

University of Chicago Law School

## Chicago Unbound

---

Journal Articles

Faculty Scholarship

---

2015

### Distributionally Weighted Cost–Benefit Analysis: Welfare Economics Meets Organizational Design

David A. Weisbach

Follow this and additional works at: [https://chicagounbound.uchicago.edu/journal\\_articles](https://chicagounbound.uchicago.edu/journal_articles)



Part of the [Law Commons](#)

---

#### Recommended Citation

David A. Weisbach, "Distributionally Weighted Cost–Benefit Analysis: Welfare Economics Meets Organizational Design," 7 *Journal of Legal Analysis* 151 (2015).

This Article is brought to you for free and open access by the Faculty Scholarship at Chicago Unbound. It has been accepted for inclusion in Journal Articles by an authorized administrator of Chicago Unbound. For more information, please contact [unbound@law.uchicago.edu](mailto:unbound@law.uchicago.edu).

# DISTRIBUTIONALLY WEIGHTED COST–BENEFIT ANALYSIS: WELFARE ECONOMICS MEETS ORGANIZATIONAL DESIGN

David A. Weisbach\*

## ABSTRACT

Approaches to cost–benefit analysis (CBA) derived from social welfare maximization conclude that it ideally should include distributional weights. Agencies using CBA, however, do not maximize welfare. They perform specialized tasks. Approaches that assume agencies should maximize welfare cannot be used to determine the design of those tasks. This article considers the design of CBA, including the pursuit of distributional goals by mapping optimal distributive systems to tasks performed by specialized agencies. It concludes that regulatory agencies using CBA cannot make distributional adjustments consistent with desirable distributive policies. Therefore, CBA should not include distributive weights.

Standard cost–benefit analysis sets the marginal cost of a project equal to its marginal benefits. This procedure maximizes the net benefits from a project and ensures efficient policies. Efficient policies, however, may reduce social welfare because of their distributional effects. For example, a policy that increases inequality by helping the wealthy and hurting the poor may make society worse off even if the monetary gains to the wealthy are greater than the losses to the poor.

Because unweighted CBA may reduce social welfare, standard economics texts, such as *The Handbook of Public Economics* (Dreze & Stern 1987), economics textbooks (e.g. (Gruber 2005)), and many modern legal scholars (e.g. (Adler 2013 and Adler & Posner 2006) conclude that the monetary costs and benefits of a policy should in theory be weighted by their distributional

\* The University of Chicago Law School, 1111 E. 60<sup>th</sup> St. Chicago, IL 606037, E-mail: d-weisbach@uchicago.edu. I thank Louis Kaplow, Jonathan Masur, Jennifer Nou, and Eric Posner, an anonymous referee, and participants at seminars at the University of Chicago and the American Law and Economics Association 2014 annual meeting for comments.

impacts.<sup>1</sup> Costs and benefits under this approach are multiplied by the marginal utility of income and by the marginal social welfare of utility of the affected individuals. Distributional weighting, as this procedure is called, ensures that projects increase social welfare, rather than merely pursuing efficiency.

There is a contrary line of literature, however, that argues that the income tax should be the sole tool used to address distributional concerns.<sup>2</sup> Based on the reasoning in [Atkinson & Stiglitz \(1976\)](#) and its extension in [Kaplow \(2006b\)](#), this literature argues that it is less efficient to redistribute through modifications to regulatory rules than to achieve the same level of redistribution directly by taxing the wealthy more or the poor less. This literature, therefore, concludes that distributional weights are undesirable.

Proponents of distributional weights often dismiss the “income tax only” argument, as it is sometimes called, as requiring restrictive assumptions (e.g. [Alder 2013](#)).<sup>3</sup> When these assumptions are relaxed, the conclusion that it is always better to use the labor income tax to address distributive concerns no longer holds. Proponents of distributional weights, therefore, conclude that weights are needed in the more general case.

My goal here is to explore when CBA should include distributional weights in the general case, when we relax the assumptions required for the income tax only result. My core claim is that to understand when distributional weights or other distributional adjustments to CBA are desirable, we have to take an explicitly institutional perspective, viewing CBA as the set of instructions we give to specialized agencies pursuing assigned tasks. The government pursues a large number of different objectives, each requiring specialization. It assigns tasks to agencies who have expertise in these tasks. The question is which agencies

- 
- 1 There is a closely related set of arguments about the extent to which legal rules should be set based in part on their distributive effects (see [Kaplow & Shavell 1994](#)). Many of the considerations are similar. Regulations set using CBA, after all, are legal rules. The legal rules literature, however, tends to focus on rules set by courts rather than by regulatory agencies. Courts and regulatory agencies have different institutional structures. Because the arguments I make here are about the institutional structure and capacities of agencies, they do not necessarily extend to courts and court-made legal rules and conversely the conclusions from the literature on legal rules may not carry over to the CBA context.
  - 2 The literature is now large ([Atkinson & Stiglitz 1976](#); [Hylland & Zeckhauser 1979](#); [Kaplow 2004](#); [Kaplow 2006a](#); [Kaplow 2008a](#); [Kaplow 2012](#); [Kaplow & Shavell 1994](#); [Kaplow & Shavell 2000](#)). Many of the arguments made here can be found in various places within this literature. Nevertheless, given the lack of engagement with these arguments by those arguing for distributional weighting, it is apparent that further exploration of the issues is warranted.
  - 3 In the economics literature, [Dreze and Stern \(1987, p. 958\)](#) make an analogous claim, arguing that because of the second-best nature of an income tax, distributional weighting is required.

should be assigned distributive tasks. The argument for distributive weights starts with the assumption that all agencies should maximize welfare, and from that assumption concludes that all agencies that use CBA should be assigned distributive tasks.

Casual observation of large institutions other than governments illustrates that individual tasks assigned to divisions are not normally the same as the overall goal. The overall goal of an automobile company is to make and sell vehicles at a profit, but no single division, on its own, does exactly that. Design, research, engineering, procurement, legal, accounting, advertising, testing, finance, grounds keeping, sanitation, security, manufacturing, sales, security, regulatory compliance, and other tasks are assigned to divisions that pursue their assigned tasks, and do not aim on their own to produce, market, and sell complete vehicles. When these tasks are properly assigned and coordinated, the divisions of the company achieve the overall goal even though none pursue it directly.

The same is true for assigning tasks to units of the government, such as regulatory agencies using CBA. For example, it likely maximizes social welfare to have a well-prepared army and air force to defend the country. It does not follow that the staff of the SEC, and the Comptroller of the Currency should fly fighter jets or hunt down terrorists. We do not want the Environmental Protection Agency determining monetary policy or the Federal Drug Administration setting trade policy. The State Department does not deliver the mail. Knowing the overall set of goals, does not tell us what any particular agency should be doing. Knowing that distribution is important does not tell us that agencies using CBA should pursue distributional goals. The task of determining what agencies should do is one of institutional designs, not one of determining the overall goals of the government. The statement that unweighted CBA does not maximize welfare says nothing about whether CBA should include distributional weights. Instead, to determine whether CBA should include distributive weights, we need to map optimal distributive policies to specialized agencies, seeking the best overall organization of government.

I will consider such a mapping here, taking it in four steps. The first step considers the standard argument for distributive weights and shows that it fails to consider even the most basic structure of government, including the presence of an income tax. It is divorced from institutional concerns, and therefore provides little or no guidance for the design of CBA. The second step reviews the literature on the design of the distributive system and shows what types of adjustments away from purely efficient policies are desirable when we relax the assumptions of the income-tax only literature.

The third part considers the mapping of the distributive system onto agencies pursuing specialized tasks and argues that specialized administrative agencies performing CBA are not good candidates to also pursue distributive tasks.

The fourth step extends the analysis by considering the interaction between agencies and legislatures. Legislatures usually set tax rates while agencies set regulatory policy. The question addressed in this last step is whether agencies should use distributive weights if the legislature fails to set reasonable tax policies.

## 1. THE ARGUMENT FOR DISTRIBUTIVE WEIGHTS AND WHY IT FAILS

Unweighted CBA is a command to maximize efficiency. CBA achieves this by setting the marginal costs of a project equal to the marginal benefits, as measured in prices. Because CBA often addresses market failures, prices may not be available. To mimic prices, CBA uses compensating variations (CVs), the payment an individual would need to offset the effect of a change in policy. Adding up CVs, and where available, prices, allows CBA to find policies that maximize efficiency.

Maximizing efficiency is not the same as maximizing welfare. Suppose that person A gains \$11 from a policy change and person B loses \$10. The gains outweigh the losses so the policy is efficient. We cannot know whether the policy change improves social welfare, however, unless we know the distribution of the gains and losses. B's loss of \$10 might reduce his utility more than A's gain of \$11 increases hers. As a result, the project may reduce social welfare even though it is efficient. Because unweighted CBA can produce policies that reduce social welfare, many argue that it is not an appropriate criteria for agencies to use when setting policies. If agencies use CBA they may make us worse off.

