The Policy Elasticity

Nathaniel Hendren[*]

October, 2015

Abstract

This paper illustrates how one can use causal effects of a policy change to measure its welfare impact without decomposing them into income and substitution effects. Often, a single causal effect suffices: the impact on government revenue. Because these responses vary with the policy in question, I term them *policy elasticities*, to distinguish them from Hicksian and Marshallian elasticities. The model also formally justifies a simple benefit-cost ratio to measure the marginal value of public funds corresponding to non-budget neutral policies. Using existing causal estimates, I apply the framework to five policy changes: top income tax rate, EITC generosity, food stamps, job training, and housing vouchers.

1 Introduction

A large and growing literature in economics focuses on estimating the causal effects of government policy changes. This rise in experimental and quasi-experimental methods have made significant advances in addressing the positive question of what policy changes do to behavior. But, translating causal effects into a normative evaluation of the policy change runs into an immediate hurdle, expressed succinctly by Goolsbee (1999): "The theory largely relates to compensated elasticities, whereas the natural experiments provide information primarily on the uncompensated effects". Rarely do policy changes hold everyone's utility constant. Thus, the prevailing wisdom is that the causal effects of a policy change are not the behavioral responses that are desired for a normative analysis of that same policy change.[1]

[1]For example, Feldstein (2012) critiqued the Mirrlees review (Mirrlees et al. (2011)), an influential analysis of tax policy influenced heavily by optimal tax theory and empirical work, for using causal instead of compensated effects of tax policies:

> *While decisions on the appropriate size of government must be left to the political process, economists can assist that decision by indicating the magnitude of the total marginal cost of increased government spending. That cost depends on the structure of taxes, the distribution of income, and the compensated elasticity of*

This paper clarifies how causal effects of policy changes can be used directly in welfare analysis of government policy changes. While it is true that traditional welfare measures, such as the marginal excess burden (MEB) require compensated behavioral responses, this paper shows that instead one can measure each individual's willingness to pay out of their own income for a given policy change. The only behavioral response required for calculating this measure of welfare is the *causal impact* of the policy – a decomposition into income effects, substitution effects, or any other mechanism is not required. Real-world policy changes are often complex; therefore, these causal effects will in general be neither a pure Hicksian nor Marshallian elasticity. Because these desired responses vary with the policies in question, I term them *policy elasticities*. These are simply the difference in behavior if the policy is undertaken relative to the counterfactual world in which the policy is not undertaken, precisely the textbook definition of the causal effect of the policy.

What types of responses need to be estimated to understand the welfare impact of a government policy change? In the broad class of models in which government taxation is the only pre-existing distortion, a single causal effect is sufficient: the causal impact of the behavioral response to the policy on the government's budget.[2] The causal effect of the policy on the government budget matters because of the envelope theorem, which implies that behavioral responses to marginal policy changes don't affect utility directly. However, to the extent to which the prices faced by individuals do not reflect their resource costs (e.g. if there are marginal tax rates on labor earnings), behavioral responses impose a resource cost on society that has no impact on the agent's utility. If the government is the only distortion between private prices and social (resource) costs, the impact of the behavioral response on the government's budget is the only behavioral response required for welfare estimation.[3]

Of course, this envelope theorem logic is not new – it underscores almost all previous literature in empirical welfare economics including Harberger (1964)'s triangle and the sufficiency of the taxable income elasticity in Feldstein (1999).[4] However, these papers often focus on calculating MEB, which relies on the compensated, not causal effect on the government budget. Here, I show that not only can one use the causal effects, but the resulting welfare framework has arguably more desirable features, such as the ability to aggregate to measures of social welfare using the social marginal utilities of income.[5]

A common feature of many government policy changes is that they are not budget neutral, at least in the short run. In such cases, one can compute a simple "benefit/cost" ratio equal to the

---

the tax base with respect to a marginal change in tax rates.

[2]To be precise, this causal effect is sufficient for all components of the second derivative of the utility function. See Footnote 26.

[3]If the government is not the sole distortion in the market, one needs to estimate the causal impact on the other externalities as well as this fiscal externality. This includes not only traditional externalities such as pollution, but also externalities on one's self caused by imperfect optimization. Even in these more general models, the causal effects are sufficient for all behavioral responses; a decomposition into income and substitution effects is not required. See Appendix C.

[4]See also Kleven and Kreiner (2005) for a nice discussion of the MEB approach and Eissa et al. (2008); Eissa and Hoynes (2011) for applications to the EITC.

[5]In contrast, aggregation of MEB measures to social welfare requires knowledge of the income effects of the policy change (Section 2.7).

marginal social welfare impact of the policy per unit of government revenue expended.[6] In practice, this marginal value of public funds (MVPF) has a very simple representation. For policies which affect taxes, transfers, or provide market goods, the benefit-cost ratio is simply $MVPF = \frac{1}{1+FE}$, where FE is the impact of the behavioral response to the policy on the government budget outlays per dollar of government expenditure – otherwise known as a "fiscal externality". Policies that have positive effects on the government budget reduce the effective cost of the policy ($FE < 0$). For the provision of public goods, one needs to adjust the formula for the ratio of the willingness to pay for the public good or service relative to the beneficiaries' income (as noted by Samuelson (1954)), yielding a slightly more complicated formula: $\frac{WTP}{1+FE}$, where WTP is the individuals' willingness to pay for the publicly-provided good out of their own income.

I illustrate the framework to study the welfare impact of changes to five U.S. policies: the top marginal income tax rate, the generosity of the earned income tax credit (EITC), food stamps (SNAP), job training programs (JTPA), and housing vouchers (Section 8). To do so, I use existing causal effects to calculate the MVPF for these policies. For example, to study the impact of raising the top marginal income tax rate, I rely on the large literature studying the behavioral responses to such increases. Saez et al. (2012) and Giertz (2009) suggest mid-range estimates that 25-50% of the mechanical revenue that is raised from increasing the top marginal income tax rate is lost due to the behavioral response to the policy. This suggests a MVPF of taxing top earners of $1.33-$2. For the EITC generosity, there is a large literature studying the impact of EITC expansions on labor earnings (Hotz and Scholz (2003)). Existing causal estimates suggest increasing EITC generosity leads to a cost that is ~14% above the mechanical cost due to behavioral responses. This suggests a MVPF of increasing EITC generosity of $0.88.

One can use the MVPF to compare across policies using Okun's leaky bucket experiment (Okun (1975)): how much resources is society willing to lose to transfer from one person to another?[7] For example, consider the MVPFs from the EITC expansions and the top marginal income tax schedule. Combining these suggests additional redistribution is desired if and only if one prefers $0.44-0.66 in the hands of an EITC beneficiary relative to $1 in the hands of the rich (earnings > $400K). From a positive perspective, the existing causal estimates of the behavioral responses to taxation suggests the U.S. tax schedule implicitly values an additional $0.44-0.66 to an EITC recipient as equivalent to $1 to someone subject to the top marginal income tax rate.

---

[6] To the best of my knowledge, this measure of the marginal value of public funds (MVPF) was initially proposed by Mayshar (1990) (see equation 9 on page 267 of Mayshar (1990); also, see Slemrod and Yitzhaki (1996, 2001) for similar definitions). Here, I show that it has the unique feature that the only behavioral responses required for their measurement are the causal effects of the non-budget neutral policy in question. This contrasts with more traditional definitions of the marginal cost of public funds (MCPF) that seek to augment the standard willingness to pay for the public expenditure (given by a Samuelson condition) with the distortionary cost of raising the required tax revenue (Ballard and Fullerton (1992)).

[7] In contrast, comparisons of MEBs across policies requires adding back in the income effects that were removed to form the MEB. This problematic feature of MEB was initially shown by Diamond and Mirrlees (1971). See also Auerbach and Hines (2002) for a simple illustration of this on page 1370, equation 3.24. In this sense, the aggregation of welfare across people is easier with the MVPF than the MEB.

Relation to Previous Literature   This paper is related to a large set of previous literature studying the marginal welfare impact of policy changes and in particular the definition of the marginal value of public funds adopted in Mayshar (1990), Slemrod and Yitzhaki (1996, 2001) and Kleven and Kreiner (2006)). Relative to this literature, the primary contribution of this paper is to illustrate why this particular definition of the MVPF does not require a decomposition of the causal effect of the policy, and why other commonly used marginal welfare definitions rely on behavioral responses other than the causal effect, such as the compensated effect (as in the MEB analysis of Kleven and Kreiner (2005); Eissa et al. (2008); Eissa and Hoynes (2011)) or the causal effect of an augmented policy that includes a hypothetical tax increase to close the government budget constraint (as in the traditional Atkinson-Stern-Stiglitz-Dasgupta definition of the MCPF in Stiglitz and Dasgupta (1971); Atkinson and Stern (1974); see also Ballard and Fullerton (1992)).[8]

The paper is also related to the literature on optimal taxation. While original optimal tax formulas generally sought to write optimal tax formulas using underlying structural price and income elasticities (e.g. Diamond and Mirrlees (1971) and Saez (2001)), more recent optimal tax literature has noted that often one need not decompose behavioral responses into underlying structural components (e.g. Piketty and Saez (2013)), which can aid both in the theoretical expression and empirical implementation of optimal tax formulas. In this sense, the paper builds on the recent literature on sufficient statistics (Chetty (2009b)) and in particular the sufficiency of the taxable income elasticity (Feldstein (1999); Chetty (2009a)). It is well known that the taxable income elasticity is no longer sufficient in cases when there are responses to the policy on multiple tax bases with different marginal tax rates (e.g. capital and labor income (Saez et al. (2012)) or intensive versus extensive margin responses (Kleven and Kreiner (2006))). However, the present analysis shows that the causal impact of the behavioral response on the government budget (e.g. tax revenue) as opposed to the tax base (e.g. taxable income) remains sufficient even in cases where the behavioral response by individuals occurs on multiple tax margins. This suggests focusing on the tax revenue impacts, as opposed to taxable income, may be the most general empirical approach for welfare analysis.

The rest of this paper proceeds as follows. Section 2 presents the model, outlines the welfare framework, and compares it to existing literature. Section 3 draws upon existing causal estimates to study the desirability of changing the top marginal income tax rate, EITC generosity, food stamps, job training, and housing vouchers. Section 4 concludes. The Online Appendix provides some derivations of the welfare formulas and also discusses extensions of the model to non-marginal welfare analysis (Appendix B), externalities (Appendix C), general equilibrium effects (Appendix D), and provides further clarification of the role of the Hicksian elasticity in previous literature (Appendix E).

2   Model

This section formalizes a general model of individual behavior to illustrate the ability to utilize causal effects, instead of compensated effects, for normative analysis of government policy changes. The

---

[8]In this sense, it is related to Auerbach (1985) who noted the different conceptual experiments underlying the MEB versus MCPF.

generality of the model captures many realistic issues faced in empirical applications and also allows the model to nest many models in previous literature. Despite its generality, the model will motivate a very simple benefit-cost ratio in equation (16) in Section 2.10. This benefit-cost ratio is simply the individuals' willingness to pay for the policy change out of their own income, normalized by the total government cost of the policy inclusive of the impact of behavioral responses on the government budget (i.e. fiscal externalities) – a term I call the marginal value of public funds,

$$MVPF = \frac{\text{Individuals' WTP for Policy Change}}{\text{Cost to Gov't}} = \frac{\text{"}Benefit\text{"}}{\text{"}Cost\text{"}}$$

When using the MVPF, the behavioral responses required for such a measure will be causal, not compensated, effects. The model clarifies below how this notion is distinct from traditional measures of the marginal cost of public funds (MCPF) and marginal excess burden (MEB), but aligns with a definition proposed by Mayshar (1990) and Slemrod and Yitzhaki (1996, 2001). I then use this formula to translate existing causal effects for five policies into welfare statements about those policies in Section 3.

2.1 Setup

More formally, there exist a set of individuals in the population, indexed by $i \in I$. They make two choices: they choose goods to consume, $\mathbf{x_i} = \{x_{ij}\}_{j=1}^{J_X}$, and labor supply activities, $\mathbf{l_i} = \{l_{ij}\}_{j=1}^{J_L}$.[9] There also exists a government that does three things: it provides a range of publicly provided goods and services to each individual, $\mathbf{G_i} = \{G_{ij}\}_{j=1}^{J_G}$, provides monetary transfers to each individual, T_i, and imposes linear taxes[10] on goods, $\tau_\mathbf{i}^\mathbf{x} = \left\{\tau_{ij}^x\right\}_{j=1}^{J_X}$ and labor supply activities, $\tau_\mathbf{i}^\mathbf{l} = \left\{\tau_{ij}^l\right\}_{j=1}^{J_L}$.

Individuals value their goods, labor supply activities, and publicly provided goods and services according to the utility function:

$$u_i\left(\mathbf{x}_i, \mathbf{l}_i, \mathbf{G}_i\right) \tag{1}$$

which is allowed to vary arbitrarily across people.[11]

To simplify the exposition, I assume a stylized model of production in which one unit of any type of labor supply produces 1 unit of any type of good under perfect competition. This means that individuals face a single linear budget constraint[12] given by

$$\underbrace{\left(1 + \tau_\mathbf{i}^\mathbf{x}\right)\mathbf{x}_i}_{\text{Spending on Goods}} \leq \underbrace{\left(1 - \tau_i^l\right)\mathbf{l}_i}_{\text{Earnings}} + \underbrace{T_i}_{\text{Transfers}} + \underbrace{y_i}_{\text{Other Income}} \tag{2}$$

---

[9]For example, j can index time so that l_{ij} is the labor supply of individual i in time j. Or, l_{i1} could be labor supplied in wage work and l_{i2} could be labor supplied in the informal (un-taxed) sector.

[10]Because I focus on marginal policy changes, the model can consider nonlinear tax settings by interpreting T_i as "virtual income" and τ_{ij}^l as the marginal tax on labor earnings.

[11]Note that these publicly provided goods could be market or non-market goods. For example, one can capture a setting where G is a market good by assuming the utility function has a form: $u_i\left(x_1, x_2, G\right) = \tilde{u}_i\left(x_1, x_2 + G\right)$, so that G and x_2 would be perfectly substitutable.

[12]All vector multiplication is the standard dot product (e.g. $\left(1 + \tau_\mathbf{i}^\mathbf{x}\right)\mathbf{x}_i = \sum_{j=1}^{J_X}\left(1 + \tau_{ij}^x\right)x_{ij}$)

The individuals expenditure on goods, inclusive of the taxes τ_i^x, must be less than the earnings from labor supply activities (plus taxes τ_i^l), transfers, T_i, and non-labor income, y_i.[13] This simplified production structure rules out many interesting features that can easily be added to a more general model, including imperfect competition (i.e. producer surplus), production externalities (e.g. spillovers), and pecuniary externalities (in which case real prices would not always be 1).[14] I assume the marginal cost to the government of producing publicly-provided goods is given by $\mathbf{c}^G = \{c_j\}_j$ for $j = 1, .., J_G$.[15]

Each individual takes taxes, transfers, non-labor income, and the provision of publicly-provided goods as given and chooses goods and labor supply activities to maximize utility. This allows individuals to obtain a level of utility V_i given by

$$V_i\left(\tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}, y_i\right) = \max_{\mathbf{x}, \mathbf{l}} u_i\left(\mathbf{x}, \mathbf{l}, \mathbf{G_i}\right)$$

$$s.t. \quad \left(1 + \tau_{\mathbf{i}}^{\mathbf{x}}\right) \mathbf{x}_i \le \left(1 - \tau_i^l\right) \mathbf{l}_i + T_i + y_i$$

where V_i depends on taxes, transfers, income, and publicly provided goods. The choices of goods and labor supply activities generated by the agent's problem are denoted $x_{ij}^m\left(\tau_{\mathbf{i}}^{\mathbf{x}}, \tau_{\mathbf{i}}^{\mathbf{l}}, T_i, \mathbf{G}_i, y_i\right)$ and $l_{ij}^m\left(\tau_{\mathbf{i}}^{\mathbf{x}}, \tau_{\mathbf{i}}^{\mathbf{l}}, T_i, \mathbf{G}_i, y_i\right)$.[16] Because the utility function is allowed to vary arbitrarily across people, it will be helpful to normalize by the individual's marginal utility of income, λ_i,

$$\lambda_i = \frac{\partial V_i}{\partial y_i}$$

which is the Lagrange multiplier from the type i maximization program. For measuring welfare, it will also be helpful to define the expenditure function, $E_i\left(u; \tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G}_i\right)$, of individual i to be the amount of income y_i required for individual i to obtain utility level u in a world with taxes, transfers, and publicly provided good $\left(\tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}\right)$.[17]

The indirect utility function provides a measure of individual i's utility; to move to social welfare, I assume there exists some set of Pareto weights, $\{\psi_i\}$, for each individual i, so that social welfare is given by

$$\underbrace{W\left(\left\{\tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}, y_i\right\}_i\right)}_{\text{Social Welfare}} = \underbrace{\int_{i \in I} \psi_i V_i\left(\tau_{\mathbf{i}}^{\mathbf{x}}, \tau_{\mathbf{i}}^{\mathbf{l}}, T_i, \mathbf{G}_i, y_i\right) di}_{\text{Weighted Sum of Individual Utilities}} \tag{3}$$

Note that this social welfare function in principle depends on a very rich set of policy choices by the government: it's an implicit function of the taxes, transfers, and publicly provided goods to every

---

[13]I allow (but do not require) taxes and transfers to be individual-specific. This allows the model to nest the standard MEB experiment that requires (potentially infeasible) individual-specific lump-sum transfers to compare a policy change to a first-best allocation.

