TEXTO PARA DISCUSSÃO

No.708

Optimal Policy Reforms

Katy Bergstrom

William Dodds

Juan Rios

DEPARTAMENTO DE
ECONOMIA
PUC RIO

Optimal Policy Reforms*

Katy Bergstrom†, William Dodds‡, and Juan Rios§

12/12/2024

Abstract

This paper develops a general framework to construct optimal policy reforms starting from a status quo set of policies. We show that if a policymaker can control how fiscal externalities are spent, then the welfare-weighted marginal value of public funds (WMVPF) is the relevant sufficient statistic for determining optimal policy reforms. If a policymaker cannot control how fiscal externalities are spent, then the welfare-weighted net social benefit (WNSB) is the relevant sufficient statistic. If a policymaker can control how a fraction of fiscal externalities are spent, then the relevant sufficient statistic is an "internal WMVPF" plus an "external correction" term. We provide a number of stylized examples to illustrate when in practice to use the WMVPF vs. the WNSB to determine optimal policy reforms.

Keywords: optimal reforms, marginal value of public funds, net social benefit, sufficient statistics

JEL: H20, H30, H50

*This paper has benefited from conversations with Nathan Hendren, Ben Sprung-Keyser, and Raj Chetty. Also, thanks to seminar participants at the NBER Public Economics Meeting and Tulane University for their comments and suggestions.

†Tulane University. Email: kbergstrom@tulane.edu

‡Tulane University. Email: wdodds@tulane.edu

§Pontifical Catholic University of Rio de Janeiro. E-mail: juanrios@econ.puc-rio.br

1 Introduction

Over the last half century, there has been a substantial amount of work measuring the welfare impacts of policy reforms across many different policy arenas. This has led to the creation and use of several welfare metrics: the marginal excess burden (MEB) (e.g., Eissa et al. (2008), Eissa and Hoynes (2011)), the net social benefit (NSB) (García et al. (2023), Olken (2007)), the benefit cost ratio (BCR) (e.g., Heckman et al. (2010), García et al. (2020), Parker and Vogl (2023)), and the marginal value of public funds (MVPF) (e.g., Hendren and Sprung-Keyser (2020)). Consequently, debates have emerged as to which of these metrics are relevant to analyze and compare policy reforms (e.g., García and Heckman (2022a), García and Heckman (2022b), Hendren and Sprung-Keyser (2022)). As such, it is still an open question as to how the government should *prioritize* modifying existing policies. For example, if the government seeks to enact a budget neutral policy reform, what welfare metric should they use to determine which policies to change?

This paper seeks to answer this question by developing a general theory for determining the optimal policy reform across an arbitrary set of policies. In doing so, we show which of the commonly used welfare metrics are relevant to solve optimal local policy reform problems. We begin by building a general model of policy reforms with a policymaker who chooses a set of policy parameters, which might represent spending on various programs, tax rates, eligibility thresholds, regulations, or subsidy amounts. Individuals make decisions over some arbitrary set of choice variables given the policies chosen by the policymaker and the policymaker aggregates individual utilities via a welfare function. We then consider an optimal local policy reform problem that takes as given a status quo set of policy parameters and seeks to determine the optimal budget neutral policy reform which both raises and spends a small amount of money.

We first need to carefully define "spends a small amount of money" and "raises a small amount of money". To begin with, we assume that "spends (or raises) a small amount of money" refers to the *net* budgetary effect of a policy on government revenue. For example, consider a policy reform that increases spending on preschool education and costs \$1,000 in an upfront expenditure but, via increased tax revenue from a more educated population, generates \$500 of government revenue in the future (this \$500 is in present-value and is referred to as the *fiscal externality* of the reform). The net cost of this policy is therefore \$500 meaning that the government could spend \$2,000 upfront on this policy (effectively borrowing against the future revenue gains) and generate \$1,000 in fiscal externalities for a total *net* cost to the government of \$1,000. Similarly, if the government decided to reduce spending on preschool education to increase *net* government revenue by \$1,000, they could only cut spending by \$667 dollars today, recognizing that this will reduce tax revenue in the future by another \$333 via fiscal externalities.

We then consider the optimal budget neutral reform which changes some policies to increase (net) revenue by \$1 and then changes other policies to increase (net) costs by \$1: our first result, Proposition 1, shows that the optimal budget neutral reform involves raising revenue via the policy with the lowest welfare-weighted marginal value of public funds (WMVPF) and spending money on the policy with the highest WMVPF. The marginal value of public funds is an observable statistic measuring the total willingness-to-pay for a policy reform divided by the net cost of that reform (Hendren and Sprung-Keyser, 2020). By focusing on *local* policy reforms, our results do not require strong assumptions about how welfare changes with policy parameters far from the status quo: we show that optimal local policy reforms can be expressed in terms of empirically observable statistics estimated under existing policy parameters. The takeaway is that the WMVPF is the sufficient statistic for determining optimal policy reforms that raise and then spend a small amount of net money.[1]

One interesting aspect of Proposition 1 is the optimal way to raise money when there are policies which increase welfare *and* increase government revenue. Hendren and Sprung-Keyser (2020) assign such policies that "pay for themselves" infinite MVPFs because it is always welfare improving to spend money (upfront) on these policies. However, we show that for the purposes of determining the optimal way to raise revenue, it is useful to think of these policies as having negative MVPFs and, in turn, negative WMVPFs because this allows us to rank and compare these policies: we show that the policy with the *most* negative WMVPF is the optimal policy to raise a small amount of net revenue.[2]

In deriving Proposition 1, we make an important implicit assumption: policymakers can optimally spend the fiscal externalities that arise from their decisions. Proposition 1 is therefore relevant for policymakers who can choose how to spend fiscal externalities, such as a federal legislature deciding on tax policy or on generosity of entitlement programs. However, in practice other sorts of policymakers may not be able to choose how to spend fiscal externalities due to bureaucratic/political-economy constraints. In particular, policymakers may not be able to choose how to spend fiscal externalities if: (1) fiscal externalities accrue to agencies outside of the control of the policymaker in question and/or (2) if fiscal externalities accrue in the future and the policymaker is unable to borrow against future revenues (e.g., due to an annual use-it-or-lose it budget constraint). For example, if the Department of Education spends \$1,000 on preschool education, this may increase tax revenue for the Treasury Department many years in the future but the Department of Education may have no control over how this future revenue

[1] The term sufficient statistic sometimes refers to the positive (empirical) object needed to solve a problem; hence we are slightly abusing this term as the WMVPF is the product of a positive object, the MVPF, and a normative object, the average welfare weight.

[2] To raise revenue from negative MVPF policies, one must increase *upfront* spending, which generates positive *net* revenue for the government.

is spent. Alternatively, an international development organization (IDO) or non-governmental organization (NGO) may spend funds on some policy for which all fiscal externalities accrue to governments so that the IDO/NGO has no choice how fiscal externalities are spent. We show that when policymakers cannot control how fiscal externalities are spent, policymakers will be constrained by *mechanical* costs as opposed to net costs. We thus consider a policymaker who wants to find the optimal budget neutral policy reform that cuts spending by \$1 mechanically (i.e., upfront) on some policies and increases mechanical spending by \$1 on other policies; any fiscal externalities that accrue as a result of this reform are assumed to be spent on some "numeraire policy" which is not controlled by the policymaker in question. Our next result, Proposition 2, shows that the optimal budget neutral policy reform that raises \$1 in mechanical revenue and has a \$1 mechanical cost, involves raising mechanical revenue from what we refer to as the welfare-weighted net social benefit (WNSB) minimizing policy and increasing mechanical spending on the WNSB maximizing policy. The WNSB captures the welfare-weighted willingness to pay for a policy reform *minus* the welfare cost of financing fiscal externalities via a numeraire policy; the WNSB is therefore a welfare-weighted version of the net social benefit (NSB) as in García and Heckman (2022a).[3] Hence, the WNSB is the relevant sufficient statistic for determining optimal policy reforms when a policymaker cannot control how fiscal externalities are spent.

We then discuss a number of extensions. We show that our results continue to hold even if there is uncertainty over the costs and benefits of policies, and we show how to construct optimal reforms that only spend money or only raise money (i.e., are not budget neutral). We also show how to determine optimal policy reforms if a policymaker is able to choose how to spend a portion of fiscal externalities: in this case, the relevant welfare metric takes the form of an "internal WMVPF" (which includes all fiscal externalities controlled by the policymaker) plus an external welfare correction term (which accounts for fiscal externalities not controlled by the policymaker).

Given these results, a key takeaway is that whether to use the WMVPF or the WNSB for optimal policy reforms depends on the constraints that the policymaker faces: if the policymaker can choose how to spend fiscal externalities then the policymaker should base policy reform decisions on the WMVPF, but if they are constrained in terms of how fiscal externalities are spent, they should base policy reform decisions on the WNSB. Our next result, Proposition 4, illustrates and measures the welfare gain that results from relaxing constraints on how fiscal externalities are spent. From a policy perspective, Proposition 4 illustrates that agencies which cannot control how fiscal externalities are spent (e.g., because they accrue to another agency or

[3]The literature typically assumes that linear income taxation is the numeraire policy used to finance fiscal externalities; our definition of the WNSB allows for any arbitrary policy to be the numeaire policy.

they accrue far in the future) will choose sub-optimal policy reforms relative to a hypothetical policymaker who only faces a constraint on the total net costs of policy choices over time and can therefore spend any fiscal externalities that accrue in the future or to other agencies. We then show how a higher-level policymaker (such as a federal legislature) can reap the welfare gains described in Proposition 4 by using policy-specific subsidies (similar in spirit to Agrawal et al. (2022)) to change the constraints of agencies, rendering it optimal for agencies to use the WMVPF rather than the WNSB for policy reforms.

We then present a series of (proof-of-concept) numerical examples showcasing how to use our results to prescribe optimal policy reforms. First, we explore a hypothetical policy reform problem for the federal legislature in the United States. We suppose that the federal legislature is considering a budget neutral policy reform among six large government policies: an expanded earned income tax credit (EITC) for adults without dependents, increased benefit generosity for disability insurance (DI), a reduction in the top income tax rate, higher education tax deductions, increased Medicaid generosity for young children, and increased spending on housing vouchers. We use data on the costs and benefits of these policies from Hendren and Sprung-Keyser (2020) and assume that welfare weights are inversely proportional to average income of beneficiaries, loosely consistent with an elasticity of marginal utility over consumption of 1 (Chetty, 2006). We show that the optimal budget neutral policy reform (within this set of 6 potential policies) is to raise revenue from the lowest WMVPF policy (housing vouchers, which has a negative WMVPF and therefore the policymaker can increase upfront spending on this policy to raise both welfare and *net* revenue) and then spend on the highest WMVPF policy (EITC expansion). Our example highlights the importance of weighting the MVPFs by welfare weights (as the highest WMVPF policy, an EITC expansion, is *not* the same as the highest MVPF policy, a top income tax cut), as well as how to rank policies that "pay for themselves". We also discuss how to augment the federal legislature's problem when some fiscal externalities accrue to state governments and are therefore not controlled by the federal legislature.

Next, we consider a policymaker at the Department of Education deciding how to allocate discretionary spending grants among a set of education policies. We assume that the policymaker faces an annual budget constraint on the amount of grant spending and that she cannot control how fiscal externalities are spent because fiscal externalities accrue far in the future to the Treasury Department via increased income tax revenue. Hence, the policymaker faces a mechanical spending constraint, meaning that optimal policy reforms should therefore be made using the WNSB. We show how to construct the optimal policy reform given a hypothetical set of policies and we then illustrate the welfare gains that could be achieved if, hypothetically, the policymaker was able to optimally spend fiscal externalities arising from her policy decisions.

In this case, the policymaker would raise money from the WMVPF minimizing policy and then spend the resulting revenue on the WMVPF maximizing policy; we show that the welfare gains from relaxing the mechanical spending constraint (and therefore using the WMVPF over the WNSB for the purposes of optimal budget neutral policy reform) are substantial in this example.

Finally, we consider a stylized example with an international development organization (IDO) considering whether to spend a grant on an unconditional cash transfer or a conditional cash transfer in a developing country. We show that the optimal spending reform should be based on the WNSB if the IDO cannot control how the government of the developing country in question spends fiscal externalities. However, we argue that the IDO should base spending decisions on the WMVPF if they can dictate how the government spends fiscal externalities and that there are welfare gains from dictating how the government spends fiscal externalities in the context of cash transfers in rural Mexico.

Relationship to the Literature: This paper relates to at least three strands of the public economics literature: optimal reforms, sufficient statistics, and empirical welfare analysis. First, to the best of our knowledge, there is relatively little literature on optimal policy reforms. There have been a number of papers on optimal tax reforms (e.g., Feldstein (1976), Guesnerie (1977), Diewert (1978), and Dixit (1979)) as well as optimal reforms to social insurance programs (Huggett and Parra (2010) and Hosseini and Shourideh (2019)). We believe this is the first general analysis of optimal reforms for arbitrary government policies. Second, our paper relates to the literature on "sufficient statistics" (e.g., Chetty (2009)) insofar as we show how the WMVPF and the WNSB are sufficient statistics for optimal policy reform problems with net revenue constraints and mechanical revenue constraints, respectively. Because we focus solely on optimal reforms (rather than characterizing optimal policy), these sufficient statistics are directly observable under existing policy, thereby avoiding the criticism of Kleven (2020) that sufficient statistics are typically functions of the unknown optimal policy. Third, this paper contributes to the burgeoning literature on empirical welfare analysis. In particular, this paper contributes to ongoing debates regarding which metrics are needed for empirical welfare analysis (Hendren and Sprung-Keyser (2022), García and Heckman (2022a), García and Heckman (2022b)): we show that for the purposes of optimal policy reform, the WMVPF is the relevant sufficient statistic unless the policymaker cannot control how fiscal externalities are spent, in which case the WNSB becomes the relevant sufficient statistic.

This paper is organized as follows: Section 2 presents the general model of optimal policy reforms and then derives our main results showing when the WMVPF or the WNSB is the relevant sufficient statistic for determining optimal policy reforms. Section 3 discusses a number of theoretical extensions as well as our relationship to the optimal tax reform literature. Section

4 then illustrates how to measure the welfare gains from relaxing mechanical cost constraints and spending fiscal externalities optimally. Section 5 presents a number of numerical examples of optimal policy reforms. Section 6 concludes.

2 General Model of Optimal Policy Reforms

2.1 Model Setup

We consider a policymaker who chooses a set of policies $\mathbf{p} = \{p_1, p_2, ..., p_N\}$. The policies p_i represent parameters that encode government programs. For instance, p_i might represent government spending on K-12 education, the size of a universal basic income transfer, the marginal income tax rate in a linear income tax schedule (or more generally, marginal tax rates and location of kink points in a piecewise-linear tax schedule), the tax rate on capital gains, the maximum duration of unemployment benefits, or an eligibility threshold for government funded healthcare.[4] Individuals differ in terms of some vector of characteristics $\mathbf{n} \in \mathbf{N}$ distributed according to a density function $f(\mathbf{n})$. Individuals make choices over a vector $\mathbf{y} = \{y_1, y_2, ..., y_M\}$ and have a utility function $U(\mathbf{y}; \mathbf{p}, \mathbf{n})$. Let us denote the schedule of choices as a function of type $\mathbf{n}$ under policy $\mathbf{p}$ by $\mathbf{y}(\mathbf{n}; \mathbf{p})$. The policymaker has a welfare function W where $\phi(\mathbf{n})$ is a welfare weight on person $\mathbf{n}$:

$$W(\mathbf{p}) = \int_{\mathbf{N}} \phi(\mathbf{n}) U(\mathbf{y}(\mathbf{n}; \mathbf{p}); \mathbf{p}, \mathbf{n}) f(\mathbf{n}) d\mathbf{n}$$

Note that our framework can accommodate two important realisms. First, we have not assumed that individuals optimize correctly: $\mathbf{y}(\mathbf{n}; \mathbf{p})$ just represents individual choices under policy $\mathbf{p}$ and does not need to maximize utility. Second, our framework can accommodate general equilibrium effects as long as the utility function is construed appropriately: for example, if changing a policy p_i impacts wages or prices or changes externality-producing choices of other agents, these effects should be included when determining how utility $U(\mathbf{y}; \mathbf{p}, \mathbf{n})$ varies with the policy p_i holding the decisions of the individual in question fixed.

In order to implement a set of policies $\mathbf{p}$, the policymaker incurs a cost. Let $C(\mathbf{p}, \mathbf{y}(\mathbf{n}; \mathbf{p}))$ represent the cost of policy $\mathbf{p}$ as a function of the policy vector $\mathbf{p}$ and the function of optimal choices $\mathbf{y}(\mathbf{n}; \mathbf{p})$. It will be helpful to separate the net cost of a policy reform from the status quo policy $\tilde{\mathbf{p}}$ to the new policy $\mathbf{p}$ into "mechanical costs" of the reform and "fiscal externalities" of the reform:

$$\underbrace{C(\mathbf{p}, \mathbf{y}(\mathbf{n}; \mathbf{p})) - C(\tilde{\mathbf{p}}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}))}_{\text{Net Cost, } NC(\mathbf{p})} = \underbrace{C(\mathbf{p}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}})) - C(\tilde{\mathbf{p}}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}))}_{\text{Mechanical Cost, } MC(\mathbf{p})} + \underbrace{C(\mathbf{p}, \mathbf{y}(\mathbf{n}; \mathbf{p})) - C(\mathbf{p}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}))}_{\text{-Fiscal Externalities, } -FE(\mathbf{p})} \quad (1)$$

[4] We have implicitly assumed that all policies are characterized by sets of scalars so that we can work in a finite dimensional setting. However, many of the results herein hold in infinite dimensional settings, see Bergstrom et al. (2024b) who explore optimal tax reforms in an infinite dimensional setting.

where $NC(\mathbf{p})$ captures the net cost of moving from $\tilde{\mathbf{p}}$ to $\mathbf{p}$, $MC(\mathbf{p})$ captures the cost of moving from $\tilde{\mathbf{p}}$ to $\mathbf{p}$ holding household behavior fixed, and $-FE(\mathbf{p})$ captures the impact on cost resulting from households changing their behavior in response to the policy change. Note, when $FE(\mathbf{p}) > 0$ this means that changes in household behavior lead to a reduction in net cost to the government. The policymaker would ideally like to choose a policy vector $\mathbf{p}$ to solve the following global reform problem which maximizes welfare while remaining budget neutral:

Problem 1 Global. *Given a status quo policy* $\tilde{\mathbf{p}}$*:*

$$\max_{\mathbf{p}} \; W(\mathbf{p}) - W(\tilde{\mathbf{p}})$$
$$\text{s.t.} \; \underbrace{C(\mathbf{p}, \mathbf{y}(\mathbf{n}; \mathbf{p})) - C(\tilde{\mathbf{p}}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}))}_{\text{Net Cost, } NC(\mathbf{p})} = 0$$

However, Problem 1 Global is written in terms of objects that are empirically difficult (impossible?) to estimate: to solve Problem 1 Global the policymaker would need to know how $W(\mathbf{p})$ and $C(\mathbf{p}, \mathbf{y}(\mathbf{n}; \mathbf{p}))$ vary with $\mathbf{p}$ across a wide variety of hypothetical policies that have never been observed in reality. Hence, Problem 1 Global is inherently divorced from empirical estimates unless we are willing to make strong assumptions about how costs and benefits of policies vary far away from the observed status quo. As a result, we now consider a more modest objective which can be tackled with existing empirical estimates of costs and benefits: finding optimal *local* policy reforms to best increase the objective function subject to a budget neutrality constraint. Loosely, given a status quo policy $\tilde{\mathbf{p}}$, we are interested in choosing a policy $\mathbf{p}$ close to $\tilde{\mathbf{p}}$ which maximizes $W(\mathbf{p}) - W(\tilde{\mathbf{p}})$ and continues to satisfy the government's budget constraint. We will assume throughout that $W(\cdot)$ is differentiable in $\mathbf{p}$.[5] Hence, by Taylor's Theorem, when $\mathbf{p}$ is close to $\tilde{\mathbf{p}}$, $W(\mathbf{p}) - W(\tilde{\mathbf{p}}) \approx \sum_{i=1}^{N} \frac{dW(\mathbf{p})}{dp_i}\Big|_{\mathbf{p}=\tilde{\mathbf{p}}}(p_i - \tilde{p}_i)$.

