The Policy Elasticity

Nathaniel Hendren*

September, 2013

Abstract

This paper applies basic price theory to study the marginal welfare impact of government policy changes. In contrast to the canonical marginal excess burden framework, the framework presented here does not require a decomposition of behavioral responses to the policy into income and substitution effects. The causal effects of the policy are sufficient. Moreover, in the broad class of models where the government is the only distortion, the causal impact of the behavioral response to the policy on the government budget is sufficient for all behavioral responses. Because these behavioral responses vary with the policy in question and are, in general, neither pure Hicksian nor Marshallian elasticities, I term them *policy elasticities*. The model yields a natural welfare measure for non-budget neutral policies: the welfare impact per dollar of government expenditure. I illustrate the framework by examining EITC generosity and changes to the top marginal income tax rate. Existing causal estimates suggest additional redistribution is desirable if and only if one prefers giving an additional \$0.44-0.66 to an EITC-eligible single mother (earning less than \$40,000) relative to an additional \$1 to a person subject to the top marginal tax rate (earning more than \$400,000).

1 Introduction

There is a long history in economics of estimating marginal deadweight loss or marginal excess burden (MEB) to study the normative implications of government policy changes. Done properly, calculation of MEB requires decomposition of the behavioral response to policy changes into income and substitution effects. Only the substitution effect is desired for such a welfare analysis.[1]

*Harvard University and NBER (e-mail: nhendren@fas.harvard.edu). I would like to thank Raj Chetty, Amy Finkelstein, Don Fullerton, Peter Ganong, Adam Guren, Louis Kaplow, Wojciech Kopczuk, Erzo Luttmer, Emmanuel Saez, and seminar participants at Chicago Booth School of Business, Brown University, The University of Chicago, the Columbia Tax Policy Workshop, and the Minneapolis Federal Reserve for helpful comments. Financial support from the NBER Health and Aging Fellowship, under the National Institute of Aging Grant Number T32-AG000186 is gratefully acknowledged.

[1] See, e.g., Harberger (1964); Mas-Colell et al. (1995); Feldstein (1999); Chetty (2009b). The resulting importance of the compensated elasticity for marginal welfare analysis is discussed in recent JEL surveys:

> *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 the tax base with respect to a marginal change in tax rates.* (Feldstein (2012))

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[2] 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.

This paper shows how causal effects can be directly used in welfare analysis of government policy changes. It does so using the canonical price theoretic framework with heterogeneous agents. In contrast to calculating MEB[3], I characterize each agent's willingness to pay out of their own income for a given policy change. The main result is that 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. Because real-world policy changes are often complex, 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 (e.g. Angrist and Pischke (2008)).

Moreover, 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.[4] 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.[5]

> *Graduate textbooks teach that the two central aspects of the public sector, optimal progressivity of the tax-and-transfer system, as well as the optimal size of the public sector, depend (inversely) on the compensated elasticity of labor supply with respect to the marginal tax rate.* (Saez, Slemrod, and Giertz (2012))

[2] In Goolsbee's case, the natural experiment was a change in top income marginal tax rates.

[3] MEB calculations compute the additional revenue the government could obtain under the policy if utility were held constant using individual-specific lump-sum transfers.

[4] The idea that a "fiscal externality" is important for welfare evaluation is not new; this is the core idea behind Harberger (1964)'s triangle and the sufficiency of the taxable income elasticity in Feldstein (1999). But in contrast to the MEB framework of Feldstein (1999), the desired response here is the causal effect, not the compensated effect.

[5] 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 Section 2.7.

With the causal effect, welfare analysis follows straightforwardly. Only two other components are required to calculate an individuals' willingness to pay for the policy change, both of which are arguably well-known. First, if a policy changes the provision of publicly provided goods, one also needs to know the net willingness to pay for these goods. This is given by the difference between individuals' marginal rates of substitution and the marginal cost of production – an insight of Samuelson (1954). Indeed, this is a term that should be interpreted broadly as the relative advantage of the government over the private market (or vice-versa) in addressing the policy change. It is positive to the extent to which the value of the provision of public goods or services exceed their resource costs. Second, one needs to know the change in net government transfers to the individual, which are valued dollar-for-dollar by the individual. These three components – (1) the causal impact of the response to the policy change on the government's budget, (2) the net willingness to pay for the change in publicly provided goods, and (3) the net transfers – fully characterize the welfare impact of marginal policy changes to an individual.

While (1)-(3) characterize the welfare impact on a given individual, aggregation across individuals involves weighting by each person's social marginal utility of income. This is useful because ratios of social marginal utilities have a simple interpretation in terms of Okun's leaky bucket experiment (Okun (1975)): how much resources is society willing to lose to transfer from on person to another? In contrast, aggregation of MEB across individuals in a manner consistent with marginal social welfare measurement requires adding back in the income effects that were removed for the construction of MEB.[6] In this sense, the aggregation of welfare across people is more easily accomplished when using the causal effects for conducting welfare as opposed to the MEB framework.

The framework can be applied to both budget-neutral and non-budget neutral policies alike. Indeed, many government policy changes are not budget neutral, at least in the short run. For dealing with non-budget neutral policies, straightforward differentiation shows that the welfare impact of two policies (e.g. tax and expenditure policies) can be added together to form a welfare analysis of a budget-neutral policy as long as the two policies sum to the policy of interest.[7] Although there is an extraordinary number of alternative definitions of the "marginal cost of public funds" in previous literature, the additivity condition motivates a particular measure of the marginal value of public funds (MVPF) suggested by (Slemrod and Yitzhaki (1996, 2001)): the marginal social welfare impact of the policy per unit of government revenue expended.[8] With this definition of the MVPF, one can compare

[6] This feature of MEB was initially derived by Diamond and Mirrlees (1971). See Auerbach and Hines (2002) for a simple derivation of this on page 1370, equation 3.24.

[7] This might seem like an obvious condition, but it is violated if, for example, one used the MEB of a tax increase to adjust the standard Samuelson condition for the cost of raising revenue to finance the public good.

[8] I use the term MVPF instead of MCPF because the policy need not be solely an expenditure or tax policy – it could be any non-budget neutral policy. Relative to Slemrod and Yitzhaki (1996, 2001), this paper shows that this definition has the unique advantage that the only behavioral responses required are the causal impacts of the policy. This feature was arguably unclear in this earlier work due to the fact that Slemrod and Yitzhaki (1996, 2001) refer to the behavioral response using the language of marginal excess burden and deadweight loss, implicitly suggesting the desired behavioral responses are compensated.

Moreover, this definition of the MVPF is particularly useful for the analysis of tax policies in a dynamic setting. Individuals taxed today may expect lower taxes in the future (i.e. a classic Ricardian equivalence generally overlooked in the MCPF literature). If individuals borrow or save against these future tax changes, they may respond in a compensated

the cost-effectiveness across policies: taking revenue from policies with low MVPF and spending on policies with high MVPF increases social welfare.

I illustrate the applicability of the model to study the desirability of additional redistribution. In particular, I consider a policy of raising the top marginal income tax rate to finance an expansion of the earned income tax credit (EITC). For the top tax rate, Saez et al. (2012) and Giertz (2009) suggest mid-range estimates that 25-50% of the mechanical revenue that is raised is lost due to behavioral distortions.[9] This suggests a MVPF of taxing top earners of \$1.33-\$2. For the EITC, 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. Aggregating using the social marginal utilities of income[10], 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.

Relation to Previous Literature This paper is of course not the first to study the types of behavioral elasticities required for normative analysis of government policies. As discussed above, previous literature has often highlighted the importance of the Hicksian (compensated) elasticity. However, Hicksian price elasticities are the causal effects of policies that are known to hold utility constant. Hence, they are insufficient in this framework for measuring the marginal welfare impact of policies that actually change utilities.

Hicksian elasticities arise in MEB calculations because the conceptual experiment compensates agents with lump-sum transfers in order to hold their utility constant. Instead of asking how much individuals are willing to pay for the policy change, MEB asks how much additional revenue the government could receive as a result of the policy change if utilities were held constant using individual-specific lump-sum transfers (Auerbach and Hines (2002)). Although MEB is a reasonable metric for evaluating marginal policy changes, it is not empirically tractable unless the empiricist can decompose the behavioral responses into income and substitution effects. In contrast, calculating individuals' marginal willingness to pay for the policy change in units of their own income relies on the causal, not compensated, effect of the policy change. Moreover, the resulting welfare measures can be aggregated using the social marginal utility of income, in contrast to MEB which requires adding back in the income effects to form the marginal social welfare impact of policy changes.

Hicksian elasticities also arise in the optimal commodity taxation program with a representative agent proposed by Ramsey (1927) and analyzed in detail by Diamond and Mirrlees (1971). In Appendix D, I illustrate how the present model can nest this result. At an optimum, the marginal welfare impact

manner (Barro (1974)). As discussed in footnote 41, the MVPF does not require knowledge of the degree to which Ricardian equivalence holds, provided one can estimate the causal effects.

[9] Saez et al. (2012) suggest a midpoint of around 20-25% while Giertz (2009) suggests a midpoint of around 50%

[10] In contrast, MEB calculations such as those in Eissa et al. (2008) for recent EITC expansions, cannot be aggregated using the social marginal utilities of income. One would first need to adjust the social marginal utilities with the income effects that were removed to calculate the MEB (see footnote 6).

of a budget-neutral policy change is zero. So, in representative agent models, optimal taxes depend on Hicksian elasticities because utility is locally constant at the optimum. More generally, however, the social welfare impact of changing commodity tax rates depends not on the Hicksian elasticity but rather on the causal impact of such policy changes.

This paper is also related to the "Stiglitz-Dasgupta-Atkinson-Stern"[11] approach to defining the marginal cost of public funds. In the language of the present framework, this tradition defines the MCPF as a particular sub-component (namely, the causal impact of the behavioral response to the policy on the government's budget) of a welfare analysis of a broader policy that increases taxes and exhausts the revenue on a public good. In contrast, the MVPF presented here is policy specific and defined for any non-budget neutral policy.

This paper is also related to the work studying the optimal design of tax and transfer systems (e.g. Mirrlees (1971) ,Diamond and Mirrlees (1971), Saez (2001) among others). This literature uses a first order condition to write the (constrained) optimal tax rates as functions of potentially estimable elasticities. By construction, these elasticities are defined locally *around the optimum*. Hence, it is important that the elasticities are "structural" so that extrapolation of estimates using local variation provides an estimate around the optimum. In contrast, estimating the welfare impact of policy changes relies solely on causal effects defined locally around the status quo.[12]

This paper is also related to 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 with nonlinear tax schedules (Kleven and Kreiner (2006))). However, I show 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 estimating the impact of the behavioral response on tax revenue, as opposed to taxable income, is most useful for welfare analysis.[13]

The rest of this paper proceeds as follows. Section 2 presents the model and provides the main result. Section 3 discusses the additivity condition and the marginal costs of public funds. Section 4 applies the framework to study the desirability of additional redistribution. Section 5 concludes. Appendix A also illustrates how to apply the model to study the Job Training Partnership Act (JTPA) analyzed in Bloom et al. (1997). Additional appendices provide proofs of the main results (Appendix B) an extension of the model to incorporate other externalities (Appendix C), and a more detailed

[11] See Ballard and Fullerton (1992). The label refers to the foundational papers by Atkinson and Stern (1974); Stiglitz and Dasgupta (1971) and Atkinson and Stern (1974). See Dahlby (2008) for a recent survey of the literature on the marginal cost of public funds.