[14]See Appendix D for a discussion of GE effects and Appendix C for a discussion of externalities.

[15]Note this nests the case of a pure public good by assuming $c_j^G = \frac{1}{N}$ and G_{ij} is constant across i.

[16]The superscript "m" refers to the fact that these are standard Marshallian demand functions.

[17]Note that the standard duality result implies:

$$E_i\left(V_i\left(\tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}, y_i\right); \tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}\right) = y_i$$

individual in the economy. By allowing for an arbitrary set of Pareto weights, one can allow for social preferences for different people in the economy (e.g. a preference for redistribution from rich to poor).

These canonical measures of individual welfare, V_i, and social welfare, W, provide a mapping from policy choices of the government to measures of well being. The remainder of the paper will seek to characterize the welfare impact of small changes to government policy.

2.2 Policy Paths and Potential Outcomes

Imagine the government makes a policy change. To describe an arbitrary policy change, I follow a setup that aligns closely with the canonical definitions of causal effects as differences of "potential outcomes" (e.g. Angrist and Pischke (2008)). For the purposes of this paper, and for comparison to earlier literature estimating MEB, I consider small policy changes. To be specific, I define a "policy path", $P(\theta)$. For any θ in a small region near 0, $\theta \in (-\epsilon, \epsilon)$, let $P(\theta)$ be the taxes, transfers, and publicly provided goods to each individual,

$$P(\theta) = \left\{ \hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta), \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta), \hat{T}_i(\theta), \hat{\mathbf{G}}_i(\theta) \right\}_{i \in I} \tag{4}$$

where the "^" indicates the policies are functions of θ. A policy path can describe a policy that increases/decreases taxes, increases spending on a public good, etc. In this sense, this path provides a method for describing a wide array of policy changes.

I make two assumptions about how the policy varies with θ. First, I normalize the value of the policy at $\theta = 0$ to be the status quo:

$$\left\{ \hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(0), \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(0), \hat{T}_i(0), \hat{\mathbf{G}}_i(0) \right\}_{i \in I} = \left\{ \tau_{\mathbf{i}}^{\mathbf{x}}, \tau_{\mathbf{i}}^{\mathbf{l}}, T_i, \mathbf{G}_i \right\}_{i \in I}$$

Second, I assume that the policy path is continuously differentiable in θ (i.e. $\frac{d\hat{\tau}_{ij}^x}{d\theta}$, $\frac{d\hat{\tau}_{ij}^l}{d\theta}$, $\frac{d\hat{T}_i}{d\theta}$, and $\frac{d\hat{G}_{ij}}{d\theta}$ exist and are continuous in θ).[18] Intuitively, $P(\theta)$ traces out a smooth path of government policies, centered around the status quo. By using this path, one can easily consider policies that vary multiple policy parameters at the same time. Given a path $P(\theta)$, I consider the welfare impact of following the path, parameterized by an increase in θ. This can be interpreted as following a policy path or evaluating a policy direction (e.g. "increasing taxes on labor earnings" would be a path with $\frac{d\tau_{labor}^l}{d\theta} > 0$, or "increasing spending on schools financed with a cigarette tax" would be a path with $\frac{d\tau_{cigarettes}^x}{d\theta} > 0$ and $\frac{dG_{schoolspending}}{d\theta} > 0$ at the same time, etc.).[19]

The normative question of interest is "*should* we follow the policy path?". Before asking this normative question, first consider the positive question of what the policy change would do to behavior.

---

[18]This does not require that the behavioral response to the policy be continuously differentiable. For notational convenience in the text, I will assume the behavioral responses are continuously differentiable. However, in the empirical application to the study of the EITC expansion in Section 3, I allow for extensive margin labor supply responses (which is a key feature of the behavioral response to EITC expansions, and is known to be an important factor in MEB estimation (Eissa et al. (2008), Eissa and Hoynes (2011))).

[19]I have not specified a scale/speed for the policy path. In practice, one can normalize the speed of the policy to one unit of a tax or one dollar of revenue raised, as illustrated in the application in Section 3.

Given a policy path, I assume individuals choose goods and labor supply activities, $\hat{\mathbf{x}}_i(\theta) = \{\hat{x}_{ij}(\theta)\}_i$ and $\hat{\mathbf{l}}_i(\theta) = \{\hat{l}_{ij}(\theta)\}_i$, that maximize their utility under policy $P(\theta)$.[20] In the now-standard language of Angrist and Pischke (2008), $\hat{x}(\theta)$ and $\hat{l}(\theta)$ are the "potential outcomes" of individual's choices of goods and labor supply activities if policy world θ is undertaken. As θ moves away from 0, $\hat{\mathbf{x}}(\theta)$ and $\hat{\mathbf{l}}(\theta)$ trace out the causal effect of the policy change on the individual's behavior..

In addition to the individual's behavior, the policy will also impact the government budget. To keep track of these effects, let $\hat{t}_i(\theta)$ denote the net government resources directed towards type i,

$$\underbrace{\hat{t}_i(\theta)}_{\text{Net Resources}} = \underbrace{\mathbf{c}^G \hat{\mathbf{G}}_i(\theta)}_{\text{Public-Provided Goods}} + \underbrace{\hat{T}_i(\theta)}_{\text{Transfers}} - \underbrace{\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta)\hat{\mathbf{x}}_i(\theta) + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta)\hat{\mathbf{l}}_i(\theta)\right)}_{\text{Tax Revenue}} \quad (5)$$

where $\mathbf{c}^G \hat{\mathbf{G}}_i(\theta)$ is the government expenditure on publicly provided goods to individual i, $\hat{T}_i(\theta)$ is the government transfers to type i, and $\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta)\hat{\mathbf{x}}_i(\theta) + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta)\hat{\mathbf{l}}_i(\theta)$ is the tax revenue collected from individual i on goods and labor supply activities.

With this definition of $\hat{t}_i$, the total impact of a policy on the government's budget is given by $\int_{i \in I} \frac{d\hat{t}_i}{d\theta} di$. The analysis does not require policies to be budget-neutral, but budget-neutrality of a policy path could be imposed by assuming the policy change has no aggregate impact on the government budget:

$$\int_{i \in I} \frac{d\hat{t}_i}{d\theta} di = 0 \quad \forall \theta$$

where $\frac{d\hat{t}_i}{d\theta}$ is obtained by taking the derivative of equation (5):

$$\frac{d\hat{t}_i}{d\theta} = \mathbf{c}^G \frac{d\hat{\mathbf{G}}_i}{d\theta} + \frac{d\hat{T}_i}{d\theta} - \frac{d}{d\theta}\left[\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta)\hat{\mathbf{x}}_i(\theta) + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta)\hat{\mathbf{l}}_i(\theta)\right]$$

The term $\mathbf{c}^G \frac{d\hat{\mathbf{G}}_i}{d\theta}$ is how much the policy changes spending on publicly provided goods; $\frac{d\hat{T}_i}{d\theta}$ is how much the policy increases direct transfers; and the last term is the impact of the policy on the net tax revenue from goods and labor supply activities.

The impact of the policy on individual behavior and on the government budget are related through the mechanical and behavioral impact of the policy on net tax revenue from goods and labor supply activities:

$$\frac{d}{d\theta}\left[\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta)\hat{\mathbf{x}}_i(\theta) + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta)\hat{\mathbf{l}}_i(\theta)\right] = \underbrace{\left(\frac{d\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}}{d\theta}\hat{\mathbf{x}}_i + \frac{d\hat{\tau}_{\mathbf{i}}^{\mathbf{l}}}{d\theta}\hat{\mathbf{l}}_i\right)}_{\substack{\text{Mechanical Impact} \\ \text{on Govt Revenue}}} + \underbrace{\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}\frac{d\hat{\mathbf{x}}_i}{d\theta} + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}\frac{d\hat{\mathbf{l}}_i}{d\theta}\right)}_{\substack{\text{Behavioral Impact} \\ \text{on Govt Revenue}}} \quad (6)$$

---

[20]These can be calculated in theory by evaluating the Marshallian demands at each θ:

$$\hat{x}_{ij}(\theta) = x_{ij}^m\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta), \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta), \hat{T}_i(\theta), \hat{\mathbf{G}}_i(\theta)\right) \quad \forall j = 1..J_X$$

$$\hat{l}_{ij}(\theta) = l_{ij}^m\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta), \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta), \hat{T}_i(\theta), \hat{\mathbf{G}}_i(\theta)\right) \quad \forall j = 1..J_L$$

but in practice they can be estimated by looking at the causal impact of the policy without knowledge of the underlying mechanisms.

The mechanical effect is the change in revenue holding behavior constant. This would be the marginal budget impact of the policy if one did not account for any behavioral responses. The behavioral impact is the effect of the behavioral response to the policy on the government's budget.

2.3 Definitions of Welfare

Moving from positive to normative analysis requires a definition of welfare. A primary aim of the paper is to argue that one need not estimate the marginal excess burden (MEB) associated with a policy change, and one can instead focus on the marginal willingness to pay for the policy change. In Section 2.6, I discuss MEB; however I begin by defining the individual's willingness to pay out of their own income to follow the policy path, as this will be the definition of welfare that corresponds to the use of causal, not compensated, behavioral responses.

Let $\hat{V}_i(\theta)$ denote the utility obtained by type i under the policy $P(\theta)$. The marginal impact of the policy on the utility of individual i is given by $\frac{d\hat{V}_i}{d\theta}|_{\theta=0}$. Normalizing by the marginal utility of income, the individual's own willingness to pay (out of their own income) for a marginal policy change is given by $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$.[21]

With this definition of individual welfare, aggregation to social welfare is straightforward by summing across individuals using their Pareto weights, $\frac{d\hat{W}}{d\theta}|_{\theta=0} = \int_{i\in I} \eta_i \frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} di$, where $\eta_i = \psi_i \lambda_i$ is the social marginal utility of income of individual i (see Saez and Stantcheva (2013)).[22] Social marginal utilities η_i can be interpreted in terms of Okun's classic bucket experiment (Okun (1975)): Society is indifferent to transferring $\frac{\eta_1}{\eta_2}$ resources to individual 2 as opposed to \$1 to individual 1. If $\eta_1 < \eta_2$, society is willing to lose resources in order to make a transfer from individual 1 to individual 2.

In principle, values of the social marginal utility of income are a subjective matter. However, Hendren (2014) provides one potential alternative that replaces these social marginal utilities of income with the "inequality deflator", which measures the marginal price of transferring resources between

---

[21] It is well-known that $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$ is equivalent to two other canonical measures of welfare for marginal policy changes. First, the equivalent variation, $EV_i(\theta)$, of policy $P(\theta)$ for type i is the amount that the consumer would be indifferent to accepting in lieu of the policy change. $EV_i(\theta)$ solves

$$V_i\left(\tau_{\mathbf{i}}^{\mathbf{1}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}, y_i + EV_i(\theta)\right) = \hat{V}_i(\theta)$$

Second, the compensating variation, $CV_i(\theta)$, of policy $P(\theta)$ for type i is the amount of money that must be compensated to the agent after the policy change to bring her back to her initial utility level. $CV_i(\theta)$ solves

$$V_i\left(\tau_{\mathbf{i}}^{\mathbf{1}}(\theta), \tau_{\mathbf{i}}^{\mathbf{x}}(\theta), T_i(\theta), \mathbf{G_i}(\theta), y_i - CV_i(\theta)\right) = \hat{V}_i(0)$$

It is straightforward to verify (e.g. Schlee (2013)) that:

$$\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} = \frac{d[EV_i]}{d\theta}|_{\theta=0} = \frac{d[CV_i]}{d\theta}|_{\theta=0}$$

[22] Note this remains true even if the welfare weights are not fixed and are functions of utility levels, since marginal policy changes do not change the welfare weights. For example, if $W = \int_{i\in I} G(V_i) di$ for a concave function G, then the social marginal utility of income would be $\eta_i = G'\left(\hat{V}_i(0)\right)\lambda_i$.

individuals using modifications to the tax schedule. Under certain conditions, the use of such weights corresponds to searching for potential Pareto improvements in the spirit of Kaldor (1939) and Hicks (1939), and does not require a subjective specification of a social welfare function. For now, we imagine the researcher has chosen a given set of social marginal utilities of income, either subjectively or using the inequality deflator of Hendren (2014).

2.4 Defining Behavioral Responses: The Policy Elasticity

In principle, the behavioral responses to the policy can be driven by any number of underlying reasons such as the impact of changing prices, changing disposable incomes, or complementary effects from changes in the provision of public goods and services. A Marshallian elasticity holds disposable income constant. A Hicksian elasticity holds utility constant. A key benefit of these well-agreed upon elasticity definitions is that they allow the economist to state clearly what is being held constant in the conceptual experiment in question.

However, in practice policy changes hold neither income nor utility constant. All aspects of the policy are varying simultaneously which could be interacting to produce behavioral responses through a potentially complicated underlying process.[23] To provide a vocabulary that allows one to make clear what is being held constant in the conceptual experiment in which the policy occurs, I define the *policy response* of x_{ij} and l_{ij} to be the local causal effect of the policy on x_{ij} and l_{ij}. Similarly, because empirical estimation often occurs in logs, I define the *policy elasticity* of x_{ij} and l_{ij} to be the local causal effect of the policy on $log(x_{ij})$ and $log(l_{ij})$.

Definition 1. The *policy response* of x_{ij} (or l_{ij}) with respect to policy $P(\theta)$ is given by $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0}$ (or $\frac{d\hat{l}_{ij}}{d\theta}|_{\theta=0}$). The *policy elasticity* of x_{ij} (or l_{ij}) is given by $\hat{\epsilon}_{ij}^x = \frac{dlog(\hat{x}_{ij})}{d\theta}|_{\theta=0}$ (or $\hat{\epsilon}_{ij}^l = \frac{dlog(\hat{l}_{ij})}{d\theta}|_{\theta=0}$)

The policy elasticity is simply the causal effect of the policy in question. Given these definitions,

---

[23]In principle, the response to a given policy could be driven by a multitude of underlying factors. To see this, note that one can write these causal effects using an individual's Marshallian or Hicksian demand functions. For example, suppose $x_{ij}^m \left(\tau_{\mathbf{i}}^{\mathbf{x}}, \tau_{\mathbf{i}}^{\mathbf{l}}, T_i, \mathbf{G}_i, y_i\right)$ is the Marshallian demand of individual i for good j. Then,

$$\frac{d\hat{x}_{ij}}{d\theta} = \underbrace{\sum_{j=1}^{J_X} \frac{\partial x_{ij}^m}{\partial \tau_{ij}^x} \frac{d\tau_{ij}^x}{d\theta} + \sum_{j=1}^{J_X} \frac{\partial x_{ij}^m}{\partial \tau_{ij}^x} \frac{d\tau_{ij}^x}{d\theta} + \frac{\partial x_{ij}^m}{\partial T} \frac{dT_i}{d\theta} + \sum_{j=1}^{J_G} \frac{\partial x_{ij}^m}{\partial G_{ij}} \frac{d\hat{G}_{ij}}{d\theta}}_{\text{Marshallian Expansion}}$$

Similarly, suppose $x_{ij}^h \left(\tau_{\mathbf{i}}^{\mathbf{x}}, \tau_{\mathbf{i}}^{\mathbf{l}}, T_i, \mathbf{G}_i, u\right)$ is the Hicksian demand of individual i for good j (evaluated at utility level u). Then,

$$\frac{d\hat{x}_{ij}}{d\theta} = \underbrace{\sum_{j=1}^{J_X} \frac{\partial x_{ij}^h}{\partial \tau_{ij}^x} \frac{d\tau_{ij}^x}{d\theta} + \sum_{j=1}^{J_X} \frac{\partial x_{ij}^h}{\partial \tau_{ij}^x} \frac{d\tau_{ij}^x}{d\theta} + \frac{\partial x_{ij}^h}{\partial T} \frac{dT_i}{d\theta} + \sum_{j=1}^{J_G} \frac{\partial x_{ij}^h}{\partial G_{ij}} \frac{d\hat{G}_{ij}}{d\theta} + \frac{\partial x_{ij}^h}{\partial u} \frac{d\hat{V}_i}{d\theta}}_{\text{Hicksian Expansion}}$$

the total impact of the policy change on government revenue in equation (6) has three representations:

$$\underbrace{\frac{d}{d\theta}\left[\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta)\hat{\mathbf{x}}_i(\theta) + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta)\hat{\mathbf{l}}_i(\theta)\right] - \left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}\frac{d\hat{\mathbf{x}}_i}{d\theta} + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}\frac{d\hat{\mathbf{l}}_i}{d\theta}\right)}_{\text{Total - Mechanical Impact on Govt Revenue}} = \underbrace{\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}\frac{d\hat{\mathbf{x}}_i}{d\theta} + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}\frac{d\hat{\mathbf{l}}_i}{d\theta}\right)}_{\substack{\text{Behavioral Impact} \\ \text{on Govt Revenue (levels)}}} = \underbrace{\left(\sum_j^{J_X} r_{ij}^x \hat{\epsilon}_{ij}^x + \sum_j^{J_L} r_{ij}^l \hat{\epsilon}_{ij}^x\right)}_{\substack{\text{Behavioral Impact} \\ \text{on Govt Revenue (logs)}}} \quad (7)$$

where the weights for the log responses, $\hat{r}_{ij}^x = \hat{\tau}_{ij}^x \hat{x}_{ij}$ (or $\hat{r}_{ij}^x = \hat{\tau}_{ij}^l \hat{l}_{ij}^x$), equal the government revenue on each good (or labor supply).