Similarly, assuming $NC(\mathbf{p})$ is differentiable, we can calculate a Taylor series approximation

[5] By the envelope theorem, sufficient conditions for differentiability of $W(\cdot)$ are that all individuals optimize rationally, that almost all individuals have a unique optimum, and that $U(\mathbf{y}; \mathbf{p}, \mathbf{n})$ is differentiable in $\mathbf{p}$. We can relax the assumption that individuals optimize rationally as long as $U(\mathbf{y}(\mathbf{n}; \mathbf{p}); \mathbf{p}, \mathbf{n})$ is (totally) differentiable in $\mathbf{p}$. Also, welfare will often be differentiable even if some individuals have multiple optima and/or $U(\mathbf{y}; \mathbf{p}, \mathbf{n})$ is not differentiable in $\mathbf{p}$ which can occur, for example, if some p_i denotes the location of a kink or notch; see Appendix A.4 of Bergstrom et al. (2024a).

for the net cost of the reform from Equation (1) as follows:[6]

$$\underbrace{\sum_{i=1}^{N} \frac{dNC(\mathbf{p})}{dp_i}\bigg|_{\mathbf{p}=\tilde{\mathbf{p}}} (p_i - \tilde{p}_i)}_{\text{Net Cost}} = \underbrace{\sum_{i=1}^{N} \frac{dMC(\mathbf{p})}{dp_i}\bigg|_{\mathbf{p}=\tilde{\mathbf{p}}} (p_i - \tilde{p}_i)}_{\text{Mechanical Cost}} - \underbrace{\sum_{i=1}^{N} \frac{dFE(\mathbf{p})}{dp_i}\bigg|_{\mathbf{p}=\tilde{\mathbf{p}}} (p_i - \tilde{p}_i)}_{\text{Fiscal Externalities}} \tag{2}$$

where $\frac{dNC(\mathbf{p})}{dp_i}\big|_{\mathbf{p}=\tilde{\mathbf{p}}} = \frac{dC(\mathbf{p},\mathbf{y}(\mathbf{n};\mathbf{p}))}{dp_i}\big|_{\mathbf{p}=\tilde{\mathbf{p}}}$ is the total derivative of cost with respect to policy p_i, $\frac{dMC(\mathbf{p})}{dp_i}\big|_{\mathbf{p}=\tilde{\mathbf{p}}} = \frac{dC(\mathbf{p},\mathbf{y}(\mathbf{n};\tilde{\mathbf{p}}))}{dp_i}\big|_{\mathbf{p}=\tilde{\mathbf{p}}}$ is the impact of changing policy p_i on costs holding individual choices constant at $\mathbf{y}(\mathbf{n};\tilde{\mathbf{p}})$, and $\frac{dFE(\mathbf{p})}{dp_i}\big|_{\mathbf{p}=\tilde{\mathbf{p}}} = -\frac{dC(\tilde{\mathbf{p}},\mathbf{y}(\mathbf{n};\mathbf{p}))}{dp_i}\big|_{\mathbf{p}=\tilde{\mathbf{p}}}$ is the impact of changing policy p_i on costs that arise as a result of individuals changing $\mathbf{y}(\mathbf{n};\mathbf{p})$. Henceforth, we will omit the arguments of $\frac{dW}{dp_i}$, $\frac{dNC}{dp_i}$, $\frac{dMC}{dp_i}$, and $\frac{dFE}{dp_i}$ for brevity, recognizing that these derivatives are functions of $\mathbf{p}$ and are evaluated at $\tilde{\mathbf{p}}$ unless otherwise stated. We will refer to $\frac{dNC}{dp_i}(p_i - \tilde{p}_i)$ as the "net cost" of a policy change from $\tilde{p}_i$ to p_i and will refer to $-\frac{dNC}{dp_i}(p_i - \tilde{p}_i)$ as the "net revenue" of a policy change from $\tilde{p}_i$ to p_i. We will refer to $\frac{dMC}{dp_i}(p_i - \tilde{p}_i)$ as the "mechanical cost" and $-\frac{dMC}{dp_i}(p_i - \tilde{p}_i)$ as the "mechanical revenue" of a policy change from $\tilde{p}_i$ to p_i. Finally, we will refer to $\frac{dFE}{dp_i}(p_i - \tilde{p}_i)$ as the fiscal externality of a policy change from $\tilde{p}_i$ to p_i recognizing that a positive fiscal externality implies that changes in behavior resulting from raising policy p_i lead to a reduction in net cost to the government.

2.2 Optimal Policy Reforms and the MVPF

Given the notation from Section 2.1, we now show how to solve an optimal local reform version of Problem 1 Global. Precisely, to ensure that the net cost constraint continues to be satisfied while only considering reforms that are close to $\tilde{\mathbf{p}}$, the policymaker seeks to find the optimal budget neutral policy which raises net revenue by a small amount (by changing certain policies) and then increases net costs by a small amount (by changing other policies):[7]

[6] What are sufficient conditions for differentiability of the net cost function? Suppose that around the status quo $\tilde{\mathbf{p}}$ we can express $NC(\mathbf{p})$ as $\int_{\mathbf{N}} c(\mathbf{p}, \mathbf{y}(\mathbf{n};\mathbf{p})) f(\mathbf{n}) d\mathbf{n}$ for some $c(\mathbf{p}, \mathbf{y}(\mathbf{n};\mathbf{p}))$ which captures the cost of providing policy $\mathbf{p}$ to individual $\mathbf{n}$. If $c(\mathbf{p}, \mathbf{y})$ is differentiable in both arguments and all individuals have a unique optima with first and second order conditions holding strictly so that we can apply the implicit function theorem to $\mathbf{y}(\mathbf{n};\mathbf{p})$, then $NC(\mathbf{p})$ will be differentiable. But these conditions can be relaxed. Loosely, the cost function will still often be differentiable even if first order conditions aren't satisfied for some individuals (due to choosing $\mathbf{y}$ where $U(\mathbf{y};\mathbf{p},\mathbf{n})$ is not differentiable), or second order conditions hold only weakly for a measure zero set of individuals, or some measure zero set of individuals have two optima. See Proposition 1 of Bergstrom and Dodds (2024) which proves government costs are differentiable with these relaxed assumptions in an infinite dimensional setting; the logic carries over almost immediately to the finite dimensional setting.

[7] Problem 1 is an optimal local reform problem corresponding to Problem 1 Global insofar as the welfare gain from solving Problem 1 equals 0 if we start from any status quo $\tilde{\mathbf{p}}$ which is a local maximum of Problem 1 Global (if not then Problem 1 would identify a budget neutral reform which improved welfare).

Problem 1. *Given a status quo policy* $\tilde{\mathbf{p}}$*:*

$$\max_{\mathbf{p}} \sum_{i=1}^{N} \frac{dW}{dp_i}(p_i - \tilde{p}_i)$$
$$\text{s.t.} \sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) > 0\right] = 1$$
$$\sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) < 0\right] = -1$$

Notationally, it is important to remember that $\frac{dW}{dp_i}(p_i - \tilde{p}_i)$ is the product of $\frac{dW}{dp_i} = \frac{dW(\mathbf{p})}{dp_i}\big|_{\mathbf{p}=\tilde{\mathbf{p}}}$ and $(p_i - \tilde{p}_i)$ and $\frac{dNC}{dp_i}(p_i - \tilde{p}_i)$ is the product of $\frac{dNC}{dp_i} = \frac{dNC(\mathbf{p})}{dp_i}\big|_{\mathbf{p}=\tilde{\mathbf{p}}}$ and $(p_i - \tilde{p}_i)$. Problem 1 aims to find a policy $\mathbf{p}$ that maximizes the first order approximation to $W(\mathbf{p}) - W(\tilde{\mathbf{p}})$ subject to the constraint that the government changes some policy parameters to raise \$1 in net revenue and then changes some other policy parameters to increase net costs by \$1. Note that the constraints in Problem 1 are *stronger* than simply mandating that the budget constraint continues to be satisfied under the new policy $\mathbf{p}$ so that $\sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0$. We need these stronger constraints to ensure that $\mathbf{p}$ remains close to $\tilde{\mathbf{p}}$ so that Problem 1 makes sense as a local version of Problem 1 Global; if we only constrain the problem so that $\sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0$, the policymaker could raise arbitrarily large sums of money from certain policies and then spend that money on other policies at which point approximating the welfare and cost impacts via Taylor series is no longer appropriate. Before we get to the solution of Problem 1, it will be first useful to define the WMVPF:

Definition 1 (WMVPF)**.** *The welfare-weighted marginal value of public funds (WMVPF) of policy i is defined as the ratio of the welfare impact of changing policy i to the net cost of changing policy i:*

$$WMVPF_i \equiv \frac{\frac{dW}{dp_i}}{\frac{dNC}{dp_i}} = \frac{\int_{\mathbf{N}} \phi(\mathbf{n}) \frac{dU(\mathbf{y}(\mathbf{n};\mathbf{p});\mathbf{p},\mathbf{n})}{dp_i} f(\mathbf{n}) d\mathbf{n}}{\frac{dNC}{dp_i}} = \bar{\eta} \underbrace{\frac{\int_{\mathbf{N}} WTP_i(\mathbf{n}) f(\mathbf{n}) d\mathbf{n}}{\frac{dNC}{dp_i}}}_{\text{MVPF}}$$

where $WTP_i(\mathbf{n}) \equiv \frac{dU(\mathbf{y}(\mathbf{n};\mathbf{p});\mathbf{p},\mathbf{n})}{dp_i} / \lambda(\mathbf{n})$ is the willingness-to-pay (WTP) for a marginal change to policy i, $\lambda(\mathbf{n})$ represents marginal utility of consumption for type $\mathbf{n}$, $\bar{\eta} \equiv \frac{\int_{\mathbf{N}} \phi(\mathbf{n})\lambda(\mathbf{n}) WTP_i(\mathbf{n}) f(\mathbf{n}) d\mathbf{n}}{\int_{\mathbf{N}} WTP_i(\mathbf{n}) f(\mathbf{n}) d\mathbf{n}}$ is an incidence-weighted average welfare weight, and $\phi(\mathbf{n})\lambda(\mathbf{n})$ is the marginal social value of consumption for type $\mathbf{n}$.[8]

[8] Note that Definition 1 does not rely on the envelope theorem (and therefore does not assume households necessarily optimize correctly) as the WMVPF is a function of $\frac{dU(\mathbf{y}(\mathbf{n};\mathbf{p});\mathbf{p},\mathbf{n})}{dp_i}$. Practitioners may nonetheless appeal to the envelope theorem which implies behavioral effects have second order utility impacts so that $\frac{dU(\mathbf{y}(\mathbf{n};\mathbf{p});\mathbf{p},\mathbf{n})}{dp_i} = \frac{\partial U(\mathbf{y};\mathbf{p},\mathbf{n})}{\partial p_i}$. We can also relax the assumption that the policymaker is a weighted-utilitarian; all of our results go through even though the second two equalities in Definition 1 no longer hold.

Proposition 1. *The solution to Problem 1 is to set:*

$$(p_i - \tilde{p}_i) = \begin{cases} \frac{1}{\frac{dNC}{dp_i}} & i = \arg\max\{WMVPF_i\} \\ \frac{-1}{\frac{dNC}{dp_i}} & i = \arg\min\{WMVPF_i\} \\ 0 & otherwise \end{cases}$$

Proof. See Appendix A.1. □

Proposition 1 establishes that the optimal way for the government to increase net costs by \$1 is to increase the WMVPF maximizing policy by $\frac{1}{\frac{dNC}{dp_i}}$ and the optimal way for the government to increase net revenue by \$1 is to decrease the WMVPF minimizing policy by $\frac{1}{\frac{dNC}{dp_i}}$. Conceptually, the WMVPF captures how much you can increase welfare by when changing a policy to have a net cost of \$1 to the government. Thus, the optimal manner for the government to increase net costs by \$1 is via spending on the WMVPF maximizing policy; similarly, the optimal way for the government to increase net revenue by \$1 is via the WMVPF minimizing policy. The key takeaway from Proposition 1 is that the optimal budget neutral reform involves increasing net revenue via the lowest WMVPF policy and increasing net costs via the highest WMVPF policy; this Proposition therefore goes against a conjecture of García and Heckman (2022a) who suggest that the MVPF is not a useful statistic for analyzing budget neutral policy reforms.[9]

2.2.1 Ranking Policies that Pay for Themselves

Hendren and Sprung-Keyser (2020) set the MVPF of policies that "pay for themselves" (i.e., policies for which $\frac{dW}{dp_i} > 0$ and $\frac{dNC}{dp_i} < 0$) to be ∞. Our definition of the WMVPF, Definition 1, does *not* set the WMVPF for such policies to be ∞, instead allowing the WMVPFs for these policies to be negative.

To see why allowing policies to have negative WMVPFs is useful for analyzing optimal reforms, note that Proposition 1 states that the optimal budget neutral reform raises \$1 in net revenue from the policy with the lowest WMVPF and spends this \$1 on the policy with the highest WMVPF. If there is a policy with a negative WMVPF, then the government should always prefer to raise net revenue via this policy than any policy with a positive WMVPF. This is because negative WMVPF policies can be used to raise revenue while simultaneously increasing welfare; hence, raising revenue from a negative WMVPF policy is always preferable to raising revenue from a positive WMVPF policy (which raises revenue but decreases welfare).[10] Finally, allowing WMVPFs to be negative has one more advantage: if there are multiple policies with

[9]Note that because **p** includes parameters governing tax rates, the optimal reform may reduce the size of government via cutting spending on a policy to raise revenue while also "spending" that revenue on tax reductions.

[10]Note, this logic also applies to policies with $\frac{dW}{dp_i} < 0$ and $\frac{dNC}{dp_i} > 0$ in which case the government can cut net spending on this program while simultaneously increasing welfare.

negative WMVPFs, Proposition 1 establishes how to *rank* these policies from the perspective of raising net revenue: the government should always raise a small amount of net revenue by increasing the policy with the *most negative* WMVPF.[11] [12]

Consider an example. Suppose that the government has four potential policies, with $\frac{dW}{dp_i} = \{1, 1, 1, 2\}$ for $i = \{1, 2, 3, 4\}$ and that the mechanical cost of each policy equals 1 so that $\frac{dMC}{dp_i} = 1 \ \forall i$. And suppose that the net costs of these policies are: $\frac{dNC}{dp_i} = \{2, 1, -1, -1\}$ for $i = \{1, 2, 3, 4\}$ so that the WMVPF's are equal to $\{\frac{1}{2}, 1, -1, -2\}$, respectively. Proposition 1 tells us that the optimal local budget neutral reform consists of raising net revenue from policy 4 (which requires increasing *mechanical* spending on policy 4 which in turn raises net revenue and also increases welfare) and then increasing net costs on policy 2, which has the highest WMVPF. Thus, we interpret policies with negative WMVPF's as the best possible way to raise net revenue as opposed to the best way to spend money (because the government does not spend money *on net* by increasing mechanical funding for these policies).

2.2.2 Multiple Time Periods

While the framework developed so far does not include an explicit discussion of time, our formulation can nonetheless be construed to include time periods. In Problem 1, for example, the set of policies $\mathbf{p} = (p_1, p_2, ..., p_N)$ do not necessarily need to take place in the same time period, the welfare function might be a discounted sum of period welfare functions over time so that $W(\mathbf{p}) = \sum_{t=0}^{T} \beta^t W_t(\mathbf{p})$ where β is a discount rate, and the cost function may be a discounted sum of period cost functions over time so that $C(\mathbf{p}, \mathbf{y}(\mathbf{n}; \mathbf{p})) = \sum_{t=0}^{T} \left(\frac{1}{1+r}\right)^t C_t(\mathbf{p}, \mathbf{y}(\mathbf{n}; \mathbf{p}))$ where r is an interest rate. Our framework therefore allows for the optimal reform to, for example, involve reducing the government deficit by decreasing spending on a current policy while increasing spending on a future policy.

2.3 Mechanical Cost Constraints and the NSB

There is an important implicit assumption that we have made when writing down Problem 1 (and Problem 1 Global): the policymaker can choose how to spend fiscal externalities that arise from their decisions. To see why Problem 1 implicitly assumes this, consider the following example of a policymaker who wants to cut net spending by \$1 on the WMVPF minimizing policy and increase net spending by \$1 on the WMVPF maximizing policy in accordance with Proposition 1. Suppose the WMVPF minimizing policy is a non-distortionary lump-sum trans-

[11] This is not to say that setting the MVPF to be ∞ whenever $\frac{dW}{dp_i} > 0$ and $\frac{dNC}{dp_i} < 0$ as in Hendren and Sprung-Keyser (2020) is wrong: defining the MVPF of these policies as ∞ is useful conceptually because it is always welfare improving to undertake these policies given they allow the government to increase welfare (assuming $\bar{\eta} > 0$) and also raise net revenue. On the other hand, allowing the MVPF to be negative enables us to compare these policies to one another to determine the optimal way to raise a small amount of net revenue.

