The Policy Elasticity

Nathaniel Hendren*

June, 2013

Abstract

This paper provides a generic framework for evaluating the welfare impact of government policy changes towards taxes, transfers, and publicly provided goods. The results show that the behavioral response required for welfare measurement is the *causal impact* of each agent's response to the policy on the government's budget. A decomposition of this response into income and substitution effects is not required. Because these desired elasticities vary with the policy in question, I term them *policy elasticities*. I also provide an additivity condition that yields a natural definition of the marginal costs of public funds as welfare impact of a policy per dollar of its cost to the government budget. Finally, I use the model, along with causal estimates from existing literature, to study the welfare impact of additional redistribution by increasing the generosity of the earned income tax credit financed by an increase in the top marginal income tax rate. I show existing causal estimates suggest additional redistribution is desirable if and only if providing an additional \$0.44 to an EITC-eligible single mother (earning less than \$40,000) is preferred to providing an additional \$1 to a person subject to the top marginal tax rate (earning more than \$400,000).

1 Introduction

There is a large and growing set of empirical work estimating the behavioral impacts of changes to government policies. This diverse literature has developed a large toolkit of structural and reduced form methods to answer the positive question of what do policy changes do to behavior. However, what is less clear is what does one need to know in order to move from a positive analysis to a normative analysis of whether or not policy changes improve social welfare.[1]

Existing literature is filled with debates about the parameters required for a normative analysis of government policy changes.[2] Often, a Hicksian (compensated) elasticity is argued to play a central role:

*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, 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., Angrist and Pischke (2010) and Nevo and Whinston (2010) for a recent discussion of reduced form and structural methods.

[2] Perhaps the most vigorous debate pertains to whether the marginal cost of public funds should rely on compensated or uncompensated elasticities (see, e.g., Allgood and Snow (1998); Atkinson and Stern (1974); Ballard et al. (1985);

> *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))
>
> *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))

As a result, it is commonplace for empirical researchers to believe that behavioral responses must be decomposed into their income and substitution (compensated) effects. Goolsbee (1999, p8) explains the resulting empirical challenge: "The theory largely relates to compensated elasticities, whereas the natural experiments provide information primarily on the uncompensated effects". Rarely do policy changes hold people's utility constant. Thus, it would appear that the causal effects of policy changes (i.e. the positive analysis of the policy) may not be exactly what is desired for a normative analysis of that same policy.

This paper revisits the debate about the types of parameters required for welfare analysis. Using a generic heterogeneous agent model with publicly provided goods, taxes, and transfers, I characterize the parameters required for welfare measurement of marginal changes to government policies. The main result is that the only behavioral response required is the impact of the behavioral response to the policy on the government's budget. This impact is the difference between the government budget with behavioral responses to the policy and a counterfactual world without any behavioral responses to the policy.

This means that neither Hicksian (compensated) nor Marshallian (uncompensated) price elasticities are generally sufficient. Rather, the causal effect of the policy in question is exactly what is needed. If a policy raises taxes and spends money on roads, one desires the impact of the simultaneous increase in taxes and increase in road spending on taxable behavior. To provide a vocabulary for the types of elasticities required for welfare analysis, and to highlight the importance of being clear about the policy experiment being studied, I term them "policy elasticities". These are simply the difference in behavior if the policy is or is not undertaken – precisely the standard definition of the causal effect of the policy (e.g. Angrist and Pischke (2008)).

The causal effect 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. In the broad class of models considered here, the government is the only distortion between private prices

Ballard (1990); Browning (1976, 1987); Fullerton (1991); Harberger (1964); Mayshar (1990); Mayshar and Slemrod (1995); Slemrod and Yitzhaki (2001, 1996); Stiglitz and Dasgupta (1971); Stuart (1984); Wildasin (1984), and summaries provided in Auerbach (1985) and, more recently, Dahlby (2008)).

and social (resource) costs; hence the impact of the behavioral response on the government's budget is the only behavioral response required for welfare estimation.[3]

In addition to the impact of the behavioral response to the policy on the government's budget, there are two other types of parameters required for welfare analysis, both of which are arguably well-known. First, if a policy changes the provision of publicly provided goods, one also needs to know the valuation of these goods. This is given by the difference between individuals' marginal rates of substitution and their marginal cost of production – an insight of Samuelson (1954). Second, to the extent to which the policy change has differential welfare impacts across people, one needs to know their social marginal utilities in order to aggregate from individual to social welfare. These three types of parameters – the impact of the behavioral response on the government budget, the net valuation of changes in publicly provided goods, and the social marginal utilities of income – fully characterize the welfare impact of marginal changes to government policies in the broad class of models considered here.

A conceptual difficulty with welfare analysis of actual government policies is that sometimes they are not budget neutral.[4] My framework naturally accommodates welfare analysis in these settings. In particular, I provide an additivity condition that yields conditions under which one can combine non-budget neutral policies to form a comprehensive budget neutral welfare analysis that recognizes the fact that the government must raise revenue to pay for its expenditure policies. The analysis yields a natural definition of the marginal cost of public funds (MCPF) as the welfare impact of a non-budget neutral policy per dollar of its budget cost. With this definition, one can compare the desirability of different policies by comparing their MCPF. If two policies have different MCPFs, then taking $1 from the low MCPF policy and using it to spend resources on the high MCPF policy will improve social welfare.[5] In contrast to traditional measures of the MCPF, this definition is policy-specific and depends not on compensated or uncompensated elasticities per se, but rather on the policy elasticities of the particular non-budget neutral policy in question.

I apply the framework 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). The model suggests additional redistribution is desirable if and only if the difference in social marginal utilities of income between rich and poor is greater than the impact of the behavioral response to the redistributive policy on the budget, normalized by the mechanical revenue raised from the rich. This is precisely the logic of Okun's leaky bucket experiment (Okun (1975)). Okun shows

[3] One can think of this behavioral response as a "fiscal externality". In Appendix B, I consider an extension of the model to a world with other externalities. If there are other distortions operating in the market that prevent social and private costs from being equated (e.g. externalities such as pollution, or even "internalities" whereby people's actions impose a welfare loss on themselves), then one also requires an estimate of the causal impact of the behavioral response to the policy on the value of the other externalities as well. Even in this more general model with other externalities, the causal effects (i.e. policy elasticities) are still the elasticities desired for welfare analysis–a decomposition of behavioral responses into income and substitution effects is not required.

[4] Non-budget neutral policies are arguably the norm rather than the exception (e.g. the 2003 creation of Medicare Part D, the Omnibus Budget Reduction Act of 1993, etc.)

[5] In this sense, my definition of the MCPF is symmetric and could equally be called a marginal "benefit" instead of "cost"; it represents the social cost of taking resources from the policy and the social benefit of spending resources on the policy. These two notions are equivalent under my definition.

one's social preference for redistribution can be stated in terms of how much money one is willing to lose in the process of redistributing money from rich to poor. My framework shows that the leaks in Okun's bucket are functions of the policy elasticities of the redistributive policy.

