NBER WORKING PAPER SERIES

EFFICIENT WELFARE WEIGHTS

Nathaniel Hendren

Working Paper 20351 http://www.nber.org/papers/w20351

NATIONAL BUREAU OF ECONOMIC RESEARCH 1050 Massachusetts Avenue Cambridge, MA 02138 July 2014, Revised October 2017

This paper is a heavily revised version of a paper that previously circulated under the title, "The Inequality Deflator: Interpersonal Comparisons without a Social Welfare Function", NBER WP No. 20351. I am deeply indebted to conversations with Louis Kaplow for the inspiration behind this paper, and to Sarah Abraham, Alex Bell, Alex Olssen, and Evan Storms for excellent research assistance. I also thank Daron Acemoglu, Raj Chetty, Amy Finkelstein, Ben Lockwood, Henrik Kleven, Patrick Kline, Kory Kroft, Matthew Notowodigdo, Jim Poterba, Emmanuel Saez, Matthew Weinzierl, Glen Weyl, Ivan Werning, and Floris Zoutman, along with seminar participants at Berkeley, Harvard, MIT, Michigan, and Stanford for very helpful comments. The opinions expressed in this paper are those of the author alone and do not necessarily reflect the views of the Internal Revenue Service or the U.S. Treasury Department. This work is a component of a larger project examining the effects of tax expenditures on the budget deficit and economic activity, and this paper in particular provides a general characterization of the welfare impact of changes in tax expenditures relative to changes in tax rates (illustrated in Section 6). The empirical results derived from tax data that are reported in this paper are drawn from the SOI Working Paper "The Economic Impacts of Tax Expenditures: Evidence from Spatial Variation across the U.S.", approved under IRS contract TIRNO-12-P-00374. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by Nathaniel Hendren. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Efficient Welfare Weights Nathaniel Hendren NBER Working Paper No. 20351 July 2014, Revised October 2017 JEL No. D6,E01,H0

ABSTRACT

How should we measure economic efficiency? The canonical measure is an unweighted sum of willingnesses to pay. In contrast, this paper provides efficient welfare weights that implement the Kaldor-Hicks tests for efficiency but account for the distortionary cost of taxation. The shape of the income distribution yields bounds on these weights that suggest it is efficient to weight surplus to the poor more than to the rich. Point estimates suggest surplus to the poor should be weighted 1.5-2x more than surplus to the rich. I illustrate how to use these weights to evaluate the efficiency of government policy changes.

Nathaniel Hendren Harvard University Department of Economics Littauer Center Room 235 Cambridge, MA 02138 and NBER nhendren@gmail.com

Efficient Welfare Weights

Nathaniel Hendren*

October, 2017

Abstract

How should we measure economic efficiency? The canonical measure is an unweighted sum of willingnesses to pay. In contrast, this paper provides efficient welfare weights that implement the Kaldor-Hicks tests for efficiency but account for the distortionary cost of taxation. The shape of the income distribution yields bounds on these weights that suggest it is efficient to weight surplus to the poor more than to the rich. Point estimates suggest surplus to the poor should be weighted 1.5-2x more than surplus to the rich. I illustrate how to use these weights to evaluate the efficiency of government policy changes.

1 Introduction

Changes to economic policies usually create winners and losers. This makes it difficult to determine whether a policy change is desirable without taking a stance on how to weigh the gains to the winners against the losses to the losers.

In their classic work, Kaldor and Hicks proposed a notion of "economic efficiency" to deal with this problem: construct a sum of individuals' willingnesses to pay. If this sum is negative, a set of individual-specific lump-sum transfers could reach an allocation that is *Pareto* superior to the alternative policy (Hicks (1940)). This means that everyone would prefer the transfers in the status quo to the alternative policy. Moreover, if the sum is positive, Kaldor (1939) argues that the winners

^{*}Harvard University, nhendren@fas.harvard.edu. This paper is a heavily revised version of a paper that previously circulated under the title, "The Inequality Deflator: Interpersonal Comparisons without a Social Welfare Function", NBER WP No. 20351. I am deeply indebted to conversations with Louis Kaplow for the inspiration behind this paper, and to Sarah Abraham, Alex Bell, Alex Olssen, and Evan Storms for excellent research assistance. I also thank Daron Acemoglu, Raj Chetty, Amy Finkelstein, Ben Lockwood, Henrik Kleven, Patrick Kline, Kory Kroft, Matthew Notowodigdo, Jim Poterba, Emmanuel Saez, Matthew Weinzierl, Glen Weyl, Ivan Werning, and Floris Zoutman, along with seminar participants at Berkeley, Harvard, MIT, Michigan, and Stanford for very helpful comments. The opinions expressed in this paper are those of the author alone and do not necessarily reflect the views of the Internal Revenue Service or the U.S. Treasury Department. This work is a component of a larger project examining the effects of tax expenditures on the budget deficit and economic activity, and this paper in particular provides a general characterization of the welfare impact of changes in tax expenditures relative to changes in tax rates (illustrated in Section 6). The empirical results derived from tax data that are reported in this paper are drawn from the SOI Working Paper "The Economic Impacts of Tax Expenditures: Evidence from Spatial Variation across the U.S.", approved under IRS contract TIRNO-12-P-00374.

could compensate the losers using individual-specific lump-sum transfers. Therefore, a modified policy that also includes compensating transfers would generate a Pareto improvement relative to the status quo. These two Pareto-based conceptual experiments underpin the definition of economic efficiency as the unweighted sum of individuals' willingnesses to pay.

In contrast to this canonical definition of efficiency, I argue that correctly measuring Kaldor and Hicks' notion of economic efficiency requires distributional weighting of surplus. This is because the redistributive experiments envisioned by Kaldor and Hicks involve distortionary costs. Individualspecific lump-sum transfers are not feasible. In originally defining their conceptual experiments, Kaldor and Hicks both suggested that they would like to incorporate these costs:

"Since almost every conceivable kind of compensation (re-arrangement of taxation, for example) must itself be expected to have some influence on production, the task of the welfare economist is not completed until he has envisaged the total effects...If, as will often happen, the best methods of compensation feasible involve some loss in productive efficiency, this loss will have to be taken into account" (Hicks (1939), p712).

The contribution of this paper is to provide a set of "efficient" social welfare weights that modifies the Kaldor-Hicks experiments to incorporate the distortionary cost of redistribution through the tax and transfer system. The efficient social welfare weight at an income level y is equal to the marginal cost of providing \$1 of welfare to individuals earning near y. In a world without distortionary taxation, this cost is \$1. But, it differs from \$1 when individual choose to change their earnings in response to a \$1 tax cut to those earning near y. For example, those earning below y might increase their incomes to y; those above y may decrease their incomes to y. By the envelope theorem, those who change their earnings will not obtain a first-order utility improvement from the transfer; but, in the presence of taxes/transfers these responses have a first-order impact on government revenue. If taxes are positive (negative), increases in incomes create positive (negative) fiscal externalities. The efficient welfare weights differ from unity because of these fiscal externalities from redistributive taxes and transfers.

Weighting individuals' willingnesses to pay by efficient welfare weights implements the Kaldor and Hicks' tests of economic efficiency. If weighted surplus is negative, then modifications to the income tax schedule can lead to all points of the income distribution being better off than would be the case under the alternative environment. In this sense, the economist can tell the policymaker of a better way to obtain the distribution of welfare gains offered by the alternative environment using a simple modification to the income tax schedule. Conversely, if the weighted surplus is positive, then a modified version of the alternative environment that includes an offsetting tax of the benefits can generate a Pareto improvement relative to the status quo. In this sense, the economist can prescribe a feasible policy that makes everyone better off. Thus, efficient welfare weights provide a toolkit for generating economic policy advice without relying on anyone's particular social preferences for redistribution.

I provide bounds and point estimates for the efficient welfare weights for the US. To do so, I first provide conditions under which these weights can be written as a function of (a) the joint distribution of taxable income and marginal tax rates and (b) the behavioral response of taxable income to changes in taxation. Because of the general lack of consensus about the size of behavioral responses to taxation (Saez et al. (2012)), I begin with bounds that do not impose strong assumptions on the size of these behavioral responses. To that aim, I show that the shape of the income distribution - in particular the local Pareto parameter of the income distribution - plays a key role in determining the extent to which the weights are above or below 1.¹ I estimate this parameter at each income level using the universe of income tax returns in the US from 2012. This reveals that the Pareto parameter rises from near -1 at the bottom of the income distribution to near 2 at the top of the income distribution, crossing zero around the 60th quantile of the income distribution (around \$43K in ordinary income). As a result, weights are generally above one for those with incomes below \$43K, and below one for those with incomes above \$43K. This means that under weak assumptions about behavioral responses, it is efficient to weight surplus to the poor more than to the rich.

To move from bounds to point estimates, I draw upon estimates of taxable income elasticities from previous literature. For the baseline specification, a \$1 tax cut to those with high incomes from a reduction in marginal tax rates costs around \$0.65. At the other end of the income distribution, the estimates suggest that expansions of the earned income tax credit (EITC) by \$1 to low earners has a fiscal cost of around \$1.15 because additional transfers cause individuals to adjust their earnings to maximize their tax credits. Thus, the results suggest that efficient welfare weights decline from around 1.15 at the bottom of the income distribution to around 0.65 at the top. In other words, \$1 to the poor can be turned into \$1.5-2 to the rich through modifications to the tax schedule. The point estimates suggest that it is efficient to value surplus to the poor 1.5-2x more than surplus to the rich.

Often, a researcher is not interested in analyzing a feasible budget-neutral alternative environment, but rather is considering a policy of increasing or decreasing government expenditure on a policy. I show that one can assess the efficiency of such a policy by using the efficient welfare weights to weight the marginal value of public funds, MVPF, of the policy. The MVPF is the ratio of the beneficiaries' willingness to pay for the policy change to the total cost of the policy change to the government inclusive of any fiscal externalities from behavioral responses to the policy (Hendren (2016)). Traditional Kaldor-Hicks efficiency would correspond to testing whether MVPF > 1, which corresponds to asking whether the beneficiaries are willing to pay the net cost of the policy to the government. In contrast, for policies whose beneficiaries have earnings near \bar{y} , I show testing for efficiency should correspond to testing whether $g(\bar{y}) MVPF > 1$, where $g(\bar{y})$ is the efficient social welfare weight for those earning near \bar{y} . I illustrate how to apply the framework by taking estimates of the MVPF from Hendren (2016) for job training programs, housing subsidies, and food stamps. In the end, weighting the MVPF by the efficient welfare weights implements an implicit comparison of the value of the policy to the value of a tax cut with the same cost and distributional incidence.

Relation to Previous Literature This is not the first paper to recognize that incorporating the distortionary cost of transfers leads to a modification of the Kaldor and Hicks compensation principle.

¹The Pareto parameter is given by $-\left(1+\frac{yf'(y)}{f(y)}\right)$ where f(y) is the density of the income distribution.

Most closely, this paper is related to Coate (2000) who proposes an "efficiency approach" theoretical framework that incorporates the costs of redistribution into the Hicks criterion by comparing the policy to feasible alternatives such as distortionary redistribution through the tax schedule. Relative to this literature, the contribution of this paper is to provide a feasible empirical method to implement these comparisons. Coate (2000) writes: "An interesting problem for further research would be to investigate whether the efficiency approach might be approximately decentralised via a system of shadow prices which convey the cost of redistributing between different types of citizens." The efficient welfare weights are the shadow prices envisioned by Coate (2000).²

The derivation of efficient social welfare weights does not require the tax schedule to be at an optimum. But, if the tax schedule is set to maximize a social welfare function, the welfare weights are the solution to the "inverse optimum" program in optimal tax: they reveal the implicit preferences of those who are indifferent to modifications in the tax schedule (Bourguignon and Spadaro (2012); Zoutman et al. (2013, 2016)).³ In this sense, the estimates suggest those indifferent to the current shape of the tax schedule in the US implicitly value an additional \$1 to the rich as equivalent to an additional \$1.5-2 to the poor.⁴ The contribution of this paper is to show that those weights can evaluate the desirability of policies beyond modifications to the tax schedule, regardless of whether one agrees with these implicit social preferences.

Because the efficient welfare weights generate indifference to modifications to the tax schedule, they do not provide guidance on the optimal degree of redistribution. Setting an optimal tax schedule or in general the optimal amount of redistribution in society still requires a social welfare function. But, conditional on setting a tax schedule, applied researchers studying the impact of policies with distributional incidence can use the efficient welfare weights to ask whether the policy being studied is efficient given the estimated costs of redistribution. For example, as illustrated in Section 7, the efficient welfare weights help an empirical researcher studying food stamp expenditures (e.g. Hoynes and Schanzenbach (2012)) identify whether it would be efficient to increase expenditure on food stamps financed by a reduction in EITC benefits (or vice versa). In this sense, efficient welfare weights implement the analyses originally envisioned by Kaldor and Hicks:

"All that economics can, and should, do in this field, is to show, given the pattern of income-distribution desired, which is the most convenient way of bringing it about? (Kaldor (1939, p552))

While many may disagree about whether this is all that economics should do, this paper shows how

 $^{^{2}}$ This paper is also related to a literature providing special cases, such as weak separability of the utility function, under which one does not need to know the distributional cost of taxation. Appendix F shows how these results can be derived as a special case of the models considered here by imposing assumptions about the behavioral responses to the policy.

 $^{^{3}}$ One notable exception to this invertibility is if the weights are estimated to be negative. In the weights are negative, it suggests there is a local Pareto improvement: by reducing taxes, one can increase government revenue. Testing whether the weights are negative implements the test of Pareto efficiency of the tax schedule provided by Werning (2007).

⁴In concurrent work, Lockwood and Weinzierl (2014) estimate the solution to the inverse optimum program in the U.S. using aggregated data from the Congressional Budget Office. An advantage of the present paper is the use of tax data to precisely measure the bounds on the shape of these weights.

to do Kaldor's desired analysis correctly by accounting for the distortionary cost of taxation and transfers.

The rest of this paper proceeds as follows. Section 2 derives the efficient welfare weights in the context of a general model setup. Section 3 illustrates how the weights implement the modified Kaldor-Hicks efficiency experiments. Section 4 provides a representation of these weights using the distribution of income, tax rates, and behavioral elasticities. Section 5 provides estimates of the joint distribution of income and tax rates and illustrates how they place bounds on the shape of the efficient welfare weights. Section 6 moves from bounds to point estimates by incorporating behavioral elasticities from previous literature. Section 7 illustrates how to apply these weights to policies using the estimated causal effects of these policies and their implied marginal value of public funds. Section 8 concludes.

2 Model

This section develops a model of utility maximization subject to nonlinear income taxation in the spirit of Mirrlees (1971) and Saez (2001, 2002). The model is used both to define economic surplus offered by an alternative environment and will also be used to describe the cost of redistribution through modifications in the tax and transfer schedule.

2.1 Setup

There exists a population of agents indexed by θ , with population size normalized to 1.⁵ There is a status quo environment and an alternative environment. In the status quo environment, agents consumption, c, and earnings, y. I allow each agent of type θ to have a potentially different utility function, $u(c, y; \theta)$, over consumption and earnings. I will not impose restrictions on the distribution of θ , so without loss of generality one can think of θ as simply indexing people in the population.

Agents choose c and y to maximize utility subject to a budget constraint,

$$c \le y - T\left(y\right) + m$$

where T(y) is the taxes paid on earnings y and m is additional income beyond earnings.⁶ With a slight abuse of notation, I let $c(\theta)$ and $y(\theta)$ denote the resulting choices of type θ .

Let $v^0(\theta)$ denote the utility level obtained by type θ in the status quo environment. And, given a utility level v, define the expenditure function $e(v; \theta)$ to be the smallest value of m that is required for a type θ to obtain utility level v in the status quo environment.⁷

The goal is to evaluate the status quo relative to an alternative environment. To allow for a general set of possible alternative environments, let $u^a(c, y; \theta)$ denote the utility function for type θ in the

⁵Formally, I assume $\theta \in \Theta$, where Θ has measure μ . I then assume (Θ, μ) is a probability space with measure μ (and the usual sigma algebra) so that the entire population is normalized to $\mu(\Theta) = 1$ and one can use the law of iterated expectations by conditioning on subsets of the population, Θ .

⁶For simplicity, I assume T(y) is the same for everyone. In the empirical implementation, I allow T to vary with individual characteristics, such as the number of dependents, and marital status. See Section 5.1.

⁷Formally, $e(v; \theta) = \inf \{ m | \sup_{c, y} \{ u(c, y; \theta) | c \leq y - T(y) + m \} \geq v \}$. Duality implies that $e(v^0(\theta); \theta) = m$.

alternative environment. Let $T^a(\circ)$ denote the tax schedule in the alternative environment so that the budget constraint is given by $c \leq y - T^a(y) + m$. Define $v^a(\theta)$ to be the level of utility obtained in the alternative environment and $e^a(v;\theta)$ is the smallest value of m that is required for a type θ to obtain utility level v in the alternative environment.⁸ Individual θ 's willingness to pay (equivalent variation) for the alternative environment is then given by

$$s(\theta) = e(v^{a}(\theta); \theta) - e(v^{0}(\theta); \theta)$$
(1)

This is the amount of additional money a type θ would need in the status quo environment to be just as well off as in the alternative environment. In addition to this equivalent variation definition of willingness to pay, one could also construct a compensating variation measure using the expenditure function in the alternative environment, $cv(\theta) = e^a(v(\theta); \theta) - e^a(v^a(\theta); \theta)$. Because the distinction between equivalent and compensating variation is second order (Schlee (2013)) and all the results below will be first order approximations, I will largely abstract from the distinction between equivalent or compensating variation in the analysis that follows.

I primarily consider alternative environments in which the willingness to pay for the alternative environment, $s(\theta)$, is homogeneous in income. With an abuse of notation, I write the willingness to pay for the alternative environment by a type θ who chooses income $y(\theta)$ in the status quo world by s(y), which is shorthand for the willingness to pay of types θ that have income level y. When willingness to pay varies conditional on income, finding Pareto improvements using modifications to the tax schedule that transfers only based on observed income is of course more difficult. I return to a discussion of this in Section 3.4.