[12] We rule out the knife-edge situation in which there is a policy with $\frac{dNC}{dp_i} = 0$.

fer so that the policymaker raises \$1 in net revenue by reducing mechanical spending on the lump-sum transfer by \$1. But suppose the WMVPF maximizing policy accrues \$0.50 (in present value) of fiscal externalities for each \$1 mechanically spent on this policy. And suppose these fiscal externalities occur a number of years in the future so that to increase net spending by \$1, the government must spend \$2 upfront on this policy and recoup \$1 (in present value) in fiscal externalities down the road. Hence, to implement the optimal local budget neutral reform, the government needs to use the \$1 they raised from the non-distortionary lump-sum transfer and also borrow \$1 against future fiscal externalities to spend \$2 upfront on the WMVPF maximizing policy, and then recoup \$1 in fiscal externalities later on to pay back their loan. Crucially, the structure of Problem 1 allows the policymaker to choose how to spend the \$1 of fiscal externalities that accrue from his/her chosen policy reform (by borrowing against future fiscal externalities and spending this \$1 of fiscal externalities on increased upfront mechanical spending on the WMVPF maximizing policy). Hence, Problem 1 is only appropriate for policymakers who can choose how to spend fiscal externalities. In practice, Problem 1 is therefore likely relevant for many important policy reform problems, such as a federal legislature deciding on tax policy or on generosity of entitlement programs. Federal legislatures can (and frequently do) borrow against future revenue streams from their policy decisions.

However, in practice other sorts of policymakers may not be able to choose how to spend fiscal externalities due to bureaucratic/political-economy constraints. There are two main reasons why this may occur: (1) fiscal externalities accrue in the future and the policymaker is unable to borrow against future revenues and/or (2) fiscal externalities accrue to agencies/departments outside of the control of the policymaker in question (as noted by García and Heckman (2022b)).[13] For example, an education initiative funded by the Department of Education may increase tax revenue many years in the future due to a more educated populace but the Department of Education may have no control over how this revenue is spent (and therefore may not be able to borrow against this future income tax revenue). Or a policy that subsidizes electric cars might be made by an environmental agency that does not directly bear the burden of reduced gas tax revenues which accrue to another part of the government (and hence cannot choose how to pay for those reduced revenues). Alternatively, an international development organization (IDO) or non-governmental organization (NGO) may spend funds on some policy for which all fiscal externalities accrue to governments, meaning the IDO/NGO has no choice how fiscal externalities are spent.[14]

[13]Problem 1 implicitly assumes that there are no fiscal externalities accruing to outside agencies (or, equivalently, if there are fiscal externalities accruing to other policymakers, then the decisions of those policymakers have no impact on the budget or welfare function of the policymaker solving Problem 1).

[14]Implicitly, these examples assume that *none* of the fiscal externalities accrue to the policymaker in question; we relax this assumption in Section 3.1.

Hence, we now consider a problem in which none of the fiscal externalities arising from policy decisions are controlled by the policymaker in question. Specifically, we consider a policymaker seeking to optimize a welfare function $W(\cdot)$ given a budget of $\$E$ but we suppose that any fiscal externalities accrue to either a different agency and/or accrue in the future (and the policymaker cannot borrow against them); one might interpret this sort of model as an agency with an annual use-it-or-lose-it budget constraint, which are fairly common in real-world governments (Liebman and Mahoney, 2017).[15] We suppose that any fiscal externalities arising from policy choices are instead spent on a "numeraire policy" p_{N+1}. If fiscal externalities accrue to another agency then p_{N+1} could represent spending on this other agency's chosen use of funds or if fiscal externalities accrue in the future and the policymaker faces a borrowing constraint, then p_{N+1} could represent spending on some policy in the future.[16] Because fiscal externalities do not impact the budget of the policymaker in question, budget neutral reforms (from the perspective of the policymaker in question) must leave *mechanical* costs (rather than net costs) unchanged.

Let us assume that the policymaker has N policies $\{p_1, p_2, ..., p_N\} = \mathbf{p}$ to choose from and that the policy p_{N+1} that is used to pay for any fiscal externalities has no impact on the budget of the policymaker in question. Thus, the budget neutrality constraint for the policymaker in question is that mechanical costs of the reform from $\tilde{\mathbf{p}}$ to $\mathbf{p}$ are zero: $MC(\mathbf{p}) \equiv C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1})) - C(\tilde{\mathbf{p}}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1})) = 0$. Additionally, the policymaker needs to incorporate the fact that all fiscal externalities that accrue due to changes in mechanical spending on $p_1, p_2, ..., p_N$ are paid for via policy p_{N+1}: hence, policy p_{N+1} changes to ensure that the net cost of the reform from $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$ to $(\mathbf{p}, p_{N+1})$ is 0: $NC(\mathbf{p}, p_{N+1}) \equiv C(\mathbf{p}, p_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, p_{N+1})) - C(\tilde{\mathbf{p}}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1})) = 0$.[17] This leads to the following optimal global reform problem:

Problem 2 Global. *Given a status quo policy* $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$*:*

$$\max_{\mathbf{p}} \; W(\mathbf{p}, p_{N+1}) - W(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$$

$$\text{s.t.} \; \underbrace{C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1})) - C(\tilde{\mathbf{p}}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1}))}_{\text{Mechanical Cost, } MC(\mathbf{p})} = 0$$

$$\underbrace{C(\mathbf{p}, p_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, p_{N+1})) - C(\tilde{\mathbf{p}}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1}))}_{\text{Net Cost, } NC(\mathbf{p}, p_{N+1})} = 0$$

[15] An annual use-it-or-lose-it budget constraint is an example of a mechanical spending constraint as long as all mechanical spending occurs upfront and all fiscal externalities accrue in the future.

[16] p_{N+1} could also represent increased spending on a *vector* of policies in the future or from outside agencies.

[17] To see why this second constraint ensures that fiscal externalities are paid for via policy p_{N+1}, subtract the first constraint from the second constraint then add and subtract $C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, \tilde{p}_{N+1}))$ to yield:

$$\underbrace{C(\mathbf{p}, p_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, p_{N+1})) - C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, \tilde{p}_{N+1}))}_{\text{Cost Changes Due to } p_{N+1}} + \underbrace{[C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, \tilde{p}_{N+1})) - C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1}))]}_{\text{-Fiscal Externalities, } -FE(\mathbf{p})} = 0$$

As with Problem 1 Global, solving Problem 2 Global requires empirical estimates of how $W(\cdot)$ and $C(\cdot)$ vary with $\mathbf{p}$ across policies that have never been observed in reality. Instead of making strong assumptions about how $W(\cdot)$ and $C(\cdot)$ vary far from the status quo, we again turn to an optimal *local* reform problem that respects the constraints of Problem 2 Global.

To ensure the mechanical spending constraint continues to be satisfied while only considering reforms that are close to $\tilde{\mathbf{p}}$, the policymaker seeks to find the optimal reform that changes policies to increase *mechanical* costs by a small amount and also increase *mechanical* revenue by the same small amount, with any resulting fiscal externalities being funded via (or spent on) p_{N+1}. This leads to the following optimal reform problem corresponding to Problem 2 Global recognizing that the $\sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i)$ is the first order Taylor series expansion of the mechanical cost constraint around $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$ and $\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i)$ is the first order Taylor series expansion of the net cost constraint in Problem 2 Global (note that all of the derivatives in Problem 2 are evaluated at $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$):[18]

Problem 2. *Given a status quo policy* $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$*:*

$$\max_{\mathbf{p}} \sum_{i=1}^{N+1} \frac{dW}{dp_i}(p_i - \tilde{p}_i)$$

$$\text{s.t.} \sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dMC}{dp_i}(p_i - \tilde{p}_i) > 0\right] = 1$$

$$\sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dMC}{dp_i}(p_i - \tilde{p}_i) < 0\right] = -1$$

$$\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0$$

Problem 2 has two key differences compared to Problem 1. First, instead of mandating that the government spend (and raise) \$1 *inclusive* of fiscal externalities, Problem 2 restricts the government to spend (and raise) \$1 *exclusive* of fiscal externalities; hence, Problem 2 seeks to find the best way for a policymaker to cut mechanical spending by \$1 on some policy and then increase mechanical spending by \$1 on another policy. Second, because this mechanical spending (and mechanical revenue raising) of \$1 nonetheless generates fiscal externalities, Problem 2 also specifies that the government spend any net fiscal externalities on the numeraire policy p_{N+1}; to see this, note that because $\sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i) = 0$, we can rewrite the constraint that $\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0$ as $\sum_{i=1}^{N} \left(\frac{dNC}{dp_i} - \frac{dMC}{dp_i}\right)(p_i - \tilde{p}_i) + \frac{dNC}{dp_{N+1}}(p_{N+1} - \tilde{p}_{N+1}) = \sum_{i=1}^{N} -\frac{dFE}{dp_i}(p_i - \tilde{p}_i) + \frac{dNC}{dp_{N+1}}(p_{N+1} - \tilde{p}_{N+1}) = 0$.[19] Towards solving Problem 2, let us define the welfare-weighted

[18] Problem 2 is an optimal local reform problem corresponding to Problem 2 Global insofar as the welfare gain from Problem 2 equals 0 if we start from any status quo $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$ which is a local maximum of Problem 2 Global (if not then Problem 2 would identify a reform that improves welfare and keeps total mechanical costs constant and pays for fiscal externalities via p_{N+1}).

[19] Problem 2 can also be construed to allow for multiple time periods. If welfare effects and/or costs occur in

net social benefit (WNSB):

Definition 2 (WNSB). *The welfare-weighted net social benefit of policy i is the welfare gain from increasing mechanical spending on policy i by \$1 and then financing fiscal externalities via policy p_{N+1}:*

$$WNSB_i^{(p_{N+1})} = \left(\frac{dW}{dp_i} + \frac{\frac{dW}{dp_{N+1}}}{\frac{dNC}{dp_{N+1}}} \frac{dFE}{dp_i} \right) \Big/ \frac{dMC}{dp_i}$$

There are four (somewhat) minor differences in how we define the WNSB compared to the literature. First, given that $\frac{dW}{dp_i} = \int_{\mathbf{N}} \phi(\mathbf{n}) \frac{dU(\mathbf{y}(\mathbf{n};\mathbf{p});\mathbf{p},\mathbf{n})}{dp_i} f(\mathbf{n}) d\mathbf{n}$, our expression includes welfare weights whereas the NSB typically does not include welfare weights.[20] Second, our definition of the WNSB depends on the numeraire policy p_{N+1} used to close the budget constraint; the literature on the net social benefit typically assumes that p_{N+1} represents linear income taxation but in principle this can be any policy (Hendren and Sprung-Keyser, 2022). Third, our definition of the $WNSB_i^{(p_{N+1})}$ measures the welfare gain from increasing mechanical spending by \$1 on policy i and then financing fiscal externalities via policy p_{N+1}; the literature often defines the NSB assuming that both fiscal externalities *and* mechanical costs are financed via policy p_{N+1}. This just changes the value of the WNSB by a constant equal to $\frac{\frac{dW}{dp_{N+1}}}{\frac{dNC}{dp_{N+1}}}$ and therefore does not impact the ranking of policies in Proposition 2. Fourth, we divide the entire expression by $\frac{dMC}{dp_i}$ because we allow p_i to represent any arbitrary policy parameter rather than mandating that p_i simply represents spending on policy i (in which case $\frac{dMC}{dp_i} = 1$). Also, it is interesting to note that in Definition 2, the term $\frac{\frac{dW}{dp_{N+1}}}{\frac{dNC}{dp_{N+1}}}$ is the WMVPF of policy p_{N+1}. When p_{N+1} represents linear income taxation, this term is often expressed as $\frac{dW}{dp_{N+1}}(1+\psi)$ where $1+\psi$ is called the "marginal cost of public funds" and represents the cost of raising \$1 via linear income taxation (Ballard and Fullerton, 1992). We can now state the solution to Problem 2:

Proposition 2. *The solution to Problem 2 is:*

$$(p_i - \tilde{p}_i) = \begin{cases} \frac{1}{\frac{dMC}{dp_i}} & i = \arg\max \left\{ WNSB_i^{(p_{N+1})} \right\} \\ \frac{-1}{\frac{dMC}{dp_i}} & i = \arg\min \left\{ WNSB_i^{(p_{N+1})} \right\} \\ 0 & \textit{otherwise} \end{cases}$$

Proof. See Appendix A.2. □

The takeaway from Proposition 2 is that the WNSB, rather than the WMVPF, becomes the sufficient statistic governing optimal policy reforms when the government cannot control how fiscal externalities are spent and therefore faces constraints on *mechanical* spending and

future time periods then the welfare function W and net cost C can be interpreted as representing discounted sums of period welfare/cost functions as in Section 2.2.2.

[20]Footnote 12 of Hendren and Sprung-Keyser (2022) notes that using the NSB for welfare analysis requires welfare weights.

mechanical revenue raised.[21] This is one of the primary takeaways of the paper: for the purposes of constructing optimal policy reforms, a policymaker should use the WMVPF if they can control how fiscal externalities are spent but should use the WNSB if bureaucratic/political-economy constraints restrict how fiscal externalities will be spent. Next, we briefly discuss problems where policymakers are able to spend a fraction of the fiscal externalities that arise from their reform decisions, larger reforms and non-budget-neutral reforms, previous work on optimal tax reforms, and uncertainty.

3 Theoretical Extensions and Discussion

3.1 Optimal Reforms When Policymakers Can Spend A Fraction of FEs

So far we have assumed that either (1) policymakers can spend all fiscal externalities optimally (Section 2.2) or (2) that policymakers cannot choose how to spend any fiscal externalities (Section 2.3); we now show how to solve optimal local reform problems when policymakers can choose how to spend a *fraction* of fiscal externalities. This sort of problem may be relevant, for example, when a federal government implements a policy reform and some of the fiscal externalities accrue to state governments who then choose how to spend their share of the fiscal externalities (so that spending of some share of the fiscal externalities may be out of the control of the federal government). Alternatively, this sort of problem may be relevant for state governments implementing a policy reform for which some of the fiscal externalities accrue to the federal government or other state governments.[22]

We consider a policymaker who chooses some set of policies $\mathbf{p} = \{p_1, p_2, ..., p_N\}$ but that only a fraction of fiscal externalities resulting from these policies accrue "internally" to the policymaker in question. Hence, the "net internal cost" of a marginal policy change from $\tilde{p}_i$ to p_i equals:[23]

$$\underbrace{\frac{dNC_I}{dp_i}(p_i - \tilde{p}_i)}_{\text{Net Internal Cost}} \equiv \underbrace{\frac{dMC}{dp_i}(p_i - \tilde{p}_i)}_{\text{Mechanical Cost}} - \underbrace{\sigma_i \frac{dFE}{dp_i}(p_i - \tilde{p}_i)}_{\text{Share of Fiscal Externalities}} \tag{3}$$

where all derivatives are evaluated at the status quo $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$ and σ_i is the share of fiscal externalities from policy p_i accruing internally to the policymaker. To ensure that the internal spending constraint continues to be satisfied while only considering reforms that are close to $\tilde{p}$, the policymaker seeks to find the optimal reform that changes policies to increase internal costs by a small amount and also increase internal revenue by the same small amount, with any resulting fiscal externalities that accrue "externally" to the agency in question being funded via

[21] Note that if no policies have fiscal externalities (i.e., $\frac{dNC}{dp_i} = \frac{dMC}{dp_i}\ \forall\ i$) then $WNSB_i = WMVPF_i\ \forall i$.

[22] Agrawal et al. (2022) discuss how state and local governments make policy decisions when some fiscal externalities accrue to the federal government or other state governments: Proposition 3 below condenses to the "local MVPF" criterion discussed in Agrawal et al. (2022) as long as the policymaker doesn't care about how fiscal externalities that accrue to outside agencies are spent so that $\frac{dW}{dp_{N+1}} = 0$.

[23] See Appendix B.1 for the global reform problem analogous to Problems 1 Global and 2 Global.