Distributional weights fix this problem and ensure that policies increase rather than reduce social welfare. The standard approach, going back to a series of papers by Robin Boadway (1974, 1975, 1976), is to derive CBA directly from a social welfare function:

$$W = W(U^1, \dots, U^j, \dots, U^m),$$

where  $U^j$  is the well-being or utility of individual  $j$ . The utility of an individual is a function of his income  $Y_j$ , which he uses to consume goods or supply factors  $X_i$ . The social welfare function  $W$  aggregates utilities according to standard criteria in welfare economics and can either be utilitarian or express a preference for helping individuals with low utility levels.

Consider a project or policy to be implemented by an agency, such as pollution regulations or new standards for financial institutions. Suppose the policy produces a small change to the quantity of good  $X_i$ . To evaluate the

effect of the policy, we need to know how it changes social welfare. To compute this, we take the derivative of  $W$  with respect to the change in  $X_i$ , which produces three terms: the change in income for a given individual when  $X_i$  changes, the change that the change in income induces in the individual's utility and the change in social welfare from the change in the individual's utility. We sum this over all individuals to find the change in welfare. In notation:

$$\frac{\partial W}{\partial X_i} = \sum_j \underbrace{\left[ \frac{\partial W}{\partial U^j} \frac{\partial U^j}{\partial Y_j} \right]}_{\substack{\text{Distributional weight} \\ \text{for individual } j}} \frac{\partial Y_j}{\partial X_i}$$

The last term is the efficiency term: it tells us how much the income changes when we change the quantity of a good. Unweighted CBA looks only at the efficiency term, adding it up over all individuals. The first two terms (grouped together in brackets) make up the distributional weight for an individual or class of individuals. It tells us how much welfare goes up if we were to give that individual or class of individuals another dollar. Multiplying the efficiency effects by distributional weights gives us the effect of a policy on social welfare. CBA with distributional weights, therefore, ensures that policies increase social welfare; without weights this may not be true.

To go back to the numerical example, we multiply A's gain of 11 by her marginal utility of income to estimate how much her utility has gone up. We also multiply this by a factor that measures how much social welfare goes up as A's utility increases. We do the same for B: we multiply his loss of \$100 by his marginal utility of income to measure how much his utility has gone down, and so forth. We then add up the changes in welfare produced by the policy. If the sum is positive, the policy increases social welfare. Following this procedure ensures that we adopt the proposed policy only if the social value of A's gains exceeds the social cost of B's losses.

This set of arguments is now standard wisdom in economics. The Handbook of Public Economics, which represents as close to official surveys as are found in economics, takes precisely this approach (Dreze & Stern 1987). It derives CBA as a procedure to maximize social welfare using the same basic derivation as Boadway. As a result, it concludes that distributional weighting is required in a complete cost–benefit analysis. Widely used (and recent) textbooks on public economics are similar (Tresch 2002; Cullis & Jones 1998; Gruber 2005). The key idea is that cost–benefit analysis is justified as a procedure that directly maximizes social welfare. Distributional weighting follows almost immediately from this statement of the problem.

Many legal scholars follow this approach.<sup>4</sup> For example, Adler (2013), in a recent survey on distributional weighting in CBA, derives distributional weights for specific social welfare functions. His derivation is identical to that outlined above: he takes the first derivative of the social welfare function and collects terms.

Adler and Posner (2006), in one of the most influential treatments of CBA in the legal literature, also take this approach. They view the correct decision criteria for agencies as one that maximizes social welfare. If there were no time or information constraints, CBA would, in their view, include distributional weights. Because agencies face time constraints and have limited information, they argue that CBA generally can be a reasonable decision criteria. It tracks welfare reasonably well and can be implemented by agencies.<sup>5</sup> They are uncertain about distributional weights, however, because of the information problems they raise. While they do not have strong conclusions on distributional weights, their underlying argument is the same as Boadway's: they view the proper decision criteria for each agency to be to maximize social welfare.

John Bronsteen, Christopher Buccafusco and Jonathan Masur (2013) also implicitly take this approach and endorse distributional weights. They suggest that CVs be replaced with the results from surveys about people's happiness because they believe that the happiness surveys approximate utility. Using utility for CBA calculations, however, is equivalent to using monetary values weighted by the marginal utility of income. This means that Bronsteen et al. implicitly assume that a version of distributional weighting is desirable.<sup>6</sup>

4 There is a substantial legal literature that is loosely connected to distributional concerns but that does not provide explicit arguments. Most prominently, Cass Sunstein argues for a pragmatic defense to CBA and believes that CBA should take note of who bears the costs and benefits (Sunstein 2002; Sunstein 2006). In some cases, Sunstein in part bases his defense of CBA on social welfare. Hahn and Sunstein (2001) take this approach and suggest that distributional effects be considered a "soft" variable.

5 Adler and Posner (2006, p. 152: "It is unclear whether the basic idea of distributive weighting is itself a feasible one. . . . If distributive weighting is not feasible, then agencies should probably use unadjusted CV's when the distribution of wealth among the winners does not differ much from the distribution among the loser." and pp. 188–89: "Cost benefit analysis should be a good proxy for overall welfare where the distribution of wealth or income among project winners is not dramatically different from its distribution among project losers."). They make a similar assumption in Adler and Posner (1999, p. 194: "the effect of a governmental project on overall well-being is a morally relevant, if not morally decisive feature of the project" and p. 216: "CBA is an appropriate means, or decision procedure for achieving the morally desirable goal of promoting overall well-being.")

6 Note that a measure that used the sum of changes in utility would exclude the term for the slope of the social welfare function,  $\partial W/\partial U^j$ . Bronsteen et al. do not discuss whether they would include this term or, if not, the reasons for excluding it. Their core argument is that happiness surveys are more accurate because the estimation of CVs is difficult. If true, however, they should then divide by the marginal utility of income to recover the monetary measure used in unweighted CV, a procedure which they do not recommend.

The problem with this approach is it does not account for the institutional setting in which CBA is used. The model does not have agencies who perform CBA. It does not even have an income tax, which has been the central distributive institution in every developed country for the last century. It attempts to determine the organizational structure of government by using a model that does not have a representation of the organizational structure.

Actual governments pursue a wide range of policies, such as regulating financial markets, pollution, drugs, transportation, education, or telecommunications, defending the country, and redistributing income. And actual governments assign these tasks to particular agencies. These agencies are not told to maximize social welfare or to pursue each and every policy that the government pursues. They are told to perform a particular task. One hopes that if the tasks are assigned appropriately and are properly coordinated, the overall outcome is to maximize welfare, but no agency is itself told to maximize welfare.

Banking regulators, for example, are concerned with the health of the banking system. They develop expertise in these issues and regulate in areas involving banks and their customers. They are not concerned with, and know little about, workplace safety, endangered species, or transportation, all areas that the government regulates to maximize welfare. Instructing the banking regulators to consider all of these policy areas would not account for the specialized nature of their tasks. Similarly, the Postal Service does not set monetary policy. Air traffic controllers do not approve drugs. The Forest Service does not set education policy.

Similarly, sub-units of other types of institutions such as corporations do not normally pursue the overall goal of the institution. As noted, large corporations have divisions—research, engineering, design, manufacturing, marketing, accounting, legal, purchasing, sales, security, and so forth—each which pursues an assigned task. Together, the divisions achieve the goals of the institution even though none pursues it directly.

In short, we cannot conclude from a claim that governments should pursue distributive goals in order to maximize welfare that each and every agency should pursue distributive goals. But this is exactly what the argument for distributive weights does. Because it does not consider the internal structure of government, it may systematically lead to wrong conclusions.

Some approaches in the legal literature do try to place CBA in its institutional setting. Most prominently, [Adler and Posner \(2006\)](#) consider the institutional limitations facing agencies when performing CBA. They begin with the assumption that agencies should ideally maximize welfare and then consider whether CBA is a reasonable decision procedure that approximates that goal. For example, agencies may lack the information needed to directly maximize welfare



or may face time pressures that force them to make rough justice decisions. In light of these constraints, they conclude that CBA may be a reasonable approach for agencies and, although they hedge, that perhaps it should be unweighted.

While they place CBA in an institutional setting, the setting does not consider the problem of task assignment, which is the central question. We need to determine which units within the government should address the regulation of pharmaceuticals, patents, crime, highway safety, distributive, or other policy. This task assignment element of institutional design is absent in Adler and Posner's treatment of the issue, exactly like it is absent in the standard, non-institutional approach exemplified by Boadway's work. That is, their assumption that agencies should maximize welfare ignores the very reason for the use of a specialized agency in the first place and, moreover, is contrary to fact: agencies are not told to maximize welfare.

The goal for the remainder of the article is to consider how distributive tasks should be assigned within a government that pursues a wide variety of policies through expert agencies. I start with a discussion of optimal distributive policy should look like, reviewing existing literature rather than deriving new results. I then consider how we should assign the resulting distributive tasks to expert agencies.

