors the rich typically fail to account for the source of funds—it makes little sense to call a policy biased if it serves the rich *and is paid for by the rich*. Moreover, conventional definitions of rich-bias can lead to unintuitive results. For example, on Zachary Liscow's popular definition of rich-bias, a policy that protects all property rights equally would be rich-biased because the rich have more property. Most fundamentally, it is not clear that analysis of bias is helpful to those seeking more redistribution. If redistribution is the goal, why distinguish between a policy that is slightly rich-biased and one that is neutral? Why not simply seek more redistribution whenever possible? A focus on bias or non-bias eliminates subtle but important gradations in redistribution.

This is not to say that all criticisms of efficiency are invalid. But specific issues with efficiency should be substantively discussed on their own terms. Attacks on it as a general theoretical concept tend to conceal more than they reveal.

To clarify the limitations of efficiency, this essay draws two important distinctions. First, we should distinguish between arguments that efficiency is a *bad tool* and arguments that efficiency is a *badly used tool*. The latter arguments are hardly contestable; efficiency can and has been cherry-picked by those attempting to advance a particular political agenda. However, this is true of virtually all popular legal concepts and hardly constitutes a strike against the core idea. In this essay, I focus on claims that efficiency itself is inappropriate or meaningless and separate these claims from frustration about its misapplication.

Second, we should distinguish between problems in *theory* and problems in *practice*. Efficiency analysis is hard; its practitioners face large and sometimes insurmountable obstacles in data and in the assumptions they must make to render the analysis tractable. It may be that these assumptions are ultimately unsupported and efficiency analysis is unworkable in practice. But these practical difficulties should be discussed without dismissing efficiency entirely. Critics of efficiency, including myself, have an even greater duty of charity to defend efficiency when it is defensible.

I. WHAT IS EFFICIENCY?

Depending on the speaker, "efficiency" can denote diverse and contradictory concepts. One common definition is Pareto efficiency. A state of affairs is Pareto efficient if no change could increase one individual's well-being without decreasing the well-being of another individual. But the Pareto criterion is a relatively harsh one: even many appealing policies will harm *someone*. For example, a regulation that prohibits a factory owner from dumping toxic waste into a lake will not be a Pareto improvement because it harms the factory owner, no matter how toxic the waste or how culpable the owner.

Kaldor-Hicks efficiency relaxes this rule by introducing a *compensation assumption*. It asks whether a change would improve the well-being of some and not harm any others, *assuming* that the winners were to compensate the losers. Because Kaldor-Hicks deals with hypothetical compensation, rather than actual compensation, Kaldor-Hicks improvements are much more common. The toxic waste regulation might be a Kaldor-Hicks improvement but not a Pareto improvement, if everyone would be at least as well off were the polluted-upon to compensate the polluter.¹

Neither Pareto nor Kaldor-Hicks analysis identifies unique efficient policies; many competing policies may each be “efficient” in a Pareto or Kaldor-Hicks sense,² or two different policy options may be incomparable.³ Pareto and Kaldor-Hicks simply attempt to identify some set of policies as efficient without providing a way to discriminate in every case. Despite these limitations, many economists prefer Pareto and Kaldor-Hicks to the maximization of social welfare because they do not explicitly require cardinal or interpersonal comparisons of utility—they only require each individual to identify which of two situations she prefers, without specifying by how much.

Wealth maximization is one attempt to implement a Kaldor-Hicks criterion in practice. Wealth maximization imagines that for any proposed policy change, each individual would pay a certain amount for the change to occur (willingness to pay) or accept a certain amount to forgo the change (willingness to accept). We could then sum all of the willingnesses to pay or

¹ Hypothetical compensation may not be paid for a variety of reasons, including transaction costs and political difficulties.

² The Kaldor-Hicks criterion has some well-known technical oddities, including the possibility of “reversals” illustrated by the Scitovsky paradox, that two different states could each be Kaldor-Hicks superior to each other. Tibor Scitovsky, *Papers on Welfare and Growth* (London: Routledge, 2003). The paradox results from differences between willingness-to-accept and willingness-to-pay. There is good reason, drawing on the behavioral economics literature, to believe that one might require greater compensation to give something up than one would have been willing to pay for that thing in the first place (“loss aversion”). Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler, “Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias,” *Journal of Economic Perspectives* 5, no. 1 (1991): 193–206. Duncan Kennedy has made an analogous point. Duncan Kennedy, “Cost-Benefit Analysis of Entitlement Problems: A Critique,” *Stanford Law Review* 33, no. 3 (1981): 387–445.

Scitovsky (and later, Samuelson, along slightly different lines) proposed amending the definition of Kaldor-Hicks superiority to exclude the possibility of reversals. Allan M. Feldman, “Kaldor-Hicks Compensation,” in Peter Newman, ed., *The New Palgrave Dictionary of Economics and the Law*, Volume 2 (London: Macmillan, 1998), 417–21. Making this amendment solves the problem but increases the likelihood that, given two possible states of the world, neither will be Kaldor-Hicks superior to the other, decreasing the usefulness of Kaldor-Hicks analysis.