[12] Moreover, although the optimality conditions provide insight into the optimal slope of the tax schedule at a given income level, the optimal level of the schedule depends on an integral of the elasticities across the entire income distribution (Piketty and Saez (2012)). Hence, the optimal amount of redistribution to the poor depends on a very complicated set of elasticities. In contrast, whether or not one wishes to redistribute through a given policy depends solely on the causal effects of that particular policy.

[13] An example of a paper that does this is Chetty et al. (2013), who estimate that behavioral responses to the EITC increase tax expenditures by 5%.

discussion of the relationship to previous literature (Appendix D).

2 Model

I consider a canonical price-theoretic model with heterogeneous agents and multiple goods, along with a government that sets taxes, transfers, and publicly provided goods. The generality captures many realistic issues faced in empirical applications and also allows the model to nest many models in previous literature. But, for simplified reading, Example 1 on page 11 illustrates the main concepts in a model with a representative agent, single taxable good, and single publicly provided good. Appendix A also applies the model to the Job Training Partnership Act (JTPA) studied in Bloom et al. (1997).

2.1 Setup

There exist a continuum of individuals of equal mass in the population, indexed by $i \in I$. These individuals make two choices: they choose a vector of J_X goods to consume, $\mathbf{x_i} = \{x_{ij}\}_{j=1}^{J_X}$, and a vector of labor supply activities, $\mathbf{l_i} = \{l_{ij}\}_{j=1}^{J_L}$.[14] There also exists a government that does three things: it provides a vector of J_G publicly provided goods to each individual, $\mathbf{G_i} = \{G_{ij}\}_{j=1}^{J_G}$, provides monetary transfers to each individual, T_i, and imposes linear taxes[15] 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 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.[16]

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. Thus, agents face a single linear budget constraint given by

$$\sum_{j=1}^{J_X}\left(1+\tau_{ij}^{x}\right)x_i \leq \sum_{j=1}^{J_L}\left(1-\tau_{ij}^{l}\right)l_{ij} + T_i + y_i \tag{2}$$

where y_i is non-labor income.[17] 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

[14] For example, l_{i1} could be labor supplied in wage work and l_{i2} could be labor supplied in the informal (un-taxed) sector.

[15] 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.

[16] 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.

[17] I allow (but do not require) taxes and transfers to be individual-specific. This allows the model to next the standard MEB experiment.

surplus), production externalities (e.g. spillovers), and pecuniary externalities (in which case real prices would not always be 1).[18] I assume the marginal cost to the government of producing publicly-provided goods, G_{ij} is given by c_j^G for $j = 1, .., J_G$.[19]

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 yields the standard indirect utility function of individual i,

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

where V_i depends on taxes, transfers, income, and publicly provided goods. The Marshallian demand functions 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 $l_{ij}^m\left(\tau_{\mathbf{i}}^{\mathbf{x}}, \tau_{\mathbf{i}}^{\mathbf{l}}, T_i, \mathbf{G_i}, y_i\right)$. Because the utility function is allowed to vary arbitrarily across people, it will be helpful to be able to normalize by an 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 $\left(\tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}\right)$. The standard duality result implies that $E_i\left(V_i\left(\tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}, y_i\right); \tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}\right) = y_i$.

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

$$W\left(\left\{\tau_{\mathbf{i}}^{\mathbf{l}}, \tau_{\mathbf{i}}^{\mathbf{x}}, T_i, \mathbf{G_i}, y_i\right\}_i\right) = \int_{i\in I} \psi_i V_i\left(\tau_{\mathbf{i}}^{\mathbf{x}}, \tau_{\mathbf{i}}^{\mathbf{l}}, T_i, \mathbf{G_i}, y_i\right) di \tag{3}$$

Note that this is an implicit function of the vector of taxes, transfers, and publicly provided goods to every type in the economy. In what follows, it will also be helpful to also consider the social marginal utility of income, $\eta_i = \psi_i \lambda_i$, which is the social welfare weight in units of the individual's own income.

2.2 Policy Paths and Potential Outcomes

The social welfare function, W, provides a theoretical metric for evaluating the desirability of government policy. In this subsection, I use this metric to evaluate the welfare impact of marginal changes to the status quo policy. To do so, I define a "policy path", $P(\theta)$. For any θ in a small region near 0,

[18] I discuss some of these extensions in Subsection 2.7 and provide a detailed discussion of externalities in Appendix C.

[19] 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.

$\theta \in (-\epsilon, \epsilon)$, let $P(\theta)$ be a vector of taxes, transfers, and publicly provided goods to each individual,

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

where the "^" indicates the policies are functions of θ. 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:

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

Second, I assume that the policy path is continuously differentiable in θ (i.e. $\frac{d\hat{\tau}^x_{ij}}{d\theta}$, $\frac{d\hat{\tau}^l_{ij}}{d\theta}$, $\frac{d\hat{T}_i}{d\theta}$, and $\frac{d\hat{G}_{ij}}{d\theta}$ exist and are continuous in θ).[20] 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.[21]

Before asking the normative question of whether the government should follow the policy path, I first consider the positive question of what the policy change would do to behavior. Given a policy path, I assume individuals choose goods and labor supply activities, $\hat{\mathbf{x}}_i(\theta) = \{\hat{x}_{ij}(\theta)\}_i$ and $\hat{\mathbf{l}}_i(\theta) = \left\{\hat{l}_{ij}(\theta)\right\}_i$, that maximize their utility under policy $P(\theta)$.[22] 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 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,

$$\underbrace{\hat{t}_i(\theta)}_{\text{Net Resources}} = \underbrace{\sum_{j=1}^{J_G} c^G_j \hat{G}_{ij}(\theta)}_{\text{Public-Provided Goods}} + \underbrace{\hat{T}_i(\theta)}_{\text{Transfers}} - \underbrace{\left(\sum_{j=1}^{J_X} \hat{\tau}^x_{ij}(\theta)\hat{x}_{ij}(\theta) + \sum_{j=1}^{J_L} \hat{\tau}^l_{ij}(\theta)\hat{l}_{ij}(\theta) \right)}_{\text{Tax Revenue}} \tag{5}$$

where $\sum_{j=1}^{J_G} c^G_j \hat{G}_{ij}(\theta)$ is the government expenditure on publicly provided goods to individual i, $\hat{T}_i(\theta)$

[20] 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 4, 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))).

[21] 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 4.

[22] These can be calculated in theory by evaluating the Marshallian demands at the policy vector for each θ:

$$\begin{aligned} \hat{x}_{ij}(\theta) &= x^m_{ij}\left(\hat{\tau}^{\mathbf{x}}_{\mathbf{i}}(\theta), \hat{\tau}^{\mathbf{l}}_{\mathbf{i}}(\theta), \hat{T}_i(\theta), \hat{\mathbf{G}}_{\mathbf{i}}(\theta)\right) \quad \forall j = 1..J_X \\ \hat{l}_{ij}(\theta) &= l^m_{ij}\left(\hat{\tau}^{\mathbf{x}}_{\mathbf{i}}(\theta), \hat{\tau}^{\mathbf{l}}_{\mathbf{i}}(\theta), \hat{T}_i(\theta), \hat{\mathbf{G}}_{\mathbf{i}}(\theta)\right) \quad \forall j = 1..J_L \end{aligned}$$

is the government transfers to type i, and $\sum_{j=1}^{J_X} \hat{\tau}_{ij}^x(\theta)\,\hat{x}_{ij}(\theta) + \sum_{j=1}^{J_L} \hat{\tau}_{ij}^l(\theta)\,\hat{l}_{ij}(\theta)$ 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[23], but budget-neutrality of a policy path could be imposed by assuming

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

where

$$\frac{d\hat{t}_i}{d\theta} = \sum_j c_j^G \frac{d\hat{G}_{ij}}{d\theta} + \frac{d\hat{T}_i}{d\theta} - \frac{d}{d\theta}\left[\sum_{j=1}^{J_X} \hat{\tau}_{ij}^x(\theta)\,\hat{x}_{ij}(\theta) + \sum_{j=1}^{J_L} \hat{\tau}_{ij}^l(\theta)\,\hat{l}_{ij}(\theta)\right]$$

The term $\sum_j c_j^G \frac{d\hat{G}_{ij}}{d\theta}$ 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.

As is well-known, 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:

$$\frac{d}{d\theta}\left[\left(\sum_{j=1}^{J_X} \hat{\tau}_{ij}^x(\theta)\,\hat{x}_{ij}(\theta) + \sum_{j=1}^{J_L} \hat{\tau}_{ij}^l(\theta)\,\hat{l}_{ij}(\theta)\right)\right] = \underbrace{\left(\sum_j^{J_X} \hat{x}_{ij}\frac{d\hat{\tau}_{ij}^x}{d\theta} + \sum_j^{J_L} \hat{l}_{ij}\frac{d\hat{\tau}_{ij}^l}{d\theta}\right)}_{\substack{\text{Mechanical Impact}\\ \text{on Govt Revenue}}} + \underbrace{\left(\sum_j^{J_X} \hat{\tau}_{ij}^x\frac{d\hat{x}_{ij}}{d\theta} + \sum_j^{J_L} \hat{\tau}_{ij}^l\frac{d\hat{l}_{ij}}{d\theta}\right)}_{\substack{\text{Behavioral Impact}\\ \text{on Govt Revenue}}} \quad (6)$$

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 Definition of Welfare

Moving from positive to normative analysis requires a definition of the normative objective. The measure of individual welfare adopted here will be the individual's willingness to pay out of their own income to follow the policy path.[24] Social welfare is then a weighted sum of individual welfare, with weights given by the social marginal utilities of income.

To be more specific, 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

[23] I do not model explicitly the source of non-budget neutrality, but one can extend the model to a world in which the government issues debt, B, and even allow B to affect behavior, $u(\mathbf{x}, \mathbf{l}, \mathbf{G}, B)$. I discuss this further in relation to the definition of the MCPF in footnote 41.