(or spent on) a single numeraire policy p_{N+1}:

Problem 3. *Given a status quo policy* $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$*:*

$$\begin{aligned}\max_{\mathbf{p}} \quad & \sum_{i=1}^{N+1} \frac{dW}{dp_i}(p_i - \tilde{p}_i) \\ \text{s.t.} \quad & \sum_{i=1}^{N} \frac{dNC_I}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC_I}{dp_i}(p_i - \tilde{p}_i) > 0\right] = 1 \\ & \sum_{i=1}^{N} \frac{dNC_I}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC_I}{dp_i}(p_i - \tilde{p}_i) < 0\right] = -1 \\ & \sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0\end{aligned}$$

Problem 3 changes some policies to raise internal revenue by \$1 and changes other policies to raise internal costs by \$1 with any fiscal externalities that accrue externally being spent on policy p_{N+1}. To see this note we can subtract $\sum_{i=1}^{N} \frac{dNC_I}{dp_i}(p_i - \tilde{p}_i) = 0$ from the constraint that $\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = \sum_{i=1}^{N+1}\left(\frac{dMC}{dp_i} - \frac{dFE}{dp_i}\right)(p_i - \tilde{p}_i) = 0$ to yield $\sum_{i=1}^{N}\left(\frac{dNC}{dp_i} - \frac{dNC_I}{dp_i}\right)(p_i - \tilde{p}_i) + \frac{dNC}{dp_{N+1}}(p_{N+1} - \tilde{p}_{N+1}) = \sum_{i=1}^{N}\left(-(1-\sigma_i)\frac{dFE}{dp_i}\right)(p_i - \tilde{p}_i) + \frac{dNC}{dp_{N+1}}(p_{N+1} - \tilde{p}_{N+1}) = 0$. Note that when all fiscal externalities accrue to the agency in question ($\sigma = 1$), Problem 3 collapses to Problem 1 (as the final constraint is then redundant). Conversely, if the only internal costs are mechanical costs (i.e., $\sigma = 0$), Problem 3 collapses to Problem 2. The solution to Problem 3 is given by Proposition 3:

Proposition 3. *The solution to Problem 3 is to set:*

$$(p_i - \tilde{p}_i) = \begin{cases} \frac{1}{\frac{dNC_I}{dp_i}} & i = \arg\max\left\{\underbrace{\frac{\frac{dW}{dp_i}}{\frac{dNC_I}{dp_i}}}_{\text{Internal WMVPF}} + \underbrace{\frac{\frac{dW}{dp_{N+1}}}{\frac{dNC}{dp_{N+1}}}\frac{(1-\sigma_i)\frac{dFE}{dp_i}}{\frac{dNC_I}{dp_i}}}_{\text{External Welfare Cost}}\right\} \\ \frac{-1}{\frac{dNC_I}{dp_i}} & i = \arg\min\left\{\underbrace{\frac{\frac{dW}{dp_i}}{\frac{dNC_I}{dp_i}}}_{\text{Internal WMVPF}} + \underbrace{\frac{\frac{dW}{dp_{N+1}}}{\frac{dNC}{dp_{N+1}}}\frac{(1-\sigma_i)\frac{dFE}{dp_i}}{\frac{dNC_I}{dp_i}}}_{\text{External Welfare Cost}}\right\} \\ 0 & \text{otherwise}\end{cases}$$

Proof. See Appendix A.3. □

The takeaway from Proposition 3 is that when policymakers can only choose how to spend a fraction of fiscal externalities that accrue from decisions, the optimal policy reform is characterized by a "augmented internal WMVPF" which equals an "internal WMVPF" plus a correction term that measures the welfare gain from external costs being funded by changing some numeraire policy p_{N+1} that is not controlled by the agency in question.[24]

3.2 Locality, Linearity, and Non-budget-neutrality

A key feature of the optimal reform problems that we study in this paper is that they are *local* to the status quo policy in the sense that they do not account for curvature in either welfare or revenue as we change the policy from $\tilde{\mathbf{p}}$ to $\mathbf{p}$. There are several points to discuss. First, while we have discussed reforms that change spending and revenue by \$1, the results immediately apply to larger reforms as well provided that both welfare and cost are approximately linear in the neighborhood around the status quo $\tilde{\mathbf{p}}$ which satisfies the problem constraints. For instance, we can change the net cost constraint in Problem 1 to be: $\sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) > 0\right] = a$ for any a such that welfare and revenue are both (approximately) linear between all $\mathbf{p}$ and $\tilde{\mathbf{p}}$ satisfying this constraint. In this case, identical logic as in the proof to Proposition 1 can be used to show that the solution to Problem 1 simply changes so that we spend $\frac{a}{\frac{dNC}{dp_i}}$ on the WMVPF maximizing instead of $\frac{1}{\frac{dNC}{dp_i}}$. Similarly, we can change the net revenue constraint in Problem 1 to equal $-b$ instead of -1 and the solution changes so that we spend $\frac{-b}{\frac{dNC}{dp_i}}$ on the WMVPF minimizing instead of $\frac{-1}{\frac{dNC}{dp_i}}$. This allows us to explore non-budget-neutral reforms with the same framework developed above (by choosing $a \neq b$). Similarly, for Problem 2, we can change the mechanical cost constraints to have larger values (instead of 1 or -1) as long as both the welfare function and government revenue function are approximately linear within the neighborhood of $\tilde{\mathbf{p}}$ satisfying these constraints. And we can allow for reforms that are not budget neutral (in terms of mechanical costs) in Problem 2 by setting the mechanical cost or mechanical revenue constraints to different values a and $-b$ respectively; this changes the solution to spend $\frac{a}{\frac{dMC}{dp_i}}$ and $\frac{-b}{\frac{dMC}{dp_i}}$ on the WNSB maximizing/minimizing policies, respectively. Thus, our results can be applied to non-infinitesimal and/or non-budget-neutral reforms as long as the reforms are sufficiently small so that there is minimal curvature in the welfare function and government cost function.

3.3 Empirical Estimates with Uncertainty

One point that we have not touched on so far is the fact that empirically the components of both the WMVPF and the WNSB are typically estimated with error. Consider a scenario where the government chooses policies $i = \{1, 2, ..., N\}$ each of which is subject to uncertainty in both the benefits and net costs. The government is constrained so that in expectation they are only allowed to spend \$1 net on these policies; however, the government still has a budget constraint that must be satisfied regardless of the outcome of the uncertainty; the government

[24] Problem 3 can be easily extended to allow a fraction of mechanical costs to accrue externally as well; this would change the definition of net internal cost from Equation (3) and the share of mechanical costs accruing externally would appear in the numerator of the external welfare cost in Proposition 3.

satisfies the budget constraint with a numeraire policy p_{N+1} which has no uncertainty in the costs or benefits.[25] We show in Appendix B.2 that Proposition 1 still holds in this world with uncertainty if we define the WMVPF as:

$$WMVPF_i \equiv \frac{\mathbb{E}\left[\frac{dW}{dp_i}\right]}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]} = \bar{\eta} \underbrace{\frac{\mathbb{E}\left[\int_{\mathbf{N}} WTP_i(\mathbf{n}) f(\mathbf{n}) d\mathbf{n}\right]}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]}}_{MVPF} \tag{4}$$

This result is perhaps initially surprising given that the government must close any discrepancy from expected and realized fiscal externalities with a numeraire policy; however, the WMVPF remains the relevant sufficient statistic because, on average, this discrepancy between expected and realized fiscal externalities washes out and we assume the government is risk-neutral. Similarly, Proposition 2 continues to hold with uncertainty (see Appendix B.2) if we define:

$$WNSB_i^{(p_{N+1})} = \left(\mathbb{E}\left[\frac{dW}{dp_i}\right] + \frac{\frac{dW}{dp_{N+1}}}{\frac{dNC}{dp_{N+1}}} \mathbb{E}\left[\frac{dFE}{dp_i}\right]\right) \Big/ \frac{dMC}{dp_i} \tag{5}$$

The takeaway is that even in a world with uncertainty about the costs and benefits of a program, the WMVPF and WNSB (defined with expectations as in Equations (4) and (5)) remain the relevant sufficient statistics for optimal policy reforms with net revenue and mechanical revenue constraints, respectively.

3.4 Relation to Previous Optimal Tax Reform Literature

While, to the best of our knowledge, there is no prior work deriving a general framework for optimal local policy reforms, there has been work on optimal *tax* reforms. For instance, a number of papers (e.g., Diewert (1978) and Dixit (1979)) consider optimal tax reform problems as in Equation (6), where $\mathbf{p}$ represents different parameters of a tax system and $\|\cdot\|_2$ is the Euclidean norm:

$$\begin{aligned} \max_{\mathbf{p}} \quad & \sum_{i=1}^{N} \frac{dW}{dp_i}(p_i - \tilde{p}_i) \\ \text{s.t.} \quad & \sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0 \text{ and } \|\mathbf{p} - \tilde{\mathbf{p}}\|_2 = 1 \end{aligned} \tag{6}$$

The solution to Equation (6), which follows from the Cauchy-Schwarz inequality, is to use gradient ascent and set $(\mathbf{p} - \tilde{\mathbf{p}}) = \frac{\nabla_{\mathbf{p}} W(\tilde{\mathbf{p}}) - \lambda \nabla_{\mathbf{p}} NC(\tilde{\mathbf{p}})}{\|\nabla_{\mathbf{p}} W(\tilde{\mathbf{p}}) - \lambda \nabla_{\mathbf{p}} NC(\tilde{\mathbf{p}})\|_2}$ where λ is a Lagrange multiplier chosen to satisfy the first constraint in Equation (6). However, there are a number of drawbacks to considering Equation (6) relative to Problem 1. First, in Problem 1 (as well as Problems 2 and

[25] We can relax the assumption of certainty over the costs and benefits of p_{N+1} if we suppose any divergence from expectation in costs of p_{N+1} are paid for via another policy p_{N+2} and any divergence from expectation in costs of p_{N+2} are paid for via p_{N+3}, and so on *ad infinitum*. In this case, Proposition 1 still holds with the WMVPF defined as in Equation (4) and Proposition 2 continues to hold as long as we replace $\frac{dW}{dp_{N+1}} \Big/ \frac{dNC}{dp_{N+1}}$ with $\mathbb{E}\left[\frac{dW}{dp_{N+1}}\right] / \mathbb{E}\left[\frac{dNC}{dp_{N+1}}\right]$ in Equation (5). We omit a proof but are happy to provide one upon request.

3), the extent to which any given policy is allowed to change is constrained by its impact on costs; hence, all policies are constrained in the same units. In contrast, Equation (6) constrains the policy reform so that the sum of squared deviations $\|\mathbf{p} - \tilde{\mathbf{p}}\|_2 = \sum_{i=1}^{N}(p_i - \tilde{p}_i)^2 = 1$. While this constraint may be somewhat sensible for tax reforms where all policies p_i are expressed as tax rates, when policies are expressed in different units, then the constraint in Equation (6) has varying levels of restrictiveness depending on the units of $\mathbf{p}$. For example, suppose we have three policies, two of which are measured in dollars and one of which is a tax rate: the optimal reform from Equation (6) will then vary depending on whether we measure the tax rate between 0 and 1 or between 0 and 100. This "units problem" also implies that the solution to Equation (6) depends in a complicated way on the set of policies available to the policymaker (via the Lagrange multiplier) and therefore there is no way to derive a metric to rank and compare policies independent of the set of policies available to the policymaker.[26] [27]

The aforementioned units problem can, however, be rectified if we change the L^2 constraint in Equation (6) to instead be in budgetary units reading $\|\nabla_{\mathbf{p}} NC \circ (\mathbf{p} - \tilde{\mathbf{p}})\|_2 = 1$ where $\circ$ denotes element-wise vector multiplication. Interestingly, after making this substitution, we show in Appendix B.4 that the WMVPF is again the sufficient statistic needed to determine the optimal local policy reform; similarly, the WNSB is the sufficient statistic for a problem with an L^2 constraint on mechanical spending and a constraint that fiscal externalities be spent on a numeraire policy p_{N+1}.[28] [29]

Ultimately, perhaps the best argument for interest in Problem 1 (as well as Problems 2 and 3) over problems like Equation (6) is that Problem 1 (as well as Problems 2 and 3) can be applied to answer real-world policy relevant questions such as "Would it be better to increase spending on policy p_1 and reduce spending on policy p_2 or increase spending on policy p_3 and reduce spending on policy p_4?" and "Should we spend a windfall on policy p_1, p_2, or p_3?" (where, as shown by Propositions 1, 2 and 3, the answer depends on the policymaker's control over fiscal externalities). On the other hand, Equation (6) can only be used to answer questions which are presumably less relevant for policymakers; for instance, if p_1, p_2, and p_3 represent spending on three different programs, then Equation (6) will answer the question "How should

[26] See Appendix B.3 for a numerical example of how the presence of a third policy can change which of two other policies is increased/decreased by more when solving Equation (6).

[27] In contrast, the solution to Problem 1 (as well as Problems 2 and 3) depends on a single welfare metric for each policy that can be used to compare and rank policies and is independent of the set of other policies available to the policymaker (insofar as if policy i has a higher WMVPF than policy j then a policymaker that can control FEs will *always* prefer to spend on policy i over policy j regardless of which other policies are available).

[28] Naturally, the *direction* of the reform is different with an L^2 constraint even though the direction is still only a function of the WMVPF or WNSB depending on whether the policymaker can control fiscal externalities. The "augmented internal WMVPF" can also be shown to still be the sufficient statistic if a policymaker can only control a fraction of fiscal externalities and faces an L^2 constraint on internal spending as in Section 3.1.

[29] Note Problems 1, 2, and 3 effectively contain L^1 constraints; for example, the two constraints in Problem 1 can be rewritten as $\sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0$ and $\|\nabla_{\mathbf{p}} NC \circ (\mathbf{p} - \tilde{\mathbf{p}})\|_1 = 2$.

we change spending on policies p_1, p_2, and p_3 in a budget neutral manner such that the sum of squared spending increase equals \$1?".

4 Welfare Gains from Using the WMVPF over the WNSB

Proposition 2 yields that the optimal way for an agency to reform a set of N policies when faced with a mechanical spending constraint and no control over how fiscal externalities are spent is to reduce mechanical spending on the WNSB minimizing policy and to increase mechanical spending on the WNSB maximizing policy. On the other hand, if we asked a policymaker who *could* control how fiscal externalities were spent how to optimally reform the same given set of N policies, then this policymaker would increase net spending on the WMVPF maximizing policy and reduce net spending on the WMVPF minimizing policy by Proposition 1. Next, we explore the welfare consequences of allowing policymakers to choose how to spend fiscal externalities by relaxing mechanical spending constraints. Specifically, we will quantify the welfare gains from conducting a budget neutral reform without a mechanical budget constraint (and therefore spending fiscal externalities optimally) as in Problem 1 relative to Problem 2.

Consider a policymaker who cannot choose how to spend fiscal externalities and therefore solves Problem 2. Let us denote the solution as $\{p_i^{WNSB}\}$ (where $\{p_i^{WNSB}\}$ is the solution to Problem 4 from Proposition 2). The proposed reform entails increasing mechanical costs for some policies by $MC^{WNSB} = 1$ and decreasing mechanical revenues from other policies by $MR^{WNSB} = 1$. The proposed reform also entails increasing net costs on some policies by an amount NC^{WNSB} and increasing net revenue from other policies by NR^{WNSB} (where $NC^{WNSB} = NR^{WNSB}$ by budget neutrality):[30]

$$\begin{aligned} NC^{WNSB} &\equiv \sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i^{WNSB} - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i^{WNSB} - \tilde{p}_i) > 0\right] \\ -NR^{WNSB} &\equiv \sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i^{WNSB} - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i^{WNSB} - \tilde{p}_i) < 0\right] \end{aligned} \tag{7}$$

Now suppose that, hypothetically, a policymaker who could choose how to spend fiscal externalities (e.g., a federal legislature) was tasked with reforming the same N policies subject only to net cost constraints as in Problem 1. Suppose this second policymaker now explores how to construct the optimal policy reform that has the same net cost, NC^{WNSB}, and same amount of net revenue raised, NR^{WNSB}, as $\{p_i^{WNSB}\}$ yet is not constrained to spend fiscal externalities on policy p_{N+1}. Hence, this second policymaker seeks to solve Problem 4:[31]

[30] Note the net cost (or net revenue) of the numeraire policy p_{N+1} is included in either NC^{WNSB} or NR^{WNSB}.

[31] Problem 4 looks for a reform that increases net costs and raises net revenue by the same amount as $\{p_i^{WNSB}\}$ so that we can make an apples-to-apples comparison of the welfare impacts. Problem 4 also assumes that policy p_{N+1} is not under control of the second policymaker and hence cannot be altered.

Problem 4. *Given a status quo policy* $\tilde{\mathbf{p}}$*:*

$$\max_{\mathbf{p}} \sum_{i=1}^{N} \frac{dW}{dp_i}(p_i - \tilde{p}_i)$$

$$\text{s.t.} \sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) > 0\right] = NC^{WNSB}$$

$$\sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) < 0\right] = -NR^{WNSB}$$

It follows from Proposition 1 and the discussion in Section 3.2 that the solution to Problem 4 is to increase net spending by NC^{WNSB} on the WMVPF maximizing policy and raise NR^{WNSB} in net revenue from the WMVPF minimizing policy. Denote the optimal policy reform from Problem 4 as $\{p_i^{WMVPF}\}$; how can we measure the gain from solving Problem 4 relative to solving Problem 2 that we began with? We propose the following:

Definition 3. *The compensating variation, CV, between $W\left(\{p_i^{WNSB}\}\right)$ and $W\left(\{p_i^{WMVPF}\}\right)$ is equal to the required amount of money spent on the WNSB maximizing policy (with any fiscal externalities being spent on policy p_{N+1}) to make up the difference in welfare:*

$$W\left(\{p_i^{WMVPF}\}\right) = W\left(\{p_i^{WNSB}\} + CV \arg\max\{WNSB_i^{(p_{N+1})}\}\right)$$

where, abusing notation, $\{p_i^{WNSB}\} + CV \arg\max\{WNSB_i^{(p_{N+1})}\}$ represents the policy that spends \$$CV$ more mechanically on the WNSB maximizing policy than $\{p_i^{WNSB}\}$.

Next, Proposition 4 proves solving Problem 4 increases welfare by more than solving Problem 2 and gives an expression for the compensating variation between $W\left(\{p_i^{WNSB}\}\right)$ and $W\left(\{p_i^{WMVPF}\}\right)$:

Proposition 4. *Suppose that p_{N+1} is not the WMVPF maximizing or minimizing policy. Then $W\left(\{p_i^{WMVPF}\}\right) \geq W\left(\{p_i^{WNSB}\}\right)$ and CV is given by the following expression:*

$$\frac{NC^{WNSB}\left[\max\{WMVPF_i\} - \min\{WMVPF_i\}\right] - MC^{WNSB}\left[\max\{WNSB_i^{(p_{N+1})}\} - \min\{WNSB_i^{(p_{N+1})}\}\right]}{\max\{WNSB_i^{(p_{N+1})}\}}$$

Proof. See Appendix A.4. □

The takeaway from Proposition 4 is that we can measure the welfare gain from removing mechanical constraints (and thereby allowing policymakers to spend fiscal externalities from their reform decisions optimally) in terms of the maximum and minimum WMVPFs and WNSBs.[32]

[32] One may question whether Problem 4 is equivalent to Problem 2 if we additionally allow the policymaker to pick the numeraire policy p_{N+1} in Problem 2; the answer turns out to be no. Picking p_{N+1} in Problem 2 allows the policymaker to choose how to spend the *net* fiscal externalities after increasing mechanical spending on some policies and decreasing mechanical spending on other policies, but it does not allow the policymaker to use fiscal externalities to finance spending or revenue raising on multiple policies. The solution to Problem 4 often requires the policymaker to use fiscal externalities to augment spending on both the WMVPF minimizing and WMVPF maximizing policies; hence, there are welfare gains from solving Problem 4 relative to solving Problem 2 as in Proposition 4 even if p_{N+1} is chosen optimally. We illustrate this with an example in Appendix B.5.

