

Understanding Different Approaches to Benefit-Based Taxation

Robert Scherf
Matthew Weinzierl

Working Paper 19-070



Understanding Different Approaches to Benefit-Based Taxation

Robert Scherf

Harvard Business School

Matthew Weinzierl

Harvard Business School

Working Paper 19-070

Copyright © 2019 by Robert Scherf and Matthew Weinzierl

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Understanding Different Approaches to Benefit-Based Taxation

Robert Scherf

Matthew Weinzierl*

March 22, 2019

Abstract

The normative principle of benefit-based taxation has exerted substantial influence on many areas of public finance, but it has been largely set aside in the modern theoretical approach to optimal income taxation, where welfarist objectives dominate. A prerequisite for that gap to close is the clarification of what benefit-based income taxation would mean, specifically in a first-best setting. This paper seeks to provide clear, accessible descriptions and novel graphical representations of four major approaches to first-best benefit-based taxation, explain how these approaches relate to each other, and apply them within the Classical Benefit-Based framework for optimal income taxation of Smith (1776).

1 Introduction

While the normative principle of benefit-based taxation has exerted substantial influence on many areas of public finance, its influence has been more limited in the field of optimal income taxation, where welfarist objectives such as utilitarianism dominate. We believe this divide is problematic, but our goal in this paper is not to overcome readers' hesitancy to upend the prevailing normative approach to income taxation and replace it with benefit-based thinking. Instead, we try to fulfill a prerequisite for benefit-based taxation even to be considered as an alternative policy objective: namely, the clarification of what benefit-based income taxation would mean, specifically in a first-best (full-information) setting.

Therefore, this paper has simple goals: to provide clear, accessible descriptions and novel graphical representations of four major approaches to first-best benefit-based taxation (Lindahl (1919), Moulin (1987), Brennan (1976*a*), and Hines (2000)); to explain how these approaches relate to each other; and to apply them within the Classical Benefit-Based framework for optimal income taxation of Smith (1776). We use the highly stylized but justly famous two-type analytical and graphical framework introduced by Lindahl, and we stress intuition-building. We hope that, by providing greater clarity on what benefit-based thinking means, we will encourage greater appreciation and study of it among those not already convinced of its relevance as a normative principle for optimal income taxation.

*Scherf: Bonn Graduate School of Economics and Harvard Business School, rscherf@hbs.edu; Weinzierl: Harvard Business School and NBER, mweinzierl@hbs.edu.

2 The benefit-based principle and optimal income taxation

Benefit-based taxation—the idea of basing tax liabilities on how much an individual benefits from the activities of the state—has appealed to tax theorists for centuries, perhaps most prominently as a normative principle in Smith (1776) and as a positive prediction in Lindahl (1919). And recent research has shown that it is also popular among the public (see Weinzierl (2016))

It is not surprising, therefore, that benefit-based thinking has both exerted substantial influence on many areas of public policy and been fertile territory for academic research working out its implementation. In particular, the benefit-based principle is the standard lens through which local taxation is viewed, and it is more generally at the core of the theory of public finance for subnational governments, where many public goods are provided.¹ For similar reasons, the study of so-called club goods has long centered on benefit-based thinking.² A large body of work has sought ways to overcome real-world obstacles to implementing benefit-based taxation, especially that an individual’s benefit from—and thus willingness to pay for—public goods is private information.³ And, related, benefit-based thinking is central to a large and active literature on the optimal level of public goods in a setting with distortionary taxes.⁴

Despite its broad influence and appeal, benefit-based taxation has been sidelined in the modern theoretical approach to optimal income taxation. The classic public finance textbook by Atkinson and Stiglitz (1980) and the modern taxation treatise by Louis Kaplow (2008) each allocate only a few (in the latter case, skeptical) pages to benefit-based taxation; moreover, benefit-based taxes are entirely absent from the textbook on the economics of taxation by Salanié (2011) and the optimal tax survey of Boadway (2012).

In part, this omission reflects specialization of effort: benefit-based taxation has been used in areas where heterogeneous *preferences* for public goods seem most important, while the optimal income tax literature spawned by Mirrlees (1971) has emphasized heterogeneity in income-earning *abilities*. One might argue that the two fields are properly disjoint. We believe, however, that optimal income tax theory can benefit from engaging with benefit-based thinking. Why? Because it is implausible to assume that the level and mix of public goods provided by a government are irrelevant for the distribution of income-earning abilities of its taxpayers or, therefore, those taxpayers’ preferences over public goods. And as recent work by one of us (Weinzierl (2017a)) has argued, these implausible assumptions can be overcome by incorporating Smith’s view of benefit-based taxation into the modern Mirrleesian optimal income tax structure.⁵ In particular, Smith’s Classical Benefit-Based Taxation directly challenges the modern specialization of effort by making income-earning abilities—and thus preferences—depend on public goods themselves.

But, again, our point in this paper is not to convince readers that benefit-based thinking is the right approach to optimal income taxation. Instead, we focus on a simpler task: clarifying what it would mean in the context of income taxation. Confusion over the many ways to specify first-best benefit-based taxation is not surprising. After all, criticism of Lindahl’s pioneering approach has led to a variety of alternatives, and the relationships among them are difficult for all but specialists to appreciate.

¹See Fischel (2001) for an especially relevant example of the voluminous literature building on the work of Tiebout (1956).

²See Sandler and Tschirhart (1980) and the references therein.

³Most prominently, the Vickrey-Clarke-Groves mechanism developed in Vickrey (1961), Clarke (1971), and Groves (1973) provided a way to elicit true preferences. Refinements included Groves and Ledyard (1977) and Hylland and Zeckhauser (1979), and a recent contribution is Buterin, Hitzig and Weyl (2018). Our focus on the first-best, normative, meaning of different approaches to benefit-based taxation allows us to set aside the complexities of its implementation explored by these papers.

⁴See, for example, Hylland and Zeckhauser (1979), Kaplow (1996), Kreiner and Verdellin (2012), and Hendren (2019).

⁵Weinzierl (2016) provides complementary positive arguments for a renewal of interest in benefit-based taxation as a guide to income taxation. At a broader level, Weinzierl (2017b) proposes a justification for the use of non-welfarist rules in optimal tax theory. Saez and Stantcheva (2016) and Weinzierl (2014) argue that conflicts between the first-best recommendations of optimal tax models and public attitudes toward taxes merit greater attention to the theory’s normative assumptions.

