

A Note on the Optimal Supply of Public Goods and the Distortionary Cost of Taxation

Louis Kaplow*

Abstract

In a recent article, I demonstrated that, under standard simplifying assumptions, it is possible to finance a public good in a manner such that a Pareto improvement results whenever the simple cost-benefit test is satisfied -- that is, without any adjustment for the "marginal cost of funds." In particular, the method of finance involves adjusting the income tax so that the combined incidence of the tax adjustment and the public good is distribution-neutral. One implication of this result is that, if the public good is financed in some other manner, the difference in outcome will be purely redistributive, so that any change in distortionary costs will be accompanied by an opposing change in redistributive benefits. I also showed how this analysis is applicable to determining the optimal level of environmental taxes.

Edgar Browning and Liqun Liu have written a critique of my article. It does not, however, disagree with any of these fundamental claims. Instead, their argument focuses on how one should interpret the term "distortion." It should not be surprising, however, that under any interpretation of the term -- including their preferred one -- my conclusions about how policy analysis should be conducted continue to be correct.

Forthcoming, National Tax Journal (March 1998)

JEL Classes H21, H23, H41, H43

* Harvard Law School and the National Bureau of Economic Research. I am grateful for comments from Steven Shavell and Al Warren. Throughout the note, I find it convenient not to restate and cite the extensive prior literature that is described in my original article. Nor do I review here the respects in which my argument has strong roots in others' prior work or the various qualifications to my main points (which are not relevant either to the prior literature or to Browning and Liu's comment).

My article advanced an unconventional view of how policy analysis of public goods (and corrective taxes)¹ should be conducted in a world with a preexisting income tax that is distortionary and redistributive. In simple terms, my claim had two related components: (1) It is possible, under standard simplifying assumptions, to finance public goods in a manner such that social welfare will rise (indeed, there will be a Pareto improvement) whenever the basic cost-benefit test is met. That is, the basic cost-benefit test, without any adjustment for the “marginal cost of funds,” suffices. The particular method of finance under which this conclusion holds is one in which the income tax is adjusted such that the incidence of the tax adjustment matches the incidence of the benefits of the public good -- a combination that is, in the aggregate, distribution-neutral.

(2) If the public good is financed in a different manner, then the conclusion just offered would have to be modified to take into account opposing distortionary costs and redistributive effects. Thus, when the method of finance increases redistribution, as in many examples offered in the literature on public goods provision, there will be a distortionary cost that should be taken into account. But proper welfare analysis would require that the redistributive benefit not be ignored. Similarly, if there were a reduction in distortionary costs, this would typically arise because the public good and method of finance, taken together, reduce the extent of redistribution, so again one would have to weigh distortionary costs and redistributive benefits before determining how, if at all, to modify the basic cost-benefit test. Also note that, as a conceptual and practical matter, these latter distortionary and redistributive effects can be seen as independent of providing the public good, for one can produce (or offset) these effects by adjusting the tax system without regard to whether the public good is provided.

These conclusions and the reasoning behind them -- both unaffected by Browning and Liu's arguments -- will be reviewed in the first section of this note. It is useful at the outset, however, to reflect on the implications of these claims. Consider the familiar examples in the public goods provision literature in which the stipulated method of finance does involve an additional distortionary cost. Or consider Browning and Liu's central claim that in some of my examples it would be possible to finance the public good in a different manner (namely, a lump-sum tax), such that there would be less distortion. Both of these results are entirely valid. But they are also misleading in two respects that I emphasized in my original article.

First, both results ignore that the identified distortionary costs or benefits are accompanied by distributive benefits or costs that run in the opposite direction. *Hence, the implicit or explicit policy prescription that one should adjust redistributive taxation in the manner that minimizes distortion -- which is only one of the two relevant components of economic welfare in this context -- is inapt.* Second, suppose that we (or, perhaps more relevant, the government decisionmaker) had a particular view concerning the appropriate balance between distortion and redistribution; moreover, assume that the view supported, say, the sort of reduction in distortion (and redistribution) that Browning and Liu seem to favor. My original analysis shows that one could reduce redistribution directly: in Browning and Liu's examples, by decreasing the preexisting labor income tax rate and increasing the lump-sum tax. *Providing more or less of the public good than indicated by the basic cost-benefit test, as they suggest, is a Pareto dominated strategy.* Indeed, this was one of the main points in my article, and it was proved in my original appendix.