[24] Alternatively, one could evaluate the marginal excess burden of the policy change – this is discussed below in Section 2.6.

marginal policy change is given by $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$.[25]

With this definition of individual welfare, aggregation to social welfare is straightforward: one can take a weighted sum of individual willingness to pay, with the weights given by the social marginal utilities of income, $\frac{d\hat{W}}{d\theta}|_{\theta=0} = \int_{i\in I} \eta_i \frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} di$.[26] 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. Intuitively, to the extent to which $\eta_1 < \eta_2$, society is willing to lose resources in order to make a transfer from individual 1 to individual 2. While $\frac{d\hat{W}}{d\theta}|_{\theta=0}$ is measured in units of social utility, it can be normalized by $\eta_{\hat{i}}$ so that it is measured in units of individual i's income. For this, I define $\frac{d\hat{W}^{\hat{i}}}{d\theta}|_{\theta=0} = \frac{\frac{d\hat{W}}{d\theta}|_{\theta=0}}{\eta_{\hat{i}}} = \int_{i\in I} \frac{\eta_i}{\eta_{\hat{i}}} \frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} di$, where the superscript $\hat{i}$ denotes the fact that social welfare is measured in units of $\hat{i}$'s income.

2.4 The Components of Welfare

With these definitions, Proposition 1 characterizes the marginal welfare gain to individual i from pursuing the policy.

Proposition 1. *The marginal welfare impact to individual i of pursuing policy path $P(\theta)$ is given by:*

$$\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} = \left(\underbrace{\frac{d\hat{t}_i}{d\theta}|_{\theta=0}}_{\textit{Net Resources}} + \underbrace{\sum_{j=1}^{J_G} \left(\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} - c_j^G \right) \frac{d\hat{G}_{ij}}{d\theta}|_{\theta=0}}_{\substack{\textit{Public Spending/} \\ \textit{Mkt Failure}}} + \underbrace{\left(\sum_j^{J_X} \tau_{ij}^x \frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} + \sum_j^{J_L} \tau_{ij}^l \frac{d\hat{l}_{ij}}{d\theta}|_{\theta=0} \right)}_{\substack{\textit{Behavioral Impact} \\ \textit{on Govt Revenue}}} \right)$$

[25] 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

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

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

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

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

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

[26] 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 $\eta_i = G'\left(\hat{V}_i(0)\right)\lambda_i$).

Proof. The proof is an application of the envelope theorem and is provided in Appendix B.1. □

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, $\sum_{j=1}^{J_G} \left(\frac{\frac{\partial u_i}{\partial G_{ij}}}{\lambda_i} - c_j^G \right) \frac{d\hat{G}_{ij}}{d\theta}|_{\theta=0}$. 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, then one needs to know the difference between the costs and benefits of its provision.

The final term in Proposition 1 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 *causal effects* of the policy on behavior locally around the status quo, $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0}$ and $\frac{d\hat{l}_{ij}}{d\theta}|_{\theta=0}$, with the weights given by the marginal tax rates.

The causal effect matters because of a fiscal externality. The envelope theorem guarantees that 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.[27] Conditional on calculating this fiscal externality, behavioral responses are not required for welfare analysis.[28]

Example 1. *Assume there is one publicly-provided good, G, called roads. There is one untaxed consumption good, x, and there is one labor supply variable, l, which has a labor tax of τ^l. Assume there is only one type of agent (drop i subscripts). Also, assume there is no lump-sum taxation, $T = 0$.*

[27] 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.

[28] For completeness, it is also important to note that a decomposition of causal effects into income and substitution effects do not 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)).

One exception is the model of Chetty (2008) who models unemployment durations with a separable effort function and a binary state. He shows that the size of the market failure (wedge between marginal utilities) is a function of the causal impact of assets on search (liquidity effect) and the causal impact of unemployment benefits on search (moral hazard). Of course, it is not a general feature of economic models that marginal utilities can be written as functions of elasticities. Generally, marginal utilities are equated to prices, and elasticities correspond to the impact of price changes.

Normalize θ to parameterize an increase in spending on roads, so that $\hat{G}(\theta) = G + \theta$ and thus $\frac{d\hat{G}}{d\theta} = 1$. To impose budget neutrality, assume the marginal tax revenue (obtained from increasing the tax on labor supply) is spent on roads,

$$\tau^l \frac{d\hat{l}}{d\theta} + \hat{l}\frac{d\tau^l}{d\theta} = \frac{d\hat{G}}{d\theta} = 1 \quad \forall\theta$$

In this environment, Proposition 1 implies that the marginal welfare impact is positive if and only if

$$\left(\frac{\frac{\partial u}{\partial g}}{\lambda} - c_g\right) \geq -\hat{\tau}^l \frac{d\hat{l}}{d\theta}|_{\theta=0} \tag{7}$$

where the LHS is the net willingness-to-pay for additional roads, $\hat{\tau}^l$ is the marginal tax rate on labor supply, and $\frac{d\hat{l}}{d\theta}|_{\theta=0}$ is the causal impact of the policy on labor supply. It is the response that would be observed if the policy were undertaken to increase G financed by an increase in τ^l.[29]

The desirability of additional roads depends on how they affect government revenue. If roads increase labor supply because they make it easier to get to work, then the policy response is smaller; if roads increase the value of leisure and decrease taxable income, this makes roads less socially desirable (not because the planner doesn't value leisure, but because the government has a stake in the labor earnings).

In practice, the intuition in equation (7) can be useful for bounding the welfare gain of a policy. For example, Baird et al. (2012) estimate that a de-worming program in Kenya led to an increase in income tax revenue (from improved health and labor supply) that was sufficient to cover the program costs (i.e. $\tau\frac{dl}{d\theta}|_{\theta=0} > c_g$). Under the mild assumption that individuals preferred being offered the de-worming program, $\frac{\partial u}{\partial g} > 0$, one can conclude the program improved welfare without fully estimating the welfare benefits of the program.

Proposition 1 shows that the type of behavioral responses required depends on the policy in question. For example, if a policy increases marginal tax rates on individual i and provides no compensation, it is an uncompensated response; if it compensates agents for their tax increase, it is a compensated response; if a policy increases tax rates to finance increased education spending, one needs to incorporate not only the impact of the increased taxes on behavior, but also incorporate the impact of the simultaneous increase in education spending on behavior that affects the government's budget. To provide terminology to distinguish the desired responses from Hicksian or Marshallian

[29] *In general, $\frac{d\hat{l}}{d\theta}|_{\theta=0}$ is neither a Marshallian nor a Hicksian response. Indeed, one can write the RHS of equation (7) using a set of Marshallian elasticities and arrive at the optimality condition provided by Atkinson and Stern (1974). Let $l^*(\tau^l, G)$ denote the solution to the agent's maximization program given taxes on labor, τ^l, and government spending G. Also, following Atkinson and Stern (1974), assume that $\tau^l l = G$, so that the government has no other spending other than on G. Then, it is easy to show that*

$$\hat{\tau}^l \frac{d\hat{l}}{d\theta}|_{\theta=0} = \frac{\epsilon^m_{l^*,\tau} + \epsilon^m_{l^*,G}}{1 + \epsilon^m_{l^*,\tau}}$$

where $\epsilon^m_{l^*,\tau}$ is the standard marshallian elasticity of labor supply with respect to the labor tax rate, holding G fixed; and $\epsilon^m_{l^*,G}$ is the elasticity of l^* with respect to G, holding τ^l fixed. So, the policy elasticity can be computed from these two marshallian elasticities. But, such a decomposition is not necessary; the policy elasticity is sufficient.

price responses, 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, I define the *policy elasticity* of x_{ij} and l_{ij} to be the local causal effect of the policy on $log\left(x_{ij}\right)$ and $log\left(l_{ij}\right)$.

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

Given these definitions, the behavioral impact term of Proposition 1 has three representations:

$$\underbrace{\frac{d}{d\theta}\left(\sum_{j=1}^{J_X}\hat{\tau}^x_{ij}\hat{x}_{ij}+\sum_{j=1}^{J_L}\hat{\tau}^l_{ij}\hat{l}_{ij}\right)-\left(\sum_{j}^{J_X}\hat{x}_{ij}\frac{d\hat{\tau}^x_{ij}}{d\theta}+\sum_{j}^{J_L}\hat{l}_{ij}\frac{d\hat{\tau}^l_{ij}}{d\theta}\right)}_{\text{Total - Mechanical Impact on Govt Revenue}}=\underbrace{\left(\sum_{j}^{J_X}\tau^x_{ij}\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0}+\sum_{j}^{J_L}\tau^l_{ij}\frac{d\hat{l}_{ij}}{d\theta}|_{\theta=0}\right)}_{\substack{\text{Behavioral Impact}\\ \text{on Govt Revenue (levels)}}}=\underbrace{\left(\sum_{j}^{J_X}r^x_{ij}\hat{\epsilon}^x_{ij}+\sum_{j}^{J_L}r^l_{ij}\hat{\epsilon}^x_{ij}\right)}_{\substack{\text{Behavioral Impact}\\ \text{on Govt Revenue (logs)}}} \tag{8}$$

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

The representations in equation (8) suggest there are multiple potential empirical strategies one can use to estimate the impact of the behavioral response to the policy on the government's budget. First, one could attempt to estimate the fiscal externality directly. If one had a counterfactual budget forecast of what the government budget would be in the absence of any behavioral responses (the "mechanical impact on government revenue" in equation (6)), one could compare the difference in the realized budget and the mechanical revenue that would have been observed in the absence of behavioral responses.[30] Second, one could estimate the micro-level behavioral changes $\mathbf{x}_i$ and $\mathbf{l}_i$ resulting from the policy and multiply by the government's stake in the behavior. In this micro approach, one can either use policy responses and marginal tax rates (levels), or using policy elasticities and government revenues on each activity (logs).

2.5 Which Policy Elasticities Are Necessary?

Proposition 1 includes the policy responses of all goods by all individuals, $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0}$, and $\frac{d\hat{l}_{ij}}{d\theta}|_{\theta=0}$. However, this requirement can be reduced in many ways depending on the setting. Clearly, one does not need to know how a policy changes the choice of untaxed goods or labor. Moreover, one can aggregate responses for goods (or labor supply activities) with the same marginal tax rate. To see this, note that if $\tau_1 = \tau_2$, then

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

In particular, if the government has only one marginal tax on all forms of taxable income and no taxes on goods, then the change in taxable income for each type i is sufficient. Moreover, one can aggregate responses across types with equal social marginal utilities of income: if $\eta_{i_1}\lambda_{i_1} = \eta_{i_2}\lambda_{i_2}$, then

[30] As discussed further in Section (4), this approach is taken by Chetty et al. (2013) who estimate the marginal incentives from the EITC schedule increase EITC expenditures by 5%.

the aggregate responses for types i_1 and i_2 (e.g. $\frac{d\left(x_{i_1 j}+x_{i_2 j}\right)}{d\theta}|_{\theta=0}$ for each j) are sufficient for each individual's response to the policy.

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. This insight was recently popularized in Feldstein (1999). I provide two clarifications to this result.[31] 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). Hence, subsequent literature tends to suggest a need for adding additional elasticities to the analysis.[32]

In contrast to these approaches, 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.[33] Of course, there are many reasons to be interested in the mechanisms driving such a response[34]; but calculating the marginal welfare impact of the policy change in question is not one of them.