To understand the policy implications of Proposition 4, consider a federal legislature which allocates annual use-it-or-lose-it budgets to a variety of federal agencies which in turn make policy decisions. Suppose that all of the fiscal externalities from these policy decisions accrue back to the federal treasury which is controlled by the federal legislature. Policymakers in charge of the agencies therefore face mechanical spending constraints; hence, policymakers in charge of the various agencies should rely on the WNSB to make policy reforms. If instead the federal legislature (who does not face a constraint on spending in a given year and only faces a constraint on total net spending over time) were to instead reform policy directly, they would rely on the WMVPF and would be able to conduct better policy reforms than the agencies themselves by Proposition 4. Practically, time and information constraints prevent federal legislatures from making all policy reform decisions; hence, one may ask whether a federal legislature could instead change the incentives of agencies to base their reform decisions on the WMVPF rather than the WNSB, thereby reaping the welfare gains from Proposition 4:

Remark 1. *A federal legislature can incentivize agencies to solve Problem 4 rather than Problem 2 by implementing policy-specific subsidies s_i on each dollar of mechanical spending equal to $s_i = 1 - \frac{dNC/dp_i}{dMC/dp_i}$. In this case the marginal cost of increasing policy p_i equals $\frac{dMC}{dp_i}(1-s_i) = \frac{dNC}{dp_i}$. This subsidy scheme, which effectively rebates fiscal externalities back to agencies, is conceptually similar to the "marginal corrective transfer" of Agrawal et al. (2022).*

The economic intuition of Remark 1 is straight-forward: if policymakers cannot control how fiscal externalities are spent, it disincentivizes them from spending on policies with large fiscal externalities because these fiscal externalities are then spent on policy p_{N+1} which may have a lower WMVPF than other policies available to the policymaker. Rebating fiscal externalities back to agency policymakers (via policy-specific subsidies) properly incentivizes them to spend on policies which generate large fiscal externalities, thereby leading to welfare gains.

5 Examples

As a proof of concept, we now discuss three stylized examples showcasing how policymakers can use the theory developed herein to make policy reform decisions.

5.1 Federal Legislature

First, we explore a hypothetical reform problem for the federal legislature in the United States. In particular, suppose that the federal legislature is considering a budget neutral policy reform among six large government policies: an expanded earned income tax credit (EITC) for adults without dependents, increased benefit generosity for disability insurance (DI), a reduction in the top income tax rate, higher education tax deductions, increased Medicaid generosity for young children, and increased spending on housing vouchers. We take estimates of the net cost and WTP for these policies from Hendren and Sprung-Keyser (2020): the EITC estimates come from the Paycheck Plus experiment, the DI estimates are estimated from a kink in the benefit formula, the top income tax estimates are from the 2013 expiration of the Economic Growth and Tax Relief Reconciliation Act of 2001, the higher education tax deduction estimates are for single filers at the phase-in region of the Tuition and Fees Deduction, Medicaid expansions to young children estimates come from state level variation in spending, and housing voucher estimates come from the experimental arm of the Moving to Opportunity experiment. Estimates of net cost, WTP, the MVPF, and the WMVPF are presented in Table 1, with welfare weights equal to $\bar{\eta} = \bar{z}^{-\gamma}$ where $\bar{z}$ represents average income of policy beneficiaries and γ is a scalar which we assume equals 1 for purposes of illustration.[33]

Table 1: WMVPFs for Selected Federal Policies

Program	Average Income	Net Cost	WTP	MVPF	$\bar{\eta}$	WMVPF
EITC Expansion to Adults without Dependents (Paycheck Plus)	10534	1074	1070	1.00	0.95	0.95
Disability Insurance Benefit Generosity	26445	1.04	1.00	0.96	0.38	0.36
Top Income Tax Reductions	737931	0.86	1.00	1.16	0.01	0.02
Higher Education Tax Deductions	28411	-460	485	-1.06	0.35	-0.37
Medicaid Expansions to Young Children	48096	-491	3681	-7.49	0.21	-1.56
Housing Vouchers (Moving to Opportunity Experiment)	16331	-9215	69601	-7.55	0.61	-4.63

Note: This Table displays WMVPFs for six policies. Data on average income, net cost, and WTP are taken from Hendren and Sprung-Keyser (2020). $\bar{\eta} = 10000 \times \text{Average Income}^{-1}$ (where the factor of 10000 is for readability). The WMVPF is calculated from Definition 1.

Given the data in Table 1, how should the federal legislature optimally construct a policy reform that raises a small amount of net revenue from some policies and then increases net costs by a small amount on some other policies? Proposition 1 tells us that the optimal local policy reform is to raise revenue from the lowest WMVPF policy (housing vouchers) and then spend on the highest WMVPF policy (EITC expansion). There are a number of points to discuss. First, note that the highest WMVPF policy (EITC expansion) is not the same as the highest MVPF policy (top income taxation): the WTP relative to net cost of reducing top income taxes is higher than the EITC expansion but the welfare weight for top income tax reductions

[33] These welfare weights capture the government's relative preferences for giving money to households with different income levels; larger values of γ correspond to more redistributive preferences. Loosely, $\gamma = 1$ is consistent with a utilitarian policymaker and utility over consumption equal to $\log(c)$ (Chetty, 2006).

is very low because it impacts only those with high incomes. This reinforces the idea that it is crucial to consider welfare weights when using the MVPF for policy reforms. Second, while higher education tax deductions, Medicaid expansions to young children, and housing vouchers all have negative net costs (these policies "pay for themselves"), the housing vouchers policy has the most negative WMVPF and is therefore the best policy to use to raise net revenue by a small amount. Third, in order to raise net revenue via housing vouchers, the government must increase upfront mechanical spending on this program and then recoup this initial outlay and more in fiscal externalities in the future from increased income tax revenue: hence, the optimal policy reform is only budget neutral over a long time horizon and implicitly requires the government to borrow against future fiscal externalities to increase upfront spending on both the EITC expansion and housing vouchers. Fourth, it is important to recognize that we assume all of the policies can be expanded/contracted on the margin and that the reform is small enough so that the marginal costs and benefits are locally constant, per the discussion in Section 3.2.

Lastly, in conducting this exercise we have implicitly assumed that all fiscal externalities of these policy reforms accrue to the federal government. However, in reality this is not entirely accurate as some share of the fiscal externalities almost certainly accrue to state governments. For example, changes in labor income that arise from policy reforms impact not only federal tax revenue but also state tax revenue. To account for this fact, we can model the federal legislatures problem as Problem 3 and appeal to Proposition 3, using the "augmented internal WMVPF" to determine optimal policy reforms assuming that the federal legislature does not get to choose how to spend fiscal externalities that accrue to state governments.[34] We present "augmented internal WMVPFs" from Proposition 3 in Table 7 in Appendix D under the assumption that states fund any fiscal externalities via a lump sum transfer (i.e., the numeraire policy p_{N+1} is a lump sum transfer). The key finding from this exercise is that accounting for fiscal externalities accruing to states has only minimal impacts on the ranking of policies from an optimal reform perspective because only a small share of fiscal externalities accrue to states for most of the given policies. The only policy which is substantially impacted by accounting for state level fiscal externalities is the higher education tax deduction because many of the costs of increased education spending that arise from higher education tax deductions are borne by state governments that fund state universities - this turns out to be a relatively large share of the costs of this tax credit and therefore leads to a $\approx 300\%$ discrepancy between the WMVPF (-0.37) and the "augmented internal WMVPF" (-1.06). We provide three additional examples

[34]Hypothetically, the federal government *could* internalize all of the fiscal externalities that accrue to state governments by simply increasing/decreasing block grants given to states in the exact amount of fiscal externalities; in this case Problem 1 again becomes the relevant optimal reform problem.

of federal policies where the WMVPF and the "augmented internal WMVPF" are substantially different (due to state-level fiscal externalities) in Appendix B.6. Overall, the takeaway is that Problem 1 is a good enough approximation to a federal government's optimal reform problem as long as the share of fiscal externalities accruing to states is relatively small.[35]

5.2 Department of Education

Next, we consider a policymaker at the Department of Education who is considering how to allocate discretionary spending grants. Hypothetically, suppose the current status quo involves the policymaker allocating funds to six education policies, listed in Table 2 (the policies chosen are higher education programs with data on costs and welfare impacts from Hendren and Sprung-Keyser (2020)). And suppose the policymaker is hoping to improve welfare by reforming their grant spending: the policymaker therefore seeks to allocate marginally more money to certain programs and allocate marginally less money to other programs.[36] However, suppose that the policymaker has a budget to spend on discretionary grants each year and that she cannot control how fiscal externalities from these policies are spent because they occur far in the future to the Treasury Department via increased income tax revenue.[37] Because the policymaker faces a constraint on annual (mechanical) spending and cannot control how fiscal externalities are spent, the relevant optimal policy reform problem for this policymaker is Problem 2. Hence, the policymaker computes the WNSB for these six policies using estimates from Hendren and Sprung-Keyser (2020), presented in the final column of Table 2, with welfare weights equal to $\bar{\eta} = \bar{z}^{-\gamma}$ where $\bar{z}$ represents average income of policy beneficiaries and γ is a scalar which we assume equals 1 for purposes of illustration. The numeraire policy, p_{N+1}, is assumed to be linear income taxation (consistent with the NSB literature); to calculate the WMVPF of linear income taxation we assume that $\bar{\eta} = \bar{z}^{-\gamma}$ where $\bar{z}$ is mean household income in 2016 and that the taxable income elasticity is 0.3 (Saez et al., 2012) (see notes for Table 2 for exact calculation).

As we see from Table 2, the optimal budget neutral policy reform when faced with Problem 2 is to raise mechanical revenue by decreasing funding to the Adams Scholarship (which is a Massachusetts-specific college financial aid program) and then spend this mechanical revenue on funding more Pell Grants in Texas. This leads to positive fiscal externalities that are then spent on reducing the linear income tax rate as in Problem 2. However, if, hypothetically, the policymaker was able to optimally spend fiscal externalities arising from her policy decisions (perhaps due to the federal government rebating fiscal externalities back to her as in Remark 1), then the

[35] In contrast, we show in Appendix B.7 that accounting for the fact that some share of fiscal externalities from state-level policies accrue to the federal government is more important for state-level policymakers.

[36] Again, this exercise should be viewed as a proof of concept: in reality, not all of these programs are currently funded by the Department of Education so for these policies, we assume that the Department of Education has a comparable program to spend on today with similar costs and benefits.

[37] While increased income tax revenue is not the only fiscal externality of education policies, it is typically the largest (Hendren and Sprung-Keyser, 2020).

Table 2: WMVPFs and WNSBs for Selected Education Policies

Program	Average Income	Net Cost	WTP	MVPF	$\bar{\eta}$	WMVPF	Mech. Cost	WNSB
District of Columbia Tuition Assistance Grant Program	22977	390	8959	22.98	0.44	**10.00**	1176	3.28
Florida Student Access Grant	21632	889	6592	7.42	0.46	3.43	891	3.31
Massachusetts Adams Scholarship	25978	1896	1364	0.72	0.38	0.28	1314	**0.24**
Expanded Admissions to Florida International University	35223	-24445	112844	-4.62	0.28	-1.31	2617	13.28
Pell Grants in Texas	32843	-17379	85737	-4.93	0.30	-1.50	1000	**28.03**
Cal Grant, GPA Threshold	36622	-2355	38575	-16.38	0.27	**-4.47**	4311	2.50

Notes: This Table displays WMVPFs and WNSBs for six education policies. Data on average income, net cost, WTP, and mechanical costs are taken from Hendren and Sprung-Keyser (2020). $\bar{\eta} = 10000 \times \text{Average Income}^{-1}$ (where the factor of 10000 is for readability). The WMVPF is calculated from Definition 1 and the WNSB is calculated from Definition 2 where the numeraire policy p_{N+1} is assumed to be linear income taxation. The WMVPF for linear income taxation equals $\bar{\eta}_{N+1} MVPF_{N+1}$. The logic of Appendix F.II of Hendren and Sprung-Keyser (2020) implies that $MVPF_{N+1} = \frac{1}{1-\frac{T'}{1-T'}\xi^u}$ where we assume $T' = 0.2$ is the average marginal tax rate (taken from Appendix G of Hendren and Sprung-Keyser (2020)) and $\xi^u = 0.3$ is the uncompensated labor supply elasticity (Saez et al., 2012). $\bar{\eta}_{N+1} = 10000 \times 97360^{-1}$ where 97360 is the average U.S. household income in 2016.

policymaker would instead solve Problem 1. In this case, the policymaker would raise money from the WMVPF minimizing policy, which in this case is the Cal Grant (a California-specific financial college financial aid program), and then spend the resulting revenue on the WMVPF maximizing policy, which is a tuition assistance program in D.C. (see Table 2). Note that the WMVPF minimizing policy, Cal Grant, has a negative WMVPF which means the policymaker spends money *mechanically* on this program so as to raise *net* revenue while simultaneously increasing welfare.

Finally, to illustrate Proposition 4, we compare the welfare gains from rebating fiscal externalities back to the policymaker as in Remark 1, thereby allowing her to optimally spend fiscal externalities (and therefore using the WMVPF over the WNSB). Suppose that we consider spending \$1,000 mechanically on the WNSB maximizing policy and raising \$1,000 in mechanical revenue from the WNSB minimizing policy: from Table 2, we see that both of these reforms *raise* net revenue so that this reform ends up raising $\approx \$19,000$ in fiscal externalities which are then used to lower linear income tax rates (so that net spending on linear income taxation increases by ≈ \$19,000). In total this raises welfare by ΔW^{WNSB}. If, hypothetically, the policymaker could instead optimally spend \$19,000 on net and raise \$19,000 on net (as in Problem 4) via the WMVPF minimizing and WMVPF maximizing policies, this raises welfare by ΔW^{WMVPF} with $\Delta W^{WMVPF} \geq \Delta W^{WNSB}$ by Proposition 4. Calculating the compensating variation as in Proposition 4 (in this example, $NC^{WNSB} = NR^{WNSB} = 19,000$), we find a large welfare gain from optimally spending fiscal externalities: the policymaker would need about \$9,000 extra to spend on the WNSB maximizing policy to make up this difference, suggesting large welfare gains from relaxing mechanical constraints and spending fiscal externalities optimally. Note, this compensating variation is high primarily because the WNSB maximizing policy generates a very large fiscal externality this is then spent on lowering linear income taxes (which is a highly sub-optimal use of fiscal externalities). The takeaway from this exercise is that relaxing mechanical constraints and allowing the policymaker to spend fiscal externalities optimally, which then entails using the WMVPF rather than the WNSB to make policy choices, can lead to very large welfare gains.

5.3 Cash Transfer Policies in Mexico

Our final stylized example considers an international development organization (IDO), such as the World Bank, considering whether to spend a \$1,000,000 grant in rural Mexico on an unconditional cash transfer (UCT) or a conditional cash transfer (CCT) that provides transfers to households only if their children attend school. We suppose that any fiscal externalities arising from the grant will accrue to the Mexican government, not the IDO. Thus, the relevant optimal local reform problem for this IDO is Problem 2.[38] To construct the WNSBs, the IDO needs estimates of the costs and benefits of these two policies. First, both programs entail the mechanical cost of transferring money to households which are valued dollar-for-dollar. Both programs also have overhead costs which are not valued at all; however, suppose that the UCT has lower overhead costs (due to no need for costly monitoring of schooling conditions). Additionally, both the UCT and CCT program increase years of schooling in the current generation with the CCT increasing schooling by more due to the conditionalities on the transfers requiring children to attend school. Hence, both programs increase incomes of the next generation and thus generate positive fiscal externalities in the next generation through increased consumption, income, and payroll tax revenue. Finally, the IDO believes that the Mexican government will spend these future fiscal externalities by lowering consumption tax rates for the future generation (i.e., policy p_{N+1} represents the consumption tax rate in Mexico).

[38] Technically the IDO's problem is a non-budget-neutral version of Problem 2 where the change in positive mechanical costs equals the size of the grant, the change in negative mechanical costs equals 0, and the change in total net costs equals the size of the grant; by the discussion in Section 3.2 the solution is still to spend on the WNSB maximizing policy.

Table 3 presents the WNSBs for both policies constructed using estimates from several papers analyzing the CCT program Progresa implemented in 1997 (see Appendix C for details). Based on these values, the IDO should invest in the UCT as this is the policy with the highest WNSB. However, now suppose the IDO can convince the Mexican government to borrow against the fiscal externalities that arise in the future and increase upfront spending by the net-present-value of these fiscal externalities on whichever program the IDO decides to fund. Now the IDO solves Problem 1.[39] Table 3 also presents the WMVPFs for the two policies assuming the Mexican government faces an 11% interest rate on additional borrowing (this was the 10-year government bond rate that the Mexican government faced during the Progresa era (OECD, 2024)). Because the CCT has a higher WMVPF, the IDO will choose to fund this policy instead if it can convince the Mexican government to use the fiscal externalities to increase upfront spending on the cash transfer.

Under this scenario, the Mexican government will no longer spend the fiscal externalities on reducing the consumption tax rate in the next generation,[40] but rather will borrow against these fiscal externalities so that mechanical (upfront) spending on the CCT is $1,449,275 ($1,000,000 coming from the IDO and $449,275 coming from future fiscal externalities discounted back to today).[41] Having the government spend the fiscal externalities on the CCT today as opposed to lowering the consumption tax tomorrow generates a larger increase in welfare: in particular, we calculate that the IDO would require a grant of $1,214,286 to be spent on the UCT (with fiscal externalities spent on lowering consumption taxes in the future) to generate the same welfare increase as the IDO spending a grant of $1,000,000 on the CCT (with fiscal externalities spent on the CCT today), i.e., CV from Definition 3 is equal to $214,286. Hence, there are sizable welfare gains to spending fiscal externalities optimally and basing policy decisions on the WMVPF as opposed to the WNSB in this stylized example.

Table 3: WMVPFs and WNSBs for CCT and UCT in Rural Mexico

Program	Net Cost	WTP	WMVPF	Mech. Cost	WNSB
CCT	0.70	0.82	1.17	1.00	0.86
UCT	0.86	0.96	1.12	1.00	0.98

Note: All values are normalized to 1 peso of mechanical spending. See Appendix C for details on how these estimates are obtained. We use a consumption tax in the second generation to close the budget constraint when calculating the WNSB.

[39]Technically, the IDO solves a non-budget-neutral version of Problem 1 where the change in positive net costs equals the size of the grant, while the change in negative net costs equals 0; by Section 3.2, the solution is nonetheless to spend on the WMVPF maximizing policy.

[40]We show in Appendix C that the consumption tax WMVPF is $0.11 \times 1.24 << 1$.

[41]From Table 3 we see that spending $1 on the CCT generates $0.31 of fiscal externalities (in present-value). Thus, spending $1,449,275 upfront on the CCT generates $449,275 of fiscal externalities making the net cost to the government $0 and the net cost to the IDO $1,000,000.