## 2. OPTIMAL DISTRIBUTIVE POLICIES

To understand how to assign distributive policies to agencies, we need to understand what good distributive policies look like. The branch of economics that addresses this question is known as optimal taxation. I review the basic tax results from the literature here. I will start with a model I call the benchmark model, which concludes that distributive policies should be pursued only through the labor income tax (the income tax only result). The benchmark model contains strong assumptions. The key claim of the benchmark model is not that the assumptions are true. Instead, the model helps us understand the implications of a more general model where the assumptions are relaxed. The second part of this section describes the more general results.

### 2.1 Benchmark Model<sup>7</sup>

Suppose that there is a population of individuals who differ in their ability to earn. Some people can command high salaries and others only much lower

---

7 Important papers developing this model and its implications include Atkinson & Stiglitz (1976); Hylland & Zeckhauser (1979); Christiansen (1981); Kaplow & Shavell (1994); Kaplow (2004); Laroque (2005); Kaplow (2006a); Kaplow (2006b); Gauthier & Laroque (2009); and Kaplow (2012).

salaries. Although the government knows the overall distribution of ability, it does not know any particular person's ability. Instead, it can only observe market transactions, such as income or spending, and make inferences from what it observes. The government must raise taxes to pay for its various functions such as national defense and it can use taxes to help redistribute. Taxes can be positive (individuals pay the government) or negative (the government makes payments to individuals). That is, the 'tax system' includes both taxes and welfare. The goal is to set taxes to maximize social welfare subject to the government's budgetary requirements.

One question of great interest that is studied using this set up is the overall progressivity of the tax system. What is the optimal rate structure? This problem was first solved by Mirrlees (1971) and involves a number of complexities not relevant here. Our question, while related, is different: we want to know sorts of instruments should be used to pursue distributive goals. Should the distributive policies be implemented only through a labor income tax or should it use other tools, such as excise taxes or subsidies on particular goods?

To help understand this latter problem, the benchmark model makes an additional simplifying assumption about how people behave. It assumes that people spend their money independently of how they earn it. Two people with the same earnings will spend their money in the same way. Called "weak separability", it means that all goods are equal complements or substitutes for labor. We can think of weak separability as meaning that people divide the choice of how much to work and earn from the choice of how to spend the earnings. They decide how much to work and earn based on the gains to additional consumption and how much they like or dislike work. For a given amount of earnings, they decide which particular goods to consume based on what sorts of goods they like and how much the goods cost.

Consider a tax on a good. The tax will increase the price of the good. An increase in the price of the good, or any good, affects both work effort and which goods individual purchase. It will reduce work effort because the amount of stuff that more earnings can purchase has gone down. It will alter which goods are purchased because the relative prices of goods have changed. For example, if we impose a tax on televisions, the gains to an additional hour of work (being able to afford an expensive television) go down so people may work less. In addition, for a given amount of earnings, the price of televisions has gone up relative to other goods, so individuals will reduce television purchases and substitute other items.

Suppose we get rid of the tax on televisions and replace it with a tax on all items that raises the same revenue. We can think of this as spreading out the tax on televisions to apply uniformly to all goods. The spread out tax would be different at each income level because people spend different portions of their

budgets on televisions at each income level. One income group, say poor people, might spend 1 percent of their income on televisions while another group, say middle-income people, might spend 3 percent. The spread out tax would reflect these amounts. It would be a different amount for people with different incomes.

If the spread out tax applies to spending, it is difficult to adjust it for people with different incomes. Sellers of goods may not know the income of buyers so they would not know what tax to charge. To impose a different tax on people with different incomes, we can impose this spread out tax on labor earnings rather than spending. If you spend what you earn, it makes no difference whether we tax you up front when you earn it or later when you spend it. The effect will be the same. By imposing the tax on earnings rather than on spending, however, the tax can be tailored to people with different income (and therefore, different spending patterns).

Therefore, compare the tax on a particular good, say televisions, with an adjustment to the tax on labor income, with adjustment at each income level reflecting the tax paid on televisions by people at that income level. By construction, tax revenue is the same at each income level. This also means that the distribution of taxes is the same. Each individual, however, is better off with the tax on labor earnings than with the tax on televisions because his choice of what to buy is no longer distorted by the tax system. The spread out tax raises the same revenue, holds distribution constant, but makes people better off. We should not tax televisions. We should just use the tax on labor earnings.<sup>8</sup>

To illustrate with an example from actual policy, imagine that we have a tax on luxury goods, goods that only people with high earnings consume, such as fancy automobiles, expensive jewelry, and private jets.<sup>9</sup> This tax will be progressive so it will improve the distribution of resources. At the same time, because it raises the cost of consuming, it will reduce labor effort. An hour's work buys less so the trade-off between work and leisure is less favorable. The goal of the tax is to make the system more progressive and more progressive

8 The heuristic argument in the text leaves out an important complication which is addressed in more formal proofs. If we replace the tax on televisions with a tax on earnings that, at least initially holds tax revenue constant, people will be better off (as suggested in the text), but this change in taxes may cause people to change their labor supply. This in turn affects revenue. The thought experiment in [Kaplow \(2006b\)](#) is to replace the tax on televisions with a tax on earnings at each income level that holds utility constant rather than holding tax revenues constant. This alternative approach ensures that there is no change in labor supply and no change in utility, but raises more revenue than the tax on televisions. This revenue can be refunded, creating a Pareto improvement. I use the heuristic explanation in the text because I think it helps generate intuitions. If labor supply uniformly goes up when we replace the tax on televisions with the tax on labor earnings, the heuristic approach in the text works.

9 The USA had a luxury tax from 1991 until 1993. See Internal Revenue Code Sections 4001–4003.

taxes will reduce the rewards from earning more. A luxury tax will also distort what types of goods people buy. They will shift away from luxuries because the price of luxuries is higher.

Now suppose that we adjust the income tax to raise rates on upper income individuals and repeal the luxury tax. For the reasons discussed above, with the assumptions in the benchmark model, we can do this without changing tax receipts or the distribution of income. With the luxury tax gone, individuals can now choose which items to buy without distortions caused by the tax system so individuals are better off. We get the same revenues and distributive effects at a lower efficiency cost using the income tax than using a luxury tax.

The argument also holds for legal rules and regulatory policies. Suppose we adjust a rule, such as a safety standard, on luxury goods to redistribute. We might make luxury goods more expensive than the efficient regulatory policy would because we care about inequality and distribution. Alternatively, we might make goods mostly consumed by the poor less expensive. Under the assumptions of the benchmark model, we can replace the redistributive regulations with efficient ones and make a corresponding adjustment to the income tax. We keep the redistributive effects but eliminate the inefficiency. Therefore, within the benchmark model, we would not want to use regulatory policy for redistributive purposes and would not want to use distributive weights.

## 2.2 General Case

The benchmark model relies on a number of assumptions including the assumption of weak separability. When we relax these assumptions, we get a broader range of taxes and subsidies. The literature examining these cases is large. I will highlight a few of the main results.<sup>10</sup>

### *Weak Separability*

The benchmark model assumes that labor and leisure are weakly separable. As noted, a heuristic for thinking about weak separability is to think of it as two-stage budgeting. Individuals decide how much to work given their earnings ability. Given these earnings, individuals decide what items to spend the earnings on. This means that if we tax bowling alleys, beer, or beauty salons, the effect on labor supply is the same. Increasing or decreasing the price of one of the goods that individuals consume has the same effect on their work effort as if the price of all goods were increased by an average amount.

---

10 Kaplow (2008a, pp. 137–145) provides a list of the results that come from relaxing the assumptions in the benchmark model.

Commentators have argued that the assumption of weak separability is unlikely to be true. When we relax the assumption of weak separability, however, we do not recover distributional weights.

Suppose that utility is not weakly separable in labor. This means that some goods are better complements or substitutes for labor, or its alternative, leisure, than others. A spread out tax on all goods no longer necessarily has the same effect on labor supply as a tax on a single good. If a good is say a complement to leisure, a tax on that good will make leisure less attractive, inducing people to work more. If a good is a substitute to leisure, subsidizing it will make leisure less attractive, again inducing people to work more.

We can use this feature—complementarity or substitutability for leisure—to improve the ability of the tax system to redistribute. If we try to redistribute through a tax on income, people will reduce their work effort, engaging in more leisure, reducing the efficacy of the tax. If we make leisure more expensive, we reduce this effect, making it easier to use the income tax to redistribute. We can make leisure more expensive by taxing goods that go with, or are complements to, leisure. We can also subsidize goods that do not go with, or are substitutes to, leisure. When we impose these taxes and subsidies, the progressive income tax works better, allowing us to do a better job reducing inequality. Therefore, once we allow goods to be relative substitutes or complements to leisure, we no longer only want to use the income. We can improve distributive policy by taxing complements to leisure and subsidizing substitutes.<sup>11</sup>

While these taxes and subsidies are designed with distributive concerns in mind, they are not directly distributive. Instead, they enhance the ability of the income tax to redistribute. They do not look pro-poor and may even appear pro-rich. When we relax the assumption of weak separability, we do not recover distributional weights.<sup>12</sup>

To illustrate, consider taxes on different types of foods (Iorwerth & Whalley 2002; Kaplow 2008a, pp. 138–139). VATs commonly exempt grocery store food but not food served in restaurants. There appears to be a distributive benefit to this because wealthy people can better afford restaurants than can poor people. Taxing restaurants hurts the rich because restaurant meals are a form of luxury.

