IS EFFICIENCY BIASED?

Zachary Liscow[*]
August 2017

ABSTRACT: The most common underpinning of economic analysis of the law has long been the goal of efficiency (i.e., choosing policies that maximize people's willingness to pay), as reflected in economic analysis of administrative rulemaking, judicial rules, and proposed legislation. Current thinking is divided on the question whether efficient policies are biased against the poor, which is remarkable given the question's fundamental nature. Some say yes; others, no.

I show that both views are supportable and that the correct answer depends upon the political and economic context and upon the definition of neutrality. Across policies, efficiency-oriented analysis places a strong thumb on the scale in favor of distributing more legal entitlements to the rich than to the poor. Basing analysis on willingness to pay tilts policies toward benefitting the rich over the poor, since the rich tend to be willing to pay more due to their greater resources. But I also categorize different types of polices and show where vigilance against anti-poor bias is warranted and where it is not, with potentially far-reaching implications for the policies that judges, policymakers, and voters should support.

Table of Contents

---

[*] Draft - comments welcome. Yale Law School, Associate Professor. Contact: zachary.liscow@yale.edu.
Thanks to Bruce Ackerman, Anne Alstott, Ian Ayres, Jake Brooks, Conor Clarke, Bob Cooter, Ed Fox, Jacob Goldin, Daniel Herz-Roiphe, Louis Kaplow, Amy Kapczynski, Max Kasy, Al Klevorick, Lewis Kornhauser, Daniel Markovits, Mitch Polinsky, Alex Raskolnikov, Susan Rose-Ackerman, David Schleicher, Alan Schwartz, Matt Stephenson, Judge Stephen Williams, Gui Woolston, and participants at the George Mason Manne Forum, Georgetown Law and Economics Workshop, Columbia Tax Policy Workshop, and American Law and Economics Association Annual Meetings for helpful comments. Thanks to Daniel Giraldo, Brian Highsmith, Quentin Karpilow, Michael Loughlin, and Kate Tian for excellent research assistance.

INTRODUCTION

Suppose that a city is considering building neighborhood parks, each of which costs $1 million to build. The residents of a rich neighborhood are willing to pay $2 million for the park, but the residents of a poor neighborhood are only willing to pay $500,000 for the park, less than the cost of building it. Suppose as well that the park increases well-being to the rich and poor by the same amount. Should the city build a park in the rich neighborhood, the poor neighborhood, both, or neither?

The typical economic analysis of law gives a clear answer: build the park in the rich neighborhood, but not the poor neighborhood. Doing so would be efficient. This is the kind of analysis deployed in torts, cost-benefit analysis, environmental justice, labor regulation, trade policy, minimum wage policy, and a host of other areas. This Article discusses whether such decision-making is biased against the poor.

The economic analysis of policy commonly has a goal of maximizing efficiency.[1] Around the modern advent of the dominance of the efficiency norm in the economic analysis of law in the 1970s and 1980s, there was vigorous debate about whether efficient policies are biased against the poor.[2] Remarkably, this foundational question about the most common goal in the economic analysis of law, if not in all analysis of law, remains unresolved. Rather, the question largely went into hibernation because it was deemed inconsequential: if the tax system achieves the appropriate distribution of income, then the distributive impacts of non-tax policies[3] do not matter. Instead, the argument goes, efficient policies should be used to grow the pie as large as possible, after which tax policy can slice the pie equitably.[4]

---

[1] For commonly used textbooks taking this view, see, for example, RICHARD A. POSNER, ECONOMIC ANALYSIS OF LAW (15-20) (9th ed. 2014) [hereinafter "POSNER, EAL"]; STEVEN SHAVELL, FOUNDATIONS OF ECONOMIC ANALYSIS OF LAW 2-3 (2004) (describing social welfare as the normative basis for analysis in law and economics, but then restricting attention to efficiency by excluding analysis on the distribution of utilities) [hereinafter "SHAVELL, FEAL"]; ROBERT COOTER & THOMAS ULEN, LAW & ECONOMICS 7-8 (6th ed. 2012) (saying that the book "will focus on efficiency rather than distribution" in analyzing the law because of the availability of the tax system for redistribution). Of course, law and economics long precedes the work of Richard Posner. See, for example, the work of Coase in the 1950s and John Commons in the 1920s.

[2] On the side arguing bias: Lucian Bebchuk, *The Pursuit of a Bigger Pie: Can Everyone Expect a Bigger Slice?*, 8 HOFSTRA L. REV. 671 (1980); C. Edwin Baker, The Ideology of the Economic Analysis of Law, 5 PHIL. & PUB. AFF. 3 (1975); Ronald Dworkin, *Is Wealth a Value?*, 9 J. LEGAL STUD. 191 (1980). On the side arguing against bias: Richard Zerbe & Tyler Scott, *(Almost) Everybody Wins: A True Pareto Justification for Practical Welfare Economics and Benefit-Cost Analysis* (U. Wash. School of Law Legal Stud. Res. Working Paper Series, 2016), *available at* https://ssrn.com/abstract=2466101; John R. Hicks, *The Rehabilitation of Consumers' Surplus*, 8 REV. ECON STUD. 108 (1941); Harold Hotelling, *The General Welfare in Relation to Problems of Taxation and of Railway and Utility Rates*, 6 ECONOMETRICA 242 (1938).

[3] Section III.A carefully discusses the distinction between non-tax and tax policies.

[4] Louis Kaplow & Steven Shavell, *Why the Legal System Is Less Efficient Than the Income Tax in Redistributing Income*, 23 J. LEGAL STUD. 667 (1994); A. MITCHELL POLINSKY, AN INTRODUCTION TO LAW AND ECONOMICS 127 (2d ed.1989).

In past decades, when income inequality was largely off the radar, this argument was especially appealing. But this Article starts from a different premise, which might be controversial to some and obvious to others: the tax system may not arrive at a welfare-maximizing distribution of income.[5] In this case, the distributive impacts of non-tax policies matters. This Article dusts the cobwebs off of the old debate and asks the question anew: are efficient policies biased against the poor? It answers: yes, severely, but only sometimes. It then turns out to the question *when* are efficient policies biased against the poor? The Article shows when to be and when not to be on guard against systematic bias against the poor in the adoption of efficient policies.

Understanding these claims requires some nuance in understanding what this "efficiency" thing is. When I ask, "Is efficiency biased?," I refer to "Kaldor-Hicks efficiency," the typical definition used in economic analysis of the law. Kaldor-Hicks ("K-H") efficiency maximizes individuals' willingness to pay for a policy change.[6] This goal is associated with scholars like Richard Posner but is a common goal for setting non-tax policies. So, when the critics say that efficient policies are biased against the poor, the view is based on the notion that efficiency is based on "willingness to pay." Since the rich have greater wealth, the view goes, they will tend to have a greater willingness to pay, and therefore policymakers maximizing efficiency will choose policies that benefit the rich over the poor.

This emphasis on efficiency in non-tax policy design presents a puzzle: Although efficiency is the typical goal for the design of non-tax policies, economists' typical ultimate social goal is to maximize not efficiency but rather "social welfare,"[7] often conceived as maximizing the sum of individuals' utilities. One key difference between the two social goals is that the distribution of income plays a very different role under each. All else equal, distributing a resource (say, a selection of vegetables) to a poor person tends to increase social welfare more than distributing the resource to a rich person because of the declining marginal utility of consumption. That is, economists usually assume that the utility of the average rich person is increased less by $1 of incremental consumption (in vegetables or otherwise) than the same $1 increase in consumption for a poor person.[8] K-H efficiency, though, pays little heed to the declining marginal utility of

---

[5] Lee Anne Fennell & Richard H. McAdams, *The Distributive Deficit in Law and Economics*, 100 MINN. L. REV. 1051, 1055 (2016); Richard S. Markovits, *Why Kaplow and Shavell's "Double-Distortion Argument" Articles Are Wrong*, 13 GEORGE MASON L. REV. 511 (2004). *See* Zachary Liscow, *Are Court Orders Sticky? Evidence on Distributional Impacts from School Finance Litigation*, 47 J. EMPIRICAL LEGAL STUD. (forthcoming 2017).

[6] For a thorough discussion of Kaldor-Hicks efficiency, sometimes confusingly called "wealth maximization," see Lewis Kornhauser, *Wealth Maximization*, NEW PALGRAVE DICTIONARY OF ECONOMICS AND THE LAW 679 (Peter Newman ed., 1998).

[7] See, for example, the long-standing standard graduate-level microeconomics textbook, ANDREU MAS-COLELL ET AL., MICROECONOMIC THEORY 117-22, 817-50 (1995) [hereinafter "MWG"]. This approach dates to Abram Bergson, *A Reformulation of Certain Aspects of Welfare Economics*, 52 Q. J. ECON. 310 (1938) and was incorporated by Paul Samuelson in his canonical textbook. PAUL SAMUELSON, FOUNDATIONS OF ECONOMIC ANALYSIS 203-56 (1947). Note that MWG does discuss potential Pareto improvements and the "compensation principle" in a brief paragraph, which the authors note can be a welfare measure with quasi-linear utility. MWG at 334. For a philosophical defense of using social welfare functions for evaluating social choices, see MATTHEW D. ADLER, WELL-BEING AND FAIR DISTRIBUTION (2012). For a brief history on the use of social welfare functions, see ADLER at 79-88.

[8] See *infra* note 80 for an analysis of the declining marginal utility of consumption.

consumption (and indeed sometimes uses it to justify lower legal entitlements to the poor). It pays attention only to the willingness to pay for the vegetables and ensuring that they go to the party willing to pay the most, which will tend to be the richer party.

How then to reconcile the recommendation of efficient non-tax policies with an ultimate goal of social welfare maximization? Tax policy. The social-welfare-maximizing means of addressing distributive goals has been taken to be tax policy: the mantra is to have efficient policies that may harm the poor, grow the economic pie as large as possible, and then redistribute to them through taxes.[9] This twostep of efficient non-tax policies and distribution through taxes will often result in the optimal policy. Even this view's most ardent defenders acknowledge that it is not always right on its own terms,[10] but that critique is not the subject of this Article.[11]

Rather, this Article asks: what if the promised redistribution through taxes does not happen? At a time of increasing inequality and after a period of decades in which income gains have gone nearly all to the top of the income distribution,[12] with little redistribution toward the bottom of the income distribution, this is a timely question.[13] And additional questions arise: What is the relationship between efficiency and welfare? Are those efficient policies biased against the poor or are they "neutral" with respect to income? In this way, potential political economy failures create an opening to return to this old debate about whether efficient policies are biased.

---

[9] *See* Kaplow & Shavell, *supra* note 4. The first mathematical statement of this general reasoning is by Aanund Hylland & Richard Zeckhauser, *Distributional Objectives Should Affect Taxes but Not Program Choice or Design*, 81 SCANDINAVIAN J. ECON. 264 (1979).

[10] Kaplow & Shavell, *supra* note 4.

[11] I and others argue elsewhere that redistributing with legal rules can be more efficient than redistributing through taxes. Zachary Liscow, Note, *Reducing Inequality on the Cheap: When Legal Rules Should Consider Equity as Well as Efficiency*, 123 YALE L.J. 2478 (2014). I call the high cost of redistributing though taxes the "one-third rule," since taxes cause about $0.33 of deadweight loss for every dollar raised. This inefficiency raises significant scope for legal rules that are more efficient at redistribution. Also, legal rules may be particularly well-equipped to extract rents. *See* Gerrit De Geest, *Removing Rents: Why the Legal System Is Superior to the Income Tax at Reducing Income Inequality* (Wash. U. in St. Louis Legal Stud. Res. Paper No. 13-10-02, 2013), https://ssrn.com/abstract=2337720. Others argue that legal rules should always redistribute at least a little since the cost of redistributing through various instruments increases with the amount of redistribution. *See* Chris Sanchirico, *Taxes versus Legal Rules as Instruments for Equity: A More Equitable View*, 29 J. LEGAL STUD. 797, 805-06 (2000). *See also* Chris Sanchirico, *Deconstructing the New Efficiency Rationale*, 86 CORNELL L. REV. 1003 (2001). Finally, it may be optimal for legal rules to redistribute because of a variety of tax-gaming responses. David Gamage, *How Should Governments Promote Distributive Justice?: A Framework for Analyzing the Optimal Choice of Tax Instruments*, 68 TAX L. REV 1 (2014). *See also* Matthew Dimick, *Should the Law Do Anything about Economic Inequality?* (SUNY Buffalo Legal Stud. Res. Paper No. 2016-011, 2016), https://ssrn.com/abstract=2723476 (listing current policy changes that could both increase efficiency and redistribute).

[12] THOMAS PIKETTY, CAPITAL IN THE TWENTY-FIRST CENTURY (2014). For example, this line of work started being published over a decade ago with Thomas Piketty & Emmanuel Saez, *Income Inequality in the United States, 1913-1998*, 118 Q. J. ECON. 1 (2003).

[13] *See e.g.*, Fennell & McAdams, *supra* note 5 at 1055 ("Omitting political failure from the analysis *requires accepting* a crucial but rarely articulated claim that we term "the invariance hypothesis": that any political failure that exists for tax-and-transfer must inevitably plague non-tax methods of distribution to at least the same degree" (emphasis added)); ADLER, WELL-BEING AND FAIR DISTRIBUTION at 565. Also, some of the issues covered in this paper are touched on in LEE ANNE FENNELL & RICHARD H. MCADAMS, FAIRNESS IN LAW AND ECONOMICS (2012).

So, to return to the old debate: On one side of the debate, still common among not only academics but also politicians and lay voters is the idea that policymakers should adopt efficient policies to grow the size of the pie.[14] For any given policy, there may be winners or losers, but over time those impacts will "even out," and the result of the efficient policies will be greater growth and greater well-being for everyone. This view should be familiar to anyone who even occasionally reads the news and is associated with comments like "a rising tide lifts all boats" and (among critics) "trickle-down economics."

On the other side of the debate are academics and politicians arguing that efficient policies do not do enough for the poor and may even be biased against them.[15] This view has resonance in the critiques of "neoliberalism" and the "Washington consensus" view that governments should adopt efficient, growth-inducing laws.[16] This view resonates with critics of so-called "trickle-down economics." Without policies aimed to help the poor, the view goes, the poor will lose out.

Which view is correct? Is efficiency biased? I excavate the foundational questions underlying these hoary academic debates before returning to the big policy issues that ultimately hinge on the answers to these questions.

To answer this question, I introduce two neutralities. By "neutrality," I mean that there is no systematic bias in the welfare criterion toward selecting policies that deliver outcomes that particularly benefit the rich or the poor. The first is "welfare weight neutrality." Welfare weight neutrality asks whether policies tend to give equal weight to the welfare of the rich and the poor. My results show – for the first time – the precise, non-neutral nature of the traditional law-and-economics efficiency norm. My analysis asks: What is implicitly maximized when efficient non-tax policies are adopted but taxes do not respond? Answering this question helps identify the degree to which political economy failings matter for welfare. Those failings matter a lot. I start by determining the specific social welfare function that is implied if an economics scholar suggests the adoption of efficient policies, but the true societal goal is maximizing a social welfare function. It is helpful to understand the answer to this question in comparison to a benchmark for social welfare maximization, the unweighted utilitarian social welfare function ("SWF"). That is, precisely how far does efficiency analysis deviate from treating the utility of everyone equally?[17] My general result is that efficiency analysis implicitly gives welfare weight to individuals in inverse proportion to their marginal

---

[14] Leslie Carbone & Jay Richards, *The Economy Hits Home: What Makes the Economy Grow?*, THE HERITAGE FOUNDATION (July 1, 2009), http://www.heritage.org/jobs-and-labor/report/the-economy-hits-home-what-makes-the-economy-grow.

[15] For example, see ARTHUR M. OKUN, EQUALITY AND EFFICIENCY: THE BIG TRADEOFF (1975). Similarly, Matthew D. Adler and Eric A. Posner note that efficiency-based cost-benefit analysis leads to a "bias in favor of wealthy people" because the wealthy generally are willing to pay more for a project. Matthew D. Adler & Eric A. Posner, *Rethinking Cost-Benefit Analysis*, 109 YALE L.J. 165, 183-87 (1999). *See also* MATTHEW D. ADLER & ERIC A. POSNER, NEW FOUNDATIONS OF COST-BENEFIT ANALYSIS 142-46 (2006) (arguing for distributional weights in cost-benefit analysis, but even then not in ways that would lead agencies to intentionally redistribute from the rich to the poor). The critique of economic analysis of the law as advantaging the wealthy has a long genealogy. *See, e.g.*, Baker, *supra* note 2.