Unfortunately, none of the preceding points are highlighted in Browning and Liu's extensive comment. Instead, they focus on the meaning of the notion of "distortion," disagreeing

with what they believe to be my understanding of the term. They show that, under their view of distortion, it is always possible to reduce distortion by using more lump-sum taxation and less distortionary income taxation. I, of course, do not disagree with this point. As I explain in the second section of this note, however, *their repeated (and obviously correct) statement that reducing the extent of redistributive income taxation will reduce distortion has no relevance to any of my substantive points.* The remainder of their discussion is largely semantic. For what it is worth, I wish to concede that my choice of language at some points in my original article does give rise to possible confusion, and thus I bear some responsibility for the misunderstanding that seems to underlie Browning and Liu's comment. Nevertheless, it should be noted that the language of "distortion," on which they focus so heavily, appears nowhere in my formal appendix, in which my main claim is rigorously stated and proved.

ANALYZING PUBLIC GOODS AND DISTORTIONARY TAXES

Central Analytical Claim

Put as precisely as possible, the central analytical claim of my original article is as follows:

If utility is weakly separable in labor (leisure) . . . , then any marginal change in [the level of the public good] g that produces positive unweighted net benefits -- i.e., for which the integral ("sum") of individuals' marginal rates of substitution (v_g/v_c) exceeds the cost -- will result in a Pareto improvement when financed through an appropriate change in the tax schedule T . (p.531; emphasis in original)²

That is, any change in the level of the public good that passes the basic cost-benefit test will result in a Pareto improvement if financed in a particular manner. This adjustment to the tax schedule, T , it turns out, is one that has an incidence that matches the incidence of the public good, so that the tax adjustment and the public good, *taken together*, are distribution-neutral.

(For example, if the public good provides a uniform monetary benefit, the tax adjustment adds a uniform dollar amount to everyone's tax bill; if the public good provides a benefit that is proportional to income, the tax adjustment consists of raising the income tax rate.)³

Why does this result follow? Much of my original article was devoted to presenting the intuition, offering examples, and, finally, presenting a formal proof, so I will not fully rehearse the analysis here. The basic idea, however, is fairly simple; it can be understood by examining the implications of financing a public good in a distribution-neutral manner. By definition, this construct involves changing the tax schedule to offset the benefit of the public good *at each level of income*. Hence, after the government provides the public good and imposes this offsetting tax adjustment, every individual -- whatever his wage (ability) level -- will get precisely the same utility from every choice of labor effort as he got before the reform. It follows that (1) labor supply will be unchanged, (2) each individual's utility will be unchanged, and (3) the government will have a surplus (or deficit) precisely equal to the excess of benefits over costs. Hence, if the simple cost-benefit test is met, there will be a surplus. The final step in the argument is to make a further adjustment in the tax schedule to rebate the surplus, say, pro rata. Now, the government's budget is balanced, and every individual has higher utility: a Pareto improvement.

Although I discuss Browning and Liu's remarks primarily in a later section, it is worth noting one of their main allegations at this point in the discussion. They state that I focus on the fact that labor supply is unchanged. This is correct, and the reason I do this should now be apparent: by showing that labor supply does not change from the distribution-neutral tax adjustment, I am able to prove that there is a government surplus whenever the cost-benefit test is met. And this, in turn, shows that distribution-neutral finance allows one to achieve a Pareto improvement, establishing my central claim.⁴

Application to Policy

It follows from my central analytical claim that, if finance is distribution-neutral, the basic cost-benefit test is correct. As I suggested in my original article, there are grounds for believing that, roughly, on average, and over the long run, government finance may be distribution neutral. To be sure, governments may change the extent of redistribution, but this can be and often is done without regard to marginal public goods decisions.

Nevertheless, distribution-neutral finance is hardly guaranteed and, in any event, it is of great interest to consider the question of what is the proper method of analysis when finance is not distribution-neutral. Accordingly, I devoted a section of my original article to this question (a section that Browning and Liu do not discuss). The main construct is described as follows:

. . . consider the prospect of supplying a public good that is to be financed by some method T (where T refers to the adjustment to the income tax and transfer payment system that provides the revenue for the project). Decompose this reform into two components:

1. Finance the public good by [a distribution-neutral] tax adjustment, T_0 .
2. Reform the tax system, changing the tax scheme from T_0 to T. (p. 519)

In principle, it should be clear that *any* public good financed by *any* method can be decomposed into this two-step reform. The first step corresponds to the method of finance described previously: one chooses the tax adjustment with the same incidence as the public good so that the combination of the two is distribution-neutral. The second step is a pure, revenue-neutral tax reform: its only effect is redistributive.

This technique of thinking about public goods provision is useful because it clarifies analysis. Analyzing the first step should now be straightforward: if a distribution-neutral tax adjustment is used to finance the public good, then the basic cost-benefit test tells the entire story.