Given s(y), the goal is to ask (a) whether benefits, s(y), can be replicated through modifications in the tax schedule (i.e. the experiment in Hicks (1940)), and (b) whether there exists a modification to the alternative environment that includes an additional redistribution through the tax/transfer system that makes everyone better off relative to the status quo (i.e. the experiment in Kaldor (1939)). The feasibility of these modifications to the tax schedule will depend on how the impact of changes in taxation on government revenue.

2.2 Measuring Government Revenue

For any tax schedule, $\tilde{T}(\circ)$, let $R\left(\tilde{T}\right)$ denote government revenue in the status quo environment with tax schedule \tilde{T} , which is given by:

$$R\left(\tilde{T}\right) = E\left[\tilde{T}\left(\tilde{y}\left(\theta;\tilde{T}\right)\right)\right]$$

where $\tilde{y}\left(\theta;\tilde{T}\right)$ denotes the earnings choice of a type θ when facing tax schedule, \tilde{T} , in the status quo environment.

⁸Analogous to the definition in the status quo environment, $e^{a}(v;\theta) = \inf \{m| \sup \{u^{a}(c,y;\theta) | c \leq y - T^{a}(y) + m\} \geq v\}.$

Individuals may make discrete changes in their earnings, $\tilde{y}(\theta; \tilde{T})$, in response to small changes in taxes (e.g. enter/exit the labor force). Therefore, I do not want to impose that behavioral responses are continuous in response to tax changes. However, it will be useful to impose that aggregate government tax revenue (aggregated across all individuals) varies smoothly in response to changes in the tax schedule. Assumption 1 states this more formally.

Assumption 1. Suppose $\tilde{T}_{\epsilon}(y) = T(y) + \epsilon \sum_{j=1}^{N} T^{j}(y)$ for some functions T^{j} . Let $\tilde{T}_{\epsilon}^{j}(y) = T(y) + \epsilon T^{j}(y)$. Then R is continuously differentiable in ϵ and

$$\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}\right) = \sum_{j=1}^{N}\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\tilde{T}_{\epsilon}^{j}\right)$$

This assumption allows for discrete changes in individual behavior in response to small tax changes (e.g. taxes can induce labor market participation or discrete changes in earnings). But, it makes the natural assumption that when aggregating across the population, government revenue varies smoothly in response to smooth changes in the tax schedule.⁹

In the alternative environment, let $R^a\left(\tilde{T}\right) = E\left[\tilde{T}\left(\tilde{y}^a\left(\theta;\tilde{T}\right)\right)\right]$ denote the revenue raised from a tax schedule $\tilde{T}(\circ)$, where $\tilde{y}^a\left(\theta;\tilde{T}\right)$ is the earnings choice of type θ in the alternative environment facing tax schedule $\tilde{T}(\circ)$. I impose the same smoothness assumptions as Assumption 1 on revenue in the alternative environment R^a .

2.3 Marginal Cost of Redistribution

Suppose individuals who earn around y^* are willing to pay s(y) = \$1 for the alternative environment. How costly is it to replicate this \$1 benefit through a tax cut in the status quo environment? Figure 1 depicts a small tax deduction to those with earnings near y^* . To be precise, let $\eta, \epsilon > 0$ and fix a given income level y^* . Consider providing an additional $\$\eta$ to individuals in an ϵ -region near y^* . Define $\hat{T}(y; y^*, \epsilon, \eta)$ by

$$\hat{T}(y; y^*, \epsilon, \eta) = \begin{cases} T(y) & \text{if } y \notin \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right) \\ T(y) - \eta & \text{if } y \in \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right) \end{cases}$$

so that \hat{T} provides η additional resources to an ϵ -region of individuals earning between $y^* - \epsilon/2$ and $y^* + \epsilon/2$.¹⁰ By the envelope theorem, individuals with earnings between $y^* - \epsilon/2$ and $y^* + \epsilon/2$ will be

 $^{^{9}}$ Because Assumption 1 is only a statement about the differentiability of the aggregate government revenue – as opposed to the tax revenue collected from any particular agent – it is weaker than most assumptions imposed in existing literature. For example, it does not require concavity (or even continuity) of any individual's utility function.

¹⁰If the original tax schedule T(y) is continuous, then the tax schedule \hat{T} will be discontinuous at $y^* + \frac{\epsilon}{2}$ and $y^* - \frac{\epsilon}{2}$. However, the main results defining the marginal cost of taxation in equation (2) below will apply regardless of whether the tax schedule to either be continuous or discontinuous. Moreover, one could alternatively define a smooth modification to the tax schedule that "bumps out" the tax schedule near y^* and obtain an identical magnitude for the marginal costs of taxation.

willing to pay η to have a tax schedule given by $\hat{T}(y; y^*, \epsilon, \eta)$ instead of T(y).¹¹

If there were no behavioral responses to changes in taxes, then the cost to the government of providing η to those earning near y^* would be η per person, so that the mechanical marginal cost in the absence of behavioral responses is 1. But, the presence of behavioral responses induce a fiscal externality on the government. To capture this, consider the derivative of revenue, $R\left(\tilde{T}(\circ; y^*, \epsilon, \eta)\right)$, with respect to the size of the tax cut, η , and evaluate at $\eta = 0$. This yields the function $\frac{d[R(\tilde{T}(\circ; y^*, \epsilon, \eta))]}{d\eta}|_{\eta=0}$, which is the marginal cost of providing an additional dollar through the tax code to individuals with earnings in an ϵ -region of y^* . Then, taking the limit as $\epsilon \to 0$, one arrives at the marginal cost to the government of providing an additional dollar of resources to an individual earning y^* :

$$g\left(y^{*}\right) \equiv \lim_{\epsilon \to 0} \frac{d\left[R\left(\tilde{T}\left(\circ; y^{*}, \epsilon, \eta\right)\right)\right]}{d\eta}|_{\eta=0}$$

$$\tag{2}$$

Equation (2) will define the efficient welfare weights.

Definition 1. For each level of income, y, the **Efficient Welfare Weight**, g(y), is the marginal cost of providing an additional \$1 to those earning near y, as defined in equation (2).

The cost embodied in g(y) is inclusive of the impact of the behavioral response to the tax cut on the government budget.¹² If income did not respond to changes in taxes, this marginal cost would be \$1 per beneficiary, g(y) = 1. Hence, g(y) can be thought of as the sum of a mechanical component of \$1 plus the impact of the behavioral response to the policy on the government budget, FE(y),

$$g\left(y\right) = 1 + FE\left(y\right)$$

These fiscal externalities reflect the impact of the tax cut on taxable income. If the increased transfer causes people to work less and thereby reduces tax revenue (FE(y) > 0); conversely, if increasing transfers to those earning near y causes people to work more and thus increases tax revenue, then the fiscal externality will be negative, FE(y) < 0. In general, the size of the fiscal externality is an empirical question and depends on the causal impact of tax changes. In Section 4, I provide additional assumptions that enable one to identify FE(y) using behavioral elasticities, the shape of the income distribution, and the shape of the tax schedule. Regardless of whether g(y) can be represented using elasticities, the general definition of g(y) is the amount it costs the government to transfer an additional \$1 to those currently earning near \$y.

¹¹More formally, this is true as long as the incidence of the tax cut falls entirely on the beneficiaries and does not result in changes in wages. For example, if firms respond to the tax cut of \$1 by lowering wages by \$0.50, then the individual would only be willing to pay \$0.50 for a \$1 tax cut. Here, I assume no general equilibrium responses, but this could be incorporated into future work. I discuss this further in Section 3.4.

¹²For simplicity, I use y to denote the argument of g(y) instead of y^* ; as shown by replacing y^* with y in equation (2), g(y) is the marginal cost of a tax cut to those earning near y.



Figure 1: Tax Cut to Those Earning Near y^*

Notes: This figure illustrates the modification to the tax schedule that provides a tax cut of \$1 to those with earnings in a region of y^* of width ϵ . To first order, those whose earnings would lie in $\left[y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right]$ will value the tax cut at \$1. But, the costs will result from both this mechanical cost and the impact of behavioral responses to the tax cut (loosely illustrated by the blue arrows). So, the total cost per unit of mechanical beneficiary will be g(y) = 1 + FE(y).

Finally, g(y) is the marginal cost of providing an additional \$1 to those earning y in the status quo environment. For the alternative environment, I define $g^a(y)$ analogously to be the cost of providing an additional \$1 to those earning y in the alternative environment.¹³

3 Using the Efficient Welfare Weights

To set the stage, consider the willingness to pay for the alternative environment, s(y). Is the policy desirable? A common method for answering this question is to weight this willingness to pay by a generalized social welfare weight Saez and Stantcheva (2016). Such weights incorporate the social preferences of the policy maker or researcher to balance the gains to the winners (e.g. for whom s(y) > 0) against the losses to the losers (e.g. for whom s(y) < 0). In contrast to imposing any particular social preference, in this section I show that it is efficient – in the sense of Kaldor and Hicks – to weight this willingness to pay by the efficient welfare weights, g(y), defined in equation (2). I

$$g^{a}\left(\boldsymbol{y}^{*}\right) \equiv \lim_{\epsilon \to 0} \frac{d\left[R\left(\tilde{T}\left(\circ;\boldsymbol{y}^{*},\epsilon,\eta\right)\right)\right]}{d\eta}|_{\eta=0}$$

¹³Formally, let $\tilde{T}(\circ; y^*, \epsilon, \eta)$ denote a modified tax schedule in the alternative environment. Then,

is the marginal cost of providing additional resources to those with earnings y^* in the alternative environment.

define efficient surplus to be the amount of resources the government needs to replicate the surplus offered by the alternative environment using modifications to the tax schedule, S = E[s(y)g(y)].

Definition 2. Efficient Surplus is given by a weighted average of surplus:

$$S = E[s(y)g(y)]$$
(3)

The next subsections illustrate how these weights and surplus measures implement the classic tests for efficiency in Kaldor and Hicks, modified to account for the distortionary cost of redistribution. I begin in Section 3.1to show that testing S > 0 provides a test of efficiency in the spirit of Hicks (1940) that replicates the surplus offered by the alternative environment using modifications to the tax schedule. In Section 3.2, I show that testing S > 0 provides a test of efficiency in the spirit of Kaldor (1939) that modifies the tax schedule in the alternative environment to search for an Pareto superior allocation that is preferred by everyone relative to the status quo.

3.1 Testing for Efficiency in the Spirit of Hicks (1940) and Coate (2000)

Can the benefits offered by the alternative environment, s(y), be more efficiently provided through modifications in the tax schedule? To assess this, imagine replacing the current tax schedule, T(y), with a new tax schedule, $\hat{T}(y) = T(y) - s(y)$, that offers a tax cut of size s(y) to those earning y. Figure 2 provides an illustration. Panel A presents a hypothetical alternative environment that is preferred by the poor but not by the rich. Panel B then modifies the tax schedule from T(y) to T(y) - s(y). To first order, the envelope theorem implies that the tax cut of s(y) is valued at s(y) by those earning y. Therefore, everyone is approximately indifferent between the alternative environment and the status quo environment with the modified tax schedule, as depicted by the dashed red line in Figure 2, Panel B. The Hicks test for efficiency asks: Is this tax modification in the status quo world feasible?

To first order, the marginal cost of providing \$1 of welfare to those earning y is given by g(y) = 1 + FE(y). Therefore, the cost of this tax cut is given by E[g(y) s(y)]. If this quantity is positive, then providing surplus s(y) through the tax schedule would not be feasible. Closing the budget constraint by raising taxes on everyone would lead to the blue line in Figure 3, Panel A. In this sense, the alternative environment would be efficient relative to what is feasible through modifications to the tax schedule in the status quo. In contrast, if S < 0, then it is possible to replicate the alternative environment through modifications to the tax schedule. Redistributing the government surplus to everyone equally leads to the blue line in Figure 3, Panel B, which is preferred by all relative to the alternative environment. In this sense, the alternative environment is efficient if and only if S > 0.

Figure 2: Hicks (1940) Efficiency Experiment



Notes: This figure illustrates the efficiency experiment of Hicks (1940) for a hypothetical alternative environment. Panel A presents the hypothetical willingness to pay for each person at different points of the income distribution. In this example, those with low incomes prefer the alternative environment, but those with higher incomes prefer the status quo. Panel B illustrates modifying the tax schedule in the status quo world to attempt to replicate the surplus offered by the alternative environment. To first order, everyone is indifferent between the alternative environment and the modified status quo with tax schedule T(y) - s(y).

Formally, these comparisons rely on the envelope theorem to ensure indifference between the modified status quo (offered by the dashed red line in Figure 2, Panel B) and the alternative environment. Proposition 1 makes these local statement more formally.

Proposition 1. For any $\epsilon > 0$ define the scaled surplus by $s_{\epsilon}(y) = \epsilon s(y)$ and $S_{\epsilon} = E[s_{\epsilon}(y)g(y)] = \epsilon S$. If S < 0, there exists an $\tilde{\epsilon} > 0$ such that for any $\epsilon < \tilde{\epsilon}$ there exists an augmentation to the tax schedule in the status quo environment that generates surplus, $s_{\epsilon}^{t}(y)$, that is uniformly greater than the surplus offered by the alternative environment, $s_{\epsilon}^{t}(y) > s_{\epsilon}(y)$ for all y. Conversely, if S > 0, no such $\tilde{\epsilon}$ exists.

Proof. See Appendix A.2.

Because these statements are only technically valid to first order, one should be cautious when making statements about very large changes in environments. The efficient welfare weights provide the right direction for adjusting for the marginal cost of redistribution, but the costs and benefits of such redistribution could change for larger movements in the tax schedule. Moving beyond this firstorder approach is an interesting direction for future work. In the meantime, testing S > 0 provides first-order guidance on how to correctly implement Hicks' original experiment in a way that accounts for the distortionary cost of redistribution.

Figure 3: Testing for Hicks Efficiency

A. Alternative Environment is Efficient, S > 0

B. Alternative Environment is Inefficient, S < 0



Notes: This figure illustrates the efficiency test of Hicks (1939). The blue line illustrates the conceptual after-tax income that is feasible through modifications to the tax schedule but has the same distributional incidence as the alternative environment. Panel A illustrates the case in which the modified status quo tax schedule would deliver lower welfare to all points of the income distribution, so that the alternative environment is efficient relative to the status quo, S > 0. In contrast, Panel B illustrates the case in which replicating the surplus offered by the alternative environment through the tax schedule leads to higher welfare for all, so that the alternative environment is inefficient.

Normative Implications: Coate (2000) versus Hicks (1940) Are efficient policies socially desirable? If S < 0, then the alternative environment is Pareto dominated by a modification to the tax schedule. In this sense, alternative environments for which S < 0 are not desirable. But, what about policies for which S > 0? Should these be pursued?

This is a point of debate in previous literature. While Hicks (1940) originally suggested yes, Coate (2000) raises the concern that moving to the alternative environment does not generate a Pareto improvement relative to the status quo; rather, it generates a Pareto improvement relative to a modified status quo that attempts to replicate the distributional incidence of the alternative environment. Actually moving to the alternative environment would generate winners and losers. Hence, S > 0 suggests it is a useful policy to consider (it's an "efficient" policy) but, given social preferences it may not be desirable relative to the status quo. Indeed, the status quo can also be an efficient environment if there are no other policies that generate Pareto improvements.

While the Hicks experiment provides weak guidance in the case of S > 0, it turns out that a different experiment – namely, that of Kaldor (1939) – will provide stronger guidance by considering modifications to the tax schedule in the alternative environment.

3.2 Finding Pareto Improvements in the Spirit of Kaldor (1939)

When can everyone be made better off relative to the status quo environment? Consider modifying the tax schedule in the alternative environment, $T^{a}(y)$, so that the winners compensate the losers,

 $T^{a}\left(y\right) \rightarrow T^{a}\left(y\right) + s\left(y\right).$

Figure 4 presents this modified tax schedule in the alternative environment. Those with incomes y are better off by s(y) relative to the status quo. The dashed red line in Figure 4 taxes back these gains. The envelope theorem suggests that to first order individuals earning y in the alternative environment are worse off by s(y) when we tax back these benefits. To first order, everyone is indifferent between the status quo environment and the alternative environment with the modified income tax schedule. Therefore, the question becomes: Is this modification to the tax schedule in the alternative environment budget feasible?



Figure 4: Testing for (Kaldor) Efficiency

Notes: This figure illustrates the test of efficiency in Kaldor (1939) that modifies the tax schedule in the alternative environment to find a Pareto improvement in the modified alternative environment relative to the status quo. The dashed red line presents the after-tax schedule that adds the surplus offered by the alternative environment to the tax schedule, T(y) + s(y). To first order, everyone is indifferent between the status quo and the modified alternative environment illustrated by the dashed red line in Panels A and B. The dash-dot blue line then illustrates the after tax income curve that results from closing the government budget constraint. Panel A illustrates the case that the alternative environment is efficient, so that after modifying the tax schedule in the alternative environment there is a Pareto improvement relative to the status quo. Panel B illustrates the case where the alternative environment is inefficient, so that after taxing back the benefits of the alternative environment and closing the budget constraint everyone is worse off relative to the status quo.

The modification will generate revenue $g^a(y) s(y)$ at each point y, where $g^a(y) = 1 + FE^a(y)$ is the cost to the government of providing \$1 to those earning near \$y in the alternative environment. Aggregating across all values of y yields $S^a = E[g^a(y) s(y)]$. If $S^a > 0$, it is feasible to find a modified alternative environment that taxes the benefits and uniformly redistributes them to make everyone better off relative to the status quo.

To search for potential Pareto improvements, one needs to weight surplus by the marginal cost of providing a tax cut in the alternative environment, $g^{a}(y)$, which can differ from the marginal cost in the status quo environment, g(y). But in practice, it will be reasonable to assume in many applications involve sufficiently small changes to the structure of the economy so that the marginal cost of taxation is similar in the status quo and alternative environments, $g^{a}(y) \approx g(y)$. I state this formally in Assumption 2.

Assumption 2. For sufficiently small $\tilde{\epsilon}$, the marginal cost of taxation, g(y), in the alternative environment is the same as in the status quo. Specifically, there exists $\tilde{\epsilon}$ such that if $\epsilon \in (0, \tilde{\epsilon})$, then (1) $y^{\epsilon}(\theta)$ is the same for all types θ that had the same income in the status quo world, $y^{\epsilon}(\theta) = y^{\epsilon}(\theta')$ iff $y(\theta) = y(\theta')$, and (2) $g(y(\theta)) = g(y^{\epsilon}(\theta))$ for all θ .