I implement this formula using existing causal estimates of EITC expansions and changes to the top marginal income tax rate. To do so, I use the additivity condition to write the redistributive policy as the sum of two policies: an increase in the top tax rate and an expansion of the EITC program. Each of these policies induce a MCPF: the welfare impact on the rich of raising \$1 by increasing the top marginal income tax rate, and the welfare impact on the poor of spending this \$1 through an increase in EITC benefits. Mid-range estimates of causal effects in each of these literatures[6] suggests a MCPF on the rich of increasing the top tax rate is roughly \$2. The MCPF on the poor of an EITC expansion is roughly \$0.88. Hence, additional redistribution is desirable as long as one prefers \$0.44 in the hands of an EITC recipient to \$1 in the hands of someone taxed at the top marginal income tax rate.[7] 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 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 for measuring the marginal welfare impact of policies that actually change utilities.

In Appendix C, I illustrate in detail how my framework nests the two classes of models for which Hicksian elasticities arise in previous literature. First, Hicksian elasticities arise in the calculation of optimal policies with representative agents, such as the optimal commodity taxation program of Ramsey (1927) and Diamond and Mirrlees (1971). At an optimum, the marginal welfare impact 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. Second, Hicksian elasticities calculate the marginal amount of additional revenue the government could collect by switching from distortionary taxation to lump-sum taxation, holding the agent's utility constant. This is the so-called marginal deadweight loss or marginal excess burden of the tax system (Mas-Colell et al. (1995)). However, because the policy does not specify what the government does with this additional revenue, the policy experiment is left incomplete: any policy aimed at reducing the marginal excess burden would have a welfare impact that depended on the causal impact of that particular policy, would not be governed by a Hicksian elasticity. In short, although Hicksian elasticities arise in these classes of models and provide measures of distortions, they are neither necessary nor sufficient elasticities for

[6] See Hotz and Scholz (2003) and Chetty et al. (2013) for a review of the EITC responses and Saez et al. (2012) for a review of responses to top marginal income tax changes.

[7] Taking broader ranges of estimates from existing literature yields estimates ranging from \$0.25-\$0.76 in the hands of the EITC recipient relative to \$1 for the rich.

welfare evaluation of actual government policy changes–instead, one needs the policy elasticities.

This paper is also related to the large literature studying the optimal design of the nonlinear income tax schedule initiated by Mirrlees (1971) and implemented by Saez (2001) among others. In contrast to focusing on "optimal" policy, my approach quantifies the welfare impact of changing policy relative to the status quo. This method may have more practical appeal for those interested in evaluating proposals to changes of existing policy.[8] Moreover, my empirical approach has less stringent requirements. Optimal tax formulas generally depend on elasticities defined locally *around the optimum*. In contrast, my approach depends on elasticities defined locally *around the status quo*. As a result, they are arguably easier to identify using existing data variation and fewer extrapolation assumptions.[9]

This paper also builds on the literature discussing the sufficiency of the taxable income elasticity for welfare analysis of tax policies (Feldstein (1999); Chetty (2009a)). My results show that the impact of the behavioral response on the government's budget remains sufficient even in cases with more than one tax rate (e.g. different capital and labor income taxes), an important case in which the aggregate taxable income elasticity is no longer sufficient. Moreover, even in a world with a single tax rate my results clarify that the desired elasticity is neither a compensated or uncompensated elasticity; it is the policy elasticity, which depends on the policy in question.

This paper is also related to the long debate over the correct definition of the MCPF.[10] This literature generally argues about whether or not the welfare cost of raising government revenue depends on compensated or uncompensated price elasticities. My approach suggests neither is exactly right. Instead, one wants the causal effects of non-budget neutral policies. As long as one can identify the causal effect of the policy, one need not be concerned with whether people respond to such policies in a compensated or uncompensated manner, which may depend on the actual (or expected) benefits they receive from the policy.[11]

My empirical application also highlights the importance of being clear about distributional in-

[8] Indeed, this methodology for welfare estimation was recommended by Feldstein (1976) because of its close connection to actual policy considerations and debates.

[9] In addition, the standard approach to optimal taxation using the Hamiltonian provides 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) which depends on the integral of a function of elasticities (Piketty and Saez (2012)). As a result, the Hamiltonian approach is well-suited to study the optimal *shape* of EITC benefits, but is more difficult to study the optimal *size* of the EITC program relative to taxes on the rich.

[10] See citations in footnote 2. Relative to this literature, my paper makes arguments similar to those in Stuart (1984), which suggests one "wishes to compare changes in utility and revenue as the economy moves from an equilibrium before a tax increase to one after the increase". It then argues that budget neutrality requires the MCPF incorporate the behavioral responses induced by government spending, along the lines of Atkinson and Stern (1974) (discussed below in Footnote 28). In contrast, I propose a definition of the MCPF for explicitly non-budget neutral policies. In this sense, my approach is also closely related to Slemrod and Yitzhaki (1996, 2001), who define the MCPF as the welfare impact of a tax increase per unit revenue raised. In contrast to this literature, I define a policy-specific MCPF and illustrate that in general it is neither Hicksian nor Marshallian elasticities, but rather the policy elasticities that are desired for its estimation.

[11] Indeed, it is commonplace to associate the causal effect with the uncompensated elasticity, as in the quote by Goolsbee (1999) referenced above. But if people who experience a tax increase receive (or expect to receive) future benefits, they may borrow against these benefits and respond in a compensated manner (the standard Ricardian equivalence logic). Fortunately, my definition of the MCPF, which depends only on the causal effects of the policies, does not require knowledge of the degree to which Ricardian equivalence holds in the economy.

cidence when calculating a MCPF. I derive two values of the MCPF for two policies aimed at two different populations (those receiving EITC benefits and those facing the top marginal income tax rate).[12] Although the precise estimates are sensitive to the range of causal effects estimated in previous literature, the results clearly illustrate that there is no single MCPF: raising resources from the poor is much cheaper ($0.88) than raising resources from the rich ($2). This does not imply one should increase taxes on the rich; rather, it suggests society would prefer to take money from the rich more than the poor.[13]

Finally, this paper is also related to the broader literature discussing the relationship between structural and reduced form methods in economics (Angrist and Pischke (2010); Nevo and Whinston (2010)) and the rise of sufficient statistic methods (Chetty (2009b)). By characterizing the set of parameters required for welfare analysis in a broad class of models, I hope the conclusion is useful: estimates of causal effects of past policy changes, or forecasts of the impact of future policy changes, can be put into a general normative framework by multiplying by the government's tax/subsidy rate on the affected behavior.

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. The Online Appendix provides proofs of the main results (Appendix A), extensions of the model to incorporate other externalities (Appendix B), and a detailed discussion of the relationship to previous literature (Appendix C).

2 Model

The model has a generic setup 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 12 illustrates the main concepts in a model with a representative agent, single taxable good, and single publicly provided good.