11 Kleven (2014) argues that Scandinavian systems are able to be highly distributive in part because they pursue precisely the policy of subsidizing complements to labor.

12 Adler (2013) recognizing that any adjustments to CBA would involve subtle factors such as complementarity with labor argues that these adjustments can still be called distributive weights. While there is no necessarily right or wrong definition of the term “distributive weights”, if adjustments of this sort are called distributive weights, one should be very clear that this is what is intended because it is not how the term is normally understood. Instead, the term normally refers to  $(\partial W/\partial U)(dU/dY)$  as discussed in the text. This is particularly true because the weights will not be distributive in the sense that they seem to directly favor the poor.

Exempting groceries similarly helps the poor. Standard distributional weights would recommend lower taxes on groceries and higher taxes on restaurants. Or we might regulate restaurants to make them more expensive or grocery stores to make them less expensive.

Groceries, however, are an input into the household preparation of meals. Home cooking is relatively inexpensive but time-consuming, so it is a complement to leisure. Relative to cooking your own meal, restaurant meals are expensive and use little time. They are a substitute for leisure. Therefore, contrary to the approach we would get using distributive weights, we should tax groceries and subsidize restaurants. Conventional “pro-poor” adjustments go in exactly the wrong direction.

The same considerations hold for luxury taxes. A naïve view might be that they are a good way of targeting the wealthy. We saw that luxury taxes do not make sense if we assume weak separability. Even when we relax the assumption of weak separability, however, luxury taxes are not consistent with optimal distributive policies. Many luxuries are expensive relative to the time that they take to consume (this is almost the definition of a luxury good). They are substitutes to leisure, which means that they should be subsidized not taxed (and we should then increase marginal tax rates on the wealthy to hold the distribution constant or perhaps increase those rates to redistribute more). Conventional distributive weights again point in exactly the wrong direction.

#### *Common Utility Functions*

A second assumption in the benchmark case is that people have the same utility function. The model assumed that they differ in their earnings ability but not otherwise. This assumption is not true. People vary widely in their tastes and preferences. In most cases, these differences will not matter to the design of distributive policies and the results from the benchmark model continue to apply. If I like rock and roll and you like hip hop, we are equally well off (except for your bad musical tastes), and we could not improve distributive policies by say, taxing rock and roll.

Suppose, however, that people differ in a way that makes some of them worse off. One possible example is that certain disabilities make people worse off even holding earnings ability constant. One reason might be that for a given level of earnings, a disabled person might have to spend more on basic needs. If this is the case, we may want to find a way to transfer resources to individuals with those disabilities. The additional personal exemption for the blind is often justified this way.

Note, however, that the adjustments are not necessarily pro-poor and do not resemble the type of adjustments to CBA produced by distributive weights. The shift in resources is toward people who are worse off not toward people who

have less income. A blind person might be quite wealthy but would (if true) be worse off than a sighted person with equal wealth.

Kaplow & Shavell (2000) illustrate this using klutziness. They posit, purely hypothetically, that some rich people are less coordinated than others so that they cause more accidents. These klutzes are less well off than their coordinated peers. In this case, we might want to find a way to redistribute toward the klutzes. For example, we might want to subsidize safety devices used on expensive automobiles and yachts or perhaps reduce tort damages on rich people.<sup>13</sup>

The sorts of policies we should adopt if we believe people are different in a way that makes them better or worse off are complex. Making the judgment that someone is worse off than their peers because of some attribute is fraught with difficulties. Policies to redistribute toward those individuals can be difficult to implement without a careful understanding of the nature of the problem. For example, optimal policies will depend on the degree to which the difference is observable and the extent to which people will be induced to be in, or pretend to be in, the favored category. The literature on how best to treat disabilities illustrates how difficult the problem is, including Weisbach (2009) and Samaha (2007). Nevertheless, for our purposes, the implications are relatively straightforward: differences in well-being due to non-income differences do not lead to distributive weights. These differences—being blind, liking hip-hop, being a klutz, and so forth—are not systematically connected to income. They may be more prevalent among the wealthy or among the poor. The required adjustments in policy, therefore, may be pro-poor but they may equally be pro-rich.

#### *In-Kind Provision*

The benchmark model implicitly uses money—taxes are paid to the government in money and transfers are payments by the government in money. While this is true for most payments and transfers, a substantial portion of transfers are in-kind or involves restrictions on the use of money that effectively make them in-kind. For example, housing vouchers and food stamps (SNAP) can only be used for specified kinds of purchases. Medicaid is an in-kind provision of healthcare. Specialized agencies, such as housing or food agencies, not the Internal Revenue Service, administer these programs.

13 Introducing heterogeneity in utility functions generates a number of complexities that are beyond the scope of the present discussion. The design of distributive policies depends on how observable the differences are, among other things (Golosov et al. 2013; Kaplow 2008b). For an illustration of how these considerations apply in the context of disability, see Weisbach (2009). A literature on “tagging” considers observable differences (Akerlof 1978; Mankiw et al. 2009).

There are a number of possible justifications for in-kind provision of goods (although in many cases, it may simply not be justified). One possibility is that in-kind provision screens for those who really need the good. It makes it less attractive to claim to be poor to obtain transfers (Nichols & Zeckhauser 1982; Besley & Coate 1992). The choice to redistribute in-kind under this rationale improves the efficiency of redistribution but it is not directly redistributive itself, much like taxes on complements to leisure.

#### *Other Modifications*

There is, now, a cottage industry examining exceptions to the benchmark case. For example, the benchmark model only considered a single time period and the results might change if we allow individuals to choose labor effort and savings for different time periods (Kocherlakota 2010). Consumption of particular types of goods might indicate hidden abilities, allowing us to use patterns of consumption to better design the income tax (Saez 2002; Bankman & Weisbach 2006; Golosov et al. 2013). Inequalities might arise from sources other than the ability to earn labor income, such as differences in the receipt of bequests (Piketty & Saez 2013).

The details of each model are beyond the scope of what is needed here. Under the general case—without the restrictive assumptions in the benchmark model—we will want commodity taxes and subsidies of various sorts. Distributive policy should use more than the labor income tax. The resulting taxes and subsidies, however, are not the conventional pro-poor redistributive policies we might have thought optimal without careful consideration. Instead, they are set based on subtle interactions with the income tax. Relaxing the assumptions in the benchmark model does not restore standard distributive weights.

### **3. MAPPING DISTRIBUTIVE POLICIES TO AGENCIES**

The goal is to apply the literature on optimal distributive policies to the design of CBA. I will assume that the government is pursuing a large number of tasks, which it assigns to agencies. The question is whether and when an agency assigned to pursue a specialized task should consider the distributive effects of its policies. We want to map optimal distributive policies to agencies. A central question, although not the only one, is whether there is a set of adjustments to CBA that are uniform across all agencies, such as distributional weights, that are desirable.

There is only a modest literature on how to assign tasks in a large entity. Most of the existing work on the internal structure of institutions focuses on the



optimal number of divisions within an institution. The idea, dating back to Adam Smith and the pin factory, is that more specialization promotes efficiency but it also imposes coordination costs. The optimal number of divisions will balance the increased benefits from specialization and the costs of coordination. These theories do not, however, tell us how to assign tasks to divisions conditional on the number of divisions. We need a theory of how to group tasks.

What limited literature there is on grouping of tasks is based on the intuitive idea that we should group “related” tasks together. There should be some efficiency gain or complementarity from the grouping, such as an economy of scale.<sup>14</sup> In particular, specialized knowledge can be helpful in performing a number of different tasks, in which case it makes sense to put them together. If there is no complementary between the tasks, it makes no sense to group them together: we don’t IRS agents flying fighter jets.

The historian of business Alfred Chandler illustrated the concept by quoting a memo from DuPont concerning a reorganization of its business: “the most efficient results are obtained at least expense when we coordinate related effort and segregate unrelated effort” (Chandler 1962, p. 69). Oliver Williamson, the Nobel Laureate economist whose work focused on transaction costs and institutions, used a similar principle in proposing the optimal design of an institution (Williamson 1975, pp. 136–141). He considered an example of a business that manufactures and sells goods in a number of different markets. His suggested grouping of activities for this business put together activities that implicitly had complementarities. For example, he suggests that the initial stage of production of the good be separated into its own division but that the intermediate and final stages of production be combined into a division with a separate division representing that combination in each market. The assignment of tasks meant to find combinations of tasks where coordination is helpful.