If Assumption 2 holds, then $S \approx S^a$ so that efficient surplus measures the amount of resources the government can obtain in the alternative environment if it taxes back the benefits offered by the alternative environment. Up to this approximation that the marginal cost of taxation is not changing, one can replicate the alternative environment through modifications to the status quo tax schedule if and only if one can modify the tax schedule in the alternative environment so that the winners compensate the losers.¹⁴ In this sense, testing whether S > 0 provides a first-order approximation to searching for potential Pareto improvements as suggested by Kaldor.

As with the Hick's experiment, using S to search for Pareto comparisons is valid to first order. Proposition 2 states the results more formally.

Proposition 2. Suppose Assumption 2 holds. For $\epsilon > 0$, let $s_{\epsilon} = \epsilon s(y)$. If S > 0, there exists $\tilde{\epsilon} > 0$ such that for any $\epsilon < \tilde{\epsilon}$, there exists an augmentation to the tax schedule in the alternative environment that delivers surplus $s_{\epsilon}^{t}(y)$ that is positive at all points along the income distribution, $s_{\epsilon}^{t}(y) > 0$ for all y. Conversely, if S < 0, then no such $\tilde{\epsilon}$ exists.

Proof. See Appendix A.3.

To first order, S > 0 implies that those who benefit from the alternative environment could compensate those who are made worse off through modifications to the income tax schedule in the alternative environment.

Summary Table 1 summarizes the main results. When efficient surplus is negative, S < 0, the alternative environment is inefficient in the sense that a feasible modification to the tax schedule in the status quo environment can lead to a Pareto superior allocation to the alternative environment. In this sense, alternative environments for which S < 0 can be rejected by the logic of Hicks (1940) and Coate (2000).

When efficient surplus is positive, S > 0, a modified alternative environment in which the winners to compensate the losers through modifications to the tax schedule can offer a Pareto superior allocation relative to the status quo. But as noted by Coate (2000), it is not clear why the alternative environment is desirable if these hypothetical compensating transfers are not actually implemented. If they are not, then the alternative environment does not itself offer a Pareto improvement relative to the status quo.

¹⁴If a particular application involves a sufficiently large change to the economy so that one worries the marginal cost of taxation has significantly changed, one can certainly estimate separately g^a and g functions and conduct the efficiency analyses using separate weights.

In this sense, whether one prefers the alternative environment depends on one's social preferences (Coate (2000)). But, one can recommend instead a different (yet feasible) alternative environment in which the winners compensate the losers through modifications to the tax schedule.

Table 1				
	S > 0	S < 0		
Hicks Experiment: Possible to replicate s(y) using tax cut in status quo?	No	Yes		
Kaldor Experiment: Possible to modify alternative environment tax schedule to make everyone better off relative to status quo?	Yes	No		

3.3 Relation to Previous Literature

3.3.1 Relation to Social Welfare Function and Inverse Optimum

Table 1 shows how efficient surplus, S, can evaluate policy changes by searching for potential Pareto improvements. In this sense one can make a recommendation about the alternative environment without resorting to a social preference that weighs the benefits to the winners against the losses to the lossers. This contrasts with a more common approach to use a social welfare function to evaluate a policy change.

While the conceptual underpinnings in the Kaldor and Hicks experiments that rely on the Pareto principle differ from the social welfare function approach, the efficient welfare weights are related to the social welfare function approach. In particular, the weights equal the implicit social welfare weights that rationalize indifference to modifications to the tax schedule. If one's social preferences were reflected by these weights, then one would prefer to make no modifications to T(y).

Appendix D provides a formal derivation. To see the logic, let $\chi(y)$ denote the social marginal utilities of income for those earning near y, so that an additional \$1 to an individual earning y has an impact of $\chi(y)$ on social welfare. Suppose one provides a small tax cut of \$1 to those earning near y. Those with incomes near y will be willing to pay \$1 for this tax cut, and it will generate a social welfare impact of $1 * \chi(y)$. But, it will have a cost of 1 + FE(y) = g(y). Hence, the marginal value of additional government spending on a tax cut to those earning near y will be given by $\frac{\chi(y)}{g(y)}$. If the tax schedule is set to maximize social welfare, then the government must be indifferent between raising \$1 from those earning y' to finance a tax cut to those earning y. In other words, $\frac{\chi(y)}{g(y)}$ must be constant for all y;

$$\frac{\chi\left(y\right)}{g\left(y\right)} = \kappa \quad \forall y$$

So, $\chi(y) = \kappa g(y)$. Since social welfare weights are only defined up to a constant, g(y) is the unique set of social welfare weights that rationalize the tax schedule as optimal. Under the assumption that the political process setting T(y) leads to the reflection of an individual's or political parties' preferences, these weights then reveal the implicit preferences of those setting the tax schedule.

In this sense, the efficient welfare weights are related to a literature estimating the inverse optimum program of optimal taxation. This literature estimates the implicit social welfare weights that rationalize the status quo tax schedule as optimal and uses them to infer properties of social or political preferences for redistribution (Bourguignon and Spadaro (2012); Zoutman et al. (2013, 2016)). Relative to this literature, the key contribution of this paper is to provide a justification based on the Pareto principle for using these weights to evaluate policies beyond modifications to the tax schedule. Although the weights generate indifference to modification to the tax schedule, they implement the Kaldor-Hicks tests of efficiency for other polices of interest. Regardless of one's own social preferences (and even if they differ from those setting the tax schedule), the efficient welfare weights characterize the existence of Pareto improvements.

3.3.2 Negative Weights and Werning (2007)

There is a special case when the weights, g(y), identify a Pareto improvement directly through modifications to the tax schedule. In particular, if one finds that for some value of y the weights are negative, g(y) < 0, then the marginal cost of taxation is positive. This means that the government can raise revenue by providing a tax cut to those earning near y. In this sense, testing whether g(y) < 0 implements a test in the spirit of Werning (2007) that searches for local Laffer effects in the tax schedule. A finding of g(y) < 0 makes a straightforward policy recommendation: cut taxes to those earning near y.¹⁵

3.4 Limitations

Efficient social welfare weights facilitate the search for Pareto comparisons. But, before turning to the empirical implementation and estimation, it is useful to highlight several of the limitations of this approach.

First, the approach relies on the envelope theorem to measure the marginal cost of taxation. If a policy change is large enough, one could potentially be concerned that the modifications required to mimic the policies through the tax schedule may lead to a change in the marginal cost of taxation. For example, an earlier draft of this paper used these weights to compare the income distributions across countries. However, as will be clear below, the basic fact that the marginal cost of taxation is higher

¹⁵Finding g(y) < 0 generates a conceptual difficulty in making Pareto comparisons between modified environments, as there would already exist a Pareto improvement within the status quo environment through modifying the tax schedule. However, in the applications and estimation below, it will turn out that g(y) > 0 for all y. So, in practice this conceptual difficulty will not arise.

for the rich than the poor is quite robust because of the general shape of the income distribution. Thus, the efficient weights are more applicable to uses such as weighting the distributional incidence of food stamp expansions and other changes to expenditures on redistributive policies.

Second, the approach assumes that tax changes have no general equilibrium or spillover effects. Targeting a \$1 tax cut to those earning near y is assumed to have a willingness to pay of \$1 for the beneficiaries of the tax cut. But, if their wages change in response to the tax cut, their willingness to pay may differ from \$1. Indeed, with spillovers and general equilibrium effects, the benefits of the tax cut may extend beyond those who are the direct target of the tax cut. Importantly, the approach does allow GE effects to drive the valuation of the alternative environment, s(y). For example, the alternative environment could be a policy that makes more land available for agriculture, which in turn lowers food prices. One can still generate individuals' willingness to pay for this alternative environment, s(y), and use the efficient welfare weights to ask whether this policy is efficient. In this sense, the efficient welfare weights, g(y), are valid even if the policy change or alternative environment has GE effects; but it has ruled out the case where changes in the tax schedule, T(y), has GE effects. I leave the incorporation of such effects for future work.

Lastly, alternative environments may generate willingness to pay that is heterogeneous conditional on income. In this case, Pareto comparisons are more difficult. To test for Hicks efficiency, one needs to construct the maximum willingness to pay at each income level, $\bar{s}(y)$, and test whether $E[\bar{s}(y)g(y)] > 0$. If it is negative, then it would be feasible for the government to replicate the surplus offered by the alternative environment and make everyone better off. Intuitively, the government can feasibly provide a tax cut that covers even the maximal willingness to pay at each income level, $\bar{s}(y)$. In this sense, the alternative environment would be inefficient. Conversely, to test for Kaldor efficiency, one needs to construct the minimum willingness to pay at each income level, $\underline{s}(y)$, and test whether $E[\underline{s}(y)g(y)] > 0$. If it is positive, then it would be feasible for the government to redistribute income in the alternative environment so that everyone prefers the modified alternative environment relative to the status quo. Appendix C provides formal statements and proofs of these claims. Often, one might find that $E[\underline{s}(y)g(y)] < 0$ and $E[\overline{s}(y)g(y)] > 0$. In this instance, the alternative environment cannot not be Pareto-ranked relative to the status quo. Nonetheless, the efficient welfare weights, g(y), continue to be the key component required to measure $E[\underline{s}(y)g(y)]$ and $E[\overline{s}(y)g(y)]$ that facilitates the search for these Pareto comparisons.

4 Representing Fiscal Externalities using Estimable Parameters

Recall that the marginal cost of providing a \$1 tax cut to those with earnings near y, as illustrated in Figure 1, is given by g(y) = 1 + FE(y), where FE(y) is the impact of the behavioral response to the tax cut on government tax revenue. This section shows how these fiscal externalities can be written as a function of behavioral elasticities and the joint distribution of income and marginal tax rates. To do so, I impose the common assumption that taxation can cause individuals to enter (or leave) the labor force, but intensive margin adjustments are continuous in the tax rate. Assumption 3 states this more precisely in the context of the general model developed in Section 2.

Assumption 3. For any $\kappa > 0$, let $B(\kappa) = [u(y(\theta) - T(y(\theta)), y(\theta); \theta) - \kappa, u(y(\theta) - T(y(\theta)), y(\theta); \theta) + \kappa]$ denote an interval of width κ near the status quo utility level. For any level of earnings y and utility level w, let $c(y; w, \theta)$ trace out a type θ 's indifference curve that is defined implicitly by:

$$u(c(y; w, \theta), y; \theta) = w$$

I assume each indifference curve, $c(y; w, \theta)$, satisfies the following conditions:

- 1. (Continuously differentiable in utility) For each $y \ge 0$, there exists $\kappa > 0$ such that $c(y; w, \theta)$ is continuously differentiable in w for all $w \in B(\kappa)$
- 2. (Convex in y for positive earnings, but arbitrary participation decision) For each y > 0, there exists $\kappa > 0$ such that $c(y; w, \theta)$ is twice continuously differentiable in y for all $w \in B(\kappa)$ and $c_{yy} > 0$ and $c_{yy} > 0$.

Part (1) imposes the standard assumption that indifference curves move smoothly with utility changes. Part (2) requires that indifference curves are convex on the region y > 0 (but not at y = 0). Importantly, it allows even small changes in the tax schedule to cause jumps between y = 0 and some positive income level (i.e. a participation response). Assumption 3 imposes very weak assumptions on utility functions and also allows for arbitrary distributions of unobserved heterogeneity, θ . ¹⁶

If the utility function satisfies Assumption 3, then three behavioral elasticities (a compensated elasticity, income elasticity, and participation elasticity) determine the response to taxation. To define these, let $\tau(y) = T'(y)$ denote the marginal tax rate faced by an individual earning y. The average intensive margin compensated elasticity of earnings with respect to the marginal keep rate, $1 - \tau(y)$, for those earning $y(\theta) = y$ is given by the percent change in earnings from a percent change in the price of consumption,

$$\epsilon^{c}\left(y\right) = E\left[\frac{1-\tau\left(y\left(\theta\right)\right)}{y\left(\theta\right)}\frac{dy}{d\left(1-\tau\right)}|_{u=u\left(c,y;\theta\right)}|y\left(\theta\right)=y\right].$$

The average income elasticity of earnings, $\zeta(y)$, is given by the percentage response in earnings to a percent increase consumption,

$$\zeta \left(y \right) = E \left[\frac{dy\left(\theta \right)}{dm} \frac{y\left(\theta \right) - T\left(y\left(\theta \right) \right)}{y\left(\theta \right)} | y\left(\theta \right) = y \right]$$

The extensive margin (participation) elasticity with respect to net of tax earnings, $\epsilon^{P}(y)$, is given by

$$\epsilon^{P}(y) = \frac{d\left[f(y)\right]}{d\left[y - T(y)\right]} \frac{y - T(y)}{f(y)}$$

¹⁶See Kleven and Kreiner (2006) for a particular utility specification that satisfies Assumption 3 and captures these features of intensive and extensive margin labor supply responses.

where f(y) is the density of income at y.

Proposition 3 shows how these three elasticities help characterize the fiscal externality, FE(y), from taxation.

Proposition 3. For any point y such that $\tau(y)$ and $\epsilon^{c}(y)$ are constant in y and the distribution of y is continuous with density f(y), the fiscal externality of providing additional resources to individuals near y is given by

$$FE(y) = -\underbrace{\epsilon_c^P(y) \frac{T(y) - T(0)}{y - T(y)}}_{Participation \ Effect} - \underbrace{\zeta(y) \frac{\tau(y)}{1 - \frac{T(y)}{y}}}_{Income \ Effect} - \underbrace{\epsilon^c(y) \frac{\tau(y)}{1 - \tau(y)} \alpha(y)}_{Substitution \ Effect}$$
(4)

where $\alpha(y) = -\left(1 + \frac{yf'(y)}{f(y)}\right)$ is the local Pareto parameter of the income distribution.

Proof. Proof provided in Appendix A The appendix also provides a generalized formula for points y such that $\epsilon^{c}(y)$ is not constant in y.¹⁷

The fiscal externality associated with providing an additional dollar resources to an individual earning y is the sum of three effects. First, people may enter the labor force. $\epsilon^{c}(y)$ measures the size of this effect. It's impact on tax revenue depends on the difference between the average taxes received at y, T(y), and the taxes/transfers received from those out of the labor force, T(0).

Second, the increased transfer may change the labor supply of those earning y due to an income effect. The size of this effect is measured by $\zeta(y)$. The impact of this change in earnings on the government budget depends on the marginal tax rate, $\tau(y)$.

Finally, people earning close to y may change their earnings towards y in order to get the transfer. The elasticity, $\epsilon^c(y)$, measures how much people move their earnings towards y in response to the tax cut. The tax ratio, $\frac{\tau(y)}{1-\tau(y)}$, captures the fiscal impact of these responses. However, the net impact on government revenue is the sum of two effects. Some people will decrease their earnings towards y; others will increase their earnings towards y, as depicted by the blue arrows in Figure 1. For positive tax rates, $\tau(y) > 0$, the former effect increases tax revenue and the latter effect decreases tax revenue. The extent to which the losses outweigh the gains depends on the elasticity of the income distribution, $\frac{yf'(y)}{f(y)}$. When $\frac{yf'(y)}{f(y)} < -1$ (as is the case with the Pareto upper tails in the US income distribution), more people increase rather than decrease their taxable earnings. This means $\alpha(y) > 0$. Conversely, if $\frac{yf'(y)}{f(y)} > -1$ (e.g. if f is a uniform distribution so that f'(y) = 0), then more people decrease than increase their earnings so that $\alpha(y) < 0$. This increases the marginal cost of the tax cut. Importantly, this shows that even if elasticities and tax rates are constant, the shape of the income distribution plays a key role in determining the marginal cost of taxation and the shape of efficient welfare weights.

¹⁷Proposition 3 is a generalization of the canonical optimal tax formula to the case of multi-dimensional heterogeneity. Consistent with the intuition provided by Saez (2001), Proposition 3 shows that the relevant empirical elasticities in the case of potentially multi-dimensional heterogeneity are the population average elasticities conditional on income. Zoutman et al. (2016) provide a derivation of equation 4 for the case when θ is uni-dimensional and distributed over an interval and for which the Spence-Mirrlees single crossing condition holds.

5 Bounds on Efficient Welfare Weights

At first glance, equation suggests one requires precise estimates of the size of behavioral responses to taxation in order to quantify the efficient welfare weights. This is potentially problematic because of the general lack of consensus on the size of behavioral responses to taxation (Saez et al. (2012)). Fortunately, under fairly plausible assumptions outlined below, it will turn out that this shape will lead to bounds on the shape of the efficient welfare weights. Regardless of the precise magnitudes of the behavioral responses to taxation, the shape of the income distribution combined with the shape of the tax schedule suggests that efficient welfare weights place higher weight on those with lower than higher incomes.

5.1 Sample and Variables

I use the universe of de-identified 2012 tax returns taken from the 2012 IRS-SOI Databank maintained under the Statistics of Income Division at the IRS. I focus on primary filers aged 25-60 and their married spouses, if applicable.¹⁸ Following Chetty et al. (2014), I restrict the sample to households with positive family income. Details of the sample and data construction are provided in Appendix B; Appendix Table I presents the summary statistics. The resulting sample has roughly 100 million filing units.

I define y to be the tax filer's ordinary income in 2012.¹⁹ This equals taxable income (f1040, line 43) minus income not subject to the ordinary income tax (long-term capital income (line 13) and qualified dividends (line 9b)). Ordinary income is primarily comprised of labor income, but subtracts deductions for things like the number of children and charitable donations.

To each tax return, I assign the marginal tax rate faced by the 2012 federal income tax schedule. The federal rate schedule on ordinary income provides the marginal tax rate for the filer as long as s/he did not have any additional tax credits, such as the earned income tax credit, and was not subject to the alternative minimum tax. If the individual was subject to the alternative minimum tax (AMT), I record their marginal tax rate at the 28% AMT level. If the filer received EITC, I add the marginal tax rate on the EITC schedule using information on the number of EITC-eligible children (reported in the tax return), filing status, and the size of the EITC benefit claimed. This provides a precise measure of the federal marginal tax rate faced by each filer on an additional dollar of ordinary income.