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

[12] Existing literature studying the causal effects of changes to the top marginal income tax rate has argued that the results imply a "marginal excess burden" of taxing the rich is approximately $0.50 per $1 (see Saez et al. (2012) and the discussion in Section 4.2.1). Because this quantity is calculated using the behavioral response to actual policies that changed the top marginal income tax rate (e.g. OBRA 1993), this turns out to be precisely the number required for computing my definition of the MCPF, (1/(1-0.5))=2. However, in contrast to the existing use of terminology, this number is not a measure of the "marginal excess burden", since the calculation of marginal excess burden requires a removal of income effects (see Mas-Colell et al. (1995); Feldstein (2012) and the discussion in Section 2.5).

[13] This provides empirical support for the theoretical analyses of Kaplow (1996, 2004, 2008) and Slemrod and Yitzhaki (2001) that suggest the welfare cost of raising revenue will be inversely related to the social marginal utility of income of those being taxed.

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 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.[15]

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.[16] This simplified production structure rules out many interesting features that can easily be added to a more general model, including imperfect competition (i.e. producer surplus), production externalities (e.g. spillovers), and pecuniary externalities (in which case real prices would not always be 1).[17] I assume the marginal cost to the government of producing publicly-provided goods, G_{ij} is given by c_j^G for $j = 1, .., J_G$.[18]

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 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

[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] 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.

[16] I allow (but do not require) taxes and transfers to be individual-specific. In practice, most policies will involve taxes and transfers will not be individual-specific, potentially due to information constraints facing the government. An advantage of allowing for individual-specific taxes in my setting is that one can consider nonlinear tax settings. In particular, one can interpret T_i as "virtual income" and τ_{ij}^l as the marginal tax on labor earnings. In this case, one must be sure that the marginal tax rate used is consistent with the segment of the budget constraint that would be chosen by the agent.

[17] I discuss some of these extensions in Subsection 2.6 and provide a detailed discussion of externalities in Appendix B.

[18] 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.

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.

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

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

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:

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

Second, I assume that the policy path is continuously differentiable in θ (i.e. $\frac{d\hat{\tau}_{ij}^x}{d\theta}$, $\frac{d\hat{\tau}_{ij}^l}{d\theta}$, $\frac{d\hat{T}_i}{d\theta}$, and $\frac{d\hat{G}_{ij}}{d\theta}$ exist and are continuous in θ).[19] Intuitively, $P(\theta)$ traces out a smooth path of government policies, centered around the status quo. Given this path, 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.[20]

Before asking the normative question of whether the government should follow the policy path, I

[19] 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 an important factor in welfare estimation (Eissa et al. (2008))).

[20] 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.

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 $P(\theta)$.[21] 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,

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

where $\sum_{j=1}^{J_G} c_j^G \hat{G}_{ij}(\theta)$ is the government expenditure on publicly provided goods to individual i, $\hat{T}_i(\theta)$ is the net 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[22], 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.

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

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

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

[22] 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 34.

activities:

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

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 Measuring Welfare

Moving from positive to normative analysis requires a definition of the normative objective. My measure of individual welfare will be the individual's willingness to pay out of their own income to follow the policy path. 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 $\hat{V}_i\left(\theta\right)$ denote the utility obtained by type i under the policy $P\left(\theta\right)$. The marginal impact of the policy on the utility of individual i is given by $\frac{d\hat{V}_i}{d\theta}|_{\theta=0}$. Normalizing by the marginal utility of income, the individual's own willingness to pay (out of their own income) for a marginal policy change is given by $\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}$. This will be my measure of individual welfare.[23] The social welfare impact of the policy will be given by the sum of the individuals' willingnesses to pay weighted by their social marginal utilities of income, $\frac{d\hat{W}}{d\theta}|_{\theta=0}=\int_{i\in I}\eta_i\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}di$.[24] This social welfare is measured in units of utility; hence, it will also be helpful to measure social welfare in units of individual $\hat{i}$'s income, $\frac{d\hat{W}^{\hat{i}}}{d\theta}|_{\theta=0}=\frac{\frac{d\hat{W}}{d\theta}|_{\theta=0}}{\eta_{\hat{i}}}=\int_{i\in I}\frac{\eta_i}{\eta_{\hat{i}}}\frac{\frac{d\hat{V}_i}{d\theta}|_{\theta=0}}{\lambda_i}di$, where the superscript $\hat{i}$ denotes the fact that social welfare is measured in units of $\hat{i}$'s income.

[23] 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, $EV_i\left(\theta\right)$, of policy $P\left(\theta\right)$ for type i is the amount that the consumer would be indifferent to accepting in lieu of the policy change. $EV_i\left(\theta\right)$ solves

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

Second, the compensating variation, $CV_i\left(\theta\right)$, of policy $P\left(\theta\right)$ 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. $CV_i\left(\theta\right)$ solves

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

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

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

[24] 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 $W=\int_{i\in I}G\left(V_i\right)di$ for a concave function G, then the social marginal utility of income would be $\eta_i=G'\left(\hat{V}_i\left(0\right)\right)\lambda_i$).

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:*

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

Proof. The proof is an application of the envelope theorem and is provided in Appendix A.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.[25] Conditional on calculating this

[25] As discussed in Appendix B, 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.

fiscal externality, behavioral responses are not required for welfare analysis.[26]

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, $T = 0$.*

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

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

where the LHS is the net willingness-to-pay for additional roads[27], $\hat{\tau}^l$ 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.[28]

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-

[26] 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)).

[27] Note that optimization implies $\lambda = \frac{\partial u}{\partial x}$.

[28] *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.

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 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 $\hat{\epsilon}_{ij}^{x} = \frac{dlog(\hat{x}_{ij})}{d\theta}|_{\theta=0}$ (or $\hat{\epsilon}_{ij}^{l} = \frac{dlog\left(\hat{l}_{ij}\right)}{d\theta}|_{\theta=0}$)

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

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

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

The 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.[29] 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).

2.4 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

[29] 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%.

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 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.[30] 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, 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. 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). In contrast, the impact of the behavioral response to the policy on the government's budget (i.e. the fiscal externality) remains sufficient even in the cases when individuals face multiple tax rates on different behaviors.

2.5 Relation to Hicksian elasticity

As discussed in the introduction, previous literature has highlighted the role of Hicksian (compensated) elasticities in the welfare evaluation of government policy changes. However, Hicksian elasticities measure the causal effects of policy changes that hold utility constant; hence Proposition 1 shows that they are not sufficient for evaluating the welfare impact of policies that actually change utilities.

Of course, my results do not contradict any mathematical results from previous literature on the role of the Hicksian elasticity in certain circumstances. Indeed, there are two prominent and distinct classes of models where a Hicksian (compensated) elasticity arises. I discuss both of these cases in detail in Appendix C, but provide a short overview here.

The first class of models in which Hicksian (compensated) elasticities arise is in the derivation of

[30] 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.

optimal taxes with a single (representative) agent. This problem was proposed by Ramsey (1927) and analyzed in detail by Diamond and Mirrlees (1971). It yields the classic "inverse elasticity" rule for commodity taxation: at the optimum, tax-weighted Hicksian price derivatives for each commodity are equated.