If there is to be anything more to say about the reform as a whole, the action must be in the second step. What is this second step? It involves no change in the level of the public good and, correspondingly, no change in government revenue. *It is a purely redistributive change.* How do we evaluate such a change? The answer is familiar: there is, in general, a tradeoff between the extent of redistribution and the magnitude of distortion, the “equity/efficiency tradeoff”. We -- or the government decisionmaker -- will wish to know the extent of the redistributive and distortionary effects and then will need to make some judgment about the balance. Conventional analysis (just like the argument of Browning and Liu), therefore, which focuses only on whether distortion rises or falls, does not even tell us the *direction* of the appropriate adjustment, if any, to the basic cost-benefit test.

A point of emphasis in my original article -- indeed, my main criticism of much of the literature on public goods provision -- was that prior work focused on the distortionary effect while ignoring the redistributive effect. Yet the primary motivation for distortionary labor income taxation (rather than, say, a uniform lump-sum tax -- a head tax) is that we care about distribution. This omission of distributive effects from policy analysis is serious. Of course, my point is not to criticize analysts who specialize in their work by studying only the distortionary costs of taxation. *Rather, when the change in distortionary costs under investigation arises directly from a change in the extent of redistribution, I argued that it is necessary to present the complete picture.*

It is worth considering briefly why it is that economists have found it so easy to focus on distortionary costs while ignoring redistributive effects in the present context. As I suggested in my article (p. 516 and n. 14), one possibility is due to the frequent use of so-called “representative individual” models, in which there is only one type of person. Here, this means,

in essence, that everyone has the same wage (ability) and hence earns the same income and enjoys the same utility. In this sort of model, of course, there cannot be any redistributive effects. But there can be distortion. Thus, as Browning and Liu are able to show, one can make identical taxpayers better off by moving from labor income taxation to uniform lump-sum taxation.

This sort of model in which all individuals are identical is useful for many purposes. But there is one purpose for which it cannot be employed -- namely, analyzing problems in which redistribution is central. As already noted, there can be no redistribution in a world where all are identical. This is why the model in my paper specifically assumed that individuals differ in their wages (and hence in their income and utility levels). In such a model, Browning and Liu's repeated suggestion of moving from labor income taxation to uniform lump-sum taxation does not make everyone better off. It makes the rich better off and the poor worse off, just as we would expect. Therefore, as I emphasized in my original article, it is extremely important *not* to use reasoning from models in which everyone's income is identical when an important question (in my step 2, the *only* question) is how to trade off distortion and redistribution. *Yet, despite my warnings, this is exactly what Browning and Liu do at almost every point in their article and in the entirety of their appendix.*

There is, however, one instance where Browning and Liu do discuss a case in which two individuals have different income levels. Yet in this discussion, they "use two individualized lump sum taxes (A's lump sum tax is twice B's)," so that redistribution can be accomplished at no distortionary cost. But, as is familiar -- and as I reminded the reader in my original article (n. 14) -- one should avoid such a "confusion about lump-sum taxation": only *uniform* lump-sum taxes are generally believed to be feasible as a practical matter. If *individualized* lump-sum taxes could be used, as they do in this example, distortionary income taxation would be wholly

unnecessary and the entire subject under discussion would be moot! Had Browning and Liu appropriately confined their discussion of the two-income-level case to uniform lump-sum taxes, as has been the practice in decades of literature on redistributive taxation, the poor individual in their example would have been made worse off by their proposed substitution of lump-sum taxation for redistributive income taxation, and, therefore, their argument that everyone would be better off would have failed.

In summary, it is most important in examining the distortionary cost of financing public goods to take into account distributive effects because, in this setting, distortionary and distributive effects go hand in hand. When finance (combined with the incidence of the public good) is distribution-neutral, the basic cost-benefit test is applicable; conversely, when the cost-benefit test is inapplicable, we know that the policy change must not be distribution-neutral, so distributive effects must be considered.⁵ (I note that this theme, which emphasizes the impropriety of focusing on distortionary costs while ignoring redistribution, is also the focus of recent work by Agnar Sandmo (1997).)

ON THE MEANING OF “DISTORTION”

The bulk of Browning and Liu’s comment is devoted to examining the definition of “distortion.” Of course, the proper definition of this term -- and the question of whether my usage was incorrect or misleading -- cannot have any relevance to the validity of my analytical points or my claims about how proper policy assessment should proceed. Definitions merely pertain to the manner in which we communicate. In addition, upon reviewing my original paper - - which frequently uses the language of “distortion” -- I noticed that the language appears nowhere in my appendix, where my claims are formally stated and proved.