³ This could happen if two different policies would result in different sets of individuals. For example, given current state-of-the-world s and potential state of the world s' that would result from a certain policy, if the policy would incentivize parents to have additional children, then s and s' are no longer comparable, because s' includes individuals (the additional children) who cannot be compensated (either actually or hypothetically) for a move to s . I hope to discuss this limitation to Pareto and Kaldor-Hicks in future work.

willingnesses to accept, and if the amount were positive, then the policy change would be wealth maximizing.⁴

Kaldor-Hicks efficiency and wealth maximization are especially important because they motivate real-world cost-benefit analysis, which attempts to quantify costs and benefits in terms of willingness to pay or accept. Although some scholars conflate wealth maximization with Kaldor-Hicks analysis,⁵ they differ in their formal definitions, and wealth maximization encounters some particular problems that Kaldor-Hicks does not.⁶

A final approach to efficiency focuses on “deadweight loss” or “excess burden.” This approach concerns welfare losses from government policies that “distort” market behavior. For example, a tax might cause a “loss of consumer (and producer) surplus when comparing the after-tax world to the before-tax world.”⁷ Deadweight loss analysis imagines that there is an efficient, undistorted, pre-tax world, and calculates the social cost of moving from that world to a less efficient, distorted, post-tax world.

Between these four concepts, Kaldor-Hicks efficiency is generally what law and economics scholars mean when they discuss efficiency simpliciter. Deadweight loss receives greater attention in some specialized fields, like taxation and antitrust. The distinctions matter: each of these definitions has its advantages and disadvantages, as we will see below.

How is efficiency used in the literature? Scholars frequently contrast efficiency with redistribution—Arthur Okun described the choice between the two as the “big tradeoff” in public policy.⁸ In legal scholarship,

⁴ Both willingness to pay and willingness to accept can be positive or negative, and each reflects a discrete and competing approach to wealth maximization. Willingness to pay is a measure of an individual’s “compensating variation,” meaning the amount of money that would leave her welfare unaltered if the change does occur, as measured using prices in the world where the change occurs. Willingness to accept is a measure of an individual’s “equivalent variation,” meaning the amount of money that would render her as well off under the status quo if the change does not occur, as measured using prices under the status quo. For a variety of reasons, including differences in price levels between the status quo and post-change world, willingness to pay and willingness to accept are different and sometimes competing measures. Lewis Kornhauser, “Wealth Maximization,” in *The New Palgrave Dictionary of Economics and the Law*, ed. Peter Newman (London; New York: Macmillan, 1998), 679.

⁵ For example, Richard Posner has sometimes referred to wealth maximization and Kaldor-Hicks interchangeably. Richard A. Posner, “The Ethical and Political Basis of the Efficiency Norm in Common Law Adjudication,” *Hofstra Law Review* 8 (1980): 491, 495.

⁶ Charles Blackorby and David Donaldson have highlighted several respects in which wealth maximization is different from, and in their view worse than, Kaldor-Hicks optimization. Increases in wealth are possible even when no Kaldor-Hicks improvement has occurred, so that an increase in wealth is a necessary but not sufficient condition for a Kaldor-Hicks improvement. Wealth maximization and Kaldor-Hicks will yield the same result if and only if price levels are the same in all scenarios under consideration, which is typically not the case. Charles Blackorby and David Donaldson, “A Review Article: The Case Against the Use of the Sum of Compensating Variations in Cost-Benefit Analysis,” *Canadian Journal of Economics* 23, no. 3 (1990): 471–94.

⁷ David A. Weisbach, “Line Drawing, Doctrine, and Efficiency in the Tax Law,” *Cornell Law Review* 84 (1999): 1650.

⁸ Arthur M. Okun, *Equality and Efficiency: The Big Tradeoff* (Washington, DC: Brookings Institution, 1975).

commenters have disagreed on whether government policies should be selected only with efficiency in mind, or with distributional effects in mind as well. In tort law, for example, we might expect that trespassers are systematically poorer than landowners, or that owners of pollution-generating factories are systematically richer than their polluted-upon neighbors. Louis Kaplow and Steven Shavell famously argued that legal rules should be chosen only on efficiency grounds, because it is systematically better to redistribute through taxes and transfers.⁹

Critics have attacked Kaplow and Shavell's argument since it was first introduced. Some scholars have argued that legal rules are efficient avenues for redistribution and provided economically grounded explanations, including both behavioral¹⁰ and neoclassical explanations.¹¹ Other scholars have argued that a preference for redistribution through taxes and transfers in a first-best world is irrelevant, because taxes and transfers are inadequate in our n^{th} -best world.¹² Still others have suggested that the assumptions undergirding Kaplow and Shavell's theoretical model are unrealistic.¹³ Many of these arguments cabin the scope of efficiency analysis in compelling and useful ways.

But a more recent line of scholarship advances an even more fundamental rebuttal. These articles cast doubt on the concept of efficiency itself—they argue that efficiency is incoherent or biased. They therefore represent a

⁹ Louis Kaplow and Steven Shavell, "Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income," *Journal of Legal Studies* 23, no. 2 (1994): 669.