The compensated response   In contrast to the policy response, one can also define the compensated response to the policy, $\frac{d\hat{x}_{ij}^c}{d\theta}$. To do so, one needs to subtract the income effects associated with the policy change.

$$\frac{d\hat{x}_{ij}^c}{d\theta} = \frac{d\hat{x}_{ij}}{d\theta} - \underbrace{\frac{\partial x_{ij}^h}{\partial u}\frac{d\hat{V}_i}{d\theta}}_{\text{Income Effect}} = \frac{d\hat{x}_{ij}}{d\theta} - \underbrace{\frac{\partial x_{ij}^m}{\partial y}\frac{\frac{d\hat{V}_i}{d\theta}}{\lambda_i}}_{\text{Income Effect}} \quad (8)$$

The compensated response to the policy P is equal to the policy response minus the portion of the behavioral response that is due to changes in income or utility. This income effect can be represented using the Hicksian response to utility, $\frac{\partial x_{ij}^h}{\partial u}\frac{d\hat{V}_i}{d\theta}$, which adjusts behavior for the change in the utility induced by the policy, $\frac{d\hat{V}_i}{d\theta}$. Alternatively, the income effect can be represented using the Marshallian response to income, $\frac{\partial x_{ij}^m}{\partial y}\frac{\frac{d\hat{V}_i}{d\theta}}{\lambda_i}$, which adjusts the policy response for the change in the income-value of the policy change, $\frac{\frac{d\hat{V}_i}{d\theta}}{\lambda_i}$.

As illustrated in Equation (8), the construction of the compensated response to a policy change is potentially much more difficult than the policy response. Not only does one need to know the behavioral response to income ($\frac{\partial x_{ij}^m}{\partial y}$) or utility ($\frac{\partial x_{ij}^h}{\partial u}$), but to scale properly one also needs to know the utility impact of the policy ($\frac{d\hat{V}_i}{d\theta}$) or the marginal willingness to pay for the policy change, $\frac{\frac{d\hat{V}_i}{d\theta}}{\lambda_i}$. As is well-known, the compensated response to a given policy is the policy response to an alternative policy, $\hat{P}^c$, that holds individuals' utilities constant via individual-specific lump-sum transfers (Auerbach and Hines (2002)).

The next section illustrates that one can construct a comprehensive welfare framework for local evaluation of government policy changes that relies solely on the policy elasticities of the policy in question, and does not require knowledge of either the compensated effect, or any of the underlying marshallian or hicksian demand functions (conditional on knowledge of the policy response).

2.5 The Marginal Willingness to Pay for the Policy Change

How much are individuals willing to pay for a policy change out of their own income. Appendix A shows this is given by

$$
\begin{aligned}
\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} &= \frac{\frac{du_i}{d\mathbf{G}_i}}{\lambda_i}\frac{d\hat{\mathbf{G}}_i}{d\theta} + \frac{dT_i}{d\theta} + \frac{d\hat{\tau}_i^x}{d\theta}\mathbf{x}_i + \frac{d\hat{\tau}_i^l}{d\theta}\mathbf{l}_i \\
&= \underbrace{\frac{d\hat{t}_i}{d\theta}|_{\theta=0}}_{\text{Net Resources}} + \underbrace{\left(\frac{\frac{du_i}{d\mathbf{G}_i}}{\lambda_i} - \mathbf{c}^G\right)\frac{d\hat{\mathbf{G}}_i}{d\theta}|_{\theta=0}}_{\substack{\text{Public Spending/}\\ \text{Mkt Failure}}} + \underbrace{\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}\frac{d\hat{\mathbf{x}}_i}{d\theta} + \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}\frac{d\hat{\mathbf{l}}_i}{d\theta}\right)}_{\substack{\text{Behavioral Impact}\\ \text{on Govt Revenue}}}
\end{aligned}
\tag{9}
$$

By the envelope theorem, the extent to which individuals respond to the policy change does not affect their utility directly. Hence, the first line of Equation 9 does not contain behavioral responses. However, behavioral responses are crucial when one attempts to account for the cost of the policy, as shown in the second line of Equation 9.

The first term, $\frac{d\hat{t}_i}{d\theta}$, is straightforward: it is the change in net government resources provided to individual i from the government, which is the difference between the change in spending on publicly provided goods and transfers and the collection of taxes on goods and labor supply activities. For budget neutral policies, recall that $\int_i \frac{d\hat{t}_i}{d\theta}di = 0$; in this sense, $\frac{d\hat{t}_i}{d\theta}$ captures the redistributive impact of the policy. These transfers increase social welfare to the extent to which those receiving the net transfer have higher values of the social marginal utility of income than those who pay for the net transfer.

The second term captures the value of any changes to publicly provided goods, $\frac{d\hat{G}_{ij}}{d\theta}|_{\theta=0}$. This is given by the difference between the willingness to pay for the publicly provided goods and their costs of production, $\left(\frac{\frac{du_i}{d\mathbf{G}_i}}{\lambda_i} - \mathbf{c}^G\right)\frac{d\hat{\mathbf{G}}_i}{d\theta}|_{\theta=0} = \sum_{j=1}^{J_G}\left(\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} - c_j^G\right)\frac{d\hat{G}_{ij}}{d\theta}|_{\theta=0}$. This component is well-known and popularized in Samuelson (1954). One can interpret this number as the size of the market inefficiency being addressed by the publicly provided goods. If the private market can efficiently supply and allocate all goods, then agents would be able to pay c_g to obtain a unit of a good that is equivalent to the publicly provided good, so that $\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} = c_j^G$. If the private market does not provide such goods as efficiently as the government (or vice-versa), then one needs to know the difference between the costs and benefits of its provision.

The final term in Equation 9 summarizes the importance of behavioral responses. It is the impact of the behavioral response to the policy on the government's budget. It is a weighted sum of the *policy responses* on behavior, $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0}$ and $\frac{d\hat{l}_{ij}}{d\theta}|_{\theta=0}$, with the weights given by the marginal tax rates.[24]

The causal effect matters because of a fiscal externality. The envelope theorem guarantees that

---

[24]Although this causal effect is the impact of a marginal change in the policy, in practice causal effects are often measured using discrete changes in policies. Appendix B provides intuitive conditions under which the non-marginal causal effects (i.e. $\hat{x}_{ij}(1) - \hat{x}_{ij}(0)$ instead of $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0}$) can be used to measure the individual's willingness to pay for the policy change.

behavioral responses do not affect utility directly; however, when prices do not reflect their resource costs (as is the case with taxation), behavioral responses impose a cost on those bearing the difference between the prices faced by the individual and their resource costs.[25] Conditional on calculating this fiscal externality, behavioral responses are not required for measuring individual's willingness to pay for the policy change.[26]

Relation to Feldstein (1999)   If there is only one tax rate on aggregate taxable income and social marginal utilities of income are the same for all types, then the aggregate taxable income elasticity is sufficient for capturing the behavioral responses required for welfare analysis.[27] This insight was recently popularized in Feldstein (1999). I provide two clarifications to this result. First, it is in general neither the Hicksian (compensated) nor the Marshallian (uncompensated) elasticity of taxable income that is desired for analyzing the welfare impact of government policy. Rather, it is the taxable income elasticity associated with the policy in question, which depends on how the revenue is spent.

Second, as is well known, the taxable income elasticity is not sufficient to the extent to which individuals face multiple tax rates. For example, if capital income is taxed at a different rate than labor income, the elasticity of the sum of these two incomes would not be sufficient (Saez et al. (2012)). If behavioral responses occur on both the participation and intensive margin, then the aggregate earnings elasticity is not sufficient (Kleven and Kreiner (2006)). Moreover, one also needs to know the extent to which policies affect consumption of subsidized goods or services (e.g. enrollment in government programs such as SSDI or unemployment insurance). While subsequent literature tends to suggest a need for adding additional elasticities to the analysis[28], the present analysis shows that if one switches the dependent variables in these analyses from the components of taxable income to aggregate tax *revenue*, such a decomposition of the mechanics of the behavioral response is not required.[29] Of course, there are many reasons to be interested in the mechanisms driving such a response; but calculating the marginal welfare impact of the policy change in question is not one of them.

---

[25]As discussed in Appendix C, if there are other externalities one also requires an estimate of the impact of the policy on those externalities as well. However, the causal effects remain the desired behavioral responses.

[26]A decomposition of causal effects into income and substitution effects does not generally help measure the size of market inefficiency, $\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} - c_j^G$. Income and price effects depend on the Hessian (2nd derivative) of the utility function, whereas the size of the market failure, $\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} - c_j^G$, depends on the first derivatives of the utility function (Mas-Colell et al. (1995)).

[27]To see this, note that if $\tau_1 = \tau_2$, then

$$\tau_1 \frac{dx_1}{d\theta}|_{\theta=0} + \tau_2 \frac{dx_2}{d\theta}|_{\theta=0} = \tau_1 \left(\frac{d(x_1 + x_2)}{d\theta}|_{\theta=0} \right)$$

[28]For example, if there are both intensive and extensive labor supply responses, one can compute both a participation elasticity that is weighted by the average tax rates and an intensive elasticity weighted by marginal tax rates (Kleven and Kreiner (2006)). If there are switches between capital and labor income, one can compute the causal impacts on each of these and weight by their respective tax rates.

[29]This approach is taken by Chetty et al. (2013) who show the behavioral responses to the marginal incentives induced by the EITC lead to a 5% increase in government expenditures.

2.6 Relation to MEB

Since Harberger (1964), the MEB framework is arguably the most common welfare framework for the evaluation of the welfare impact of changes to tax and transfer policies.[30] As discussed in the introduction, calculating welfare changes with this approach requires the compensated (Hicksian) elasticity. However, the compensated (Hicksian) elasticity is the causal effect of a policy that holds utility constant; under the present framework, it is therefore not sufficient for calculating the welfare impact of policies that actually change utilities.

This section clarifies the difference between the marginal willingness to pay measure of welfare – which relies on the policy responses – and the marginal excess burden – that is known to rely on compensated effects. I begin by defining the MEB in this environment.

Let $P(\theta)$ be a policy path. To compute the marginal excess burden to individual i from the policy $P(\theta)$, let $\mathbf{v} = (v_i)$ be a set of pre-specified utilities. Most commonly, $\mathbf{v}$ is chosen to be the set of status quo utilities. This corresponds to the "equivalent variation" measure of MEB.[31] Now, define the *compensated* policy path, $P^{\mathbf{v}}(\theta)$, such that $P^v = \left\{ \hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta), \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta), \hat{T}_i(\theta) + \hat{C}_i(\theta; v), \hat{\mathbf{G}}_{\mathbf{i}}(\theta) \right\}_i$ where $\hat{C}_i(\theta; u)$ is a compensation provided to individual i such that $V_i\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}(\theta), \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}(\theta), \hat{T}_i(\theta) + \hat{C}_i(\theta; v), \hat{\mathbf{G}}_{\mathbf{i}}(\theta), y_i \right) = v_i$. Intuitively, $P^v(\theta)$ is the same as the proposed policy path, $P(\theta)$, with the addition of individual specific lump-sum transfers, $\hat{C}_i(\theta; v)$, that hold agent i's utility constant at v_i.

Now, let $\hat{t}_i^v$ denote the net government resources allocated to individual i under the compensated policy $P^v(\theta)$. Following the textbook definitions of Auerbach and Hines (2002), the class of marginal excess burden measures are defined as

$$MEB_i^{v_i} = \frac{d\hat{t}_i^v}{d\theta}|_{\theta=0} \qquad (10)$$

This measures the amount of additional resources the government must give to individual i in order to maintain individual her utility constant at v_i while the policy change is implemented.[32] If the policy

---

[30]For example, Eissa et al. (2008) and Eissa and Hoynes (2011) apply the MEB framework to study the welfare impact of recent expansions of the earned income tax credit in the US.

[31]See Auerbach and Hines (2002). Choosing $\mathbf{v}$ to be the utilities obtained in the hypothetical first-best world with no economic distortions yields the "compensating variation" measure of MEB. Of course, the distinction between CV and EV measures of MEB depend on whether one is starting from the perspective of the first best or from the status quo.

[32]An alternative definition of marginal excess burden is given in the handbook chapter of Auerbach (1985) that preceded Auerbach and Hines (2002). In this chapter, the equivalent variation MEB is defined as the marginal willingness to pay for a hypothetical policy that is the same as the original policy but for which the budget constraint is closed using individual-specific lump-sum taxation. To express this definition of MEB in the present framework, define an augmented policy path

$$P^{85} = \left\{ \left\{ \hat{\tau}_{ij}^l(\theta) \right\}_j, \left\{ \hat{\tau}_{ij}^x(\theta) \right\}_j, \hat{T}_i(\theta) - \hat{t}(\theta), \hat{\mathbf{G}}_{\mathbf{i}}(\theta) \right\}_i$$

where individual is forced to pay for net resources, $\hat{t}_i(\theta)$, provided to her by the policy path. Given this, the equivalent variation MEB from Auerbach (1985) is

$$MEB_i^{85} = \frac{\frac{d\hat{V}_i^{P^{1985}}}{d\theta}|_{\theta=0}}{\lambda_i}$$

which depends on compensated elasticities (since the individual must pay for the resource cost), but it is straightforward to verify that these are not "fully compensated" Hicksian elasticities since the transfers are not guaranteed to hold utility constant.

change is not desirable to individual i, she must be compensated to hold her utility constant (so MEB is positive); conversely if the policy change is good for individual i, the government must take away resources to hold her utility constant (so MEB is negative).

If $\mathbf{v}$ is the status quo set of utilities (i.e. the EV measure), then MEB is related to $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$ through the income effects that were removed to construct the MEB policy experiment. Let $\hat{x}_{ij}^c$ and $\hat{l}_{ij}^c$ denote the compensated choices of goods and labor supply activities under policy path $P^v(\theta)$. Then, the income effect component of the response to the policy on x_{ij} is the difference between the causal and compensated response: $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} - \frac{d\hat{x}_{ij}^c}{d\theta}|_{\theta=0}$. Then, MEB is related to $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$ through the impact of the behavioral response to the compensation on the government budget:

$$MEB_i^{v_i} = \underbrace{\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} - \frac{d\hat{t}_i}{d\theta}}_{\substack{\text{WTP Above} \\ \text{Resource Cost}}} - INC_i \tag{11}$$

where

$$INC_i = \underbrace{\left(\tau_i^x \left(\frac{d\hat{\mathbf{x}}_i}{d\theta}|_{\theta=0} - \frac{d\hat{\mathbf{x}}_i^c}{d\theta}|_{\theta=0} \right) + \tau_i^l \left(\frac{d\hat{\mathbf{l}}_i}{d\theta}|_{\theta=0} - \frac{d\hat{\mathbf{l}}_i^c}{d\theta}|_{\theta=0} \right) \right)}_{\text{Income Effects on Government Budget}}$$

Also, if $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} = 0$, then no marginal compensation is provided to individual i, so that $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} - \frac{d\hat{x}_{ij}^c}{d\theta}|_{\theta=0} = 0$ and $MEB_i^{v_i} = 0$ (and vice-versa).[33]

Equation (11) illustrates that MEB and the marginal willingness to pay are related to each other through the income effects induced by the lump-sum taxation that hold individuals' utility constant in the MEB experiment. In the special case for which there are no income effects, $INC_i = 0$ so that MEB is given by the individual's willingness to pay above the resource costs for the policy. More generally, calculation of the MEB requires removing the portion of the causal effects that are due to the income effects associated with the policy change.