3 Four approaches to benefit-based taxation

Most students of benefit-based taxation would, we venture, agree that it is the idea that taxes are to be seen—and assessed—as payment for beneficial government activities, and thus each taxpayer’s liability should correspond to that taxpayer’s benefit. Turning those words into a rigorous algorithm has proven more difficult, with a range of suggestions for how to calculate the income-equivalent value of an individual’s benefit from the activities of the state. All four of the suggested solutions we discuss below implement taxes that exactly cover the total cost of the efficient level of government spending on public goods, as defined by the Samuelson (1954) rule, but they differ on how to allocate those taxes across individuals.⁶

In this section, we first provide concise descriptions of the four leading approaches to benefit-based taxation with a limited use of formalism and a focus on intuition. We then present (what we believe are) novel complementary graphical representations of these approaches that substantially facilitate comparisons of them. In both the verbal and graphical explanations, we build on the elegant two-taxpayer analysis of Lindahl (1919).

3.1 Conceptual comparisons

To ease comparison across the four cases, we consider an economy with two individuals L, R , and a level of public goods spending G with constant marginal cost equal to 1.

3.1.1 Lindahl’s marginal benefit approach

Lindahl’s (1919) method measures benefit with an individual’s marginal willingness to pay for the activities of the state. It funds the efficient level of public goods by charging more from individuals who are willing to pay more for marginal public goods. Roughly stated, the concept behind this method is that those who would demand greater activities of the state should pay for them.

Formally, Lindahl’s model has the tax authority assign shares of the total tax burden τ_L to L and $\tau_R = (1 - \tau_L)$ to R . Individuals maximize utility, which depends on private consumption and the quantity of public goods provided, subject to the budget constraint that their incomes must pay for their private consumption and their share of the public goods expenditure. Denote with $G_L(\tau_L)$ and $G_R(\tau_R)$ the quantities of the public goods that would maximize the utilities of L and R , respectively, given their specified tax shares; that is, the quantities L and R would demand at their personalized "prices."

Lindahl’s equilibrium is defined by a level of public goods provision G^* and tax shares τ_L^*, τ_R^* such that $G_L(\tau_L^*) = G_R(\tau_R^*) = G^*$. That is, the tax shares are chosen so that L and R both demand the same level of public goods provision. Because each individual’s utility is maximized at this level of public goods and tax shares, we know that the tax shares they pay correspond to their marginal rates of substitution between public and private goods. And because these shares sum to one, this correspondence means that the sum of these individuals’ marginal rates of substitution equals the social marginal cost of public goods, so that the G^* on which they agree satisfies the Samuelson (1954) rule for the efficient level.

Lindahl’s method is strongly supported by Aaron and McGuire (1970), who argue that it constitutes "the logically correct" method of calculating the income value of public goods because of its strong resemblance to how the income values of private goods are typically calculated. That is, since we calculate the income

⁶As noted in the Introduction, we focus on first-best settings in this analysis. The translation of these approaches to second-best settings, where individuals hold private information on their benefit from public goods, is not relevant for understanding how these approaches differ from a normative perspective.

value of a private goods as the product of the individual's marginal rate of substitution (the price) and the quantity consumed, so too should we calculate the income value of public goods. As each individual pays exactly this valuation in the Lindahl equilibrium, Aaron and McGuire argue that Lindahl's solution involves "zero redistribution in income-equivalent terms." In other words, they claim that the introduction of public goods financed by Lindahl taxes matches benefit to cost for each individual, satisfying the central idea of benefit-based taxation.

3.1.2 Brennan's equal shares proposal

Brennan (1976a) offered a challenge to the Lindahl approach, arguing that the assignment of different "prices" for public goods to different persons is inconsistent with matching benefits to taxes. His key conceptual point is that only if we charge individuals the same price for a good can we be sure that the good affects all individuals' opportunities equally.

Brennan uses an analogy to private goods to make his point. Suppose only apples are available, and two persons purchase the same number of apples, but we then introduce oranges. If our goal is to have the introduction of oranges be distribution-neutral, should we charge a higher price for oranges to the person who has a relative preference for oranges? Brennan says we should not, and he extends the idea to public goods for which, as we have discussed above, Lindahl's solution would charge more from those persons who relatively prefer them. To Brennan, that solution distributes away from those with a relative taste for public goods.

Brennan infers from this analogy that "the most appropriate income-equivalent measure" of the benefit from public goods is "probably" to tax each individual an equal unit price (in our two-person example, equal to 0.5) for the provision of G^* . After all, if we had a common market price for the services provided by the public goods, all individuals would purchase them up to the point where their marginal rates of substitution were equal to that price. If, instead, we follow Lindahl and charge according to marginal rates of substitution at a fixed quantity, we will be charging more from those who have a stronger preference for public goods.

To Brennan, the Lindahl solution errs because it forces those who desire more public goods provision to pay more simply because they are paired with persons who want less public goods provision. Might we not, in fact, feel that the former are being denied benefit they would obtain were they to determine the role of the state? As we will see, Hines (2000) pursues this line of thinking further.

Aaron and McGuire (1976) respond to this critique, and Brennan responds to their response in Brennan (1976b). At the crux of their disagreement is whether, as Aaron and McGuire state, "the proper framework for considering the valuation of public goods is one in which the publicness constraint is recognized for what it is, *sui generis*." In other words, can public goods be thought of as private goods (with a different cost), as Brennan suggested they could? Or, given that all individuals must consume the same quantity of public goods, is Lindahl correct in using willingness to pay as our best guide to individuals' benefit from them?

3.1.3 Moulin's focus on inframarginal benefits

Moulin (1987) critiques Lindahl's approach from another angle, arguing that we should focus not on individuals' marginal valuations of public goods but on their *inframarginal* valuations. Moulin's key conceptual innovation is to have individuals pay taxes tied to the benefit they've obtained from all the activities of the state, not just the benefit they would obtain (and thus be willing to pay for) from the "marginal" activity. Moulin thus breaks Lindahl's link between tax shares and marginal rates of substitution, but in exchange he can claim to have matched taxes to a more informative measure of benefits.