In addition to federal taxes, I account for state and local taxes as follows. For state taxes, I assume a constant tax rate of 5% and account for the fact that state taxes are deductible from federal tax liability when calculating the total marginal tax rate. For Medicare and sales taxes, I follow Saez et al. (2012) and assume a 2.9% tax rate for Medicare and a 2.3% sales tax rate. Finally, some states provide additional EITC benefits. To account for this, I assume a 10% "top-up" EITC rate for EITC

¹⁸I exclude individuals below age 25 because of the likelihood they still live at home and are part of another household. I exclude people above 60, the age at which many begin exiting the labor force and begin collecting unearned income such as social security income or savings withdrawals.

¹⁹Because ordinary income determines the federal tax, it is the notion of income that most closely aligns with the theory.

filers.²⁰ In the end, this generates a marginal tax rate, $\tau(y)$, faced by each filer on an additional dollar of income.

5.2 Estimation of Shape of Income Distribution, $\alpha(y)$

I estimate $\alpha(y)$ for each level of income separately for those facing different tax schedules. The key additional complication that one needs to address is how to deal with the fact that two tax filers with the same ordinary income may face different tax schedules. In practice, there are many factors that lead two filers with the same taxable income to face different tax rates, such as filing status, EITC eligibility, and subjectivity to the AMT. The combination of filing status, EITC status (which includes marital status plus number of qualifying dependents), and AMT subjectivity leads to more than 100 different tax schedules, each with numerous different marginal tax rates along the income distribution. With multiple tax schedules, the average fiscal externality at each level of income requires estimating the average value of the substitution effect, $\epsilon \frac{\tau(y)}{1-\tau(y)} \alpha(y)$, separately for each tax schedule, $\tau(y)$, and then averaging across all individuals at each income level, y. To construct this, I separately estimate $\alpha(y)$ for each tax schedule using the information in the tax returns on filing status and other determinants of the tax schedule. The details of this procedure are provided in Appendix B.

²⁰Choosing alternative values for the state tax rates or the EITC rates do not significantly alter the results; as discussed below, the primary primary driver of the shape of the weights is the Pareto parameter combined with the assumption of a constant elasticity, not the shape of tax rates, $\tau(y)$.





Notes: This figure presents the estimates of the average $\alpha(y)$ for each quantile of the income distribution. This function is given by $\alpha(y) = -\left(1 + \frac{yf'(y)}{f(y)}\right)$, where f(y) is the density of the income distribution. For values of y below the 60th quantile, $\alpha(y) < 0$ so that the substitution effect in equation (4) raises the marginal cost of taxation. In contrast, for values of y above the 61st quantile, $\alpha(y) > 0$ so that the substitution effect lowers the marginal cost of taxation.

Figure 5 provides a picture of the mean value of $\alpha(y)$ at each quantile of the ordinary income distribution. Consistent with findings in existing literature, the average $\alpha(y)$ reaches around 2.5 in the upper regions of the income distribution. This corresponding to an elasticity of the income distribution of 1.5, consistent with findings in previous literature (Diamond and Saez (2011) and Piketty and Saez (2013)).

Implications for Bounds on g(y) For the purposes of quantifying g(y), the key point on Figure 5 is that $\alpha(y)$ exhibits considerable heterogeneity across the income distribution. It is negative below the 60th percentile of the income distribution, $\frac{yf'(y)}{f(y)} > -1$, so that the substitution effect increases the marginal cost of an additional tax cut (assuming a positive elasticity).²¹ Conversely, it crosses zero around the 60th percentile, and is then positive. This means that $\frac{yf'(y)}{f(y)} < -1$ for values of y above the 60th quantile. For those earning more than about \$43K in ordinary income, the substitution effect reduces the cost of providing a tax cut.

Remark. (Shape of Substitution Effect) Suppose $\epsilon^{c}(y) \geq 0$ and $\tau(y) \geq 0$. Then, the substitution effect,

 $^{^{21}}$ This is consistent with the findings of Werning (2007) who estimates the marginal cost of taxation using the SOI public use file.

 $-\epsilon^{c}(y) \frac{\tau(y)}{1-\tau(y)} \alpha(y)$, in equation (4) is positive for incomes below \$43K (60th quantile of 2012 ordinary income) and negative for incomes above \$43K.

The shape of the income distribution suggests that even if tax rates and elasticities were constant across the income distribution, the substitution effect resulting from providing \$1 of surplus through the tax schedule will be lower at the top of the income distribution than at the bottom of the income distribution.

In addition to the substitution effect, it is also possible to put bounds on the natural shape of the impact of the participation effect on the government budget. For those with low incomes, the EITC offers transfers for those who enter the labor force; this renders T(y) < 0 so that those who enter the labor force in response to an increased tax cut actually increase the budgetary cost because they obtain the EITC benefits. In contrast, for higher values of y individuals contribute positive tax revenue so that T(y) > 0; thus any increase in labor force participation for those at higher income levels will result in a positive fiscal externality. These patterns suggest the participation effect is naturally declining in y.

Finally, for income effects most empirical works suggests these effects are either small (Gruber and Saez (2002); Saez et al. (2012)) or declining in income (Cesarini et al. (2015)). Under these assumptions, one has a natural bound on the shape of the efficient welfare weights.

Remark. (Shape of Efficient Welfare Weights) Suppose $\epsilon^c(y) \ge 0$ and $\tau(y) \ge 0$. And, suppose that the participation and income effects are weakly declining in income, y. Then, the efficient welfare weights are higher for those with incomes below \$43K than for those with incomes above \$43K.

In sum, the shape of the income distribution and tax schedules suggest the efficient welfare weights will tend to place greater weight on those with low incomes than those with high incomes.

6 Quantifying Efficient Welfare Weights for the U.S.

6.1 Elasticity Calibration

Moving from bounds to point estimates of g(y) requires estimates of the behavioral responses to taxation. For those subject to the EITC, I draw upon Chetty et al. (2013) who calculate elasticities of 0.31 in the phase-in region (income below \$9,560) and 0.14 in the phase-out region (income between \$22,870 and \$43,210). Using the income tax return data, I assign these elasticities to EITC filers in these regions of the income distribution. Second, for filers subject to the top marginal income tax rate, I assign a compensated elasticity of 0.3. This is consistent with the midpoint of estimates estimated from previous literature studying the behavioral response to changes in the top marginal income tax rate, I assign a compensated elasticity of 0.3, consistent with Chetty (2012) who shows such an estimate can rationalize the large literature on the response to taxation.

In addition to these intensive margin responses, there is also significant evidence of extensive margin behavioral responses, especially for those subject to the EITC. This literature suggests EITC expansions are roughly 9% more costly to the government due to extensive margin behavioral responses.²² Therefore, I assume the participation effect in equation (4) is equal to 0.09 for income groups subject to the EITC. Above the EITC range, there is mixed evidence of participation responses to taxation. Liebman and Saez (2006) find no statistically significant impact of tax changes on women's labor supply of women married to higher-income men. Indeed, higher tax rates can reduce participation from a price effect but increase participation due to an income effect. As a result, I assume a zero participation elasticity for those not subject to the EITC.

Lastly, I assume away intensive margin income effects, consistent with a large literature suggesting such effects are small (Gruber and Saez (2002); Saez et al. (2012)). Recent work by Cesarini et al. (2015)does evidence of income effects using Swedish lotteries; however a large portion of these effects are driven by extensive margin responses and arguably already captured by the EITC responses measured above. However, Appendix E reports the robustness of the results to an alternative specification that incorporates income effects assuming that the estimates from Cesarini et al. (2015) are entirely along the intensive margin and correspond to an elasticity of $\zeta = 0.15$. As shown in Appendix Figure 3, income effects tend to increase the marginal cost of taxation at all income levels; but in contrast to the compensated elasticity they do not affect the relative difference in the weights to low versus high income individuals.

Given the heterogeneity in $\alpha(y)$, equation (4) shows that a key factor determining the shape of FE(y) will be the compensated elasticity, $\epsilon^{c}(y)$. To assess robustness to the specification of this elasticity, I also consider a high elasticity specification of $\epsilon^{c}(y) = 0.5$ and a low compensated elasticity of $\epsilon^{c}(y) = 0.1$ for all y.

6.2 Results

I use equation (4) to combine the estimates of the shape of the income distribution, marginal tax rates, and elasticity calibrations, which generates an estimate of FE(y) for each filer. I then bin the income distribution into 100 quantile bins and construct the mean fiscal externality, FE(y), for each quantile of income. The efficient social welfare weight at each income quantile is then given by g(y) = 1 + FE(y). Figure 6 presents the resulting estimates for g(y). Figure 7 presents the results for the alternative calibrations of the compensated elasticity of $\epsilon^c = 0.1$ and $\epsilon^c = 0.5$.

The weights have several key features. First, consistent with the bounds shown in the previous section, the results suggest it is efficient to place higher weight on surplus to the poor than to the rich. Under the baseline specification, these weights fall from around 1.15 for those at the bottom of the income distribution to 0.65 for those at the top. Transferring \$1 from the top of the distribution can generate around 0.65/1.15 = \$0.57 of welfare to someone at the bottom of the distribution. Conversely, transferring \$1 from the bottom of the income distribution can generate around 1.15/0.65 = \$1.77 of welfare to the those at the top of the income distribution.

Second, although the weights place more weight on low versus high income individuals, the weights never differ by more than a factor of 2. In other words, $\left|\frac{g(y)}{g(y')}\right| < 2$ for all y and y'. This means that it is

 $^{^{22}}$ See Hotz and Scholz (2003) for a summary of elasticities and Hendren (2016) for the 9% calculation.



Figure 6: Efficient Welfare Weights, g(y)

Notes: This figure presents the baseline estimates of the efficient welfare weights, g(y), estimated using equation (4) for each quantile of the income distribution.

not efficient to discount surplus more than 50%, regardless of where it falls in the income distribution. This contrasts with the "consumer surplus standard" often applied to merger policies (e.g. Nocke and Whinston (2014)) that provides no weight to producer surplus on distributional grounds.²³

Third, while the weights generally decline in income, there is an increase in the top 1%. For the baseline specification, it is cheaper to provide additional transfers to the upper middle class than to the top 1%. However, Figure 7 illustrates that this non-monotonicity is not robust to plausible assumptions about how elasticities might change across the income distribution. In particular, if the elasticity moves from 0.3 to 0.5 as one goes from the top 2% to the top 1%, the weights would again be monotonically declining in income.

Fourth, all the weights are positive, g(y) > 0 for all y for the baseline and alternative specifications. This means that it is always costly to provide a tax cut. This implements the test suggested by Werning (2007) for local Laffer effects, and in particular it suggests there are no Pareto improvements solely from modifying the tax schedule.

Lastly, as foreshadowed by the bounding exercise in the previous Section, there is a similarity

²³It is important to remember that the efficient welfare weights depend on the marginal cost of taxation and are thus not invariant to the setting. For example, the weights may differ across countries. Indeed, some have argued that the top tax rate in France is at or above the revenue-maximizing level (Bourguignon and Spadaro (2012)). This would mean that $FE(y) \approx -1$ and $g(y) \approx 0$. If this is the case, it would suggest the efficient welfare weights in France may be close to zero in the upper regions of the income distribution. In France, a consumer surplus standard for merger approval may be optimal.



Figure 7: Robustness to Alternative Elasticities

Notes: This figure presents the baseline specification for the efficient welfare weights alongside with estimates under alternative constant compensated elasticity scenarios of $\epsilon^{c}(y) = 0.1$ and $\epsilon^{c}(y) = 0.5$.

between the estimates of $\alpha(y)$ in Figure 5 and the shape of the efficient welfare weights, g(y). Higher elasticities, $\epsilon^{c}(y)$, increase the difference between the weights on the low- versus high-income individuals. But, they do not affect the general conclusion that g(y) > 1 for those with low incomes and g(y) < 1 for those with high incomes.

7 Applying the Weights to Analysis of Non-Budget Neutral Policy Changes

7.1 Setup

Often one does not want to analyze a feasible alternative environment, but rather the researcher is faced with a policy that increases government expenditure on a policy. In this case, willingnesses to pay for the policy may all be positive, but one needs to account for the fact that the policy involves a cost to the government. This section shows how one can weight the policy's marginal value of public funds, MVPF, by the efficient welfare weights to assess whether it is efficient to increase government expenditure on the policy.

To begin, suppose for simplicity that the policy in question affects those earning near \bar{y} .²⁴ Assume

 $^{^{24}}$ I discuss the generalization to the case when the policy affects individuals at more than one income level in Footnote 26 below.

the researcher has measured (1) individuals' willingnesses to pay for the policy change, \bar{s} , and (2) the net cost of the policy per beneficiary, c. Importantly, c should incorporate any fiscal externalities from the policy change.²⁵

To see how the efficient welfare weights can test whether this policy is efficient, consider the Hicks' experiment in which the government tries to replicate this willingness to pay through modifications to the tax schedule. This would cost $\bar{s}g(\bar{y})$. Therefore, it would be cheaper to replicate this surplus through the tax schedule if and only if

$$\bar{s}g\left(\bar{y}\right) \ge c \tag{5}$$

If equation (5) holds, it is more efficient to provide a tax cut to those earning near \bar{y} than it is to increase spending on the policy.²⁶

Conversely, one can consider the Kaldor experiment that combines the policy with a modification to the tax schedule to search for a Pareto improvement. Taxing back the benefits lowers taxes by $-\bar{s}$ to those earning near \bar{y} . Under Assumption 3 (which assumes the policy does not significantly change the marginal cost of taxation), this generates a cost to the government of $g(\bar{y})(-\bar{s})$. So, the total cost of the combined policy is $-g(\bar{y})\bar{s} + c$. If equation (5) holds, then this total cost is negative. This means there are sufficient resources for the winners to compensate the losers through modifications to the income tax schedule. In this sense, equation (5) characterizes whether increased spending on the policy is efficient.

Equation (5) has an interpretation of comparing the MVPF of the policy to the MVPF of a tax cut with a similar distributional incidence. The MVPF of a policy is defined as the beneficiaries willingness to pay divided by the net cost to the government, $MVPF = \frac{\bar{s}}{c}$. An MVPF of $\frac{\bar{s}}{c}$ means that \$c of net government expenditure on the policy generates \bar{s} units of welfare to the beneficiaries in terms of their willingness to pay. Re-writing equation (5) yields

$$MVPF = \frac{\bar{s}}{c} \ge \frac{1}{g\left(\bar{y}\right)} \tag{6}$$

Equation (6) compares this MVPF to a tax cut to those earning near \bar{y} . Every \$1 of tax cut to those earning near \bar{y} costs the government $g(\bar{y})$ and delivers \$1 of welfare. So, it generates $\frac{1}{g(\bar{y})}$ units of welfare to the beneficiaries. Kaldor-Hicks efficiency involves comparing the MVPF of the policy to the MVPF of a tax cut to the policy's beneficiaries.

$$E\left[g\left(y\right)s\left(y\right)|y\in Y\right] = g\left(\bar{y}\right)\bar{s} + \underbrace{\left(E\left[g\left(y\right)|y\in Y\right] - g\left(\bar{y}\right)\right)\bar{s}}_{\text{Nonlinearity in }g\left(y\right)} + \underbrace{\cot\left(g\left(y\right),s\left(y\right)|y\in Y\right)}_{\text{Cov of WTP with }g\left(y\right)}$$

²⁵For example, if it's a policy that increases spending on schools, it includes increased tax revenue from future earnings increases of the children.

²⁶If a policy affects those at multiple income levels, $y \in Y$, let s(y) denote their willingness to pay. In this case, one should replace $\bar{s}g(\bar{y})$ in equation (5) with $E[g(y)s(y)|y \in Y]$, where the expectation is taken over the set of beneficiaries, $y \in Y$. The bias from using $\bar{s}g(\bar{y})$ instead of $E[g(y)s(y)|y \in Y]$ will introduce bias from nonlinearities in g(y) and from any covariance between g(y) and s(y) amongst the beneficiaries,

These biases are small when the income of the target population is concentrated around a particular \bar{y} or if the efficient welfare weights are relatively constant within the beneficiary population.

Relation to "The" Marginal Cost of Funds Testing equation (6) differs from the common approach that would test whether the beneficiaries' willingness to pay, \bar{s} , exceeds $(1 + \phi)c$, where $1 + \phi$ is "the" marginal cost of public funds. This marginal cost of funds is often assumed to take on a number such as 1.3, and is applied regardless of the distributional incidence of the policy. Equation (6) shows that it is inefficient and misguided to apply the same threshold to a benefit-cost ratio regardless of the incidence of the policy. Intuitively, it's costly for society to redistribute to the poor. So, if one finds a more efficient method, this is efficient from a Kaldor Hicks perspective. Conversely, it is cheap to redistribute additional resources to the rich. Therefore a policy must offer significant benefits per dollar of government expenditure to be more efficient than what can be offered through a tax cut to high-income individuals. In the end, the Kaldor-Hicks logic embodied in equation (6) tests whether the policy is a more efficient method of redistribution than modifications to the tax schedule with similar distributional incidence.²⁷

7.2 Applying to three policy changes

I employ equation (6) to study the efficiency of three policies for which Hendren (2016) computes their MVPF. These policies include Section 8 housing vouchers, food stamps, and job training. Hendren (2016) provides details on the construction of the MVPF for each of these policies. The MVPF for Section 8 vouchers draws upon estimates of the fiscal externality inferred from the causal effects in Jacob and Ludwig (2012) on adult labor supply. The MVPF for food stamps draws upon estimates of the labor supply fiscal externalities from Hoynes and Schanzenbach (2012). And, the MVPF for job training draws upon estimates of the impact of the Job Training Partnership Act (JTPA) on labor earnings and benefit substitution from Bloom et al. (1997).