Assignment of tasks to agencies in the federal government generally follows this approach. The EPA, for example, regulates environmental externalities. It has experts in environmental science, environmental economics, and environmental law on its staff. It does not regulate securities, the spectrum, or education or try to prevent child abuse because its specialized knowledge will not help it perform these tasks. The FDA approves drugs but does not run monetary policy or tax policy. Grouping the approval of different types of drugs together into a single agency makes sense, but grouping drug approval with monetary policy and tax policy does not. In response to the terrorist attacks in 2001, the government regrouped tasks related to security into a single agency, the

---

14 For a survey of literature, see Weisbach & Nussim (2003).

Department of Homeland Security based on the idea that there were complementarities among these activities.

Task assignment to agencies in an institution as large as the federal government is, of course, more complex than this simple picture. Similar tasks are assigned to more than one agency, tasks are poorly defined, some tasks are (mis)assigned to agencies with no seeming expertise in the area, and agencies sometimes closely coordinate their actions when single tasks are split across agencies. Tasks may be assigned haphazardly or based on political considerations. These complications create opportunities for learning more about the best set of instructions and task assignment for agencies. I will examine a relatively simple setting here, where we consider agencies that are assigned groups of related tasks based on complementarity of the tasks.

I start with a relatively straightforward insight from applying the task assignment approach to welfare economics: unweighted CBA is the right criteria for fixing market failures. To the extent an agency's task is to correct a market failure, it should use CBA. CBA need not maximize social welfare for this logic to hold. CBA, in this case, is not the tenth-best criteria given information and time constraints on agencies as legal scholars such as Adler, Posner suggest. It is the first-best criteria if the problem is to fix a market failure. After discussing this justification for CBA, I turn to the more difficult issue of whether CBA should be adjusted to take distributive concerns into account.

### 3.1 CBA is the Correct Instruction for Solving Market Failures

An important function of regulations is to correct market failures. Circular A-4, the guidance document governing cost-benefit analysis in the federal government highlights this, listing three market failures (externalities, information asymmetries, and monopoly power), as central reasons for regulation ([Office of Management and Budget 2003](#)). Numerous regulations fit this description. The most central case is the regulation of pollution. Externalities are also an important reason for bank regulation: bank runs potentially impose externalities on other banks or on the market as a whole. Other agencies, such as those that focus on health and safety regulation or consumer use of financial instruments, deal with informational asymmetries and, therefore, can also be thought of as correcting market failures. By no means are all regulations aimed at correcting market failures, but many are.

Suppose that an agency's task is to correct a market failure. Unweighted CBA provides the correct criteria for regulatory action in this case. When a market is functioning (i.e. there is no market failure), prices and quantities equilibrate so that marginal costs equal marginal benefits. If not, there would be a trade that makes both the buyer and seller of a good better off. If we regulate because of a

market failure, choosing a policy that sets marginal costs equal to marginal benefits puts the regulated market in the same place as a well-functioning market. Unweighted CBA provides precisely these criteria, so it is the correct tool to use if the task is to fix a market failure. It fixes the problem at hand.

There are two immediate consequences to this observation: first, when we use CBA to correct market failures, monetary values are the right measure of costs and benefits. In well-functioning markets, marginal costs will equal marginal benefits in monetary terms. To fix market failures, CBA should also use monetary values rather than an estimate of utility or some other value.

Second, using distributive weights would mean that CBA would not correct the market failure. In a perfect market, actors set marginal costs equal to marginal benefits without regard to distributive concerns. To produce the same result, CBA should not be adjusted for distributive concerns either.

Using monetary values unadjusted for distributive concerns means using CVs even though they seem to favor the wealthy. The wealthy might be willing to pay a lot for some trivial good because it represents a small portion of their wealth. The poor may not be able to pay very much for critical goods simply because they do not have very much money. As a result, an approach that relies on CVs would seem to suggest that trivial goods for the wealthy are worth more than critical goods for the poor. Can it really be the case that some trinket, say a watch or fancy vehicle that sells for tens or hundreds of thousands of dollars, is worth more than medical care, clean water, or schooling for a poor person? CVs would suggest that the answer is yes.

This effect, however, is exactly what happens in the market. The market provides goods based on consumers' willingness to pay for them. It will provide yachts, luxury automobiles, carbon fiber bicycles, expensive clothing, and fancy watches far beyond what many would say is needed. We rely on distributive instruments to reduce inequality and thereby moderate this effect, but markets produce goods for those who can purchase them. Using CVs in cost–benefit analysis does exactly the same thing. It fixes the market failure. If distributive concerns are an issue, they are an issue with functioning markets generally. To be clear, these effects are not necessarily desired, but they are general effects, not effects limited to regulated markets.

This is most dramatically illustrated for the value of life. Wealthy people can afford to pay more for safety than can poor people. It is not that in some intrinsic sense that they want to live more. They can pay more simply because they have more money. As a result, the market will produce goods and services that reflect a higher value of life for the wealthy than the poor. Expensive cars, for example, may have safety features that cheap cars do not. Wealthy people will pay for these features and poor people will not (because they cannot).

CBA should do the same. We may deplore wealth differentials and do all we can to reduce inequality, but to the extent there is inequality, markets will provide goods and services to those who can pay for them. This includes the value of life. How to fix such inequalities is a central question of public policy, but it is hard to see why choosing to fix it in markets that happen to be regulated rather than choosing based on effectiveness is the right solution.

An important corollary to the argument that unweighted CBA is the right tool for correcting market failures is that if an agency does not seek to correct a market failure, CBA may not be the correct criteria. If the goal of an agency is to do something other than set marginal costs equal to marginal benefits, CBA will tell the agency to do something other than pursue its goal. For example, Martha Nussbaum asks whether CBA can address whether Amish children should be required to go to public school (Nussbaum 2000, p. 1025). CBA likely has little to add to such a question because the agency addressing this question is not seeking to fix a market failure. Similarly, if an agency is given a task of redistributing resources to a favored group, CBA would not provide good criteria for this task. For example, if a regulation is designed to give farmers inefficient subsidies, CBA cannot be used to determine the extent of the subsidies. CBA tells the agency to be efficient while the regulation is, by design, inefficient.<sup>15</sup> And if an agency's goal is to provide an in-kind good such as housing, food, or medical care, as part of the overall distributive system, its goal is by design not efficient. CBA would inhibit the agency's pursuit of its assigned task.

### **3.2 Standard Distributive Weights are not Desirable; there are no Uniform Adjustments**

The discussion above assumed that the sole goal of an agency was to correct a market failure. It assumed away the problem of whether agencies should also pursue distributive goals. Suppose now that we have an agency that is applying CBA because it is correcting a market failure. The question is whether the agency should also consider distributive effects.

A straightforward conclusion from the optimal tax literature is that agencies should not use standard distributive weights. Once we have an income tax, we do not use standard distributional weights even under general models without restrictive assumptions. If agencies are to consider the distributional effects of their policies, it will not involve adjustments to policies based on the simple intuition that regulations should be adjusted to be pro-poor. Instead, agencies

---

15 Sunstein argues that even in this context, CBA may be useful as a device to generate information and to prevent behavioral biases (Sunstein 2006; Sunstein 2000). To evaluate these claims in this context, we would have to compare the benefits Sunstein lists with the costs of giving an agency instructions that are contrary to its assigned goal.

would have to make the sorts of adjustments to goods and factor markets that were discussed above, such as increasing the price of complements to leisure.

### 3.3 Other Adjustments are Unlikely to be Desirable

Should agencies use the more subtle types of distributional considerations discussed above when they perform CBA? There are three problems with trying to use CBA this way.

#### *Wrong Markets*

The first is that the markets that we should use for distributive purposes are not necessarily the markets that we regulate. The two—the places where there are market failures and the markets that are desirable to use for distributive purposes—are effectively unrelated which means that CBA in regulated markets will only randomly be able to provide the types of market adjustments that are desirable. In many cases, CBA will be performed in markets where we should not make distributional adjustments even if it appears that there are substantial distributive effects of the regulation. For example, when we relax the assumption of weak separability from the benchmark model, we want to intervene in markets where the goods are complements to or substitutes for labor. Whether there is a market failure bears no relationship to whether the good produced in that market a good is a complement or substitute for labor. We cannot make the determination by observing whether the regulation has large distributive effects.

To illustrate, suppose there are three goods, good 1, good 2, and everything else, aggregated into good 3, and an agency is regulating good 1 because there is a market failure. The market for good 2 functions perfectly and is not regulated. If the agency uses unweighted CBA to set its regulatory policy, it will fix the market failure. Good 1 will now be on the same footing as good 2: in both cases, marginal costs will equal marginal benefit.

Now consider adjusting one or more of the markets to take distributional considerations into account. We might, for example, want to increase the price of either good 1 or good 2 if it is a complement to leisure or reduce its price if it is a substitute for leisure (both relative to good 3). There is no connection between there being a market failure in good 1 and it being a complement to or substitute for leisure. It is just as likely that good 2 is the right good to use for distributive purposes. Merely because we happen to regulate the market for good 1 does not mean it is a desirable good or market to use for distributive policy.