Moulin's method uncovers what might be called the "excess benefit puzzle." The total value (across all individuals) of the inframarginal benefit Moulin measures typically will be greater than the total cost of the public goods. How should this excess be allocated under a benefit-based tax system? (See Kaplow (2008) for a related discussion) Note that Lindahl avoids this problem by ignoring inframarginal valuations: after all, the Samuelson rule that determines the efficient quantity of public goods and the Lindahl equilibrium are both based on marginal valuations.

To implement his modification of Lindahl and solve this excess benefit puzzle, Moulin's "egalitarian equivalent solution" relies on the identification of a crucial quantity of public goods: g^* . Moulin measures an individual's benefit from the activities of the state as the total (inframarginal) benefit from the last $(G^* - g^*)$ units of public goods, where g^* is chosen so that the sum of this benefit across all individuals exactly equals the cost of G^* . In other words, he allocates the excess benefit from the activities of the state by giving everyone their benefit from g^* units of public goods "tax free." By taxing away the income-equivalent value of each individual's inframarginal benefits from the last $(G^* - g^*)$ units of public goods consumption, Moulin's "egalitarian" solution thereby gives each individual the same level of utility as if they had freely received the hypothetical level g^* of public goods.

Another way to summarize Moulin's approach is that it measures the benefit from public goods as the money amount needed to compensate an individual for receiving only g^* instead of G^* for free, where g^* is chosen so the sum of individuals' benefits equals the total cost of the public goods. Thinking in terms of compensation will be useful in our discussion of Hines (2000).

Moulin's modification of Lindahl can have substantial effects on the assignment of taxes if individuals' benefits from the activities of the state decline at different rates, especially around G^* . In particular, Moulin motivates his modification of Lindahl by noting that a decrease in the marginal cost of public goods provision could theoretically leave one individual worse off under Lindahl taxation. Such a decrease in cost would increase the equilibrium level G^* and could thus increase one individual's tax share if the other person's marginal benefit drops dramatically after G^* . Moulin's solution avoids this result, as lower marginal costs result in a higher g^* and thus higher utility for both individuals.

3.1.4 Hines's inclusion of extramarginal benefits

Hines (2000) offers an alternative solution that, like Brennan's, gives priority to charging all individuals a common price for public goods. As markets charge a uniform price across individuals, Hines seeks a way to do the same with public goods. His key conceptual innovation accomplishes this goal by funding public goods in a way that leaves all individuals exactly as well off as they would be if they could purchase public goods in a private market at a common price. That is, even though the services of the public goods are provided publicly, there is no redistribution (in terms of well-being) from the hypothetical state in which the same services were provided privately.

More formally, analogous to Moulin's key quantity g^* , Hines's method relies on the identification of a key price ρ at which individuals could (hypothetically) purchase public goods. Hines measures the benefit of public goods provision as the increase between each individual's well-being if public goods could be purchased at price ρ and that individual's well-being if the level of public goods G^* were produced tax-free (to the individual). Hines then taxes each individual an amount equal to that benefit, choosing ρ so that the tax revenue funds the total cost of G^* .

One way to understand how Hines's approach differs from Moulin's, with which it shares important

features⁷, is how they answer a central question of benefit-based taxation: Benefit relative to what? Moulin's approach measures benefit relative to a setting in which the public goods in question are not provided; Hines's approach measures benefit relative to a setting in which the individual's preferred level of public goods are provided. Moulin is led to this approach because the central problem he wants to solve is "selecting an equitable production level and cost sharing of a single public good," so it is natural for him to focus on "an agent's global increment of utility between the optimal level and the (fictitious) egalitarian-equivalent level" as a guide to that individual's share of the costs. Hines, in contrast, wants to "allocate taxes in a manner akin to market prices for public services," so it is natural for him to focus on achieving the same distribution of utilities as would result in the hypothetical private market for the activities of the state.

Despite its achievements, Hines's method produces a result that even he found potentially troubling: tax shares may be *lower* for those who more prefer public goods. But this result has a straightforward intuition behind it and can be seen as an extension of Brennan's thinking. In a private market, some individuals might have demanded, at price ρ , more public goods than G^* . When public goods provision is limited to G^* , these individuals lose out on benefit they would have obtained in the private market, and Hines's approach thus includes compensation for these foregone "extra-marginal" benefits. In other words, Hines may tax those who most value public goods *less* than he taxes others because the former are harmed by the low level of public goods agreed to in equilibrium.

3.2 Graphical comparison

Graphical representations of these approaches to benefit-based taxation can reveal, with more simplicity than can words or equations, how they relate to each other. In this section, we extend Lindahl's graphical approach to the four approaches described above, uncovering what we believe are novel visual connections (and distinctions) between them.

As in the previous section, we assume an economy with two individuals L, R , and a level of public goods provision G with constant marginal cost equal to 1. We plot the level of public goods provision on the vertical axis. On the horizontal axis, we plot the share of the costs of public goods paid for by individual L , or equivalently the share not paid for by individual R (whose share can be read starting from the right end of the horizontal axis).

The central feature of these figures is a pair of demand curves, one for each individual, that trace out the individuals' preferred levels of public goods provision across the range of cost shares. We will focus on the intersection of these curves as well as the areas of various shapes created by them to help illustrate the four approaches to benefit-based taxation.

The figures are substantially simpler to interpret if we assume that individual demand for public goods provision is linear (and decreasing) in the individual's cost share. We use the following demand functions:

$$\begin{aligned} G_L &= -5\tau_L + 5 \\ G_R &= -6\tau_R + 4 = 6\tau_L - 2 \end{aligned}$$

Linearity of demand is restrictive, and we will relax it in the next section where we provide some rough quantitative illustrations, but it is has substantial pedagogical value.

⁷Both, for example, are in the core and satisfy the appealing condition that reductions in the cost of public goods provision cannot make anyone worse off. See Moulin (1987) and Hines (2000) for details.

3.2.1 Lindahl's marginal benefit approach

Figure 1 shows Lindahl's (1919) original illustration of how to distribute the cost of public goods. The intersection of the two curves is the Lindahl equilibrium: at the tax share $\tau_L^{\text{Lindahl}} = 0.64$, both individuals demand the same level of public goods provision $G^* = 1.82$.