6 Conclusion

This paper has developed a framework to construct general optimal policy reforms starting from a status quo set of policies. We conclude by discussing how the theoretical results and examples provided herein translate into real-world policy recommendations. First, when deciding which welfare metric to use to construct an optimal local policy reform, policymakers need to consider what sort of constraints they face. For instance, a federal legislature which can borrow against future revenues that is considering reforming policies for which all fiscal externalities accrue to the federal government should use the WMVPF to determine optimal policy reforms. In contrast, a government agency determining education or environmental policy with an annual use-it-or-lose-it budget and where fiscal externalities accrue to the Treasury Department (rather than the agency in question) should use the WNSB to determine optimal policy reforms. Similarly, NGOs/IDOs enacting policy reforms for which fiscal externalities accrue to country governments should use the WNSB to determine optimal reforms unless they can dictate how fiscal externalities are spent (at which point they should use the WMVPF). In situations where only a share of fiscal externalities are controlled by the policymaker in question (e.g., a federal policy where a chunk of fiscal externalities accrue to state governments), policymakers should use the "augmented internal WMVPF" to determine optimal policy reforms. Finally, constraints on how fiscal externalities are spent, which render it optimal to use the WNSB or the "augmented internal WMVPF" rather than the WMVPF for policy reforms, can generate substantial inefficiences. These inefficiences can be remedied by rebating fiscal externalities back to the policymaker rather than spending them on a numeraire policy; this allows policymakers to spend more on policies that generate large fiscal externalities and thereby construct better policy reforms by using the WMVPF. In practice, this may entail the federal legislature endogenously increasing the budgets of agencies which spend on policies with large fiscal externalities or NGOs/IDOs dictating how fiscal externalities arising from their programs should be spent.

References

David R. Agrawal, William H. Hoyt, and Tidiane Ly. A new approach to evaluating the welfare effects of decentralized policies. *Working Paper*, 2022. URL http://dx.doi.org/10.2139/ssrn.4009872.

Pierre Bachas, Lucie Gadenne, and Anders Jensen. Informality, consumption taxes, and redistribution. *Review of Economic Studies*, Sep 2023. doi: 10.1093/restud/rdad095.

Charles L. Ballard and Don Fullerton. Distortionary taxes and the provision of public goods. *Journal of Economic Perspectives*, 6(3):117–131, September 1992. doi: 10.1257/jep.6.3.117. URL https://www.aeaweb.org/articles?id=10.1257/jep.6.3.117.

Katy Bergstrom and William Dodds. The targeting benefit of conditional cash transfers. *Journal of Public Economics*, 194:104330, Feb 2021. doi: 10.1016/j.jpubeco.2020.104330.

Katy Bergstrom and William Dodds. A general theory of inverse welfare weights. *Working Paper*, 2024.

Katy Bergstrom, William Dodds, and Juan Rios. Welfare Analysis of Changing Notches: Evidence from Bolsa Família. *American Economic Journal: Economic Policy (Forthcoming)*, 2024a. URL https://www.aeaweb.org/articles?id=10.1257/pol.20230024.

Katy Bergstrom, William Dodds, and Juan Rios. Optimal Non-Linear Tax Reforms. *Working Paper*, 2024b.

Raj Chetty. A new method of estimating risk aversion. *American Economic Review*, 96(5): 1821–1834, 2006.

Raj Chetty. Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics*, 1(1):451–488, Sep 2009. doi: 10.1146/annurev.economics.050708.142910.

W.E. Diewert. Optimal tax perturbations. *Journal of Public Economics*, 10(2):139–177, Oct 1978. doi: 10.1016/0047-2727(78)90034-8.

Avinash Dixit. Price changes and optimum taxation in a many-consumer economy. *Journal of Public Economics*, 11(2):143–157, Mar 1979. doi: 10.1016/0047-2727(79)90001-x.

Nada Eissa and Hilary Hoynes. Redistribution and tax expenditures: The earned income tax credit. *National Tax Journal*, 64(2.2):689–729, Jun 2011. doi: 10.17310/ntj.2011.2s.10.

Nada Eissa, Henrik Jacobsen Kleven, and Claus Thustrup Kreiner. 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, Apr 2008. doi: 10.1016/j.jpubeco.2007.08.005.

Martin Feldstein. On the theory of tax reform. *Journal of Public Economics*, 6(1-2):77–104, 1976.

Jorge Luis García and James Heckman. *On criteria for evaluating social programs*, Apr 2022a. doi: 10.3386/w30005.

Jorge Luis García and James Heckman. *Three criteria for evaluating social programs*, Sep 2022b. doi: 10.3386/w30507.

Jorge Luis García, James J. Heckman, Duncan Ermini Leaf, and María José Prados. Quantifying the life-cycle benefits of an influential early-childhood program. *Journal of Political Economy*, 128(7):2502–2541, Jul 2020. doi: 10.1086/705718.

Jorge Luis García, Frederik Bennhoff, and Duncan Ermini Leaf. *The dynastic benefits of early childhood education: Participant benefits and family spillovers*, Aug 2023. doi: 10.3386/w31555.

Roger Guesnerie. On the direction of tax reform. *Journal of Public Economics*, 7(2):179–202, Apr 1977. doi: 10.1016/0047-2727(77)90045-7.

James J. Heckman, Seong Hyeok Moon, Rodrigo Pinto, Peter A. Savelyev, and Adam Yavitz. The rate of return to the highscope perry preschool program. *Journal of Public Economics*, 94(1–2):114–128, Feb 2010. doi: 10.1016/j.jpubeco.2009.11.001.

Nathaniel Hendren and Ben Sprung-Keyser. A unified welfare analysis of government policies*. *The Quarterly Journal of Economics*, 135(3):1209–1318, Mar 2020. doi: 10.1093/qje/qjaa006.

Nathaniel Hendren and Ben Sprung-Keyser. *The case for using the MVPF in empirical welfare analysis*, May 2022. doi: 10.3386/w30029.

Roozbeh Hosseini and Ali Shourideh. Retirement financing: An optimal reform approach. *Econometrica*, 87(4):1205–1265, 2019. doi: https://doi.org/10.3982/ECTA15088. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15088.

Mark Huggett and Juan Carlos Parra. How well does the u.s. social insurance system provide social insurance? *Journal of Political Economy*, 118(1):76–112, 2010. ISSN 00223808, 1537534X. URL http://www.jstor.org/stable/10.1086/651513.

Henrik Kleven. Sufficient statistics revisited. *Annual Reviews*, May 2020. doi: 10.3386/w27242.

Jeffrey B. Liebman and Neale Mahoney. Do expiring budgets lead to wasteful year-end spending?

evidence from federal procurement. *American Economic Review*, 107(11):3510–49, November 2017. doi: 10.1257/aer.20131296. URL https://www.aeaweb.org/articles?id=10.1257/aer.20131296.

OECD, 2024. URL https://fred.stlouisfed.org/series/IRLTLT01MXM156N. Retrieved from FRED, accessed July 8th, 2024.

Benjamin A. Olken. Monitoring corruption: Evidence from a field experiment in indonesia. *Journal of Political Economy*, 115(2):200–249, Apr 2007. doi: 10.1086/517935.

Susan W Parker and Tom Vogl. Do conditional cash transfers improve economic outcomes in the next generation? evidence from mexico. *The Economic Journal*, 133(655):2775–2806, Jul 2023. doi: 10.1093/ej/uead049.

Emmanuel Saez, Joel Slemrod, and Seth H Giertz. The elasticity of taxable income with respect to marginal tax rates: A critical review. *Journal of Economic Literature*, 50(1):3–50, 2012. doi: 10.1257/jel.50.1.3.

World Bank. World Bank Open Data [dataset]. pip.worldbank.org [Accessed September 4, 2024], 2024.

A Appendix: Proofs

A.1 Proof of Proposition 1

Proof. Define $q_i = \frac{dNC}{dp_i} \times (p_i - \tilde{p}_i)$ with $\mathbf{q} = (q_1, q_2, ..., q_N)$. Let $\mathbf{q}^+ = (q_1 \mathbb{1}[q_1 > 0], q_2 \mathbb{1}[q_2 > 0], ..., q_N \mathbb{1}[q_N > 0])$ and $\mathbf{q}^- = -(q_1 \mathbb{1}[q_1 < 0], q_2 \mathbb{1}[q_2 < 0], ..., q_N \mathbb{1}[q_N < 0])$, noting that $\mathbf{q} = \mathbf{q}^+ - \mathbf{q}^-$. Changing variables from p_i to q_i^+ and q_i^-, we can rewrite Problem 1 as follows:

$$\max_{\mathbf{q}^+, \mathbf{q}^-} \left\{ \sum_{i=1}^{N} \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} q_i^+ - \sum_{i=1}^{N} \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} q_i^- \right\}$$
$$\text{s.t.} \ \sum_{i=1}^{N} q_i^+ = 1 \text{ and } \sum_{i=1}^{N} q_i^- = 1$$

By linear separability of the objective we can write this problem as:

$$\begin{aligned} &\max_{\mathbf{q}^+} \sum_{i=1}^{N} \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} q_i^+ + \max_{\mathbf{q}^-} \sum_{i=1}^{N} -\frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} q_i^- \\ &\text{s.t.} \ \sum_{i=1}^{N} q_i^+ = 1 \text{ and } \sum_{i=1}^{N} q_i^- = 1 \end{aligned} \tag{8}$$

Recognizing that the first constraint in Equation (8) applies only to the first maximization and the second constraint applies only to the second maximization, we effectively have two separate maximization problems. For the first problem we have that:

$$\begin{aligned} &\sum_{i=1}^{N} \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} q_i^+ \leq \sum_{i=1}^{N} \max_i \left\{ \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} \right\} q_i^+ \\ &= \max_i \left\{ \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} \right\} \sum_{i=1}^{N} q_i^+ = \max_i \left\{ \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} \right\} \end{aligned}$$

where the first inequality follows because $q_i^+ \geq 0 \ \forall i$. We can achieve this maximum by setting $q_i^+ = 1$ if $i = \arg\max \left\{ \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} \right\} = \arg\max \{WMVPF_i\}$ and $q_i^+ = 0$ otherwise.

Identical logic implies that the solution to the second maximization problem from Equation (8) is to set $q_i^- = 1$ if $i = \arg\max \left\{ -\frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} \right\} = \arg\min \left\{ \frac{dW}{dp_i} \Big/ \frac{dNC}{dp_i} \right\} = \arg\min\{WMVPF_i\}$ and $q_i^- = 0$ otherwise.

Using the definition of q_i^+ and q_i^-, we have that the solution to Equation (8) is to set:

$$(p_i - \tilde{p}_i) = \begin{cases} \frac{1}{\frac{dNC}{dp_i}} & i = \arg\max\{WMVPF_i\} \\ \frac{-1}{\frac{dNC}{dp_i}} & i = \arg\min\{WMVPF_i\} \\ 0 & \text{otherwise} \end{cases}$$

□

A.2 Proof of Proposition 2

Proof. The net cost constraint from Problem 2 can be rewritten as:

$$\begin{aligned}\frac{dNC}{dp_{N+1}}(p_{N+1}-\tilde{p}_{N+1}) &= -\sum_{i=1}^{N}\frac{dNC}{dp_i}(p_i-\tilde{p}_i)\\ &= -\sum_{i=1}^{N}\left(\frac{dNC}{dp_i}-\frac{dMC}{dp_i}\right)(p_i-\tilde{p}_i)\\ &= \sum_{i=1}^{N}\frac{dFE}{dp_i}(p_i-\tilde{p}_i)\end{aligned} \tag{9}$$

where the second equality results because $\sum_{i=1}^{N}\frac{dMC}{dp_i}(p_i-\tilde{p}_i)=0$ by the first two constraints of Problem 2.

Next, define $q_i = \frac{dMC}{dp_i}(p_i-\tilde{p}_i)$. Let $\mathbf{q}^+ = (q_1\mathbb{1}[q_1>0], q_2\mathbb{1}[q_2>0], ..., q_N\mathbb{1}[q_N>0])$ and $\mathbf{q}^- = -(q_1\mathbb{1}[q_1<0], q_2\mathbb{1}[q_2<0], ..., q_N\mathbb{1}[q_N<0])$, noting that $\mathbf{q}=\mathbf{q}^+-\mathbf{q}^-$. Plugging in the expression for $(p_{N+1}-\tilde{p}_{N+1})$ from Equation (9) and changing variables from p_i to q_i^+ and q_i^-, we aim to maximize:

$$\begin{aligned}&\max_{\mathbf{q}^+,\mathbf{q}^-}\ \sum_{i=1}^{N}\left(\frac{dW}{dp_i}+\frac{dW}{dp_{N+1}}\frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}}\right)\Big/\left(\frac{dMC}{dp_i}\right)(q_i^+-q_i^-)\\ &\text{s.t. } \sum_{i=1}^{N}q_i^+=1 \text{ and } \sum_{i=1}^{N}q_i^-=1\end{aligned}$$

Using identical logic as in Proposition 1, the solution to the above problem is to set $q_i^+=1$ when $i=\arg\max\{WNSB_i^{(p_{N+1})}\}$ and $q_i^+=0$ otherwise; similarly, we set $q_i^-=1$ when $i=\arg\min\{WNSB_i^{(p_{N+1})}\}$ and $q_i^-=0$ otherwise. Given the definition of q_i^+ and q_i^- in terms of $(p_i-\tilde{p}_i)$, the solution stated in the Proposition follows. □

A.3 Proof of Proposition 3

Proof. The net cost constraint from Problem 3 can be rewritten as:

$$\begin{aligned}\frac{dNC}{dp_{N+1}}(p_{N+1}-\tilde{p}_{N+1}) &= -\sum_{i=1}^{N}\frac{dNC}{dp_i}(p_i-\tilde{p}_i)\\ &= -\sum_{i=1}^{N}\left(\frac{dNC}{dp_i}-\frac{dNC_I}{dp_i}\right)(p_i-\tilde{p}_i)\\ &= \sum_{i=1}^{N}(1-\sigma_i)\frac{dFE}{dp_i}(p_i-\tilde{p}_i)\end{aligned} \tag{10}$$

where the second equality results because $\sum_{i=1}^{N}\frac{dNC_I}{dp_i}(p_i-\tilde{p}_i)=0$ by the first two constraints of Problem 3.

Next, define $q_i = \frac{dNC_I}{dp_i}(p_i-\tilde{p}_i)$. Let $\mathbf{q}^+ = (q_1\mathbb{1}[q_1>0], q_2\mathbb{1}[q_2>0], ..., q_N\mathbb{1}[q_N>0])$ and $\mathbf{q}^- = -(q_1\mathbb{1}[q_1<0], q_2\mathbb{1}[q_2<0], ..., q_N\mathbb{1}[q_N<0])$, noting that $\mathbf{q}=\mathbf{q}^+-\mathbf{q}^-$. Plugging in the expression for $(p_{N+1} - \tilde{p}_{N+1})$ from Equation (10) and changing variables from p_i to q_i^+ and q_i^-, we aim to maximize:

$$\max_{\mathbf{q}^+,\mathbf{q}^-} \sum_{i=1}^{N} \left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{(1-\sigma_i)\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \Big/ \left(\frac{dNC_I}{dp_i} \right) (q_i^+ - q_i^-)$$

$$\text{s.t. } \sum_{i=1}^{N} q_i^+ = 1 \text{ and } \sum_{i=1}^{N} q_i^- = 1$$

Using identical logic as in Proposition 2, the solution to the above problem is to set $q_i^+ = 1$ when $i = \arg\max \left\{ \left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{(1-\sigma_i)\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \Big/ \left(\frac{dNC_I}{dp_i} \right) \right\}$ and $q_i^+ = 0$ otherwise; similarly, we set $q_i^- = 1$ when $i = \arg\min \left\{ \left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{(1-\sigma_i)\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \Big/ \left(\frac{dNC_I}{dp_i} \right) \right\}$ and $q_i^- = 0$ otherwise. Given the definition of q_i^+ and q_i^- in terms of $(p_i - \tilde{p}_i)$, the solution stated in the Proposition follows. □

A.4 Proof of Proposition 4

Proof. Consider the following general problem:

Problem 5. *Given a status quo policy $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$:*

$$\max_{\mathbf{p}, p_{N+1}} \sum_{i=1}^{N+1} \frac{dW}{dp_i}(p_i - \tilde{p}_i)$$

$$\text{s.t. } \sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i) \mathbb{1}\left[\frac{dMC}{dp_i}(p_i - \tilde{p}_i) > 0\right] = MC^{WNSB}$$

$$\sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i) \mathbb{1}\left[\frac{dMC}{dp_i}(p_i - \tilde{p}_i) < 0\right] = -MR^{WNSB}$$

$$\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0$$

$$\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) \mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) > 0\right] = NC^{WNSB}$$

$$\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) \mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) < 0\right] = -NR^{WNSB}$$

where $MC^{WNSB} = MR^{WNSB} = 1$ and $NC^{WNSB} = NR^{WNSB}$ by budget neutrality.[42] Note that the solution to Problem 5 is the same as the solution to Problem 2 because the final two constraints are satisfied under the solution to Problem 2, $\{p_i^{WNSB}\}$, by the definition of NC^{WNSB} and NR^{WNSB}, see Equation (7). But we know that the value of the objective function from solving Problem 5 (and therefore Problem 2) must (weakly) be lower than the following Problem which has fewer constraints:

[42] Note, in Problem 5 we allow the policymaker to pick p_{N+1}, but because they must change policies $p_1, ..., p_N$ to satisfy the first two mechanical spending constraints, policy p_{N+1} is pinned down by the third constraint that net costs of the reform are 0.

Problem 6. *Given a status quo policy* $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$*:*

$$\max_{\mathbf{p}, p_{N+1}} \sum_{i=1}^{N+1} \frac{dW}{dp_i}(p_i - \tilde{p}_i)$$
$$\text{s.t.} \sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) \mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) > 0\right] = NC^{WNSB}$$
$$\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) \mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) < 0\right] = -NR^{WNSB}$$

But note that Problem 6 is essentially equivalent to Problem 1 with the net cost and net revenue constraints replaced with NC^{WNSB} and $-NR^{WNSB}$ instead of 1 and -1. From here, it is straight-forward to show (using identical logic as in the proof to Proposition 1, see Section 3.2) that the solution to this problem is therefore to spend NC^{WNSB} on the WMVPF maximizing policy and raise NR^{WNSB} revenue via the WMVPF minimizing policy so that:

$$(p_i^{WMVPF} - \tilde{p}_i) = \begin{cases} \frac{NC^{WNSB}}{\frac{dNC}{dp_i}} & i = \arg\max\{WMVPF_i\} \\ \frac{-NR^{WNSB}}{\frac{dNC}{dp_i}} & i = \arg\min\{WMVPF_i\} \\ 0 & \text{otherwise} \end{cases}$$

But note that Problem 4 is the same as Problem 6 but restricted to only spend on policies $i = 1, 2, ..., N$; hence, the solutions to these problems coincide as long as policy p_{N+1} is not the highest or lowest WMVPF policy, which we assume. Given that we denote welfare under the solution to Problem 4 as $W\left(\{p_i^{WMVPF}\}\right)$ and welfare under the solution to Problem 5 as $W\left(\{p_i^{WNSB}\}\right)$ we know that $W\left(\{p_i^{WMVPF}\}\right) \geq W\left(\{p_i^{WNSB}\}\right)$, which is the first statement of the proposition.