This formula involves compensated responses because, with a single agent, a necessary condition for taxes to be at an optimum is that small budget-neutral changes to taxes do not affect utility. Hence, around the optimum, the causal effects are Hicksian responses (because utility is not changing at the optimum). Thus, the optimal commodity taxes in this setup depend on Hicksian elasticities defined locally around the optimum.[31] But evaluating the welfare impact of policies that change commodity taxes require the causal effect of that policy. If policy is at an optimum and there is a single agent, then the policy elasticity for budget-neutral policies is a Hicksian elasticity and the marginal welfare impact of a policy change is zero.

The second class of models in which Hicksian (compensated) elasticities arise is in the calculation of the marginal deadweight loss or "excess burden" from taxation (see, e.g., Feldstein (1995, 1999); Chetty (2009a)). Such a calculation involves a hypothetical comparison between distortionary taxation of goods and labor supply activities and individual-specific lump-sum taxation. In particular, the compensating variation (CV) measure of deadweight loss asks how much additional revenue the government could obtain if it switched from distortionary taxation to individual-specific lump-sum taxation, while holding utility constant (Mas-Colell et al. (1995)). Because the hypothetical policy involves holding utilities constant, the causal effect of the policy experiment defining marginal deadweight loss is governed by a Hicksian price response.[32]

While the marginal excess burden of taxation provides a potential measure of the welfare cost of distortionary taxation, the hypothetical policy experiment is left incomplete: it calculates the additional revenue obtained by the government but does not specify what the government does with this revenue. Any spending would affect utilities and behavior. Indeed, the welfare impact of policies aimed at reducing the distortionary burden of taxation depend on the causal effects of such policies, which will not in general be Hicksian price responses. Of course, one may wish to know the welfare cost of raising government revenue, but as I show in Section 3, the Hicksian price response will not in general be required.

2.6 Extensions

Although the model allowed for considerable heterogeneity across individuals, it assumed a stylized model of production with perfect competition and fixed resource prices.[33] This rules out many phe-

[31] Diamond and Mirrlees (1971) also consider a model with heterogeneous agents and derive their tax rules in such settings. With heterogeneous agents, the equation of compensated revenues is no longer desirable so that 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.

[32] Appendix C also discusses the equivalent variation (EV) measure of deadweight loss. This measure also depends on a form of compensated elasticities, but they are not technically Hicksian elasticities.

[33] 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

nomena 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 sufficient, 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 B 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 estimate of the individual's marginal rate of substitution between pollution and income, analogous to the requirement for the provision of un-priced 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.

Finally, the model presented here is not explicitly dynamic. To be sure, one can think of j indexing time, but in these cases one must then consider policy paths which specify not only current but also future policies. In practice, it is easier to consider responses to current policy changes without needing to account for any potential future policy changes. A natural path forward is to think of the model as static but then consider policies that are not budget neutral in the short run. One can then ask whether pursuing such non-budget neutral policies are worth their costs imposed on the government that (at least eventually) must satisfy its budget constraint. It is this type of an approach that arguably motivates the large literature on the marginal cost of public funds, to which I now turn.

3 Additivity and the Marginal Costs of Public Funds

The previous section shows that in general it is sufficient to consider the causal impact of a policy on behavior. However, many government policies are not budget neutral. So, it is often desirable to adjust the welfare analysis of non-budget neutral policies for the welfare cost of policies needed to raise their required revenue. Such an adjustment is the the motivation for the large literature on the marginal cost of public funds (MCPF).

This section provides a condition that allows welfare impacts of policies to be added together. This condition leads to a natural definition of the marginal costs of public funds of non-budget neutral

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 B).

policies as the welfare impact of policies per dollar of government spending.

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, $P(\theta)$:

$$(P(\theta) - P(0)) = (P_{Tax}(\theta) - P(0)) + (P_{Exp}(\theta) - P(0)) \tag{9}$$

Condition (9) 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.[34]

Proposition 2. *Suppose $P(\theta)$, $P_{Tax}(\theta)$, and $P_{Exp}(\theta)$ satisfy equation (9). Then, the marginal welfare impact of the comprehensive policy on type i, denoted $\frac{\frac{\partial \hat{V}_i^P}{\partial \theta}|_{\theta=0}}{\lambda_i}$, is given by*

$$\frac{\frac{\partial \hat{V}_i^P}{\partial \theta}|_{\theta=0}}{\lambda_i} = \frac{\frac{\partial \hat{V}_i^{P_{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda_i} + \frac{\frac{\partial \hat{V}_i^{P_{Exp}}}{\partial \theta}|_{\theta=0}}{\lambda_i} \tag{10}$$

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}.

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 $\theta = 0$. □

[34] 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.

Given a welfare estimate of an expenditure policy, $\frac{\frac{\partial V_i^{P_{Exp}}}{\partial \theta}|_{\theta=0}}{\lambda_i}$, Proposition (2) shows how one can add a welfare estimate of a tax policy, $\frac{\frac{\partial V_i^{P_{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda_i}$, in order to analyze the welfare impact of the budget-neutral policy, $\frac{\frac{\partial V_i^{P}}{\partial \theta}|_{\theta=0}}{\lambda_i}$. The key requirement in equation (9) is straightforward: the expenditure policy and the tax policy must sum to the total policy of interest.

Despite being straightforward, equation (9) is not innocuous. For example, it is violated by welfare analyses of non-budget neutral policies that use marginal deadweight loss or marginal excess burden to adjust for the welfare cost of raising revenue. As discussed in Appendix C, the standard (equivalent variation) version of marginal deadweight loss corresponds to a budget-neutral policy, P_{Tax}, that imposes a tax and returns the revenue lump-sum to the agent. Equation (9) shows that marginal deadweight loss cannot be combined with a non-budget neutral (e.g. expenditure) policy to form a budget-neutral comprehensive policy.[35]

The motivation for creating a notion of the marginal cost of public funds (MCPF) is to provide an adjustment to the welfare analysis of non-budget neutral policies. Hence, a particularly natural definition of the MCPF is the welfare impact per dollar change in the government budget. To be specific, let P denote a non-budget neutral policy. I define the marginal cost of public funds in units of individual $\hat{i}$'s income to be

$$MCPF_P^{\hat{i}} = \frac{\int_i \frac{\eta_i}{\eta_{\hat{i}}} \frac{\frac{d\hat{V}_i^P}{d\theta}|_{\theta=0}}{\lambda_i} di}{\int_{i \in I} \frac{d\hat{t}_i^P}{d\theta} di} \tag{11}$$

which is the sum of the welfare impact on each individual, $\frac{\frac{\partial \hat{V}_i^P}{\partial \theta}|_{\theta=0}}{\lambda_{\hat{i}}}$, normalized in units of dollars to individual $\hat{i}$.[36] Then, given a budget neutral policy, P, that can be decomposed into two policies, P_{Tax} and P_{Exp}, the additivity condition implies

$$\frac{d\hat{W}}{d\theta} = \eta^{\hat{i}} \left(MCPF_{P_{Exp}}^{\hat{i}} - MCPF_{P_{Tax}}^{\hat{i}} \right) \tag{12}$$

so that policy P_{Exp} provides a benefit of $MCPF_{P_{Exp}}^{\hat{i}}$ per dollar of government revenue and a cost of $MCPF_{P_{Tax}}^{\hat{i}}$ per dollar of government revenue. Intuitively, whether the comprehensive policy increases welfare depends on whether the expenditure policy has a greater benefit per unit of government revenue than the cost imposed by the tax policy of raising the revenue. 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*

[35] For equation (9) to hold, P_{Exp} must be a budget-neutral expenditure policy financed using lump-sum taxation. Intuitively, one would need to add a third policy that would impose a lump sum tax to raise the revenue spent on the expenditure policy.