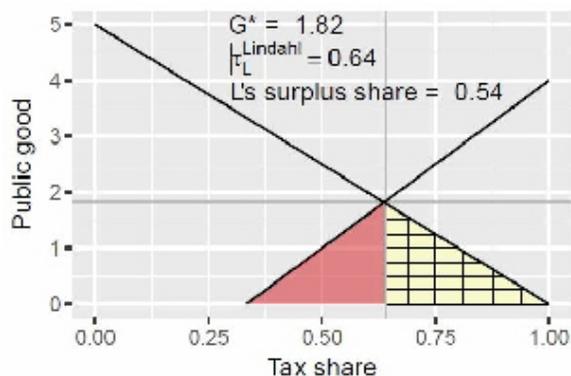


Figure 1. Lindahl's equilibrium

To compare the four approaches, we will examine not only the equilibrium tax shares τ_L and τ_R but also what we will call the "surplus" for each individual. Conceptually, an individual's surplus is simply her benefit from public goods less her share of the total cost. Graphically, in Figure 1, we measure surplus with the area "below" an individual's demand curve—which traces the marginal benefit the individual receives from each increment of public goods—and "above" her equilibrium tax share, where below and above refer to the standard directions in the analysis of consumer surplus (for L , the triangle with cross-hatching in Figure 1; for R , the shaded triangle). We then calculate the ratio of these shapes' areas to their sum, yielding the result that L obtains 54% of the total available surplus in this scenario's Lindahl equilibrium.

3.2.2 Brennan's equal shares proposal

Figure 2 depicts Brennan's proposal, where all individuals equally share in the cost of the efficient level of public goods, so $\tau_L = \tau_R = 0.5$. As a result, the individual with a stronger taste for the public goods (L) pays a smaller tax share (of the same $G^* = 1.82$) than in the Lindahl equilibrium.

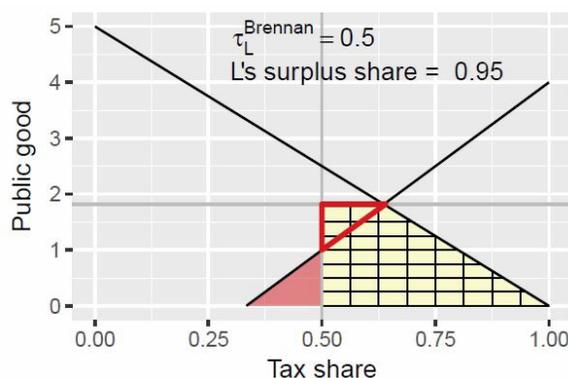


Figure 2. Brennan's equilibrium

We indicate individual surplus largely as in the previous figure, but Brennan's equilibrium concept generates one important difference. Individual R incurs *negative* surplus—its marginal benefit falls below its cost share of 0.5—once the level of public goods provision rises above 1.0. We depict this negative surplus with the bold-outlined small triangle, and R 's surplus is the shaded triangle less this bold-outlined triangle. Note that this bold-outlined triangle is also cross-hatched because it is part of the surplus obtained by L . In other words, R is subsidizing L in Brennan's solution, such that R receives only 5% of the total surplus.

3.2.3 Moulin's focus on inframarginal benefits

Figures 3a and 3b depict Moulin's egalitarian equivalent solution, where taxes are set equal to each individual's inframarginal benefits from the last $(G^* - g^*)$ units of public goods and g^* is pinned down by the requirement that these taxes exactly pay for G^* .

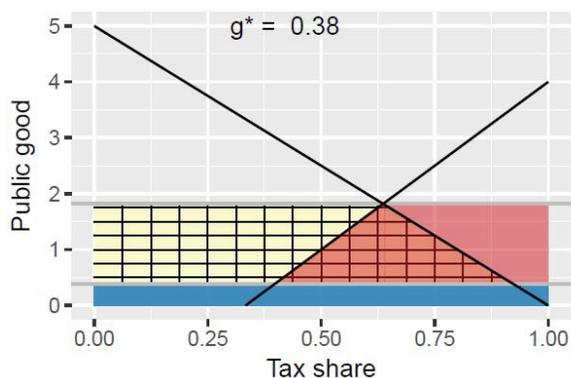


Figure 3a. Moulin's equilibrium, measuring benefit

In Figure 3a, we cross-hatch L 's and shade R 's inframarginal benefit from the last $(G^* - g^*)$ units of the public good as we shaded net surplus of benefit over cost in previous figures (note that these trapezoids do not indicate surplus of benefit over cost, as we have yet to determine cost shares, so they are just the areas below the individuals' marginal benefit curves). These trapezoids overlap in the shaded and cross-hatched triangle, an area that is received as tax revenue *twice*, once from L and once from R . This triangle is the "excess benefit" identified in the previous section as a puzzle for Moulin's approach, and as discussed above Moulin's solution was to distribute this excess benefit equally to all through a tax-free grant of g^* public goods. We indicate this grant with the dark-shaded rectangle in Figure 3a, inspection of which reveals that it has the same area as the triangle of benefit that both individuals obtain. Equating these two areas is what pins down, in this example, the level $g^* = 0.38$.

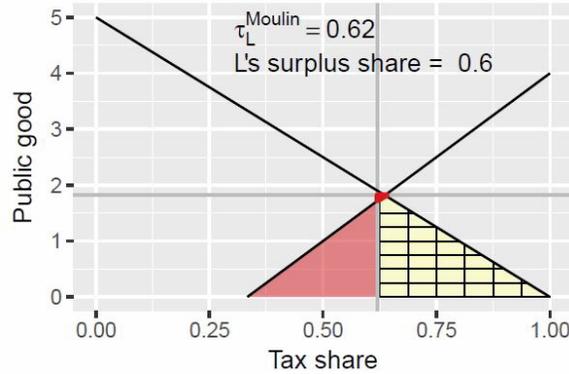


Figure 3b. Moulin’s equilibrium, tax and surplus shares

In Figure 3b we calculate the tax and surplus shares that correspond to Moulin’s approach in this example. Moulin has taxes correspond to inframarginal benefits from the last $(G^* - g^*)$ units of public goods, so tax shares equal the shares of total inframarginal benefits. In graphical terms, τ_L equals the ratio of the cross-hatched trapezoid in Figure 3a to the sum of the cross-hatched and shaded trapezoids. In our numerical example, $\tau_L = 0.62$. With this tax share, we can calculate and shade the areas of surplus in Figure 3b as we did in Figures 1 and 2. As in Figure 2, because $\tau_L < \tau_L^{Lindahl}$, R incurs a (very) small triangle of negative surplus (outlined in bold). In the end, L claims 60% of the total surplus.