One reason MEB is a mainstay in the welfare analysis toolkit is perhaps because it is a fundamental input into the optimal commodity taxation analysis initiated by Ramsey (1927) and studied in detail in Diamond and Mirrlees (1971). Their results show that, in a model with a representative agent, the marginal excess burdens across commodities are equated. This yields the classic "inverse elasticity" rule for commodity taxation: at the optimum, tax-weighted compensated price derivatives for each commodity are equated. However, as shown in Appendix E, this optimality formula involves compensated responses because a necessary condition for taxes to be at an optimum is that small budget-neutral changes to taxes do not affect utility. Hence, around the optimum, the policy responses are compensated responses (i.e. $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} = \frac{d\hat{x}_{ij}^c}{d\theta}|_{\theta=0}$ because utility is not changing at the

---

[33]If $\mathbf{v}$ is not the status quo utilities, no such relationship is guaranteed between MEB and $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$ because λ_i and $\frac{\partial u_i}{\partial G_{ij}}$ need to be computed in the alternative world for which V_i is the utility level specified in the MEB experiment.

optimum).[34]

2.7 Aggregating to Social Welfare

Heretofore, I have defined measures of the welfare impact of a policy change on an individual, as measured by willingness to pay (Equation 9) and MEB (Equation (10)). Here, I illustrate how these measures aggregate across individuals to measures of social welfare. As noted in Section 2.3, the marginal willingness to pay for a policy change can be aggregated to social welfare using the social marginal utilities of income, η_i.

$$\frac{dW}{d\theta} = \int_i \eta_i \frac{\frac{d\hat{V}_i^P}{d\theta}|_{\theta=0}}{\lambda_i} di \tag{12}$$

As shown in Equation 9, the construction of each term, $\frac{\frac{d\hat{V}_i^P}{d\theta}|_{\theta=0}}{\lambda_i}$, does not require decomposing behavioral responses into underlying components such as income and substitution effects. Of course, one must specify social marginal utilities of income, η_i, in order to provide such aggregation. As is well-known, ratios of social marginal utilities of income, $\frac{\eta_j}{\eta_i}$, measure one's willingness to pay to transfer resources from individual i to individual j and nest many forms of social preferences (Saez and Stantcheva (2013)).

As mentioned above, a potentially more objective method of aggregation of welfare impacts across the income distribution is to weight welfare by the marginal price of transferring resources across the income distribution using modifications to the income tax schedule, as suggested by Hendren (2014). Regardless of how one does the aggregation, the marginal willingness to pay measure of welfare, $\frac{\frac{d\hat{V}_i^P}{d\theta}}{\lambda_i}$, has the useful feature that it aggregates to social welfare using these social marginal utilities of income, η_i.

In contrast, aggregation of MEB across individuals is often more difficult, even for budget neutral policies. Using equation (11), one can express the aggregate impact ed by Diamond and Mirrlees (1971), the aggregate impact on social welfare for a budget neutral policy is given by the sum of the net transfers to individuals , the marginal excess burden, and the income effects:

$$\frac{dW}{d\theta} = \int_i \eta_i \left(\frac{d\hat{t}_i}{d\theta} + MEB_i + INC_i \right) di$$

In other words, if one wishes to use the social marginal utilities to aggregate to social welfare, one needs to add back in the income effects that were removed from the policy response in order to form the compensated response.

The difficulty of aggregating MEB across individuals was arguably well-documented by Diamond and Mirrlees (1971). In the case of a representative agent, the optimality condition for policy to be set optimally is $\frac{dV}{d\theta} = 0$, so that the income effects associated with the policy change are zero at an

---

[34]Moreover, away from an optimum, the causal effects from policies that change commodity taxes continue to provide information on the desirability of changing commodity tax rates. In contrast, compensated elasticities defined not around the optimum will not necessarily provide information about the optimal commodity tax rate, as this would require an assumption that the compensated elasticities are constant.

optimum (see Appendix E for further details). But, with heterogeneous individuals, the marginal welfare formulae no longer depend on compensated responses.[35] This is because small budget-neutral policy changes does not hold the agents' utilities constant at the optimum when there are heterogeneous agents. Some agents are better off; others are worse off. And, the aggregation of MEB across individuals requires knowledge of the income effects associated with the policy change (Diamond and Mirrlees (1971)).[36] In this sense, aggregation of MEB across individuals is arguably more difficult than aggregation of marginal willingness to pay, $\frac{\frac{d\hat{V}_i^P}{d\theta}|_{\theta=0}}{\lambda_i}$, since the latter does not require knowledge of the income effects associated with the policy change.

2.8   An MVPF for non-budget neutral policies

If a policy is budget neutral, then the aggregate impact of the policy on social welfare (Equation (12) and the components in Equation 9) provides a natural measure for the total welfare impact of the policy change. Such a measure depends on the policy responses or policy elasticities, and does not require decomposition of behavioral responses into income and substitution effects, for example.

In reality many policy changes are not budget neutral, at least in the short run (i.e. $\int_{i \in I} \frac{d\hat{t}_i^P}{d\theta} di \neq 0$). For such policies, one needs to adjust the welfare framework to account for the total cost of the policy. Individuals will generally be willing to pay to obtain a transfer.

In this section, I show that the framework provides a natural justification for computing a simple marginal value of public funds (MVPF) as the welfare cost per dollar of government budget expended. I use the term "MVPF" to be distinct from the traditional definition of the MCPF which generally does not rely on the causal effect of the policy in question, as discussed in Section 2.9.

Normalizing social welfare into units of individual $\hat{i}$'s income, the MVPF is given by:

$$MVPF_P^{\hat{i}} = \frac{\int_i \frac{\eta_i}{\eta_{\hat{i}}} \frac{\frac{d\hat{V}_i^P}{d\theta}|_{\theta=0}}{\lambda_i} di}{\int_{i \in I} \frac{d\hat{t}_i^P}{d\theta} di} = \frac{\text{"Benefit"}}{\text{"Cost"}} \tag{13}$$

which is the sum of the welfare impact on each individual, $\frac{\frac{\partial \hat{V}_{\hat{i}}^P}{\partial \theta}|_{\theta=0}}{\lambda_{\hat{i}}}$, weighted by their social marginal utilities of income, η_i, and normalized in units of dollars to individual $\hat{i}$.[37] There are many different ways of constructing welfare measures for non-budget neutral policies (Fullerton (1991); Auerbach and Hines (2002); Dahlby (2008)). This particular definition of the welfare impact of non-budget neutral policies in equation (13) was initially proposed by Mayshar (1990) and also by Slemrod and Yitzhaki (1996, 2001) and Kleven and Kreiner (2006).

The MVPF has two key advantages. First, the needed behavioral responses depend solely on the causal, not compensated, effects of the non-budget neutral policies in question. This contrasts with

---

[35]See Section VII, page 268 of Diamond and Mirrlees (1971).

[36]See Auerbach and Hines (2002) for a simple derivation of this on page 1370, equation 3.24.

[37]The $\hat{i}$ notation denotes the units of income used in the definition; it is not the welfare impact on type $\hat{i}$. It is the welfare impact on all types measured in units of $\hat{i}$'s income.

both the MEB and also the traditional definition of the marginal cost of public funds, as discussed below in Section 2.9.

Second, comparisons of MVPFs across policies correspond to comparisons of the social welfare impacts of policies. Given any two non-budget neutral policies, P_{Tax} and P_{Exp}, let $P(\theta)$ denote a policy that increases spending by \$$\theta$ on P_{Exp} financed by a decrease in spending by \$$\theta$ on P_{Tax}. Then, the welfare impact of this combined policy is given by:

$$\frac{d\hat{W}^P}{d\theta} = \eta^{\hat{i}} \left(MVPF^{\hat{i}}_{P_{Exp}} - MVPF^{\hat{i}}_{P_{Tax}} \right) \tag{14}$$

so that policy P_{Exp} provides a benefit of $MVPF^{\hat{i}}_{P_{Exp}}$ per dollar of government revenue and a cost of $MVPF^{\hat{i}}_{P_{Tax}}$ per dollar of government revenue. If $MVPF^{\hat{i}}_{P_{Exp}}$ is greater (less) than $MVPF^{\hat{i}}_{P_{Tax}}$, then taking resources from the tax (expenditure) policy and using it to finance the expenditure (tax) policy will improve social welfare. Identifying heterogeneity in the MVPF across different policies is equivalent to identifying welfare-improving budget neutral policies.

2.9 Relation to the MCPF

As mentioned, the definition of the MVPF is based on welfare measures of Mayshar (1990) and Slemrod and Yitzhaki (1996, 2001); Kleven and Kreiner (2006). However, it differs from the canonical definition of the marginal cost of public funds. Following the papers of Stiglitz and Dasgupta (1971) and Atkinson and Stern (1974), this literature seeks a number that can be used to adjust the standard Samuelson (1954) condition for the welfare cost of raising the resources to finance the public expenditure (Ballard and Fullerton (1992)). This definition is the impact of the behavioral response to the policy on the government's budget of both increasing taxes and spending resources on the public good (i.e. the final term in Equation 9 for a policy that raises taxes and increases spending on G).

Suppose P_{Exp} is non-budget neutral a policy that increases spending on a public good G. The MCPF is intended to provide an adjustment to the welfare measurement of the value of the public spending to account for the distortionary cost of raising tax revenue.

For simplicity, consider a stylized environment with one individual and a linear tax labor supply, l. Now, let P be a budget-neutral policy that simultaneously increases spending on G but that is financed by an increase in taxation on labor supply. The aggregate welfare impact is given by

$$\frac{d\hat{V}^P}{d\theta} = \underbrace{\left(\frac{\frac{\partial u}{\partial G}}{\lambda} - c^G \right) \frac{d\hat{G}}{d\theta}}_{\text{Samuelson Condition}} + \underbrace{\tau \frac{d\hat{l}^P}{d\theta}}_{\text{"MCPF"}}$$

The first term is the canonical Samuelson condition for the provision of public goods, which measures the difference between the willingness to pay and the per-person cost for providing the public good. The second term is generally referred to as the marginal cost of public funds (Fullerton (1991),Dahlby

(2008)).[38] It adjusts the Samuelson condition to account for the distortionary cost of raising tax revenue.

By construction, the MCPF does not depend on the policy response to the non-budget neutral policy in question, P_{Exp}. Rather, it depends on the policy response to the augmented policy, P, which also includes a modification to the tax schedule to account for the cost of raising tax revenue. Indeed, this definition of the MCPF in principle varies with how the revenue is spent. In practice, many papers assume that the expenditure has a separable impact on utility and hence does not have a causal impact on labor supply or the distortionary cost of taxation (Ballard and Fullerton (1992)). This is violated in many realistic policy settings, such as job training programs and education more generally, where perhaps a primary motivation for the government expenditures is to generate positive tax impacts from labor supply responses.

A more fundamental difficulty with this method for defining "the" MCPF is that there is an implicit assumption that a single measure of the cost of raising revenue can be applied across different expenditure settings. In practice, revenue can be obtained not only from the tax schedule but also from a reduction in expenditure on alternative public goods and services.

In contrast, the MVPF is not a component of a broader welfare calculation but rather it is the total welfare impact of the policy per unit of government expenditure. By computing the MVPF for a range of policies, the government can improve social welfare by moving resources from policies with low to high MVPF policies, regardless of whether they are "tax" or "expenditure" policies, or combinations of both. More importantly, in contrast to the MCPF or the MEB, the only behavioral responses required for constructing the MVPF are the policy responses to the non-budget neutral policy in question.

2.10   Simplifications to the MVPF

At first glance, the MVPF in equation (13) may seem complicated to estimate in practice. However, a couple of reasonable simplifications allow for a straightforward implementation. First, one may wish to assume that the social marginal utility of income varies across beneficiaries of different policies (e.g. the EITC versus the top marginal income tax rate), but may be willing to assume that the beneficiaries of a given policy change have the same social marginal utility of income (e.g. EITC beneficiaries have similar social marginal utilities of income and those subject to the top marginal income tax rate have similar social marginal utilities of income).

With this assumption, one can define the MVPF in units of the income of the beneficiaries and make comparisons using the MVPF and the social marginal utilities of income. For example, suppose P_1 and P_2 are two policies that affect two different populations with different social marginal utilities of income, η_1 and η_2. For example, P_1 could be an EITC expansion and η_1 is the social marginal utility of income of EITC beneficiaries; P_2 could be a decrease in the top marginal income tax rate and η_2 is the social marginal utility of income of those subject to the top marginal income tax rate.

---

[38]An alternative tradition – the so-called Pigou-Harberger-Browning tradition (Pigou (1947); Harberger (1964); Browning (1976, 1987)) – uses the MEB as the measure of the MCPF.

Then, increasing spending on P_1 financed by a decrease in P_2 increases social welfare if and only if

$$\frac{\eta_1}{\eta_2} \geq \frac{MVPF_{P_2}^2}{MVPF_{P_1}^1} \tag{15}$$

Intuitively, the value $\frac{\eta_1}{\eta_2}$ measures society's willingness to pay to transfer resources from beneficiaries of P_2 to beneficiaries of P_1; the ratio $\frac{MVPF_{P_2}^2}{MVPF_{P_1}^1}$ measures the marginal cost of transferring such resources through a reduction of P_2 and increase of P_1.[39]

Second, in practice many policies only affect either market goods (e.g. taxes and transfers) or publicly provided goods. In the case that the policy changes only market goods, the envelope theorem (shown in the first line in Equation 9) shows that individuals value these policy changes dollar-for-dollar. So, letting θ parameterize the increase in mechanical spending, the MVPF in units of the beneficiaries own income is given by

$$MVPF = \frac{1}{\frac{1}{|I|}\int_{i\in I}\frac{d\hat{t}_i^P}{d\theta}di} = \frac{1}{1+FE}$$

where the numerator, 1, is the average marginal benefit of the policy and $\frac{1}{|I|}\int_{i\in I}\frac{d\hat{t}_i^P}{d\theta}di = 1+FE$ is the average marginal cost of the policy, which incorporates the causal impact of the behavioral response to the policy on the government budget outlays – i.e. the fiscal externality FE. If the behavioral response to policies increase government expenditures, then $FE > 0$. Conversely, if the behavioral response to the policy decreases government expenditures, then $FE < 0$.

More generally, if a policy increases spending on a publicly provided good, G, the MVPF is given by

$$MVPF = \underbrace{\frac{\frac{\partial u}{\partial G}}{\lambda}}_{\text{WTP for }G\text{relative to income}} * \underbrace{\frac{1}{\frac{1}{|I|}\int_{i\in I}\frac{d\hat{t}_i^P}{d\theta}di}}_{\text{Fiscal cost per dollar raised}} \tag{16}$$

Intuitively, the MVPF is the same as for market goods, but now one needs to adjust for the individual's willingness to pay for \$1 in the publicly provided good, G, in units of her income. This is given by the marginal rate of substitution, $\frac{\frac{\partial u}{\partial G}}{\lambda}$. The second term remains the same as the case with market goods: it captures the total cost of providing \$1 of public goods. The key is that this total cost incorporates the causal impact of the behavioral response to the policy on the government budget.

3   Applications

This section draws on the large literature estimating the causal effects of policy changes and places the in this normative framework by calculating their MVPF in equation (16). In particular, I study

---

[39]Equation (15) does not hold if one were to instead use MEB instead of the MVPF. As can be seen in the analysis of Diamond and Mirrlees (1971), such comparisons require using modified social welfare weights that re-incorporate the income effects that were removed for the MEB calculation.

changes involving the top marginal income tax rate, the generosity of the earned income tax credit (EITC), food stamps (SNAP), job training programs (JTPA), and housing vouchers (Section 8). The calculations are presented in Table 1.[40]

Table 1: MVPF for Various Policy Changes

| Policy | Welfare Components Per Dollar of Net Resources | | | |
	WTP for publicly provided good versus income	Fiscal cost per dollar collected	MVPF	Source
Top Tax Rate	N/A	1.33 to 2	1.33 to 2	Saez, Slemrod, and Giertz (2012); Giertz (2009)
EITC Expansion	N/A	0.88	0.88	Hotz and Scholz (2003), Chetty et al. (2013), Appendix A Calculations
Food Stamps	0.8 to 0	0.66[1]	0.53 to 0.66[1]	Whitmore (2002); Hoynes and Schanzenbach (2012), Appendix A Calculations
Job Training	0-1.22	1.52	0 to 1.85	Bloom et al (1997), Appendix A Calculations
Housing Vouchers	0.83	0.95	0.79	Reeder (1985); Jacob and Ludwig (2012); Appendix A Calculations

[1]Estimate for labor supply response in Hoynes and Schanzenbach (2012) is not significantly different from zero, so that the MVPF is statistically indistinguishable from 0.8 to 1, depending on the valuation of the public spending.

3.1 Top Tax Rate Increase

There is a large literature estimating the causal effect of changes to the top marginal income tax rate and the impact of such behavioral responses on the government's budget (see Saez et al. (2012) for a recent review). Such estimates generally come from variation induced by two policy reforms: the Omnibus Budget Reconciliation Act of 1993 (a.k.a. OBRA-93 or the Clinton tax increases) and the Tax Reform Act of 1986 (a.k.a. TRA-86 or the Reagan tax cuts).

Translating the causal estimates from the literature into impacts on the government's budget requires a couple of assumptions. First, I assume that the policy has no spillover effects, so that the response to the top marginal income tax rate is zero amongst those whose earnings are below $\bar{l}$. This is commonly assumed in existing literature (e.g. Feldstein (1999)), as lower income groups are used as controls for macroeconomic effects argued to be unrelated to the tax policy. Of course, this assumption could be relaxed if one had an estimate of the causal effect of the policy on taxable behavior of those earning below the top income tax threshold.