The welfare gain from the solution to Problem 4 equals (to first order, which is WLOG given that we are studying small reforms):

$$\begin{aligned} W\left(\{p_i^{WMVPF}\}\right) - W(\tilde{\mathbf{p}}, \tilde{p}_{N+1}) &\approx \sum_{i=1}^{N} \frac{dW}{dp_i}(p_i^{WMVPF} - \tilde{p}_i) \\ = \sum_{i=1}^{N} \Bigg(\frac{dW}{dp_i} \frac{NC^{WNSB}}{\frac{dNC}{dp_i}} \mathbb{1}[i = \arg\max\{WMVPF_i\}] &- \frac{dW}{dp_i} \frac{NR^{WNSB}}{\frac{dNC}{dp_i}} \mathbb{1}[i = \arg\min\{WMVPF_i\}] \Bigg) \\ = NC^{WNSB} \max\{WMVPF_i\} - NR^{WNSB} \min\{WMVPF_i\} \end{aligned} \tag{11}$$

where the first equality plugs in the definition of $\{p_i^{WMVPF}\}$ from above.

To first order, the welfare gain from solving Problem 2 can be written as follows by the definition of $W\left(\{p_i^{WNSB}\} + CV \arg\max\{WNSB_i^{(p_{N+1})}\}\right)$ given in Definition 3:

$$
\begin{aligned}
& W\left(\{p_i^{WNSB}\} + CV \arg\max\{WNSB_i^{(p_{N+1})}\}\right) - W(\tilde{\mathbf{p}}, \tilde{p}_{N+1}) \\
& \approx \sum_{i=1}^{N+1} \frac{dW}{dp_i}(p_i^{WNSB} - \tilde{p}_i) + \frac{CV}{\frac{dMC}{dp_i}} \left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \mathbb{1}[i = \arg\max\{WNSB_i^{p_{N+1}}\}] \\
& = \sum_{i=1}^{N} \left\{ \left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \Big/ \left(\frac{dMC}{dp_i} \right) \mathbb{1}[i = \arg\max\{WNSB_i^{p_{N+1}}\}] \right. \\
& \left. - \left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \Big/ \left(\frac{dMC}{dp_i} \right) \mathbb{1}[i = \arg\min\{WNSB_i^{p_{N+1}}\}] \right\} \\
& + CV \max\{WNSB_i^{(p_{N+1})}\} \\
& = \max\{WNSB_i^{(p_{N+1})}\} - \min\{WNSB_i^{(p_{N+1})}\} + CV \max\{WNSB_i^{(p_{N+1})}\}
\end{aligned}
\tag{12}
$$

where the first approximation uses the fact that the policymaker is able to spend CV dollars on the highest WNSB policy (allowing them to raise this policy by $\frac{CV}{\frac{dMC}{dp_i}}$) and any resulting fiscal externalities are spent on policy p_{N+1}. The first equality then follows by the expression for the welfare gain from the optimal policy $\{p_i^{WNSB}\}$ given in the proof to Proposition 2 in Appendix A.2. Combining Equations (11) and (12) yields the second statement of the proposition, recognizing that by budget neutrality $NC^{WNSB} = NR^{WNSB}$ and $MC^{WNSB} = MR^{WNSB} = 1$.

□

B Appendix: Additional Results

B.1 Problem 3 Global

Consider a policymaker who would like to solve the following global problem:

Problem 3 Global. *Given a status quo policy* $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$*:*

$$
\begin{aligned}
\max_{\mathbf{p}} \quad & W(\mathbf{p}, p_{N+1}) - W(\tilde{\mathbf{p}}, \tilde{p}_{N+1}) \\
\text{s.t.} \quad & \underbrace{C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1})) - C(\tilde{\mathbf{p}}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1}))}_{\text{Mechanical Cost, } MC(\mathbf{p})} + \\
& \underbrace{S(\mathbf{p}, [C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, \tilde{p}_{N+1})) - C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1}))])}_{\text{Internal Share of Fiscal Externalities, } S(\mathbf{p}, FE(\mathbf{p}))} = 0 \\
& \underbrace{C(\mathbf{p}, p_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, p_{N+1})) - C(\tilde{\mathbf{p}}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1}))}_{\text{Net Cost, } NC(\mathbf{p}, p_{N+1})} = 0
\end{aligned}
$$

Problem 3 Global contains two constraints. The first constraint ensures that the policymaker's internal budget constraint is unchanged by the policy reform and mandates that the sum of mechanical costs plus the share of fiscal externalities that accrue internally, S, equals zero. Note, the net internal cost of the reform is equal to the mechanical cost of the reform plus the share of fiscal externalities of the reform accruing internally: $NC_I(\mathbf{p}) = MC(\mathbf{p}) + S(\mathbf{p}, FE(\mathbf{p}))$.[43] The second constraint mandates that the total net cost of the reform is unchanged and ensures that the fiscal externalities that accrue externally to the policymaker in question are funded by changing the numeraire policy p_{N+1}. To see this, subtracting the first constraint from the second constraint, and adding and subtracting $C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, \tilde{p}_{N+1}))$ yields:

$$\begin{aligned}&\underbrace{C(\mathbf{p}, p_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, p_{N+1})) - C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, \tilde{p}_{N+1}))}_{\text{Cost Changes Due to } p_{N+1}} \\ &+ \underbrace{[C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, \tilde{p}_{N+1})) - C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1}))]}_{\text{Fiscal Externalities, } FE(\mathbf{p})} \\ &- \underbrace{S(\mathbf{p}, [C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \mathbf{p}, \tilde{p}_{N+1})) - C(\mathbf{p}, \tilde{p}_{N+1}, \mathbf{y}(\mathbf{n}; \tilde{\mathbf{p}}, \tilde{p}_{N+1}))])}_{\text{Internal Share of Fiscal Externalities, } S(\mathbf{p}, FE(\mathbf{p}))} = 0\end{aligned} \tag{13}$$

In order to argue that Problem 3 Global is a global version of Problem 3, it should be the case that the maximum value of Problem 3 is zero if we start from a status quo that is a local maximum of Problem 3 Global. We can show this by contradiction: suppose that we start from a status quo that is a local maximum of Problem 3 Global and the maximum value of Problem 3 is non-zero (and therefore positive because if there is a policy reform that generates a negative welfare impact then the negative of that reform generates a positive welfare impact by our differentiability assumptions). Because $\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i)$ is a first-order Taylor series approximation around $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$ for the net cost in Problem 3 Global's first constraint, we know that (to first order) the net cost constraint must also be satisfied under the reformed policy vector from Problem 3. Moreover, the first order Taylor series approximation around $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$ for the internal budget in Problem 3 Global's second constraint is:

$$\sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i) - \sum_{i=1}^{N} \left[\frac{\partial S}{\partial FE}\frac{dFE}{dp_i} + \frac{\partial S}{\partial p_i}\right](p_i - \tilde{p}_i)$$

If we define $\sigma_i \equiv \frac{\partial S}{\partial FE} + \frac{\partial S}{\partial p_i}\frac{1}{\frac{dFE}{dp_i}}$, then we can rewrite the above as:

$$\sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i) - \sum_{i=1}^{N} \left[\sigma_i \frac{dFE}{dp_i}\right](p_i - \tilde{p}_i)$$

and we know that this constraint is also satisfied under the solution to Problem 3. Hence, with properly defined σ_i, if the solution to Problem 3 is not zero then we have a found a policy vector arbitrarily close to the status quo that improves welfare yet still satisfies the constraints (to first order) of Problem 3 Global. If we were to form a Lagrangian for Problem 3 Global, then the solution to Problem 3 would increase the value of this Lagrangian (by leaving the constraints unchanged up to first order and increasing welfare), which (by the Lagrange multiplier theorem) contradicts the fact that we assumed that the status quo was a local maximum of Problem 3

[43] We assume that any changes in p_{N+1} (that ensue to pay for any fiscal externalities that accrue externally) do not impact the internal costs of the agency in question.

Global.

As with Problem 1 Global and Problem 2 Global, Problem 3 Global is very difficult to relate to empirical estimates, which is why we focus our attention on instead solving the local version of Problem 3 Global given by Problem 3.

B.2 Uncertainty in Costs and Benefits

We first prove an analogue for Proposition 1. Consider the following problem where the government chooses policies $\{p_1, p_2, ..., p_N\} = \mathbf{p}$ each of which is subject to uncertainty in both the benefits and net costs. The government is constrained so that in expectation they are only allowed to spend \$1 net on these policies; however, the government still has a budget constraint that must be satisfied under all realizations of the uncertainty and the government satisfies this budget constraint with a numeraire policy p_{N+1} which has no uncertainty in the costs or benefits:

Problem 7. *Given a status quo policy $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$ and a numeraire policy p_{N+1} for which the welfare effects and cost have no uncertainty:*

$$\max_{\mathbf{p}} \sum_{i=1}^{N+1} \mathbb{E}\left[\frac{dW}{dp_i}(p_i - \tilde{p}_i)\right]$$

$$\text{s.t.} \sum_{i=1}^{N} \mathbb{E}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) > 0\right]\right] = 1$$

$$\sum_{i=1}^{N} \mathbb{E}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dNC}{dp_i}(p_i - \tilde{p}_i) < 0\right]\right] = -1$$

$$\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0$$

Proposition 5. *The solution to Problem 7 is to set:*

$$(p_i - \tilde{p}_i) = \begin{cases} \frac{1}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]} & i = \arg\max\{WMVPF_i\} \\ \frac{-1}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]} & i = \arg\min\{WMVPF_i\} \\ 0 & \textit{otherwise} \end{cases}$$

where $WMVPF_i$ is defined as in Equation 4.

Proof. Define $q_i = \mathbb{E}\left[\frac{dNC}{dp_i} \times (p_i - \tilde{p}_i)\right] = \mathbb{E}\left[\frac{dNC}{dp_i}\right] \times (p_i - \tilde{p}_i)$. Let $\mathbf{q}^+ = (q_1\mathbb{1}[q_1 > 0], q_2\mathbb{1}[q_2 > 0], ..., q_N\mathbb{1}[q_N > 0])$ and $\mathbf{q}^- = -(q_1\mathbb{1}[q_1 < 0], q_2\mathbb{1}[q_2 < 0], ..., q_N\mathbb{1}[q_N < 0])$, noting that $\mathbf{q} =$

$\mathbf{q}^+ - \mathbf{q}^-$. A change of variables therefore allows us to rewrite Problem 7:

$$\max_{\mathbf{q}^+,\mathbf{q}^-} \sum_{i=1}^{N} \mathbb{E}\left[\frac{dW}{dp_i}(q_i^+ - q_i^-)\right] \Big/ \mathbb{E}\left[\frac{dNC}{dp_i}\right] + \frac{dW}{dp_{N+1}}\mathbb{E}\left[p_{N+1} - \tilde{p}_{N+1}\right]$$
$$\text{s.t. } \sum_{i=1}^{N} q_i^+ = 1 \text{ and } \sum_{i=1}^{N} q_i^- = 1$$
$$\frac{dNC}{dp_{N+1}}(p_{N+1} - \tilde{p}_{N+1}) = -\sum_{i=1}^{N} \frac{dNC}{dp_i} \frac{q_i^+ - q_i^-}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]}$$

Plugging in the budget constraint (and using the fact that there is no uncertainty in the welfare or budgetary impacts of policy p_{N+1}):

$$\max_{\mathbf{q}^+,\mathbf{q}^-} \sum_{i=1}^{N} \frac{\mathbb{E}\left[\frac{dW}{dp_i}(q_i^+ - q_i^-)\right]}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]} - \frac{\frac{dW}{dp_{N+1}}}{\frac{dNC}{dp_{N+1}}}\mathbb{E}\left[\sum_{i=1}^{N} \frac{dNC}{dp_i} \frac{q_i^+ - q_i^-}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]}\right]$$
$$\text{s.t. } \sum_{i=1}^{N} q_i^+ = 1 \text{ and } \sum_{i=1}^{N} q_i^- = 1$$

However, noting that because $q_i^+ - q_i^- = q_i = \mathbb{E}\left[\frac{dNC}{dp_i} \times (p_i - \tilde{p}_i)\right]$, we have that: $\mathbb{E}\left[\frac{dNC}{dp_i}(q_i^+ - q_i^-)\right] = \mathbb{E}\left[\frac{dNC}{dp_i}\right](q_i^+ - q_i^-)$. Thus:

$$\mathbb{E}\left[\sum_{i=1}^{N} \frac{dNC}{dp_i} \frac{q_i^+ - q_i^-}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]}\right] = \sum_{i=1}^{N}(q_i^+ - q_i^-) = 0$$

Hence, the problem reduces to:

$$\max_{\mathbf{q}^+,\mathbf{q}^-} \sum_{i=1}^{N} \frac{\mathbb{E}\left[\frac{dW}{dp_i}\right]}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]} q_i^+ + \sum_{i=1}^{N} -\frac{\mathbb{E}\left[\frac{dW}{dp_i}\right]}{\mathbb{E}\left[\frac{dNC}{dp_i}\right]} q_i^-$$
$$\text{s.t. } \sum_{i=1}^{N} q_i^+ = 1 \text{ and } \sum_{i=1}^{N} q_i^- = 1$$

From here, identical logic as in the proof to Proposition 1 (starting from Equation (8)) shows that the statement of the Proposition holds. □

Next, we prove an analogue for Proposition 2 when there is uncertainty in benefits and costs. We consider the following augmented version of Problem 2 where there is uncertainty in benefits and costs of each policy other than the numeraire policy (and we assume that there is no uncertainty about the mechanical costs of each policy):

Problem 8. *Given a status quo policy* $(\tilde{\mathbf{p}}, \tilde{p}_{N+1})$ *and a numeraire policy* p_{N+1} *for which the* *welfare effects and cost have no uncertainty:*

$$\max_{\mathbf{p}} \sum_{i=1}^{N+1} \mathbb{E}\left[\frac{dW}{dp_i}(p_i - \tilde{p}_i)\right]$$

$$\text{s.t.} \sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dMC}{dp_i}(p_i - \tilde{p}_i) > 0\right] = 1$$

$$\sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i)\mathbb{1}\left[\frac{dMC}{dp_i}(p_i - \tilde{p}_i) < 0\right] = -1$$

$$\sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0$$

Proposition 6. *The solution to Problem 8 is to set:*

$$(p_i - \tilde{p}_i) = \begin{cases} \frac{1}{\frac{dMC}{dp_i}} & i = \arg\max\{WNSB_i^{(p_{N+1})}\} \\ \frac{-1}{\frac{dMC}{dp_i}} & i = \arg\min\{WNSB_i^{(p_{N+1})}\} \\ 0 & \textit{otherwise} \end{cases}$$

where $WNSB_i$ is defined as in Equation (5).

Proof. As in Appendix A.2, the net cost constraint from Problem 8 can be rewritten as:

$$\frac{dNC}{dp_{N+1}}(p_{N+1} - \tilde{p}_{N+1}) = \sum_{i=1}^{N} \frac{dFE}{dp_i}(p_i - \tilde{p}_i) \tag{14}$$

Next, define $q_i = \frac{dMC}{dp_i}(p_i - \tilde{p}_i)$. Let $\mathbf{q}^+ = (q_1\mathbb{1}[q_1 > 0], q_2\mathbb{1}[q_2 > 0], ..., q_N\mathbb{1}[q_N > 0])$ and $\mathbf{q}^- = -(q_1\mathbb{1}[q_1 < 0], q_2\mathbb{1}[q_2 < 0], ..., q_N\mathbb{1}[q_N < 0])$, noting that $\mathbf{q} = \mathbf{q}^+ - \mathbf{q}^-$. Doing a change of variables from p_i to q_i^+, q_i^- and plugging in the net cost constraint, we can rewrite Problem 8 as (where we distribute the expectation operator and use the fact that there is no uncertainty in q_i^+, q_i^-, $\frac{dMC}{dp_i}$, $\frac{dW}{dp_{N+1}}$, or $\frac{dNC}{dp_{N+1}}$):

$$\max_{\mathbf{q}^+,\mathbf{q}^-} \sum_{i=1}^{N} \left(\mathbb{E}\left[\frac{dW}{dp_i}\right] + \frac{dW}{dp_{N+1}} \frac{\mathbb{E}\left[\frac{dFE}{dp_i}\right]}{\frac{dNC}{dp_{N+1}}}\right) \frac{q_i^+ - q_i^-}{\frac{dMC}{dp_i}}$$

$$\text{s.t.} \sum_{i=1}^{N} q_i^+ = 1 \text{ and } \sum_{i=1}^{N} q_i^- = 1$$

From here, identical logic as in the proof to Proposition 2 shows that the statement of the Proposition holds. □

B.3 Numerical Example with L^2 Constraint

Suppose we have three policies p_1, p_2, and p_3 with $\frac{dW}{dp_i} = \{1, 2.1, 2\}$ and $\frac{dNC}{dp_i} = \{0.5, 1, 0.2\}$. Solving Problem 6 assuming the policymaker chooses p_1 and p_2, we have the following problem (where we assume the status quo has $\tilde{p}_1 = \tilde{p}_2 = \tilde{p}_3 = 0$ to reduce notation):

$$\begin{aligned} \max_{p_1,p_2} \; & p_1 + 2.1p_2 \\ \text{s.t. } & 0.5p_1 + p_2 = 0 \text{ and } p_1^2 + p_2^2 = 1 \end{aligned} \tag{15}$$

The solution to the Equation (15) is to set $(p_1, p_2) = (-0.894427, 0.447214)$ so that the policymaker reduces policy p_1 more than it reduces policy p_2 (and actually increases policy p_2). But now, consider the problem where the policymaker gets to choose p_1, p_2, and p_3. Now, the policymaker solves:

$$\begin{aligned} \max_{p_1,p_2,p_3} \; & p_1 + 2.1p_2 + 2p_3 \\ \text{s.t. } & 0.5p_1 + p_2 + 0.2p_3 = 0 \text{ and } p_1^2 + p_2^2 + p_3^2 = 1 \end{aligned} \tag{16}$$

The solution to Equation (16) is to set $(p_1, p_2, p_3) = (-0.10436, -0.144614, 0.983969)$ so that the policymaker reduces policy p_1 by *less* than p_2, in reverse to the solution of Equation (15).