Readers may note that Lindahl and Moulin’s approaches yield quite similar results in this case. This is generally the case when the individuals’ relative marginal benefits at the efficient level of public goods G^* are similar to their relative marginal benefits at lower levels of public goods provision, as is the case with these linear demand curves.

3.2.4 Hines’s inclusion of extramarginal benefits

Figures 4a and 4b depict Hines’s solution, where benefit is measured by the difference between an individual’s surplus from receiving G^* for free and from being able to purchase the public goods at a common price ρ in a private market. As with Moulin’s g^* , ρ is chosen such that the sum of these individual benefits equals the total cost of G^* .

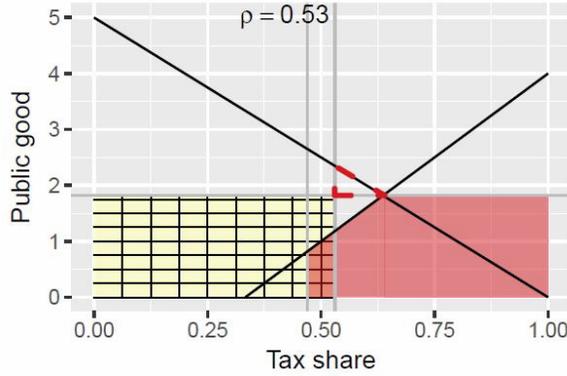


Figure 4a. Hines's equilibrium, measuring benefit

In Figure 4a, we want to visually identify inframarginal benefits as before, but the appropriate areas differ substantially from those in the other figures because Hines's approach differs so substantially from the others. Consider L 's benefit. In our example, $\rho = 0.53$. L would purchase 2.35 units of G at this price, as shown by the intersection of L 's demand curve and the vertical line at $\tau_L = 0.53 = \rho$. Thus, the amount that L would consume in this private market is *greater* than $G^* = 1.82$. Receiving G^* for free rather than at price ρ results in a benefit of $\rho G^* = 0.97$, shown as the cross-hatched rectangle. Hines' solution, however, also *compensates* L for the "extra-marginal" benefit from those 0.53 ($= 2.35 - 1.82$) units of public goods that L would have purchased in a private market *in excess of* the equilibrium quantity G^* . This compensation is shown as the triangle outlined in a dashed bold line. Thus, L 's total benefit under Hines's definition equals the cross-hatched rectangle less the triangle outlined in a dashed bold line. R 's case is simpler, as R would consume less than G^* in a private market with unit price $\rho = 0.53$ (as shown by the intersection of R 's demand curve and the vertical line at $\tau_L = 0.47$; $\tau_R = \rho = 0.53$). Hence, R requires no compensation for extramarginal units of the public goods, and R 's total benefit of receiving G^* for free rather than at price ρ is depicted by the shaded pentagon.

The most novel piece of Hines's solution is ρ . Its value is set so that the sum of the individuals' benefits (that is, the sum of the cross-hatched rectangle and shaded pentagon minus the triangle outlined in a dashed bold line) equals the total cost of the public goods G^* . Another way to see this in Figure 4a is that the small trapezoid in which the cross-hatched rectangle and the shaded pentagon overlap have the same area as that of the triangle outlined in a dashed bold line *plus* its twin, adjacent empty triangle. These two triangles are, respectively, counted against and not counted as part of L 's benefit, despite being under L 's demand curve, because L would have liked to purchase a greater level of public goods than is being provided in equilibrium. If we were to increase ρ , these triangles would shrink (because L would purchase less at a higher price), and the area of overlap would expand. When the level of ρ is found at which these areas cancel out, the benefit-based tax precisely funds G^* .

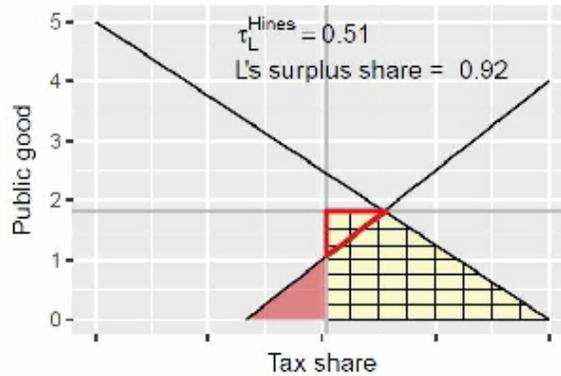


Figure 4b. Hines's equilibrium, tax and surplus shares

In Figure 4b, we show how this solution translates into tax shares and, thus, the distribution of surplus. As with Moulin, tax shares under Hines equal the shares of total benefits. In graphical terms, τ_L equals the ratio of the area of the cross-hatched rectangle less the triangle outlined in a dashed bold line in Figure 4a to the sum of that area and the shaded pentagon's area. In our numerical example, $\tau_L = 0.51$, as shown in Figure 4b. With this tax share, we can calculate and shade the areas of surplus, concluding that L claims a remarkable 92% of the total surplus even though L has a stronger taste for public goods. This result is due to the compensation Hines's approach provides L for the "extra-marginal" benefit from the units beyond G^* that L would have consumed in a private market.

In the online appendix to this paper, the reader can find code (written in R) that will re-create Figures 1 through 4 for different numerical examples.

4 Application to Smith's "Classical" Benefit-Based Taxation

Benefit-based taxation is a natural approach to the funding of marginal public goods, such as a single infrastructure project. But a persistent concern has been how to apply it beyond the case of incremental public goods to the full suite of the state's activities.