Second, I assume that the rich have no income shifting across tax bases with different nonzero tax rates. This rules out the program having an impact on capital gains, for example. Again, this

---

[40]The literature studying the causal effects of these policies focuses on a wide range of outcomes including many aspects of taxable behavior. However, in most cases, these studies do not construct a comprehensive measure of the revenue impact of the behavioral responses. Therefore, I construct such a measure using the causal effects on various components of taxable behavior. These details are discussed in Appendix F.

assumption could be relaxed with additional empirical work estimating the causal effect of raising the top income tax rate on tax revenue from capital gains.

With these assumptions, the MVPF of raising revenue from the rich through an increase in the top marginal tax rate is given by

$$MVPF_{P^{Tax}}^{Rich} = \frac{1}{1 + FE^{top}}$$

where FE^{top} is the fraction of mechanical ordinary income tax revenue lost from behavioral responses to the tax increase,

$$FE^{top} = \frac{\int_{i \in Rich} \tau_i^l \frac{d\hat{l}_i^{Tax}}{d\theta}|_{\theta=0} di}{\int_{i \in Rich} \frac{d\hat{\tau}_{Rich}^{Tax}}{d\theta}|_{\theta=0} \left(\hat{l}_i^{Tax} - \bar{l} \right) di} \tag{17}$$

Here, $\hat{l}_i$ is the taxable ordinary income of the rich and $\frac{d\hat{l}_i^{Tax}}{d\theta}|_{\theta=0}$ is the response of taxable ordinary income to a policy that raises the top marginal tax rate and uses the finances to raise government revenue.[41] Note $FE^{top} < 0$ if behavioral responses lower tax revenue.

Fortunately, there is a large literature focused on estimating FE in equation (17). Generally, this parameter is referred to as the "marginal excess burden" of the change in the top tax rate (Mirrlees et al. (2011); Saez et al. (2012)). However, as noted in earlier handbook chapters (Auerbach (1985); Auerbach and Hines (2002)), such an interpretation is technically incorrect in the presence of income effects. Such a point was raised by Feldstein (2012) in his critique of the ambitious and widely-celebrated Mirrlees Review (Mirrlees et al. (2011)) analyzing British tax policy:

> *Despite the centrality of the concept of excess burden, the Mirrlees Review fails to provide a clear explanation that the excess burden is the difference between the loss to taxpayers caused by the tax (e.g., the amount that taxpayers would have to receive as a lump sum to be as well off as they were before the imposition of the tax) and the revenue collected by the government. There are instead several alternative definitions at different points in the text, some of which are vague and some of which are simply wrong. For example, the Mirrlees Review states "it is the size of this revenue loss that determines the 'excess burden' of taxation" (61). That is not correct since the excess burden depends only on the substitution effects while revenue depends also on the income effects. (Feldstein (2012))*

Because this literature generally does not remove income effects, Feldstein (2012) is technically correct that it is not estimating the marginal excess burden. However, taking FE to be the causal impact of the policy (without removing income effects), one obtains precisely the desired parameter for calculating

---

[41]To see this, note that

$$\frac{\frac{\partial \hat{V}_{P^{Tax}}^{Rich}}{\partial \theta}|_{\theta=0}}{\lambda^{Rich}} = -\frac{\int_{i \in Rich} \frac{d\tau^{Rich}}{d\theta}|_{\theta=0} \left(\hat{l}_i - \bar{l} \right) di}{\int_{i \in Rich} \tau_i^l \frac{d\hat{l}_i^{Tax}}{d\theta}|_{\theta=0} di + \int_i \left(\sum_j \tau_{ij}^x \frac{d\hat{x}_{ij}^{Tax}}{d\theta}|_{\theta=0} + \tau_{ij}^l \frac{d\hat{l}_{ij}^{Tax}}{d\theta}|_{\theta=0} \right) di}$$

$$= \frac{1}{1 + FE^{top}}$$

the MVPF in the present framework.[42]

While there is a wide range of existing causal estimates, Saez et al. (2012) and Giertz (2009) suggest mid-points for FE^{top} ranging from 25-50% – i.e., roughly 25-50% of the mechanical revenue that is raised from the tax is lost due to behavioral distortions. This implies $MVPF_{PTax}^{Rich}$ is between 1.33 and 2, as reported in the first row of Table 1.[43]

3.2   EITC Expansion

There is a large literature estimating the causal effects of EITC expansions, especially impacts on single mothers. Unfortunately, there is no study that estimates the impact of the behavioral response to EITC expansions on government expenditures directly. Here, I construct such a causal estimate by taking the causal impacts on earnings and labor supply estimated in previous literature.

To do so, I make several assumptions commonly made in the empirical literature. First, I assume the policy has no effect on groups ineligible for the expansion. This assumes no response amongst individuals not receiving EITC after the expansion and no impact on the choice to become single mothers to become EITC eligible. Support for this latter assumption is found in Hotz and Scholz (2003) who summarize the empirical literature as finding little or no effects on marriage and family formation. Both of these assumptions could easily be relaxed with precise estimates of the fiscal impact of the behavioral responses of these groups to EITC expansions.

For EITC eligibles, I assume that the only behavioral impact of the program that affects tax revenue is through ordinary taxable (labor) income. Although capital income is less of an issue for EITC recipients, this assumption also rules out fiscal externalities of the EITC expansion on other social program take-up, such as SSDI or food stamps. Such impacts are likely to be present, as significant earnings generally disqualifies eligibility for such programs. To the extent to which an EITC expansion crowds out take-up other government services, the analysis will underestimate the social desirability of expanding the EITC.

With these assumptions, one obtains an expression analogous to the change in the top income tax policy:

$$MVPF_{PEITC}^{Poor} = \frac{1}{1 + FE^{EITC}}$$

where FE^{EITC} is the fraction of the mechanical revenue distributed that is increased due to behavioral

---

[42]Saez et al. (2012) show also how FE^{top} can be incorporated into the calculation of the optimal top income tax rate. However, the optimal top tax rate depends on r defined locally *around the optimum*; hence one must assume that FE is constant as the tax rate changes towards the optimum. In contrast, estimating $MVPF_{PTax}^{Rich}$ relies on local estimates of FE^{top} for variation in taxes around the status quo.

[43]There are many caveats to this calculation. For example, it assumes all of the reduction in taxable income is a social cost; in practice some avoidance behavior that reduces taxable income might be socially beneficial (e.g. if people increase charitable giving in response to higher tax rates) or even privately beneficial if individuals are not optimally choosing their degree of avoidance as in Chetty (2009a). If higher tax rates increase charitable giving or causes other positive spillovers, then the MVPF will be lower to the extent to which society values these induced transfers. Conversely, if tax rate increases cause negative spillovers or "trickle-down" general equilibrium effects as in the model of Scheuer and Rothschild (2013), the MVPF will be higher.

distortions,

$$FE^{EITC} \quad = \quad \frac{\int_{i \in Poor} \tau_i^l \frac{d\hat{l}_i^{EITC}}{d\theta}|_{\theta=0} di}{\int_{i \in Poor} \left(\frac{d\hat{T}_i^{EITC}}{d\theta}|_{\theta=0} + \frac{d\hat{\tau}_{ij}^{EITC}}{d\theta}|_{\theta=0} l_i \right) di}$$

There is a large literature focused on estimating the causal effects of EITC expansions on taxable behavior, such as labor supply. For my purposes, these studies would have ideally looked at the impact on tax revenue/expenditure in order to form an aggregate estimate of FE^{EITC}. Without this estimate, I instead take estimates of the extensive and intensive margin labor supply response to the EITC to construct an estimate of FE^{EITC}.

In Appendix F.1, I generalize the model to allow for extensive margin (i.e. discontinuous) responses in labor supply. I show that estimates of causal effects summarized in Hotz and Scholz (2003) suggest that the cost of EITC expansions are ~9% larger than their mechanical cost due to extensive margin behavioral responses.

In addition to extensive margin responses, recent literature has also found evidence that the EITC induces distortions on the intensive margin as well but that these effects may take a while to fully be realized. Using variation in knowledge about the marginal tax rates induced by the EITC, Chetty et al. (2013) estimate that the cost of the EITC program is 5% higher due to behavioral responses. If responses to the marginal incentives primarily govern intensive margin responses and extensive margins are primarily about the average EITC rebate, this suggests these two estimates can be summed so that the total cost of EITC expansions are 14% higher due to behavioral responses. The estimate of $FE^{EITC} = 14\%$ suggests that raising \$1 in general government revenue through a reduction in EITC spending would only require a reduction in benefits of 1/1.14=\$0.88. Hence, the marginal value of public funds of the EITC policy is roughly $MVPF_{PEITC}^{Poor} = 0.88$.[44]

3.3 Food Stamps (SNAP)

The Supplemental Nutrition Assistance Program (a.k.a. SNAP or "food stamps") provides financial assistance to low-income households for the purchase of food and is one of the largest transfer programs in the United States. Indeed, 1 in 4 children received benefits in July 2011 (Ganong and Liebman (2013)). Despite the program's size, there have been relatively few quasi-experimental studies analyzing its impact on behavior, perhaps due to the fact that it is a national program (Hoynes and Schanzenbach (2012)). A notable exception is the recent work of Hoynes and Schanzenbach (2012)

---

[44]The calculation of $MVPF_{PEITC}^{Poor}$ ignored the potential of general equilibrium effects of the policy expansion. This is easily incorporated if one has estimates of the causal impact of the policy on prices. For example, recent research suggests beneficiaries may only capture \$0.73 for every dollar of EITC spending (Rothstein (2010)). This suggests the marginal value of the program should be a weighted average of 0.73 for the beneficiaries and .27 for those who benefit from the reduced wages. To the extent to which the reduction in wages increases firm profits, one would then wish to add this fiscal externality into the benefit of the EITC expansion. But, to the extent to which those benefiting from the wage reduction have lower social marginal utilities of income, this will reduce the MVPF for the MVPF for the EITC expansion. For example, in the extreme case where society had zero value for the beneficiaries of the wage reduction and the wage change did not induce any additional fiscal externality, the MVPF would be 0.88 - 0.27 = 0.51.

who exploit county-level variation in the introduction of food stamps in the 1960s and 70s to estimate its impacts on labor supply. They find significant but noisy reductions in intensive labor supply and large but noisily estimated reductions in labor earnings (with a 95% confidence interval that includes zero impact). Appendix F.2 translates their point estimates into an estimate of the implied fiscal externality. The results suggest that although the program cost was roughly $1,153.25 per household, the reductions in labor supply led to a reduced state and federal income tax receipt of $588. This implies that the total cost to the government is roughly $1,153.25 + $588 = $1,741.25.[45]

Because the program benefits must be used to purchase food, these benefits may not be valued dollar-for-dollar by beneficiaries. However, many studies have shown that in general food stamp receipt does not significantly alter purchase decisions in the U.S. (Smeeding (1982)) and Puerto Rico (Moffitt (1989)). However, Whitmore (2002) uses a slightly different modeling approach and experimental data to arrive at an estimate of roughly 0.8 for some subsamples of recipients. The former studies suggest the transfer of $1,153.25 in food stamps is valued dollar-for-dollar by beneficiaries, whereas the latter suggests they are valued at only $922.60 by beneficiaries. Dividing by the net cost to the government yields an MVPF for food stamps ranging from 0.53 to 0.64, depending on whether one takes the Smeeding (1982) estimate versus the Whitmore (2002) estimate for the cash-value of the food stamps to beneficiaries.

There are a couple of important caveats to keep in mind in interpreting these results. First, the estimated behavioral responses correspond to a 1970s world with very different tax rates and extrapolation to present day may be problematic. Indeed, the presence of the EITC changes the government's stake in labor force participation and labor earnings. Second, food stamps may have significant benefits on children that are not perfectly incorporated into the utility function of the parents. For example, Almond et al. (2012) estimate that food stamps led to improvements in birth outcomes, such as increased birth weight. These potential externalities on newborn children are of course not captured in the current MVPF calculation. If one wished to add such effects, one could take the causal effects from Almond et al. (2012) and multiply by the valuation of the externality along the lines discussed in Appendix C.[46]

3.4 Job Training

Job training programs have the potential to generate positive fiscal externalities by increasing labor supply and taxable earnings. For example, the Job Training Partnership Act (JTPA) of 1982 provided job training to economically disadvantaged youth and adults with the attempt of promoting entry into the labor force. Bloom et al. (1997) reports results from a randomized controlled trial of the program, focusing not only on earnings impacts but also on budget-relevant variables such as welfare and tax receipt. For brevity, I focus here on their results for adult women presented at the top of Table 8 on p573 of Bloom et al. (1997).

---

[45]One should be cautious in interpreting the exact magnitude of this coefficient since it is not statistically significantly different from zero.

[46]For example, one could translate the increase in birth weight to an implied increase in quality-adjusted life years (QALY) and multiply by an assumed value of a QALY.

The program increased earnings on adult women by \$1,683, which led to an increased tax collection of \$236 per enrollee and also a \$235 reduction in welfare expenditures (AFDC). Assuming these capture the behavioral response impact on the budget, the total impact is \$471. The mechanical cost to the government of providing the program to an adult female enrollee is \$1,381.[47] Adding back the fiscal externality of \$471 suggests that the net cost to the government is \$910. This means that $\frac{1}{1+FE} = 1.52$.[48] If individuals valued the job training at its cost (i.e. each individual who had obtained the job training would have been willing to pay \$1,381 for it), this would suggest the MVPF is 1.52. However, if $\frac{\frac{\partial u}{\partial G}}{\lambda} \neq 1$ so that individuals don't value \$1 spent on training as \$1 in cash, then one must adjust the MVPF accordingly.

Calculating the extent to which the individual would be willing to pay for the job training program in excess of its mechanical cost is difficult. In their cost-benefit analysis, Bloom et al. (1997) implicitly assume that the earnings increase of the beneficiaries is a welfare benefit. For this to be the case, one needs to assume that this earnings increase was the result of a positive externality imposed on the beneficiaries (e.g. an increase in their productivity that was incurred with no cost to the beneficiary) and not the result of their increased labor effort. Under this assumption, the individuals willingness to pay for the program in excess of its cost is \$1,683-\$1,381 = \$302 per enrollee. This implies a MVPF of \$1,683/\$910 = 1.85 on the job training beneficiaries.

However, the envelope theorem suggests caution in this calculation. In the canonical model with no distortions besides government intervention, people who are induced into the labor force were, to first order, indifferent to working. Hence, the increase in earnings provides no direct evidence on the willingness to pay for the program. One potential assumption one can make is that the government has no comparative advantage or disadvantage in the provision of job training (e.g. there's an equally good private training program that can also provide similar job training for \$1,381). In this case, the MVPF would be \$1,381/\$910=1.52. A more pessimistic assumption is that the enrollees had virtually no value for the program ($\frac{\partial u}{\partial G} = 0$), so that the MVPF is zero. Indeed, revealed preference of the participants in the JTPA program only guarantees that the willingness to pay was nonnegative for beneficiaries: $\frac{u_G}{\lambda} \geq 0$. An important direction for future work is in characterizing the extent to which government job training programs are valued above or below their cost by beneficiaries.[49]

---

[47]\$1,227 of this is the administrative cost; \$154 is the cost of a wage subsidy associated with the program.

[48]The report also indicates women reduced their spending on private training programs by \$56 and considers this a benefit of the program. But by the envelope theorem, such crowd-out estimates are not welfare relevant to first order.

[49]An important caveat to this MVPF calculation is that it does not account for general equilibrium effects of such policies. For example, Crepon et al. (2012) find evidence that a job placement program in France had an increase in employment among beneficiaries but was offset by a decrease in employment by non-beneficiaries. In this case, even if the beneficiaries had a willingness to pay for the program, it might be perfectly offset by negative impacts on those not enrolled in the program. Incorporating such general equilibrium effects would reduce the estimate of the MVPF, as illustrated in Appendix C. For example, if the program was simply causing sorting within a fixed labor market and non-beneficiaries have equal social marginal utilities of income to beneficiaries of the program, one would find a MVPF of zero.

3.5 Housing Vouchers (Section 8)

With roughly 2 million beneficiaries, Section 8 housing vouchers constitute the largest low-income housing program in the US (Rice and Sard (2006)). Such vouchers provide rent assistance to low-income households. Because they are means-tested, they induce potentially significant labor earnings distortions.

Often, voucher applications exceed supply, which leads to rationing. Jacob and Ludwig (2012) exploit the randomness in the allocation of vouchers to excess applications in Chicago. They analyze not only the impact of voucher receipt on labor supply, but also on other governmental program participation such as TANF (cash welfare), Medicaid, and SNAP. They find housing vouchers lead to a significant reductions in labor supply – on both the intensive and extensive margin – and an increase in participation in other welfare programs.

Appendix F.3 translates these calculations into the impact of the behavioral response to the policy on the government's budget. Although Jacob and Ludwig (2012) estimate significant behavioral responses to housing vouchers, they are arguably modest relative to the cost of the voucher and correspond to \$432 per \$8,400 voucher, so that the fiscal externality comprises roughly 5% of the total cost of the program, or $\frac{1}{1+FE} = 0.95$.