To illustrate the construction of the MVPF, consider the food stamp application. Hoynes and Schanzenbach (2012) estimate the program cost of the food stamp program was \$1,153 per household. But, reductions in labor supply led to reduced federal and state income taxes of \$588. Hence, the net cost of the program was \$1,741. In contrast, assuming the food stamps are primarily infra-marginal to the budget of the beneficiaries (consistent with Whitmore (2002)), the envelope theorem suggests beneficiaries value food stamps at their program cost of \$1,153. Hence, the MVPF of the food stamp program is given by

$$MVPF^{FoodStamps} = \frac{\bar{s}}{c} = \frac{1,153}{1,153+588} = 0.66 \tag{7}$$

As is noted in Hendren (2016), there is considerable statistical and model uncertainty around each of the point estimates of the MVPF for each of these policies. Indeed, the estimate of the fiscal impact of \$588 from Hoynes and Schanzenbach (2012) is not statistically distinguishable from zero. Moreover,

²⁷Testing equation (6) is also related to the literature that has identified a weak separability assumption on the utility function under which one can compare the unweighted willingness to pay for a policy to its mechanical cost, excluding any fiscal externalities (Hylland and Zeckhauser (1979); Kaplow (1996, 2004, 2008)). Appendix F shows that equation (6) reduces to exactly the tests suggested in these papers if the weak separability assumption holds. This is because, under weak separability, the fiscal externality associated with the policy expenditure is proportional to the fiscal externality from the tax schedule, FE(y). As a result, these two cost impacts cancel out so that equation (6) reduces to testing whether \bar{s} exceeds the mechanical cost of the policy before incorporating any of its fiscal externalities.



Figure 8: Testing the Efficiency of Policy Changes

Notes: This figure illustrates the use of the efficient welfare weights for assessing the efficiency of government policy changes. The line presents the value of $\frac{1}{g(y)}$, which represents the amount of welfare that can be delivered to each portion of the income distribution per dollar of government spending. The dots present estimates of the marginal value of public funds (MVPF) for three policy examples: the job training partnership act (JTPA) from Bloom et al. (1997), food stamps from Hoynes and Schanzenbach (2012), and Section 8 housing vouchers from Jacob and Ludwig (2012). The vertical axis presents the estimated MVPF from Table 1 of Hendren (2016); the horizontal axis presents the estimated income quantiles of the beneficiaries of each policy (normalized to 2012 income using the CPI-U). An MVPF that falls above (below) the the Income/EITC line correspond to policies that can(not) generate Pareto improvements.

measuring willingness to pay for a policy is always difficult. As a result, the exercise in this section should be interpreted as an illustration of how to conduct welfare analysis, as opposed to a definitive conclusion about the efficiency of these particular policies.

Figure 8 plots the $MVPF = \frac{\bar{s}}{c}$ for each of these policies taken directly from Table 1 of Hendren (2016). The horizontal axis corresponds to the quantile corresponding to the mean income, \bar{y} , of the policy beneficiaries.²⁸ The point estimates suggest that food stamps and housing vouchers may be less efficient forms of redistribution than modifications to the income tax schedule. Put differently, the beneficiaries of these policies would prefer the government instead spend the same amount of money on a tax cut (e.g. EITC expansion). In contrast, the estimates suggest the JTPA may be a more efficient policy than a tax cut. As shown in Hendren (2016), spending on the JTPA generated large fiscal externalities through increased taxable income and reductions in other social programs, which effectively lowers the net cost of the policy to the government, c. However, there is considerable uncertainty surrounding beneficiaries' willingnesses to pay for this program, \bar{s} , discussed in Hendren (2016).²⁹

In the end, Figure 8 provides a graphical method for assessing the efficiency of policies by comparing their MVPF to the MVPF of a tax cut with a similar distributional incidence. In particular, it is not efficient to ask whether low-income populations are willing to pay the cost of the policy. Rather, it is efficient to ask whether they prefer additional spending on the policy relative to spending the same amount on a tax cut targeted to the same individuals. Because it is costly to provide transfers to the poor and cheap to provide them to the rich, the results suggest one should use a threshold for the MVPF that is below 1 for low-income beneficiaries and above 1 for high income beneficiaries.³⁰

8 Conclusion

In their original work, Kaldor and Hicks hoped to provide a method to avoid the inherent subjectivity involved in resolving interpersonal comparisons. Weighting surplus using efficient welfare weights measures the economic efficiency of an alternative environment (or policy change) using the Kaldor-Hicks redistributive experiments but accounting for the distortionary cost of taxation. Estimates for the US suggest that redistribution from rich to poor is more costly than from poor to rich. Thus, it is efficient to place greater weight on the poor than on the rich. Regardless of one's own social preferences, surplus to the poor can be turned into greater welfare for everyone than surplus to the rich. The shape of the efficient welfare weights is largely driven by the shape of the income distribution, as opposed to assumptions about behavioral responses to taxation. As a result, the broad conclusion of declining

 $^{^{28}\}mathrm{Incomes}$ are deflated to 2012 dollars using the CPI-U.

²⁹Figure 8 presents the baseline estimates of the MVPF for the JTPA that assumes individuals value the program at its program cost. Hendren (2016) discusses several alternative methods for recovering this willingness to pay; on the one hand, beneficiaries saw large increases in earnings, which suggests they may have substantial willingness to pay for the program above its program cost. On the other hand, enrollees were allowed to enter for free, so that revealed preference would not rule out arbitrarily small willingnesses to pay.

³⁰Moreover, recall that the conclusion about whether g(y) is above or below one does not depend on the exact specification of $\epsilon^{c}(y)$, but rather depends largely on the shape of the income distribution, $\alpha(y)$. In this sense, the conclusion is robust to a wide range of assumptions about the behavioral impact of taxes.

efficient welfare weights is robust to a wide range of assumptions about behavioral elasticities.

There are many important directions for future work. By assuming a \$1 tax cut is valued at \$1 by those targeted by the tax cut, I assumed no spillover effects from tax cuts or impacts of taxes on wages. Future work could incorporate potential general equilibrium effects of taxation. Additionally, the analysis largely assumed that willingness to pay for a policy does not vary conditional on income. If willingness to pay varies conditional on income, modifications to the income tax will generally be insufficient to render Pareto comparisons. In this sense, it might be useful to estimate efficient welfare weights for individuals based on observable characteristics beyond solely their income (e.g. medical spending, etc.). Finally, I provide estimates of efficient welfare weights for the U.S., and I discuss estimates of the marginal cost of taxation in other countries (e.g. France as in Bourguignon and Spadaro (2012)). Going forward, it would be valuable to compile and compare estimates from other countries and settings.

In the end, reasonable people and economists will always disagree about the optimal degree of redistribution in society. But, such debates need not paralyze analysis of economic policies. Hopefully the efficient welfare weights can be a tool to help generate greater consensus about the desirability of economic policies.

References

- Bargain, O., M. Dolls, D. Neumann, A. Peichl, and S. Siegloch (2011). Tax-benefit systems in europe and the us: between equity and efficiency. *CESifo Working Paper No. 3534*. D
- Bloom, H., L. L. Orr, S. H. Bell, G. Cave, F. Doolittle, W. Lin, and J. M. Boss (1997). The benefits and costs of jtpa title ii-a programs: Key findings from the national job training partnership act study. *The Journal of Human Resources* 32(3), 549–576. 7.2
- Blundell, R., M. Brewer, P. Haan, and A. Shephard (2009). Optimal income taxation of lone mothers: An empirical comparison of the uk and germany. *The Economic Journal 119*, 101–121. D
- Bourguignon, F. and A. Spadaro (2012). Tax-benefit revealed social preferences. Journal of Economic Inequality 10, 75–108. 1, 3.3.1, 23, 8, D
- Cesarini, D., E. Lindqvist, M. J. Notowidigdo, and R. Östling (2015). The effect of wealth on individual and household labor supply: evidence from swedish lotteries. Technical report, National Bureau of Economic Research. 5.2, 6.1, E
- Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica* 80(3), 969–1018. 6.1
- Chetty, R., J. Friedman, and E. Saez (2013). Using differences in knowledge across neighborhoods to uncover the impacts of eitc on earnings. *American Economic Review (Forthcoming)*. 6.1

- Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity: The geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics* 129(4), 1553–1623. 5.1
- Coate, S. (2000). An efficiency approach to the evaluation of policy changes. *Economic Jour*nal 110(463), 437–455. 1, 3.1, 3.2
- Currie, J. and F. Gahvari (2008). Transfers in cash and in-kind: Theory meets the data. Journal of Economic Literature 46(2), 333–83. 36
- Diamond, P. A. and E. Saez (2011). The case for a progressive tax: From basic research to policy recommendations. *Journal of Economic Perspectives* 25(4), 165–90. 5.2, B
- Gruber, J. and E. Saez (2002, April). The elasticity of taxable income: evidence and implications. Journal of Public Economics 84(1), 1–32. 5.2, 6.1
- Hendren, N. (2016). The policy elasticity. Tax Policy and the Economy 30. 1, 22, 7.2, 7.2, 29
- Hicks, J. R. (1939). The foundations of welfare economics. Economic Journal 49, 696-712. 1
- Hicks, J. R. (1940). The valuation of social income. *Economica* $\gamma(26)$, 105–124. 1, 2.1, 3.1, 3.2
- Hotz, J. and K. Scholz (2003). The earned income tax credit. In R. A. Moffit (Ed.), *Tax Policy and the Economy*. University of Chicago Press. 22
- Hoynes, H. W. and D. W. Schanzenbach (2012). Work incentives and the food stamp program. *Journal* of Public Economics 96(1-2), 151–162. 1, 7.2, 7.2
- Hylland, A. and R. Zeckhauser (1979). The efficient allocation of individuals to positions. The Journal of Political Economy, 293–314. 27, F
- Jacob, B. A. and J. Ludwig (2012). The effects of housing assistance on labor supply: Evidence from a voucher lottery. *American Economic Review* 102(1), 272–304. 7.2
- Kaldor, N. (1939). Welfare propositions of economics and interpersonal comparisons of utility. The Economic Journal 49(195), 549–552. 1, 1, 2.1, 3.1
- Kaplow, L. (1996). The optimal supply of public goods and the distortionary cost of taxation. National Tax Journal 49, 513–534. 27, F
- Kaplow, L. (2004). On the (ir)relevance of the distribution and labor supply distortion to government policy. Journal of Economic Perspectives 18, 159–175. 27, F
- Kaplow, L. (2008). The Theory of Taxation and Public Economics. Princeton University Press. 27, F
- Kleven, H. J. and C. T. Kreiner (2006). The marginal cost of public funds: Hours of work versus labor force participation. *Journal of Public Economics*, 1955–1973. 16

Liebman, J. and E. Saez (2006). Earnings responses to increases in payroll taxes. Working Paper. 6.1

- Lockwood, B. B. and M. Weinzierl (2014). The evolution of revealed social preferences in the united states and the costs of unequal growth and recessions. *Working Paper*. 4
- Mirrlees, J. A. (1971). An exploration into the theory of optimal income taxation. The Review of Economic Studies 38(2), 175–208. 2
- Nocke, V. and M. Whinston ((2014)). Merger policy with merger choice. *American Economic Review* (Forthcoming). 6.2
- Piketty, T. and E. Saez (2013). Optimal labor income taxation. Handbook of Public Economics 5, 391. 5.2, B
- Saez, E. (2001). Using elasticities to derive optimal income tax rates. Review of Economic Studies 68, 205–229. 2, 17, 33, B
- Saez, E. (2002). Optimal income transfer programs: Intensive versus extensive labor supply responses. Quarterly Journal of Economics 117(3), 1039–1073. 2
- Saez, E., J. Slemrod, and S. H. Giertz (2012). The elasticity of taxable income with respect to marginal tax rates: A critical review. Journal of Economic Literature 50(1), 3–50. 1, 5, 5.1, 5.2, 6.1
- Saez, E. and S. Stantcheva (2016). Generalized social marginal welfare weights for optimal tax theory. American Economic Review 106(1), 24–45. 3, D
- Schlee, E. E. (2013). Radner's cost-benefit analysis in the small: An equivalence result. Working Paper. 2.1
- Werning, I. (2007). Pareto efficient income taxation. MIT Working Paper. 3, 3.3.2, 21, 6.2
- Whitmore, D. (2002). What are food stamps worth? Princeton University Industrial Relations Section Working Paper #468. 7.2
- Zoutman, F. T., B. Jacobs, and E. L. W. Jongen (2013). Optimal redistributive taxes and redistributive preferences in the netherlands. *Mimeo.* 1, 3.3.1, D
- Zoutman, F. T., B. Jacobs, and E. L. W. Jongen (2016). Revealed social preferences of dutch political parties. Working Paper. 1, 3.3.1, 17

Appendix (Not for Publication)

A Proofs

A.1 Proof of Proposition 3

Recall $\mu(\theta)$ is the probability over the type space Θ and F(y) is the cumulative distribution of income in the status quo, $F(x) = \int 1\{y(\theta) \le x\} d\mu(\theta)$. The function $\hat{T}(y; y^*, \epsilon, \eta)$ is given by

$$\hat{T}(y; y^*, \epsilon, \eta) = \begin{cases} T(y) & \text{if } y \notin \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right) \\ T(y) - \eta & \text{if } y \in \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right) \end{cases}$$

so that \hat{T} provides η additional resources to an ϵ -region of individuals earning between $y^* - \epsilon/2$ and $y^* + \epsilon/2$. Fix ϵ and let $\hat{q}(y^*, \epsilon, \eta)$ denote the net government resources expended under $\hat{T}(y; y^*, \epsilon, \eta)$. Given the tax schedule, $\hat{T}(y; y^*, \epsilon, \eta)$, let $\hat{y}(\theta; y^*, \epsilon, \eta)$ denote the individual θ 's choice of earnings, y. The net resources expended is given by:

$$\hat{q}\left(y^{*},\epsilon,\eta\right) = \frac{-1}{F\left(y^{*}+\frac{\epsilon}{2}\right) - F\left(y^{*}-\frac{\epsilon}{2}\right)} \int_{\theta} \hat{T}\left(\hat{y}\left(\theta;y^{*},\epsilon,\eta\right);y^{*},\epsilon,\eta\right) d\mu\left(\theta\right)$$

One needs to evaluate this derivative with respect to η at $\eta = 0$. WLOG, I assume $\hat{q}(y^*, \epsilon, 0) = 0$ so that the status quo tax schedule is budget neutral.

A.1.1 Continuous Responses Only

To begin, I assume there is no participation margin response. Specifically, I assume that preferences are convex in consumption-earnings space so that $\hat{y}(\theta; y^*, \epsilon, \eta)$ is continuously differentiable in η . Below, I add back in extensive margin responses that allow types θ to move to/from 0 and a point of interior earnings, y > 0, in response to a change in the size of the tax cut, η .

A key source of complexity is that individuals may have different curvatures of their utility function. To capture this, define $c(y;\theta)$ to be the individual θ 's indifference curve in consumption-earnings space at the baseline utility level. Given an agent θ 's choice $y(\theta)$ facing the baseline tax schedule T(y), the indifference curve solves

$$u(c(y;\theta), y;\theta) = u(T(y(\theta)) - y(\theta), y(\theta);\theta)$$

Note that the individual's first order condition requires:

$$c'\left(y\left(\theta\right);\theta\right) = -\frac{u_y}{u_c} = 1 - T'\left(y\left(\theta\right)\right) \tag{8}$$

so that the slope of this indifference curve equals the marginal keep rate, 1 - T'.

In addition, the curvature of this indifference curve governs the size of the fraction of people who

change their behavior in order to obtain the transfer, η . Let $k(\theta) = c''(y(\theta); \theta)$ denote the curvature of the indifference curve of type θ in the status quo world. First, consider those whose baseline income is just above $y^* + \frac{\epsilon}{2}$ but the opportunity to obtain the η transfer induces them to drop their income down to $y^* + \frac{\epsilon}{2}$. For individuals with curvature k, a second-order expansion of c (i.e. first order expansion of c') shows that anyone between $y^* + \frac{\epsilon}{2}$ and $y^* + \frac{\epsilon}{2} + \gamma(\eta; k)$ will choose incomes at $y^* + \frac{\epsilon}{2}$, where $\gamma(\eta; k)$ solves

$$\frac{\left(\gamma\left(\eta;k\right)\right)^{2}}{2}k = \eta$$

or

$$\gamma\left(\eta;k\right) = \sqrt{\frac{2\eta}{k}}$$

Similarly, for individuals with curvature k, those with incomes between $y^* - \frac{\epsilon}{2} - \gamma(\eta; k)$ and $y^* - \frac{\epsilon}{2}$ will choose to increase their incomes to $y^* - \frac{\epsilon}{2}$.

Given these definitions, one can write the budget cost as the sum of four terms:

$$\int_{\theta} \hat{T} \left(\hat{y} \left(\theta; y^*, \epsilon, \eta \right); y^*, \epsilon, \eta \right) d\mu \left(\theta \right) = A + B + C + D + o\left(\eta \right)$$

where $\lim_{\eta\to 0} \frac{o(\eta)}{\eta} = 0$ (so that $\frac{do}{d\eta}|_{\eta=0} = 0$, so that one can ignore this term in the calculation of $\frac{d\hat{q}}{d\eta}|_{\eta=0}$).

The first term, A is the mechanical cost that must be paid to all those who receive the η transfer.

$$A = \eta \int 1\left\{ y\left(\theta\right) \in \left(y^* - \frac{\epsilon}{2} - \sqrt{\frac{2\eta}{k\left(\theta\right)}}, y^* + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k\left(\theta\right)}}\right) \right\} d\mu\left(\theta\right)$$

The second term is the cost from those with baseline earnings above $y^* + \frac{\epsilon}{2}$ who drop their income down to $y^* + \frac{\epsilon}{2}$,

$$B = \int \left(T\left(y^* + \frac{\epsilon}{2}\right) - T\left(y\left(\theta\right)\right) \right) 1\left\{ y\left(\theta\right) \in \left(y^* + \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k\left(\theta\right)}}\right) \right\} d\mu\left(\theta\right)$$

And, conversely, the third term is from those with baseline earnings below $y^* - \frac{\epsilon}{2}$ who increase their incomes to $y^* - \frac{\epsilon}{2}$,

$$C = \tau \left(y - \frac{\epsilon}{2} \right) \int \left[y \left(\theta \right) - \left(y^* - \frac{\epsilon}{2} \right) \right] 1 \left\{ y \left(\theta \right) \in \left(y^* - \frac{\epsilon}{2} - \sqrt{\frac{2\eta}{k \left(\theta \right)}}, y^* - \frac{\epsilon}{2} \right) \right\} d\mu \left(\theta \right)$$

and finally the fourth term is the income effect on earnings for those with baseline earnings in the ϵ -region near y^* ,

$$D = \int \left[T\left(\hat{y}\left(\theta; y^*, \epsilon, \eta\right)\right) - T\left(y\right) \right] \mathbf{1} \left\{ y\left(\theta\right) \in \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right) \right\} d\mu\left(\theta\right)$$

The remaining term, $o(\eta)$, captures the bias from approximating the B and C terms using the second-

order expansion for $c(y; \theta)$.