Smith (1776) provided an early solution to this problem by arguing that an individual's benefit from the activities of the state is given by the state's impact on her ability to earn income. Smith's first maxim of income taxation states that individuals should be taxed "as near as possible, in proportion to their respective abilities; that is in proportion to the revenue which they respectively enjoy under the protection of the state." Musgrave (1959) called this Smith's (1776) "classical" view of benefit-based taxation, and it provides a natural logic for the application of benefit-based thinking to the optimal income tax problem.⁸

Moreover, Smith's approach can be seen as a way to address an implausible assumption of modern optimal income tax theory: namely, that the level and mix of public goods provision is irrelevant for the distribution of income-earning abilities that is at the center of Mirrleesian analysis. If we instead adopt the view that the activities of the state matter to individuals primarily because of their effects on the individuals' productive

⁸The benefit-based principle may be one of a set of normative criteria used to determine optimal taxation. Other criteria, such as the utilitarianism conventionally used in Mirrleesian optimal income taxation, would typically recommend more net redistribution (that is, where some taxpayers receive more in transfers than they pay in taxes) than would benefit-based taxation. This idea of "normative diversity" is discussed in Weinzierl (2014) and Saez and Stantcheva (2016).

ability, rather than because of individuals' heterogeneous "tastes" for public goods (conditional on income), the application of benefit-based thinking to income taxation becomes natural.

Smith did not formalize his maxim, leaving room for multiple approaches to how classical benefit-based taxation could be operationalized, so in this section we will apply the different benefit-based solutions described above to Smith's classical benefit-based taxation. We follow Weinzierl (2017a) and model individual income-earning ability as a function of both innate talent and spending on public goods. The benefit that an individual derives from the activities of the state is then given by the extent to which expenditure on public goods increases her ability to earn income.

Specifically, consider an economy of two individuals $i = L, R$ who differ only in their innate talents $a_i \in \{a_L, a_R\}$ with $a_L > a_R$. An individual of type a_i has preferences

$$u_i = c_i - v\left(\frac{y_i}{f(a_i, G)}\right),$$

where y_i and c_i denote i 's pre-tax and after-tax income, i 's ability to earn income $f(a_i, G)$ is a function of innate ability a_i and spending on public goods G , labor effort is therefore $\frac{y_i}{f(a_i, G)}$, and the disutility of labor function $v(\cdot)$ is such that $v'(\cdot) > 0$ and $v''(\cdot) > 0$. We also consider a special case of these preferences:

$$u_i = c_i - \frac{1}{\sigma} \left(\frac{y_i}{a_i G^\gamma}\right)^\sigma, \quad (1)$$

where the disutility of labor function is iso-elastic and the ability production function takes the Cobb-Douglas form. This will allow us to calculate a numerical comparison of the different benefit-based solutions. G is produced at unit cost, and the efficient level of government spending, denoted G^* , is defined by the Samuelson rule that the sum of marginal benefits from the public goods should equal the marginal cost of the public goods:

$$\sum_{i=L,R} \frac{y_i f_G(a_i, G^*)}{f(a_i, G^*)^2} v'\left(\frac{y_i}{f(a_i, G^*)}\right) = 1.$$

If income-earning ability takes the Cobb-Douglas form in (1), the Samuelson rule implies that total spending on G should equal

$$G^* = \gamma \sum_{i \in \{L, R\}} y_i^*,$$

that is, a share γ of total income. Given G^* , individuals will choose to earn income y_i^* satisfying

$$\frac{1}{f(a_i, G^*)} v'\left(\frac{y_i^*}{f(a_i, G^*)}\right) = 1,$$

or, given our more specific functional form in (1):

$$y_i^* = (a_i G^{*\gamma})^{\frac{\sigma}{\sigma-1}}.$$

We can now apply to this economy all four of the approaches to benefit-based taxation—that is, to determining how G^* will be funded through taxes on individuals $i = L, R$ —that we discussed in the previous sections.

4.1 Analytical results

First, we derive expressions for the tax shares under each approach.

4.1.1 Lindahl's marginal benefit approach

Lindahl's solution would assign to each individual a share of the cost of G^* equal to that individual's marginal benefit from the last dollar of public spending. The resulting tax shares in this setting are given by

$$\tau_i^{Lindahl} = \frac{f_G(a_i, G^*)}{f(a_i, G^*)} y_i^*.$$

As shown in Weinzierl (2017a), whether Lindahl's taxes are progressive in this setting depends on the extent to which innate talent and government spending are complements in determining individual income-earning ability. If the ability production function is Cobb-Douglas as in (1), then Lindahl's solution assigns to each individual a flat income tax rate equal to γ :

$$\tau_i^{Lindahl} G^* = \gamma y_i^*. \quad (2)$$

4.1.2 Brennan's equal shares proposal

Brennan's proposal would split the total tax burden equally between all individuals:

$$\tau_L^{Brennan} = 0.5, \quad i = L, R. \quad (3)$$

Equal tax shares yield a regressive income tax schedule.

4.1.3 Moulin's focus on inframarginal benefits

Moulin's solution sets each individual's tax burden equal to that individual's inframarginal benefit from the last $G^* - g^*$ units of public goods spending, leaving each individual as well off as if g^* were provided tax-free. Stating this approach formally, individual tax shares in this setting are given by:

$$\tau_i^{Moulin} = \frac{1}{G^*} \int_{g^*}^{G^*} y_i^* \frac{f_G(a_i, G)}{f(a_i, G)^2} v' \left(\frac{y_i^*}{f(a_i, G)} \right) dG, \quad (4)$$

where g^* is chosen such that the sum of these benefits equals the cost of the public goods: $\tau_L^{Moulin} + \tau_R^{Moulin} = 1$. To set the stage for the upcoming numerical calculations, the expression for i 's tax share in (4) can be rewritten as follows using our more specific functional form in (1):

$$\begin{aligned} \tau_i^{Moulin} &= \frac{1}{G^*} \left[\int_{g^*}^{G^*} \gamma \left(\frac{y_i^*}{a_i} \right)^\sigma \frac{1}{G^{1+\gamma\sigma}} dG \right] \\ &= \frac{1}{G^*} \left[\gamma \left(\frac{y_i^*}{a_i} \right)^\sigma \left[-\frac{1}{\gamma\sigma G^{\gamma\sigma}} \right]_{g^*}^{G^*} \right] \\ &= \frac{1}{G^*} \left[\frac{1}{\sigma} \left(\frac{y_i^*}{a_i} \right)^\sigma \left(\frac{1}{g^{*\gamma\sigma}} - \frac{1}{G^{*\gamma\sigma}} \right) \right] \\ &= \frac{1}{G^*} \left[\left(y_i^* - \frac{1}{\sigma} \left(\frac{y_i^*}{a_i G^{*\gamma}} \right)^\sigma \right) - \left(y_i^* - \frac{1}{\sigma} \left(\frac{y_i^*}{a_i g^{*\gamma}} \right)^\sigma \right) \right]. \quad (5) \end{aligned}$$

In other words, (5) shows that the individual tax payments under Moulin's solution are equal to the difference between the individual's utilities from receiving G^* and g^* for free. The first term inside square brackets is the utility for i when G^* is provided tax-free; the second term is the utility when g^* is provided tax-free. By taxing i an amount $(\tau_i G^*)$ equal to the difference between these terms, Moulin's solution leaves each individual exactly as well off as if g^* were provided tax-free.