Nevertheless, it may be helpful to review their semantic arguments. I will follow Browning and Liu and consider the case they emphasize from my original paper: a public good that has benefits that are proportional to income (say, the more income one has to buy goods, the more valuable is police protection against theft). In this case, distribution-neutral finance would call for a tax that was proportional to income, a uniform increase in the labor income tax rate.

My claim about this case (which follows from the earlier discussion) was that *if such a finance mechanism is used, then there will be a welfare gain -- indeed, a Pareto improvement -- if and only if the basic cost-benefit test is met*. They never question this claim. Instead, they object to the manner in which I express some of the intuition offered in support of it.

First, viewing the labor income tax component *in isolation from the public good*, they argue -- at great length -- that the tax rate increase is distortionary. Naturally, I agree. But this fact is irrelevant to my argument, which was concerned entirely with the *combined* effects of a tax *and* the public good financed by the tax.

Second, they do ultimately comment on the labor supply effects of the public good. They acknowledge that this particular public good, which has benefits that increase with income, will tend to induce an increase in labor supply, absent an offsetting income effect (which, it turns out, is present in this case). To elaborate, suppose that theft reduced consumption by two percent and that increased expenditures on police would eliminate theft. Then, each dollar of earnings will buy about two percent more in effective consumption when police protection is provided. In other words, *the price of effective consumption falls* when the public good is provided.

Browning and Liu agree with all of this, but object to my calling this phenomenon a “substitution effect” because, they say, no “price” has changed. Although there is some merit in their characterization, it is hardly the only plausible one, since the price per unit of effective

consumption has changed.⁶ In any event, I would have been happy to avoid the “substitution effect” language and simply let the point stand with regard to the ultimate effect on labor supply, which is all that mattered for my argument.

Their third point returns to the fact that my postulated distribution-neutral financing scheme may involve raising the tax rate of a distortionary labor income tax. They agree with my point that there will be no labor supply effect. (As noted above, this suffices for all of my conclusions to be valid -- namely, that the basic cost-benefit test is applicable.) But they argue that one could substitute a uniform lump-sum tax for the augmented labor income tax rate, and this would reduce distortion. Hence, they conclude, my claim that the distribution-neutral scheme is “nondistortionary” is wrong, for whenever there exists a less distortionary finance mechanism, the one under contemplation cannot be understood as nondistortionary. This, in fact, is Browning and Liu’s major argument: because we can reduce the labor income tax rate, making up the difference with uniform lump-sum taxation, we can reduce distortion. And this, they suggest, is desirable.

This policy prescription, however, is invalid, as I explained in my original article and restated in the preceding section. The reduction in distortion is -- in a proper model with differences in income -- at the cost of redistribution. Thus, the real question is not whether we can reduce distortion (which we obviously can whenever we are using a redistributive income tax), but whether reducing the extent of redistribution is desirable. The purpose of my article was to emphasize this question, showing the inadequacy of focusing on distortion alone. That is, their central argument entails precisely the error that was the focus of my original article.

In addition, if it is desirable to reduce the extent of redistribution in the manner they suggest, one can accomplish this independently of any change in public goods provision. That is,

one can simply reduce the labor income tax rate and impose a uniform lump-sum tax (or, more practically, reduce the level of income that is exempt from the income tax and cut welfare payments to produce the equivalent result).⁷ Moreover, their implicit suggestion that it would be desirable to pursue a reform that combined this reduction in redistribution with the provision of a public good that fails the cost-benefit test is a poor one, because my main argument demonstrated that this is a Pareto dominated strategy.⁸

Before concluding, I wish to acknowledge another aspect of their semantic argument, one that occupies much of their comment. As just discussed, they object to my calling the combination of a public good and tax adjustment “nondistortionary” when the tax adjustment involves a tax that is distortionary and could be replaced by a nondistortionary lump-sum tax. Even though the choice of terminology has no relevance for the substantive analysis, I do wish to note that their semantic point has some merit. Their writing of their comment suggests that some of my passages may mislead some readers; were I to write the article anew, I would be more careful and precise in my language choice, as I was in the formal appendix.⁹

Observe, however, that Browning and Liu’s own preferred characterization of what should count as “distortion” -- whether there exists “an alternative method of finance that would leave taxpayers even better off” -- is problematic in its own way. To illustrate, under their view, even a lump-sum tax cannot be called nondistortionary in our present world. The reason is simply that we have many inefficient provisions in our income tax code and elsewhere. If we finance a public good, say, by closing an inefficient tax loophole or repealing an inefficient subsidy, it would be possible that total distortion would be even less than if we used a lump-sum tax. Yet this observation does not constitute a good reason for economists to abandon their current use of language.