2.6 Relation to MEB

In contrast to the calculating the individuals willingness to pay to pursue the policy path, $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$, a common method for evaluating government policy changes is to calculate its marginal excess burden.[35] To compute the marginal excess burden to individual i, let $\mathbf{v} = (v_i)$ be a vector of pre-specified utilities. Most commonly, $\mathbf{v}$ is chosen to be the set of status quo utilities, which will correspond to the "equivalent variation" measure of MEB.[36] Now, define the *compensated* policy path, $P^{\mathbf{v}}(\theta)$, such that

[31] These clarifications are distinct from the insight of Chetty (2009a) who shows that the aggregate taxable income elasticity is not sufficient if the private marginal cost of tax avoidance is not equal to its social marginal cost.

[32] 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.

[33] Indeed, this is the approach taken in 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.

[34] For example, one may wish to forecast causal effects of policies not yet undertaken.

[35] For a recent application of this methodology to the study of EITC expansions in the US, see Eissa et al. (2008).

[36] 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

$P^v = \left\{ \hat{\tau}^{\mathbf{x}}_{\mathbf{i}}(\theta), \hat{\tau}^{\mathbf{l}}_{\mathbf{i}}(\theta), \hat{T}_i(\theta) + \hat{C}_i(\theta; v), \hat{\mathbf{G}}_{\mathbf{i}}(\theta) \right\}_i$ where $\hat{C}_i(\theta; u)$ is a compensation provided to agent i such that $V_i\left(\hat{\tau}^{\mathbf{x}}_{\mathbf{i}}(\theta), \hat{\tau}^{\mathbf{l}}_{\mathbf{i}}(\theta), \hat{T}_i(\theta) + \hat{C}_i(\theta; v), \hat{\mathbf{G}}_{\mathbf{i}}(\theta), y_i\right) = v_i$. Intuitively, $P^v(\theta)$ is the same as the proposed policy path 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 $\hat{t}^v_i$ 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

$$MEB_i^{v_i} = \frac{d\hat{t}^v_i}{d\theta}|_{\theta=0} \tag{9}$$

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. If the policy 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). Although this paper argues this is not the most empirically tractable measure, it is an intuitive welfare concept.

If $\mathbf{v}$ is the status quo vector 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 $\hat{x}^c_{ij}$ and $\hat{l}^c_{ij}$ 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: $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} - \frac{d\hat{x}^c_{ij}}{d\theta}|_{\theta=0}$.[37] 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:

$$MEB_i^{v_i} = \frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i} - \underbrace{\left(\sum_j^{J_X} \tau^x_{ij} \left(\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} - \frac{d\hat{x}^c_{ij}}{d\theta}|_{\theta=0} \right) + \sum_j^{J_L} \tau^l_{ij} \left(\frac{d\hat{l}_{ij}}{d\theta}|_{\theta=0} - \frac{d\hat{l}^c_{ij}}{d\theta}|_{\theta=0} \right) \right)}_{\text{Income Effects on Government Budget}}$$

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 $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} - \frac{d\hat{x}^c_{ij}}{d\theta}|_{\theta=0} = 0$ and $MEB_i^{v_i} = 0$ (and vice-versa).[38]

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

EV measures of MEB depend on whether one is starting from the perspective of the first best or from the status quo. Hence, some papers switch these two definitions around.

[37] The Slutsky equation guarantees that $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} - \frac{d\hat{x}^c_{ij}}{d\theta}|_{\theta=0} = \frac{dx^m_{ij}}{dy_i}\frac{d\hat{C}_i}{d\theta}$ where x^m_{ij} is the marshallian demand, $\frac{dx^m_{ij}}{dy_i}$ is the response to income, and $\frac{d\hat{C}_i}{d\theta}$ is the amount of compensation required to hold utility constant.

[38] 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.

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 D, 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.[39] Hence, around the optimum, the causal effects are compensated responses (i.e. $\frac{d\hat{x}_{ij}}{d\theta}|_{\theta=0} = \frac{d\hat{x}^c_{ij}}{d\theta}|_{\theta=0}$ because utility is not changing at the optimum). 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.

2.7 Extensions

Although the model allowed for considerable heterogeneity across individuals, it assumed a stylized model of production with perfect competition and fixed resource costs.[40] This rules out many phenomena that may be important for real-world welfare estimation but can easily be incorporated into the model. For example, by assuming real prices are always 1, the model ruled out general equilibrium effects and pecuniary externalities. 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 impact of the policy on the after-tax wage faced by individual i on her jth labor supply activity. These additional impacts can simply be added to the resource transfer term, $\frac{d\hat{t}_i}{d\theta}|_{\theta=0}$, in Proposition 1. 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 resource transfers in Proposition 1 accordingly. The causal effects are still the desired responses, but one needs to also know the general equilibrium effects of government policies.

Policy analysis becomes slightly more difficult when there are non-pecuniary externalities. Appendix C provides an extension of the model to the case where there is a variable (e.g. pollution) affecting the individuals utility that is a function of other individuals' behavior. In these cases, one requires the causal effect of the policy on the level of pollution; but in addition, one requires an esti-

[39] Diamond and Mirrlees (1971) also consider a model with heterogeneous agents and derive their tax rules in such settings. With heterogeneous agents, the formulae no longer depend on compensated responses (see Section VII, page 268). 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. Their optimal formula incorporate income effects into the social marginal utility weightings (See Auerbach and Hines (2002) for a simple derivation of this on page 1370, equation 3.24). Intuitively, Diamond and Mirrlees (1971) are adding back in the income effects that were taken out of the causal effects of feasible budget neutral changes to the commodity tax structure.

[40] Note that the aggregate impact of the policy on the value of production (i.e. GDP) does not enter the welfare calculation. This is not because of the stylized model of production per se. At the optimum, individuals trade off their private benefit from production (their after-tax wage) with their private cost of production (their disutility of labor supply activities). If production increases because of the policy, this envelope condition suggests individuals were privately indifferent to the change. Hence, such changes to production matters for welfare only through the impact on the government budget. However, if there are spillovers or externalities in the production process, one would need to account for the impact of the policies on these externalities in a manner analogous to the impact on the fiscal externality (see Appendix C).

mate of the individual's marginal rate of substitution between pollution and income, analogous to the net willingness to pay required to value the provision of publicly provided goods. Welfare analysis in these models is more complicated because of the difficulty in valuing the externality, but the policy elasticities continue to be the required behavioral responses.

3 Additivity and the Marginal Value of Public Funds

Many government policies are not budget neutral, at least in the short run. Naturally, one desires a coherent way of analyzing these non-budget neutral policies. This section provides a condition that allows the welfare impacts of policies to be added together. It leads to natural definition of the marginal value of public funds as the social welfare impact of the policy per dollar of government revenue expended.

To begin, suppose one is interested in characterizing the marginal welfare impact of a policy path, $P(\theta)$. Suppose that two policy paths, $P_{Tax}(\theta)$ and $P_{Exp}(\theta)$, sum to the policy path of interest, $P(\theta)$:

$$(P(\theta) - P(0)) = (P_{Tax}(\theta) - P(0)) + (P_{Exp}(\theta) - P(0)) \tag{10}$$

Condition (10) requires that the movement from the initial policy position, $P(0)$ towards $P(\theta)$ can be written as the sum of two movements: first in the direction of $P_{Tax}(\theta)$ and second in the direction of $P_{Exp}(\theta)$ (or vice-versa). This equality must hold for all components of the policy (taxes, transfers, and public provision of goods). For example, $P_{Exp}(\theta)$ could be a policy path that spends money from the government budget on a public good; $P_{Tax}(\theta)$ could be a policy that raises government revenue through increasing the labor tax rate. In this case, $P(\theta)$ would be a policy that simultaneously increases the labor tax rate and spends the resources on the public good.[41]

If equation (10) is satisfied, it is straightforward to show[42] that the marginal welfare impact of the

[41] The non-budget neutral policies, P_{Tax} and P_{Exp}, implicitly change government debt obligations. Intuitively, when the government implements non-budget neutral policies, it is either borrowing resources from its own citizens or from abroad (in an open economy). I do not explicitly model such borrowing, but it is important to note that one can augment the model to allow the level of government debt or obligations, B, to affect the agents' behavior, $u_i(\mathbf{x_i}, \mathbf{l_i}, \mathbf{G_i}, B)$. In this case, non-budget neutral policies can increase B; but when considering the sum of two non-budget neutral policies that sum to a budget neutral policy, one can ignore the impact of each individual policy on B, since on aggregate B remains unchanged in any budget neutral policy experiment.

[42]

Proof. Let ∇V_i denote the gradient of V_i, so that $\frac{\partial \hat{V}_i^P}{\partial \theta} = \nabla V_i^P \frac{dP}{d\theta}$, where $\frac{dP}{d\theta}$ is the vector of policy changes. Note that

$$\begin{aligned}
\frac{\partial V_i^P}{\partial \theta} &= \nabla V_i \frac{dP}{d\theta} \\
&= \nabla V_i \left(\frac{dP_{Tax}}{d\theta} + \frac{dP_{Exp}}{d\theta} \right) \\
&= \nabla V_i \frac{dP_{Tax}}{d\theta} + \nabla V_i \frac{dP_{Exp}}{d\theta} \\
&= \frac{\partial V_i^{P_{Tax}}}{\partial \theta} + \frac{\partial V_i^{P_{Exp}}}{\partial \theta}
\end{aligned}$$

where all derivatives are evaluated at $\theta = 0$. □

comprehensive policy on type i, denoted $\frac{\frac{\partial \hat{V}_i^P}{\partial \theta}|_{\theta=0}}{\lambda_i}$, is given by the sum of the two welfare impacts:

$$\frac{\frac{\partial \hat{V}_i^P}{\partial \theta}|_{\theta=0}}{\lambda_i} = \frac{\frac{\partial \hat{V}_i^{P_{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda_i} + \frac{\frac{\partial \hat{V}_i^{P_{Exp}}}{\partial \theta}|_{\theta=0}}{\lambda_i} \tag{11}$$

where $\frac{\frac{\partial \hat{V}_i^{P_{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda_i}$ and $\frac{\frac{\partial \hat{V}_i^{P_{Exp}}}{\partial \theta}|_{\theta=0}}{\lambda_i}$ denote the marginal welfare impact of the component policies, P_{Tax} and P_{Exp}.