[36] 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.

revenue through a tax on labor supply, $\hat{\tau}(\theta)$.[37] *The marginal welfare impact of this policy is*

$$\frac{\frac{\partial \hat{V}^{P_{Tax}}}{\partial \theta}|_{\theta=0}}{\lambda} = -1 + \tau \frac{d\hat{l}_{P_{Tax}}}{d\theta}|_{\theta=0} \tag{13}$$

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 MCPF 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 MCPF of the tax policy is given by

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

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,

$$\frac{\frac{\partial \hat{V}^{P_{Exp}}}{\partial \theta}|_{\theta=0}}{\lambda} = \left(\frac{\frac{\partial u}{\partial g}}{\frac{\partial u}{\partial x}} - c_g \right) + 1 + \tau \frac{d\hat{l}_{P_{Exp}}}{d\theta}|_{\theta=0} \tag{14}$$

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.

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

Combining equations (13) *and* (14),

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

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)*.*

The definition of the marginal cost of public funds in equation (11) has several features. First, it is defined separately for any non-budget neutral policy, $\int_{i\in I}\frac{d\hat{t}_i^P}{d\theta}di \neq 0$. So, one can talk not only of the MCPF of increased taxation but also the MCPF of reduced education spending or other potential financing mechanisms.

Second, identifying heterogeneity in the MCPF across different policies is equivalent to identifying welfare-improving budget neutral policies. Given two non-budget neutral policy paths, P^1 and P^2, there exists a welfare improvement for financing more of policy P^1 and less of P^2 if and only if $MCPF^{\hat{i}}_{P^1} > MCPF^{\hat{i}}_{P^2}$ (note this comparison does not depend on the choice of income units, $\hat{i}$, as long as the same units are used for the MCPF of P^1 and P^2). This suggests analyses of expenditure policies can simply report their MCPF as a measure of "cost-effectiveness" of the government spending, as opposed to embedding external estimates of a MCPF from other studies. Intuitively, social welfare is improved when resources are allocated from less cost-effective policies to more cost-effective policies. This definition of the MCPF provides a formal and generic toolkit for making such comparisons.

Third, although previous literature has debated extensively about whether the MCPF should be constructed using compensated or uncompensated elasticities, my definition in general does not rely on either. It depends on the causal effect of the non-budget neutral policy in question.[38]

Fourth, the definition of the MCPF requires the researcher to be clear about whose income is being used to measure welfare, $\hat{i}$. One can easily move from $\hat{i}$ to $\hat{j}$'s income by multiplying by the ratio $\frac{\eta_{\hat{j}}}{\eta_{\hat{j}}}$. But, one must be explicit about whose income units are being used for the construction of the MCPF, a point previously made by Slemrod and Yitzhaki (1996, 2001) and Kaplow (2008). I return to this issue in the analysis of redistribution in Section 4.

[38]This causal effect may be either a compensated response, an uncompensated response, or neither. For tax policies, 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 34, comparisons of the values of MCPF are implicitly constructing budget neutral policies (e.g. $MCPF^{\hat{i}}_{P_1} - MCPF^{\hat{i}}_{P_2}$ 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.

Relation to Previous Literature Of course, this is not the first paper to propose a definition of the marginal cost of public funds. There are many different conceptual definitions[39]. Equation (10) perhaps provides an explanation for why: there is an infinite number of ways to split a comprehensive policy into a tax and expenditure policy. Each tax policy simply imposes a different requirement on the expenditure policy and vice-versa. For example, one could use the marginal deadweight loss as a measure of the marginal cost of public funds, as in the so-called Pigou-Harberger-Browning tradition (Pigou (1947); Harberger (1964); Browning (1976, 1987)). But, then the expenditure policy must finance the policy using lump-sum taxation (and one must incorporate the associated income effects as discussed in footnote 35). Because expenditure policies rarely finance the expenditure with lump-sum taxation, such a decomposition adds an empirical burden for identifying the welfare impact of comprehensive policies. In contrast, my definition of the MCPF aligns the hypothetical policy experiment with the actual policy experiment, so that welfare analysis depends solely on the policy elasticities; a decomposition of causal effects into income and substitution effects are not required.

My definition of the MCPF is closely related to Slemrod and Yitzhaki (1996, 2001), who also define a MCPF for tax policies as the per-dollar welfare impact of the tax change and generally interpret the desired behavioral responses to be uncompensated elasticities. My definition of the MCPF has several distinctions. Most importantly, it is applicable to any non-budget neutral policy, including hybrids of tax and expenditure policies. This allows it to closely align with actual government policy proposals. More technically, my construction of the MCPF shows that the desired behavioral responses are in general neither compensated nor uncompensated responses; one instead needs the policy elasticities. Hence, the MCPF can be estimated using the causal effects of actual non-budget neutral policy changes.

Finally, at a more abstract level, my approach contrasts with the broader spirit of the MCPF literature in that I do not attempt to define a single MCPF; rather, identifying heterogeneity in the MCPF across policies is equivalent to identifying welfare improving policies, as shown in equation (12). If there were a single MCPF for all policies, the marginal cost/benefit of government spending on all policies would be the same.

4 Redistribution

I apply the framework to study the desirability of additional redistribution. 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.[40] 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. However, the desirability

[39] See the citations in Footnote 2 and Dahlby (2008) for a recent overview. Also, Fullerton (1991) provides evidence that these definitions matter: he shows heterogeneity across three estimates of the marginal excess burden are rationalized by the different conceptual definitions employed in each of the papers, not by different data usage or empirical methods.

[40] 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.

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 rich.

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[41]. 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 3. $\frac{d\hat{W}_P}{d\theta}|_{\theta=0} \geq 0$ *if and only if*

$$\underbrace{\frac{\eta^{Rich} - \eta^{Poor}}{\eta^{Poor}}}_{\substack{\textit{Social Preference} \\ \textit{for Add'l Redistribution}}} \leq \underbrace{\frac{\int_i \left(\sum_j \tau_{ij}^x \frac{dx_{ij}}{d\theta}|_{\theta=0} + \tau_{ij}^l \frac{dl_{ij}}{d\theta}|_{\theta=0} \right) di}{\int_{i \in Rich} \frac{d\hat{\tau}^{Rich}}{d\theta}|_{\theta=0} \left(\hat{l}_i - \bar{l} \right) di}}_{\textit{Leaks in Okun's Bucket}} \tag{15}$$

Proof. The proof is a straightforward application of Proposition 1, and is provided in Appendix A.3. □

The LHS of equation (15) measures the marginal benefit to social welfare of transferring money from rich to poor. The RHS of equation (15) is the fraction of the mechanical revenue raised by taxing the rich that is lost due to behavioral responses.