Observing that the distributive effects of unweighted CBA in the good 1 market are large and bad tells us nothing about whether it is desirable to use that market for distributive purposes. Like with the luxury tax and restaurant meals, we cannot know by looking only at distributive effects whether the

market is a desirable one to use for distributive purposes. We may want to subsidize good 1 because of how that subsidy interacts with the income tax even if the distributive effects on their own seem to go in the wrong direction.

One reason that people may view CBA as an appropriate tool for distributive goals is the availability heuristic. An agency issuing a regulation is doing something that people can see. Our attention is drawn to this action. If it has undesirable distributive effects, it is an action some person or agency is taking that seems to be causing those effects. On the surface, we might think that the person or agency should not take that action. If the market causes identical effects, there is no apparent actor and no one to blame. Yet the effects are the same and the tools we use to address them should be the same.

An alternative way to frame the argument is to suppose that CBA should be adjusted for distributive effects and to consider the implications. Unless we think that the best markets to use for distributive purposes happen to be the markets we regulate using CBA, there would be no reason to limit distributionally adjusted CBA to the markets we are already regulating. Systematic use of CBA for distributive purposes would require that agencies regulate for distributive purposes even in perfectly functioning markets because considering distribution is now part of their mandate. Agencies should *create* market failures to improve the distribution of income. In our example, the agency responsible for the perfectly functioning good 2 market should consider issuing regulations which create market failures but which have good distributive effects. This, after all, is the same thing as deviating from CBA when we just so happen to be regulating.

For example, suppose the EPA has jurisdiction over a market where there is no externality because pollution is well controlled via market mechanisms. Perhaps because of the way that property rights are defined or the way that the market works, firms in the industry have to pay to dispose of their waste products. The market is akin to the good 2 market above. If we want the EPA to consider distributional issues, we might want the EPA to *require* excess pollution if the distributive effects are good. This is identical to asking the EPA to regulate less because of distributional concerns in markets where there is too much pollution.

#### *Uncertain Size and Direction*

The second implication of the optimal tax literature is that agencies are not likely to know the size or even the direction of the appropriate adjustments to make to take distributive concerns into account. Conventional distributive weights are a function of the marginal utility of income  $dU^j/dY_i$  and the social marginal welfare of utility,  $\partial W/\partial U^j$ . They are pro-poor because marginal utility is declining in income and if the social welfare function is egalitarian,

marginal social welfare is declining in utility. Even if agencies cannot determine the precisely value of these variables, they know the direction: regulations should generally be made more pro-poor to account for distributive concerns.

Adjustments due to optimal distributive policies, however, depend on subtle factors such as whether a good is a substitute or complement to labor or an indicator of ability. They would not necessarily be pro-poor and might appear to be pro-rich. The example of restaurant food compared to grocery store food illustrates. The adjustments—exempting expensive restaurants from tax and having additional taxes on grocery store food—on their surface appear to be pro-rich. They improve the distribution of income because of how they interact with the income tax: they allow the income tax to increase the extent of distribution by lowering the distortions from progressivity.

Agencies are not likely to be able to determine the direction of the adjustments to make in this context. If they are going to take distributive concerns into account, they will have to determine whether a particular type of consumption associated with the regulated market is a complement to labor or a substitute for labor, is an indicator of ability, or for some other reason is good to use for distributive purposes. This is a subtle determination and is not one that agencies have any particular ability to evaluate. Crude guesses, such as exempting grocery store food and taxing restaurant food, may go in the wrong direction.

### *Wrong Tools*

The third implication from the optimal tax literature is that, conditional on regulating in the right market and on knowing the direction and possibly the size of the adjustments to market prices to make, agencies will not have the best tools for making the adjustments.<sup>16</sup> The reason is that it is likely to be more efficient to use a price mechanism than a regulatory mechanism. Regulations effect product attributes. They determine how safe a product must be, what sort of information must be provided to users, how the product may be used, and so forth. To increase or reduce the price of the good, the regulation would alter the product attributes. A tax or subsidy on a good lets the market determine the right mix of attributes and, therefore, avoids this source of inefficiency.

To illustrate, suppose we are regulating in a market that is an appropriate market for distributive purposes. Suppose also that we know the direction and size of the appropriate adjustment. Consider two alternatives for making the adjustment: adjusting the regulation so that the attributes of the product are changed in a way that makes it more or less expensive in the correct amount and taxing or subsidizing the good and letting the market determine the product

---

<sup>16</sup> The appendix in Kaplow and Shavell (1994, pp. 679–680) makes this point.

attributes. Changing the product attributes adds an additional inefficiency that is not required by distributive concerns; a tax can achieve those concerns without altering the product attributes.

For example, libraries are complements to leisure and, therefore, are potential markets for intervening for distributive purposes. We could regulate them to increase their price, say by forcing them to carry only unpopular books, having dingy interiors, or limiting the hours that they are open. Alternatively, we could impose a fee and let the market determine which books, interiors, and hours are the best. The fee can raise the price of libraries in the same amount as the regulation but without creating the secondary problems.

#### *Summary*

To return to the basic framing of the problem of one of task assignments, we can ask when it is desirable to assign part of the distributive task to an agency which is otherwise correcting a market failure. Most of the time, it will not be appropriate to assign an agency part of the task of redistributing. The agency would have to be regulating in a market in which it is desirable to adjust the price for distributive purposes. The agency would have to know which direction to make the adjustment and could not just make conventionally pro-poor adjustments. It would have to have some estimate of the appropriate size of the adjustments. And it would have to be better to make the adjustment via regulation of product attributes than via a tax or subsidy.

It is not impossible that these conditions will be met, so we cannot rule out a version of distributive adjustments to CBA based on a priori reasoning. We can say, however, that it is unlikely that they are met often. Most importantly, we can rule out conventional distributive weights or uniform weights of any sort.

## **4. LEGISLATURES**

The argument so far has considered the allocation of tasks to agencies. There is a more subtle and more complex problem, which is how agencies interact with legislatures. This is a broad topic. I focus here on the claim that the tax and transfer system will not redistribute sufficiently because of political constraints, so agencies need to use distributive weights as a substitute.

To see the basis of the claim, recall that in the benchmark model, we could eliminate price distortions in particular markets for goods or factors, and make an offsetting adjustment to the labor income tax to keep the distribution the same but make the tax system more efficient. Getting rid of the luxury tax and adjusting the marginal rate schedule under the income tax retains the distributive benefits of the luxury tax but reduces the resulting inefficiencies. If the tax and transfer system does not respond this way, the argument does not go



through. Just getting rid of the luxury tax without adjusting the income tax may not be desirable.<sup>17</sup> Perhaps, in this case, we should use conventional distributional weights for CBA.

We can divide this claim into what I will call naïve versions and sophisticated versions. The naïve versions have a number of elements. The primary element is an assertion that the tax system does not, or is unlikely to, change to address distributive concerns.

This claim is flat out contradicted by the facts. As [Graetz \(1995\)](#) discusses, tax systems are constantly changed, most often with great focus on distributional issues. In the USA, the changes are made in fantastic detail to respond to particular groups that are aggrieved. Credits, deductions, elections, alternative tax schedules, and detailed tax rules are changed to ensure that various groups are given the distributive treatment that Congress deems appropriate. For example, in 2008, the National Taxpayer Advocate reported that in the seven years between the start of 2001 and 2008, there were more than 3,250 changes to the tax law and there were more than 500 changes in 2008 alone. *Forbes Magazine* estimated in 2013 that there had been 5,000 since 2001.<sup>18</sup> Senator Baucus, at the time, Chairman of the Finance Committee, claimed in 2012 that since 1986, we had made more than 15,000 changes to the tax code.<sup>19</sup> Even if only a modest portion of these changes were distributive in nature and even if the counts are exaggerations, the number of adjustments would be high. Blanket assertions that the tax system does not respond to distributional concerns are false.

Moreover, the argument that the tax system does not respond requires a regulator to know this at the time the regulator issues a regulation. The regulator under this approach uses distributive weights only when the tax system does not respond. If the tax code never responded, it would be easy for a regulator to know always to use distributive weights. If, however, it sometimes responds, as is the case, it is hard to see how a regulator would know when the tax system will respond and when it will not. The tax system cannot respond before a regulator has issued the regulation, so at the time the regulator is considering the regulation, there is no way to make this determination.

Commentators may assert that the tax law does not adjust because the tax law does not respond immediately to a given regulation and because the effects of particular regulations are most often not cited in the legislative history of tax changes. We often cannot tie the regulatory change and tax law

---

17 When the USA repealed the luxury tax in 1993, however, it raised the top marginal tax bracket, so the example shows that in fact these sorts of adjustments do happen, at least in a rough fashion.

18 <http://tinyurl.com/p22ok5q>

19 <http://tinyurl.com/cyrbyj7>

changes together. We see a regulation with bad distributive effects but cannot easily find the offsetting tax adjustment.