Despite being straightforward in the present framework, equation (10) is not innocuous from the perspective of the MEB framework. For example, suppose one were to take a MEB calculation for a tax increase from existing literature as P_{Tax}. This hypothetical policy involves the government collecting or providing individual-specific lump-sum transfers in a manner that holds utility constant. Hence, for the additivity condition to hold there are two options depending on whether one seeks a comprehensive MEB estimate or a $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$ estimate. For a MEB calculation, the expenditure policy must also hold utility constant while raising taxes via individual-specific lump-sum to finance the expenditure. For a calculation of individuals' willingness to pay, $\frac{\frac{\partial \hat{V}_i^P}{\partial \theta}|_{\theta=0}}{\lambda_i}$, one must consider an expenditure policy that not only provided the expenditure but also removed the lump-sum transfers that were provided in the tax policy. In both cases, the causal effects of the expenditure policy are not sufficient for the behavioral responses required to compute the welfare impact of the comprehensive policy, even conditional on knowing the MEB of the tax policy. In contrast, if one uses the measures of welfare in the present framework, $\frac{\frac{\partial \hat{V}_i^{P_{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda_i}$ and $\frac{\frac{\partial \hat{V}_i^{P_{Exp}}}{\partial \theta}|_{\theta=0}}{\lambda_i}$, the causal effects of the tax and expenditure policies are sufficient.[43]

The additivity condition in equation (11) suggests a natural method for dealing with non-budget neutral policies. One can simply compute the welfare cost per dollar of government budget expended, which captures a measure of the marginal value of public funds (MVPF). Normalizing social welfare into units of individual $\hat{i}$'s income, the MVPF is given by:

$$MVPF_P^{\hat{i}} = \frac{\int_i \frac{\eta_i}{\eta_{\hat{i}}} \frac{\frac{d\hat{V}_i^P}{d\theta}|_{\theta=0}}{\lambda_i} di}{\int_{i \in I} \frac{d\hat{t}_i^P}{d\theta} di} \tag{12}$$

[43] In the context of tax policies, it is interesting to note that causal effect of tax increases may be either a pure Hicksian response, an uncompensated response, or neither. If agents expect the increased revenue to be returned through future transfers or publicly provided goods and then borrow against these in capital markets (i.e. Ricardian equivalence holds), then the behavioral response may be similar to a compensated response. In contrast, the uncompensated approach may describe behavior if people do not expect future tax revenue or do not borrow against these future benefits. Indeed, whether or not the policy response is compensated or uncompensated arguably depends the degree to which Ricardian equivalence holds and how people respond to government debt. Of course, I do not explicitly model government debt. But, as eluded to in Footnote 41, comparisons of the values of MVPF are implicitly constructing budget neutral policies (e.g. $MCPF_{P_1}^{\hat{i}} - MCPF_{P_2}^{\hat{i}}$ is the welfare impact of taking \$1 along policy path P_2 and using it to increase spending along policy path P_1). Hence, the combined policy is budget neutral so that one need not isolate the particular impact of government debt on behavior and utility.

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 $\hat{i}$.[44] There is an extraordinary number of different definitions for the MCPF in previous literature (Fullerton (1991); Auerbach and Hines (2002); Dahlby (2008)). The particular definition in equation (12) was initially proposed as a marginal cost/benefit of public funds in Slemrod and Yitzhaki (1996).[45] The the key advantage of this definition of the MVPF is that the behavioral responses depend solely on the causal, not compensated, effects of the non-budget neutral policies in question.[46]

Given any two policies, P_{Tax} and P_{Exp}, satisfying equation (10), the additivity condition implies

$$\frac{d\hat{W}^{P}}{d\theta} = \eta^{\hat{i}} \left(MVPF_{P_{Exp}}^{\hat{i}} - MVPF_{P_{Tax}}^{\hat{i}} \right) \tag{13}$$

so that policy P_{Exp} provides a benefit of $MVPF_{P_{Exp}}^{\hat{i}}$ per dollar of government revenue and a cost of $MVPF_{P_{Tax}}^{\hat{i}}$ per dollar of government revenue. If $MVPF_{P_{Exp}}^{\hat{i}}$ is greater (less) than $MVPF_{P_{Tax}}^{\hat{i}}$, 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.

I illustrate this definition using Example 1.

Example. *(Example 1 Continued) Consider the welfare cost a policy $P_{Tax}(\theta)$ that raises θ units of revenue through a tax on labor supply, $\hat{\tau}(\theta)$.*[47] *The marginal welfare impact of this policy is*

$$\frac{\frac{\partial \hat{V}^{P_{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda} = -1 + \tau \frac{d\hat{l}_{P_{Tax}}}{d\theta}|_{\theta=0} \tag{14}$$

where the "-1" arises from the net negative transfer, and $\frac{d\hat{l}_{P_{Tax}}}{d\theta}|_{\theta=0}$ is the behavioral response to the tax policy that increases government revenue. Recall there is a single agent so that the MVPF does not depend on the choice of income units, $\hat{i}$. Moreover, $\frac{d\hat{t}}{d\theta} = -1$ because the policy raises θ units of revenue. So, the MVPF of the tax policy is given by

$$\begin{aligned} MVPF_{P_{Tax}} &= \frac{\frac{\frac{\partial \hat{V}^{P_{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda}}{\frac{d\hat{t}}{d\theta}} \\ &= \frac{-1 + \tau \frac{d\hat{l}_{P_{Tax}}}{d\theta}|_{\theta=0}}{-1} \\ &= 1 - \tau \frac{d\hat{l}_{P_{Tax}}}{d\theta}|_{\theta=0} \end{aligned}$$

[44] Note that the $\hat{i}$ notation makes clear 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.

[45] A version of this approach was also implemented in Immervoll et al. (2007) using quasilinear utility, so that the distinction between causal and compensated effects was important.

[46] This was perhaps unclear in Slemrod and Yitzhaki (1996) because they refer to the behavioral responses using deadweight loss and excess burden terminology.

[47] For simplicity, I normalize the speed of the path so that $\frac{d\hat{t}}{d\theta} = -1$

Intuitively, the marginal cost of public funds is given by one plus the causal impact of the response to taxation on the government's budget constraint.

Now, let $P_{Exp}(\theta)$ denote a policy that spends $\hat{G}(\theta) = G + \theta$ on additional roads. Then,

$$\frac{\frac{\partial \hat{V}^{P_{Exp}}}{\partial \theta}|_{\theta=0}}{\lambda} = \left(\frac{\frac{\partial u}{\partial g}}{\frac{\partial u}{\partial x}} - c_g\right) + 1 + \tau \frac{d\hat{l}_{P_{Exp}}}{d\theta}|_{\theta=0} \tag{15}$$

and, since $\frac{d\hat{t}}{d\theta} = 1$,

$$MCPF_{P_{Exp}} = \left(\frac{\frac{\partial u}{\partial g}}{\frac{\partial u}{\partial x}} - c_g\right) + 1 + \tau \frac{d\hat{l}_{P_{Exp}}}{d\theta}|_{\theta=0}$$

where $\left(\frac{\frac{\partial u}{\partial g}}{\frac{\partial u}{\partial x}} - c_g\right)$ is the net willingness to pay for the roads and "1" arises from the net positive transfer. The last term, $\tau \frac{d\hat{l}_{P_{Exp}}}{d\theta}$ is the impact of the behavioral response to the increased expenditure on roads on the government's budget. This term would be positive if roads increased labor supply; negative if it caused people to take more vacations and reduce labor earnings.

Combining equations (14) *and* (15),

$$\begin{aligned}
\frac{\frac{\partial \hat{V}^{P}}{\partial \theta}|_{\theta=0}}{\lambda} &= MVPF_{P_{Exp}} - MVPF_{P_{Tax}} \\
&= \left(\frac{\frac{\partial u}{\partial g}}{\frac{\partial u}{\partial x}} - c_g\right) + \tau \left(\frac{d\hat{l}_{P_{Tax}}}{d\theta}|_{\theta=0} + \frac{d\hat{l}_{P_{Exp}}}{d\theta}|_{\theta=0}\right) \\
&= \left(\frac{\frac{\partial u}{\partial g}}{\frac{\partial u}{\partial x}} - c_g\right) + \tau \frac{d\hat{l}_{P}}{d\theta}|_{\theta=0}
\end{aligned}$$

where $\frac{d\hat{l}_P}{d\theta}|_{\theta=0}$ is the joint effect of the expenditure and taxation policy on labor supply. Hence, $\frac{\frac{\partial \hat{V}^{P}}{\partial \theta}|_{\theta=0}}{\lambda}$ is precisely equal to the total welfare impact given in equation (7)*. If $MVPF_{P_{Exp}}$ is greater (less) than $MVPF_{P_{Tax}}$, then increasing (decreasing) taxes through the tax policy P_{Tax} and decreasing (increasing) expenditures through the P_{Exp} policy will increase social welfare.*

Relation to previous definitions As mentioned, the definition of the MVPF is proposed by Slemrod and Yitzhaki (1996, 2001). However, it is conceptually distinct from the two main traditions in the MCPF literature (see Dahlby (2008) for a recent overview and Fullerton (1991) for evidence these conceptual differences lead to different numerical estimates). The so-called Pigou-Harberger-Browning tradition (Pigou (1947); Harberger (1964); Browning (1976, 1987)) uses the MEB as the measure of the MCPF. As mentioned above, this requires the expenditure policy to be financed using lump-sum taxation and that one incorporates these income effects. The so-called Stiglitz-Dasgupta-Atkinson-Stern tradition 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 Proposition 1 for a policy that raises taxes and increases spending on G). In practice, many papers estimating the marginal cost of public funds assume that the expenditure has a separable impact on utility and hence does not have an associated fiscal externality (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 these expenditures is to capture fiscal externalities.[48]

But a more general conceptual difference is that it is unclear why one desires a single measure of the cost of raising revenue to finance projects. 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 to the Stiglitz-Dasgupta-Atkinson-Stern definition of the MCPF, the MVPF proposed here 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. Moreover, the behavioral responses required to answer such questions are precisely the causal effects of each of the policies in question.

4 Redistribution

Many if not all government policy changes involve distributional tradeoffs. In this section, I illustrate how the framework can be used to frame such policy questions using the classic bucket experiment (Okun (1975)).

To be more specific, I consider the welfare impact of a policy that would increase the generosity of the earned income tax credit (EITC) to poor single mothers financed by an increase in the top marginal income tax rate.[49] A benefit of applying the framework to the case of pure redistribution through taxation is that I do not need to estimate the value of any changes to publicly provided goods.[50] However, the desirability of redistributing from rich to poor will depend on their relative social welfare weights. Because many may disagree about such parameters, I will not solve directly for the social welfare impact of the policy. Rather, I solve for the set of implicit social marginal utilities of income that rationalize the status quo amount of redistribution as optimal, $\frac{d\hat{W}_P}{d\theta} = 0$. If one's own social preferences are more (less) redistributive than these implicit weights, then one would prefer a more (less) redistributive policy. From a positive perspective, the approach will illustrate how much the current U.S. income tax structure implicitly values money in the hands of the poor relative to the

[48] Appendix A calculates the MVPF for the Job training Partnership Act, which had significant impacts on tax receipt and reductions in welfare payments (Bloom et al. (1997)).

[49] The EITC program program provides benefits to groups other than single mothers. However, most previous literature has focused on the causal impact of expansions to the EITC program on single mothers. To align my policy with these existing causal estimates, I consider an expansion of the program targeted solely to single mothers.