¹⁰ Christine Jolls argues from the perspective of behavioral economics that redistribution through legal rules may in fact cause less distortion than redistribution through taxation. Christine Jolls, "Behavioral Economic Analysis of Redistributive Legal Rules," *Vanderbilt Law Review* 51, no. 6 (1998): 1656–67.

¹¹ Chris Sanchirico argues that legal rules can facilitate efficient redistribution based on "ability," for which labor income is a proxy. Conceptually, if one's likelihood of incurring tortious liability were also an imperfect proxy for ability, it could be optimal to enact redistribution through legal rules in order to better redistribute on ability. Chris William Sanchirico, "Taxes Versus Legal Rules as Instruments for Equity: A More Equitable View," *Journal of Legal Studies* 29, no. 2 (2000): 802–3. Brian Galle argues that local tort rules can redistribute more effectively than local tax law, because local taxation incurs additional distortions that tort rules do not (primarily due to taxpayer mobility and choice of residence). Brian Galle, "Is Local Consumer Protection Law a Better Retributive Mechanism than the Tax System," *NYU Annual Survey of American Law* 65 (2010): 525.

¹² See, for example, Richard S. Markovits, "Why Kaplow and Shavell's 'Double-Distortion Argument' Articles Are Wrong," *George Mason Law Review* 13, no. 3 (2005): 521; Lee Anne Fennell and Richard H. McAdams, "The Distributive Deficit in Law and Economics," *Minnesota Law Review* 100, no. 2 (2016): 1051. Zachary Liscow also argues that because "individuals silo their policy views," we cannot practically offset less redistribution through legal rules with more redistribution through taxes and transfers. Zachary Liscow, "Redistribution for Realists," *University of Iowa Law Review* 107, no. 2 (2022): 495.

¹³ See, for example, Zachary Liscow, "Reducing Inequality on the Cheap: When Legal Rule Design Should Incorporate Equity as Well as Efficiency," *Yale Law Journal* 123, no. 7 (2015): 2478; Richard L. Revesz, "Regulation and Distribution," *New York University Law Review* 93, no. 6 (2018): 1489–1578; Markovits, "Why Kaplow and Shavell's 'Double-Distortion Argument' Articles Are Wrong"; Fennell and McAdams, "The Distributive Deficit in Law and Economics."

more ambitious project, not just to erode the priority of tax-and-transfer at the margins but to jettison it entirely.

I argue that these arguments all reduce to familiar disputes over the empirical assumptions underlying efficiency analysis, and that efficiency is neither biased nor incoherent. While many of the objections to specific uses of efficiency analysis are valid, they should be addressed on their own terms and not through broad condemnations of the concept of efficiency.

II. WHY EFFICIENCY IS NOT INCOHERENT

One prominent critique of efficiency is that it is simply *incoherent*—that efficiency contains so many problematic assumptions that it is meaningless. Neil Buchanan and Michael Dorf have recently made this point with particular persuasiveness.

Here is the argument: efficiency must be analyzed with respect to a baseline of property rights. However, property rights are *themselves* the products of legal rules.¹⁴ The selection of optimal property rights will therefore change the optimal set of other legal rules; and any changes to the optimal set of other legal rules will change the optimal set of property rights. Thus, efficiency is circular. In the absence of some “natural” baseline set of legal rules, it is impossible to assess which legal rules maximize efficiency.¹⁵

Consider a specific example that Buchanan and Dorf provide: sport-utility vehicles (SUVs). They argue that SUVs exist as a class of vehicle in the United States primarily to meet federal fuel efficiency standards from the 1980s. It is therefore incoherent to ask what the “efficient” quantity of SUVs would be. “The supply curve only exists at all because the government made a choice decades ago that it did not have to make (but for which there was no default baseline).”¹⁶

Buchanan and Dorf seem to suggest that supply and demand cannot be considered at all when designing legal rules. But the alternatives hardly seem appealing: pure formalism, perhaps, or rules selected solely on the

¹⁴ “The rules of property law ... are fundamental in determining what people can and cannot do in the marketplace.” Neil H. Buchanan, “The Role of Economics in Tax Scholarship,” in David A. Brennan, Karen B. Brown, and Darryl Jones, eds., *Beyond Economic Efficiency in United States Tax Law* (New York: Wolters Kluwer Law and Business, 2013), 11.

¹⁵ Other scholars have made analogous arguments, perhaps most famously Liam Murphy and Thomas Nagel. Liam B. Murphy and Thomas Nagel, *The Myth of Ownership: Taxes and Justice* (New York: Oxford University Press, 2002).

¹⁶ Neil H. Buchanan and Michael C. Dorf, “A Tale of Two Formalisms: How Law and Economics Mirrors Originalism and Textualism,” *Cornell Law Review* 106 (2021): 20. In an essay that Buchanan published several years earlier, he offered as another example the property law rule of adverse possession, which awards property rights to an adverse possessor who maintains “open and notorious” possession of property against the property’s nominal owner. Buchanan argued that there is no natural or proper baseline property right that we can assess efficiency against here, because the concept of adverse possession is itself a property right. Neil H. Buchanan, “The Role of Economics in Tax Scholarship,” 12-13.

basis of their distributional effects, eschewing any consideration of the rules' incentive effects. Thus, Buchanan and Dorf prove too much—even critics of efficiency generally acknowledge that incentive effects matter, even if they should not be the sole considerations in crafting legal rules.