There may, however, be hundreds of regulations which have distributive effects, all of which have to be considered when designing the tax system. The net effect may be very different than the effect of any single regulation. The distributive effects of the tax laws are determined using data on the existing distribution. The net effect of all prior regulations will be reflected in the aggregate data whether those regulations are referenced or not ([Joint Committee on Taxation 1993](#); [Cronin 1999](#)). Congress then considers whether and how to adjust the progressivity of the tax system in light of this information. Failure to find a direct link for a particular regulation is not evidence that the tax law does not reflect the distributive effects of that regulation.<sup>20</sup>

The second element of the naïve version is to assert that if the tax system does not change, regulators should then be free to make distributive adjustments on their own. They would, under this version, make distributive determinations without making any reference to (1) the allocation of power in a given government and (2) whether the regulator has or is likely to have the relevant information necessary to make the adjustment. That is, the naïve claim is often an acontextual, non-institutional claim that in the absence of adjustments to the tax system, regulators are free to do as they wish.

Here is an admirably bald statement of the claim from [Adler \(2013, p. 28\)](#):

Imagine that a decisionmaker has the power to choose between P and P\*. If the tax code does not change, her morally preferred [social welfare function] favors P over P\*. Although P\* bundled with a change to the tax code is Pareto-superior to P, the decisionmaker does not believe the legislature will make this tax change. She is therefore morally justified in picking P.

The argument is devoid of any theory of government. It makes no reference to which entities or individuals in a particular type of government are allocated which responsibilities. Decision makers are instead morally justified in overruling legislative outcomes simply because they disagree. The argument does not even make it a requirement that the social welfare function used by the

---

20 Another reason people may make the assertion that the tax system does not adjust is that the tax system may not meet the distributive preferences of an individual commentator or scholar. If this is true, it shows only that the particular commentator's views about redistribution deviate from the equilibrium produced by the electoral process.

decision maker be generally accepted by the population or be justified in any particular way other than it be morally preferred by a particular person. It does not require the decision maker to have any particular expertise in distributive issues or have moral views that reflect a general consensus of voters.

The statement supports actions by regulators to make the overall system regressive. A regulator may believe that the legislature redistributes far too much. The regulator, implementing his preferred social welfare function, would adjust his regulations to help the rich. Adler would presumably support this action because the regulator is morally justified in picking his preferred social welfare function.

For the statement to be credible, it needs a theory of the allocation of responsibility in a structure of government which allows regulators to overturn legislators. This is not impossible. Some nations, possibly including the USA, might have such a structure.<sup>21</sup> The arguments, however, would have to be institutional and based on claims about the best assignment of tasks to parts of the government given disagreement about the desirable amount of redistribution. (And the argument would have to actually be made!) The normal course for Western democracies is for elected legislatures to be allocated the power to make the primary distributive judgments. In a government with separation of powers, such as the USA, perhaps there are arguments that this power is split up among branches. Regardless, one might hope for at least some story or argument or example for the proposed structure that allows conclusions like that quoted above. A blanket statement, devoid of institutional context, makes little sense.

The more sophisticated version of the “tax system doesn’t respond” argument considers the institutional capacities of different parts of the government. It starts with the task assignment considerations discussed above and then considers the possibility that a particular government at a particular moment in time for a particular policy might have features that cause us to want to deviate from the otherwise desirable assignment. Lee [Fennell and Richard McAdams \(2013\)](#) hint at such a possibility in a short essay on the subject, but they do not develop the argument in any detail.

---

21 One might tell a story about presidential systems like the USA that allows regulators to make distributive judgments. Regulatory agencies report to an elected president. The elected president is entitled to instruct them to reflect his distributive preferences. Congress is free to override those preferences through changes in the tax law. Congress ultimately decides on distributive issues but the President is allowed to move first, subject to override. Whether this is true, however, would depend on the particular constitutional structure of a given nation.

To begin a sketch of the argument, imagine that there is a demand for some amount of redistribution among the electorate. Let us say that ideally those demanding redistribution want to satisfy this demand as cheaply as possible because by doing so, they can get more redistribution or otherwise use the saved resources (as compared to pursuing less-efficient redistributive policies). Suppose, however, that the political system is such that a minority-blocking coalition can prevent such a policy. Those demanding redistribution might pursue a second or third best policy that cannot be blocked. Overall welfare might be improved conditional on not being able to dislodge the blocking coalition.

Several questions arise from this sketch. An initial and central problem is that in this story, like the naïve one discussed above, at the time they issue their regulations, regulators would have no idea whether the tax system is going to adjust. The timing of the story does not work.

Second, if the regulator knows in advance the tax system will not adjust, it does not follow that it is welfare improving to use traditionally weighted CBA. Traditional distributional weights are derived from a model that does not include an income tax. Even if our tax system does not adjust optimally, it exists and it is progressive. To determine the appropriate weights, we need to write down a model that includes the tax system and solve for the welfare-maximizing action for an agency. The sorts of interactions highlighted above such as complementarity with leisure and the considerations in the optimal progressivity literature stemming from [Mirrlees \(1971\)](#) will arise in this model. We cannot go directly from a claim that the tax system does not adjust to a claim that we can therefore use distributional weights without serious (or really any) analysis.

The third question is why the blocking coalition faced with the inevitability of inefficient redistribution would not then bargain to have a like amount of redistribution done efficiently. That is, there is a political version of the Coase theorem in the background, and a model of this type of redistribution has to explain why it fails ([Acemoglu 2003](#)).

A fourth question, also arising from a version of a political Coase theorem, is whether attempts to use CBA to redistribute are futile. Suppose that we consider an executive branch agency that is subject to different electoral pressures than the legislature and therefore wants more (or just as likely, less) redistribution than the legislature. The executive agency cannot change the tax code, so it pursues a second best policy of distributionally adjusted CBA (or the opposite, making regulations regressive). The question is, if the executive pursues this strategy, whether the legislature can undo it by changing the progressivity of the tax code. Said another way, there is some overall demand for redistribution and the ultimate amount is determined via bargaining in the political process. A single actor, such as an executive agency, may not be able to undo

that bargain. Or, at a minimum, a claim that an agency can change that outcome needs a story explaining why and how.<sup>22</sup> For example, status quo biases, veto power, and other details of the policy-making process may allow changes to the equilibrium by one actor in the system.

Specifying the exact story and determining when it applies is important. If the executive uses inefficient regulation to pursue distributive goals and the legislature responds with offsetting tax adjustments, we end up with the same overall distribution of income or wealth but will have inefficient regulations. We will be regulating pollution, safety, information, and so forth too much or too little because of distributive goals, but because of the tax adjustment there is no overall distributive effect. Individuals, possibly including the poor, are overall worse off.

Finally, there are a set of very difficult questions about how we overlay normative evaluation on a story that is at its core, political. Once we embed the task assignment arguments in politics, we can no longer claim that the actors are pursuing the maximization of social welfare. After all, the story posits that the legislature does not use the tax system optimally. Instead, actors are responding to some mix of their constituent's demands and their own preferences. We cannot simply posit a welfare maximizing agency using distributionally weighted CBA and a legislature failing to maximize welfare because it, say, is focused on serving special interests. Given a consistent story about the motivations of the various actors, we might then ask what sort of institutional setup best maximizes welfare.

It is possible that one can put together a convincing story that embeds the task assignment claims in politics and recommends that an agency *should* sometimes use distributional weights and that these weights are the pro-poor weights of the sort conventionally imagined. Any such story, however, would be based on the particular circumstances: a particular type of regulation in a particular political context. It would not support the use of general distributive weights in CBA. For example, if the legislature redistributes too much, the story would flip and regulators should weight regulations to be pro-rich to maximize social welfare. If after the next election the legislature redistributes too little, regulators would have to flip the distributional adjustments. And so on. The "tax system does not respond" argument is not an excuse to throw out considerations of optimal distributive policy in favor of conventional distributional weights.

---

22 The status quo may change if an agency uses a regulation to redistribute and a presidential veto may enable him to preserve the status quo when he could not otherwise force the legislature to change the prior level of distribution. Some legislative changes, however, may not require a veto. For example, the legislative can refuse to change the law when the economic system changes the level of inequality that it produces.

## 5. CONCLUSION

CBA is part of the set of instructions we give agencies performing assigned tasks, such as regulating pollution, policing the borders, or ensuring the safety of the banking system. We cannot determine how CBA should be conducted without reference to the specialized nature of the tasks assigned to agencies. We cannot go directly from the overall goal, maximize social welfare, to the individual tasks given to particular agencies.

To understand how distributive tasks should be assigned, we need to start with the design of distributive systems. The central tool of distributive systems is the income tax. Adjustments away from efficiency in product or factor markets arise because of their interaction with the income tax. The adjustments do not look like generic distributive weights.

When we assign the tasks of pursuing this distributive policy to units within the government, we do not get conventional distributionally adjusted CBA. Moreover, it is difficult to imagine that we would want very many, if any, of the more subtle distributive tasks assigned to specialized agencies regulating in particular markets. There is no reason to believe that these agencies are regulating in markets that are good to use for distributive purposes, that these agencies would know the size or even the direction of the appropriate adjustments, or that these agencies would have the right set of tools to use (price mechanisms) to make the adjustments.