[50] As shown in Appendix A, welfare analysis is more difficult for public programs like job training programs because one must recover the willingness to pay for the public goods or services in excess of their costs.

rich.[51]

4.1 Setup

Let $P(\theta)$ denote the policy where θ dollars are raised from the rich through an increase in the top marginal tax rate on ordinary income that are then transferred to poor single mothers through the an increase in the size of the EITC.[52] Let $\hat{l}_i(\theta)$ denote the taxable income of individual i subject to the standard income tax rate (i.e. $\hat{l}_i$ excludes dividends) and let $\bar{l}$ denote the threshold above with this income is taxed at the top rate, $\hat{\tau}^{Rich}(\theta)$. It will be helpful to classify individuals, i, into two (non-exhaustive) groups: $i \in Rich$, for whom $\hat{l}_i(0) \geq \bar{l}$, where $\bar{l} \approx \$400K$, and $i \in Poor$, who are low-income single mothers currently eligible for EITC benefits, generally $\hat{l}_i(0) \leq \$40K$.

Importantly, I allow the social marginal utility of income to differ between rich and poor. However, within the set of rich and poor, I make the simplifying assumption that the social marginal utilities of income are the same. Let η^{Poor} denote the social marginal utility of income for a poor individual and let η^{Rich} denote the social marginal utility of income for a rich individual. Let $\hat{W}_P(\theta)$ denote the social welfare under the policy $P(\theta)$. Under these two simplifications, the desirability of redistribution is characterized in the following proposition.

Proposition 2. $\frac{d\hat{W}_P}{d\theta}|_{\theta=0} \geq 0$ *if and only if*

$$\underbrace{\frac{\eta^{Rich} - \eta^{Poor}}{\eta^{Poor}}}_{\substack{\textit{Social Preference} \\ \textit{for Add'l Redistribution}}} \leq \underbrace{\frac{\int_i \left(\sum_j \tau^x_{ij} \frac{dx_{ij}}{d\theta}|_{\theta=0} + \tau^l_{ij} \frac{dl_{ij}}{d\theta}|_{\theta=0} \right) di}{\int_{i \in Rich} \frac{d\hat{\tau}^{Rich}}{d\theta}|_{\theta=0} \left(\hat{l}_i - \bar{l} \right) di}}_{\textit{Leaks in Okun's Bucket}} \tag{16}$$

Proof. The proof is a straightforward application of Proposition 1, and is provided in Appendix B.2. □

The LHS of equation (16) measures the marginal benefit to social welfare of transferring money from rich to poor. The RHS of equation (16) is the fraction of the mechanical revenue raised by taxing the rich that is lost due to behavioral responses.

[51] This 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.

[52] To align the precise EITC policy with past EITC expansions, I consider an increase in 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 intuition in equation (16) is Okun's classic leaky bucket experiment (Okun (1975)): one's preference for redistribution can be stated as how much resources one is willing to lose in order to take from the rich and give to the poor. Equation (16) is also a generalization of the standard Baily-Chetty formula for the optimal amount of social insurance (Baily (1978); Chetty (2006)). At the optimum, the value of transferring money from rich to poor (given by the difference in social marginal utilities) is equated to its cost (given by Okun's bucket). Equation (16) depends on the causal, not compensated, effects of the redistributive policy. As a result, the MEB of the redistributive policy does not provide a social welfare-based guide into the question of whether additional redistribution is desirable without adding back in the income effects.

Equation (16) also differs from approaches in previous literature studying optimal taxation with heterogeneous agents beginning with Mirrlees (1971) and pioneered empirically by Saez (2001). Although the first order conditions of the Hamiltonian provide insight into the optimal *slope* of the tax schedule, the optimal *level* of the schedule is identified from the transversality condition (i.e. budget constraint). Hence, the optimal level of redistribution to the poor depends on an integral of elasticities across the income distribution, evaluated at their optimized levels (Piketty and Saez (2012)). Such an integral is difficult to estimate in practice. As a result, optimal tax formulas often only focus on optimal tax rates, not the level of redistribution (Piketty and Saez (2012)). In contrast, Equation (16) does not provide information about the optimality of the entire tax structure. However, it does yield a fairly simple formula that characterizes whether the level of redistribution to the poor should be increased through the particular redistributive policy in question.

4.2 Empirical implementation: A MVPF Approach

The leaks in Okun's bucket in equation (16) depend on the policy elasticities for a policy that simultaneously increases EITC benefits and raises the top marginal income tax rate on ordinary income. In practice, the causal effects studied in the literature focus on each policy independently. Therefore, I use the additivity condition to write the comprehensive policy as the sum of two policies: an increase in EITC generosity by \$1, P^{EITC}, that is financed out of government revenue; and a raising of the top marginal income tax rate, P^{Tax}, that is used to increase government revenue by \$1.

Both of these non-budget neutral policies induce a marginal cost of public funds. To raise \$1 in tax revenue from taxes on the rich, one imposes a welfare loss on the rich given by

$$MVPF^{Rich}_{P^{Tax}} = \frac{\frac{\frac{\partial \hat{V}^{Rich}_{P^{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda^{Rich}}}{\int_{i \in I} \frac{d\hat{t}^{P^{Tax}}}{d\theta} di}$$

which does not depend on social marginal utilities of income because we've assumed these are constant amongst the rich. Similarly, to raise \$1 in tax revenue through a reduction in EITC benefits, one

imposes a welfare loss on the poor given by

$$MVPF^{Poor}_{P^{EITC}} = \frac{\frac{\frac{\partial \hat{V}^{Poor}_{P^{EITC}}}{\partial \theta}|_{\theta=0}}{\lambda^{Poor}}}{\int_{i \in I} \frac{d\hat{t}^{P^{EITC}}}{d\theta} di}$$

which again does not depend on social marginal utilities of income because we have assumed these are constant amongst the poor.

Using equation (13) and the ratio of social marginal utilities of income, $\frac{\eta^{Rich}}{\eta^{Poor}}$, to translate $MCPF^{Rich}_{P^{Tax}}$ into units of income to the poor, the welfare impact of additional redistribution is given by

$$\frac{d\hat{W}_P}{d\theta}|_{\theta=0} = \eta^{Poor} MVPF^{Poor}_{P^{EITC}} - \eta^{Rich} MVPF^{Rich}_{P^{Tax}}$$

which yields the following Corollary to Proposition (2).

Corollary 1. $\frac{d\hat{W}_P}{d\theta}|_{\theta=0} \geq 0$ *if and only if*

$$MVPF^{Poor}_{P^{EITC}} - \frac{\eta^{Rich}}{\eta^{Poor}} MVPF^{Rich}_{P^{Tax}} \geq 0 \tag{17}$$

Whether additional redistribution is desirable depends on whether the marginal value of the expenditure, given by $MVPF^{Poor}_{P^{EITC}}$, is greater than the cost, given by $MVPF^{Rich}_{P^{Tax}}$. Since the MVPF of the tax increase on the rich is defined as the willingness to pay out of income of the rich, one needs to multiply the social marginal utility of income of the poor, $\frac{\eta^{Rich}}{\eta^{Poor}}$, so that the tax policy is evaluated in units of income to the poor. I consider the calculation of $MVPF^{Rich}_{P^{Tax}}$ and $MVPF^{Poor}_{P^{EITC}}$ in turn.

4.2.1 Tax Increase on Rich

There is a large literature estimating the causal effect of changes to the top marginal income tax rate (see Saez et al. (2012) for a recent review). To construct an estimate of the impact of the behavioral response to such tax rate increases on the government's budget, I make several assumptions that are common in this empirical literature. 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 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

$$MVPF^{Rich}_{PTax} = \frac{1}{1+r}$$

where r is the fraction of mechanical ordinary income tax revenue lost from behavioral responses to the tax increase,

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

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.[53] Note $r < 0$ if behavioral responses lower tax revenue.

There is a large literature focused on estimating r by studying the impact of changes in the top marginal income tax rate, such as the Omnibus Budget Reconciliation Act of 1993 (a.k.a. the Clinton tax increases). For example, Saez et al. (2012) show how r 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 r is constant as the tax rate changes. In contrast, estimating $MVPF^{Rich}_{PTax}$ relies on local estimates of r for variation in taxes around the status quo. Moreover, this literature also often attempts to decompose responses into income and substitution effects.[54] Indeed, the parameter r is sometimes referred to as the "marginal excess burden" of the income tax (Saez et al. (2012); Feldstein (1999)).

In contrast, for my analysis I prefer the value of r corresponding to the causal effect of the policy that changes the top marginal tax rate. Hence, the estimates of r in previous literature, which are derived from causal effects of policies that vary the top marginal income tax rate, are arguably better suited for my welfare framework than for estimating the marginal excess burden.

Saez et al. (2012) and Giertz (2009) note that there is a wide range of estimates of the fraction of mechanical revenue that is lost due to behavioral responses to changes in the top marginal income tax rate. However, they suggest mid-points ranging from 25-50%, which implies $MVPF^{Rich}_{PTax}$ is between 1.33 and 2.[55]

[53] To see this, note that

$$\begin{aligned}
\frac{\frac{\frac{\partial \hat{V}^{Rich}_{PTax}}{\partial \theta}|_{\theta=0}}{\lambda^{Rich}}}{\int_i \frac{d\hat{\iota}^{PTax}}{d\theta} di} &= -\frac{\int_{i \in Rich} \frac{d\tau^{Rich}}{d\theta}|_{\theta=0} \left(\hat{l}_i - \bar{l}\right) di}{\int_{i \in Rich} \tau_i^l \frac{d\hat{l}_i^{Tax}}{d\theta}|_{\theta=0} di + \int_i \left(\sum_j \tau_{ij}^x \frac{d\hat{x}_{ij}^{Tax}}{d\theta}|_{\theta=0} + \tau_{ij}^l \frac{d\hat{l}_{ij}^{Tax}}{d\theta}|_{\theta=0}\right) di} \\
&= \frac{1}{1+r}
\end{aligned}$$

[54] For example, Gruber and Saez (2002) assume constant income and price elasticities across the income distribution so that heterogeneous responses to tax rate changes can separately identify the income and substitution effects.

[55] Identifying the behavioral responses of top earners to taxation is a difficult empirical exercise (but perhaps easier than decomposing these responses into income and substitution effects). As a result, I will also show how the conclusions differ for a broader range of parameters.

4.2.2 EITC Expansion

There is also 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. So, 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 (1) individuals above the income eligibility threshold and (2) low-income women choosing to become single mothers to become EITC eligible. Support for (1) is found in Chetty et al. (2013) who find minimal effects of behavioral responses in the so-called "phase-out" region of earnings above the refund-maximizing earnings level. Support for (2) is found in Hotz and Scholz (2003) who summarize the empirical literature as finding little or no effects on marriage and family formation.[56] Both of these assumptions could easily be relaxed with precise estimates of the impact of the behavioral responses of these groups to EITC expansions on its budgetary cost.

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. To the extent to which an EITC expansion crowds out take-up other government services, the analysis will underestimate the social desirability of increasing funding of the EITC.