[16] ROBERTO UNGER, WHAT SHOULD LEGAL ANALYSIS BECOME? (1996).

[17] Among the gamut of potential SWFs, this is fairly moderate position to take; for example, the maximin SWF (incorporated into the work of John Rawls and also discussed below) places zero weight on anyone but the poorest members of society.

utility of income. Since economists usually assume that poorer people tend to have a higher marginal utility of consumption, efficiency analysis effectively places less welfare weight on the poor because they are poor. So, this result provides grist to the efficiency critics.

But maybe "welfare weight neutrality" is not the right subject. Consider instead "legal entitlement neutrality:" the tendency of a type of policy to assign a larger or smaller amount of legal entitlements (e.g., an entitlement to clean air or road safety) to different individuals on the basis of their income. I do not mean here some platonic ideal of neutrality, a term that is fraught.[18] I do not select this because people think that it is the ultimate end of government policy-making, but rather because many may think that it is a requirement that government policy-making should satisfy. And, in particular, since many hold the view that certain branches of government (often the courts and administrative agencies) should not "redistribute,"[19] redistribution being the exclusive province of the legislature, if efficient rules really do redistribute downward, then such a view would lose its force.

Whether a policy satisfies legal entitlement neutrality depends upon context. I divide policies into three categories: rich-biased (which allocate more legal entitlements to the rich than the poor), neutral (which exhibit no bias), and poor-biased (which allocate more legal entitlements to the poor than the rich). I formally define the categories and offer examples of each. The rich are willing to pay more for some legal entitlements, leading to rich-biased policies. But for some legal entitlements, the rich are not willing to pay more. In particular, everyone has the same willingness to pay for one dollar in increased or decreased income because of economic changes induced by increased international trade: one dollar. A policy that distributes dollars is neutral and exhibits no bias. For a small number of policies (like public bus routes), despite their lower wealth, the poor are willing to pay more, leading to poor-biased policies.

Overall though efficiency analysis places a heavy thumb on the scale in favor of rich-biased policies, since the rich—because of their greater wealth—tend to be willing to pay more for the things that legal entitlements confer. For these policies, rather than allocating resources to the poor, who are most in need—and who have the highest utility increase from a given gain in resources—efficient policies tend to do the *opposite* of what would maximize aggregate utility, allocating resources to the rich, who are willing to pay the most. Since entitlements have value, efficient policies will therefore tend to allocate more valuable legal entitlements to the rich: more spending on parks in rich neighborhoods, more spending on roads and airports that the well-off use rather than public transportation the poor use, cleaner air in rich places than in poor ones. I call this principle the accession principle of law and economics, invoking the property law principle in which ownership of some unclaimed resource is assigned based upon ownership of a related resource. In effect, efficient polices tend imply that greater ownership of wealth entitles individuals to a larger allocation of policy entitlements—

---

[18] Herbert Wechsler, *Toward Neutral Principles of Constitutional Law*, 73 HARV. L. REV. 1 (1959).

[19] *See, e.g.*, Kaplow & Shavell, supra note 4 at 675 (regarding courts); Richard A. Posner, *The Ethical and Political Basis of the Efficiency Norm in Common Law Adjudication*, 8 HOFSTRA L. REV. 487, 503-04 (regarding courts); Posner & Adler, *supra* note 15 at 186 (describing the typical view of the purpose of cost-benefit analysis in regulatory agencies as "separat[ing] out the distributional issue").

even if the rich do not pay for it. That is, rich-biased policies give disproportionate legal entitlements to the rich for free, exacerbating inequality.

But, beyond the overall impacts of efficient policies,. the Article emphasizes how the division of policies into rich-biased, neutral, and poor-biased can guide policymaking. If one has a goal of legal entitlement neutrality in judicial administrative, or even legislative rule-making, the results imply differential vigilance depending on the nature of the rule at hand. For policies on dollars (like trade-induced changes in incomes), policymakers could hold different views on the necessity of compensating losers on a policy-by-policy basis, but there is a supportable argument that policy impacts could even out over time: as a matter of methodology, there is not a bias. For rich-biased policies, however, there is an inherent legal entitlement bias. If one holds the view that policy should maximize welfare but that taxes may not achieve welfare-maximizing distributional impacts, then this result should concern you, since efficient rich-biased policies will tend to increase inequality over time. This reasoning suggests that paying attention to the *institution* producing the policy—and the likelihood that a policy's impacts will result in compensating offsets—may be an important consideration in policy analysis. This Article is primarily descriptive, largely leaving implications for policymakers and economic policy analysts for future work, but this Article does describe the conditions under which bias against the poor is more likely to be present, and set out an agenda for better-understanding those conditions.

This Article begins in Sections I and II with foundations on the meaning of social welfare and efficiency. This material may be (or at least seem) obvious to some, but a return to basics is required for the precision needed to make progress on these foundational questions. Section III links the traditional goal of having efficient non-tax policies with having tax and non-tax policies that jointly maximize social welfare. The Article then departs from that conventional view in Section IV by supposing that in reality taxes do not maximize social welfare, making the distributive impacts of efficient non-tax policies an essential question. Section V shows that efficient policies violate welfare weight neutrality in a precise way. Section VI shows when legal entitlement neutrality is violated and offers examples. The conclusion sketches out implications for future work on distributional concerns in economic policy analysis.

I. SOCIAL WELFARE

Economic analysis usually begins with a social welfare function.[20] I follow the approach taken in the leading graduate microeconomics textbook, *Microeconomic Theory*, by Andreu Mas-Collell, Michael Whinston, and Jerry Green.[21] (Those well-

[20] One famous economics article to invoke the need for a social welfare function to claim "policy implications" is Charles Tiebout's article on local government. He argued, "On the usual economic welfare grounds, municipal integration is justified only if more of any service is forthcoming at the same total cost and without reduction of any other service. A general reduction of costs along with a reduction in one or more of the services provided cannot be justified on economic grounds unless the social welfare function is known." Charles M. Tiebout, *A Pure Theory of Local Expenditures*, 64 J. POL. ECON. 416, 423 (1956).

[21] As of May 2016, Google Scholar counts 9,560 cites of MWG. GOOGLE SCHOLAR, https://scholar.google.com/scholar?q=whinston+green+microeconomic&btnG=&hl=en&as_sdt=0%2C7. I have yet to find someone who received a PhD in economics in the last two decades who did not use this book in the first year of the PhD program; just say "MWG" and virtually any of these economists will know what you are talking about. For a recent through description and defense of social welfare functions, see

versed can skim, but I nevertheless want to be clear about the basics.) The economics social-welfare-function approach typically sets out "a benevolent central authority," or "social planner."[22] In its most general form, a social welfare function considers a society of I individuals facing a set of social alternatives (i.e., combinations of policies) $X = \{x_1, x_2, x_3, \dots, x_N\}$, which includes all the various possible arrangements of wealth, policies, and other "inputs" that affect utility. (Policies are indexed by subscript n.) For example, one social alternative may be adopting President Obama's Affordable Care Act and another may be adopting President Trump's American Health Care Act. Each individual *i* (among the *I* individuals in the society) has a utility function u_i that converts the policy chosen from this set of possible rules and arrangements (x_n) into an individualized measure of utility.[23] Social welfare (W) then depends upon these utilities, aggregated across the individuals in society—which in turn depend upon the policy alternative *x* that is chosen, can be expressed as:[24]
$W(x_n) = W\big(u_1(x_n), \dots, u_I(x_n)\big)$.

I take as my baseline normative goal the standard "unweighted utilitarian" social welfare function, $W(x) = \sum_i u_i(x_n) = u_1(x_n) + u_2(x_n) + \dots + u_I(x_n)$.[25] Under this formulation, social welfare is conceived of as the sum of utility across all individuals (determined by the unique welfare function of each), with equal weight given to the utility measure of each individual in society. The goal then is to choose the alternative *x* that maximizes this sum of individual utilities, where each person's preferences are valued equally.[26]

Note how capacious this social welfare function is.[27] It can take many different forms, and other forms abound. For example, one conventional alternative is the "Maximin" SWF, where social welfare is conceptualized not as the sum of every individual's utility, but rather is measured by reference to the utility of the individual with the lowest (i.e., "Min") utility: $W(u) = Min\{u_1, \dots, u_I\}$.[28] In the Maximin SWF, which

---

Matthew D. Adler, *A Better Calculus for Regulators: From Cost-Benefit Analysis to the Social Welfare Function* (Duke Law School Pub. Law & Legal Theory Series No. 2017-19, 2017), https://ssrn.com/abstract=2923829.

[22] MWG, *supra* note 7 at 117. Interestingly, the dominant textbook before MWG, *Microeconomic Analysis*, begins with the "compensation principle." HAL VARIAN, MICROECONOMIC ANALYSIS 404 (1978). See also Varian on the first and second theorems of welfare economics at 323-36 and welfare analysis at 221.

[23] Put more precisely, $u_i: X \to R$, where *R* denotes a real number. That is, a utility function takes features of the social alternative and yields an amount of utility.

[24] I take a social welfare function similar to that in MWG, *supra* note 7 at 117. I make a slight generalization, applying the SWF to not only wealth but also legal rules, so that non-financial impacts of legal rules can also be considered. I use lecture notes from John Roemer, available from the author upon request.

[25] MWG, *supra* note 7 at 119.

[26] Sometimes "utilitarian" is used interchangeably with "welfarist" (i.e., using a social welfare function). The conventional practice in economics is to follow Jeremy Bentham and John Stuart Mill in restricting "utilitarianism" to the simple adding up of utilities. Since the conventional practice in philosophy is to refer to utilitarianism as the goal of adding up utilities across all members of society (which may be weighted or unweighted, depending on the underlying philosophical theory), I will refer to this utility function as "unweighted utilitarian" to avoid confusion.

[27] Though this particular interpretation is likely outside the central tendency of economics, Sen's interpretation of the utilities as basic capabilities is consistent with this math.

[28] MWG, *supra* note 7 at 827. Emmanuel Saez, for example, has used a similar welfare function in his work. Though economists sometimes refer to this SWF as "Rawlsian," I will refrain from doing so because

some think of as a simplified representation of Rawls's famous difference principle, this one individual's utility is all that matters for social welfare.

Another is the "generalized utilitarian" SWF: $\sum_i g_i(u_i)$.[29] This is often simplified to $\sum_i \omega_i u_i$, which I will call a "weighted utilitarian" SWF. This SWF is similar to the original formulation in that it, unlike the Maximin SWF, conceives of social welfare as some aggregation of individual utilities across society—but represents a departure by assigning different weights to individuals' utility during this aggregation. It thus provides a simple way to weight different individuals' utility based upon whatever factors matter to a social planner. The "welfare weights" ω_i vary by features of the individual i with whom they are associated, based on the individual's income, social status, desert, or whatever other feature the social planner deems appropriate. Section V determines what weights are implied by efficiency analysis.[30]

The reason that the utilitarian welfare function suggests a redistribution of income from the rich to the poor is the widely-held belief in a declining marginal utility of income. To some, the intuition that an individual on the edge of starvation has a much higher marginal utility from an extra dollar to buy a bowl of rice than a billionaire who spends another dollar on his hundredth sports car is sufficient to convince them that the marginal utility of income declines. Indeed, for any given individual, it is relatively easy to observe that the marginal utility of income declines. In particular, the widely-observed phenomenon of risk aversion largely reflects a declining marginal utility of consumption for given individuals. Take the example of insurance against one's house burning down. Why do people take out such insurance? Largely because of the declining marginal utility of consumption.[31] Suppose that you have a $400,000 house. When would you rather have an extra $400,000 from an insurance company: when the house has not burned down or after the house burns down, leaving you $400,000 poorer? Of course, you would rather have the extra $400,000 when your wealth has just been dramatically reduced—your utility of money is higher when you are poorer.

But there is a hitch. Though a declining marginal utility of consumption is widely-assumed in economics, it is not an unproblematic assumption. In particular, making comparisons across individuals requires *interpersonal* comparisons of utility.[32] And,

---

Rawls has two principles of justice, and the first (maximum liberty compatible with equal liberty for all) is lexically prior. Also, the second principle (the "difference principle") sometimes represented as the Maximin SWF also has to coexist with fair equality of opportunity and is a function of "primary goods," rather than utilities.

[29] MWG, *supra* note 7 at 828.

[30] One might be concerned that these models involve unmeasurable parameters, like distributional preferences and concerns about fairness. However, there is growing empirical work on fairness and distributional preferences, making what seemed intractable problems tractable. *See, e.g.*, Ilyana Kuziemko, et al. *How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments*, 105 AM. ECON. REV. 1478 (2015); Emmanuel Saez & Stefanie Stantcheva, *Generalized Social Marginal Welfare Weights for Optimal Tax Theory*, 106 AM. ECON. REV. 24 (2016); Raymond Fisman, Pamela Jakiela, & Shachar Kariv, *The Distributional Preferences of Americans* (Nat'l Bureau of Econ. Research Working Paper No. 20145, 2014), http://www.nber.org/papers/w20145.pdf. Furthermore, plenty of work in normative law and economics involves quantities that are not easily measured. There is little reason to apply the critique that things are difficult to measure to desire for fairness in torts law, but not to the desire for clean air in torts law, if both are important to social goals.

[31] There are of course other partial explanations, including liquidity constraints and consumption commitments.

[32] Utility representations are only unique up to a positive, affine transformation, complicating interpersonal

saying that each individual has a declining marginal utility of income does not necessarily imply that the same is true *across* individuals. The problem is that utility is not observable. So, without further assumptions, it could be the case that, though a typical poor person and a typical rich person each has a declining marginal utility of income, the marginal utility of income is still higher for a particular rich person than a particular poor person if the individual rich person loves consuming so much. However, with restrictions on the utility function, choices can imply a declining marginal utility of income. For example, a declining marginal utility of consumption within individuals implies a declining marginal utility of consumption *between* people if there are not systematic differences in utility functions across income levels (as is the case, for example, if everyone has the same utility function).[33]

Apart from risk aversion, another argument for an interpersonal declining marginal utility of wealth comes from hedonic surveys. These hedonic surveys more closely reassemble an older conception of utility dating back to the 19th century, before the formalization of the connection between choices and utility in the 20th century, in which utility represented well-being rather than an expression of choices and preferences. The recent Nobel Laureate in economics Angus Deaton offers evidence that measures of "life satisfaction . . . increase[s] linearly with the logarithm of income."[34] Of course, nothing here is to say that subjective well-being is the same thing as utility, only to suggest that it is evidence in favor of a declining marginal utility of income. With this brief defense of the assumption of declining marginal utility of income, I will make such an assumption throughout the rest of the paper.

II. EFFICIENCY

Broadly, there are two concepts of "efficiency" used in economics and in law and economics, though the two are often not clearly distinguished. The two are Pareto

---

utility comparisons.

[33] This approach dates back at least to William Vickery in 1945. William Vickrey, *Measuring Marginal Utility by Reactions to Risk*, 13 ECONOMETRICA 319 (1945). And, though many may find it implausible, we cannot rule out systematic differences in utility across the income distribution. Other steps in the logic chain are contestable too. For example, Matthew Rabin has argued that risk attitudes may not come only from the utility-of-wealth function, but rather from behavioral anomalies, at least over modest stakes. Matthew Rabin, *Risk Aversion and Expected-Utility Theory: A Calibration Theorem*, 68 ECONOMETRICA 1281, 1286-87 (2000).

[34] *See* Angus Deaton, *Income, Aging, Health and Well-Being around the World: Evidence from the Gallup World Poll*, *in* RESEARCH FINDINGS IN THE ECONOMICS OF AGING 235, 238 (David A. Wise ed., 2010). Similarly, Betsey Stevenson and Justin Wolfers use data worldwide to measure this relationship. Betsey Stevenson & Justin Wolfers, *Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox*, 2008 BROOKINGS PAPERS ON ECON. ACTIVITY 1 (2008). Their main goal is showing that, contrary to prior cross-country analysis, that greater income increases actually is associated with some discernable increase in happiness *at all*. Their results imply "a logarithmic effect of GDP on subjective well-being with a semi-elasticity of around 0.2 to 0.4." Comment by Alan Krueger, 2008 Brookings Papers on Econ. Activity at 100. That is, with a 10 percent increase in income (i.e., approximately a 1 percent increase in the logarithm of income), subjective well-being increases by between 0.2 percent and 0.4 percent, a very small amount. *See also* Richard Layard, Stephen Nickell & Guy Mayraz, *The Marginal Utility of Income*, 92 J. PUB. ECON. 1846, 1846 (2008) (finding similar results, with an elasticity of happiness with respect to income of approximately -1.26, meaning that, when income doubles, the marginal utility of income goes down by substantially more than half).

efficiency (or Pareto optimality)[35] and Kaldor-Hicks efficiency. A policy is Pareto efficient if there is no alternative policy that makes someone better off without making anyone worse off. Though a Pareto efficient policy is not necessarily the welfare-maximizing policy, since it does not incorporate all the features of a social welfare function (in particular issues of equity), a policy that is Pareto efficient is viewed by economists as an improvement on the status quo. However, Pareto efficiency has often been seen as a criterion that is not very helpful, since for most policies, making *no one* worse off is impossible because of the large number of people involved.[36] Also, and less frequently noted, Pareto efficiency is rarely anyone's theory of justice.