[41] More specifically, the precise EITC policy expansion I consider is 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 causal impacts (policy elasticities) would not be too sensitive to the precise design of the phase-out of the program.

The intuition in equation (15) 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 (15) 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 (15) differs from approaches in previous literature studying optimal taxation with heterogeneous agents. These approaches generally use Hamiltonian-based methods introduced by 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 all elasticities across the income distribution, evaluated at their optimized levels (Piketty and Saez (2012)). Such an integral is difficult to estimate in practice, and hence studies of optimal taxation often only comment on tax rates, not the level of redistribution (Piketty and Saez (2012)). In contrast, Equation (15) 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 benefits to the poor should be increased through the particular redistributive policy in question.

4.2 Empirical implementation: A MCPF Approach

The leaks in Okun's bucket in equation (15) 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 tend to 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

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

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

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

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

Using equation (12) and the ratio of social marginal utilities of income, $\frac{\eta^{Rich}}{\eta^{Poor}}$, to translate $MCPF^{Rich}_{PTax}$ into units of income to the poor, the welfare impact of additional redistribution is given by

$$\frac{d\hat{W}_P}{d\theta}|_{\theta=0} = \eta^{Poor} MCPF^{Poor}_{PEITC} - \eta^{Rich} MCPF^{Rich}_{PTax}$$

which yields the following Corollary to Proposition (3).

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

$$MCPF^{Poor}_{PEITC} - \frac{\eta^{Rich}}{\eta^{Poor}} MCPF^{Rich}_{PTax} \geq 0 \tag{16}$$

Whether additional redistribution is desirable depends on whether the marginal value of the expenditure, given by $MCPF^{Poor}_{PEITC}$, is greater than the cost, given by $MCPF^{Rich}_{PTax}$. Since the MCPF 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 $MCPF^{Rich}_{PTax}$ and $MCPF^{Poor}_{PEITC}$ 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 MCPF of raising revenue from the rich through an increase in the top marginal tax rate is given by

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

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

the tax increase,

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

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.[42] 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). Although this literature estimates r using causal effects of changes to top marginal tax rates, there is often a goal of attempting to subsequently decompose these behavioral responses into income and substitution effects. Indeed, the parameter r is sometimes referred to as the "marginal excess burden" of the income tax (Saez et al. (2012); Feldstein (1999)). However, r is only a measure of the marginal excess burden if the individuals taxed are compensated for the impact of the policy change, so that $\frac{d\hat{l}_{ij}^{Tax}}{d\theta}|_{\theta=0}$ is a compensated response (Mas-Colell et al. (1995)). 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, 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.[43]

Saez et al. (2012) note that there is a wide range of estimates of the taxable income elasticity estimated as causal effects from tax policies, but they suggest a measure of 0.5 as relatively middle-of-the-road. Given their estimates of the income distribution, this implies that 50% of the mechanical revenue is lost due to behavioral responses, $r = -0.5$. Put differently, this suggests the marginal cost of raising \$1 in government revenue from an increase in marginal income tax rates on the rich imposes a \$2 welfare loss on those subjected to the tax increase., so that $MCPF_{P^{Tax}}^{Rich} = 2$.

There is of course considerable disagreement about the response to changes in the top marginal tax rate. Indeed, precise identification of the responses of top earners to taxation is a difficult empirical exercise. As a result, I will also consider the range of plausible values of $-r$ from 20-70%, as suggested by Giertz (2009).

[42] To see this, note that

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

[43] Of course, by taking estimates of r derived from previous literature, I am assuming that the policy elasticities of the top tax rate change today would be similar to the response to historical policy changes. If one knew how specific economic factors might change these policy elasticities relative to historical experience, one could adjust the causal estimates accordingly.

4.2.2 EITC Expansion

There is also a large literature estimating the causal effects of EITC expansions, especially impacts on single mothers. To conform existing causal estimates into an estimate of $MCPF_P^{Poor}$, 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.[44] 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, my analysis will underestimate the social desirability of increasing funding of the EITC.

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

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

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. These causal effects provide estimates of the policy elasticities necessary for computing p. 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

[44] 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)).

the poor single mothers that are in the labor force. With this notation, the impact of the behavioral response to the policy by the poor on the government's budget is given by:

$$-\int_{i\in Poor}\tau_i^l\frac{d\hat{l}_i^{EITC}}{d\theta}|_{\theta=0}di = \underbrace{\left(\tau^l_{i^{LFP}(0)}l_{i^{LFP}(0)}\right)\frac{di^{LFP}}{d\theta}|_{\theta=0}}_{\text{Extensive Margin}} - \underbrace{\int_{i<\hat{i}}\tau_i^l\frac{d\hat{l}_i^{EITC}}{d\theta}|_{\theta=0}di}_{\text{Intensive Margin}} \quad (17)$$

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. 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.[45]

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 (17) by normalizing θ to parameterize an additional unit of the mechanical subsidy[46] 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 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 $\frac{E\left[\tau^l_{\hat{i}}l_{\hat{i}}\right]}{E\left[\left(1-\tau^l_{\hat{i}}\right)l_{\hat{i}}\right]}$, 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%.

[45] 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.

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

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.

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.

If knowledge of the average subsidy of the EITC generates extensive margin responses and knowledge of the kink schedule (as proxied by the presence of self-employed bunchers) generates intensive margin responses, then 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 inChetty et al. (2013) is along the extensive margin responses. 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 $MCPF_{PEITC}^{Poor} = 0.88$. 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 Combining the midpoint estimates suggests 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)

Table 1 shows how the welfare analysis varies with different assumptions about r and p. Taking estimates from the range of previous literature ($-r \in [0.2, 0.7]$ as suggested by Giertz (2009) and the fraction of EITC revenue lost due to responses of 5-16%) yields a range of $\frac{\eta^{Rich}}{\eta^{Poor}} \in [0.26, 0.76]$. Hence, at the upper range of the elasticity estimates, the results suggest roughly 75% of the mechanical revenue gain would be lost due to included behavioral distortions, so that redistribution is desirable only if one prefers \$0.26 in the pocket of an EITC recipient relative to \$1 in the pocket of someone subject to the top marginal tax rate. On the other end of existing estimates, redistribution is desirable as long as one prefers providing \$0.76 to an EITC beneficiary over an additional \$1 to someone subject to the top marginal tax rate.

Table 1: MCPF_Poor / MCPF_Rich

		r				
		0.00	-0.20	**-0.50**	-0.80	-1.00
p	MCPF	1.00	1.25	**2.00**	5.00	N/A
0.00	1.00	1.00	0.80	**0.50**	0.20	0.00
0.05	0.95	0.95	0.76	**0.48**	0.19	0.00
0.14	**0.88**	**0.88**	**0.70**	**0.44**	**0.18**	**0.00**
0.16	0.86	0.86	0.69	**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 a 56% marginal reduction in economic activity 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. 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. The results show the causal effect of the policy in question on behavior that affects the government budget are the key behavioral responses required for welfare analysis. Because these responses vary with the policy in question, they are in general neither Hicksian nor Marshallian price elasticities; I term them policy elasticities.