With these assumptions, one obtains an expression analogous to the tax policy:

$$MVPF^{Poor}_{P^{EITC}} = \frac{1}{1+p}$$

where p is the fraction of the mechanical revenue distributed that is increased due to behavioral distortions,

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

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 p. Short of this, I take estimates of the extensive and intensive margin labor supply response to the EITC to construct an estimate of the associated fiscal externality. To begin, I need to generalize the model slightly to incorporate extensive margin (i.e. discontinuous) responses.[57]

[56] A further defense of this assumption is found in the EITC papers using single women without children as a control group (e.g. Eissa and Liebman (1996); Chetty et al. (2013)).

[57] See Kleven and Kreiner (2006) for a generalization of MEB calculations to account for extensive margin responses to the EITC.

Extensive margin responses 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:

$$-\int_{i\in Poor} \tau_i^l \frac{d\hat{l}_i^{EITC}}{d\theta}|_{\theta=0} di \quad = \quad \underbrace{\left(\tau^l_{i^{LFP}(0)} l_{i^{LFP}(0)}\right) \frac{di^{LFP}}{d\theta}|_{\theta=0}}_{\text{Extensive Margin}} - \underbrace{\int_{i<\hat{i}} \tau_i^l \frac{d\hat{l}_i^{EITC}}{d\theta}|_{\theta=0} di}_{\text{Intensive Margin}} \tag{18}$$

where $\tau^l_{i^{LFP}(0)} l_{i^{LFP}(0)}$ 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.[58] The cost 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. [59]

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, $E\left[\left(1-\tau^l_{\hat{i}}\right) l_{\hat{l}}\right]$, with estimates ranging from 0.69-1.16.

I translate this elasticity into equation (18) by normalizing θ to parameterize an additional unit of the mechanical subsidy[60] and writing:

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

where $\epsilon^{LFP}_{E\left[\left(1-\tau^l_{\hat{i}}\right) l_{\hat{i}}\right]}$ is the elasticity of the labor force participation rate with respect to the after tax wage rate and $\frac{E\left[\tau^l_{\hat{i}} l_{\hat{i}}\right]}{E\left[\left(1-\tau^l_{\hat{i}}\right) l_{\hat{i}}\right]}$ is the size of the subsidy as a fraction of after tax income for the marginal

[58] This formula is conceptually similar to that of Eissa et al. (2008) who simulate the MEB of recent EITC expansions using estimates of compensated labor supply elasticities on both the extensive and intensive margin.

[59] Because my 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.

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

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. However, Eissa and Hoynes (2011) report that the average recipient obtains a subsidy equal to 9.2% of gross income in the 2004 SOI; 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. 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 policy elasticity would be $p = 0.09$, ranging between 0.07 and 0.11.

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 $p = 0.09 + 0.05 = 14\%$ with a range of 0.12-0.16 taking the range of extensive margin labor supply responses. However, 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.[61] Therefore, I also consider the case that the 0.05 figure in Chetty et al. (2013) captures all of the EITC response (so that $p = 0.05$). 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 $p = 0.11 + 0.05 = 16\%$.

The estimate of $p = 14\%$ 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 cost of raising public funds from poor single mothers through a reduction in their EITC benefits is $MVPF^{Poor}_{P^{EITC}} = 0.88$. Put differently, the causal estimates from previous literature suggest that the government could lower EITC benefits mechanically by \$0.88 to obtain an extra \$1 because of the reduction in behavioral distortions.

Combining EITC and Tax Policy If one takes the mid-range estimate of $MVPF^{Rich}_{P_{Tax}} = 2$, additional redistribution is desirable iff

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

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 estimate of $MVPF^{Rich}_{P^{Tax}} = 1.33$, 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.

Table 1 shows how the welfare analysis varies with different assumptions about r and p. At the upper range of the top elasticity estimates are $r = 0.8$ (Giertz (2009)), additional redistribution is desirable if and only if one prefers \$0.18 in the hands of the poor relative to the rich. Therefore, the estimates are definitely sensitive to the size of the revenue lost from behavioral responses amongst the rich. However, the estimates are generally less sensitive to the range of estimates for $MVPF^{Poor}_{P^{EITC}}$.

[61] See Meyer and Rosenbaum (2001) for this debate.

Table 1: MVPF_Poor / MVPF_Rich

		r					
		0.00	-0.10	**-0.25**	**-0.50**	-0.80	-1.00
p	MVPF	1.00	1.11	**1.33**	**2.00**	5.00	N/A
0.00	1.00	1.00	0.90	**0.75**	**0.50**	0.20	0.00
0.05	0.95	0.95	0.86	**0.71**	**0.48**	0.19	0.00
0.14	**0.88**	**0.88**	**0.79**	**0.66**	**0.44**	**0.18**	**0.00**
0.16	0.86	0.86	0.78	**0.65**	**0.43**	0.17	0.00

Individual vs. social marginal utility of income Whether one wishes to redistribute at these costs is ultimately a matter of social preference. Okun suggested 60% leakage (Okun (1975)) was tolerable to himself, which implies he values \$0.40 in the hands of the poor equally with \$1 in the hands of the rich.

Many may regard 56% or even 34% leakage to be too high a cost to pay for additional redistribution. But it is also perhaps helpful to compare this to standard measures of within-person preferences towards lotteries over income. In particular, suppose individuals have CRRA utility over consumption with coefficient of relative risk aversion, σ. Then, the social marginal utilities of income can be written as

$$\frac{\eta^{Rich}}{\eta^{Poor}} = \frac{\psi^{Rich} u'\left(c^{Rich}\right)}{\psi^{Poor} u'\left(c^{Poor}\right)} = \frac{\psi^{Rich}}{\psi^{Poor}} \left(\frac{c^{Poor}}{c^{Rich}}\right)^{\sigma}$$

where ψ^{Rich} and ψ^{Poor} are the relative planner weights on marginal utilities.

Assuming the rich consume at least 5 times that of the poor, and assuming $\sigma \geq 1$, one can derive the bound: $\left(\frac{c^{Poor}}{c^{Rich}}\right)^{\sigma} \geq 20\%$. Hence, a utilitarian planner for which $\psi^{Rich} = \psi^{Poor}$, should be willing to lose at least 80% of the mechanical revenue in order to redistribute from rich to poor. Thus, even if redistribution would entail a loss of 75% of the mechanical revenue, a utilitarian planner should prefer additional redistribution for the range of baseline estimates presented here. Put differently, under the additional assumption of CRRA utility with $\sigma \geq 1$, the range of existing policy elasticities suggest the implicit social welfare weights that rationalize the status quo policy are regressive: $\psi^{Rich} > \psi^{Poor}$.

5 Conclusion

This paper provides a general framework for evaluating the marginal welfare impact of government policy changes that values them using individuals' marginal willingness to pay out of their own income. The behavioral response required for such welfare measurement is the causal, not compensated, impact of the policy. Moreover, 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 is sufficient for all behavioral responses.

I hope the framework, and the clarification relative to alternative frameworks, is useful for papers conducting analysis of causal effects of policy changes and seeking a welfare framework to evaluate the

normative aspects of the policy change. Indeed, translating such causal estimates into their implicit MVPF would seem particularly promising with the potential to create a volume of estimates for different policies and a more comprehensive analysis of the desirability of potential government policy changes.

References

Angrist, J. and J.-S. Pischke (2008). *Mostly Harmless Econometrics*. Princeton University Press. 1, 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. 4.2.2

Atkinson, A. B. and N. H. Stern (1974, January). Pigou, taxation and public goods. *The Review of Economic Studies 41*(1), 119–128. 11, 29

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. 6, 1, 36, 2.6, 39, 3

Baily, M. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics 10*, 379–402. 4.1

Baird, S., J. Hicks, M. Kremer, and E. Miguel (2012). Worms at work: Long-run impacts of child health gains. *Working Paper*. 1

Ballard, C. L. and D. Fullerton (1992). Distortionary taxes and the provision of public goods. *Journal of Economic Perspectives 6*(3), 117–131. 11, 3

Barro, R. J. (1974). Are government bonds net wealth? *Journal of Political Economy 82*, 1095–1117. 8

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. 1, 2, 48, A

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. 3

Browning, E. (1987). On the marginal welfare cost of taxation. *American Economic Review 77*, 11–23. 3, 51

Browning, E. and W. R. Johnson (1984). The trade-off between equality and efficiency. *The Journal of Political Economy 92*(2), 175–203. 51

Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics 90*, 1879–1901. 4.1

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy 116*(2), 173–234. 28

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, 31, 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)*. 13, 30, 33, 52, 4.2.2, 56, 4.2.2

Dahlby, B. (2008). *The Marginal Cost of Public Funds*. The MIT Press. 11, 3, 3

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. 6, 1, 2.6, 39, D, D, D

Eissa, N. and H. Hoynes (2011). Redistribution and tax expenditures: The earned income tax credit. *National Tax Journal 64*, 689–730. 4.2.2

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. 10, 20, 35, 58

Eissa, N. and J. Liebman (1996). Labor supply responses to the earned income tax credit. *The Quarterly Journal of Economics 111*(2), 605–637. 56, 4.2.2

Feldstein, M. (1999). Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics 81*(4), 674–680. 1, 4, 1, 2.5, 4.2.1

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

Fullerton, D. (1991). Reconciling recent estimates of the marginal welfare cost of taxation. *American Economic Review 81*, 302–08. 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, 9, 4.2.1, 4.2.2

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

Gruber, J. and E. Saez (2002, April). The elasticity of taxable income: evidence and implications. *Journal of Public Economics 84*(1), 1–32. 54

Harberger, A. (1964). The measurement of waste. *American Economic Review 54*, 58–76. 1, 4, 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. 4.2.2, 4.2.2

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. 45

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, 32, 57

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995, June). *Microeconomic Theory*. Oxford University Press, USA. 1, 28

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. 61

Mirrlees, J. A. (1971). An exploration into the theory of optimal income taxation. *The Review of Economic Studies 38*(2), 175–208. 1, 4.1

Okun, A. (1975). *Equality and Efficiency*. Brookings Institution Press. 1, 2.3, 4, 51, 4.1, 4.2.2

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. 3

Piketty, T. and E. Saez (2012). Optimal labor income taxation. *NBER Working Paper No. 18521*. 12, 4.1

Ramsey, F. P. (1927, March). A contribution to the theory of taxation. *The Economic Journal 37*(145), 47–61. 1, 2.6, D

Saez, E. (2001). Using elasticities to derive optimal income tax rates. *Review of Economic Studies 68*, 205–229. 1, 4.1

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, 9, 1, 2.5, 4.2.1

Salanie, B. (2003). *The Economics of Taxation*. MIT Press. 65

Samuelson, P. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics 36*, 387–389. 1, 2.4, 3, A

Schlee, E. E. (2013). Radner's cost-benefit analysis in the small: An equivalence result. *Working Paper*. 25

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. 1, 8, 3, 46, 3

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*. 1, 8, 3

Stiglitz, J. E. and P. Dasgupta (1971, April). Differential taxation, public goods, and economic efficiency. *The Review of Economic Studies 38*(2), 151–174. 11

Online Appendix: Not for Publication

A Illustration Using Job Training Partnership Act

This Appendix illustrates how the model can be used to think about cost benefit analysis of the Job Training Partnership Act (JTPA). This act was studied using a randomized controlled trial, the results of which were presented in Bloom et al. (1997). Fortunately, they look not only at the impact of the program on earnings but also on budget-relevant variables such as welfare and tax receipt.