The First and Second Fundamental Welfare Theorems pertain to Pareto efficiency. Though they are foundational parts of economics, they have surprisingly little to do with modern welfare economics, which generally uses social welfare maximization. The First Fundamental Welfare Theorem states that the outcome of a competitive market will be Pareto efficient, assuming some things like perfect competition, complete information, and rational actors.[37] The Second Fundamental Welfare Theorem states that any Pareto optimal outcome can be achieved if costless lump-sum transfers of wealth are available, along with further more technical assumptions.[38] Unfortunately, lump-sum transfers are in practice almost never available, since the realistic tools available to redistribute money are distortionary—that is, they induce the subjects of taxation or recipients of transfers to change how hard they work or how much they save (or alter another behavior, like where one lives), thereby reducing efficiency, all else equal, and leading to something other than a Pareto optimal outcome. This reality radically limits the value of the Second Fundamental Welfare Theorem. Thus, Pareto efficiency is not a criterion for choosing the best or "optimal" policies, but rather a minimum criterion used by economists who are uncomfortable with the greater assumptions required by the social welfare function approach. Since it is not a comprehensive measure of all that a social planner would care about, it is not viewed by economists as suggestive of the correct policies, but rather is viewed as a guide for policy analysis. note

Recognizing this limitation, economists have developed an alternative criterion, Kaldor-Hicks efficiency.[39] Rather than aggregating essentially unobservable utilities as a SWF does, K-H efficiency uses something more observable: willingness to pay. It is a strand sometimes used for policy analysis—but it is the main strand used in law and economics, and is the primary subject of this Article. I refer to K-H efficiency when I refer to "efficiency" or "efficiency analysis," I am referring to K-H efficiency unless otherwise noted. K-H efficiency seeks the arrangement of goods, services, and

---

[35] Pareto superior changes are those that benefit at least someone while harming no one. A Pareto optimal or Pareto efficient outcome is one that has no more Pareto superior changes left to make.

[36] *See, e.g.*, Guido Calabresi, *The Pointlessness of Pareto: Carrying Coase Further*, 100 YALE L.J. 1211 (1991).

[37] MWG, *supra* note 7 at 308. Another important, but more esoteric, assumption is that of "compete markets," which requires that all goods (including all possible insurance contracts) are available and available in perfectly competitive markets.

[38] MWG *supra* note 7 at 308. In particular, in addition to the assumptions needed for the First Fundamental Theorem, household preferences and firm production technologies must have certain properties (they must be "convex").

[39] The criterion may be described a variety of ways. DANIEL MARKOVITS, CONTRACT LAW AND LEGAL METHODS 25-26 (2012).

externalities that the free market would achieve, taking the current wealth distribution as given.[40] In that way, efficiency analysis seeks to maximize total consumer and producer surplus.

To put it more precisely, the following notation defines K-H efficiency. Just as with a social welfare function, there are I individuals indexed by i. Again there is a set of social alternatives $X = \{x_1, x_2, x_3, \dots, x_N\}$.[42] K-H efficiency asks which of these policies creates the greatest amount of aggregate "social surplus," denoted for each individual by $s_i(x)$. K-H efficiency adds up the surplus for each individual and chooses the policy x_n with the highest sum; that is, it chooses the policy that satisfies $\max_x \sum_i s_i(x)$. The criterion is also sometimes called "potential Pareto efficiency" because it is viewed as identifying changes that increase overall surplus and thus have the "potential" to be Pareto efficient after transfers from those who gain from the policy change to those who lose from it.

But what exactly is surplus? Surplus measures how much a policy increases the "expenditure function," denoted $e_i(v)$, which measures the the smallest amount that an individual needs to spend to achieve utility level v. In other words, the expenditure function, and thus surplus, is a measure of how much people are willing to pay for a policy change. For example, suppose that, under the status quo policy, (1) an individual achieves a utility level of 13 and spends \$100 to achieve that, (2) he would achieve a utility level of 15 with a new government policy x, and (3) it would take \$150 of spending to achieve a utility level of 15. Then $e_i(13) = \$100$, $e_i(15) = \$150$, and $s_i(x) = \$50$. Surplus thus measures how the expenditure function changes under different possible policies.[44] It measures the amount that individual i would have to pay or be paid to make him indifferent between the status quo and the alternative policy. In this example, the individual would be willing to pay \$50 to shift to the new policy. This amount will be unique to each individual, as determined by his utility function and income. Note that I have said nothing about how social surplus relates to social welfare; Section IV explains that connection.

Put a different way, efficiency analysis seeks to maximize willingness to pay and thereby achieve the allocation of goods, services, and externalities that would result if everyone who lost from the new policy were compensated (as they would be in a frictionless free market, where all changes are Pareto improvements), whether or not the compensation actually takes place. That is, even if one party is harmed by a policy change, if another party would be willing to pay more to have the policy change than the harmed party would be willing to pay to avoid the policy change, the policy change is

---

[40] One oddity of Kaldor-Hicks efficiency is the so-called Scitovsky paradox, in which the efficient outcome depends upon whether the wealth distribution used is that before or after a change in legal rules. *See* Tibor de Scitovsky, *A Note on Welfare Propositions in Economics*, 9 REV. ECON. STUD. 77 (1941). Posner acknowledges this paradox. EAL, *supra* note 1 at 15.

[42] I follow the notation of Nathaniel Hendren, *The Inequality Deflator: Interpersonal Comparisons without a Social Welfare Function* at 12 (Nat'l Bureau of Econ. Res. Working Paper No. 20351, 2014), http://www.nber.org/papers/w20351.pdf. For one commonly-used exposition, see ANGUS DEATON & JOHN MUELBAUER, ECONOMICS AND CONSUMER BEHAVIOR (1980).

[44] More precisely, surplus compares the expenditure function evaluated at the utility under the alternative policy $v_i^a(x)$ with the expenditure function evaluated at the status quo utility level v_i^0, or $s_i = e_i(v_i^a) - e_i(v_i^0)$.

determined to be worth adopting—regardless of whether there is actually a transfer from the beneficiary to the harmed party. That way, the total amount that people are willing to pay in aggregate for policies in the world has increased—and, therefore, "wealth" has increased.

With the weighted utilitarian SWF, there are a variety of ways that social welfare maximization can differ from efficiency analysis.[45] For my purposes here, the most important way is that allocating money, goods, entitlements, or other forms of wealth to individuals with low wealth may increase utility more than allocating these forms of wealth to individuals with high wealth because of the declining marginal utility of wealth, a conventional assumption in economics. Efficiency analysis, in contrast, does not directly consider the declining marginal utility of wealth. Similarly, a social planner may place different welfare weights ω_i on individuals depending upon their wealth, with either more or less weight depending on how wealthy an individual is. As I will show below, since wealthier individuals tend to be able to pay more for desirable things, under some circumstances, efficiency will tend to allocate *more*—rather than less—to wealthier individuals.

III. THE ROOT OF THE LAW AND ECONOMICS APPROACH: EFFICIENCY AS WELFARE MAXIMIZATION

Law and economics has generally taken the approach of maximizing Kaldor-Hicks efficiency rather than social welfare.[46] The main argument in favor of efficiency used in law and economics is that efficient policies actually do promote social welfare maximization.[47] Essentially all the leading law and economics textbooks make an argument along these lines.[48]

---

[45] For example, some argue that "horizontal equity" across individuals of the same income is a concern that can be reflected in the welfare weights. One example of horizontal equity is that someone may "deserve" compensation because he has been harmed, generating a higher welfare weight for that individual.

[46] In a thoughtful discussion aimed at first-year law students, Daniel Markovits notes that economic analysis of the law using economic efficiency seeks to "avoid the quagmire of resolving moral pluralism on the merits" by appealing to preference satisfaction. Markovits, *supra* note 39 at 21. For a very helpful introduction to economic efficiency, see Markovits, *supra* note 39 at 19-28.

[47] Richard Posner provides another important argument for the efficiency norm in law and economics. In the late 1970s, Judge Posner articulated a distinct reason for seeking to maximize efficiency with legal rules: that efficiency (i.e., wealth maximization) is itself a normative goal. Richard A. Posner, *Utilitarianism, Economics, and Legal Theory*, 8 J. LEGAL STUD. 103 (1979). Posner found it to be an appealing mix of Pareto efficiency and utilitarianism, without the downsides of either. It lacked the practical difficulties of the Pareto criterion. Likewise, to Posner, wealth maximization lacked the problems like "utility monsters" and failure to consider individual autonomy present in utilitarianism. Yet, the argument went, wealth maximization retained the positive elements of Pareto efficiency and utilitarianism. By simulating what the market would produce, wealth maximization respected individual autonomy and was a practical means of analysis based on individuals' willingness to pay. I will not revisit the debate on the merits of the Posner's justification for wealth maximization, in part because that has already been extensively argued. *See* Anthony T. Kronman, *Wealth Maximization as a Normative Principle*, J. LEGAL STUD. 227 (1980); Dworkin, *supra* note 2. As well, it is not clear how much Posner himself supports the argument anymore. *See* Richard A. Posner, *Wealth Maximization Revisited*, 2 NOTRE DAME J.L. ETHICS & PUB. POL'Y 85 (1985). For example, Shavell claims that Posner "has since adopted instead other social goals (which he labels pragmatic)." SHAVELL, FEAL at 667.

[48] *See* COOTER & ULEN, *supra* note 1 at 8; A. MITCHELL POLINSKY, AN INTRODUCTION TO LAW AND ECONOMICS 10 (4th ed. 2011) ("efficiency should be the principal criterion for evaluating the legal system [since] it is often impossible to redistribute income through the choice of legal rules and that, even when it

A. Tax vs. Non-Tax Policies

Before explaining, I need to address a basic definitional question that may cause confusion: what are "taxes" and what are "non-tax policies," when the discussion is one about whether "non-tax policies" (sometimes also called "legal rules") should be efficient? The appropriate definition can be obscured in the debates, partly because the breadth of the claim of efficiency's supremacy has rarely been made apparent. As well, sometimes scholars make claims about the appropriateness of deviating from "efficient" policies in "private law"[49] or in arenas outside of the legislature.[50] However, the debate in which I am engaging is not one that depends upon the institution considering the efficient policy: the considerations of bias apply in legislatures, administrative agencies, and courts—concerning public and private law. Of course, the institution under consideration could matter tremendously for the appropriateness of adopting an efficient rule—indeed, that's a key point in this paper—but the only question I am asking at this point is: are efficient non-tax policies adopted on their own biased or not?

In the context of the debate about efficiency then, the appropriate definition must be driven by the claim about the supremacy of efficiency. Of course, the appropriate definition of these terms would be different in a different context. For example, if one is asking what the Internal Revenue Service is best-equipped to handle, different definitions may be appropriate. And, even in this specific context, defining the two terms is not easy. I start by abstracting away from taxes on corporate profits, consumption, wealth, or savings, and focus only on the personal income tax, which is normally the tax considered as the alternative to non-tax policies. I also abstract away from issues of defining the tax-filing unit relating to the treatment of individuals, families, and children, which are not relevant for the present discussion.

If one goes back to the original theory, the claim about efficiency, as explained in the next subsection, is that all policies outside of cash transfers based on income, should be efficient.[51] This claim is very broad. By this claim, government provision of health care, carbon taxes, torts policies, minimum wage laws, and the mortgage interest deduction should all be efficient. So, in general, "taxes" are *cash* payments[52] between the government and tax-paying units that depend solely upon *income*. "Transfers" are also a type of taxes because some government payments like welfare (e.g., Supplemental Security Income and Aid to Families with Dependent Children) are cash payments on the basis of income. For the purposes of this paper, taxes are used to do two things: 1) transfer income between individuals of different incomes (i.e., establish the optimal income tax before a proposed policy change) and 2) offset the consequences of a policy change (relative to the status quo ante) across the income distribution. Of course, such a definition leads to many difficult line-drawing problems. Indeed, programs like Aid to

---

is possible, redistribution through the government's tax and transfers system may be cheaper and is likely to be more precise."). *See supra* note 1.

[49] COOTER & ULEN, supra note 1 at 7-8.

[50] *Id*; Aditi Bagchi, *Distributive Injustice and Private Law*, 60 Hastings L.J. 105 (2008) (describing distributive justice concerns in the sphere of private law).

[51] Hylland & Zeckhauser, *supra* note 9 at 264 (1979); RICHARD MUSGRAVE, THE THEORY OF PUBLIC FINANCE (1959).

[52] Technically, a tithe on a farmer could be the equivalent of cash. This is in fact a subject of dispute in the Supreme Court. *See Horne v. Dep't of Agric.*, 135 S. Ct. 2419 (2015).

Families with Dependent Children depends not only on income but also on having dependent children. But, as noted, I am abstracting away from some problems by ignoring family-related issues, like "marriage penalties" and "marriage bonuses," and from other bases of taxation like wealth, which can affect eligibility for some cash payments for low-income individuals.

Putting aside ambiguous issues though, anything that is non-cash is a non-tax policy. For example, spending on a public park, antitrust policy, and trade policy are non-tax policies. As well, spending through the tax code is generally best-considered a non-tax policy. The mortgage interest deduction is best-considered a non-tax policy, not a tax policy, since it is a payment that, although cash, depends upon how much one spends on a mortgage and is not designed to compensate for the distributional effects of a policy change. Similarly, a carbon tax is a non-tax policy, since how much one pays does not depend solely upon one's income and is not intended to compensate for a policy change; rather, it depends upon how much carbon one emits. To see this a different way, consider that a carbon tax can have more or less equivalent impacts on pollution to a cap-and-trade scheme and to a command-and-control scheme limiting pollution. All three forms of the policy should be treated the same for the purpose of this debate about the appropriateness of deviating from "efficient" non-tax policies.[53] Similarly, spending on education is a non-tax policy, since it is not a cash transfer to a tax-paying unit. The minimum wage also is a non-tax policy, since it does not involve cash transfers with the government.

B. Assumptions Underlying the Conventional Logic

The idea that policies other than taxes should not take into account distributional effects is longstanding and has an impressive list of proponents, including Nobel laureate Paul Samuelson[54], foundational scholar of modern public finance Richard Musgrave[55], and leading law and economics scholars Louis Kaplow[56] and Steven Shavell.[57] The classic argument for this result in law and economics comes from Kaplow and Shavell. They argue that it is more efficient for redistribution to take place through taxes. In their model, they consider the case of a jury deciding how much a defendant should pay a plaintiff after causing harm—and, in particular, whether that amount should depend upon the relative income of the parties. They argue that the amount should not depend on

---

[53] This is not to say that there are not gray areas, such as the child tax credit, which can be viewed alternatively as an encouragement for having more children or as a way of merely transferring money to families with more mouths to feed. Under the first conception, it is a non-tax policy; under the latter, it is a tax policy.

[54] Paul Samuelson, *The Pure Theory of Public Expenditure*, 36 REV. OF ECON. & STAT. 387 (1954).

[55] MUSGRAVE, *supra* note 51.

[56] For example, Louis Kaplow has written:

> It is usually best to use instruments that are most directly related to the matter in question. In the case of redistribution, the income tax system (including cash transfer) is that instrument . . . . Policy tools other than the income tax system tend to be advantageous in pursuing distributive and revenue objectives only when they are able to address particular shortcomings of a more direct approach, such as by mitigating evasion or reducing the labor-leisure distortion in various subtle ways. . . . Although administrative and political considerations often require substantial deviations from what would otherwise be ideal, it is difficult to reach a sensible accommodation without first obtaining a reasonably comprehensive view of the relevant options.

LOUIS KAPLOW, THEORY OF TAXATION AND PUBLIC FINANCE (2011) at 21.