There is little work studying the willingness to pay for the housing voucher. To my knowledge, the only known study is Reeder (1985) who estimates \$1 in Section 8 vouchers are valued at \$0.83 by the beneficiaries (see also the discussion in Jacob and Ludwig (2012)). Taking this estimate, one arrives at an MVPF of $0.83\frac{1}{1+0.05} = 0.78$. For every dollar of cost to the government for the vouchers, individuals would be willing to pay \$0.78.

There are several caveats to this estimate that are worth mentioning. First, Jacob and Ludwig (2012) provides novel identification of the impact of voucher receipt on labor supply and public program take-up; but, they do not estimate the impact of increased voucher generosity on ex-ante labor supply decisions. Indeed, people may decrease their labor supply to become eligible in the first place, which further increases the fiscal cost. Second, there is some recent evidence that suggests a portion of housing vouchers (e.g. \$0.13) may be captured by landlords instead of tenants (Collinson and Ganong (2013)). If the social marginal utilities of landlords and tenants were the same, this would not affect the MVPF; but if landlords have lower social marginal utilities of income, one would need to adjust for the fact that 13% of the mechanical subsidy falls in the hands of landlords and weight the MVPF accordingly.

Finally, the calculation of the MVPF would include all behavioral responses impacting the government budget. This includes other responses by the individuals or other parties. For example, recent analysis of the Moving to Opportunity (MTO) program suggests that young children whose families were required to move to lower-poverty census tracts in order to obtain a section 8 voucher moved to better neighborhoods earn more as adults and pay more taxes, significantly reducing the effective cost of the voucher (Chetty et al. (2015)). An appropriate calculation of the welfare impact of MTO would require accounting for the impact of this tax revenue in the denominator of the MVPF. While the fiscal externality logic hopefully clarifies the estimand required for welfare, this example highlights the many

difficulties obtained in practice; for example, if policies have impacts on childrens' lifetime earnings, then one needs to sum up the present discounted value of the associated tax revenue to calculate the total cost to the government of these policies. For the present purposes, the MVPF calculated here should be seen as the welfare impact of vouchers in settings where they do not have impacts on their childrens' earnings.

3.6 Combining Policies Using Okun's Bucket Experiment

While the previous sections constructed estimates of the MVPF for each of these policies, the beneficiaries for each policy are different. Hence, the MVPF are not directly comparable across policies. To make such comparisons, one needs to aggregate using the social marginal utilities of income.

To illustrate this, consider a budget neutral policy of increasing EITC generosity financed by an increase in the top marginal income tax rate.[50] Following equation (15), let η^{Rich} and η^{Poor} denote the social marginal utilities of income on the rich with incomes above \$400K who are subject to the top tax rate under the status quo and the poor single mothers earning less than \$40K who are eligible for the EITC, which I assume to be constant within each group for simplicity.

If one takes the upper range estimate of $MVPF_{P_{Tax}}^{Rich} = 2$, additional redistribution is desirable iff

$$0.88 - 2\frac{\eta^{Rich}}{\eta^{Poor}} \geq 0$$

or

$$\eta^{Rich} \leq 0.44\eta^{Poor}$$

Additional redistribution is desirable if and only if one prefers \$0.44 in the pocket of an EITC recipient relative to \$1 in the pocket of an individual subject to the top marginal tax rate (i.e. with income above ~\$400K). Similarly, if one takes the lower estimate of $MVPF_{P_{Tax}}^{Rich} = 1.33$, additional redistribution is desirable if and only if one prefers \$0.66 in the pocket of an EITC recipient relative to \$1 to someone subject to the top marginal tax rate.

Ratios of the other MVPF reveal other implicit ratios of social marginal utilities of income that rationalize the status quo as optimal. For example, the EITC MVPF of 0.88 and the housing voucher MVPF of 0.78 suggest society values money in the hands of Section 8 voucher-holders more than EITC beneficiaries. From a more normative perspective, if housing voucher beneficiaries have equal social marginal utilities of income as EITC beneficiaries, then the estimates suggest social welfare would be improved by increasing EITC funding financed by a decrease in Section 8 housing vouchers.

---

[50]This exercise is similar to Browning and Johnson (1984) who simulate the marginal reduction in resources from an increased demogrant at the bottom of the income distribution. For their baseline simulation, additional redistribution is desirable if one prefers \$0.29 to the poor relative to \$1 to the rich. Because Browning and Johnson (1984) simulate the causal impacts of the redistributive policy, the desirability of pursuing the policy depends on the social marginal utilities of income, and hence have an interpretation in terms of Okun's bucket (Okun (1975)). In contrast, if one were to take the MEB estimates for increasing tax rates from Browning (1987), one would need to add back in the income effects before interpreting the results using the social marginal utilities of income.

In more recent work, Immervoll et al. (2007, 2011) also consider a "critical values" approach to additional redistribution that solves for what the social welfare weights must be in order to prefer a policy change.

4 Conclusion

This paper illustrates how one can conduct welfare analysis of policy changes using the causal effects of those policy changes – no decomposition into income and substitution effects is required. In the broad class of models in which the government is the only distortion, the causal impact of the behavioral response to the policy on the government budget – i.e. the "fiscal externality" – is sufficient for all behavioral responses. For non-budget neutral policies, the framework motivates the use of a simple benefit/cost ratio as a measure of the marginal value of public funds (MVPF): individuals' willingness to pay for the policy change per unit government revenue. In short, one can use causal effects to conduct welfare analysis by focusing on the implications of those causal effects for government revenue. With the surge of methods to estimate causal effects, calculating these fiscal externalities and implied MVPFs holds the potential to generate a volume of welfare measurements of government policy changes.

References

Almond, D., H. Hoynes, and D. W. Schanzenbach (2012). Inside the war on poverty: The impact of food stamps on birth outcomes. *The Review of Economics and Statistics 93*(2), 387–403. 3.3

Angrist, J. and J.-S. Pischke (2008). *Mostly Harmless Econometrics*. Princeton University Press. 2.2, 2.2

Athreya, K., D. Reilly, and N. Simpson (2010). Earned income tax credit recipients: Income, marginal tax rates, wealth, and credit constraints. *Federal Reserve Bank of Richmond Economic Quarterly 96*, 229–258. F.1

Atkinson, A. B. and N. H. Stern (1974, January). Pigou, taxation and public goods. *The Review of Economic Studies 41*(1), 119–128. 1, 2.9

Auerbach, A. (1985). The theory of excess burden and optimal taxation. Handbook of public economics, Elsevier. 8, 32, 3.1

Auerbach, A. J. and J. R. Hines (2002). Chapter 21 taxation and economic efficiency. In *Handbook of Public Economics*, Volume Volume 3, pp. 1347–1421. Elsevier. 7, 2.4, 2.6, 31, 32, 2.8, 36, 3.1

Ballard, C. L. and D. Fullerton (1992). Distortionary taxes and the provision of public goods. *Journal of Economic Perspectives 6*(3), 117–131. 6, 1, 2.9

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. 3.4

Bovenberg, A. L. and R. A. de Mooji (1994). Environmental levies and distortionary taxation. *American Economic Review 84*(4), 1085–1089. C

Browning, E. (1976). The marginal cost of public funds. *Journal of Political Economy 84*, 283–98. 38

Browning, E. (1987). On the marginal welfare cost of taxation. *American Economic Review 77*, 11–23. 38, 50

Browning, E. and W. R. Johnson (1984). The trade-off between equality and efficiency. *The Journal of Political Economy 92*(2), 175–203. 50

Chetty, R. (2009a). Is the taxable income elasticity sufficient to calculate deadweight loss? the implications of evasion and avoidance. *American Economic Journal: Economic Policy 1*(2), 31–52. 1, 43, C

Chetty, R. (2009b). Sufficient statistics for welfare analysis: A bridge between structural and Reduced-Form methods. *Annual Review of Economics 1*(1), 451–488. 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).* 29, 3.2, F.1, F.1, 57

Chetty, R., N. Hendren, and L. F. Katz (2015). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. Technical report, National Bureau of Economic Research. 3.5

Collinson, R. and P. Ganong (2013). Incidence and price discrimination. *Working Paper.* 3.5

Crepon, B., M. Gurgand, R. Rathelot, and P. Zamora (2012). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *Working Paper.* 49

Dahlby, B. (2008). *The Marginal Cost of Public Funds.* The MIT Press. 2.8, 2.9

Diamond, P. A. and J. A. Mirrlees (1971, June). Optimal taxation and public production II: tax rules. *The American Economic Review 61*(3), 261–278. 7, 1, 2.6, 2.7, 35, 39, E, E, E

Eissa, N. and H. Hoynes (2011). Redistribution and tax expenditures: The earned income tax credit. *National Tax Journal 64*, 689–730. 4, 1, 18, 30, F.1

Eissa, N., H. J. Kleven, and C. T. Kreiner (2008). Evaluation of four tax reforms in the united states: Labor supply and welfare effects for single mothers. *Journal of Public Economics 92*(3-4), 795–816. 4, 1, 18, 30

Eissa, N. and J. Liebman (1996). Labor supply responses to the earned income tax credit. *The Quarterly Journal of Economics 111*(2), 605–637. F.1

Feldstein, M. (1999). Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics 81*(4), 674–680. 1, 1, 2.5, 3.1

Feldstein, M. (2012, September). The mirrlees review. *Journal of Economic Literature 50*(3), 781–790. 1, 3.1

Fullerton, D. (1991). Reconciling recent estimates of the marginal welfare cost of taxation. *American Economic Review 81*, 302–08. 2.8, 2.9

Ganong, P. and J. Liebman (2013). The decline, rebound, and further rise in snap enrollment: Disentangling business cycle fluctuations and policy changes. *NBER Working Paper No. 19363*. 3.3

Giertz, S. H. (2009). The elasticity of taxable income: Influences on economic efficiency and tax revenues, and implications for tax policy. In A. Viard (Ed.), *Tax Policy Lessons from the 2000s*, pp. 101–136. AEI Press. 1, 3.1

Goolsbee, A. (1999, January). Evidence on the High-Income laffer curve from six decades of tax reform. *Brookings Papers on Economic Activity 1999*(2), 1–64. 1

Goulder, L. H. (1995). Environmental taxation and the double dividend: A reader's guide. *International Tax and Public Finance 2*(2), 157–183. C

Harberger, A. (1964). The measurement of waste. *American Economic Review 54*, 58–76. 1, 2.6, 38

Hendren, N. (2014). The inequality deflator: Interpersonal comparisions without a social welfare function. *Working Paper*. 2.3, 2.7

Hicks, J. R. (1939). The foundations of welfare economics. *Economic Journal 49*, 696–712. 2.3

Holahan, J. and S. McMorrow (2012). Medicare, medicaid, and the deficit debate. Technical report, Urban Institute. F.3

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. 1, 3.2, F.1, 56

Hoynes, H. W. and D. W. Schanzenbach (2009). Consumption responses to in-kind transfers: Evidence from the introduction of the food stamp program. *American Economic Journal: Economic Policy 1*(4), 109–139. F.2

Hoynes, H. W. and D. W. Schanzenbach (2012). Work incentives and the food stamp program. *Journal of Public Economics 96*(1-2), 151–162. 3.3, F.2

Immervoll, H., H. J. Kleven, C. T. Kreiner, and E. Saez (2007). Welfare reform in european countries: A microsimulation analysis. *The Economic Journal 117*, 1–44. 50

Immervoll, H., H. J. Kleven, C. T. Kreiner, and N. Verdelin (2011). Optimal tax and transfer programs for couples with extensive labor supply responses. *Journal of Public Economics 95*, 1485–1500. 50

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. 3.5, F.3

Kaldor, N. (1939). Welfare propositions of economics and interpersonal comparisons of utility. *The Economic Journal 49*(195), 549–552. 2.3

Kleven, H. J. and C. T. Kreiner (2005). Labor supply behavior and the design of tax and transfer policy. *Danish Journal of Economics 143*, 321–358. 4, 1

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. 1, 2.5, 28, 2.8, 2.9

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995, June). *Microeconomic Theory*. Oxford University Press, USA. 26

Mayshar, J. (1990). On measures of excess burden and their applications. *Journal of Public Economics 43*(3), 263–89. 6, 1, 2, 2.8, 2.9

Meyer, B. D. and D. T. Rosenbaum (2001). Welfare, the earned income tax credit, and the labor supply of single mothers. *The Quarterly Journal of Economics 116*(3), 1063–1114. 57

Mirrlees, J. A., S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles, and J. Poterba (Eds.) (2011). *Tax by Design: the Mirrlees Review*, Chapter Chapter 3: The Taxation of Labour Earnings. Oxford University Press. 1, 3.1

Moffitt, R. (1989). Estimating the value of an in-kind transfer: The case of food stamps. *Econometrica 57*(2), 385–409. 3.3

of Representatives. Committee on Ways, U. H. and Means (2004). 2004 green book: Background material and data on programs within the jurisdiction of the committee on ways and means. Technical report, 111th Congress. F.3

Okun, A. (1975). *Equality and Efficiency*. Brookings Institution Press. 1, 2.3, 50

Parry, I. W. H. (1995). Pollution taxes and revenue recycling. *Journal of Environmental Economics and Management 29*(3), S64–S77. C

Pigou, A. (1947). *A Study in Public Finance*. Macmillian. 38

Piketty, T. and E. Saez (2013). A theory of optimal inheritance taxation. *Econometrica 81*(5), 1851–1886. 1

Ramsey, F. P. (1927, March). A contribution to the theory of taxation. *The Economic Journal 37*(145), 47–61. 2.6, E

Reeder, W. J. (1985). The benefits and costs of the section 8 existing housing program. *Journal of Public Economics 26*(3), 349–377. 3.5

Rice, D. and B. Sard (2006). President's 2007 budget renews same number of housing vouchers funded in 2006. Center on Budget and Policy Priorities. 3.5

Rothstein, J. (2010). Is the eitc as good as an nit? conditional cash transfers and tax incidence. *American Economic Journal: Economic Policy 2*(1), 177–208. 44

Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies 68*, 205–229. 1

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, 1, 2.5, 3.1, 3.1, 42

Saez, E. and S. Stantcheva (2013). Generalized social marginal welfare weights for optimal tax theory. *NBER Working Paper No. 18835*. 2.3, 2.7

Salanie, B. (2003). *The Economics of Taxation*. MIT Press. 52

Samuelson, P. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics 36*, 387–389. 1, 2.5, 2.9

Scheuer, F. and C. Rothschild (2013). Redistributive taxation in the roy model. *Quarterly Journal of Economics 128*, 623–668. 43

Schlee, E. E. (2013). Radner's cost-benefit analysis in the small: An equivalence result. *Working Paper*. 21

Slemrod, J. and S. Yitzhaki (1996). The social cost of taxation and the marginal cost of funds. *International Monetary Fund Staff Papers 43*(1), 172–98. 6, 1, 2, 2.8, 2.9

Slemrod, J. and S. Yitzhaki (2001). Integrating expenditure and tax decisions: The marginal cost of funds and the marginal benefit of projects. *National Tax Journal*. 6, 1, 2, 2.8, 2.9

Smeeding, T. M. (1982). Alternative methods for valuing selected in-kind transfer benefits and measuring their effect on poverty. *Technical Paper No. 50*. 3.3, F.2

Stiglitz, J. E. and P. Dasgupta (1971, April). Differential taxation, public goods, and economic efficiency. *The Review of Economic Studies 38*(2), 151–174. 1, 2.9

Whitmore, D. (2002). What are food stamps worth? *Princeton University Industrial Relations Section Working Paper #468*. 3.3, F.2

Online Appendix: Not for Publication

A Derivation of Willingness to Pay

I first characterize $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$. Taking the total derivative of V_i with respect to θ, I have

$$\frac{d\hat{V}_i}{d\theta} = \frac{dV_i\left(\hat{\tau}_\mathbf{i}^\mathbf{l}, \hat{\tau}_\mathbf{i}^\mathbf{x}, \hat{T}_i, y_i, \hat{G}_i\right)}{d\theta} = \frac{\partial V_i}{\partial T_i}\frac{d\hat{T}_i}{d\theta} + \sum_{j=1}^{J_G}\frac{\partial V_i}{\partial G_i}\frac{d\hat{G}_i}{d\theta} + \sum_{j=1}^{J_X}\frac{\partial V_i}{\partial \tau_{ij}^x}\frac{d\hat{\tau}_{ij}^x}{d\theta} + \sum_{j=1}^{J_L}\frac{\partial V_i}{\partial \tau_{ij}^l}\frac{d\hat{\tau}_{ij}^l}{d\theta}$$

Applying the envelope theorem from the agent's maximization problem and evaluating at $\theta = 0$ implies

$$\frac{\partial V_i}{\partial \tau_{ij}^x} = -x_{ij}\lambda_i$$
$$\frac{\partial V_i}{\partial \tau_{ij}^l} = -l_{ij}\lambda_i$$
$$\frac{\partial V_i}{\partial T_i} = -\lambda_i$$
$$\frac{\partial V_i}{\partial G_i} = \frac{\partial u_i}{\partial G_i}$$