4.1.4 Hines's inclusion of extramarginal benefits

Hines suggests taxing individuals on the difference between the benefit they would obtain from the tax-free provision of G^* and the benefit they would obtain from purchasing public goods at the private-market price ρ . Formally, tax shares in this setting would be given by:

$$\tau_i^{Hines} = \frac{1}{G^*} \left(\rho G_i^\rho + \int_{G_i^\rho}^{G^*} y_i^* \frac{f_G(a_i, G)}{f(a_i, G)^2} v' \left(\frac{y_i^*}{f(a_i, G)} \right) dG \right), \quad (6)$$

where G_i^ρ is what individual i would demand if public goods could be purchased privately at a unit price ρ :

$$G_i^\rho = \arg \max_G y_i^* - \rho G - v \left(\frac{y_i^*}{f(a_i, G)} \right),$$

and ρ is chosen such that the sum of the benefits equals the cost: $\tau_L^{Hines} + \tau_R^{Hines} = 1$.

As with Moulin's solution, the individual benefit according to Hines's solution can be broken into two parts. First, each individual gains an amount ρG_i^ρ that they would have paid to purchase public goods in the private market. Second, while R receives inframarginal benefit from the additional $G^* > G_R^\rho$ units that the government provides, L loses the "extra-marginal" benefit from the $G_L^\rho - G^*$ units that L would have demanded in the private market. Using once again our specific functional forms in (1), we can transform expression (6) into a form that makes these two parts clear:

$$\begin{aligned} \tau_i^{Hines} &= \frac{1}{G^*} \left[\rho G_i^\rho + \int_{G_i^\rho}^{G^*} \gamma \left(\frac{y_i^*}{a_i} \right)^\sigma \frac{1}{G^{1+\gamma\sigma}} dG \right] \\ &= \frac{1}{G^*} \left[\rho G_i^\rho + \gamma \left(\frac{y_i^*}{a_i} \right)^\sigma \left[-\frac{1}{\gamma\sigma G^{\gamma\sigma}} \right]_{G_i^\rho}^{G^*} \right] \\ &= \frac{1}{G^*} \left[\rho G_i^\rho + \frac{1}{\sigma} \left(\frac{y_i^*}{a_i} \right)^\sigma \left(\frac{1}{(G_i^\rho)^{\gamma\sigma}} - \frac{1}{(G^*)^{\gamma\sigma}} \right) \right] \\ &= \frac{1}{G^*} \left[\left(y_i^* - \frac{1}{\sigma} \left(\frac{y_i^*}{a_i (G^*)^\gamma} \right)^\sigma \right) - \left(y_i^* - \rho(G_i^\rho) - \frac{1}{\sigma} \left(\frac{y_i^*}{a_i (G_i^\rho)^\gamma} \right)^\sigma \right) \right]. \quad (7) \end{aligned}$$

The first term inside square brackets in expression (7) is the utility for i when G^* is provided tax-free; the second term is the utility when i can purchase public goods at the hypothetical price ρ . By taxing i an amount $(\tau_i G^*)$ equal to the difference between these terms, Hines's solution leaves each individual exactly as well off as if they could obtain public goods at price ρ in the private market.

4.2 Numerical comparison of results

To make these applications more tangible, we simulate these results under specific parameter values and the specification of preferences in (1). We set $\sigma = 5$ to give a labor supply elasticity of 0.25. Recall that with

this preference specification the efficient size of government according to the Samuelson rule is given by a constant share γ of the total output. We set $\gamma = 0.2$, as this roughly corresponds to the ratio of government expenditures to GDP in the United States. We further choose a_L and a_R so that the ratio of L 's income to R 's equals four, approximately the ratio of average incomes of the top versus bottom half of the U.S. income distribution according to data from the Congressional Budget Office. Given these preferences, we can derive the individual demand curves for government spending $G_i = \left[\left(\frac{y_i^*}{a_i} \right)^\sigma \frac{\gamma}{\tau_i} \right]^{\frac{1}{1+\gamma\sigma}}$, which are log-linear in τ_i and which exhibit constant elasticity of demand $\frac{1}{1+\gamma\sigma}$.

In Table 1, we present the four different approaches' solutions—for our functional form and parameter specifications—as stated in the numbered equations above.

Table 1: Numerical results for classical benefit-based taxation

Lindahl's solution		
G^*	10.00	Calculated with Samuelson rule
y_L^*	40.00	Equilibrium income for L
y_R^*	10.00	Equilibrium income for R
G^*/Y	0.20	G/Y ratio, should equal γ
$\tau_L^{Lindahl}$	0.80	L's tax share
$\tau_R^{Lindahl}$	0.20	R's tax share
$U_L^{Lindahl}$	24.00	L's utility
$U_R^{Lindahl}$	6.00	R's utility
Moulin's solution		
g^*	5.00	Hypothetical "egalitarian" level of G
τ_L^{Moulin}	0.80	L's tax share
τ_R^{Moulin}	0.20	R's tax share
U_L^{Moulin}	24.00	L's utility
U_R^{Moulin}	6.00	R's utility
Brennan's solution		
$\tau_L^{Brennan}$	0.50	L's tax share
$\tau_R^{Brennan}$	0.50	R's tax share
$U_L^{Brennan}$	27.00	L's utility
$U_R^{Brennan}$	3.00	R's utility
Hines's solution		
ρ	0.56	Hypothetical market price for G
G_L^ρ	11.95	G demanded by L at price ρ
G_R^ρ	5.98	G demanded by R at price ρ
τ_L^{Hines}	0.54	L's tax share
τ_R^{Hines}	0.47	R's tax share
U_L^{Hines}	26.60	L's utility
U_R^{Hines}	3.30	R's utility