[57] Kaplow & Shavell, *supra* note 4.

relative incomes—and that the amount should equal the monetary value of the harm caused to the plaintiff—because of what they call "the double distortion." Consider what would happen if the plaintiff were poor and the defendant rich and the jury required compensation beyond the monetary value of the harm suffered. First, future well-off defendants would be "too cautious," knowing that, if ever sued, they would have to pay greater than the harm that they actually caused; this would reduce efficiency, since the monetary value of the distortion to potential defendants' behavior would exceed the benefit gained by potential plaintiffs who would not have to suffer the harm. Second, since this variation from the actual harm caused is a result of relative income levels, the entire population is effectively subject to an income tax. The more income that one makes, the less he would be able to recover as a plaintiff and the more he would be forced to pay out as a defendant. That is a non-tax policy with the effect of an income tax, which would have the same distorting effect as a tax imposed through the federal Internal Revenue Code. So, given the goal of redistributing, it is more efficient to do so through the tax code, which would result in just the second distortion, without distorting tort-causing behavior.

To lay observers of policy, a more familiar example of this argument comes from trade policy. The long-time refrain from economists of (nearly) all stripes has been that countries should adopt free trade, notwithstanding potentially negative impacts on the poor since trade increases the size of the economic pie, and those gains can be redistributed to the poor through taxes and transfers. The Kaplow-Shavell torts example is the same thing: everyone can be made better off, the argument goes, through efficient non-tax policies, plus taxes and transfers.

An implicit political economy assumption obviously underlies this analysis: that those taxes and transfers actually happen, indeed that the political system will achieve an optimal distribution of income (what I call "distributive optimality").[58] Distributive optimality itself consists of two parts.[59] First is the "tax-offset assumption"—that taxes adjust to compensate those who lose from policy changes.[60] For example, an efficient non-tax policy change that creates a lot of income for the rich, but reduces income for the poor, could cause a lot of harm to welfare. Because of the declining marginal utility of income, overall efficiency can increase at the same time as welfare goes down if income decreases among those with a high marginal utility of income. If taxes and transfers do not respond after a change in a non-tax policy to reoptimize the social welfare function in light of the policy's distributional consequences, then social welfare may not be improved, let alone optimized, by the adoption of an efficient policy. And, to be clear, an assumption does not need to be perfect for it to be a good one. For example, though

---

[58] Implicitly assuming distributive optimality has a long lineage. As Richard Musgrave wrote in one of the foundational texts of public economics, "Our normative model of public economy is not designed to be realistic in the sense of describing what goes on in capitals of the world. Rather, it is designed to show what would go in if an optimal result were achieved." MUSGRAVE, *supra* note 51.

[59] Note that the distributive optimality assumption is a necessary condition for the general claim that legal rules should *always* be efficient to maximize welfare. In any particular case, the assumption may not be required. For example, if there is one policy option, then that policy is most efficient, least efficient, welfare-maximizing, and welfare-minimizing all at once.

[60] The tax-offset assumption is related to but distinct from the "invariance hypothesis" of Fennell and McAdams, "that any political failure that exists for tax-and-transfer must inevitably plague non-tax methods of distribution to at least the same degree." Fennell & McAdams, *supra* note 5 at 1055.

short run distributional impacts matter, it need not be the case that the distribution of income is reoptimized immediately. So, putting aside the argument that social welfare matters in the short run as policies adjust, not seeing immediate or even fast adjustments need not give serious reservations about efficiency-based analysis of policies.[61] The second part of the distributive optimality is the "tax optimality assumption"—that taxes are set optimally before the policy change. If those are not set optimally, then using other tools to redistribute can increase welfare.

To be clear, few explicitly assert that the distributive optimality assumption actually is true. The more common explicit claim in canonical texts is that taxes *should* be used, rather than that they *are* used—a normative claim rather than a positive claim.[62] But law and economics analysis that recommends efficient policies de facto makes that assumption implicitly; if the tax-offset assumption does not hold, then the efficient policy would no longer be the optimal (i.e., welfare-maximizing) recommendation. Sometimes law-and-economics analysis merely describes the efficient policy, without recommending its adoption, and thus technically does not make the distributive optimality assumption. But even there, law-and-economics scholars leave a lot of the empirical and theoretical machinery of economics—that could be used to promote welfare—on the table.

IV. A DIFFERENT POLITICAL ECONOMY ASSUMPTION

The tax-offset assumption represents a prediction about how the political system will respond to changes in non-tax policies.[63] Efficiency analysis imagines, either implicitly or explicitly, a single planner (or a majority coalition), choosing frictionlessly from among infinite policy alternatives to optimize according to some given welfare function, which is sometimes theorized as approximating voter preferences through the democratic process. But as a number of scholars have pointed out[64], this may be an inaccurate description of our politics. For example, Lee Anne Fennell and Richard McAdams have recently drawn attention to the "political action costs" of policy change—the diverse and extensive set of obstacles and impediments that parties and political actors encounter in attempting to achieve their desired distributive outcomes, a concept akin to "transaction costs."[65] As Jacob Hacker and Paul Pierson have noted: "In our fragmented

---

[61] As Louis Kaplow notes, "There may exist a sort of political equilibrium regarding the extent of redistribution. Thus, there may be a tendency for policies—perhaps not individually, but taken as a whole over a period of time—to be implemented in a distribution-neutral fashion." Both are political economy conditions, but I do not frame the assumption in terms of comparative "political failures," but rather whether in terms of whether taxes respond empirically to policy changes. He gives the example of the Tax Reform Act of 1986. KAPLOW, *supra* note 56 at 32.

[62] See references at *supra* note 1.

[63] *See* Fennell & McAdams, *supra* note 5 at 1071: "[U]niversal tax superiority can logically coexist with true distributive agnosticism only if one assumes that any distributive pattern that is achievable at all can actually be achieved through the tax system"; Markovits, *supra* note 5 at 597 (noting that the K-S efficiency claims rely on the assumption that "the tax policy that they argue is superior is politically available").

[64] *See, e.g.*, Markovits, *supra* note 5 at 597 ("KS' recommendation that government decision makers not adjust legal rules in an economically-inefficient way to effectuate distributional norms that could be effectuated less-economically-inefficiently and more desirably through taxes is based on the assumption that the superior tax policies in question are politically available.")

[65] Fennell & McAdams, *supra* note 5 at 1109 ("Just as positive transaction costs introduce the possibility of divergence between the existing allocation of resources and the efficient allocation, positive political

political system . . . [s]truggles over policy—over what the government actually does for and to its citizens—are usually long, hard slogs."[66] Indeed, the wide and growing literature documenting the interaction of rising party polarization[67] with the multiplicity of veto points in our political system[68] suggests that gridlock—rather than the perfect responsiveness assumed under typical law-and-economics analysis—is likely to more accurately describe legislative dynamics we observe in practice.[69] And it is hardly controversial to say that some policies are more politically feasible than others, regardless of the policy merits.

It is well beyond the scope of this Article to "prove" one way or another whether the political system gravitates toward the optimal distribution of income. For my purposes, it is enough to suggest that the government may not arrive at an optimally distributional tax system. For example, in one piece of evidence inconsistent with arriving at an optimally distributional system of taxes and spending, I show how state court orders to increase school funding, largely for the poor, change the distribution of taxes and spending even decades afterwards.[70] They do not reflect any "optimal" distribution of taxes and spending toward which legislatures move. Rather, the institution of the judiciary matters.

Another piece of evidence comes from tax rates, which estimates suggest are far below the utility-maximizing rate. For the utilitarian social welfare function, the most rigorous estimates, by leading economist Emmanuel Saez, show that, with a utilitarian welfare function, tax rates are not nearly redistributive enough.[71] Saez's analysis

---

action costs introduce the possibility of divergence between the existing distributive results and those that would maximize welfare."). *See also id*. at 1056 ("Welfarists working in law and economics should give the role of political action costs in sustaining distributive deficits attention on a par with that already given to the role of transaction costs in impeding efficient results").

[66] Jacob S. Hacker & Paul Pierson, *Winner-Take-All Politics: Public Policy, Political Organization, and the Precipitous Rise of Top Incomes in the United States*, POL. & SOC'Y 152, 172.

[67] *See* VOTEVIEW, *The Polarization of the Congressional Parties* (Mar. 2015), http://voteview.com/political_polarization_2014.htm ("By this measure [DW-NOMINATE scores] polarization is now at a post-Reconstruction high in the House and Senate"). *See also* Hacker & Pierson, *supra* note 66 at 171 (describing the "increasing polarization of the two major political parties, which has fostered partisan stalemate even on issues that once featured cross-party bipartisan coalitions").

[68] *See* Jack M. Balkin, *The Last Days of Disco: Why the American Political System is Dysfunctional*, 94 B.U. L. Rev. 1159, 1200 (2014) ("Today . . . there are more blocking points than in the original constitutional design. Perhaps more important, when combined with today's highly polarized political parties, veto points that once promoted bargaining and compromise now produce intransigence and gridlock."); *see also* Sarah A. Binder, *Polarized We Govern? in* GOVERNING IN A POLARIZED AGE: ELECTIONS, PARTIES, AND POLITICAL REPRESENTATION IN AMERICA 223 (Alan S. Gerber & Erich Schickler eds., 2017) ("Divided government—aided by parties' influence over the content of the floor agenda—empowers the opposition party to block agenda issues they oppose").

[69] *See* Binder, *supra* note 68. (finding that levels of legislative gridlock have steadily risen over the past half-century).

[70] Liscow, *supra* note 5.

[71] Emmanuel Saez, *Using Elasticities to Derive Optimal Income Tax Rates*, 68 REV. ECON. STUD. 205 (2001). *See also* Thomas Piketty, Emmanuel Saez & Stefanie Stantcheva, *Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities*, 6.1 AM. ECON. J.: ECON. POL. 230 (2014) (conducting a similar exercise for top incomes, but adding in bargaining over compensation, which increases optimal top income tax rates yet further). *But see* Alan J. Auerbach, Laurence J. Kotlikoff & Darryl R. Koehler, *U.S. Inequality, Fiscal Progressivity, and Work Disincentives: An Intragenerational Accounting*, (Nat'l Bureau of Econ. Res. Working Paper No. 22032, 2016) (finding that across the lifecycle, there is less inequality in spending

suggests that, if the goal is utilitarian, the highest marginal tax rate on income should be between 69 and 81 percent, depending on how much labor supply is distorted by taxes.[72] (Incidentally, most Americans agree that the wealth distribution is unfair and favor heavy taxes on the rich to rectify the situation.[73]) Of course, the work of one prominent economist does not mean that taxes are insufficiently redistributive. That work depends upon a particular normative criterion (an unweighted utilitarian social welfare function with a particular set of utility functions) that is not universally shared. But, given this analysis—and the political economy arguments suggesting that the evidence makes sense[74]—it seems at least a supportable assumption that taxes are not optimally redistributionary.

Another example that has recently received a lot of attention is increasing post-tax income inequality.[75] Analyzing these changes over time does not require the assumption of a particular social welfare function, like Saez's simulations. What has happened from the early 1980s to the present is that both taxes and pre-tax incomes have become more regressive: with a larger share of income earned by the rich and less of it taxed away with time. Even studies that account for in-kind transfers like health care still find substantially increasing inequality over time, though somewhat less than studies ignoring in-kind transfers.[76] This widely-discussed phenomenon is difficult to square with the notion that society gravitates toward a certain optimal distribution of taxes and spending, since the distribution has changed so much over time. It is possible that preferences for redistribution have changed dramatically over time, justifying the big shift. And it is possible that the nature of the economy has changed, making it more difficult to tax the rich. But they are also consistent with a different story: preferences for redistribution have been relatively stable,[77] but politics has intervened. Institutions may reflect long lags and great inertia. Americans may have become particularly hostile to taxes as a means of redistribution. One could tell many stories—and in particular many stories consistent with the idea that the distribution of non-tax policy affects the ultimate distribution of taxes and spending.

In light of this evidence, there is at minimum a plausible case that taxes won't achieve an optimal degree of redistribution—that progressive or regressive distributive impacts of a non-tax policy will not be offset and that in any case the overall system may be out of alignment with optimal income. I do not intend take a stand either way in this debate about whether or not taxes are "optimal," but rather to note that there is good

---

power across the income distribution than within a year).

[72] Saez, *supra* note 71 at 225. I take the numbers for the optimal non-linear tax, incorporating income effects, Table 2: Panel B.

[73] Frank Newport, *Americans Continue to Say U.S. Wealth Distribution Is Unfair* (May 4, 2015), http://www.gallup.com/poll/182987/americans-continue-say-wealth-distribution-unfair.aspx (finding that 63% in 2015 found the distribution of wealth unfair and 52% favoring redistributing the wealth with heavy taxes on the rich).

[74] See especially the work of Fennell & McAdams, *supra* note 5.

[75] Some of these facts of course have been laid out elsewhere. *See, e.g.*, Fennell & McAdams, *supra* note 5 at 1079-80.

[76] Richard V. Burkhauser et al., *Recent Trends in Top Income Shares in the USA: Reconciling Estimates from March CPS and IRS Tax Return Data*, 94 REV. ECON & STAT. 371 (2012).

[77] Vivekinan Ashok, Ilyana Kuziemko, & Ebonya Washington, *Support for Redistribution in an Age of Rising Inequality: New Stylized Facts and Some Tentative Explanations*, BROOKINGS PAPERS ON ECON. ACTIVITY (2015).

reason to believe that the tax-offset assumption may not hold and that taxes may not be optimal. The rest of the Article works with this at least plausible political economy constraint.

V. WELFARE NEUTRALITY

In this section, I ask: if distributive optimality fails, but an efficient policy is still adopted, what does that mean for welfare? This provides an answer to the question of whether efficient non-tax policies satisfy welfare neutrality. It turns out that there is a precise formulation for this non-neutrality. This bias implies that typical law-and-economics models may lead to quite perverse policy recommendations. I first adopt an alternative assumption. Instead of taxes responding completely to the distributional consequences to changes in non-tax policies, I assume that they do not change at all. I then model the consequences for social welfare of adopting efficient policies, as law-and-economics scholars generally advocate, under this alternative assumption. I am, of course, not the first to question this assumption[78], but I believe that I am the first to formalize the implications of making this different assumption.

In particular, if social welfare maximization is society's goal, it is instructive to see what the welfare weights would be if K-H efficiency is the criterion used to develop policy recommendations but a weighted sum of utilities is the ultimate goal. I start by developing some economic machinery using basic economic theory. This result is very general: it pertains to any policy. I start with a population of individuals who consume subject to a budget constraint; that is, they have to consume less than the amount that they earn, minus taxes that depend upon their income and plus an individual-specific transfer. I then compare the utility for each consumer under two potential policies and the amount that the consumer would have to spend to achieve each of those two levels of utility. Comparing these two amounts gives the "equivalent variation," or the monetized gain for switching policies. And, maximizing the equivalent variation is precisely what maximizing K-H efficiency means. Having found (at a very abstract level) the policy that is K-H efficient, I next find the welfare weights in a weighted utilitarian SWF such that the K-H efficient policy would be that chosen by a SWF-maximizing social planner. Thus far, I have been discussing how a social planner can maximize K-H efficiency. I now find the implied welfare weights for each individual associated with this policy. We tend to think about welfare weights as coming exogenously to the model—being specified by a social planner's preferences or reflecting society's collective preferences. In this case, by contrast, I derive the welfare weights that are *implied* by a social planner's decision to maximize efficiency, with a proof of surprising simplicity—so simple in fact that I have it directly in the main text.

Recall the definition of efficiency from the setup in Section II. Maximizing surplus across individuals maximizes efficiency, and the Kaldor-Hicks criterion requires that, for any marginal policy change, the sum of equivalent variation (that is, "surplus") across individuals is weakly positive: $\sum_i s_i \geq 0$. Since the equivalent variation is equal to the change in the value of the expenditure function accompanying a policy change,[79] the Kaldor-Hicks criterion is the same as the requirement that

[78] See especially Fennell & McAdams, *supra* note 5.

[79] This statement is literally true only for marginal policy changes. For a non-marginal policy change, the expenditure function changes due to wealth effects. For utility functions without wealth effects, the

$$\sum_i de_i \geq 0, \quad (1)$$

where $de_i$is the change in the expenditure function resulting from the policy change.