Moreover, not every type of efficiency analysis requires a baseline. Buchanan and Dorf purport to discuss Pareto and Kaldor-Hicks efficiency, as is standard in law and economics.¹⁷ But their argument seems most apt as a critique of deadweight loss. Of the types of efficiency discussed above, Pareto and Kaldor-Hicks efficiency do not require a baseline; only deadweight loss analysis and wealth maximization do, for different reasons.

Pareto requires only that we compare each individual's utility in a world with and without SUVs, which is a conceptually feasible task. Kaldor-Hicks is somewhat more complicated because it requires us to imagine compensation is paid. However, the compensation can be paid in whatever goods exist in the particular hypothetical world in question; in SUVs, sedans, bananas, dollars, or anything else. Hypothetical compensation does not need to be paid in any particular coin. Thus, again, Kaldor-Hicks efficiency requires no baseline and avoids charges of incoherency.¹⁸

Put differently, Pareto and Kaldor-Hicks require us to compare two different worlds with comprehensively specified packages of background property rights and other legal rules, and then compare individuals' welfare levels between these two worlds. Pareto and Kaldor-Hicks are inherently relational and thus require no uniquely privileged comparison world to make assessments of efficiency.

Buchanan and Dorf's argument is most apt with deadweight loss analysis, focusing specifically on supply, demand, and the deadweight loss "triangle." This formula does, indeed, suffer from a baseline problem. And so does wealth maximization, albeit for reasons that Buchanan and Dorf do not discuss.

First, wealth maximization is not well-defined because *wealth* is not well-defined between different states of the world. For example, Rolls-Royces might be more expensive, and Toyotas cheaper, in a society with more wealth inequality. Changes to entitlement regimes therefore change relative prices, so that money is not a consistent numeraire, and wealth in one hypothetical world cannot be directly compared with wealth in another.¹⁹ This problem has been widely known and discussed among economists for

¹⁷ Buchanan and Dorf, "A Tale of Two Formalisms," 602.

¹⁸ Kaldor-Hicks has its own problems, which are well-known and have been discussed for decades. See note 2 above, discussing the Scitovsky paradox and some potential solutions to the paradox.

¹⁹ Kaplow and Shavell note that "wealth maximization is not a well-defined concept; to compute wealth, one must know the prices of different goods and services, yet there is no natural set of prices to use. . . . [T]he absence of a natural set of prices is not a problem that can be resolved by a simple price index adjustment, similar to adjustments for pure inflation, because relative prices differ." Louis Kaplow and Steven Shavell, *Fairness Versus Welfare* (Cambridge, MA and London: Harvard University Press, 2002), 36.

decades.²⁰ It may, however, be relatively insignificant in practice, since most policies will have little impact on overall relative prices.

Relatedly, wealth maximization requires a baseline in the sense that individuals' willingness to accept or pay is determined by their existing endowments. As Duncan Kennedy has explained, if two individuals are stranded in the desert with a life-saving glass of water, wealth maximization will leave the glass with whichever individual started with it, because no amount of money could persuade the glass's current possessor to give it up. These "wealth effects" make wealth maximization dependent on the choice of baseline and, therefore, arbitrary.²¹

If deadweight loss minimization or wealth maximization were the only goal, these problems would be insuperable—which is one of many reasons why they are *not* the only goal. Richard Posner's onetime suggestion that wealth maximization might be an end in itself was roundly criticized within the legal academy, with Posner eventually conceding that wealth maximization is merely a pragmatic means to the ultimate end of welfare maximization.²² Among legal economists in general, welfare maximization is the standard fundamental justification for wealth maximization and cost-benefit analysis.²³

Crucially, even if deadweight loss minimization and wealth maximization are not globally well-defined, they may still be *locally* well-defined; that is, we might still identify which moves minimize deadweight loss or maximize wealth *holding all other factors constant*. Under this approach, all background legal rules, other than the specific legal rules under consideration, are assumed as a baseline. This simplifies analysis but comes with a cost—it may make deadweight loss minimization and wealth maximization imperfect proxies for welfare maximization, if that is the ultimate goal. But using these assumptions, deadweight loss and wealth maximization remain at least *coherent* concepts as used in practice.

Consider again Buchanan and Dorf's SUV example. Imagine that the National Highway Traffic Safety Administration (NHTSA) is considering whether to require electronic anti-rollover protection, to prevent a kind of accident to which SUVs are especially prone. The NHTSA conducts

²⁰ Lewis A. Kornhauser, "Wealth Maximization," 679-83.

²¹ Kennedy, "Cost-Benefit Analysis of Entitlement Problems: A Critique," 422-44.