Clearly,

$$\frac{d\left[\int_{\theta} \hat{T}\left(\hat{y}\left(\theta; y^{*}, \epsilon, \eta\right); y^{*}, \epsilon, \eta\right) d\mu\left(\theta\right)\right]}{d\eta}|_{\eta=0} = \frac{dA}{d\eta}|_{\eta=0} + \frac{dB}{d\eta}|_{\eta=0} + \frac{dC}{d\eta}|_{\eta=0} + \frac{dD}{d\eta}|_{\eta=0}$$

I characterize each of these terms. After doing so, one can divide by $F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)$ and take the limit as $\epsilon \to 0$ to arrive at the expression for $\lim_{\epsilon \to 0} \frac{d\hat{q}}{d\eta}|_{\eta=0}$.

Characterizing $\frac{dA}{d\eta}|_{\eta=0}$ First, I show that $\frac{dA}{d\eta}|_{\eta=0} = F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)$.

To see this, first write A by conditioning on $k(\theta)$. Formally, recall that $\mu(\theta)$ is the measure on the type space. Let $\mu_{\theta|k}(\theta|k)$ denote the measure of θ conditional on having curvature k (i.e. $c''(y(\theta)) = k$) and let $\mu_k(k)$ denote the measure of those having curvature k.³¹ Then,

$$A = -\eta \int_{k} \int_{\theta|k} 1\left\{ y\left(\theta\right) \in \left(y^{*} - \frac{\epsilon}{2} - \sqrt{\frac{2\eta}{k}}, y^{*} + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}\right) \right\} d\mu_{\theta|k}\left(\theta|k\left(\theta\right) = k\right) d\mu_{k}\left(k\right)$$

Taking a derivative yields

$$\frac{dA}{d\eta} = -F\left(y^* + \frac{\epsilon}{2} - \sqrt{\frac{2\eta}{k}}\right) + F\left(y^* - \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}\right) - \int_k \eta \left[f_{y|k}\left(y^* + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}|k\right)\sqrt{\frac{1}{2\eta k}} - f_{y|k}\left(y^* - \frac{\epsilon}{2} - \sqrt{\frac{2\eta}{k}}|k\right)\sqrt{\frac{1}{2\eta k}}\right] d\mu_k(k)$$

where $f_{y|k}(y|k)$ is the density of $y(\theta)$ given $k(\theta)$. Note that one can re-write the second term in a manner that makes it clear that it is proportional to $\sqrt{\eta}$:

$$\frac{dA}{d\eta} = -F\left(y^* + \frac{\epsilon}{2} - \sqrt{\frac{2\eta}{k}}\right) + F\left(y^* - \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}\right) - \sqrt{\eta} \left[\int_k \left[f_{y|k}\left(y^* + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}|k\right)\sqrt{\frac{1}{2k}} - f_{y|k}\left(y^* - \frac{\epsilon}{2} - \sqrt{\frac{2\eta}{k}}|k\right)\sqrt{\frac{1}{2k}}\right] d\mu_k\left(k\right)\right]$$

Therefore, evaluating at $\eta = 0$ yields

$$\frac{dA}{d\eta}|_{\eta=0} = -\left[F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)\right]$$

Characterizing $\frac{dB}{d\eta}|_{\eta=0}$ To see this, note that

$$\begin{aligned} \frac{dB}{d\eta} &= \frac{d}{d\eta} \int_{k} \int_{\theta|k} \left(T\left(y^{*} + \frac{\epsilon}{2}\right) - T\left(y\left(\theta\right)\right) \right) \mathbf{1} \left\{ y\left(\theta\right) \in \left(y^{*} + \frac{\epsilon}{2}, y^{*} + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}\right) \right\} d\mu_{\theta|k}\left(\theta|k\left(\theta\right) = k\right) d\mu_{k}\left(k\right) \\ &= \int_{k} \left(T\left(y^{*} + \frac{\epsilon}{2}\right) - T\left(y^{*} + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}\right) \right) \sqrt{\frac{1}{2\eta k}} f_{y|k}\left(y^{*} + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}|k\right) d\mu_{k}\left(k\right) \end{aligned}$$

³¹In other words, for any function of the type space and level of curvature, $r(\theta, k(\theta))$, one has

$$\int \int r(\theta, k(\theta)) d\mu_{\theta|k}(\theta|k(\theta)) d\mu_{k}(k(\theta)) = \int r(\theta, k(\theta)) d\mu(\theta)$$

so that one can either integrate over θ (RHS) or one can first condition on curvature (and integrate over θ given curvature $k(\theta)$) and then integrate over curvature, $k(\theta)$.

which follows from differentiating at the upper endpoint $y^* + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}$ after conditioning on curvature k. Re-writing yields

$$\frac{dB}{d\eta} = \int_{k} \frac{T\left(y^{*} + \frac{\epsilon}{2}\right) - T\left(y^{*} + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}\right)}{\sqrt{\frac{2\eta}{k}}} \frac{1}{k} f_{y|k}\left(y^{*} + \frac{\epsilon}{2} + \sqrt{\frac{2\eta}{k}}|k\right) d\mu_{k}\left(k\right)$$

Now, evaluating as $\eta \to 0$, yields

$$\frac{dB}{d\eta}|_{\eta=0} = -T'\left(y^* + \frac{\epsilon}{2}\right) \int_k \frac{f_{y|k}\left(y^* + \frac{\epsilon}{2}|k\right)}{k} d\mu_k\left(k\right)$$
$$= -T'\left(y^* + \frac{\epsilon}{2}\right) E\left[\frac{1}{k\left(\theta\right)}|y\left(\theta\right) = y^* + \frac{\epsilon}{2}\right] f\left(y^* + \frac{\epsilon}{2}\right)$$

so that tax revenue is decreased by individuals decreasing their income down to $y^* + \frac{\epsilon}{2}$ in order to get the η transfer.

Characterizing $\frac{dC}{d\eta}|_{\eta=0}$ Analogous to the calculation for $\frac{dB}{d\eta}|_{\eta=0}$, it is possible to show that

$$\frac{dC}{d\eta}|_{\eta=0} = T'\left(y^* - \frac{\epsilon}{2}\right) E\left[\frac{1}{k\left(\theta\right)}|y\left(\theta\right) = y^* - \frac{\epsilon}{2}\right] f\left(y^* - \frac{\epsilon}{2}\right)$$

so that tax revenue is increased because individuals move from below $y^* - \frac{\epsilon}{2}$ up to $y^* - \frac{\epsilon}{2}$ in order to get the η transfer.

Characterizing $\frac{dD}{d\eta}|_{\eta=0}$ Finally, I show that

$$\frac{dD}{d\eta}|_{\eta=0} = \left[F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)\right]E\left[\frac{dy}{d\eta}T'\left(y\left(\theta\right)\right)|y\left(\theta\right) \in \left[y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right]\right]$$

so that $\frac{dD}{d\eta}|_{\eta=0}$ is proportional to the average income effects near y^* .

To see this, note that

$$\frac{dD}{d\eta} = \frac{d}{d\eta} \int \left[T\left(\hat{y}\right) - T\left(y\right)\right] \mathbf{1}\left\{y\left(\theta\right) \in \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right)\right\} dF\left(\theta\right)$$

Note that for these individuals in the ϵ region near y^* they only receive an income effect from the policy change. Therefore, we have

$$\frac{dD}{d\eta}|_{\eta=0} = \int T'\left(y\left(\theta\right)\right) \frac{d\hat{y}}{d\eta}|_{\eta=0} 1\left\{y\left(\theta\right) \in \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right)\right\} dF\left(\theta\right)$$

where $\frac{d\hat{y}}{d\eta}|_{\eta=0}$ is the effect of an additional dollar of after-tax income on labor supply. One can define

the income elasticity by multiplying by the after-tax price,

$$\zeta\left(\theta\right) = \left(1 - T'\left(y\right)\right) \frac{d\hat{y}}{d\eta}|_{\eta=0}$$

so that

$$\frac{dD}{d\eta}|_{\eta=0} = \int \frac{T'(y(\theta))}{1-T'(y(\theta))} \zeta(\theta) \, 1\left\{y(\theta) \in \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right)\right\} dF(\theta)$$

Taking $\epsilon \to 0$ Now, to take the limit as $\epsilon \to 0$, note that

$$\frac{dB}{d\eta}|_{\eta=0} + \frac{dC}{d\eta}|_{\eta=0} = -T'\left(y^* + \frac{\epsilon}{2}\right)E\left[\frac{1}{k\left(\theta\right)}|y\left(\theta\right) = y^* + \frac{\epsilon}{2}\right]f\left(y^* + \frac{\epsilon}{2}\right) + T'\left(y^* - \frac{\epsilon}{2}\right)E\left[\frac{1}{k\left(\theta\right)}|y\left(\theta\right) = y^* - \frac{\epsilon}{2}\right]f\left(y^* - \frac{\epsilon}{2}\right)$$

so that

$$\lim_{\epsilon \to 0} \left(\frac{\epsilon}{F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)} \right) \left(\frac{-T'\left(y^* + \frac{\epsilon}{2}\right) E\left[\frac{1}{k(\theta)} | y(\theta) = y^* + \frac{\epsilon}{2}\right] f\left(y^* + \frac{\epsilon}{2}\right) + T'\left(y^* - \frac{\epsilon}{2}\right) E\left[\frac{1}{k(\theta)} | y(\theta) = y^* - \frac{\epsilon}{2}\right] f\left(y^* - \frac{\epsilon}{2}\right)}{\epsilon} \right) = \frac{1}{\epsilon} \left(\frac{\epsilon}{F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)} \right) \left(\frac{-T'\left(y^* + \frac{\epsilon}{2}\right) E\left[\frac{1}{k(\theta)} | y(\theta) = y^* + \frac{\epsilon}{2}\right] f\left(y^* + \frac{\epsilon}{2}\right) + T'\left(y^* - \frac{\epsilon}{2}\right) E\left[\frac{1}{k(\theta)} | y(\theta) = y^* - \frac{\epsilon}{2}\right] f\left(y^* - \frac{\epsilon}{2}\right)}{\epsilon} \right)$$

or

$$\lim_{\epsilon \to 0} \frac{\frac{dB}{d\eta}|_{\eta=0} + \frac{dC}{d\eta}|_{\eta=0}}{F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)} = \frac{1}{f\left(y^*\right)} \left(-\frac{d}{dy}|_{y=y^*} \left[T'\left(y\right)E\left[\frac{1}{k\left(\theta\right)}|_{y}\left(\theta\right) = y\right]f\left(y\right)\right]\right)$$

Now, note also that

$$\lim_{\epsilon \to 0} \frac{-\frac{dA}{d\eta}|_{\eta=0}}{F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)} = 1$$

and

$$\lim_{\epsilon \to 0} \frac{-\frac{dD}{d\eta}|_{\eta=0}}{F\left(y^* + \frac{\epsilon}{2}\right) - F\left(y^* - \frac{\epsilon}{2}\right)} = -T'\left(y^*\right)E\left[\frac{d\hat{y}}{d\eta}|_{\eta=0}|y\left(\theta\right) = y^*\right]$$

which is given by the average income effect at y^* multiplied by the marginal tax rate.

Combining,

$$\lim_{\epsilon \to 0} \frac{d\hat{q}(y^*, \epsilon, \eta)}{d\eta}|_{\eta=0} = 1 + \frac{1}{f(y^*)} \frac{d}{dy}|_{y=y^*} \left[T'(y) E\left[\frac{1}{k(\theta)}|y(\theta) = y\right] f(y) \right] - \frac{T'(y)}{1 - T'(y)} E\left[\zeta(\theta)|y(\theta) = y\right] d\theta$$

Replacing curvature with compensated elasticity Now, note that the curvature, k, is related to the compensated elasticity of earnings. To see this, note that

$$c'(y(\theta);\theta) = 1 - \tau$$

where τ is the marginal tax rate faced by the individual, $\tau = T'(y(\theta))$. Totally differentiating with respect to one minus the marginal tax rate yields

$$c''\left(y\left(\theta\right)\right)\frac{dy^{c}}{d\left(1-\tau\right)} = 1$$

where $\frac{dy^c}{d(1-\tau)}$ is the compensated response to an increase in the marginal keep rate, $1-\tau$. Re-writing,

$$\frac{dy^{c}}{d\left(1-\tau\right)} = \frac{1}{c''\left(y\left(\theta\right)\right)}$$

Intuitively, the size of a compensated response to a price change is equal to the inverse of the curvature of the indifference curve.

Now, by definition, the compensated elasticity of earnings is given by

$$\epsilon^{c}\left(\theta\right) = \frac{dy^{c}}{d\left(1-\tau\right)}\frac{\left(1-\tau\right)}{y\left(\theta\right)} = \frac{1}{c''\left(y\right)}\frac{1-\tau}{y}$$

or

$$\frac{1}{k\left(\theta\right)} = \epsilon^{c}\left(\theta\right) \frac{y\left(\theta\right)}{1 - T'\left(y\left(\theta\right)\right)}$$

where $\epsilon^{c}(\theta)$ is the compensated elasticity of type θ defined locally around the status quo tax schedule.

Replacing $\frac{1}{k(\theta)}$ in the main equation yields

$$\lim_{\epsilon \to 0} \frac{d\hat{q}\left(y^{*}, \epsilon, \eta\right)}{d\eta}|_{\eta=0} = 1 + \frac{1}{f\left(y^{*}\right)} \frac{d}{dy}|_{y=y^{*}} \left[T'\left(y\right) E\left[\epsilon^{c}\left(\theta\right) \frac{y\left(\theta\right)}{1 - T'\left(y\left(\theta\right)\right)}|y\left(\theta\right) = y\right] f\left(y\right)\right] - \frac{T'\left(y\right)}{1 - T'\left(y\right)} E\left[\zeta\left(\theta\right)|y\left(\theta\right) = y\right] f\left(y\right) = \frac{T'\left(y\right)}{1 - T'\left(y\right)} E\left[\zeta\left(\theta\right)|y\left(\theta\right) = \frac{T'\left(y$$

A.1.2 Adding a Participation Margin

Heretofore, I have ignored the potential for extensive margin responses. Put differently, I assumed everyone's intensive margin first order condition (equation (8)) held. Now, I show how one can overlay participation margin responses for people who move in and out of the labor force in response to changes in the tax schedule.

For simplicity, consider an alternative world where y = 0 was removed from individuals' feasibility set. Let $y^{P}(\theta)$ denote the earnings choice of type θ in this restricted world. Clearly, $y^{P}(\theta)$ solves

$$y^{P}(\theta) = argmax_{y>0} u (y - T(y), y; \theta)$$

For all types in the labor force in the status quo world, $y^P(\theta) = y(\theta)$. For those out of the labor force, $y(\theta) = 0$. I retain the assumption that preferences are convex over the region y > 0. Therefore, $y^P(\theta)$ is continuously differentiable in response to changes in the tax schedule, T. So, I allow for discrete moves between 0 and y > 0, but do not allow discrete moves across two different labor supply points in response to small changes in the tax schedule.

Given $y^{P}(\theta)$, let $c^{P}(\theta)$ denote the consumption level required by type θ to enter into the labor force to earn $y^{P}(\theta)$:

$$u\left(c^{P}\left(\theta\right), y^{P}\left(\theta\right); \theta\right) = u\left(y - T\left(0\right), 0; \theta\right)$$

Given $y^{P}(\theta)$ and $c^{P}(\theta)$, one can define the labor force participation rate at each point along the income distribution. Note that an individual of type θ chooses to work whenever

$$c^{P}\left(\theta\right) \leq y^{P}\left(\theta\right) - T\left(y^{P}\left(\theta\right)\right)$$

For any consumption and income level, (c, y), let LFP(c, y) denote the fraction of individuals with $y^{P}(\theta) = y$ who choose to work, $y(\theta) = y$:

$$LFP(c, y) = \int 1\left\{c \ge c^{P}(\theta)\right\} d\mu\left(\theta|y^{P}(\theta) = y\right)$$

With this definition, one can write

$$\hat{q}\left(y^{*},\epsilon,\eta\right) = A + B + C + D + P + o\left(\eta\right)$$

where P is the cost resulting from non-marginal changes in labor supply and $\frac{dP}{d\eta}|_{\eta=0,\epsilon=0}$ is given by

$$\begin{aligned} \frac{dP}{d\eta}|_{\eta=0} &= \frac{d}{d\eta}|_{\eta=0} \int_{y^{P}(\theta)\in\left[y^{*}-\frac{\epsilon}{2},y^{*}+\frac{\epsilon}{2}\right]} \left[\left(T\left(y^{P}\left(\theta\right)\right)-\eta-T\left(0\right)\right) LFP\left\{y^{P}-T\left(y^{P}\right)+\eta,y^{P}\left(\theta\right)\right\} dF\left(\theta\right) \right] \\ &= \int_{y(\theta)\in\left[y^{*}-\frac{\epsilon}{2},y^{*}+\frac{\epsilon}{2}\right]} \left[\left(T\left(y\left(\theta\right)\right)-T\left(0\right)\right) \frac{dLFP\left\{y\left(\theta\right)-T\left(y\left(\theta\right)\right),y\left(\theta\right)\right\}}{dc} dF\left(\theta\right) \right] \end{aligned}$$

so that

$$\begin{split} \lim_{\epsilon \to 0} \frac{1}{\epsilon} \frac{dP}{d\eta}|_{\eta=0} &= E\left[\left(T\left(y \right) - T\left(0 \right) \right) \frac{dLFP\left(y \right)}{dc} | y^P\left(\theta \right) = y \right] LFP\left(y \right) \\ &= \frac{T\left(y \right) - T\left(0 \right)}{y - T\left(y \right)} \hat{\epsilon} \left(y^P \right) \\ &= \frac{T\left(y \right) - T\left(0 \right)}{y - T\left(y \right)} \epsilon_c^{LFP}\left(y \right) \end{split}$$

where $\hat{\epsilon}_{c}^{LFP}(y)$ is the semi-elasticity of labor force participation at y off of the base of all potential people who have $y^{P}(\theta)$ as their most preferred earnings point. To align with A-D, we need to replace the distribution of y^{P} with the distribution of y, so that we must divide by LFP. Dividing by LFP(y), this is equal to the elasticity of labor force participation at $y^{P}(\theta)$

$$\epsilon_{c}^{LFP}\left(y\right) = \frac{1}{LFP\left(y - T\left(y\right), y\right)} \frac{\partial LFP\left(y - T\left(y\right), y\right)}{\partial c}$$

Therefore, we have

$$\lim_{\epsilon \to 0} \frac{d\hat{q}(y^*, \epsilon, \eta)}{d\eta}|_{\eta=0} = 1 + \frac{1}{f(y^*)} \frac{d}{dy}|_{y=y^*} \left[\frac{T'(y)}{1 - T'(y)} \epsilon^c(y) y f(y) \right] - \frac{T'(y)}{1 - T'(y)} \zeta(y) + \frac{T(y) - T(0)}{y - T(y)} \epsilon^{LFP}(y) q^{LFP}(y) q^$$

where

$$\epsilon^{c}(y) = E\left[\epsilon^{c}(\theta) | y(\theta) = y\right]$$

and

$$\zeta(y) = E[\zeta(\theta) | y(\theta) = y]$$

A.2 Proof of Proposition 1

Statement of Proposition For any $\epsilon > 0$ define the scaled surplus by $s_{\epsilon}(y) = \epsilon s(y)$ and $S_{\epsilon} = E[s_{\epsilon}(y)g(y)] = \epsilon S$. If S < 0, there exists an $\tilde{\epsilon} > 0$ such that for any $\epsilon < \tilde{\epsilon}$ there exists an augmentation to the tax schedule in the status quo environment that generates surplus, $s_{\epsilon}^{t}(y)$, that is uniformly greater than the surplus offered by the alternative environment, $s_{\epsilon}^{t}(y) > s_{\epsilon}(y)$ for all y. Conversely, if S > 0, no such $\tilde{\epsilon}$ exists.