What's the parallel requirement for a social welfare function? Consider the weighted social welfare function $\sum_i \omega_i u_i$, where ω_i is the welfare weight for each individual. This function aggregates utility across all individuals in society, after applying the varying weights given to each individual. A marginal change in policy increases social welfare when $\sum_i \omega_i du_i \geq 0$, where du_i is the marginal change in utility for an individual because of a policy change. Multiplying by de_i/de_i yields the requirement that

$$\sum_i \omega_i \frac{du_i}{de_i} de_i \geq 0. \quad (2)$$

We can readily see that (1) and (2) are the same for any set of de_i's if and only if $\omega_i \frac{du_i}{de_i} = 1$. That is, $\omega_i = 1/\frac{du_i}{de_i}$. Since $\frac{du_i}{de_i}$ is the marginal utility of consumption (i.e., how utility changes with a given change of consumption), we have the following result:

*Efficiency-maximizing policies are consistent with implicit welfare weights in a weighted utilitarian social welfare function equal to the inverse of the marginal utility of income facing each individual***.**

Assuming that the marginal utility of income declines as incomes go up,[80] this result implies that to rationalize using efficiency as a criterion for public policy, *poorer individuals must receive less welfare weight*.

To put a fine point on it, suppose that utility is logarithmic in income, something that scholars of subjective (i.e., survey-based) measures of well-being find plausible and also a common assumption in economics.[81] Data show a 99th percentile income of roughly $400,000 a median income of roughly $50,000, and a 20th percentile income of roughly $20,000.[82] In the special case of logarithmic utility—though one that is arguably a reasonable approximation of utility—individuals are weighted in the social welfare function precisely in proportion to their income. So, since the 99th percentile income is 20 times the 20th percentile income, the 99th percentile earner is weighted 20 times as much, as shown in Table 1.

---

approximation of discussing marginal policy changes does not matter, since the expenditure function does not change with more wealth.

[80] See the discussion in Section II about the declining marginal utility of income.

[81] See the discussion in Section II well-being surveys.

[82] For the last two, U.S. CENSUS BUREAU, INCOME AND POVERTY IN THE UNITED STATES: 2014 tbl. A-2 (2015), https://www.census.gov/hhes/www/income/data/historical/inequality/Table%20IE-1.pdf.

Table 1: Implicit Welfare Weights at Various Incomes with Efficient Policies and No Tax Offset

Income percentile	Income level	Ratio of social welfare weight versus the 20th percentile earner
20th	$20,000	1
50th	$50,000	2.5
99th	$400,000	20

Notes: Assumes that utility is logarithmic.

The intuition is simple: a welfare function with equal weights leads to policies that distribute resources toward the poor because the poor have a higher marginal utility of income, making the marginal contribution of a dollar toward the poor (all else equal) more welfare-enhancing. Yet, the Kaldor-Hicks criterion ignores this declining marginal utility of income (while paying attention to the efficiency of allocating resources to those who value them the most). So, the only rationalization for a social welfare function that effectively "ignores" the declining marginal utility of income is to have welfare weights that are the inverse of the marginal utility of income.

VI. LEGAL ENTITLEMENT NEUTRALITY

The last section related efficiency and welfare maximization in a precise way and showed that efficiency is not neutral with respect to welfare. It is systematically biased in favor of the wealthy in a specific way. But others may think that welfare neutrality is not the is not the proper inquiry and that non-tax policies, at least those in the hands of judges or administrative rule-makers, ought not be concerned with redistribution. That is, some might want judges and administrative rule-makers to be neutral with respect to the rich and the poor. For this, I ask whether efficient non-tax policies satisfy "legal entitlement neutrality:" whether a policy has a tendency assign a more legal entitlements (including goods, services, or money) to different individuals on the basis of their income. That is, for which type of legal rules is there a tendency to allocate legal entitlements that are worth more to the rich? Some may find this an important minimum threshold that courts and agencies should satisfy because if tax policy may not gravitate toward an optimal distribution, then systematically regressive policies would exacerbate income inequality. This section shows that the answer to this question turns crucially on the type of policy under consideration.

I start with a mathematical definition of "legal entitlement neutrality." I define "legal entitlement neutrality" as meaning that, as one's income increases, efficiency-maximizing policies are no more likely to systematically endow one with legal entitlements. In other words, legal entitlement neutrality means that the sign of:

$$\frac{\partial\left(\frac{\partial s}{\partial x}\right)}{\partial c}$$

equals zero, where (as described above) s is surplus (the measure of Kaldor-Hicks efficiency), x is some policy variable (conceptualized here as the provision of a good or service), and c is one's consumption or income (which are equivalent and thus interchangeable in this one-period model, since people consume all their income). That is, does providing good or service x provide more surplus s as one's income c increases?

If so, then that is a "rich-biased" good: since the rich are willing to pay more for it, efficiency-oriented analysis will endow well-off individuals with more of it than poor individuals. In contrast, if the poor are willing to pay more for it, so that surplus for the provision of x increases as income *decreases*, then the good is "poor-biased." In between are "neutral" goods, where income does not impact the surplus from provision of the good.

As the Appendix shows, with some simplifications, I get the following result:

A policy is rich-biased depends if and only if:

$$\frac{\partial^2 u}{\partial x \partial c} > \frac{\partial^2 u}{\partial c^2}$$

The formula then boils down to a simple comparison: whether, as an individual's consumption increases, the marginal utility from policy x decreases more rapidly than that of income. If the marginal utility of policy x decreases less rapidly than the utility of income, then the good is rich-biased—for example, if the utility of the good (say, clean air) stays constant with income $\left(\frac{\partial^2 u}{\partial x \partial c} = 0\right)$ and there is a declining marginal utility of income $\left(\frac{\partial^2 u}{\partial c^2} < 0\right)$. If the inequality goes the other way, the good is poor-biased. And, if the two terms are equal, the good is neutral. In other words, goods that poor people really want—that they get so much utility from that they are willing to pay more for them than rich people, despite their greater poverty—are poor-biased. Otherwise, goods are rich-biased or neutral. The intuition for the result is as follows: Efficiency-based legal entitlement allocations are based on willingness to pay. Suppose for simplicity that the utility of something stays constant with income (e.g., both rich and poor people may value clean air the same amount in utility terms)—in other words, the first term equals 0. Then the willingness to pay for something increases with income as long as the utility from yet an extra dollar of consumption goes down as income increases—that is, as long as there is a declining marginal utility of consumption (i.e., the second term is negative). But other times, one's utility from having something does not stay constant with income and instead declines with income. In those cases, when the utility of the having the thing declines rapidly enough, willingness to pay can stay constant or even decline with income. As this section explains, efficiency analysis places a heavy thumb on the scales in favor of likely rich-biased policies. This section offers examples of each type of policy in turn and then returns to the generalization of legal entitlement neutrality.

A. Rich-Biased Policies

Rich-biased policies endow more of a legal entitlement to the rich than to the poor. As this subsection shows, these policies are pervasive. The example here is from environmental justice—the government is allocating clean air across rich and poor communities—but the key feature is how the government policy is valued by rich and poor people: it is valued *the same*. That is, it is valued the same *in utility terms*. In dollar terms, though, the rich value it more and are willing to pay more for it. As a result, when the rich and the poor value something the same amount in utility terms—and even when the poor value something considerably more than the rich—the policy is rich-biased. Such policies risk exacerbating income inequality.

In particular, suppose that a policymaker is deciding where to shut down some polluting factories. As might happen in a situation like this, there is no practical way to compensate those who are harmed by pollution with the tax-and-transfer system. Suppose that there are two communities of an equal number of individuals that are identical except that those in Richtown each have $9 of income and those in Poortown have only $1 of income.[83] Suppose further that each has the utility function $u = \log(x) + \log(c)$, where c is the amount that individuals consume and x is how clean the environment is. This logarithmic utility function is a standard assumption in the economics public finance literature and receives support from hedonic surveys of income and happiness.[84]

This environment is rich-biased. The reason is that clean air is equally valuable to rich and poor people $\left(\frac{\partial^2 u}{\partial x \partial c} = 0\right)$ and there is a declining marginal utility of consumption $\left(\frac{\partial^2 u}{\partial c^2} < 0\right)$. These are sensible assumptions in this context: that rich and poor people get roughly the same utility gains from clean air and that an extra dollar of consumption is worth more to the poor than to the rich. This example walks through how a policymaker maximizing K-H efficiency would allocate clean air under these circumstances.

Suppose that the policymaker has 10 units of "cleanliness" (i.e., the opposite of pollution) to allocate because of a new technological development. The "status quo" policy is that Richtown and Poortown have 1 unit of cleanliness. (Initially, the environment is very polluted.) Applying the K-H methodology described earlier of maximizing surplus yields an allocation of 0 units of cleanliness for the poor and all 10 units of cleanliness for the rich.[85] This is the allocation that maximizes parties' willingness to pay. Consumption has a declining marginal utility because of the logarithmic way consumption enters the utility function; so does cleanliness. And, since the residents of Richtown do not value the marginal unit of consumption very much (because they are already consuming so much), and they have significant financial resources, they are willing and able to buy all of the clean air.[86]

Following the general result in the previous section, suppose instead that the policymaker implemented the solution to maximize efficiency, but that the actual goal was a weighted social welfare function $\sum_i \omega_i u_i$. What would the implicit welfare weights be? It turns out that the only way to rationalize the unequal distribution of pollution is to put at least nine times more weight on residents of Richtown than residents of Poortown. (Of course, even more relative weight on the rich would also justify this allocation.) As I showed above generally, the welfare weights are inversely proportional to the marginal utility of consumption, which means that the weights are proportionate to income in this specific case of the logarithmic functional form. In other words, to take a social welfare function with one representative individual from each community, the social welfare function implicitly used by efficiency analysis is $W = 9u_{Richtown} + 1u_{Poortown}$. This

[83] Also assume that individuals are immobile.
[84] See Deaton, *supra* note 34; Stevenson & Wolfers, *supra* note 34.
[85] This is the allocation that maximizes the sum of transfers that makes parties indifferent between the world with and without the pollution cleanup.
[86] An alternative way of setting up the problem would be to consider the willingness to pay for a fixed amount of pollution, which could reside in either community. A similar result would obtain; Richtown would be willing to pay more to avoid the pollution, and the pollution would be sent to Poortown.

social welfare function is not inconsistent with any principle of economics, but represents a theory of justice to which I suspect few people, economists or otherwise, would subscribe.

Thus far in this subsection, I have inferred social welfare weights from a chosen efficient outcome. Now, I reverse the reasoning and consider four allocations—the efficient allocation in addition to three others—in light of an unweighted utilitarian SWF $\sum_i u_i$, which treats the utility of everyone identically. First, consider the outcome if the policymaker is maximizing this SWF and no trading in pollution is allowed. In this case, the pollution would be evenly split between the two communities. This is because the rich and the poor each have the same utility function and the same initial levels of pollution at the start of the thought experiment, so pollution has the same effect on the utility of both types of individuals. Thus, any additional unit of pollution on any individual already subject to the same level of pollution affects all the individuals the same.

In even greater contrast with the case of efficiency, consider the allocation of pollution if the social goal is an unweighted utilitarian SWF and the pollution rights can be traded in a Coasean fashion, switching from the implicit assumption earlier that the rights could not be traded. (The assumption did not matter for the efficiency analysis, since no trading would take place after the allocation anyway, being a condition of K-H efficiency.) Now, those units of cleanliness are convertible into money and the marginal utility of income starts to matter. With the logarithmic utility function, the marginal utility with respect to consumption is $1/c$, meaning that the marginal utility of a dollar of income for the poor person is 1 versus just 1/9 for the rich person. In other words, the marginal utility of income is nine times as high for the residents of Poortown as for Richtown. That ratio is the inverse of the welfare weights implicit in the efficiency-maximizing policymaker's behavior. It turns out that allocating 9.8 units of cleanliness to the poor and 0.2 to the rich maximizes the unweighted utilitarian SWF. With a price of \$0.83 per unit of cleanliness (see the Appendix for the derivation), the residents of Poortown sell 4.8 units of their entitlement to cleanliness to the residents of Richtown for \$4, yielding complete equality in cleanliness (5 units each) and in consumption (also \$5 each).

Finally, consider an even allocation with trading. By fiat, each person receives 5 units of cleanliness. Because the poor have so little consumption, they are willing to trade some of their cleanliness to the rich at \$0.83 a unit. As a result, the poor end up with 3 units of cleanliness and \$3 of consumption, and the rich end up with 7 units of cleanliness and \$7 in consumption.

Table 2: Total Utility with Various Allocations of Cleanliness

	Allocation of Cleanliness			Veil of Ignorance:
	Poor	Rich	Total Utility	% WTP to Avoid Efficient Allocation
Efficient allocation	0	10	2.00	0%
SWF-maximizing allocation (no trading)	5	5	2.51	45%
SWF-maximizing allocation (with trading)	9.8	0.2	2.95	67%
Even allocation (with trading)	5	5	2.80	61%

Table 2 lists the sum of utilities under the four options discussed: an efficient allocation of pollution, the SWF-maximization allocation with no trading, the SWF-maximizing allocation with trading, and the even allocation with trading. It shows how perverse the efficient policy can be if the goal is utilitarian and there are no tax-and-transfer-offsets. In each of the four policies, the table lists the total utility, which is the sum of individual utilities for the unweighted utilitarian social welfare function. While utility does not have cardinal meaning, there are large differences in total utility among the three options. The efficient allocation has the lowest utility, since both consumption and cleanliness are highly unequal, and the individuals have a declining marginal utility from both—meaning that (holding total cleanliness and consumption fixed) moving either consumption or cleanliness to the less-well-off party increases utility. Utility increases with the unweighted utilitarian SWF-maximizing solution without trading because at least the distribution of cleanliness becomes equal. And it increases further with the SWF-maximizing solution with trading because both cleanliness and pollution are equally distributed. Even under the even allocation with trading—something not explicitly "redistributionary"—the total utility is substantially higher than under the efficient allocation, since at least the high-marginal-utility party is receiving an even share of the cleanliness. Of course, taxes and transfers could be used to achieve similar results, but this example shows what happens if taxes and transfers are not used, as plausibly may be the case.[87]

The rightmost column gives a cardinal meaning to these differences in utility. I suppose that each person is behind a veil of ignorance with $5 of income each and ask how much each person would be willing to pay to have a 50 percent chance of being rich and a 50 percent chance of poor in each of the three alternatives instead of the efficient allocation.[88] The differences are large; an efficient allocation is not a good approximation of the SWF-maximizing allocation. The individuals behind the veil of ignorance would be willing to pay 45 percent of their income to be certain to have an equal share of cleanliness regardless of their income, 67 percent of their income for equality in income and cleanliness as a result of a disproportionate endowment to the poor party, and 61 percent for an even allocation with trading allowed.[89]

---

[87] Also, note that distributing legal entitlements can be a more efficient way of reducing inequality than using taxes, so that total utility can be higher using legal rules rather than taxes to redistribute. *See* Zachary Liscow, Note, *Reducing Inequality on the Cheap*, 123 YALE L.J. 2134 (2014).

[88] Specifically, I solve for w in $\log(5) - \log(5 - w) = EU_{alternative} - EU_{efficient}$. That is, I solve for the w that constitutes what someone behind the veil of ignorance would be willing to pay to have the expected utility under an alternative regime ($EU_{alternative}$) instead of the expected utility of the efficient regime ($EU_{efficient}$).

[89] These large results partly result from the curvature of the logarithmic utility function. Marginal utility may decline more or less rapidly than this. I use this utility function because studies in happiness show that the effect of money on happiness may have a steep curvature, as money appears to stop having much impact on happiness above a certain level, making logarithm a good approximation. As well, the separability of the utility function into two parts—the logarithm of cleanliness plus the logarithm of consumption—means that the effect of cleanliness on utility does not interact with income. If, for example, high-income individuals like to play golf more than low-income individuals, requiring more time outside and making clean air more important for the utility of the rich than for the poor, then income would interact with the effect of cleanliness on utility. In that case, the effects on utility of the efficient legal rule relative to the SWF-maximizing rule may be reduced. For example, taking the extreme case in which the poor do not value a clean environment at all (and assuming that cleanliness is not tradable), then there would be no

The example illustrates a key point of the paper: policies distribute entitlements (like the right to reduce pollution) that have value.[90] If taxes and transfers do not respond to the adoption of an efficient non-tax policy, then the efficient non-tax policy may not be neutral. The efficient allocation misses an opportunity to use legal entitlements to address existing disparities, as we see in the case of tradability. But, more importantly, when this good is allocated, not only is the declining marginal utility of income ignored, but also the fact that the wealthy tend to have a higher willingness to pay for the good will lead systematically to *more* allocation of the good to the well-off. It actually exacerbates existing inequalities and leads to lower total utility than a "neutral" distribution (like the even split of cleanliness, especially with tradability). So, for this policy, government cost-benefit analyses that follow the efficiency criterion, and are not offset by changes through taxes, will systematically choose policies that increase the utility of the rich more than the utility of the poor.