²² Richard A. Posner, "Wealth Maximization and Tort Law: A Philosophical Inquiry," in David G. Owen, ed., *Philosophical Foundations of Tort Law* (Oxford; New York: Oxford University Press, 1995) 99-100.

²³ See, for example, Matthew D. Adler and Eric A. Posner, *New Foundations of Cost-Benefit Analysis* (Cambridge, MA and London: Harvard University Press 2006), 6 ("Our argument is that CBA is best defended as a *welfarist* decision procedure."); Kaplow and Shavell, *Fairness Versus Welfare*, 37 ("[M]aximization of wealth (defined, perhaps, with respect to current prices) may in fact reasonably approximate maximization of social welfare in many contexts. Thus, under welfare economics, although wealth is not in itself deemed to be valuable, analysis that assesses policies based on their aggregate impact on wealth will often prove useful."). See also *ibid.*, 36, making an analogous point about efficiency in general.

conventional cost-benefit analysis and estimates that the benefits of the regulation would exceed its costs—in other words, that the regulation is wealth maximizing, or that it minimizes deadweight loss. The cost-benefit analysis is necessarily a *ceteris paribus* exercise—it holds all other factors constant, including the existence of SUVs and by extension the existence of the regulatory environment that encourages consumers to buy SUVs.

This may or may not be the best method to use, but it is at least a *well-defined, unambiguous* method. Costs and benefits are unambiguously quantifiable within this framework—the calculations will involve assumptions, but the assumptions are definite and tractable. If the regulations encouraging consumers to purchase SUVs are not wealth-maximizing, then under the tax-and-transfer paradigm, those regulations should be removed. And if they are removed, the cost-benefit analysis regarding anti-rollover protection would become superfluous as well. Wealth maximization is not a single, global analysis. In theory and in practice, it is a series of piecemeal local assessments.

Again, neither wealth maximization nor deadweight loss minimization are ends in themselves. They are merely techniques that are *instrumentally useful* in attempting to reach some other goal, welfare maximization being the usual candidate. If these tools do not maximize welfare, they should be discarded—and there are many reasons to believe that they may not in fact be welfare-maximizing, especially if society inadequately redistributes through taxes and transfers. (As noted above, the general theoretical argument that we should maximize regulatory efficiency and then redistribute elsewhere has come up against legitimate criticism on empirical grounds.) However, this is a step away from efficiency being *incoherent*, and our doubts about the efficacy of efficiency should be confronted and discussed substantively rather than assuming them away.

In fact, Pareto and Kaldor-Hicks arguably also operate better as means rather than ends. Pareto and Kaldor-Hicks efficiency do not identify unique optima—many different policies might be Pareto or Kaldor-Hicks efficient in the sense that there are no available Pareto or Kaldor-Hicks improvements. Thus, we need some other criterion (like welfare maximization) to break the tie between these many possible “efficient” policies.²⁴

If efficiency analysis, in whatever form, does not satisfy this other criterion, it should be discarded. But the question whether it does in fact satisfy this criterion is also a cognizable question. Of course, efficiency is complicated, and the calculation process is not easy. It is assuredly difficult; but it is not *incoherent*.

²⁴ Of course, this argument will be unpersuasive to theorists that favor wealth maximization, Kaldor-Hicks, or Pareto over welfarism because they eschew interpersonal or cardinal comparisons of utility. The argument instead targets those who theoretically subscribe to welfarism but find it too demanding in practice.

III. WHY EFFICIENCY IS NOT BIASED

Zachary Liscow advances another influential argument against efficiency by contending that it is generally “rich-biased” in a manner that will “give disproportionate legal entitlements to the rich for free, exacerbating inequality.”²⁵ In Liscow’s view, any analysis (like cost-benefit analysis) that uses willingness to pay as its measure of efficiency will generally favor the rich—the rich will generally be willing to pay more, because they have more money. (Other authors have made similar arguments, including Edwin Baker, Matthew Adler, and Eric Posner,²⁶ but Liscow’s is a particularly recent and powerful rendition of the critique.)

Which version of efficiency does Liscow discuss? His discussion and mathematical proofs suggest he has in mind what I call “wealth maximization,” although he calls it “Kaldor-Hicks efficiency.” The move to equate Kaldor-Hicks with wealth maximization has precedents in law and economics, and is not a fatal flaw in any case—wealth maximization is one attempt to instantiate Kaldor-Hicks, subject to the issues with relative prices and wealth effects noted above, which are incidental to Liscow’s analysis.

In a neat piece of analytical symmetry, Liscow’s argument also assumes the existence of a neutral baseline. He establishes the baseline by noting that “[a]n efficient policy is neutral if, as one gets richer, efficient legal rules do not change one’s legal entitlements.”²⁷ In contrast, a policy “is poor-biased if, as one gets richer, one gets fewer legal entitlements from efficient legal policies.”²⁸ And a policy “is rich-biased if, as one gets richer, one tends to get more legal entitlements from efficient policies.”²⁹ He notes as an example that “everyone has the same willingness to pay for one dollar in increased or decreased income.”³⁰ Thus, a policy that provided every individual with a flat amount of money would be neutral. In contrast, a policy that tended to provide richer citizens with more money would be rich-biased.