Proof Suppose S < 0. Then,

$$\int s(y) g(y(\theta)) d\mu(\theta) < 0$$

so that

$$\int s_{\epsilon}(y) g(y(\theta)) d\mu(\theta) = \epsilon \int s(y) g(y(\theta)) d\mu(\theta) < 0$$

For any tax schedule \hat{T} , let $y(\theta; \hat{T})$ denote the choice of earnings by type θ facing tax schedule \hat{T} . Given these choices, total tax revenue is given by

$$R\left(\hat{T}\right) = \int \hat{T}\left(y\left(\theta;\hat{T}\right)\right) d\mu\left(\theta\right)$$

Now, consider an augmented tax schedule. Let $P = \{P_j\}_{j=1}^{N_P}$ denote a partition of the income distribution into intervals and let η_j^P denote transfers provided to each such region of the income distribution:

$$\hat{T}_{\epsilon}^{P}\left(y\right) = T\left(y\right) - \epsilon \sum_{j=1}^{N_{P}} \eta_{j}^{P} \mathbb{1}\left\{y \in P_{j}\right\}$$

and let

$$\hat{T}_{\epsilon}^{j}(y) = T(y) - \epsilon \eta_{j}^{P} \mathbb{1}\left\{y \in P_{j}\right\}$$

I assume that

$$\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}^{P}_{\epsilon}\right) = \sum_{j=1}^{N_{P}} \frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}^{j}_{\epsilon}\right)$$
(10)

This assumption is satisfied for most common forms of preferences. It would be violated if, for example, there were a mass of agents perfectly indifferent between three earnings points. Then, providing additional transfers to one of these two points would both induce movement from the other point and thus the sum of the two tax movements would be larger than the combined tax movement. In practice, it is satisfied if agents have at most two points of indifference in their utility function (since then their movement is affected by the relative price changes to be in these two locations).

Note that each partition can be represented as

$$P_j = [y_j^* - \epsilon_j, y_j^* + \epsilon_j)$$

so that

$$\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}^{j}\right) = -\eta_{j}^{P}\frac{d}{d\eta}|_{\eta=0}\hat{q}\left(y_{j}^{*},\epsilon_{j},\eta\right)\Pr\left\{y\left(\theta\right)\in P_{j}\right\}$$

where $\Pr\left\{y\left(\theta\right)\in P_{j}\right\}=\mu\left\{y^{-1}\left(P_{j}\right)\right\}.$

Now, define η_j^P as

$$\eta_{j}^{P} = \sup\left\{s\left(y\right)|y \in P_{j}\right\} - \frac{S^{ID}}{2}\bar{g}$$

where $\bar{g} = E[g(y)]$ is the average value of the marginal cost of taxation. Let $s_{\epsilon}^{t}(y)$ denote the surplus the individual earning y obtains when facing tax schedule \hat{T}_{ϵ}^{j} . By the envelope theorem (and the assumption of no externalities / GE effects), there exists $\tilde{\epsilon}$ such that for all $\epsilon < \tilde{\epsilon}$, the individual obtains surplus at least as large as $\epsilon (\sup \{y | y \in P_j\})$

$$s_{\epsilon}^{t}(y) > \epsilon \left(\sup \left\{ y | y \in P_{j} \right\} \right)$$
(11)

for all $\epsilon \in (0, \tilde{\epsilon})$ (to see this, note that the tax augmentation not only gives people surplus s^y but also provides $-\frac{S}{2} > 0$; so this inequality is made strict).

Now, consider the marginal cost of the policy. By construction

$$\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}^{P}\right) = -\sum_{j=1}^{N_{P}}\eta_{j}^{P}\frac{d}{d\eta}|_{\eta=0}\hat{q}\left(y_{j}^{*},\epsilon_{j},\eta\right)\Pr\left\{y\left(\theta\right)\in P_{j}\right\}$$

and taking the limit as partition widths go to zero,

$$\lim_{width(P)\to 0} \frac{d}{d\epsilon}|_{\epsilon=0} R\left(\hat{T}_{\epsilon}^{P}\right) = -\lim_{width(P)\to 0} \sum_{j=1}^{N_{P}} \left(\left(\sup\left\{ y|y\in P_{j}\right\} \right) - \frac{S}{2} \right) \frac{d}{d\eta}|_{\eta=0} \hat{q}\left(y_{j}^{*},\epsilon_{j},\eta\right) \Pr\left\{ y\left(\theta\right)\in P_{j}\right\} \right) + \sum_{j=1}^{N_{P}} \left(\sum_{j=1}^{N$$

Note that the terms inside the sum have limits that exist and are unique (because g(y) is assumed to be continuous and the mean surplus function is assumed to be continuous). Note in principle this limit existing does not require continuity of either the surplus function or the marginal cost function g(y) – some suitable integrability condition would work – but this is sufficient. So,

$$\begin{split} \lim_{width(P)\to 0} \frac{d}{d\epsilon}|_{\epsilon=0} R\left(\hat{T}^P_{\epsilon}\right) &= -\int s\left(y\left(\theta\right)\right) g\left(y\left(\theta\right)\right) d\mu\left(\theta\right) + \frac{S}{2}\bar{g} \\ &= -S\bar{g} + \frac{S}{2}\bar{g} \\ &= -\frac{S}{2}\bar{g} \end{split}$$

which is positive. Therefore, there exists $\epsilon^* < \tilde{\epsilon}$ such that for all $\epsilon \in (0, \epsilon^*)$ we have $R\left(\hat{T}^P_{\epsilon}\right) > 0$ and $s^t_{\epsilon}(y) > s(y)$ for all y.

Converse Now suppose S > 0. Then,

$$\int s(y(\theta)) g(y(\theta)) d\mu(\theta) > 0$$

And, suppose for contradiction that some $\tilde{\epsilon}$ exists so that there are a set of tax schedules, \hat{T}_{ϵ} , that deliver greater surplus along the income distribution, $\frac{d}{d\epsilon}|_{\epsilon=0}s^t_{\epsilon}(y) \ge s(y)$. I will show that this implies the tax schedule modification is not budget neutral for sufficiently small ϵ .

Note that the envelope theorem implies $\frac{d}{d\epsilon}|_{\epsilon=0}\hat{T}_{\epsilon}(y) = \frac{d}{d\epsilon}|_{\epsilon=0}s^{t}_{\epsilon}(y)$ for all y. For any $\epsilon > 0$ and $\gamma > 0$, one can approximate the revenue function using a partition $P^{\gamma} = \left\{P_{j}^{\gamma}\right\}_{j=1}^{N_{P^{\gamma}}}$ and a step tax function $T_{\epsilon}^{P^{\gamma}}$ that provides exactly $E\left[s_{\epsilon}^{t}(y(\theta)) | y(\theta) \in P_{j}^{\gamma}\right]$ units of tax reduction. Therefore, the marginal cost of the policy is approximated by $\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}^{P^{\gamma}}\right)$,

$$\left|\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}\right) - \frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}^{P^{\gamma}}\right)\right| < \gamma$$

where

$$\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}^{P^{\gamma}}\right) = -\sum_{j=1}^{N_{P^{\gamma}}}\frac{d}{d\epsilon}|_{\epsilon=0}E\left[s_{\epsilon}^{t}\left(y\left(\theta\right)\right)|y\left(\theta\right)\in P_{j}^{\gamma}\right]\frac{d}{d\eta}|_{\eta=0}\hat{q}\left(y_{j}^{*},\epsilon_{j},\eta\right)\Pr\left\{y\left(\theta\right)\in P_{j}\right\}\right\}$$

where $P_j^{\gamma} = [y_{\gamma,j}^* - \epsilon_{\gamma,j}, y_{\gamma,j}^* + \epsilon_{\gamma,j})$. For sufficiently small ϵ we know that $E\left[s_{\epsilon}^t\left(y\left(\theta\right)\right) | y\left(\theta\right) \in P_j^{\gamma}\right] > E\left[s_{\epsilon}\left(y\left(\theta\right)\right) | y\left(\theta\right) \in P_j^{\gamma}\right]$. Therefore,

$$\frac{d}{d\epsilon}|_{\epsilon=0}E\left[s_{\epsilon}^{t}\left(y\left(\theta\right)\right)|y\left(\theta\right)\in P_{j}^{\gamma}\right] > \frac{d}{d\epsilon}|_{\epsilon=0}E\left[s_{\epsilon}\left(y\left(\theta\right)\right)|y\left(\theta\right)\in P_{j}^{\gamma}\right] - \gamma$$

since both have values of zero when $\epsilon = 0$.

Therefore,

$$\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}^{P}\right) < -\sum_{j=1}^{N}\frac{d}{d\epsilon}|_{\epsilon=0}E\left[s_{\epsilon}\left(y\left(\theta\right)\right)|y\left(\theta\right)\in P_{j}^{\gamma}\right]\frac{d}{d\eta}|_{\eta=0}\hat{q}\left(y_{j}^{*},\epsilon_{j},\eta\right)\Pr\left\{y\left(\theta\right)\in P_{j}^{\gamma}\right\}+\gamma$$

and taking the limit as the partition widths converge towards zero (so that $\gamma \to 0$), we arrive as

$$\frac{d}{d\epsilon}|_{\epsilon=0}R\left(\hat{T}_{\epsilon}\right) \leq -SE\left[g\left(\theta\right)\right] < 0$$

so that the policy is not budget neutral.

Discussion The proof relied on two key assumptions. First, I assume that providing a small amount of money through modifications in the tax schedule generates surplus of at least the mechanical amount of money provided in the absence of any behavioral response. This follows from the envelope theorem,

combined with the assumption that infinitesimal tax changes in one portion of the income distribution do not affect the welfare of anyone at other points of the distribution. This was implicitly assumed by writing the utility function as a function of one's own consumption and earnings, and not a function of anyone else's choices of labor supply or earnings. For example, if taxing the rich caused them to reduce their earnings which in turn increased the wages of the poor, then equation (11) would no longer hold, since individuals outside of the intended target of the tax transfers would have surplus impacts. Accounting for such general equilibrium effects is an interesting and important direction for both theoretical and empirical work.

Second, I assume that the revenue function is continuously differentiable and additive in modifications to the tax schedule. This is primarily a technical assumption that rules out types that are indifferent to many points along the income distribution (which would cause them to be double-counted as costs in equation (10)).

A.3 Proof of Proposition 2

Statement of Proposition Suppose Assumption 2 holds. For $\epsilon > 0$, let $s_{\epsilon} = \epsilon s(y)$. If S > 0, there exists $\tilde{\epsilon} > 0$ such that for any $\epsilon < \tilde{\epsilon}$, there exists an augmentation to the tax schedule in the alternative environment that delivers surplus $s_{\epsilon}^{t}(y)$ that is positive at all points along the income distribution, $s_{\epsilon}^{t}(y) > 0$ for all y. Conversely, if S < 0, then no such $\tilde{\epsilon}$ exists.

Proof I provide the brief sketch here that does not go through the formality of defining the partitions as in the proof above, but one can do so analogously to the proof of Proposition 1. Let $y(\theta)$ continue to denote the choice of income of a type θ in the status quo environment, which may differ from their choice of y in the alternative environment. To capture this, let $y_{\epsilon}^{\alpha}(y)$ denote the choice of income in the alternative environment made by those who chose y in the status quo environment. Per Assumption 2, this function is a bijection. Given the surplus function, s(y), consider a modification to the income distribution that taxes away all but $\epsilon \frac{S}{2}$ of this surplus to those earning y in the status quo (i.e. those earning $y_{\epsilon}^{\alpha}(y)$ in the ϵ -alternative environment). If \tilde{T}_{ϵ} is the tax schedule in the ϵ -alternative environment, then the modified tax schedule is

$$\hat{T}_{\epsilon}\left(y\right) = \tilde{T}_{\epsilon}\left(y\right) + \epsilon \left(s\left((y_{\epsilon}^{\alpha})^{-1}\left(y\right)\right) - \frac{S}{2}\right)$$

Let $s_{\epsilon}^{t}(y)$ denote the surplus of the tax-modified ϵ -alternative environment with tax schedule $\hat{T}_{\epsilon}(y)$. For sufficiently small ϵ , the off-setting transfer ensures everyone is better off relative to the status quo (note this relies on the fact that S > 0, so that there is aggregate surplus to spread around). Hence, $s_{\epsilon}^{t}(y) > s_{\epsilon}(y) - E[s_{\epsilon}(y)]$ for sufficiently small ϵ ; and taking the expectation conditional on $y(\theta) = y$ yields

$$E\left[s_{\epsilon}^{t}\left(y\right)\right] > 0 \quad \forall y$$

Now, one needs to show that, for sufficiently small ϵ , the cost of the modification to the tax

schedule is not budget-negative. Note that for each y, the tax modification provides a transfer of $s_{\epsilon}\left((y_{\epsilon}^{\alpha})^{-1}(y)\right)$. Note that Assumption 2, the marginal cost of implementing these surplus transfers is the same as in the status quo environment.

$$\frac{dR}{d\epsilon}|_{\epsilon=0} = \int s\left(y\right)g\left(y\right)dF\left(y\right) - \frac{S}{2} = \frac{S}{2} > 0$$

so that the transfer scheme is feasible for sufficiently small $\epsilon.$

Marginal Federal Tax Rate	Number of Filers	Percent of Filers	Mean Ordinary Income	Mean Family Income
(1)	(2)	(3)	(4)	(5)
-35.0%	589.750	0.6%	79	11.622
-30.0%	1 413 594	1 4%	229	11 598
-25.0%	6.604	0.0%	16.004	34,459
-24.0%	1.366.424	1.4%	346	9,435
-19.0%	10.905	0.0%	15.088	30.532
2.4%	2.738.247	2.7%	170	6.320
7.4%	18,321	0.0%	19,049	35,938
10.0%	19,355,436	19.3%	(7,140)	18,754
15.0%	35,974,064	35.9%	32,510	52,689
17.4%	620	0.0%	58,739	65,642
17.7%	2,736,669	2.7%	1,120	11,585
22.7%	927	0.0%	12,936	16,493
25.0%	19,217,503	19.2%	74,974	97,153
26.0%	3,035,148	3.0%	6,416	25,402
28.0%	6,266,531	6.3%	193,243	234,677
31.0%	1,078,152	1.1%	15,822	32,212
31.1%	4,868,642	4.9%	4,622	28,992
33.0%	225,353	0.2%	243,716	234,349
35.0%	400,306	0.4%	1,977,424	1,154,254
36.1%	795,541	0.8%	16,523	38,355
Other	549	0.0%	15,934,759	54,534
Total	100,099,286	100.0%	46,110	64,745

Appendix Table I Summary Statistics

Notes. This table presents summary statistics for the universe of income tax returns in 2012 for U.S. citizens aged 25-60. This sample is used to construct the elasticity of the density of the income distribution (i.e. alpha) for each marginal tax rate. The table presents the number of filers, mean ordinary income, and mean family income by each federal marginal tax rate. The mean federal tax rate is the effective marginal tax rate each filer would face on an additional dollar of income. This equals the tax on ordinary income for most filers, but includes additional tax rates generated by the earned income tax credit (EITC) for EITC filers and the alternative minimum tax (AMT) for filers subject to the AMT.

B Summary Statistics and Estimation Details

Appendix Table I presents the summary statistics of the sample used to construct the estimates of the shape of the income distribution conditional on the marginal income tax rate. Overall, there are roughly 100M filers aged 25-60 used in the analysis, with mean family incomes of roughly \$65M, and mean ordinary incomes of \$46M. The negative tax rates are generated by the various phase-in regions of the EITC for filers of different marital status and combinations of children and EITC-eligible children.³²

To estimate the Pareto parameter of the income distribution, I proceed as follows. First, for computational simplicity, I define 1000 equally sized bins of ordinary income. I then collapse the data to generate counts of returns in each of these 1000 bins separately for returns facing different tax schedules, j. I generate these groups as the intersection of filing status, EITC status (marital status + number of qualified EITC dependents), and those subject to the alternative minimum tax rate.

Given these groupings, I estimate the shape of the income distribution, α , in a manner that allows it to vary with the marginal tax rate for a majority of the population. Let j index the set of tax schedules. For tax schedules with at least at least 500,000 observations with earnings between the 10th and 99th percentile of the income distribution, I estimate the elasticity of the income distribution separately for each filing characteristic, which I denote α_j (y). ³³ To do so, I construct the log density of the income distribution measuring the number of households in each bin divided by the width of the bin. I then regress this on a fifth order polynomial of log income in the bin (where income is the mean income within the bin). The estimated slope at each bin generates an estimate of α_j for each income bin in tax group j. I verify that the results are virtually identical when increasing or decreasing the number of bins or changing the number of polynomials in the regression.