Figure 1 provides a graphical representation that helps explain what drives these results. Figure 1A shows the relationship between an individual's utility and income—a curve that flattens out as one's income increases. This pollution example involves two types of individuals with different levels of income, each of whom receives the same utility gains from an improvement in environmental quality. But, even if the two types of people have the same utility gains, it would take quite different amounts in dollars to achieve these same utility gains. The y-axis shows equal utility gains for the rich and the poor groups. With dashed lines, the figure then shows the dollar gains that it would take to achieve that level of utility gains for each group. Because of the declining marginal utility of income (i.e., because of the curved line), the amount of income it would take the rich to achieve the same utility gain is much larger. Dollars are "cheap" to the rich, since they already have so many of them; thus, the rich need to receive a lot of dollars for a given utility gain. And this is precisely what drives the results in the example: the rich have a higher "willingness to pay" in dollar terms for the pollution reduction because dollars are cheap to them. As a result, efficiency analysis allocates the pollution reduction to the rich because as Figure 1B shows, the willingness to pay for an allocation of goods goes up with income.

---

difference between the efficient and SWF-maximizing outcomes.

[90] See, e.g., Guido Calabresi & Douglas Melamed, *Property Rules, Liability Rules, and Inalienability: One View of the Cathedral*, 85 HARV. L. REV. 1089 (1972) (making a similar point about the distributional impacts of allocating entitlements).

Figure 1: Relationships with Rich-Biased Policy

Figure 1A: Utility vs. Income

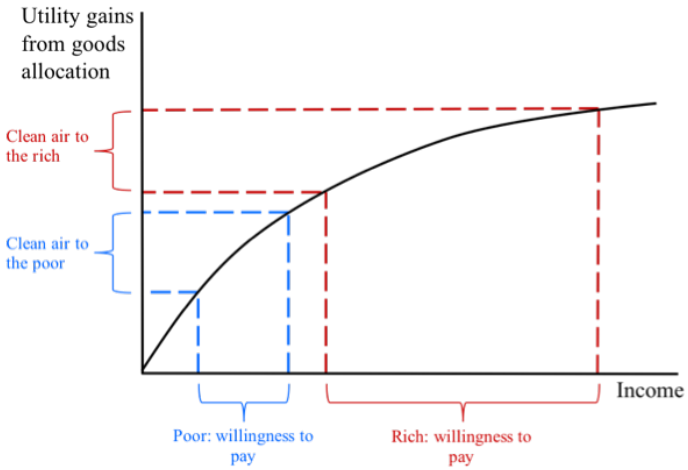

Figure 1B: WTP vs. Income

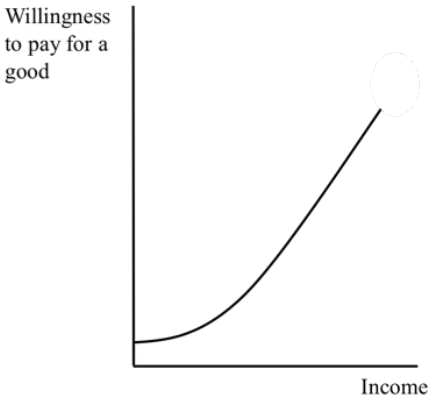

B. Neutral Policies

In contrast to rich-biased policies, neutral policies exhibit no bias toward the rich (or the poor). These policies affect things that the rich and the poor value equally in dollar terms: typically, money itself. For example, such policies might affect the income that someone receives. For groups of these neutral policies, the efficiency-supporters' argument that "growing the pie" holds out the hope of increasing the pie without systematic bias across policies proves tenable. The intuition is that everyone has the same willingness to pay for a dollar: one dollar. So efficient policies that effectively dole out dollars to one group and take them from another group exhibit no systematic bias because there will be no tendency for either the rich or the poor to value these policies more.

My neutral good example is drawn from international trade, which (in the example) operates only by affecting people's income. There are rich people and poor people, and trade affects their incomes differently, either benefitting the rich or the poor. If efficient policies were systematically biased against the poor, as they were in the rich-biased example above, then efficiency analysis would systematically assign greater income gains to the rich than to the poor. However, in contrast to efficient rich-biased policies, this policy is neutral with respect to legal entitlements: a policy that increases incomes for some, decreases incomes for others, but increases incomes overall is no more likely to pass the Kaldor-Hicks criterion if the gains go to the rich than if the gains go to the poor.

In the example, the rich and the poor are identical, except that the rich have greater income and the two are differentially affected by an increase in international trade. Their utility functions are specified as:

$$U_{rich} = \log(c_{rich} + T_{rich})$$
$$U_{poor} = \log(c_{poor} + T_{poor})$$

The rich and the poor have the same utility functions—again, logarithmic as in the previous example. One need not understand the math here to understand the main point: there is still a declining marginal utility of income. The rich and the poor differ in only two ways. First, the poor have less income than the rich: the rich have income c_{rich} and the poor have incomec_{poor}, which is a lower income than that of the rich. Second, there

is a trade policy that changes the incomes of the rich and the poor: changing that of the rich by T_{rich} and that of the poor by T_{poor}. These impacts of trade can be either positive or negative.

The Appendix shows that such a utility function satisfies the definition of a neutral policy. The policy itself just provides income for consumption; thus, consumption and the policy are the *same thing*, and the marginal utility from consumption and the marginal utility from the policy must be the same.

Consider the efficiency-maximizing policy under these circumstances as the gains to trade (T_{rich} plus T_{poor}) change. It should be readily apparent that as long as the net gains to the rich and the poor are positive (that is, $T_{rich} + T_{poor} > 0$), then the policy is efficient. Efficiency asks whether the policy change creates a change that people on net would be willing to pay to achieve, and as long as the gains to one group are larger than the losses to the other, this condition is satisfied. For example, suppose that the policy change benefits the rich by \$100 ($T_{rich} = 100$) and hurts the poor by \$10 ($T_{poor} = -10$). How much would each party be willing to pay for this policy change, and how much would the two parties be willing to pay on net? The rich would be willing to pay \$100 here for the \$100 gain and the poor would require a transfer of \$10. On net then, efficiency increases by \$100 - \$10 = \$90. So this policy increases Kaldor-Hicks efficiency.

Now suppose that the impacts of trade are reversed. That is, suppose that, under the new trade policy, the income of the poor would go up by \$100 and the income of the rich would go down by \$10. Then the same reasoning applies: the willingness to pay of the poor and the rich would be \$100 and -\$10 respectively, and the net efficiency gains are \$90. So, this policy treats the rich and the poor symmetrically.

The neutrality can perhaps be most easily seen by considering the case in which the gains and losses are even and just different from even. Suppose, for example, that the rich gain \$50 and the poor lose \$50 from the trade deal. Then there are no efficiency gains. But suppose that the rich gain \$51 and the poor lose \$49; then the trade deal is efficient. Or suppose that it is the poor who gain \$50 and the rich who lose \$50. There are no efficiency gains then. But, if the poor gain \$51 and the rich lose \$49, then there are efficiency gains. That is, there is no inherent bias against the poor when dealing with a trade deal that impacts incomes. The reason is that the object of government policy-making, incomes, is a matter of dollars. And the poor and the rich value a dollar the same amount: a dollar is worth a dollar.

The same result would not have been true in the case of pollution. Suppose that there are two types of areas, rich and poor, and the government is considering a policy that would clean up pollution in one type of area and create more in the other type. In this circumstance, the same type of neutrality would not hold. A policy that cleaned up pollution by 3 units in the rich area and increased it by 2 in the poor area would be efficient—because the rich are willing to pay more to reduce pollution than the poor would be willing to accept for an increase of 2. But the same outcome may not be efficient if the setup is reversed: a 3-unit pollution reduction in the poor areas and a 2-unit increase in the rich areas. In this case, since the rich are willing to pay more per unit of pollution, it may not be the case that the policy is efficient. With this rich-biased policy of pollution, efficiency is not neutral. Rather, it is "legal entitlement" biased against the poor.

How can one reconcile this distinction between rich-biased policies and neutral policies with efficiency's consistent violation of welfare neutrality? The explanation returns to what efficiency analysis asks, which is: how much is someone willing to pay for a policy change? For a good like environmental quality, the rich are willing to pay a lot for improvements in environmental quality, so the government allocates (for free!) more to them. Put another way, it takes a lot more dollars for the rich to achieve a given utility gain than for the poor to achieve the same utility gain. Dollars are cheap to the rich, and goods policies reflect that. But the same is not true for dollars. For those policies, the rich and the poor value them the same: in dollar terms, a dollar is always worth a dollar. Efficiency's failure to shift (or "redistribute") resources toward the poor makes it biased against the poor in welfare terms. But, for legal entitlement neutrality, neutral policies have no intrinsic bias against the poor.

To understand the results graphically, compare Figure 1 above, where the utility gains to a given policy are the same for the rich and the poor, to Figure 2, but which has the utility gains not from gaining a good but rather from gaining dollars as in the trade example. Figure 2A compares the utility gains for a given dollar gain between the rich and the poor. Because of the declining marginal utility of income, a given dollar gain will result in a smaller utility gain to the rich than to the poor, as reflected on the y axis. However, the same dollar gain will result in the exact same willingness to pay between the rich and the poor, as reflected on the x axis. Thus, as shown in Figure 1B, for a given dollar gain (say, $100), the income of the person receiving the income does not vary the willingness of the person to pay for the dollar gain: a dollar is worth a dollar to everyone.

Figure 2: Relationships with Neutral Policy

Figure 2A: Utility vs. Income

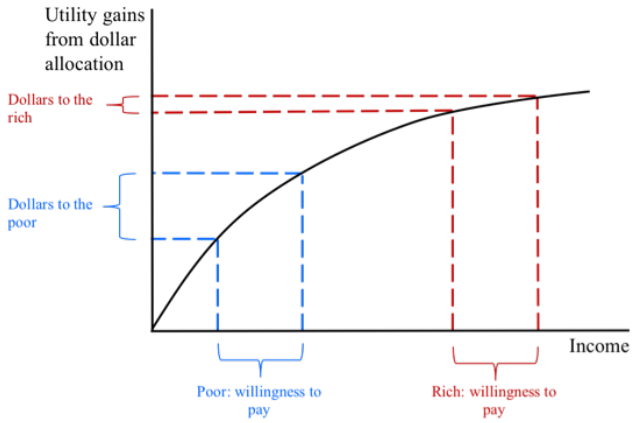

Figure 2B: WTP vs. Income

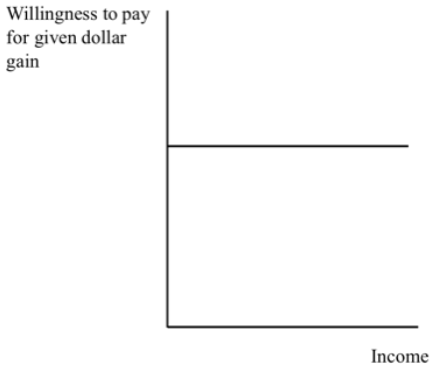

Before turning to poor-biased policies, it is important to note a couple of things about neutral policies. First, how the dollars are spent must be unconditional. For example, allocating vouchers for school may seem like a neutral policy, it is not because the money must be spent on education, and rich and poor people can have a differential willingness to pay for education.

Similarly, just because the result of a legal entitlement is dollar-denominated does not make it a neutral policy; rather, the entitlement itself must be to distribute dollars. For

example, damage payments to those who receive slow responses from 911 calls may be in dollars, but they may also give more dollars to rich than to poor survivors because the legal entitlement is not unconditional cash but rather a right to a timely response to an emergency call, or money compensating one for losses after a failed response. And, since the rich may gain more economic surplus from a fast response (e.g., because their wages are higher and thus face higher economic losses in the case of injury resulting from a slow response time), this policy may be rich-biased.[91] In both the voucher and the 911 call examples, the money was not distributed unconditionally, and thus the policy was not neutral.

C. Poor-Biased Policies

It is not the case that willingness to pay always increases or stays even with income. Rather, some goods become more valuable (in dollar terms) to people as income goes *down*. For these poor-biased polices, efficiency analysis is also biased, but toward *poor* individuals. An example of a poor-biased policy could be spending on bus-based public transit. The logic is relatively straightforward: dollars are still cheaper to the rich than to the poor, but additional income makes an inferior good so much less desirable that—even though the well-off have more money to spend on the good, they don't want to spend the money. For bus-based public transit, the rich gain very little utility from public provision, since they are unlikely to use the service.

One can think of poor-biased policies graphically as those for which there is such an enormous difference in the utility gained by a rich and a poor person that the difference overcomes the declining marginal utility of income—in other words, it overcomes the fact that it is a lot more expensive to pay off a rich person than a poor person in dollars for a given utility gain because of the declining marginal utility of consumption. Figure 3 shows this case. Another way to think about this figure is as the continuum of the spectrum from Figure 1 on rich-biased goods, with equal utility gains for the rich and the poor, through Figure 2 on neutral goods, with much larger utility gains for the poor for the policy but just enough that they compensate for the declining marginal utility of income, and Figure 3, with yet smaller relative utility gains for the rich.

[91] Similar reasoning applies to damages from a tortious car accident. Damages are in dollars, but courts may still award more money to a rich person for having been injured than for a poor person because the lost wages may be higher for the rich and the poor person.

Figure 3: Relationships with Poor-Biased Policies

Figure 3A: Utility vs. Income

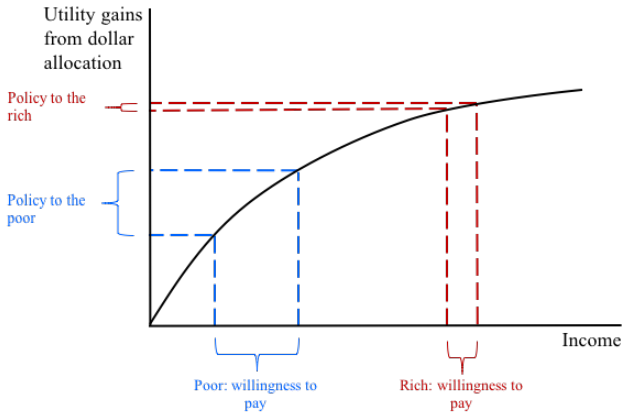

Figure 3B: WTP vs. Income

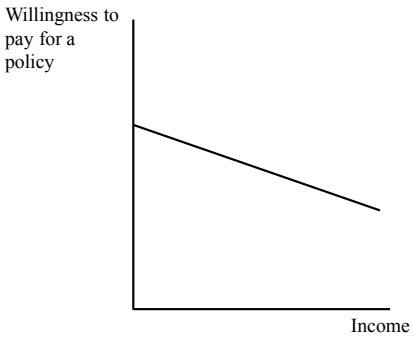

D. Example Policies

If the distribution of income is not optimal or the impacts of policy changes may not be offset—and if neutrality in the sense defined here matters—then these results have important implications for cost-benefit analysis of rules in courts, agencies, and legislatures. In particular, there should be an asymmetry between the treatment of different categories of policies.

If the distribution of taxes and spending may not be optimal, then institutions matter. A court ruling, administrative ruling, or legislative statute that is inefficient but promotes desirable distributional outcomes can improve welfare. How economic policy analysis could proceed under these circumstances is a subject for a different Article. For now, though, it is worth drawing out which types of policies are at risk of having a systematic bias as a matter of methodology. Table 3 gives examples of rich-biased, neutral, and poor-biased policies for each branch of government: executive (through agency rule-making), courts, and legislature.

The table helps clarify the different kinds of policies that are or are not characterized by legal entitlement bias. All of these illustrative policies are far more complex than I imply, but they at least illustrate the point schematically. I take first the rich-biased policies. As described at length above, agencies may decide on siting polluting facilities, creating an inherent bias against low-income individuals. Similarly, local governments choosing which community to spend the most money on for law enforcement or personnel or machines to make voting quicker face an inherent bias: rich people are willing to spend more for a marginal decrease in crime or to spend less time voting. Courts also decide on rich-biased policies—for example, on nuisance suits in which the rich are paid more due to lost wages than the poor are paid. And legislatures decide on many such issues. Legislatures decide whether to build public infrastructure like libraries and may choose to do so in richer neighborhoods that have a higher willingness to pay. Examples also include the parks that begin the paper: the efficient policy is to build more parks in rich neighborhoods than in poor ones. Legislatures also decide between funding bus mass transit (often used by lower-income individuals, who have a lower willingness to pay) or highways (used on average by higher-income individuals): efficient policies will systematically spend more on what rich people want than what poor people want, even if a given amount of spending would increase well-

being for the rich and the poor by the same amount. And Congress decides how much to fund research and development in life-extending cures, which the well-off are willing to pay more for than other priorities that benefit lower-income families.