Replacing terms, I have

$$\frac{d\hat{V}_i}{d\theta}|_{\theta=0} = \lambda_i\left(\frac{d\hat{T}_i}{d\theta} + \sum_{j=1}^{J_G}\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i}\frac{d\hat{G}_{ij}}{d\theta} - \sum_{j=1}^{J_X}x_{ij}\frac{d\hat{\tau}_{ij}^x}{d\theta} - \sum_{j=1}^{J_L}l_{ij}\frac{d\hat{\tau}_{ij}^l}{d\theta}\right)$$

Now, I use equation 5 to replace the total transfers, $\frac{d\hat{T}_i}{d\theta}$, with the net government budgetary position, $\frac{d\hat{t}_i}{d\theta}$, which yields

$$\frac{d\hat{V}_i}{d\theta}|_{\theta=0} = \lambda_i\left(\sum_{j=1}^{J_G}\left(\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} - c_j^G\right)\frac{d\hat{G}_{ij}}{d\theta} + \frac{d\hat{t}_i}{d\theta} + \frac{d}{d\theta}\left[R\left(\hat{\tau}_\mathbf{i}^\mathbf{x}, \hat{\mathbf{x}}_\mathbf{i}, \hat{\tau}_\mathbf{i}^\mathbf{l}, \hat{\mathbf{l}}_\mathbf{i}\right)\right] - \sum_{j=1}^{J_X}x_{ij}\frac{d\hat{\tau}_{ij}^x}{d\theta} - \sum_{j=1}^{J_L}l_{ij}\frac{d\hat{\tau}_{ij}^l}{d\theta}\right)$$

Finally, note that equation 6 shows I can replace the difference between the total revenue impact, $\frac{d}{d\theta}\left[R\left(\hat{\tau}_\mathbf{i}^\mathbf{x}, \hat{\mathbf{x}}_\mathbf{i}, \tau_\mathbf{i}^\mathbf{l}, \hat{\mathbf{l}}_\mathbf{i}\right)\right]$, and the mechanical revenue effect, $\sum_{j=1}^{J_X}x_{ij}\frac{d\hat{\tau}_{ij}^x}{d\theta} + \sum_{j=1}^{J_L}l_{ij}\frac{d\hat{\tau}_{ij}^l}{d\theta}$, with the behavioral impact of the policy on the government budget constraint, yielding

$$\frac{dV_{ij}}{d\theta}|_{\theta=0} = \lambda_i\left(\sum_{j=1}^{J_G}\left(\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} - c_j^G\right)\frac{d\hat{G}_{ij}}{d\theta} + \frac{d\hat{t}_i}{d\theta} + \left(\sum_{j}^{J_X}\tau_{ij}^x\frac{d\hat{x}_{ij}}{d\theta} + \sum_{j}^{J_L}\tau_{ij}^l\frac{d\hat{l}_{ij}}{d\theta}\right)\right)$$

B Non-Marginal Welfare Analysis

In reality, policy changes are not always small. In such cases, one might be worried that the use of the difference in potential outcomes may not reflect a local derivative, $\frac{dx}{d\theta}$. Here, I provide conditions under which one can use the difference in causal effects to construct a measure of the (non-marginal) equivalent variation of the policy change. Heuristically, one can use the framework to estimate equivalent variation as long as the policy does not induce a significant effect on the marginal utility of income.

Equivalent variation, $EV(\theta)$, of the policy at point θ from the initial point $\theta = 0$ is given by the implicit equation:

$$V(P, y + EV(\theta)) = \hat{V}(\theta)$$

where $V(P, y)$ is the utility obtained under policy P with non-labor income $y + EV(\theta)$. Differentiating yields:

$$EV'(\theta) = \frac{\frac{d\hat{V}(\theta)}{d\theta}}{\lambda(P, y + EV(\theta))} = \underbrace{\frac{\lambda\left(\hat{P}(\theta), y\right)}{\lambda(P, y + EV(\theta))}}_{\text{MU of Income Adjustment}} \underbrace{\frac{\frac{d\hat{V}}{d\theta}}{\lambda\left(\hat{P}(\theta), y\right)}}_{\text{Std Measure}}$$

where $\frac{\frac{d\hat{V}}{d\theta}}{\lambda\left(\hat{P}(\theta),y\right)}$ relies on the local causal effects of the policy at $P(\theta)$. Expanding yields:

$$EV(1) = \int_0^1 \frac{\lambda\left(\hat{P}(\theta), y\right)}{\lambda(P, y + EV(\theta))}\left[\left(\frac{\frac{\partial\hat{u}}{\partial G}}{\lambda\left(\hat{P}(\theta), y\right)} - c_G\right)\frac{d\hat{G}}{d\theta} + \frac{d\hat{t}}{d\theta} + \sum_j \hat{\tau}_j \frac{d\hat{x}_j}{d\theta}\right]d\theta$$

Conditions for Global = Local If two conditions are satisfied, global and local conditions are equivalent. Suppose that:

(a) the marginal utility of income does not vary for the policy relative to the income effects: $\lambda\left(\hat{P}(\theta), y\right) = \lambda(P, y + EV(\theta))$

(b) the causal effects are linear in θ (i.e. $\frac{d\hat{x}_j}{d\theta} = \hat{x}_j(1) - \hat{x}_j(0)$ and $\frac{d\hat{l}_j}{d\theta} = \hat{l}_j(1) - \hat{l}_j(0)$ for all θ.

Note that (a) is implied by quasilinear utility, but is far less restrictive. Also, (b) is commonly imposed in empirical applications. To derive the total equivalent variation for the policy, let $D_j = \int_0^1 \left(\frac{\frac{\partial\hat{u}}{\partial G_j}}{\lambda(\hat{P}(\theta),y)} - c_j^G\right)d\theta$ denote the average willingness to pay above cost for the publicly provided goods. Then, if (a) and (b) hold, one can show that:

$$EV(1) = \underbrace{\sum_j \Delta\hat{G}_j * D_j}_{\text{Public Goods}} + \underbrace{\Delta\hat{t}}_{\text{Net Transfer}} + \underbrace{\sum_j \bar{\tau}_j^x \Delta\hat{x}_j + \sum_j \bar{\tau}_j^l \Delta\hat{l}_j}_{\text{Behavioral Reponse}}$$

where $\Delta\hat{G}_j = \hat{G}_j(1) - \hat{G}_j(0)$ is the change in publicly provided good j, $\Delta\hat{t}$ is the change in net resources, and $\Delta\hat{x}_j = \hat{x}_j(1) - \hat{x}_j(0)$ is the difference in potential outcomes in policy world $\theta = 1$ relative to $\theta = 0$

(i.e. $\Delta \hat{x}_j$ is the **non-marginal** causal effect of the policy on $\hat{x}_j$).

C   Appendix: Externalities (and Internalities)

The fact that the causal effect does not need to be decomposed into income and substitution effects extends to a more complex environment with internalities and externalities.

To see this, now suppose that the agents' utility function is given by

$$u_i \left(\mathbf{x_i}, \mathbf{l_i}, \mathbf{G_i}, E_i \right)$$

where the externality imposed on agent i, E_i, is produced in response to the consumption choices of all agents in the economy,

$$E_i = f_i^E \left(\mathbf{x} \right)$$

where $\mathbf{x} = \{\mathbf{x_i}\}_\mathbf{i}$ is the consumption decisions made by the agent (one could generalize this easily to incorporate l). I assume that there is no market for E_i and that agents do not take E_i into account when conducting their optimization. Note that I allow E_i to interact arbitrarily with the utility function, but I assume it is taken as given in the agents' maximization problem. Thus, E_i could represent a classical externality (e.g. pollution) or a behavioral "internality". An internality could be welfare costs of smoking that are not incorporated into their maximization program, or could incorporate "optimization frictions" of the form used by Chetty (2009a) where taxpayers over-estimate the costs of tax sheltering so that the marginal utility of tax sheltered income is not equal to the marginal utility of taxable income.

The value function is now given by

$$
\begin{aligned}
V_i \left(\tau_\mathbf{i}^\mathbf{l}, \tau_\mathbf{i}^\mathbf{x}, T_i, y_i, \mathbf{G_i}, E_i \right) \quad &= \quad \max_{\mathbf{x}, \mathbf{l}} u_i \left(\mathbf{x}, \mathbf{l}, \mathbf{G_i}, E_i \right) \\
&s.t. \quad \sum_{j=1}^{J_X} \left(1 + \tau_{ij}^x \right) x_{ij} \leq \sum_{j=1}^{J_L} \left(1 - \tau_{ij}^l \right) l_{ij} + T_i + y_i
\end{aligned}
$$

Given each agent's solution to this program, $\mathbf{x_i}$, I construct $E_i = f_i^E \left(\mathbf{x} \right)$ and $\mathbf{x}$ is the vector of solutions to each agents optimization program.

All other definitions from Section 2 are maintained. In particular, policy paths are defined as in equation 4.[51] Proposition 2 presents the characterization of the marginal welfare impact of a policy evaluated at $\theta = 0$.

---

[51]Note that I do not allow the government to directly affect the level of E. This would be duplicating the role of publicly provided goods, as I could specify G to be provision of goods which mitigate the externality (either directly or through their effect on agents' choices of x).

Proposition 1. *The welfare impact of the marginal policy change to type i is given by*

$$
\frac{dV_i}{d\theta}\big|_{\theta=0} = \left(\underbrace{\sum_{j=1}^{J_G} \left(\frac{\frac{\partial u_i}{\partial G}}{\lambda_i} - c_j^G \right) \frac{d\hat{G}_{ij}}{d\theta}}_{\text{Public Spending /}\atop\text{Market Failure}} + \underbrace{\frac{d\hat{t}_i}{d\theta}}_{\text{Transfer}} + \underbrace{\left(\sum_j^{J_X} \tau_{ij}^x \frac{d\hat{x}_{ij}}{d\theta} + \sum_j^{J_L} \tau_{ij}^l \frac{d\hat{l}_{ij}}{d\theta} \right)}_{\text{Impact on Govt}} + \underbrace{\frac{\frac{\partial u_i}{\partial E_i}}{\lambda_i} \frac{d\hat{E}_i}{d\theta}}_{\text{Impact on Externality}} \right)
$$

where

$$
\frac{d\hat{E}_i}{d\theta} = \left(\sum_i \sum_j^{J_X} \frac{\partial f_i^E}{\partial x_{ij}} \frac{d\hat{x}_{ij}}{d\theta} \right)
$$

is the net marginal impact of the policy on the externality experienced by type i.

Proof. Taking the total derivative of V_i with respect to θ, I have

$$
\frac{dV_i \left(\hat{\tau}_{\mathbf{i}}^{\mathbf{l}}, \hat{\tau}_{\mathbf{i}}^{\mathbf{x}}, \hat{T}_i, y_i, \hat{\mathbf{G}}_{\mathbf{i}}, \hat{E}_i \right)}{d\theta} = \frac{\partial V_i}{\partial T_i} \frac{d\hat{T}_i}{d\theta} + \sum_{j=1}^{J_G} \frac{\partial V_i}{\partial G_{ij}} \frac{d\hat{G}_{ij}}{d\theta} + \sum_{j=1}^{J_X} \frac{\partial V_i}{\partial \tau_{ij}^x} \frac{d\hat{\tau}_{ij}^x}{d\theta} + \sum_{j=1}^{J_L} \frac{\partial V_i}{\partial \tau_{ij}^l} \frac{d\hat{\tau}_{ij}^l}{d\theta} + \frac{\partial V_i}{\partial E_i} \frac{d\hat{E}_i}{d\theta}
$$

Applying the envelope theorem from the agent's maximization problem and evaluating at $\theta = 0$ implies

$$
\begin{aligned}
\frac{\partial V_i}{\partial \tau_{ij}^x} &= -x_{ij}\lambda_i \\
\frac{\partial V_i}{\partial \tau_{ij}^l} &= -l_{ij}\lambda_i \\
\frac{\partial V_i}{\partial T_i} &= -\lambda_i \\
\frac{\partial V_i}{\partial G_{ij}} &= \frac{\partial u_i}{\partial G_{ij}} \\
\frac{\partial V_i}{\partial E_i} &= \frac{\partial u_i}{\partial E_i}
\end{aligned}
$$

Replacing terms, I have

$$
\frac{dV_i}{d\theta}\big|_{\theta=0} = \lambda_i \left(\frac{d\hat{T}_i}{d\theta} + \sum_{j=1}^{J_G} \frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} \frac{d\hat{G}_{ij}}{d\theta} - \sum_{j=1}^{J_X} x_{ij} \frac{d\hat{\tau}_{ij}^x}{d\theta} - \sum_{j=1}^{J_L} l_{ij} \frac{d\hat{\tau}_{ij}^l}{d\theta} + \frac{\partial u_i}{\partial E_i} \frac{d\hat{E}_i}{d\theta} \right)
$$

Now, I use equation 5 to replace the total transfers, $\frac{d\hat{T}_i}{d\theta}$, with the net government budgetary position,

$\frac{d\hat{t}_i}{d\theta}$, which yields

$$\frac{dV_i}{d\theta}|_{\theta=0} = \lambda_i \left(\sum_{j=1}^{J_G} \left(\frac{\frac{\partial u_i}{\partial \hat{G}_{ij}}}{\lambda_i} - c_j^G \right) \frac{d\hat{G}_{ij}}{d\theta} + \frac{d\hat{t}_i}{d\theta} + \frac{d}{d\theta}\left[R\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}, \hat{\mathbf{x}}_{\mathbf{i}}, \hat{\tau}_{\mathbf{i}}^{\mathbf{l}}, \hat{\mathbf{l}}_{\mathbf{i}} \right) \right] - \sum_{j=1}^{J_X} x_{ij} \frac{d\hat{\tau}_{ij}^x}{d\theta} - \sum_{j=1}^{J_L} l_{ij} \frac{d\hat{\tau}_{ij}^l}{d\theta} + \frac{\partial u_i}{\partial E_i} \frac{d\hat{E}_i}{d\theta} \right)$$

Finally, note that equation 6 shows I can replace the difference between the total revenue impact, $\frac{d}{d\theta}\left[R\left(\hat{\tau}_{\mathbf{i}}^{\mathbf{x}}, \hat{\mathbf{x}}_{\mathbf{i}}, \tau_{\mathbf{i}}^{\mathbf{l}}, \hat{\mathbf{l}}_{\mathbf{i}} \right) \right]$, and the mechanical revenue effect, $\sum_{j=1}^{J_X} x_{ij} \frac{d\hat{\tau}_{ij}^x}{d\theta} + \sum_{j=1}^{J_L} l_{ij} \frac{d\hat{\tau}_{ij}^l}{d\theta}$, with the behavioral impact of the policy on the government budget constraint, yielding

$$\frac{dV_i}{d\theta}|_{\theta=0} = \lambda_i \left(\sum_{j=1}^{J_G} \left(\frac{\frac{\partial u_i}{\partial \hat{G}_{ij}}}{\lambda_i} - c_j^G \right) \frac{d\hat{G}_{ij}}{d\theta} + \frac{d\hat{t}_i}{d\theta} + \left(\sum_j^{J_X} \tau_{ij}^x \frac{d\hat{x}_{ij}}{d\theta} + \sum_j^{J_L} \tau_{ij}^l \frac{d\hat{l}_{ij}}{d\theta} \right) + \frac{\partial u_i}{\partial E_i} \frac{d\hat{E}_i}{d\theta} \right)$$

And, note that I can expand $\frac{d\hat{E}_i}{d\theta}$ by taking a total derivative of $E_i = f_i^E(\mathbf{x})$ across all goods and types, yielding

$$\frac{d\hat{E}_i}{d\theta} = \sum_i \sum_{j=1}^{J_X} \frac{\partial f_i^E}{\partial x_{ij}} \frac{d\hat{x}_{ij}}{d\theta}$$

which concludes the proof. □

With externalities, I must know the net causal effect of behavioral response to the policy on the externality, $\frac{dE_i}{d\theta} = \left(\sum_j^{J_X} \frac{\partial f_i^E}{\partial x_{ij}} \frac{d\hat{x}_{ij}}{d\theta} \right)$, along with the the marginal willingness to pay for the externality, $\frac{\frac{\partial u_i}{\partial E_i}}{\lambda_i}$. Therefore, the welfare loss from a behavioral response that reduces government revenue may be counteracted by the welfare gain from any reduction on the externality imposed on other individuals. Thus, financing government revenue using so-called "green taxes" that also reduce externalities may deliver higher government welfare than policies whose financing schemes do not reduce externalities.[52] This is the so-called "double-dividend" highlighted in previous literature (Bovenberg and de Mooji (1994); Goulder (1995); Parry (1995)). But even in this world, the causal effect of the policy on behavior, i.e. the policy elasticity, continue to be the behavioral elasticities that are relevant for estimating welfare impact of the policy.

D   General Equilibrium Effects

By assuming one unit of goods are produced with one unit of labor supply, the model ruled out general equilibrium effects (i.e. that the policy change affects prices). However, such effects are easily incorporated into the model by adding the implied transfers to the net resources term, $\frac{d\hat{t}_i}{d\theta}$.