My results suggest estimates of causal impact of policies can readily be translated into a general welfare framework. In particular, I hope the framework is useful for those employing the rise in reduced-form methods for evaluating changes to government policies which generally do not have the modeling structure provided by structural methods (which readily yield welfare statements). Indeed, translating the growing literature that uses reduced-form methods to estimate the causal effects of government policies and using them to derive their implicit MCPF would seem particularly promising. It has the potential for the creation of a new volume of estimates of the MCPF for different policies and a more comprehensive analysis of the desirability of potential government policy changes.

References

Allgood, S. and A. Snow (1998). The marginal cost of raising tax revenue and redistributing income. *Journal of Political Economy 106*, 1246–1273. 2

Angrist, J. and J.-S. Pischke (2008). *Mostly Harmless Econometrics.* Princeton University Press. 1, 2.2

Angrist, J. and J.-S. Pischke (2010). The credibility revolution in empirical economics: How better research design is taking the con out of econometrics. *Journal of Economic Perspectives 24*(2), 3–30. 1, 1

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. 2, 10, 28

Auerbach, A. (1985). The theory of excess burden and optimal taxation. Handbook of public economics, Elsevier. 2, 50

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. (1990). Marginal welfare cost calculations. *Journal of Public Economics 41*, 263–276. 2

Ballard, C. L., J. B. Shoven, and J. Whalley (1985, March). General equilibrium computations of the marginal welfare costs of taxes in the united states. *The American Economic Review 75*(1), 128–138. 2

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

Browning, E. (1976). The marginal cost of public funds. *Journal of Political Economy 84*, 283–98. 2, 3

Browning, E. (1987). On the marginal welfare cost of taxation. *American Economic Review 77*, 11–23. 2, 3

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

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, 30, 2.5, B

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)*. 6, 29, 41, 4.2.2, 44, 4.2.2

Dahlby, B. (2008). *The Marginal Cost of Public Funds*. The MIT Press. 2, 39, C.2

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. 1, 2.5, 31, C.1, C.1, C.1

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. 19

Eissa, N. and J. Liebman (1996). Labor supply responses to the earned income tax credit. *The Quarterly Journal of Economics 111*(2), 605–637. 44, 4.2.2

Feldstein, M. (1976). On the theory of tax reform. *Journal of Public Economics 6*, 77–104. 8

Feldstein, M. (1995). The effect of marginal tax rates on taxable income: A panel study of the 1986 tax reform act. *Journal of Political Economy 103*(3), 551–572. 2.5

Feldstein, M. (1999). Tax avoidance and the deadweight loss of the income tax. *Review of Economics and Statistics 81*(4), 674–680. 1, 2.4, 2.5, 4.2.1

Feldstein, M. (2012, September). The mirrlees review. *Journal of Economic Literature 50*(3), 781–790. 1, 12

Fullerton, D. (1991). Reconciling recent estimates of the marginal welfare cost of taxation. *American Economic Review 81*, 302–08. 2, 39

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. 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, 11

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

Harberger, A. (1964). The measurement of waste. *American Economic Review 54*, 58–76. 2, 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. 6, 4.2.2, 4.2.2

Kaplow, L. (1996). The optimal supply of public goods and the distortionary cost of taxation. *National Tax Journal 49*, 513–534. 13

Kaplow, L. (2004). On the (ir)relevance of the distribution and labor supply distortion to government policy. *Journal of Economic Perspectives 18*, 159–175. 13

Kaplow, L. (2008). *The Theory of Taxation and Public Economics*. Princeton University Press. 13, 3, 51

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995, June). *Microeconomic Theory*. Oxford University Press, USA. 1, 12, 26, 2.5, 4.2.1, C.2

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

Mayshar, J. and J. Slemrod (1995). Dalton-improving tax reform. *American Economic Review 85*, 4. 2

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, C.2

Nevo, A. and M. D. Whinston (2010). Taking the dogma out of econometrics: Structural methods and credible inference. *Journal of Economic Perspectives*, 69–82. 1, 1

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

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*. 9, 4.1

Ramsey, F. P. (1927, March). A contribution to the theory of taxation. *The Economic Journal 37*(145), 47–61. 1, 2.5, C.1

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, 6, 12, 4.2.1

Salanie, B. (2003). *The Economics of Taxation*. MIT Press. 48

Samuelson, P. (1954). The pure theory of public expenditure. *The Review of Economics and Statistics 36*, 387–389. 1, 2.3

Schlee, E. E. (2013). Radner's cost-benefit analysis in the small: An equivalence result. *Working Paper*. 23

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. 2, 10, 3, 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*. 2, 10, 13, 3, 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. 2

Stuart, C. (1984). Welfare costs per dollar of additional tax revenue in the united states. *American Economic Review 74*(3), 352–362. 2, 10

Wildasin, D. (1984). On public good provision with distortionary taxation. *Economic Inquiry 22*(2), 227–43. 2

Online Appendix: Not for Publication

A Appendix: Proofs

A.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

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

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

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

Replacing terms, I have

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

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

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

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

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

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.

A.2 MCPF with Hicksian Elasticities

Given the prominence of the Hicksian elasticity and role of MDWL in previous theoretical and empirical literature, it is perhaps helpful to explain when it could be used as a measure of the marginal cost of public funds. Let $P_{Tax}(\theta)$ denote a policy that increases the tax rate and returns money lump-sum in a manner such that utility is held constant. Clearly, $\frac{dV^{P_{Tax}}}{d\theta} = 0$ by construction. However, the tax policy is not budget neutral; the government runs a marginal deficit because of the lump-sum transfers needed to hold the agent's utility constant. The deficit is equal to:

$$\frac{d\hat{t}_{P_{Tax}}}{d\theta} = -\tau\frac{d\hat{l}_{P_{Tax}}}{d\theta}|_{\theta=0}$$

where $\frac{d\hat{l}^{P^{Tax}}}{d\theta}$ is the hicksian response to raising 1 unit of revenue through labor taxation.

In order for the additivity condition to hold, the corresponding expenditure policy must use lump-sum taxation to finance not only the roads but also the budget deficit induced by the tax policy.

Hence, $P_{Exp}(\theta)$ *must raise a surplus*,

$$\frac{d\hat{t}^{P_{Exp}}}{d\theta} = \tau \frac{d\hat{l}^{P_{Tax}}}{d\theta}$$

and provide one unit of roads, $\frac{dG}{d\theta} = 1$. Thus, the welfare impact of the expenditure policy (which equals the welfare impact of the combined policy) is given by

$$\frac{\frac{\partial V}{\partial \theta}^{P_{Exp}}|_{\theta=0}}{\lambda} = \left(\frac{\frac{\partial u}{\partial G}}{\frac{\partial u}{\partial x}} - c_g\right) + \tau \frac{d\hat{l}^{P_{Tax}}}{d\theta} + \tau \frac{d\hat{l}^{P_{Exp}}}{d\theta} \tag{18}$$

where $\frac{d\hat{l}^{P_{Exp}}}{d\theta}$ is the behavioral response to a policy that simultaneously increases government spending on public goods and raises more revenue than the cost of the public goods via lump-sum taxation. Therefore, one can use MDWL (CV) as a measure of the welfare cost of distortionary taxation; but one must then be able to estimate the impact of an expenditure policy that not only finances the public goods, but also imposes a lump-sum tax of a greater amount than the cost of the public goods. Intuitively, the behavioral response required from the expenditure policy simultaneously adds back in the income effects from lump-sum taxation and effects from the road expenditure.