The middle column illustrates neutral policies. Bureaucrats might decide whether to spend limited research money (e.g., National Science Foundation grants) on research projects that increase the income of lower-wage workers (e.g., making solar panels cheaper and thereby increasing labor demand among low-income workers, or making low-income home health care aides more desirable) or increasing the income among owners of capital (e.g., making offshore wind energy or pharmaceuticals, both capital-intensive activities, more attractive). One can think about this research policy (or at least a simple understanding of it) as one of transferring income from the wealthy owners of companies to lower-income workers. Courts consider similar neutral policies. Employment and labor law decisions help allocate money between the incomes of relatively low-income workers and the profits of the companies owned by relatively high-income individuals. Legislatures decide trade issues, which can allocate money between richer and poorer individuals. These are all neutral policies; there is no systematic bias as an analytical matter toward the rich when adopting them. There is a plausible case that the net effects could even out across these many policies.

The third column contains poor-biased policies. The goal is to think of policies in which an efficient choice would endow a larger legal entitlement to poorer groups than to richer groups. An example for agencies could be a department of transportation deciding which community should receive more spending on busses. Although the well-off might use them occasionally, they are essential for poor people, so an agency might spend more on busses in poorer communities. For courts, it might be efficient to provide subsidies for class-action lawsuits, given the common action problems involved in bringing them. Such subsidies would tend to result in costly avoidance of behaviors by (presumably relatively rich) defendants, so that they can avoid paying damages in class-action lawsuits, which would be easier to bring by the (presumably relatively poor) plaintiffs. For legislatures, adopting an efficient policy on building public swimming pools might spend more in lower-income communities if the wealthy would be reluctant to use the pools, perhaps because the well-off can opt to build their own pools at their homes or in their subdivisions.

Note how poor-biased policies are sub-categories of larger categories from rich-biased policies. Busses are a sub-category of transportation infrastructure, and pools are a sub-category of public amenities infrastructure. That is not a coincidence. In any broad category of efficiency-oriented policy, like transportation infrastructure or public amenities, the rich are likely willing to pay more virtue of their greater income. For example, the rich may not use busses, but they use roads, bridges, and airports—and are willing to pay quite a lot to commute to and from their well-paying jobs and travel on vacations. The poor might be willing to spend more on busses, but not on infrastructure overall. This result makes sense: the rich have more money to spend. Across all categories of goods, they must be willing to pay more—or else the rich would spend less than the poor, which is clearly false. But, the narrower the sub-category of overall consumption, the more likely there will be *something* that the poor are actually willing to spend more on, even if when more aggregated, there are few categories for which that is true. The implications of this phenomenon are explored in the next subsection.

Table 3: Examples of Efficient Policies by Institutional Context

	Rich-biased	Neutral	Poor-biased
Agency	Siting polluting facilities (more on poor, less on rich); Law enforcement or investment in voting (spending on publicly-provided goods like safety in rich vs. poor neighborhoods)	R&D policy on technologies benefitting different kinds of workers	Allocation of spending on busses across communities (lower-income communities may get more)
Court	Nuisance suits (higher lost wages from the rich than the poor)	Employment and labor law (shifting money between workers' incomes and companies' profits)	Subsidies for class-action lawsuits (high-income defendants must adopt costly behavior, low-income defendants benefit)
Legislature	Infrastructure siting (where to place public parks, libraries, etc.); Infrastructure choices (mass transit for poor or highways for the rich); R&D appropriations (healthcare cures the well-off are willing to pay more for)	Trade (shift in incomes of rich vs. poor)	Building public pools (lower-income communities may get more, if wealthy can opt to build private pools)

Policies can have rich-biased, neutral, and poor-biased aspects. For example, workplace safety rules cost (relatively wealthy) employers money (a feature of neutral policies), but benefit (relatively poor) workers through the rich-biased good of workplace safety. And, no doubt, in many of the examples listed above, one could argue with the categorization. The examples here are meant to illustrate how to apply these three categories to the wide—indeed, essentially infinite—range of non-tax policies, as policy analysts, judges, legislators, agency rule-makers, and voters choose whether deviating from efficient policies are likely to be welfare-maximizing or whether there is an inherent bias against the poor that they may find problematic. In particular, the neutral policies are the ones where—if there is a wide range of policies big or small—there is some hope that the distributive consequences may "even out." And poor-biased policies have a thumb on the scale in favor of the poor. On the other hand, the rich-biased policies are subject to an inherent bias against the poor that won't even out across many policies.

E. Unifying Framework: The Accession Principle of Law and Economics

To know whether efficiency is legal-entitlement biased, one can know what kind of policy is being considered: that concerning a rich-biased, neutral, or poor-biased policy. But what can be said about bias in aggregate—across all policies considered by the government? This subsection shows that efficient policies place a large thumb on the scale in favor of rich-biased policies, systematically providing the rich with more legal entitlements than the poor.

I begin by returning to the definition of whether an efficient policy is rich-biased:

$$\frac{\partial\left(\frac{\partial s}{\partial x}\right)}{\partial c} > 0$$

which in turn is determined by whether the following is satisfied:

$$\frac{\partial^2 u}{\partial x \partial c} > \frac{\partial^2 u}{\partial c^2}$$

In other words, the bias of the policy depends upon the comparison of how utility changes with income and the marginal utility of income. The figures portrayed this comparison by showing on the y-axis with brackets the different utility gains of rich and poor people for a policy change and then using the curve for the marginal utility of income to translate those utility gains from a policy change into a willingness to pay.

This inequality makes clear that efficient policies are tilted in favor of rich-biased policies. The reason is simple: the rich get a higher utility from some policies, and poor people get a higher utility from other policies. And, if the question were who gets a higher utility, then policies might be roughly split between those that are rich-biased and poor-biased. But that is not the question. Instead, for a policy to be poor-biased, the extent to which the poor gain more utility than the rich must surpass a big hurdle: the rate at which the utility from the policy goes down with increased income must be even faster than the rate at which utility from income itself goes down with increased income, a high threshold to cross.

To see this this graphically, consider Figure 4, which beneath the axis shows who (the rich or the poor) gains more utility from a policy change and above the axis shows the direction of the bias. On the left half of the figure, the rich gain more utility than the poor: those policies are rich-biased, as the scale above the axis shows. In the middle ("Rich = poor"), utility gains are equal. At that point, the policy exhibits pro-rich bias. To the right of that point, the poor gain more utility than the rich, but the pro-rich bias continues until the utility gains from the policy decline at the same rate as the marginal utility of consumption declines, at which point the policy is neutral. It is only to the right of that point—a narrow portion of the overall spectrum—that there is pro-poor bias. The nature of efficiency is such that it tends to produce outcomes that favor the rich.

Figure 4: Bias and Utility Gains from Policy

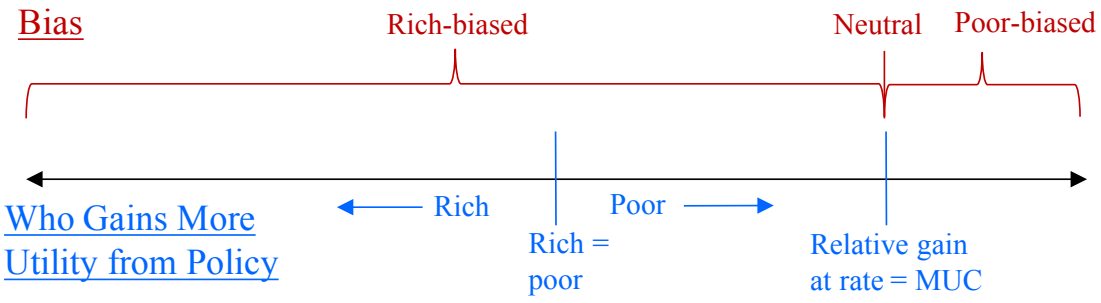

Another way to understand why relatively few policies are poor-biased is that, if most consumption were not of goods that the rich were willing to pay more for, the poor would consume more than the rich, which is not the case. Of course, in principle it is possible that—among government *policies*—there is a disproportionate share of inferior goods, in which case efficient policies would endow the poor with more legal entitlements than the rich. It is impossible to know, but I find this possibility vanishingly unlikely. There are government policies that pertain to goods that the rich are willing to pay less for, of course, like public transportation on busses. But examples are much harder to come by than normal goods, which likely include health care, military defense, parks, environmental amenities, highways, etc. One could dispute these characterizations. For example, perhaps the rich do not need public parks because they have private ones.

Of course, many non-tax policies disproportionately benefit the poor. For example, any means-tested program, like Medicaid, does so. But means-tested programs like Medicaid are typically not driven by efficiency goals, but instead explicitly redistributive ones, and this Article focuses on *efficient policies*. Rather, the relevant policy in a sector like health care for the current discussion would be a program like universal government-provided health insurance (a policy which could in principle be justified on efficiency grounds). That policy would very likely be rich-biased, since a person making $10,000 a year would likely be willing to pay far less for health insurance than someone making $200,000 a year for the same reasons that underlie any rich-biased policy: a rich person gets a lot less utility from a dollar of consumption than a poor person and thus will be willing to spend more money on health insurance.

Yet another way of understanding why poor-biased policies are rare is by thinking about the utility functions that would justify such an outcome. It turns out that ones that economists are familiar with tend to be rich-biased. As the Appendix shows, any "separable" utility function (in which utility consists adding together some function of the policy and some function of consumption) is rich-biased; this follows from the fact that the rich and the poor would have the same utility gains from the policy, thereby making it rich-biased. As the Appendix also shows, Cobb-Douglas ($U = x^{\alpha}c^{1-\alpha}$) and "constant elasticity of substitution" ($U = (\alpha x^{r} + (1-\alpha)c^{r})^{1/r}$) utility functions are also rich-biased. Economists use these utility functions not only because they are relatively convenient but also because they conform with consumer behavior: like declining willingness to pay as quantities increase and a preference for diversity.

Yet another way of understanding the tendency toward pro-rich bias is that people tend to have a declining marginal utility of consuming something: as they gain each extra unit of something, it adds less utility than the previous unit added. It is intuitive that commonly having an extra unit of something causes a bigger decline in marginal utility for *the thing itself* than in the marginal utility of something else (e.g., a government-provided good).

Thus, efficient polices without distributional offsets are systematically regressive in the distribution of entitlements.[92] If there is no wealth transfer that goes along with the adoption of efficient policies, government policies will tend to systematically advantage the well-off. Maximizing efficiency tends to lead to "accession rules." The principle of accession is a doctrine in property law in which "[o]wnership of some unclaimed or contested resource is assigned to the owner of some other resource that has a particularly prominent relationship to the unclaimed or contested resource."[93] For example, by the doctrine of accretion, if land accumulates along a riverbank, the owner of the riverbank generally receives ownership of the land. As Thomas Merrill has pointed out, these doctrines can lead to magnification of wealth inequalities.[94] After all, parties that already own the property that entitles them to the new property are more likely to be well-off. Generalizing this point, efficiency-oriented analysis will tend to reward those that already have income and wealth, leading to what one may call the "accession principle of law and economics": the more you have in income and wealth, the more you get in legal entitlements. The "accession principle" demonstrates the benefit and the costs of assigning new resources based upon existing resources. The benefit is that it may be monstrously inefficient to allocate accreted land along a riverbank to poor individuals who live far away. But, if other policies systematically benefitted the landowners along the river more than the poor landowners far away, then efficient policies would increase inequality.

Of course, if policy changes harm the poor, but the tax-and-transfer system compensates, then that could be the welfare-maximizing regime. However, with some good reason, I am assuming otherwise. If the assumption is indeed false, then the result is that the efficient policies often recommended by law-and-economics analysis are implicitly based on welfare weights that may be consistent with the moral intuitions of very few Americans.[95]

If the proper social goal is actually to maximize an unweighted utilitarian social welfare function and the tax-offset assumption does not hold, then a key question concerns the distribution of costs and benefits from a policy change.[96] And the empirical

---

[92] Also, the result that welfares are weighted precisely by income will depend upon the utility function.

[93] THOMAS W. MERRILL & HENRY E. SMITH, PROPERTY: PRINCIPLES AND POLICIES 161-68 (2007). They describe several doctrines including increase, the doctrine of accession, and accretion.

[94] Thomas W. Merrill, *Accession and Original Ownership*, 1 J. LEGAL ANALYSIS 459 (2009).

[95] For an attempt to infer changes in Americans' distributional preferences from changes in the tax system, see Ben Lockwood & Matthew Weinzierl, *Positive and Normative Judgments Implicit in U.S. Tax Policy, and the Costs of Unequal Growth and Recessions*, 77 J. MONETARY ECON. 30 (2016).

[96] One potential criticism of the analysis is that it seems static—that is, it considers only efficiency at a point in time rather than considering impacts on economic growth. *See, e.g.*, Robert Cooter & Aaron Edlin, *Law and Growth Economics: A Framework for Research* (Berkeley Program in Law and Econ. Working Paper Series, 2011), http://eprints.cdlib.org/uc/item/50t4d0kt. I am engaging in the standard practice with this analysis, and the goal if this paper is not to question that standard practice. Nevertheless, I will touch upon four reasons that the analysis stands even when considering growth. The first is the most basic. A

and theoretical analysis just serves to suggest that there may be an even more important space than previously realized for analysis of the distributive consequences of policy changes in law-and-economics analysis[97] when the goal is utilitarian if taxes and transfers do not offset the distributive consequences of policies.

CONCLUSION: LAW AND ECONOMICS IN AN AGE OF INEQUALITY

When Richard Posner published the *Economic Analysis of the Law* in 1972, law-and-economics scholar Mitchell Polinsky called the book a "potentially defective product," in that "even a valuable product is subject to misuse if proper instructions are not included."[98] In particular, the distributive consequences of policies had to be considered. Despite Polinsky's warning, economic analysis of the law has long been guided by the assumption that the distributive consequences of non-tax policies do not matter, since taxes should respond to take care of distributive considerations. But there is little evidence that taxes in fact do respond. This paper draws out the distributive implications of adopting efficient non-tax policies when taxes are not optimal.

I show that, under many circumstances, efficient non-tax policies are not merely neutral with respect to the distribution of income—rather, systematically efficient policies tend to distribute legal entitlements to the rich, exacerbating income inequalities and possibly leading to multiplication across time. At a time of rising income inequalities and growing concern with these inequalities, as shown by the response to the work of Thomas Piketty, it may be time to consider analysis of non-tax policies that reduce efficiency but have desirable distributional outcomes. Such analysis may be more difficult, but—lacking evidence of an optimally redistributive tax and transfer system—such policies might be what truly maximizes welfare.

If scholars and policymakers are to engage in analysis of policies that considers distributional consequences, this Article suggests the importance of considering context

---

simple redefinition of the question leads to the same results. Instead of thinking about the policy question as pertaining to one point in time, one think about policies across, say, a 100-year period, including economic growth over that period. Then the same analysis goes through, and the same biases exist, just including the growth impacts. Second, even if such a definition were not possible, to benefit the poor, the benefits have to reach the poor—and as much recent research has shown, that has not happened for the most part in decades, either through rising wages or changing economic policies. Thomas Piketty, Emmanuel Saez, & Gabriel Zucman, *Distributional National Accounts: Methods and Estimates for the United States* (Nat'l Bureau of Econ. Res. Working Paper No. 22945, 2016), http://www.nber.org/papers/w22945.pdf, Fig. 2. Third, there is some evidence that people's well-being is in significant part experienced relative to their neighbors. RICHRAD LAYARD, HAPPINESS: LESSONS FROM A NEW SCIENCE (2005). This "keeping up with the Jones's" aspect of utility further limits the scope for growth in income that primarily benefits the rich to increase overall well-being. Finally, growing evidence suggests that inequality itself reduces economic growth, potentially making everyone worse off. In fact, even the International Monetary Fund recently issued a report to that effect. ERA DABLA-NORRIS ET AL., CAUSES AND CONSEQUENCES OF INCOME INEQUALITY: A GLOBAL PERSPECTIVE (2015), *available at* https://www.imf.org/external/pubs/cat/longres.aspx?sk=42986.0.

[97] One response to this paper might be that it gives additional reason for taxes to change to offset the distributional consequences of changes in legal rules. However, I am making a point about the economic analysis of legal rules. That analysis was based on a foundation without any evidence—and now based on a foundation with some contrary evidence. Yes, if taxes responded, it would often be better to have efficient legal rules and offsetting taxes, but alas it appears that this may not be what happens.