For example, if the policy increases the price of i's labor supply activity j, then she will obtain a resource benefit of $l_{ij} \frac{dw_{ij}}{d\theta}|_{\theta=0}$, where $\frac{dw_{ij}}{d\theta}|_{\theta=0}$ is the causal impact of the policy on the after-tax wage faced by individual i on her jth labor supply activity. These additional impacts are valued

---

[52] As is well-known (e.g. Salanie (2003)), if taxes are initially near their optimal levels, then at the margin it is not clear that an additional green tax will be any more desirable than a tax on any other good.

dollar-for-dollar and can simply be added to the resource transfer term, $\frac{d\hat{t}_i}{d\theta}|_{\theta=0}$. Hence, when policies have general equilibrium effects, one also needs to track the causal impact of the policy on prices, and adjust the size of the transfers, $\frac{d\hat{t}_i}{d\theta}$ accordingly. The causal effects are still the desired responses, but one needs to also know the general equilibrium effects of government policies.

E   Optimal Commodity Taxation and the "Inverse Elasticity" Rule

Ramsey (1927) proposes the question of how commodities should be taxed in order to raise a fixed government expenditure, $R > 0$. Diamond and Mirrlees (1971) provide a formal modeling of this environment and show that, at the optimum, the tax-weighted Hicksian price derivatives for each good are equated. Here, I illustrate this result and relate it to the framework provided in this paper.

Assume there is a representative agent and drop i subscripts. A necessary conditions for tax policy to be at an optimum is given by

$$\frac{d\hat{V}_P}{d\theta} = 0$$

for all feasible policy paths, P. With a representative agent, the optimal tax would be lump-sum of size R. However, the optimal commodity tax program proposed by Ramsey (1927) makes the assumption that the government cannot conduct lump-sum taxation. Hence, the only feasible policies are those that raise and lower tax rates in a manner that preserves the budget constraint.

Consider a policy, $P(\theta)$, that lowers the tax on good 1 and raises the tax on good 2. The optimality condition is given by

$$\sum_k \hat{\tau}_k \frac{d\hat{x}_k}{d\theta} = 0 \tag{18}$$

Equation (18) suggests more responsive goods should be taxed at lower rates, thereby nesting the standard "inverse elasticity" argument (higher $\frac{d\hat{x}_k}{d\theta}$ should be associated with lower $\hat{\tau}_k$). The optimal tax attempts to replicate lump-sum taxes by taxing relatively inelastic goods.

Diamond and Mirrlees (1971) further note that, because $\frac{d\hat{V}_P}{d\theta} = 0$ at the optimum, one can expand the behavioral change using the Hicksian demands, x_k^h,

$$\frac{dx_k}{d\theta} = \frac{\partial x_k^h}{\partial \tau_1}\frac{d\tau_1}{d\theta} + \frac{\partial x_k^h}{\partial \tau_2}\frac{d\tau_2}{d\theta}$$

where, in general, there would be the additional term, $\frac{\partial x_k^h}{\partial u}\frac{dV_p}{d\theta}$, but this vanishes at the optimum. Hence, that the optimality condition is given by

$$\sum_k \tau_k \frac{\partial x_k^h}{\partial \tau_1}\frac{d\tau_1}{d\theta} = \sum_k \tau_k \frac{\partial x_k^h}{\partial \tau_2}\left(-\frac{d\tau_2}{d\theta}\right) \tag{19}$$

so that the tax-weighted Hicksian responses are equated across the tax rates – precisely the classic result in Diamond and Mirrlees (1971) (see equation 38).[53]

---

[53]Under the additional assumption that compensated cross-price elasticities are zero, one arrives at the classic inverse

However, note that one never relied on compensated elasticities to test the optimality condition in equation (18). Compensated elasticities arise only because of the assumption that policy is at the optimum. One could consider any budget-neutral policy that simultaneously adjusts two commodity taxes and test equation (18) directly. Conditional on knowing the causal effects of such a policy, one would not need to know whether income or substitution effects drive the behavioral response to commodity taxes. The policy elasticities would be sufficient.

F Application Details

F.1 EITC

This section outlines the welfare analysis of an EITC expansion. To correspond with the causal effects analyzed in much previous literature, the marginal expansion of the EITC program can best be thought of as increasing the maximum benefit level in a manner that maintains current income eligibility thresholds and tax schedule kink points (but raises the phase-in and phase-out rates in order to reach the new maximum benefit). However, the results from Chetty et al. (2013) suggest the phase-out slope of the EITC has only a minor impact on labor supply (most of the response is from individuals below the EITC maximum benefit level choosing to increase their labor supply). This suggests the impact on the behavioral response on the government budget would not be too sensitive to the precise design of the phase-out of the program.

The effects documented in previous literature consist of both intensive and extensive labor supply responses. With extensive margin responses, $\frac{d\hat{l}_i^{EITC}}{d\theta}$ may not exist for all i, as individuals make discrete jumps in their choice of labor supply. However, this is easily accommodated into the model. To see this, normalize the index of the Poor to be the unit interval, $i \in Poor = [0,1]$. Then, order the index of the poor population such that $\hat{l}_i(\theta) > 0$ implies $\hat{l}_j(\theta) > 0$ for $j < i$ and all $\theta \in (-\epsilon, \epsilon)$. With this ordering, there exists a threshold, $i^{LFP}(\theta)$, such that $i < i^{LFP}(\theta)$ indicates that i is in the labor force and $i > i^{LFP}(\theta)$ indicates that i is not in the labor force. Hence, $i^{LFP}(\theta)$ is the fraction of the poor single mothers that are in the labor force. With this notation, the impact of the behavioral response to the policy by the poor on the government's budget is given by:

$$-\int_{i \in Poor} \tau_i^l \frac{d\hat{l}_i^{EITC}}{d\theta}|_{\theta=0} di = \underbrace{\left(\tau_{i^{LFP}(0)}^l l_{i^{LFP}(0)}\right) \frac{di^{LFP}}{d\theta}|_{\theta=0}}_{\text{Extensive Margin}} - \underbrace{\int_{i<\hat{i}} \tau_i^l \frac{d\hat{l}_i^{EITC}}{d\theta}|_{\theta=0} di}_{\text{Intensive Margin}} \quad (20)$$

where $\tau_{i^{LFP}(0)}^l l_{i^{LFP}(0)}$ is the average taxable income (or loss) generated by the marginal type entering the labor force and $\frac{di^{LFP}}{d\theta}$ is the marginal rate at which the policy induces labor force entry. The cost

elasticity rule:

$$\frac{\tau_2}{\tau_1} = \frac{\frac{\partial x_1^h}{\partial \tau_1} \frac{d\tau_1}{d\theta}}{\frac{\partial x_2^h}{\partial \tau_2} \frac{d\tau_2}{d\theta}}$$

so that optimal tax rates are inversely proportional to their compensated (Hicksian) demands.

resulting from extensive margin responses is given by the impact of the program on the labor force participation rate, multiplied by the size of the average subsidy to those entering the labor force.[54]

There is a large literature analyzing the impact of the EITC expansion on labor force participation of single mothers, beginning with Eissa and Liebman (1996). These approaches generally estimate the causal effect of EITC receipt on behavior using various expansions in the generosity of the EITC program. Hotz and Scholz (2003) summarize this literature and find consistency across methodologies in estimates of the elasticity of the labor force participation rate of single mothers, $\hat{i}$, rate with respect to the average after-tax wage, $E\left[\left(1-\tau_{\hat{i}}^l\right)l_{\hat{i}}\right]$, with estimates ranging from 0.69-1.16.

I translate this elasticity into equation (20) by normalizing θ to parameterize an additional unit of the mechanical subsidy[55] and writing:

$$\left(\tau_{iLFP(0)}^l l_{iLFP(0)}\right)\frac{di^{LFP}}{d\theta}|_{\theta=0} = \frac{\left(\tau_{iLFP(0)}^l l_{iLFP(0)}\right)}{\left(\left(1-\tau_{iLFP(0)}^l\right)l_{iLFP(0)}\right)}\epsilon_{E\left[\left(1-\tau_{\hat{i}}^l\right)l_{\hat{i}}\right]}^{LFP}$$

where $\epsilon_{E\left[\left(1-\tau_{\hat{i}}^l\right)l_{\hat{i}}\right]}^{LFP}$ is the elasticity of the labor force participation rate with respect to the after tax wage rate and $\frac{E\left[\tau_{\hat{i}}^l l_{\hat{i}}\right]}{E\left[\left(1-\tau_{\hat{i}}^l\right)l_{\hat{i}}\right]}$ is the size of the subsidy as a fraction of after tax income for the marginal labor force entrant. For the elasticity of labor force participation, I choose an estimate of 0.9, equal to the midpoint of existing estimates (Hotz and Scholz (2003)). For $\frac{E\left[\tau_{\hat{i}}^l l_{\hat{i}}\right]}{E\left[\left(1-\tau_{\hat{i}}^l\right)l_{\hat{i}}\right]}$, one desires the after tax wages and subsidies for marginal entrants into the labor force. While such parameters could be identified using the same identification strategies previous papers have used to estimate the labor supply impact of the EITC, to my knowledge no such estimates of the marginal wages and subsidies exist. Using the 2004 SOI, Eissa and Hoynes (2011) report that the average subsidy is $1,806 per beneficiary, which corresponds to 9.2% of a $20,000 gross income for EITC beneficiaries. Athreya et al. (2010) report the average recipient obtains a subsidy equal to 11.7% of gross income in the 2008 CPS. I therefore take the approximate midpoint of 11%.

These calculations suggest the extensive margin impact on the government budget is given by:

$$\underbrace{E\left[\tau_{\hat{i}}^l l_{\hat{i}}\right]\frac{d\hat{i}}{d\theta}}_{\text{Extensive Margin}} = \frac{0.11}{1+0.11}*0.9 = 0.09$$

so that the EITC is 9% more costly to the government because of extensive margin labor supply responses.[56]

---

[54]Because the model assumed individuals face linear tax rates, the distinction between the average and marginal tax rate is not readily provided, but it is straightforward to verify that the fiscal externality imposed by those entering the labor force is given by the size of the subsidy they receive by entering the labor force, not by the marginal tax or subsidy they face if they were to provide an additional unit of labor supply.

[55]This normalizes $\int_{i\in Poor}\left(\frac{d\hat{T}_i^{EITC}}{d\theta}|_{\theta=0} + \frac{d\hat{\tau}_{ij}^{EITC}}{d\theta}|_{\theta=0}l_i\right)di = 1$

[56]Taking elasticity estimates in the 0.69-1.12 range reported by Hotz and Scholz (2003), yields estimates of the extensive margin impact ranging from 0.07 to 0.11. Hence, if one assumed only extensive margin responses were operating, the

Intensive margin responses   Until recently, there was little evidence that the EITC had intensive margin impacts on labor supply. However, the recent paper by Chetty et al. (2013) exploits the geographic variation in knowledge about the marginal incentives induced by the EITC, as proxied by the local fraction of self-employed that bunch at the subsidy-maximizing kink rate. Using the universe of tax return data from EITC recipients, their estimates suggest that the behavioral responses induced by knowledge about the marginal incentives provided by the EITC increase refunds by approximately 5% relative to what they would be in the absence of behavioral responses, with most of these responses due to intensive margin adjustments. What is particularly useful about this study is that it uses tax expenditures as an outcome variable, and hence can compute the associated fiscal externality directly.

The downside of Chetty et al. (2013) is that the policy path in question is the degree of "knowledge about the shape of the EITC schedule". While this policy path provides guidance on the size of the distortions induced by these marginal incentives, one could imagine that even in places with no knowledge of the EITC schedules the existence of the EITC generates extensive margin responses.

To account for this, I make the baseline assumption that the knowledge of the average EITC subsidy generates extensive margin responses and knowledge of the shape of the EITC schedule generates intensive margin responses. With this assumption, the results of Chetty et al. (2013) should be added together with the extensive margin responses found in previous literature to arrive at the total impact of an EITC expansion. This yields an estimate of $FE^{EITC} = 0.09 + 0.05 = 14\%$ with a range of 0.12-0.16 taking the range of extensive margin labor supply responses.[57]

F.2   Food Stamps

Using variation induced in the introduction of food stamps in the 1960s and 70s Hoynes and Schanzenbach (2012) estimate that food stamps led to a significant reduction in labor supply, especially among female headed households. They estimate a fairly imprecise and large reduction in labor hours (-658 hours per year, with a 95% CI of [-1186 , 130]; see Column (2) of Table 2 on page 157). They also estimate a large and imprecise change in annual earnings of -$2,943 (95% CI of [-10,169 , 4,284]). Corresponding to the tax rates operating around 1970, I assume a linear marginal tax rate of 20% on earnings, consistent with the absence of an EITC program during this time period. I arrive at 20% using the 14% bottom tax bracket for federal taxes and a 6% state tax assumption. With this assumption, the net change tax revenue collected due to behavioral responses to food stamps is $2,943*.2=$588.60. It is important to note that this estimate is not statistically significantly different from zero.

In contrast, the food stamp program provided an average monthly benefit of $26.77 per per-

policy elasticity would be $FE^{EITC} = 0.09$, ranging between 0.07 and 0.11.

[57]This is potentially an overestimate of the net effect of behavioral responses because some of the responses found in Chetty et al. (2013) is along the extensive margin and is more amenable to the potential critique that the earlier literature could not effectively separate the impact of EITC expansions from the impact of the decrease in welfare generosity (see Meyer and Rosenbaum (2001) for this debate). Therefore, I also consider the case that the 0.05 figure in Chetty et al. (2013) captures all of the EITC response (so that $FE^{EITC} = 0.05$). This arguably provides a lower bound of the impact of the policy. For an upper bound, I consider the upper range of extensive margin response can be added to Chetty et al. (2013), so that $FE^{EITC} = 0.11 + 0.05 = 16\%$.

son in 1978[58], which corresponds to \$321.24 per person per year. Hoynes and Schanzenbach (2009) estimate a mean household size of 3.59 in their sample, which implies a household-level transfer size of \$1,153.25. Hence, the total cost to the government of providing the food stamps policy is \$1,153.25+\$588.60=\$1,741.85.

For the net valuation of food stamps, Smeeding (1982) estimates that food stamps are valued dollar-for-dollar. In contrast, Whitmore (2002) estimates that every dollar of food stamps is valued at ~\$0.80 by the beneficiaries. In the absence of behavioral responses this estimate suggests the MVPF would be 0.8. Placing this into the context of the size of the transfers, the estimate suggests that the mechanical transfer of \$1,153.25 is valued by beneficiaries at only \$922.60.

F.3   Section 8 Housing Vouchers

Jacob and Ludwig (2012) study the impact of obtaining a housing voucher on labor supply (intensive + extensive), Medicaid receipt, TANF receipt, and SNAP receipt. For the extensive margin labor supply response, I use the 11% tax rate assumption from the EITC section. For the intensive margin response, Jacob and Ludwig (2012) report a marginal tax rate of 24% for the treatment group that includes phase-out of government benefits in addition to marginal income tax rates. For the change in TANF and SNAP use, I use the Green Book (2004) and compute average costs per household in 2002, normalized to 2007 dollars using the CPI-U to be consistent with Jacob and Ludwig (2012). For the change in Medicare enrollment, I use costs compiled by Holahan and McMorrow (2012). Table A1 reports the calculations.

---

[58]www.fns.usda.gov/pd/SNAPsummary.htm+&cd=3&hl=en&ct=clnk&gl=us&client=safari

Table A1: Housing Voucher Calculation

Item	Value	Source
Voucher Cost	$8,400	Jacob and Ludwig (2012)
Extensive Margin		
Extensive Margin Change	0.036	
Mean earnings	22,232	Mean earnings of $5557 per quarter (Jacob and Ludwig (2012))
Assumed Avg Tax Rate	-11%	Avg EITC subsidy (see EITC section)
Fiscal Impact	-88.0	
Intensive Margin		
% Intensive	0.592	Jacob and Ludwig (2012)
Earnings change	910	Jacob and Ludwig (2012) report quarterly change of 227.54
Assumed Tax Rate	24%	Jacob and Ludwig estimate 18% tax rate in control group and 24% tax rate in treatment group that includes phase-out of services
Fiscal Impact	129.3	
TANF		
% Increase	0.017	Jacob and Ludwig (2012)
Avg Cost	254	Mean monthly TANF from Green Book (2007 dollars deflated using CPI-U as in Jacob and Ludwig (2012))
Fiscal Impact	4.3	
Medicaid		
% Increase	0.058	Jacob and Ludwig (2012)
Avg Cost per enrollee	6192	Halahan and McMorrow (2012) Appendix Table 1 reports costs for 2002
Fiscal Impact	359.1	
Food Stamps		
% Increase	0.076	
Avg Cost per family	364	Family of 4 assumption
Fiscal Impact	27.7	
Total Program Cost	$8,832.4	
Total Behavioral Impact	$432.4	
% of cost	0.05	
Valuation of $1 of Voucher	0.83	Reeder (1985)
Value of $8,400 Voucher	6,972	
Net Valuation	-1,428	
% of cost	-0.16	
MVPF	$0.79	