Equations (**??**) and (18) present two different potential approaches to estimating the welfare impact of the same policy. If one knew the Hicksian elasticity or the MDWL, one could attempt to identify the corresponding expenditure policy. But, since most expenditure policies are not financed out of lump-sum taxation, it is arguably more natural to define the marginal costs of public funds as the welfare costs of policies that raise government revenue.

A.3 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

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

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

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

B 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 $E_i = f_i^E\left(\mathbf{x}\right)$ and $\mathbf{x}$ is the vector of solutions to each agents optimization program.

All other definitions from Section 2 are maintained. In particular, policy paths are defined as in equation 4.[47] Proposition 2 presents the characterization of the marginal welfare impact of a policy evaluated at $\theta = 0$.

Proposition 4. *The welfare impact of the marginal policy change to type i is given by*

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

[47] 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).

where

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

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

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

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

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

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

Replacing terms, I have

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

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.[48] 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.

C Appendix: Relation to the Hicksian (Compensated) Elasticity

As discussed in the introduction, previous literature has highlighted the role of Hicksian (compensated) elasticities in the welfare evaluation of government policy. This appendix outlines the two classic cases where the Hicksian elasticity arises: (1) optimal commodity taxation and the "inverse elasticity" rule and (2) marginal deadweight loss from distortionary taxation.

C.1 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.

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.

[48] 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.

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).[49]

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 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.

C.2 Tax Distortions and Marginal Deadweight Loss

It is also well known that compensated elasticities measure the marginal excess burden (MEB) or deadweight loss (MDWL) of the tax system. To illustrate this, again focus on one representative agent and drop i subscripts. Assume there is just one good, $J_X = 1$, and one type of labor supply, $J_L = 1$. Let τ denote the marginal tax on labor supply. Let $P(\theta)$ denote a policy which increases the tax rate $\hat{\tau}(\theta) = \tau + \theta$. Assume the policy does not change publicly provided goods. Thus, the

[49] 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.

marginal welfare impact of the tax increase is given by

$$\frac{\frac{\partial \hat{V}}{\partial \theta}}{\lambda_i}|_{\theta=0} = \frac{d\hat{t}}{d\theta}|_{\theta=0} + \tau \frac{d\hat{l}}{d\theta}|_{\theta=0}$$

where for now I refrain from specifying the impact on net transfers, $\frac{d\hat{t}}{d\theta}|_{\theta=0}$.

Previous literature defines two measures of MDWL: a "compensated variation" (CV) and equivalent variation (EV) measure (see Mas-Colell et al. (1995) or Dahlby (2008) for a discussion). I discuss each of these in turn.

CV Measure of DWL The compensating variation (CV) measure of deadweight loss is the increased revenue a government could obtain switching from the distortionary taxation to lump-sum taxation while holding the agent's utility constant. Thus, marginal CV is the value of $\frac{d\hat{t}}{d\theta}|_{\theta=0}$ that solves the equation $\frac{\frac{\partial \hat{V}}{\partial \theta}|_{\theta=0}}{\lambda} = 0$, or

$$MDWL_{CV} = \frac{d\hat{t}}{d\theta}|_{\theta=0} = -\tau \frac{d\hat{l}}{d\theta}|_{\theta=0}$$

Because utility is held constant in this policy experiment, $\frac{d\hat{l}}{d\theta}|_{\theta=0}$ is precisely the compensated (Hicksian) response to the policy. Thus, the Hicksian elasticity captures the additional revenue a government could collect from agents if they switched from distortionary to lump-sum taxation but held utilities constant.[50]

While the CV measure of DWL provides a measure of the welfare cost of distortionary taxation, the hypothetical policy experiment is left incomplete: it does not specify how the additional revenue is spent by the government. Since this spending would affect utilities and behavior, it would need to be incorporated into a welfare analysis of an actual policy aimed at reducing the distortionary burden of taxation.[51] Welfare analysis of such a policy would no longer require the Hicksian elasticity; rather, it would require the causal effect of that particular policy – i.e. the policy elasticity.

EV Measure of DWL The second measure of deadweight loss proposed in the literature is the equivalent variation (EV) measure of deadweight loss. This is the agent's willingness to pay for switching from distortionary taxation to individual-specific lump-sum taxation in a manner that holds the government's budget constraint constant. So, the EV measure of marginal deadweight loss is the value of $-\frac{\frac{\partial \hat{V}}{\partial \theta}|_{\theta=0}}{\lambda_i}$ when the marginal revenue from the tax is returned lump-sum, so that $\frac{d\hat{t}}{d\theta}|_{\theta=0} = 0$. Hence,

$$MDWL_{EV} = -\frac{\frac{\partial \hat{V}}{\partial \theta}|_{\theta=0}}{\lambda} = -\tau \frac{d\hat{l}}{d\theta}|_{\theta=0}$$

[50] This interpretation of the CV measure of DWL and the Hicksian elasticity is well-known (e.g. Auerbach (1985)).

[51] This criticism of the CV measure of DWL is related to the point raised by Kaplow (2008) that a policy experiment should completely specify all features of the policy.

Here, $\frac{d\hat{l}}{d\theta}|_{\theta=0}$ is not the Hicksian elasticity. Rather, it is the behavioral response to a policy which increases marginal taxes, τ, and returns the revenue to the agent through lump-sum transfers. Because of the distortions, the lump-sum transfers are insufficient to hold the agents' utility constant.[52] Thus, $\frac{d\hat{l}}{d\theta}|_{\theta=0}$ is a type of compensated elasticity (the tax revenue is returned to the agent), but it is not a "fully" compensated (Hicksian) elasticity.

In contrast to the CV measure of deadweight loss, the EV measure does correspond to a budget-neutral policy experiment. However, the corresponding policy assumes that the revenue is returned to precisely the same agent from whom it is taxed. To the extent to which distortionary taxation arises because of the in-feasibility of individual-specific lump-sum transfers (Mirrlees (1971)), this policy is not only unlikely to be observed in practice, but is infeasible given the information constraints imposed on the government. In sum, the results suggest the Hicksian elasticity is not particularly useful for marginal welfare analysis.

[52] It is straightforward to show that

$$\frac{\frac{\partial \hat{V}}{\partial \theta}|_{\theta=0}}{\lambda} = \frac{\tau^l \frac{\partial l^h}{\partial \tau^l}}{1 - \tau^l \frac{\partial l^m}{\partial y}}$$

where $\frac{\partial l^h}{\partial \tau^l}$ is the slope of the Hicksian demand function and $\frac{\partial l^m}{\partial y}$ is the (Marshallian) income effect.