Given the flexibility of task assignment and the number of considerations involved, we cannot and does not rule out the possibility that a particular agency working in a particular market at a given point of time should consider distributive concerns. Such examples, however, should not obscure the more general point, which is that there should be any generic adjustments to cost-benefit analysis to account for distributive concerns.

## REFERENCES

- Acemoglu, Daron. 2003. Why Not a Political Coase Theorem? Social Conflict, Commitment, and Politics. *J. Comp. Econ.* **31**, 620–652.
- Adler, Matthew. 2013. *Cost-Benefit Analysis and Distributional Weight: An Overview*. Duke Environmental and Energy Economics Working Paper Series EE 13-04.
- Adler, Matthew D., & Eric A. Posner. 1999. Rethinking Cost-Benefit Analysis. *Yale L. J.* **109**, 165–247.
- Adler, Matthew, & Eric Posner. 2006. *New Foundations of Cost-Benefit Analysis*. Cambridge, MA: Harvard University press.

- Akerlof, George A. 1978. The Economics of 'Tagging' as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning. *Am. Econ. Rev.* **68**, 8–19.
- Atkinson, Anthony B., & Joseph E. Stiglitz. 1976. The Design of Tax Structure: Direct versus Indirect Taxation. *J. Pub. Econ.* **6**, 55–75.
- Bankman, Joseph, & David A. Weisbach. 2006. The Superiority of an Ideal Consumption Tax over an Ideal Income Tax. *Stan. L. Rev.* **58**, 1413–1456.
- Besley, Timothy, & Stephen Coate. 1992. Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty-Alleviation Programs. *Am. Econ. Rev.* **82**, 249–261.
- Boadway, Robin. 1974. The Welfare Foundations of Cost-Benefit Analysis. *Econ. J.* **84**, 926–939.
- . 1976. Integrating Equity and Efficiency in Applied Welfare Economics. *Q. J. Econ.* **90**, 541–556.
- Boadway, Robin W. 1975. Cost-Benefit Rules in General Equilibrium. *Rev. Econ. Stud.* **42**, 361–374.
- Bronsteen, John, Christopher Buccafusco, & Jonathan Masur. 2013. Well-Being Analysis vs. Cost-Benefit Analysis. *Duke L. J.* **62**, 1603–1689.
- Chandler, Alfred D. 1962. *Strategy and Structure: Chapters in the History of the Industrial Enterprise*. Cambridge, MA: MIT Press.
- Christiansen, Vidar. 1981. Evaluation of Public Projects under Optimal Taxation. *Rev. Econ. Stud.* **48**, 447–457.
- Cronin, Julie Anne. 1999. *U.S. Treasury Distributional Analysis Methodology*; Office of Tax Analysis Paper 85. <http://www.treasury.gov/resource-center/tax-policy/tax-analysis/Documents/ota85.pdf>.
- Cullis, John, & Philip Jones. 1998. *Public Finance and Public Choice*, 2nd edn. New York: Oxford University Press.
- Dreze, Jean, & Nicholas Stern. 1987. The Theory of Cost-Benefit Analysis. In Alan J. Auerbach, Martin Feldstein, & Michael D. Intriligator (eds.), *Handbook of Public Economics*. New York: North-Holland, pp. 909–990.
- Fennell, Lee Anne, & Richard H. McAdams. 2013. Introduction. In Lee Anne Fennell, & Richard H. McAdams (eds.), *Fairness in Law and Economics*. Cheltenham, UK: Edward Elgar.
- Gauthier, Stéphane, & Guy Laroque. 2009. Separability and Public Finance. *J. Pub. Econ.* **93**, 1168–1174.
- Golosov, Mikhail, Maxim Troskin, Aleh Tsyvinski, & Matthew Weinzierl. 2013. Preference Heterogeneity and Optimal Capital Income Taxation. *J. Pub. Econ.* **97**, 160–175.
- Graetz, Michael. 1995. Paint-by-Numbers Tax Lawmaking. *Col. L. Rev.* **95**, 609–682.

- Gruber, Jonathan. 2005. *Public Finance and Public Policy*. New York, NY: Worth Publishers.
- Hahn, Robert W., & Cass R. Sunstein. 2001. A New Executive Order for Improving Federal Regulation? Deeper and Wider Cost-Benefit Analysis. *Univ. Pen. L. Rev.* **150**, 1489–1552.
- Hylland, Aanund, & Richard Zeckhauser. 1979. Distributional Objectives Should Affect Taxes but Not Program Choice or Design. *Scand. J. Econ.* **81**, 264–284.
- Iorwerth, Aled Ab, & John Whalley. 2002. Efficiency Considerations and the Exemption of Food from Sales and Value Added Taxes. *Can. J. Econ. /Revue Canadienne D'économie* **35**, 166–182.
- Joint Committee on Taxation. 1993. *Methodology and Issues in Measuring Changes in the Distribution of Tax Burdens*. Washington, DC: U.S. Government Printing Office.
- Kaplow, Louis. 2004. On the (Ir)Relevance of Distribution and Labor Supply Distortion to Government Policy. *J. Econ. Persp.* **18**, 159–175.
- . 2006a. Public Goods and the Distribution of Income. *Eur. Econ. Rev.* **50**, 1627–1660.
- . 2006b. On the Undesirability of Commodity Taxation Even When Income Taxation Is Not Optimal. *J. Pub. Econ.* **90**, 1235–1250.
- . 2008a. *The Theory of Taxation and Public Economics*. Princeton: Princeton University Press.
- . 2008b. Optimal Policy with Heterogeneous Preferences. *BE J. Econ. Anal. Pol.* **8**
- . 2012. Optimal Control of Externalities in the Presence of Income Taxation. *Int. Econ. Rev.* **53**, 487–509.
- Kaplow, Louis, & Steven Shavell. 1994. Why the Legal System Is Less Efficient than the Income Tax in Redistributing Income. *J. Legal Stud.* **23**, 667–681.
- . 2000. Should Legal Rules Favor the Poor - Clarifying the Role of Legal Rules and the Income Tax in Redistributing Income. *J. Legal Stud.* **29**, 821–835.
- Kleven, Henrik Jacobsen. 2014. How Can Scandinavians Tax So Much? *J. Econ. Persp.* **28**, 77–98.
- Kocherlakota, Narayana R. 2010. *The New Dynamic Public Finance*. Princeton, NJ: Princeton University Press.
- Laroque, Guy R. 2005. Indirect Taxation Is Superfluous under Separability and Taste Homogeneity: A Simple Proof. *Econ. Lett.* **87**, 141–144.
- Mankiw, N. Gregory, Matthew Weinzierl, & Danny Yagan. 2009. Optimal Taxation in Theory and Practice. *J. Econ. Persp.* **23**, 147–174.
- Mirrlees, James. 1971. An Exploration in the Theory of Optimum Income Taxation. *Rev. Econ. Stud.* **38**, 175–208.



- Nichols, Albert L., & Richard J. Zeckhauser. 1982. Targeting Transfers through Restrictions on Recipients. *Am. Econ. Rev.* **72**, 372–377.
- Nussbaum, Martha C. 2000. The Costs of Tragedy: Some Moral Limits of Cost-Benefit Analysis. *J. Legal Stud.* **29**, 1005–1036.
- Office of Management and Budget. 2003. *Circular A-4*. [www.whitehouse.gov/omb](http://www.whitehouse.gov/omb).
- Piketty, Thomas, & Emmanuel Saez. 2013. A Theory of Optimal Inheritance Taxation. *Econometrica* **81**, 1851–1886.
- Saez, Emmanuel. 2002. The Desirability of Commodity Taxation under Non-Linear Income Taxation and Heterogeneous Tastes. *J. Pub. Econ.* **83**, 217–230.
- Samaha, Adam M. 2007. What Good Is the Social Model of Disability? *Univ. Chicago L. Rev.* **74**, 1251–1308.
- Sunstein, Cass. 2002. *The Cost-Benefit State, The Future of Regulatory Protection*. USA: American Bar Association.
- Sunstein, Cass R. 2000. Cognition and Cost-Benefit Analysis. *J. Legal Stud.* **29**, 1059–1103.
- . 2006. The Availability Heuristic, Intuitive Cost-Benefit Analysis, and Climate Change. *Climatic Change* **77**, 195–210.
- Tresch, Richard W. 2002. *Public Finance, A Normative Theory*, 2nd edn. United States: Academic Press.
- Weisbach, David A. 2009. Toward a New Approach to Disability Law. *Univ. Chi. Legal F.* **2009**, 47–102.
- Weisbach, David, & Jacob Nussim. 2003. The Integration of Tax and Spending Programs. *Yale L. J.* **113**, 955–1028.
- Williamson, Oliver. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications*. Philadelphia: University of Pennsylvania.