We learn several lessons from Table 1. First, the Lindahl and Moulin solutions coincide under this specification; as discussed earlier, these approaches will often yield similar results when relative demand for public goods is similar at different levels of G , as is the case when we assume the Cobb-Douglas form in (1). Second, using Lindahl or Moulin, Smith's classical benefit-based taxation implies a proportional tax in this case: that is, each individual pays the same constant marginal (and thus average) tax rate, equal to $\gamma = 20\%$. In contrast, both Brennan's and Hines's solutions result in regressive taxes. For Brennan, the reason is simply that his approach levies uniform taxes in levels. As noted earlier, for Hines the reason is that his approach compensates L for the "extra-marginal" benefit from the additional public goods L would like society to provide but must forego (note that L would have bought nearly twenty percent more G than G^* at price ρ). We can see these different perspectives reflected in the tax shares under each approach, which

are equal to the shares of the total benefit for each individual. While Lindahl and Moulin assign benefits proportional to income, Brennan and Hines assign (precisely and roughly, respectively) equal benefits to L and R despite L having four times the income of R . The utility levels under each approach capture succinctly the distributional implications of these different perspectives.

Smith's writings suggest that his intuition was closest, conceptually and quantitatively, to Moulin's. He emphasized the state as a protector of property rights and assets, and he endorsed proportional taxation above a minimum. See Weinzierl (2017a) for more on Smith's maxim and its interpretation.

5 Conclusion

How we might choose to apply benefit-based thinking to the theory of optimal income taxation reflects what we see as the heart of its appeal. If benefit-based taxation is appealing because it charges people for what the state has made possible for them, we are likely to find Moulin's refinement of Lindahl most appropriate. If benefit-based taxation is appealing because it gets us as close as possible to funding public goods the way private goods are purchased, we are likely to side with Hines's development of Brennan's approach. Our simple hope with this paper is to clarify these different views and, thereby, encourage a renewal of interest in and the study of benefit-based taxation in the context of optimal income tax theory, where we believe its application is natural. Many questions remain to be studied and answered, but given benefit-based taxation's substantial conceptual, historical, and empirical appeal, we believe the effort is well justified by the potential benefit.

Disclosure statement

The authors have no financial arrangements that might give rise to conflicts of interest with respect to the research reported in this paper.

References

- Aaron, Henry, and Martin McGuire.** 1970. "Public goods and income distribution." *Econometrica*.
- Aaron, Henry, and Martin McGuire.** 1976. "Reply to Geoffrey Brennan, "The Distributional Implications of Public Goods"." *Econometrica*.
- Brennan, Geoffrey.** 1976a. "The Distributional Implications of Public Goods." *Econometrica*.
- Brennan, Geoffrey.** 1976b. "Public Goods and Income Distribution: A Rejoinder to the Aaron-McGuire Reply." *Econometrica*.
- Buterin, Vitalik, Zoe Hitzig, and E. Glen Weyl.** 2018. "Liberal Radicalism: A Flexible Design for Philanthropic Matching Funds." *working paper*.
- Clarke, E.H.** 1971. "Multipart pricing of public goods." *Public Choice*.
- Fischel, William A.** 2001. "Homevoters, municipal corporate governance, and the benefit view of the property tax." *National Tax Journal*.
- Groves, Theodore.** 1973. "Incentives in Teams." *Econometrica*.
- Groves, Theodore, and John Ledyard.** 1977. "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem." *Econometrica*.

- Hendren, Nathaniel.** 2019. “Efficient Welfare Weights.” *NBER Working Paper 20351*.
- Hines, James R.** 2000. “What is Benefit Taxation?” *Journal of Public Economics*.
- Hylland, Aanund, and Richard Zeckhauser.** 1979. “The Efficient Allocation of Individuals to Positions.” *Journal of Political Economy*.
- Kaplow, Louis.** 1996. “The optimal supply of public goods and the distortionary cost of taxation.” *National Tax Journal*.
- Kaplow, Louis.** 2008. *The Theory of Taxation and Public Economics*. Princeton.
- Kreiner, Claus Thustrup, and Nicolaj Verdelin.** 2012. “Optimal provision of public goods: A synthesis.” *Scandinavian Journal of Economics*.
- Lindahl, Erik.** 1919. “Just taxation—a positive solution, in Classics in the theory of public finance.” 168–176.
- Mirrlees, James.** 1971. “An Exploration in the Theory of Optimal Income Taxation.” *Review of Economic Studies*.
- Moulin, Herve.** 1987. “Egalitarian-Equivalent Cost Sharing of a Public Good.” *Econometrica*.
- Saez, Emmanuel, and Stefanie Stantcheva.** 2016. “Generalized Social Marginal Welfare Weights for Optimal Tax Theory.” *American Economic Review*.
- Samuelson, Paul.** 1954. “The Pure Theory of Public Expenditure.” *Review of Economics and Statistics*.
- Sandler, Todd, and John T. Tschirhart.** 1980. “The Economic Theory of Clubs: An Evaluative Survey.” *Journal of Economic Literature*.
- Smith, Adam.** 1776. *Wealth of Nations*. Prometheus Books.
- Tiebout, Charles M.** 1956. “A pure theory of local expenditure.” *Journal of political economy*.
- Vickrey, William.** 1961. “Counterspeculation, Auctions, and Competitive Sealed Tenders.” *Journal of Finance*.
- Weinzierl, Matthew.** 2014. “The Promise of Positive Optimal Taxation: Normative Diversity and a role for Equal Sacrifice.” *Journal of Public Economics*.
- Weinzierl, Matthew.** 2016. “Popular Acceptance of Inequality due to Brute Luck and Support for Classical Benefit-Based Taxation.” *Journal of Public Economics*.
- Weinzierl, Matthew.** 2017a. “Revisiting the Classical View of Benefit Based Taxation.” *Economic Journal*.
- Weinzierl, Matthew.** 2017b. “A Welfarist Role for Nonwelfarist Rules.” *NBER WP 23587*.