CONCLUSIONS

My original article emphasized two conclusions. First, if public goods are financed in a manner that is distribution-neutral (combining the effects of the public good and of the tax adjustment), then the basic cost-benefit test is appropriate. Second, if the method of finance is not distribution-neutral, then any modification to the initial conclusion is entirely concerned with matters of redistribution. If an identified method of finance involves greater distortion (as was true in much of the prior literature), this is precisely because that method of finance involves greater redistribution. And if a method of finance involves less distortion (as in Browning and Liu's proposed substitution of uniform lump-sum taxation for labor income taxation), this is precisely because there is less redistribution. We already know that more distortion is inefficient and that we can reduce distortion if we are willing to cut back on redistribution.

Browning and Liu, despite their strong semantic disagreements, do not dispute my central analytical claim, which was in fact proved in the appendix to my original article. Nor do they dispute -- or even address -- my argument about how policy analysis should be conducted in the present context. Instead, they focus entirely upon the distortionary costs of the income tax, ignoring that the *raison d'être* of redistributive income taxation is to redistribute income.

REFERENCES

Browning, Edgar K. and Liqun Liu. “The Optimal Supply of Public Goods and the Distortionary Cost of Taxation: Comment.” *National Tax Journal* (forthcoming, March 1998).

Kaplow, Louis. “The Optimal Supply of Public Goods and the Distortionary Cost of Taxation.” *National Tax Journal* 49 (1996): 513-533.

Sandmo, Agnar. “Redistribution and the Marginal Cost of Public Funds.” Institute of Economics, Norwegian School of Economics and Business Administration, Discussion Paper 8/97 (1997).

ENDNOTES

1. Because Browning and Liu do not specifically address the portion of my article on environmental taxation, I will not consider it further here. (Note, however, that their points and mine are applicable to both contexts.)
2. At one point, Browning and Liu suggest that I do not address the question of how the marginal rate of substitution is measured. Yet, as is clear from the quote in the text, I employ an explicit measure, v_g/v_c , and it is the standard one in this context: the ratio of the marginal utility of the public good, g , to the marginal utility of the aggregate consumption good, c .
3. Note that an implication of this result (see my original appendix, n. 6) is that if one is providing more of a public good than is indicated by the basic cost-benefit test, then it is possible to reduce the level of the public good and, if the tax schedule is adjusted to rebate the savings in a manner that preserves the preexisting distribution, a Pareto improvement will result.
4. It is not the case that I am suggesting that there is something generally optimal or magical about keeping labor supply constant. Rather, it just so happens that this construction is useful in making a rigorous argument about the merits of the basic cost-benefit test when finance is distribution-neutral.
5. It is possible, of course, for there to be types of distortions that are unrelated to conventional income redistribution (e.g., tax expenditures that favor specific industries), but such sources of distortion are not the focus of the public goods provision literature, my article, or Browning and Liu's comment.
6. Their position is actually quite extreme, as they assert that a public good "*never* produces a substitution effect." Obvious examples, like constructing cheap mass transit for commuters, as well as the more subtle example in the text, show this claim to be false. I would concede, however, their semantic claim about the appropriateness of "substitution effect" language may be more apt with respect to some of my other examples or explanations.

In any event, I find their main way of putting their point about substitution effects to be confusing. They state, for example, that "The essential point is that if the same [public good] is financed by a lump sum tax instead of an increase in the income tax rate, the quantity of labor supplied would increase even more -- and taxpayers would be better off in this case." But the same could be said of providing a wage subsidy rather than a public good. By this logic, the substitution effect of a 1% wage subsidy should not (contrary to their later assertion) be thought of as offsetting the substitution effect of a 1% wage tax increase -- and, I guess, should not even be thought of as a substitution effect -- because, after all, we could always replace the wage tax increase with a lump-sum tax. Even as a purely semantic point, I do not find this formulation appealing.
7. In my original article (pp. 520-21), I comment briefly on the political feasibility of different courses of action. It is not my understanding, however, that Browning and Liu are concerned with politics in advancing their argument, so I do not pursue the point here.

8. See note 3.

9. The phrasings in my original article, which Browning and Liu criticize, reflected an attempt to track the language that was most common in the writings on and discussions of this subject. My main point about prior work was that it typically examined combinations of public goods and financing that, taken together, were *not* distribution-neutral, and, moreover, the “distortions” identified in prior work were attributable to this fact. Hence, distribution-neutral finance would indeed eliminate these “distortions.”