For the remaining smaller tax groups (~25% of the sample) with fewer than 500,000 returns, I impose the assumption that the elasticity of the income distribution is the same across these less-populated tax schedules at a given level of income.³⁴ I then take advantage of the fact that the aggregate elasticity can be written as a weighted average of the elasticities of the income distribution for each marginal tax rate, α_j . So, I estimate the elasticity of the aggregate income distribution and then construct the implied elasticity for these smaller groups as the population weighted difference between the total elasticity and the elasticities of the larger tax groups. To estimate the elasticity of the aggregate income distribution, I regress the log density on a tenth order polynomial in log income for each bin (again, results are nearly identical if one includes additional polynomials) and compute the slope at each bin.

The advantage of this estimation approach is that it allows the elasticity of the income distribution

³²As is well-known, EITC dependents is technically different from the number of dependents one can claim for baseline 1040 deductions. Fortunately, the tax data contains information on both the number of EITC-eligible dependents and the number of 1040 dependents, allowing me to precisely identify the marginal tax rate.

³³I do not include the returns below the 10th quantile of the income distribution because of the large fraction of returns posting exactly \$3k in ordinary income, which introduces significant nonlinearities in some of these groups. Above the 99th percentile, I follow a strategy from Saez (2001) described below.

 $^{^{34}}$ This 500,000 threshold is chosen for computational simplicity on the remaining groups, but the results are similar to lowering it to 250,000.



Appendix Figure 1: Estimation of Shape of Income Distribution

Notes: Panel A of this figure presents estimates of the average value of $\alpha(y)$ by income quantile. Panel B presents estimates using an alternative method of estimating $\frac{E[Y|Y \ge y] - y}{y}$ in each quantile.





Notes: This figure presents the average tax rate by income quantile. Each tax rate is the sum of the federal income tax rate, state taxes, Medicare, sales taxes, and EITC top-up, as discussed in Section 5.1.

to vary non-parametrically with the tax rates for ~75% of the sample. Hence, I allow for significant correlation between the shape of the income distribution and the marginal tax rate, as is potentially required for accurate estimation of the substitution effect in the presence of multiple tax schedules.

For individuals near the top of the income distribution, the local calculation of the elasticity of the income distribution becomes difficult and potentially biased because of endpoint effects. Intuitively, the binning of incomes into 1,000 bins ignores the fact that the U.S. income distribution has a fairly thick upper tail. Fortunately, it is well documented that the upper tail of the income distribution is Pareto, and hence has a constant elasticity so that $\alpha(y) = \frac{E[Y|Y \ge y] - y}{y}$ (Saez (2001)). Hence, I also compute an "upper tail" value of α given by $\frac{E[Y|Y \ge y] - y}{y}$ for each income bin. Appendix Figure 1 (left panel) plots the average local estimate of α (using the fifth order polynomial) across the income distribution and Appendix Figure 1 (right panel) plots both this estimate and the upper tail value of α , $\frac{E[Y|Y \ge y] - y}{y}$, for the upper decile of the income distribution.

For the upper regions of the income distribution, the value of $\frac{E[Y|Y \ge y] - y}{y}$ converges to around 1.5, consistent with the findings of Diamond and Saez (2011) and Piketty and Saez (2013). Conversely, the local estimate of the elasticity of the income distribution arguably becomes downwardly biased in the upper region because the fifth order polynomial does not capture the size of the thick tail in the top-most income bucket. Hence, for incomes in this upper region with earnings above \$250,000, I assign the maximum value of these two estimates.

C Heterogeneity

If two people earning the same income, $y(\theta)$, have different surplus, $s(\theta)$, then undoing the distributional incidence through the tax schedule will necessarily make one of the two people strictly better off.³⁵ Fortunately, with a slight modification of the surplus function, one can use the efficient welfare weights to characterize the existence of local Pareto improvements.

Given the surplus function $s(\theta)$ of interest, I define the min and max surplus at each point of the income distribution. First, for any \hat{y} let $\underline{s}(\hat{y}) = \inf \{s(\theta) | y(\theta) = \hat{y}\}$ be the smallest surplus obtained by a type θ that earns \hat{y} (note this number may be negative). Second, let $\overline{s}(\hat{y}) = \sup \{s(\theta) | y(\theta) = \hat{y}\}$ be the largest surplus obtained by a type θ that earns \hat{y} . The search for local Pareto improvements involves weighting not actual surplus, $s(\theta)$, but rather these min and max surplus functions conditional on income. In particular, let

$$\underline{S} = \int \underline{s}(y) g(y(\theta)) d\mu(\theta)$$

and

$$\overline{S}=\int\overline{s}\left(y
ight)g\left(y\left(heta
ight)
ight)d\mu\left(heta
ight)$$

If $\overline{S} < 0$, then there exists a modification to the existing tax schedule such that everyone locally prefers the modified status quo to the alternative environment.

Proposition 4. Suppose $\overline{S} < 0$. Then, there exists an $\tilde{\epsilon} > 0$ such that, for each $\epsilon < \tilde{\epsilon}$ there exists a modification to the income tax schedule that delivers a Pareto improvement relative to $s_{\epsilon}(\theta)$. Conversely, if $\overline{S} > 0$, there exists an $\tilde{\epsilon} > 0$ such that for each $\epsilon < \tilde{\epsilon}$ any budget-neutral modification to the tax schedule results in lower surplus for some θ relative to $s_{\epsilon}(\theta)$.

Proof. The proof follows immediately by providing surplus $\overline{s}_{\epsilon}(y) = \sup \{s_{\epsilon}(\theta) | y(\theta) = y\}$ instead of $E[s_{\epsilon}(\theta) | y(\theta) = y]$ in the proof of Proposition 1.

When $\overline{S} < 0$, a change in the tax schedule within the status quo locally Pareto dominates the alternative environment. Clearly, $\overline{S} \ge S$ so that this is a more restrictive test of whether the status quo should be preferred to the alternative environment.

Conversely, using Assumption 2, one can test whether the alternative environment, modified with a change to the tax schedule, provides a local Pareto improvement relative to the status quo.

Proposition 5. Suppose Assumption 2 holds. Suppose $\underline{S} > 0$. Then, there exists an $\tilde{\epsilon} > 0$ such that, for each $\epsilon < \tilde{\epsilon}$ there exists a modification to the income tax schedule in the alternative environment such that the modified alternative environment delivers positive surplus to all types relative to the status quo, $s_{\epsilon}^{t}(\theta) > 0$ for all θ .

³⁵For another example, suppose an alternative environment offers a surplus of \$20 to one person earning \$40K and a surplus of -\$10 to another person also earning \$40K. Then efficient surplus would ask whether one can modify the tax schedule to provide \$5 of surplus to those earning \$40K. Of course, this \$5 would not sufficiently compensate the individual with -\$10 in surplus and hence testing whether efficient surplus is positive would not correspond to a potential Pareto improvement.

Proof. The proof follows immediately by providing surplus $\underline{s}_{\epsilon}(y) = \inf \{s_{\epsilon}(\theta) | y(\theta) = y\}$ instead of $s_{\epsilon}(y)$ in the proof of Proposition 2.

In general, it can be the case that $\overline{S} > 0 > \underline{S}$, so that the potential Pareto criterion cannot lead to a sharp comparison between the status quo and the alternative environment.

Corollary 1. Suppose $s(\theta)$ does not vary with θ conditional on income, $y(\theta)$ (i.e. $s(\theta) = \tilde{s}(y(\theta))$). Then, $\overline{S} = \underline{S} = S$.

Dealing with Heterogeneity in Practice When surplus is heterogeneous conditional on income, it may be the case that $\overline{S} > 0 > \underline{S}$. In this case, there does not exist a modification to the tax schedule in the alternative or status quo environment that can render a Pareto comparisons between the status quo and alternative environment. Here, there are several options. First, one could bias the status quo, choosing the alternative environment iff $\underline{S} > 0$. Of course, this might be overly conservative. Second, one can use average surplus, $S = E[s(\theta) g(y(\theta))]$, and decide if the alternative environment brings sufficient benefits to each point of the income distribution to warrant the lack of Pareto improvement. This approach of course violates the Pareto principle, but may be a useful application in cases with important sources of heterogeneity conditional on income.

Third, one could consider additional compensation instruments, such as capital taxation, commodity taxation, Medicaid eligibility, etc. Intuitively, when $\overline{S} > 0 > \underline{S}$, the income tax alone is too blunt an instrument to conduct compensating transfers. For example, if surplus is a function of both health and income, one could imagine making compensating transfers through modifications to both income and Medicaid / Medicare generosity and eligibility. Here, one requires estimates of $FE(\mathbf{X})$ (e.g. if $\mathbf{X} = (y, m)$ where m is Medicaid expenditures m, one requires the causal effect of the behavioral response to a transfer directed towards those not only with income y but also with Medicaid expenditures m. The key requirement is empirical estimation of the fiscal externalities.

Finally, one can consider policies that have smaller variations in surplus conditional on income. Intuitively, it is likely easier to find Pareto improvements for policies of the form "approve mergers of type X" as opposed to policies of the form "approve merger X", since the willingness to pay can be thought of as ex-ante to the set of mergers that will be approved. Efficient surplus is well-suited to addressing comparisons where the key source of heterogeneity is income.

D Inverse Optimum Derivation

Efficient social welfare weights correspond to weights used in the literature on the optimal income taxation. Up until this point, motivation for using the weights has rested solely on the Pareto principle – they apply regardless of the shape of one's own social preferences (or that of an external researcher or policymaker). But, it is straightforward to show that g(y) equals the average marginal utilities of income that rationalize the status quo tax schedule as optimal: If one's own social preferences are equal to g(y), then one would be indifferent to modifications in the nonlinear income tax schedule. In this sense, the weights are the solution to the "inverse optimum" program in optimal taxation (Bourguignon and Spadaro (2012); Blundell et al. (2009); Bargain et al. (2011); Zoutman et al. (2013)).

To see this, let $\chi(\theta)$ denote the social marginal utility of income of individual θ , so that the marginal impact on social welfare of providing an additional \$1 of resources to type θ is $\chi(\theta)$, which is normalized so that $E[\chi(\theta)] = 1$. Ratios of social marginal utilities of income, $\frac{\chi(\theta_1)}{\chi(\theta_2)}$, characterize the social willingness to pay to transfer resources from θ_2 to θ_1 and provide a generic local representation of social preferences (Saez and Stantcheva (2016)).

Proposition 6. Suppose the income tax schedule in the status quo, T(y), maximizes social welfare and let $\chi(\theta)$ denote the local social marginal utilities of income. Then, the efficient welfare weights g(y), equals the average social marginal utilities of income for those earning $y(\theta) = y$,

$$g(y) = E[\chi(\theta) | y(\theta) = y]$$

Proof. Given a tax function $\hat{T}(y; y^*, \epsilon, \eta)$, let $\hat{v}(\theta, \epsilon, \eta)$ denote the utility to type θ . By the envelope theorem, we have

$$\frac{d\hat{v}}{d\eta}|_{\eta=0} = \begin{cases} 0 & \text{if } y \notin \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right) \\ \frac{\partial v(\theta)}{\partial m} & \text{if } y \in \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right) \end{cases}$$

so that the impact on the social welfare function is $\int \chi(\theta) 1\left\{y(\theta) \in \left(y^* - \frac{\epsilon}{2}, y^* + \frac{\epsilon}{2}\right)\right\} d\mu(\theta)$, where $\chi(\theta)$ equals $\frac{\partial v(\theta)}{\partial m}$ multiplied by the local social welfare weight. Taking the limit as $\epsilon \to 0$, we have that the benefit of a small increase in η is $E[\chi(\theta) | y(\theta) = y]$; moreover, by definition the cost of a small increase in η is g(y). Optimality of the tax code implies that the welfare benefit per unit cost is equated for all y:

$$\frac{E\left[\chi\left(\theta\right)|y\left(\theta\right)=y_{1}\right]}{E\left[\chi\left(\theta\right)|y\left(\theta\right)=y_{2}\right]}=\frac{g\left(y_{1}\right)}{g\left(y_{2}\right)}$$

Finally, note that $g(y) = \frac{E[\chi(\theta)|y(\theta)=y]}{E[\chi(\theta)|y(\theta)=y_2]}g(y_2)$, so that $E[g(y)] = \frac{E[\chi(\theta)]}{E[\chi(\theta)|y(\theta)=y_2]}g(y_2)$. Now, by construction E[g(y)] = 1 and $E[\chi(\theta)] = 1$, so replacing notation of y_2 with y yields $g(y) = E[\chi(\theta)|y(\theta)=y]$.

Appendix Figure 3: Incorporating Income Effects



Notes: This figure presents the efficient social welfare weights using both the baseline specification (solid blue line) and a modified specification that incorporates an income effect (dashed red line). To calculate the modified specification with the income effect, I assume a constant elasticity of labor supply with respect to income of -0.15, similar to the estimate in Cesarini et al. (2015).

E Income Effects

The baseline specification assumes no income effects on labor supply. This section illustrates how income effects increase the marginal cost of taxation, g(y), but do so similarly at all points of the income distribution (assuming a constant elasticity). To illustrate, Appendix Figure 3 presents the baseline specification for g(y) combined with an alternative specification that incorporates income effects. For simplicity, I approximate the income effect as $\zeta(y) \frac{\tau(y)}{1-\tau(y)}$ where $\tau(y)$ is the average marginal tax rate for those in each quantile of ordinary income. For $\zeta(y)$, I take an estimate of 0.15 from Cesarini et al. (2015) who study the impact of winning the lottery in Sweden on labor supply.

As shown in Appendix Figure 3, incorporating income effects raises the marginal cost of taxation at all income levels. But, in contrast to the substitution effect and the compensated elasticity, it does not differentially affect the marginal cost of taxation at different income levels. In this sense, the broad set of conclusions that one should apply greater weight to surplus to the poor than to the rich remains true if one incorporates income effects into the analysis.

F Testing Weak Separability: Relation to Kaplow (2000, 2006, 2008, etc.) and Hylland and Zeckhauser (1976?)

There is an long debate about whether or not one should weight the willingness to pay for publicly provided goods for the poor differently than the rich. Most influentially, Hylland and Zeckhauser (1979) followed by Kaplow (1996, 2004, 2008) provide a weak separability assumption on the utility function that, if satisfied, implies that additional spending on the publicly provided good increases utility if and only if the sum of individuals' willingness to pay exceeds the mechanical cost of the publicly provided good.

To nest this result in the present framework, consider a policy of spending \$1 per capita on a publicly provided good, G. This will have a net cost to the government of c that may differ from \$1 because of any fiscal externalities from behavioral responses to the provision of the public good, $FE^G = c - 1$. Assume that individuals of income level y are willing to pay s(y) for this additional expenditure so that the average willingness to pay is E[s(y)]. Individuals are thus willing to pay the mechanical cost of the expenditure if and only if $E[s(y)] \ge 1$. In contrast, equation (6) (generalized to the case of willingness to pay that varies with y) suggests that additional spending on G is efficient if and only if $E[g(y) s(y)] \ge c$. How are these different?

It turns out that the weak separability assumption in Hylland and Zeckhauser (1979) and Kaplow (1996, 2004, 2008) has a particular restatement in the present context. Intuitively it assumes that the behavioral response to \$1 of a tax cut to those earning near y should be the same as the behavioral response to a policy that provides \$1's worth of additional G to those earning near y. Mathematically, the assumptions of Hylland and Zeckhauser (1979) imply

$$FE^G = E\left[s\left(y\right)FE\left(y\right)\right] \tag{12}$$

where s(y) is the willingness to pay of individuals with income y for the additional spending on Gand FE(y) is the fiscal externality associated with the tax cut. Hence, the total cost of the additional spending on G is equal to

$$c = 1 + FE^{G}$$

= 1 + E [s (y) FE (y)]

Hence, testing whether $E[s(y) g(y)] \ge c$ is equivalent to testing whether

$$E\left[s\left(y\right)g\left(y\right)\right] \ge c \iff E\left[s\left(y\right)\left(1+FE\left(y\right)\right)\right] \ge 1+E\left[s\left(y\right)FE\left(y\right)\right] \iff E\left[s\left(y\right)\right] \ge 1$$

So, the test for efficiency reduces to

$$E\left[s\left(y\right)\right] \geq^{?} 1$$

which asks whether the aggregate willingness to pay exceeds the mechanical cost of the policy (that does not include the fiscal externalities). Under the condition that In this sense, if equation (12) holds

for the policy change in question, one need not know either the efficient welfare weights, g(y), or the fiscal externalities induced by the policy change, FE^G . One can simply compare unweighted aggregate willingness to pay to the mechanical cost of the policy.

Of course, imposing equation (12) imposes a very strong assumption – indeed, there has been great debate in the literature about whether the proposal by Kaplow to focus on aggregate mechanical surplus is reasonable. This debate has largely hinged on critiques of the strength of the weak separability assumption. This paper provides a path forward hopefully by showing how one can translate this theoretical debate about the validity of weak separability assumptions into empirical statements about the size of fiscal externalities. For example, it seems reasonable from the estimates of the JTPA that the increase in labor supply from \$1 of government spending may have been higher than the increase in labor supply by an equivalent-sized tax cut. Instead of arguing about whether the policy spending on training was a complement to labor supply, here one can attempt to actually translate empirical estimates into the precise comparison offered in equation (12) directly. In this sense, one can turn theoretical debates into empirical debates.³⁶

³⁶Some in-kind transfer policies also are also targeted towards increasing labor earnings and generating fiscal externalities. However, Currie and Gahvari (2008) note that the empirical literature on the provision of in-kind transfers often finds minimal impacts of such transfers on labor supply. This is then interpreted as suggesting that in-kind redistribution may not be superior to redistribution through the tax schedule. However, this is somewhat misguided, as the relevant benchmark is not no fiscal externality, but rather the fiscal externalities associated with redistributive taxation. Because it's costly to provide money to those at the bottom of the income distribution because of fiscal externalities, the absence of fiscal externalities suggests FE^G is less than E[s(y) FE(y)]. This means that one may actually prefer a policy even if its aggregate surplus is negative (E[s(y)] < 1) but because it has better fiscal externalities than modifications to the tax schedule.