To simplify the analysis, assume the social marginal utilities of income are constant amongst the beneficiaries of the JTPA and consider the marginal welfare impact on these beneficiaries from the program. For simplicity, I follow the results for adult women at the top of Table 8 on p573 of Bloom et al. (1997).

Behavioral Response Impact on Budget: $471

They estimate that the program increased earnings of adult women by $1,683, which led to an increased tax collection of $236 per enrollee and also a $235 reduction in welfare benefit expenditures (AFDC). Summing, the total impact of the behavioral response to the program on the government budget is $471.

Net Transfer: $910

They estimate that the marginal cost of providing the program to an adult female enrollee is $\$1,227$ plus a marginal wage subsidy of $154 per enrollee. This leads to a marginal cost of $c_G = \$1,381$. Subtracting this from the $471 resource transfer provided to the government yields a net resource transfer to the enrollee of $910.[62]

Market Failure: -$1,381 to $302 depending on assumptions

Finally one needs to calculate the extent to which an individual would be willing to pay for the job training program in excess of its marginal cost, $\frac{u_G}{\lambda} - c_G$. This is always the most difficult component because it requires pricing a non-market good. One pathway forward is to value the program by assuming the program increased earnings without any change in effort (i.e. the impact on earnings is a pure productivity increase, and hence equivalent to a transfer). Indeed, this is implicitly the view taken by Bloom et al. (1997) who estimate an impact on earnings of $1,683 and include this as a benefit of the program, $\frac{u_G}{\lambda} = \$1,683$. With this assumption, the size of the market failure is $1,683-$1,381 = $302 per enrollee and the MVPF is $1,683/$910 = 1.85.

However, the envelope theorem suggests caution in interpreting the full impact on earnings as the welfare increase. Although the fact that individuals voluntary enrolled in the program suggests $\frac{u_G}{\lambda} > 0$, it could be the case that job training program increased the benefits of working slightly above its costs. In this case $\frac{u_G}{\lambda}$ could be arbitrarily close to zero. In this case, the MVPF would be close to zero.

[62] 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, this term is irrelevant for welfare calculations in the present framework.

One benchmark assumption could be to assume that the government has no comparative advantage or disadvantage in providing the job training program, in which case the programs value would equal its cost of \$1,381.[63] This would imply a MVPF of \$1,381/\$910=1.52, which is above 1 because of the fiscal externalities associated with the program.

In short, the application of the model to the provision of a publicly provided good or service requires being able to value this good or service in the manner proposed by Samuelson (1954). But, the causal effects are sufficient. If instead one were to try to conduct an estimate of the MEB of the job training program, one would still need the valuation of the publicly provided serve, but one would also need to decompose the \$471 behavioral response into income and substitution effects – only the compensated response would be desired.

B Appendix: Proofs

B.1 Proof of Proposition 1

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

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

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

$$\begin{aligned}
\frac{\partial V_i}{\partial \tau^x_{ij}} &= -x_{ij}\lambda_i \\
\frac{\partial V_i}{\partial \tau^l_{ij}} &= -l_{ij}\lambda_i \\
\frac{\partial V_i}{\partial T_i} &= -\lambda_i \\
\frac{\partial V_i}{\partial G_i} &= \frac{\partial u_i}{\partial G_i}
\end{aligned}$$

Replacing terms, I have

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

[63] Of course, the natural concern with this is that the value on those who select the job training program from the government could be lower than the value experienced by those who choose to receive training in the private market.

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

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

Now, I show that $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$ is equal to the marginal equivalent variation and marginal compensating variation. Recall that $EV_i(\theta)$ solves

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

Thus, differentiating with respect to θ and evaluating at $\theta = 0$ yields

$$\frac{\partial V_i}{\partial y_i} \frac{d[EV_i]}{d\theta}|_{\theta=0} = \frac{d\hat{V}_i}{d\theta}|_{\theta=0}$$

or

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

Similarly, recall $CV_i(\theta)$ solves

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

Differentiating with respect to θ and evaluating at $\theta = 0$ yields

$$\frac{d\hat{V}_i}{d\theta}|_{\theta=0} - \frac{d[CV_i]}{d\theta}|_{\theta=0} \frac{\partial V_i}{\partial y_i} = 0$$

or

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

so that $\frac{\frac{dV_i}{d\theta}}{\lambda_i}$ is equal to the marginal equivalent variation and marginal compensating variation of the program.

B.2 Proof of Okun's Leaky Bucket Formula

The marginal welfare impact of the policy on a rich individual, i, is given by

$$\begin{aligned}\frac{\frac{dV_i^{Rich}}{d\theta}}{\lambda_i^{Rich}} &= \frac{d\hat{t}_i^{Rich}}{d\theta} - \left(\sum_j^{J_X} \tau_{ij}^x \frac{d\hat{x}_{ij}}{d\theta} + \sum_j^{J_L} \tau_{ij}^l \frac{d\hat{l}_{ij}}{d\theta}\right) \\ &= -\frac{d\tau^{Rich}}{d\theta}\left(\hat{l}_i - \bar{l}\right)\end{aligned}$$

where the second line follows from the fact that the net revenue taken from the rich excludes what was lost from their behavioral response. Hence the welfare impact on the rich is purely the mechanical impact of the tax change on their income.

To simplify the analysis, I assume that the marginal social utility of income is equated among the rich and given by η^{Rich}. Hence, the aggregate welfare impact of raising this revenue from the rich is given by

$$\eta^{Rich} \int_{i \in Rich} \frac{\frac{dV_i^{Rich}}{d\theta}}{\lambda_i^{Rich}} di = -\eta^{Rich} \int_{i \in Rich} \frac{d\tau^{Rich}}{d\theta}\left(\hat{l}_i - \bar{l}\right) di$$

For the poor single mothers, the marginal welfare increase is given by

$$\frac{\frac{dV_i^{Poor}}{d\theta}}{\lambda_i^{Poor}} = \frac{d\hat{t}_i^{Poor}}{d\theta} - \left(\sum_j^{J_X} \tau_{ij}^x \frac{d\hat{x}_{ij}}{d\theta} + \sum_j^{J_L} \tau_{ij}^l \frac{d\hat{l}_{ij}}{d\theta}\right)$$

Again, to simplify the analysis, I assume that the social marginal utility of income is constant amongst EITC recipients. With this assumption, the aggregate welfare impact on the poor is given by

$$\eta^{Poor} \int_{i \in Poor} \left(\frac{\frac{dV_i^{Poor}}{d\theta}}{\lambda_i^{Poor}}\right) di \quad = \eta^{Poor} \int_{i \in Poor} \left[\frac{d\hat{t}_i^{Poor}}{d\theta} - \left(\sum_j^{J_X} \tau_{ij}^x \frac{d\hat{x}_{ij}}{d\theta} + \sum_j^{J_L} \tau_{ij}^l \frac{d\hat{l}_{ij}}{d\theta}\right)\right] di$$

Now, budget neutrality implies that the net transfer to the poor is given by the mechanical revenue raised minus the behavioral responses from all non-poor types:

$$\int_{i \in Poor} \frac{d\hat{t}_i^{Poor}}{d\theta} di = \int_{i \in Rich} \frac{d\tau^{Rich}}{d\theta}\left(\hat{l}_i - \bar{l}\right) di - \int_{i \notin Poor} \left(\sum_j \tau_{ij}^x \frac{dx_{ij}}{d\theta} + \tau_{ij}^l \frac{dl_{ij}}{d\theta}\right) di$$

so that the welfare impact on the poor is given by

$$\eta^{Poor} \int_{i \in Poor} \left(\frac{\frac{dV_i^{Poor}}{d\theta}}{\lambda_i^{Poor}}\right) di = \eta^{Poor} \left[\int_{i \in Rich} \frac{d\tau^{Rich}}{d\theta}\left(\hat{l}_i - \bar{l}\right) di - \int_i \left(\sum_j \tau_{ij}^x \frac{dx_{ij}}{d\theta} + \tau_{ij}^l \frac{dl_{ij}}{d\theta}\right) di\right]$$

Combining, the impact of the policy on social welfare is given by

$$\frac{dW}{d\theta} = \eta^{Poor}\left[\int_{i\in Rich}\frac{d\tau^{Rich}}{d\theta}\left(\hat{l}_i - \bar{l}\right)di - \int_i\left(\sum_j \tau_{ij}^x\frac{dx_{ij}}{d\theta} + \tau_{ij}^l\frac{dl_{ij}}{d\theta}\right)di\right] - \eta^{Rich}\int_{i\in Rich}\frac{d\tau^{Rich}}{d\theta}\left(\hat{l}_i - \bar{l}\right)di$$

So that the policy increases social welfare if and only if

$$1 - \frac{\eta^{Rich}}{\eta^{Poor}} = \underbrace{-\frac{\int_i\left(\sum_j \tau_{ij}^x\frac{dx_{ij}}{d\theta} + \tau_{ij}^l\frac{dl_{ij}}{d\theta}\right)di}{\int_{i\in Rich}\frac{d\tau^{Rich}}{d\theta}\left(\hat{l}_i - \bar{l}\right)di}}_{\text{Leaks in Okun's Bucket}}$$

C Appendix: Externalities (and Internalities)

The analysis assumes individuals maximize their welfare without imposing any externalities on others or internalities on themselves. While researchers may debate the extent of externalities or internalities, my result that the causal response to the policy is required for policy analysis readily extends to a world 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 vector of all 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

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

Given each agent's solution to this program, $\mathbf{x_i}$, I construct $E_i = f_i^E(\mathbf{x})$ 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.[64] Proposition 2 presents the characterization of the marginal welfare impact of a policy evaluated at $\theta = 0$.

Proposition 3. *The welfare impact of the marginal policy change to type i is given by*

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

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

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

[64] 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).

Replacing terms, I have

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

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

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

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

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

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.[65] This is the so-called "double-dividend" highlighted in previous literature (Bovenberg and de Mooji (1994); Goulder (1995); Parry (1995)). My results show that 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.

[65] 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.

D 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

$$\sum_k \hat{\tau}_k \frac{d\hat{x}_k}{d\theta} = 0 \tag{19}$$

Equation (19) 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

$$\sum_k \tau_k \frac{\partial x_k^h}{\partial \tau_1}\frac{d\tau_1}{d\theta} = \sum_k \tau_k \frac{\partial x_k^h}{\partial \tau_2}\left(-\frac{d\tau_2}{d\theta}\right) \tag{20}$$

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).[66]

However, note that one never relied on compensated elasticities to test the optimality condition in equation (19). Compensated elasticities arise only because of the assumption that policy is at the

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

optimum. One could consider any budget-neutral policy that simultaneously adjusts two commodity taxes and test equation (19) 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.