There are several problems with this choice of baseline. First, it only considers public spending, without considering sources of public revenue. This is a classic trope in tax policy: the distributional effect of public

²⁵ Zachary Liscow, “Is Efficiency Biased?” *University of Chicago Law Review* 85, no. 7 (2018): 1656.

²⁶ C. Edwin Baker, “The Ideology of the Economic Analysis of Law,” *Philosophy and Public Affairs* 5 (1975): 3, 16-19, 47-48; Posner and Adler, *New Foundations of Cost-Benefit Analysis*, 16-19; Eric A. Posner and Matthew D. Adler, “Rethinking Cost-Benefit Analysis,” *Yale Law Journal* 109 (1999): 184.

²⁷ Liscow, “Is Efficiency Biased?” 1655.

²⁸ *Ibid.*, 1655.

²⁹ *Ibid.*, 1655.

³⁰ *Ibid.*, 1655.

spending cannot be properly analyzed without understanding how that spending is funded, and vice versa.³¹

For example, imagine a new demogrant that gave each citizen one dollar. This policy would be neutral within Liscow's framework, neither rich-biased nor poor-biased. But the actual distributional effect of the demogrant depends on how it is to be funded. If the demogrant were funded by our federal income tax—which is generally progressive, so that the rich pay a larger percentage of their incomes than the poor—then such a demogrant would substantially redistribute from the rich to the poor. In fact, the tax rate would not even need to be progressive in order to effect redistribution. It is a well-known result within the literature on universal basic income that a demogrant plus a *flat* tax can effect substantial redistribution.³²

To take the strongest version of Liscow's argument, consider one of his examples: public parks. A government that maximizes wealth outside of tax policy will build parks in line with residents' willingness to pay rather than in line with the utility that residents actually derive from the parks.³³ Because rich residents are able to pay more, this means that the government will in general spend more on parks for the rich than parks for the poor. In Liscow's terms, this policy is therefore rich-biased.

But this characterization hangs on Liscow's definition of neutrality, which crucially includes only spending and not taxing. Once taxes are taken into account, the park system might significantly redistribute from the rich to the poor. This would be the case if public works were generally funded through progressive taxes. To put it more concretely, imagine that the National Parks Service is deciding whether to fund a new national park, using the proceeds of federal income taxes. The average taxpayer within the bottom 50 percent of the income distribution paid approximately \$694 in federal income taxes in 2017. The average taxpayer within the top 10 percent of the income distribution paid approximately \$132,168 in federal income taxes in 2017.³⁴ It might be that the new national park would benefit the rich more than the poor—but would it benefit them 190 times more?³⁵

³¹ As Daniel Shaviro has noted: "The distinction between taxes and spending ... depends on pure form." Daniel Shaviro, "Rethinking Tax Expenditures and Fiscal Language," *Tax Law Review* 57 (2004): 191.

³² N. Gregory Mankiw, Matthew Weinzierl, and Danny Yagan, "Optimal Taxation in Theory and Practice," *Journal of Economic Perspectives* 23, no. 4 (2009): 147–74.

³³ Liscow, "Is Efficiency Biased?" 1650–51.

³⁴ Erica York, "Summary of the Latest Federal Income Tax Data, 2020 Update," <https://taxfoundation.org/summary-of-the-latest-federal-income-tax-data-2020-update>. These are the most recently available data on tax collections stratified by income level; 2018 data will likely show somewhat less progressivity, due to 2017 tax reform.

³⁵ The situation is more complicated for local parks funded by local property taxes. Many of the poorest taxpayers rent, rather than own, so do not directly pay any property tax at all. However, the incidence of the property tax may still partly fall on poor taxpayers, who would pay higher rent as a result. On the other hand, to the extent that property tax is capitalized into the value of real estate, and particularly to the extent that a Georgist tax on unimproved land theoretically taxes pure economic rent and therefore falls solely on capital-owners, the incidence on the poor may be minimal. Finally, property taxes are typically flat, and property

This alternative baseline emphasizes the actual source of funds and the actual alternative to contemplated policies. If the Parks Service decided not to fund any parks, would the poor taxpayer benefit from the elimination of this “rich-biased” policy? If she bore little of the cost of the parks, she would likely prefer to keep them.

The essential distinction between my alternative baseline and Liscow’s baseline is that mine assumes that public spending is funded through tax revenues. Moreover, it assumes that this is also true on the *margins*—that any contemplated additional spending will generally be funded through taxation with roughly the same distributional profile. Liscow instead assumes that distributional effects are “sticky,” such that a regressive policy generally will not be offset by *any* change to the tax system.³⁶ This is likely true when considering policies in isolation: no congressperson would observe an Environmental Protection Agency (EPA) regulation that benefits the rich and conclude that we must therefore raise the top tax rate by 0.001 percent. But on a wider scale, this assumption becomes difficult to defend. New public spending will ultimately be paid for *somehow*; and as a rough approximation, it seems most reasonable to assume that it will be paid for in the same manner as existing public spending.³⁷

Under Liscow’s baseline, rich-biased policies are exceedingly common; but taking the source of funds into account, rich-biased policies become exceedingly rare. (Most taxes in the United States are progressive; sales taxes are a key exception, but they make up a relatively small percentage of overall government revenue, and there is no national general sales tax.³⁸) If tax rates were merely flat, a policy would be distributionally neutral if willingness to pay increased as a simple linear function of income. But so long as taxes are progressive, policies will only be distributionally

ownership as a percentage of income or wealth likely declines the richer one gets, so that property taxes may ultimately be flat or regressive in the upper-income range.