[98] A. Mitchell Polinsky, *Economic Analysis as a Potentially Defective Product: A Buyer's Guide to Posner's "Economic Analysis of Law"*, 87 HARV. L. REV. 1655, 1681 (1974).

in deciding whether deviating from the efficient rule will best maximize welfare. A first contextual question is one of political economy: is this a context in which the distributive consequences are likely to be offset? And, is the existing tax code optimal? If either of these is answered in the negative, then one must worry about distributive consequences. Then ask: is this a rich-biased, a neutral, or a poor-biased policy? Since neutral policies may have distributive consequences that even out in the long-term, while rich-biased policies do not, the case for deviating from the efficient rule is stronger for the rich-biased policies.

How policymakers should respond is a longer-term project. One can imagine two extremes: (1) efficient policies and (2) policies that ignore willingness to pay differences driven by income and instead maximize welfare policy-by-policy. But both have unappealing features. Without optimal taxes, efficient policies have a systematic bias against the poor. On the other hand, if there *are* some taxes and transfers, then maximizing welfare policy-by-policy and ignoring willingness to pay driven by income—for example, investing less in airports (used largely by the rich) and more in mass transit (used largely by the poor)—has the problem of missing opportunities in which combining efficient policies with taxes and transfers can make everyone better off. The best policy choice probably lies between these two extremes, and depends upon context.

These contextual elements do powerfully motivate elements of future scholarly research. A first implication is empirical in nature: On what types of policies is tax-offset likely to happen and when is it not? The more promising the prospects for tax offset, the more traditional efficiency analysis is appropriate. A second empirical element of the agenda is determining what policies have significant distributive consequences. My results show no problem with efficient policies that do not have (systematic) distributive consequences. So, where a policy does not have significant distributive consequences, similarly, efficiency analysis is relatively more appropriate. A third empirical task is determining which policies can benefit the poor the most while causing the least loss in efficiency, which is important for maximizing total utility if taxes are not responding to changes in policies.

Likewise, the results raise the urgency of actually implementing offsetting taxes and transfers, which would improve welfare in concert with efficient policies and, if the tax efficiency assumption is accurate, be first-best policies. Failing increased tax offset though, the results suggest a greater scope for law-and-economics analysis that trades off equity and efficiency to complement existing efficiency-minded law-and-economics analysis.

Appendix 1: Legal Entitlement Neutrality

The goal is to sign $\frac{\partial(\partial s/\partial x)}{\partial c}$. I start by supposing that there are two goods, policy variable x and numeraire consumption good c. Further suppose that the policy variable x is entirely government-provided, so that endowment $y = c$ (since the only thing to spend money on is c), so that we can discuss the marginal utility of consumption c and the

marginal utility of income y.[99] Thus, $\frac{\partial(\partial s/\partial x)}{\partial y} = \frac{\partial(\partial s/\partial x)}{\partial c}$. I suppress taxes, since that just adds terms without adding insight. I also suppress prices because I only have a government-provided good and a numeraire good.

To start, recall that s is defined as follows: $s = e(v_i^1)$- $e_i(v_i^0)$. But, we know that $v_i^1 = u(c_i, x_1)$ and $v_i^0 = u(c_i, x_0)$. That is, we can replace utility achieved with utility functions.

Thus, turning first to $\frac{\partial s}{\partial x}$, we know that (for $x_1 > x_0$), for a high-income (h) individual,

$$\frac{\partial s}{\partial x} = e\big(u(c_h, x_1)\big) - e\big(u(c_h, x_0)\big) \cong \left(\frac{de_h}{du} \cdot \frac{\partial u_{c_h}}{\partial x}\right)(x_1 - x_0)\ (1)^{100}$$

and for a low-income (l) individual (for $c_h > c_l$),

$$\frac{\partial s}{\partial x} = e\big(u(c_l, x_1)\big) - e\big(u(c_l, x_0)\big) \cong \left(\frac{de_l}{du} \cdot \frac{\partial u_{c_l}}{\partial x}\right)(x_1 - x_0).\ (2)$$

Subtracting (2) from (1), we get (for a small change in c)

$$\frac{\partial(\partial s/\partial x)}{\partial c} \cong (x_1 - x_0)\left[\frac{de\big(u(c_h, x)\big)}{du} \cdot \frac{\partial u(c_h, x)}{\partial x} - \frac{de\big(u(c_l, x)\big)}{du} \cdot \frac{\partial u(c_l, x)}{\partial x}\right]$$

$$= (x_1 - x_0)\left[\frac{\frac{de\big(u(c, x)\big)}{du} \cdot \frac{\partial u(c, x)}{\partial x}}{\partial c}\right]\ (3)^{101}$$

$$= (x_1 - x_0)\left[\frac{d^2e}{du^2} \cdot \frac{\partial u}{\partial c} \cdot \frac{\partial u}{\partial x} + \frac{de}{du} \cdot \left(\frac{\partial^2 u}{\partial c \partial x}\right)\right].$$

What matters is the sign of $\frac{\partial(\partial s/\partial x)}{\partial c}$, so we can drop $(x_1 - x_0)$ without loss of generality, since we know $x_1 - x_0 > 0$. Thus, we know:

$$sign\left(\frac{\partial(\partial s/\partial x)}{\partial c}\right) = sign\left(\frac{d^2e}{du^2} \cdot \frac{\partial u}{\partial c} \cdot \frac{\partial u}{\partial x} + \frac{de}{du} \cdot \frac{\partial^2 u}{\partial c \partial x}\right).\ (4)$$

Turning away from this line of argument for a moment, we know from price theory that $e(v) = e\big(u(c)\big) = c$. That is, total expenditure will equal one's endowment, which in this case is equal to c. Totally differentiating $e\big(u(c, x)\big) = c$ with respect to c yields

$$\frac{de}{du} \cdot \frac{\partial u}{\partial c} = 1.$$

Totally differentiating with respect to c a second time yields:

---

[99] It could also be the case that x is a variable over which individuals could optimize to be more general, but I leave that case to future work.

[100] For simplicity, throughout the analysis, I suppress the fact that these derivatives are evaluated at $x = x_0$.

[101] Again, for simplicity, I suppress the fact that these derivatives are evaluated at $c = c_l$ and $x = x_0$.

$$\frac{d^2e}{du^2}\cdot\frac{\partial u}{\partial c}\cdot\frac{\partial u}{\partial c}+\frac{de}{du}\cdot\frac{\partial^2 u}{\partial c^2}=0$$

Rearranging gives:

$$\frac{d^2e}{du^2}=-\frac{\frac{de}{du}\cdot\frac{\partial^2 u}{\partial c^2}}{\left(\frac{\partial u}{\partial c}\right)^2} \quad (5)$$

Substituting (5) into (4) gives (after simplification):

$$sign\left(\frac{\partial(\partial s/\partial x)}{\partial c}\right)=sign\left(\frac{\partial^2 u}{\partial c\partial x}-\frac{\partial^2 u}{\partial c^2}\cdot\frac{\partial u/\partial x}{\partial u/\partial c}\right) \quad (6)$$

Formula (6) is the key formula, so I will spend time discussing it. It compares two things: (1) how more consumption c changes the marginal utility of good x $\left(\frac{\partial^2 u}{\partial c\partial x}\right)$ with (2) the slope of the marginal utility of consumption $\left(\frac{\partial^2 u}{\partial c^2}\right)$, or how more consumption changes the marginal utility of consumption, multiplied by the ratio of the marginal utility of good x divided by the marginal utility of consumption. We know that, under typical assumptions, the first term $\left(\frac{\partial^2 u}{\partial c\partial x}\right)$ has an unclear sign, but the second set of terms $\left(\frac{\partial^2 u}{\partial c^2}\cdot\frac{\partial u/\partial x}{\partial u/\partial c}\right)$ is negative, since $\frac{\partial^2 u}{\partial c^2}<0$ by the declining marginal utility of income and $\partial u/\partial x>0$ and $\partial u/\partial c>0$ by nonsatiation.

One way to understand the ratio $\frac{\partial u/\partial x}{\partial u/\partial c}$ is as a normalization. For example, consumption c could be denominated in dollars or cents, and policy x could similarly be denominated in big or small units. This ratio thus provides a normalization of the marginal utility of consumption such that, when multiplied by this ratio, it is in the same units as the cross-partial term $\frac{\partial^2 u}{\partial c\partial x}$.

Broadly, there are three ways that a good is likely to be rich-biased. First, as one would intuitively expect, when income has a more positive effect on the marginal utility of good x, good x is more likely to be rich-biased. Second, when the marginal utility of consumption is diminishing very rapidly (i.e., it is strongly negative), the policy is more likely to be rich-biased because it will take a large money transfer to make up for the utility gains from the policy. Third, when there is a high ratio of utility gains from the policy x versus consumption c, the good is more likely to be rich-biased. Again, this is intuitive, since it will take more money to compensate for the gain of x if the marginal utility of income is lower relative to the marginal utility of the good x.[102]

[102] Interestingly, if we know that the consumer optimizes over z and x before the government intervention, then we know that $\partial u/\partial x=\partial u/\partial c$ (if the price of x and c are the same), as a condition of optimization. I do not think that I can use this because I am assuming that x is entirely government-provided. But, if it were true, then (6) would reduce to $\frac{\partial^2 u}{\partial c\partial x}-\frac{\partial^2 u}{\partial c^2}$.

To simplify exposition in the text, I focus on the case where the marginal utility from the publicly-provided good is at least as large as the marginal utility of consumption. We can typically expect this to be the case over ranges where expansion in good *x* are considered (e.g., starting from no provision); after all, why invest in the policy if its product is not at least as good as income? Thus, in the main text I use the formula the formula $\frac{\partial^2 u}{\partial c \partial x} - \frac{\partial^2 u}{\partial c^2}$.

Appendix 2: Functional Forms and Legal Entitlement Neutrality

Rich-Biased Policies:

Separable utility functions: For any separable utility function, $\frac{\partial^2 u}{\partial c \partial x} = 0$. As a result, the sign of (6) is positive and thus there is pro-rich bias.

Cobb-Douglas: For utility functions of the form $u = Ac^{\alpha}x^{1-\alpha}$, for $0 < \alpha < 1$ and $A > 0$, we know that policies are rich-biased because $\frac{\partial^2 u}{\partial c \partial x} = A\,\alpha(1-\alpha)c^{\alpha-1}x^{-\alpha} > 0$, since every term is positive. Thus, equation (6) is positive, and this utility function is rich-biased.

Constant Elasticity of Substitution: Constant elasticity of substitution utility functions of the form $u = (\alpha c^r + (1-\alpha)x^r)^{1/r}$, for $0 < \alpha < 1$ and $-\infty < r < 1$. For these, we know $\frac{\partial^2 u}{\partial c \partial x} = \alpha(1-\alpha)(1-r)(\alpha c^r + (1-\alpha)x^r)^{\frac{1}{r}-2}c^{r-1}x^{r-1}$ and $\frac{\partial^2 u}{\partial c^2} = \alpha^2(1-r)(\alpha c^r + (1-\alpha)x^r)^{\frac{1}{r}-2}c^{2r-2} + \alpha(1-r)(\alpha c^r + (1-\alpha)x^r)^{\frac{1}{r}-1}c^{r-2}$. Thus, since $\frac{\partial^2 u}{\partial c^2} \cdot \frac{\partial u/\partial x}{\partial u/\partial c} = \alpha(1-\alpha)(1-r)(\alpha c^r + (1-\alpha)x^r)^{\frac{1}{r}-2}c^{r-1}x^{r-1} + (r-1)(1-\alpha)(\alpha c^r + (1-\alpha)x^r)^{\frac{1}{r}-1}x^{r-1}c^{-1})$, the first half of which is equal to $\frac{\partial^2 u}{\partial c \partial x}$, we know $\frac{\partial^2 u}{\partial c \partial x} - \frac{\partial^2 u}{\partial c^2} \cdot \frac{\partial u/\partial x}{\partial u/\partial c} = (1-r)(1-\alpha)(\alpha c^r + (1-\alpha)x^r)^{\frac{1}{r}-1}x^{r-1}c^{-1} > 0$. We know that this result is positive because every term is positive as a result of $\alpha, r < 1$. So, equation (6) is positive, and these functions are rich-biased.

Neutral Policies:

Determining the sign of (6) is trivially easy when the "good" is the same thing as the numeraire good, or money. Then $c = x$, and (6) reduces to

$$\frac{\partial(\partial s/\partial x)}{\partial c} = \frac{\partial^2 u}{\partial c^2} - \frac{\partial^2 u}{\partial c^2} \cdot \frac{\partial u/\partial c}{\partial u/\partial c} = 0$$

which meets the definition of a wealth-neutral policy.

Poor-Biased Policies:

Take the utility function $U = \ln\left(x - \frac{1}{2}\right) - 2 \cdot \ln(10 - c)$, which is a member of a class of utility functions for which x is an inferior good.[103] Variable c is private consumption. Suppose that there are two individuals, a poor one with an endowment of 6 and a rich one with an endowment of 9. Suppose that initially the quantity of x available is 1. Now the government is considering providing an extra 0.5 units of the inferior good x. Going through the calculations to calculate surplus shows that the poor receive surplus of \$0.45, while the rich receive surplus of \$0.38, so that, if the good costs, say, \$0.40 to produce, the poor would receive it, but the rich would not. This result is just a function of the good being inferior: the poor demand a higher quantity of it than the rich do.

Appendix 3: Rich-Biased Policy Example

Optimal Allocation with Tradable Pollution Rights

The poor maximize $U_p = \log(1 + x_p) + \log(c_p)$ s.t. $c_p = y_p + g(\bar{x}_p - x_p)$ and the rich maximize $U_r = \log(1 + x_r) + \log(c_r)$ s.t. $c_r = y_r + g(\bar{x}_r - x_r)$ by choosing cleanliness units x_p and x_r, respectively, given price for cleanliness g, initial allocations of cleanliness $\bar{x}_p$ and $\bar{x}_r$, and income allocations y_p and y_r. The social planner wants to choose $\bar{x}_p$ and $\bar{x}_r$ so as to maximize

$$U_r + U_p \text{ (1)}$$

There are ten units of cleanliness total so

$$\bar{x}_p + \bar{x}_r = 10 \text{ (2) and}$$
$$x_p + x_r = 10 \text{ (3).}$$

The initial endowments are $y_p = 1$ and $y_p = 9$.

To solve for the initial allocations, the final allocations, and the price of a unit of cleanliness, I will first solve for the cleanliness demand curves of the rich and the poor. Rewriting the utility functions in terms of cleanliness yields $U_p = \log(1 + x_p) + \log(y_p + g\bar{x}_p - gx_p)$ and $U_r = \log(1 + x_r) + \log(y_r + g\bar{x}_r - gx_r)$. Taking the first order conditions with respect to cleanliness gives the demand curves

$$x_p = \frac{y_p + g\bar{x}_p - g}{2g} \text{ (4) and } x_r = \frac{y_r + g\bar{x}_r - g}{2g} \text{ (5)}$$

Combining the demand curves (4-5), equation (2), and the social welfare function (1), we get SWF $= \log\left(\frac{y_p + g\bar{x}_p + g}{2g}\right) + \log\left(\frac{y_p + g\bar{x}_p + g}{2}\right) + \log\left(\frac{y_r + g(10 - \bar{x}_p) + g}{2g}\right) + \log\left(\frac{y_r + g(10 - \bar{x}_p) + g}{2}\right)$. Solving the first-order condition for $\bar{x}_p$ and substituting in the values of the endowments gives:

[103] See Rein Naagsma, *A Convenient Utility Function with Giffen Behavior*, ISRN Economics (2012), *available at* https://www.hindawi.com/journals/isrn/2012/608645/.

$$\bar{x}_p = \frac{5g + 4}{g}$$

This implies that $\bar{x}_r = \frac{5g-4}{g}$, $x_r = \frac{9+4g-4}{2g}$, and $x_p = \frac{1+4g+4}{2g}$. Combining the final allocations with equation (3) reveals that the price of cleanliness is $g = \frac{5}{6}$. With this price we can solve for all other values. Thus, $\bar{x}_p = \frac{49}{5}, \bar{x}_r = \frac{1}{5}$, $x_p = x_r = c_p = c_r = 5$.

Even Allocation with Trading

If the initial allocation of cleanliness is 5 for both rich and the poor, then substituting into the demand curves (4 and 5) along with the initial endowments gives $x_p = \frac{1+4g}{2g}$ and $x_r = \frac{9+4g}{2g}$

Combining this with equation (3) gives that the price is again $g = \frac{5}{6}$ and the final allocations are $x_p = \frac{13}{5}$ and $x_r = \frac{37}{5}$.