³⁶ Liscow, “Is Efficiency Biased?” 1662–66. Liscow describes circumstances under which offsets would be more or less likely, and emphasizes that if offsets occur then rich-biased policies may be justified; nevertheless, in contrast to this essay, he appears to focus on specific offsets rather than the potential that expenditures would be implicitly offset through the general budget process.

³⁷ Because the federal government may freely borrow, it is possible that new spending will not be paid for until far into the distant future. But because money is fungible, it is generally not possible to distinguish which expenditures are paid for now versus later, and given the difficulty of predicting future tax rates, current tax rates are likely our best approximation of the taxing side for distributional analysis.

³⁸ Payroll taxes are also regressive, but they are theoretically earmarked to fund specific government programs like unemployment insurance, Social Security, and Medicare, whose overall effect (taking into account both taxes and spending) is progressive. As a result, it seems more plausible that marginal expenditures are funded by income taxes rather than payroll taxes. This may not always be the case given that unemployment insurance, Social Security, and Medicare are pay-as-you-go programs with historically underfunded accounts to pay for future expenses. However, we are now reaching the inflection point where expenses for these programs begin to exceed revenues, making it less likely that general social programs can be funded from payroll taxes in the future.

rich-biased if the rich are willing to pay more *as a percentage of their income*, and if their willingness to pay scales-up more quickly than rates of tax. Thus the dramatic difference between average tax revenues from the rich and poor taxpayer mentioned above. The rich taxpayer pays an average federal income tax of 21.5 percent, whereas the poor taxpayer pays an average federal income tax of 4 percent. This is compounded by the fact that the rich taxpayer has a substantially higher income, thus the result that the rich taxpayer pays 190 times more than the poor one.

The rich might value some public goods more as a percentage of their income, compared to poor taxpayers. Poor taxpayers can reap higher marginal benefits from spending on private goods—at the margin, they are more likely to spend on necessities, while the rich might spend on luxuries. Having satiated her interest in personal material comforts, the rich taxpayer might conceivably be willing to pay for some public goods a percentage of her income 5 times greater than the poor taxpayer's percentage, which corresponds to an amount 190 times greater in raw expenditure.³⁹ But it is not obvious what public goods this would apply to. Moreover, it seems the exception rather than the rule, since many of the most important government-provided public goods—national defense, clean air and water—are basic necessities rather than luxuries. After taking both spending and taxing into account, these policies would seem to favor the poor, even though they meet Liscow's criterion for rich-bias.

I am *not* arguing that present policy should be altered to favor the rich. I only mean to say that efficiency can be meaningfully described as poor-biased just as easily as it can be described as rich-biased. The finding of bias is thus heavily dependent on the neutral baseline that one chooses, and there is reason to doubt Liscow's choice.

But what if legal rules are in question rather than public spending? In these cases, there is no issue about the source of funds. Liscow tends to emphasize rich-bias in public spending,⁴⁰ presumably because these cases are intuitive and common. And he acknowledges that, "[i]n practice, torts typically do not offer any relief to those whose quality of life is harmed by worse health but who suffer no financial harm, such as compensation for pain and suffering, in ways that give more compensation to the rich than to the poor."⁴¹ However, it seems likely that all legal rules that protect property in a broad fashion are systematically rich-biased under Liscow's criterion. This is so because the rich tend to have more property than the poor, so they will benefit more from such rules in dollar terms.

One hypothetical example of a neutral or poor-biased property rule would be one written narrowly to protect the sort of property that the poor

³⁹ This calculation uses the *average* tax rate for the rich and the poor; it is possible that using *marginal* tax rates could give different results.

⁴⁰ He cites as particular examples government spending on research, roads, law enforcement, voting, parks, and transportation. Liscow, "Is Efficiency Biased?" 1674–77.

⁴¹ Liscow, "Is Efficiency Biased?" 1672 n. 80.

disproportionately own. So we could, say, enact theft prevention laws that protect mopeds but not Rolls Royces. Alternatively, we could effect redistribution within legal proceedings by explicitly changing standards of proof or levels of damages to benefit poorer litigants. For example, poor plaintiffs might be required only to meet the preponderance-of-evidence standard in civil proceedings (the standard under current law), whereas rich plaintiffs might be required to meet the beyond-a-reasonable-doubt standard.⁴²

So, within Liscow's framework, a tort rule that broadly requires arsonists to pay compensation for property they destroy is rich-biased; tweaking the rule to benefit poor arsonists and disadvantage rich ones could render it neutral.⁴³ But this is a highly counterintuitive understanding of legal neutrality. In the examples above, rich-bias seems more akin to what most people would consider a neutral baseline, given that the legal rule was not selected to benefit the rich and applies in a nondiscriminatory fashion.

Again, I am not arguing that we adopt an alternative neutral baseline. I instead question whether we should frame our analysis in terms of rich-bias or poor-bias at all. Neutrality is important in Liscow's framework because neutral wealth-maximizing policies are presumptively acceptable and should be retained. In contrast, rich-biased wealth-maximizing policies are acceptable only if their distributional consequences will actually be offset by additional taxes and transfers.⁴⁴

But if current tax-and-transfer is insufficient and if offsets will not occur, there is no reason to distinguish between neutral and rich-biased policies. If policies are to be selected with redistribution in mind, we should also favor poor-biased policies over neutral policies, and very poor-biased policies over policies that are only mildly poor-biased. The designation of some policies as poor-biased and others as rich-biased should not affect our analysis at all; only the fundamental tradeoff between efficiency and redistribution should matter.

IV. WHY EFFICIENCY IS NOT BIASED: A MORE GENERAL CRITIQUE

Another way that efficiency might be rich-biased is in comparison to normative theories based on fairness or equity that produce results more favorable to the poor. Bias in this sense usually means bias in comparison to the speaker's preferred alternative theory.

This formulation of bias is obviously problematic, since it begs the question by assuming the speaker's preferred normative theory as the neutral

⁴² Alternatively, a multiplier could be applied to civil or criminal penalties imposed on the rich. This could be analogous to Finland's well-known system of income-adjusted fines.

⁴³ The specific tweak to legal standard proposed above may be insufficient to render the rule neutral; or it could be excessive, rendering the rule poor-biased. However, there is theoretically *some* level of advantage that could be given to the poor in order to render the rule neutral overall, whether it is a 55 percent burden of proof, 60 percent, and so forth.

⁴⁴ Liscow, "Is Efficiency Biased?" 1695.

baseline. The neutral baseline might be more convincing if there were some compelling and universally agreed-upon alternative to efficiency—but there is not. Philosophical theories that emphasize fairness over efficiency diverge from each other in a variety of ways. A critical legal theorist's vision of fairness may favor the poor relative to efficiency analysis; but a minimal-state libertarian's vision of fairness may favor the rich relative to efficiency analysis.

A more minimal formulation of this claim might be that Kaplow and Shavell's elevation of efficiency is specifically rich-biased. Here, the alternative is more obvious. Kaplow and Shavell proposed that redistribution should not be conducted through legal policy; so the alternative is to redistribute through legal policy. But this formulation, too, merely buries the main disagreement. A core claim by Kaplow and Shavell is that focusing on efficiency through legal policy leaves everyone better off, including the poor. This might not be true for a variety of empirical reasons, but any rich-bias is not by design. Any claim of bias therefore must substantively engage with the assumptions of Kaplow and Shavell's model, rather than broadly dismissing it as biased.

Again, allegations that efficiency is biased assume a neutral baseline that tends to hide assumptions rather than clarify them. Instead of focusing on whether a particular framework is or is not biased, we should ask whether it substantively fails or succeeds. We can consider the many important and well-known objections to the promiscuous invocation of efficiency in law, without throwing out efficiency analysis entirely.

V. CONCLUSION

There is a palpable sense of frustration in writing by some critics of efficiency. Many have argued that supposedly "efficient" rules and regulations increasingly disfavor the poor without any compensation through more progressive taxes and transfers. To skeptics of efficiency, the idea that tax-and-transfer should compensate for redistribution through legal rules feels increasingly like a bait-and-switch.

I sympathize with these concerns; in my own view, there would be little to lose and much to gain by raising taxes on the rich. But disappointment about the political decline of progressive taxation does not impugn the concept of efficiency. Defenders of the tax-and-transfer system are careful to note that they endorse the concept of redistribution itself, but merely disagree on the appropriate methods.⁴⁵ The failure to impose more progressive tax rates may be better understood as a failure of politics or democracy.

Another way to frame criticisms of efficiency is as a meta-intellectual move away from the supposed objectivity of economic analysis, and toward a more honest discussion of political power. On this view, the appeal of

⁴⁵ See, for example, Kaplow and Shavell, *Fairness Versus Welfare*, 35.

efficiency is not analytical but *rhetorical*. Few people, even within academia, have both the training and the interest to engage substantively with the trade-off between efficiency and redistribution. Somewhat like the Coase theorem, the pursuit of efficiency remains a convenient and memorable default rule, even though at this point the exceptions might swallow the rule. Consequently, substantive analysis that chips away at the borders of efficiency's empire may not be enough. A bolder, more sweeping approach is needed to finally dispatch it once and for all.

Critiques of bias and incoherency may be meta-intellectually useful if they discourage excessive reliance on efficiency. Nonetheless, these critiques fail on purely analytical grounds. The priority of tax and transfer can and should be confronted while still acknowledging that efficiency is a useful concept.

Law, University of Minnesota Law School