B.4 L^2 Constraints on Budgetary Impacts

We now consider the following augmented version of Equation 6 where the L^2 norm constraint is now on the vector of budgetary impacts of the policy reform:

$$\begin{aligned} \max_{\mathbf{p}} \; & \sum_{i=1}^{N} \frac{dW}{dp_i}(p_i - \tilde{p}_i) \\ \text{s.t. } & \sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0 \text{ and } \|\nabla_{\mathbf{p}} NC \circ (\mathbf{p} - \tilde{\mathbf{p}})\|_2 = 1 \end{aligned} \tag{17}$$

Assuming that $\frac{dNC}{dp_i} > 0 \; \forall i$ we can do a change of variables setting $q_i = \frac{dNC}{dp_i}(p_i - \tilde{p}_i)$ so that the above problem can be rewritten with $\mathbf{q} = (q_1, q_2, ..., q_N)$:

$$\begin{aligned} \max_{\mathbf{q}} \; & \sum_{i=1}^{N} \frac{\frac{dW}{dp_i}}{\frac{dNC}{dp_i}} q_i \\ \text{s.t. } & \sum_{i=1}^{N} q_i = 0 \text{ and } \|\mathbf{q}\|_2 = 1 \end{aligned} \tag{18}$$

Appending the constraint $\sum_{i=1}^{N} q_i = 0$ onto the objective with a Lagrange multiplier λ and appealing to the Cauchy-Schwarz inequality, we can immediately deduce that the solution to Equation (18) is to set:

$$q_i = \frac{\frac{dW}{dp_i} / \frac{dNC}{dp_i} - \lambda}{\left\| \left[\frac{dW}{dp_i} / \frac{dNC}{dp_i} - \lambda \right] \right\|_2}$$

where $\left[\frac{dW}{dp_i} / \frac{dNC}{dp_i} - \lambda\right]$ represents the vector formed with each element i given by $\frac{dW}{dp_i} / \frac{dNC}{dp_i} - \lambda$. Note that q_i represents the total amount of spending on policy i due to the reform; hence, the optimal local reform to spending on policy i is a function only of the WMVPF of policy i along with the Lagrange multiplier λ which is chosen to satisfy the budget-neutrality constraint.

We can also consider a similar L^2 norm problem for the case where the policymaker does not control how fiscal externalities are spent and therefore fiscal externalities are spent on some numeraire policy p_{N+1}, recalling that $\mathbf{p} = (p_1, p_2, ..., p_N)$:

$$\begin{aligned} \max_{\mathbf{p}} \quad & \sum_{i=1}^{N+1} \frac{dW}{dp_i}(p_i - \tilde{p}_i) \\ \text{s.t.} \quad & \sum_{i=1}^{N} \frac{dMC}{dp_i}(p_i - \tilde{p}_i) = 0 \\ & \left\| \nabla_{\mathbf{p}} MC \circ (\mathbf{p} - \tilde{\mathbf{p}}) \right\|_2 = 1 \\ & \sum_{i=1}^{N+1} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) = 0 \end{aligned} \tag{19}$$

Problem (19) attempts to maximize the welfare gain from the policy reform subject to a (mechanical) budget-neutrality constraint, an L^2 norm constraint on the size of mechanical cost changes, and a total budget-neutrality constraint which, as discussed in Section 2.3, ensures that all fiscal externalities are funded via policy p_{N+1}. Using identical logic as in Appendix A.2, the net cost constraint from Equation (19) can be rewritten as (using the fact that mechanical costs sum to zero):

$$\begin{aligned} \frac{dNC}{dp_{N+1}}(p_{N+1} - \tilde{p}_{N+1}) &= -\sum_{i=1}^{N} \frac{dNC}{dp_i}(p_i - \tilde{p}_i) \\ &= -\sum_{i=1}^{N} \left(\frac{dNC}{dp_i} - \frac{dMC}{dp_i} \right)(p_i - \tilde{p}_i) \\ &= \sum_{i=1}^{N} \frac{dFE}{dp_i}(p_i - \tilde{p}_i) \end{aligned} \tag{20}$$

Next, let us substitute in the value of $(p_{N+1} - \tilde{p}_{N+1})$ determined by $\mathbf{p} = (p_1, p_2, ..., p_N)$ via Equation (20) and assume that $\frac{dMC}{dp_i} > 0 \ \forall i$ so we can do a change of variables setting $q_i = \frac{dMC}{dp_i}(p_i - \tilde{p}_i)$. Hence, the above problem can be rewritten with $\mathbf{q} = (q_1, q_2, ..., q_N)$:

$$\begin{aligned} \max_{\mathbf{q}} \quad & \sum_{i=1}^{N} \left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \Big/ \left(\frac{dMC}{dp_i} \right) q_i \\ \text{s.t.} \quad & \sum_{i=1}^{N} q_i = 0 \text{ and } \left\| \mathbf{q} \right\|_2 = 1 \end{aligned} \tag{21}$$

Again appending the constraint $\sum_{i=1}^{N} q_i = 0$ onto the objective with a Lagrange multiplier λ and appealing to the Cauchy-Schwarz inequality, we can immediately deduce that the solution to Equation (21) is to set:

$$q_i = \frac{\left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \Big/ \left(\frac{dMC}{dp_i} \right) - \lambda}{\left\| \left[\left(\frac{dW}{dp_i} + \frac{dW}{dp_{N+1}} \frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}} \right) \Big/ \left(\frac{dMC}{dp_i} \right) - \lambda \right] \right\|_2}$$

where $\left[\left(\frac{dW}{dp_i}+\frac{dW}{dp_{N+1}}\frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}}\right)\Big/\left(\frac{dMC}{dp_i}\right)-\lambda\right]$ represents the vector with element i given by $\left(\frac{dW}{dp_i}+\frac{dW}{dp_{N+1}}\frac{\frac{dFE}{dp_i}}{\frac{dNC}{dp_{N+1}}}\right)\Big/\left(\frac{dMC}{dp_i}\right)-\lambda$. Note that q_i represents the total amount of spending on policy i due to the reform; hence, the optimal reform to spending on policy i is a function only of the WNSB of policy i along with the Lagrange multiplier λ which is chosen to satisfy the budget-neutrality constraint.

For brevity, we omit the proof showing that the same logic implies that the augmented internal WMVPF is still the sufficient statistic for optimal policy reforms with partial control of fiscal externalities and an L^2 constraint on internal spending as in Section 3.1; the logic is nearly identical to the above.

B.5 Example with Optimal Choice of p_{N+1}

Consider solving Problem 2 with four possible policies listed in Table 4, but suppose we additionally allow the policymaker to choose which policy p_{N+1} they use to close the budget. One can simply set each of the four policies as p_{N+1} and then use Proposition 2 to determine the solution given this choice of p_{N+1}. Then one just picks the solution which generates the highest welfare gain. For this example problem, the solution is to set policy 4 as the numeraire policy. Direct computation shows that the solution to Problem 2 with policy 4 as the numeraire policy is to increase mechanical spending by \$1 on policy 3 (yielding a welfare gain of 5 and no fiscal externalities) and reduce mechanical spending by \$1 on policy 2 (yielding a welfare gain of -1 and fiscal externalities of 4, which are spent on policy 4 generating a further welfare gain of 24). The total welfare gain from this reform is $5-1+24=28$. In total, the government ends up spending \$5 on net (\$1 on policy 3 and \$4 on policy 4) and raises \$5 in net revenue (reducing mechanical spending by \$1 on policy 2 and then getting an additional \$4 back in fiscal externalities).

Table 4: Example Policies

Policy	$\frac{dW}{dp_i}$	$\frac{dNC}{dp_i}$	$WMVPF_i$	$\frac{dMC}{dp_i}$	FEs
1	5	-1	-5	1	2
2	1	5	$\frac{1}{5}$	1	-4
3	5	1	5	1	0
4	6	1	6	1	0

But now consider Proposition 4 that says we can always do (weakly) better by using the WMVPF maximizing/minimizing policies to spend/raise the same amount of net money. In this example, this would amount to finding the optimal manner to raise \$5 in net revenue and then spend \$5 in net revenue. Proposition 1 tells us that the optimal way to do this is by raising

\$5 in revenue from the minimum WMVPF policy (policy 1) and then spending \$5 (net) on the maximum WMVPF policy (policy 4), which increases welfare by $-5 \times -5 + 5 \times 6 = 55$, which is higher than can be achieved solving Problem 2 even when choosing p_{N+1}. The idea is that the solution from Proposition 1 to raise \$5 in net revenue and then spend \$5 in net revenue entails raising \$10 in fiscal externalities from policy 1 and effectively borrowing against this \$10 in fiscal externalities to *both* spend \$5 upfront on the WMVPF maximizing policy (policy 4) and then spending \$5 upfront on the WMVPF minimizing policy (policy 1). Because this solution effectively entails spending some fiscal externalities on both the WMVPF maximizing policy and the WMVPF minimizing policy, this cannot be achieved via Problem 2 even when choosing p_{N+1}.

B.6 Federal Policies with Large Differences in WMVPF vs. "Augmented Internal WMVPF"

This Appendix discusses three federal programs which have a large difference in their WMVPF (computed assuming the federal legislature controls all fiscal externalities arising from the policy reform) and the "augmented internal WMVPF" (computed assuming the federal legislature does not control fiscal externalities accruing to states). We present estimates of internal costs (mechanical costs plus fiscal externalities accruing to the federal government), external costs (fiscal externalities accruing to state governments), WMVPFs, and "augmented internal WMVPFs" for these three policies in Table 5. Note that if any of these three were included in the choice set of the policymaker in our example problem in Section 5.1, then accounting for state-level fiscal externalities *would* change the policies included in the optimal reform direction by applying Proposition 3 rather than Proposition 1.

First, consider a reform to the Head Start program. The net cost of this policy reform assuming all fiscal externalities accrue to the federal government is $-\$447$, which is equal to the mechanical cost of \$4230 minus the fiscal externality of increased tax revenue arising from increased earnings of \$4677. However, assuming a cumulative average tax rate of 20% consisting of a 2.6% state income tax rate and a 17.4% tax rate for federal income taxes as in Hendren and Sprung-Keyser (2020), this implies that only $\frac{17.4}{20}$ of this \$4677 accrues to the federal government; hence, the internal cost to the federal government is \$4230-\$4677$\times\frac{17.4}{20}$ = \$161 and there is a \$608 fiscal externality accruing to state governments. Thus, from the federal government's perspective (recognizing that about 13% of the tax revenue increase goes to state governments), this program does not pay for itself; as a result, the "augmented internal WMVPF" is 44 whereas the WMVPF is -16.

For the Hope Tax Credit (HTC) and American Opportunity Tax Credit (AOTC), there are three fiscal externalities: increased tax revenue from higher educational attainment, increased government education costs from increased enrollment, and changes in government education costs among the enrolled. Hendren and Sprung-Keyser (2020) assume the average tax rate for these households is 18.9% of which 2.6% goes to state governments. Hence, for the HTC (AOTC) $\frac{2.6}{18.9}$ of the \$0.12 (\$1.24) of increased tax revenue accrues to the state. For the enrollment costs and changes in costs among the enrolled, we assume that 53% of marginal government higher-education costs are borne by states (taken from the fraction of state higher education expenditures in 2015 from the Delta Cost Project). Hence, for the HTC (AOTC), we assume 53% of the \$0.069 (\$0.795) enrollment costs accrue to state governments. Similarly, for the HTC (AOTC), we assume 53% of the \$0.209 (-\$0.082) cost increases among the enrolled accrue to state governments. Adding the federal government share of these costs to the mechanical cost of the HTC (AOTC) of \$0.064 (\$0.202) we arrive at the internal cost in Table 5 (external costs are calculated by adding up the aforementioned fiscal externalities accruing to state governments). Accounting for the fact that some fiscal externalities accrue to state governments leads to a relatively large difference between the WMVPF of 0.69 and the "augmented internal WMVPF" of 1.51 for the HTC. Similarly, the WMVPF for the AOTC is -4.86 whereas the "augmented internal WMVPF" is -2.93.

Table 5: Augmented Internal WMVPFs for Selected Federal Policies

Program	Average Income	Net Cost	WTP	$\bar{\eta}$	WMVPF	Internal Cost	External Cost	Augmented Internal WMVPF
Hope Tax Credit, Joint Filers at Phase End	34497	0.22	0.52	0.29	0.69	0.09	0.13	1.51
American Opportunity Tax Credit, Single Filers at Phase Start	29876	-0.33	4.72	0.33	-4.86	-0.53	0.21	-2.93
Head Start Introduction	26760	-447	18708	0.37	-16	161	-608	44

Note: This Table displays the "augmented internal WMVPF" from Proposition 3 for three policies. Data on average income, net cost, and WTP are taken from Hendren and Sprung-Keyser (2020). Data on costs accruing to the federal government (internal costs) and costs accruing to state governments (external costs) are calculated by the authors using data in Hendren and Sprung-Keyser (2020). We assume $\bar{\eta} = 10000 \times \text{Average Income}^{-1}$ (where the factor of 10000 is for readability). We assume the numeraire policy, p_{N+1}, for states is a lump-sum transfer which generates no behavioral responses. The WMVPF for this lump-sum transfer equals $\bar{\eta}_{N+1} MVPF_{N+1}$ where $MVPF_{N+1} = 1$. We assume $\bar{\eta}_{N+1} = 10000 \times 97360^{-1}$ where 97360 is the average U.S. household income in 2016.

B.7 State Level Policies

In this Appendix we briefly illustrate how a state-level policymaker (e.g., a state legislature) might use our framework to devise optimal policy reforms recognizing that some share of fiscal externalities arising from policy reforms accrue externally (i.e., to the federal government or to other state governments). We assume that the state legislature can control how fiscal externalities that accrue to their state are spent but cannot control how fiscal externalities accruing to the federal government are spent (we assume there are no cross-state spillovers for simplicity). Moreover, we assume the state legislature only cares about choosing policies to benefit residents of their state. Hence, if state-level policies generate fiscal externalities for the federal government, we assume the state legislature only cares about the impact that this federal spending has on residents of their state.

Suppose a state legislature is considering a budget-neutral policy reform of three education policies: community college tuition reductions, K-12 school expenditures, and spending on public universities. The state legislature cannot control the share of fiscal externalities that accrue to the federal government so they solve Problem 3. We present costs and benefits for these three policies in Table 6 from Hendren and Sprung-Keyser (2020). We separate out the "internal costs" of the reform, which consist of the state-level mechanical costs as well as the fiscal externalities which accrue to the state and the "external costs", which consist of fiscal externalities that accrue to the federal government. For these three policies there are two fiscal externalities: increased income tax revenue and increased educational expenditures. For the income tax revenue increases, we follow Hendren and Sprung-Keyser (2020) and assume that the state-level income tax rate is 2.6% on average; we use this tax rate and the estimated causal impact on incomes reported in Hendren and Sprung-Keyser (2020) to calculate state-level income tax revenue increases. Federal tax revenue increases are computed using the average federal tax rates and the estimated causal impact on incomes in Hendren and Sprung-Keyser (2020). Finally, we assume that 53% of marginal government spending on universities is borne by the state and 47% is borne by the federal government (taken from the fraction of state higher education expenditures in 2015 from the Delta Cost Project).

The key takeaway from Table 6 is that for a state-level policymaker, accounting for the fact that only a fraction of fiscal externalities are under their control can have a substantial impact on the direction of optimal policy reform. For instance, the K-12 school expenditure policy pays for itself when accounting for all federal fiscal externalities but does not pay for itself from the perspective of the state legislature; similarly, the "augmented internal WMVPF" is substantially lower than the WMVPF for the community college tuition reduction because most of the fiscal externalities accrue to the federal government. Thus, while a hypothetical social planner that internalized all fiscal externalities would reform these three policies by raising revenue from K-12 school expenditures (which simultaneously increases welfare because this program pays for itself) and then spending this money on community college tuition reductions, a state-level legislature that cannot control how fiscal externalities accruing to the federal government are spent would instead reduce spending on public universities and increase spending on K-12 expenditures.

Table 6: Augmented Internal WMVPFs for Selected State-Level Policies

Program	Average Income	Net Cost	WTP	$\bar{\eta}$	WMVPF	Internal Cost	External Cost	Augmented Internal WMVPF
Community College Tuition Changes	21215	27	9322	0.47	**165**	2037	-2012	2.16
K-12 School Expenditures	52246	-1.03	8.78	0.19	**-1.63**	0.72	-1.75	**2.35**
Spending at Colleges from State Appropriations	32017	793	3173	0.31	1.25	1110	-317	**0.89**

Notes: This Table displays WMVPFs and augmented internal WMVPFs for three state-level education policies. Data on average income, net cost, WTP, and mechanical costs are taken from Hendren and Sprung-Keyser (2020). $\bar{\eta} = 10000 \times \text{Average Income}^{-1}$ (where the factor of 10000 is for readability). The numeraire policy p_{N+1} for the federal government is assumed to be linear income taxation. The WMVPF for linear income taxation equals $\bar{\eta}_{N+1} MVPF_{N+1}$. The logic of Appendix F.II of Hendren and Sprung-Keyser (2020) implies that $MVPF_{N+1} = \frac{1}{1-\frac{T'}{1-T'}\xi^u}$ where we assume $T' = 0.2$ is the average marginal tax rate (taken from Appendix G of Hendren and Sprung-Keyser (2020)) and $\xi^u = 0.3$ is the uncompensated labor supply elasticity (Saez et al., 2012). $\bar{\eta}_{N+1} = 10000 \times 97360^{-1} \times 0.05$ where 97360 is the average U.S. household income in 2016 and the factor of 0.05 assumes that the state policymaker only cares about residents of their state and that their state makes up 5% of the country's population. Data on costs accruing to the state government (internal costs) and costs accruing to the federal government (external costs) are calculated by the authors using data in Hendren and Sprung-Keyser (2020).

C Appendix: Calculations for Table 3

Note in this section, "pesos" refers to 2010 Mexican pesos.

C.1 CCT Estimates

Our estimates for the CCT WMVPF and WNSB combine empirical results from a number of papers exploring the impacts of the CCT program Progresa in rural Mexico. First, to obtain mechanical costs, Parker and Vogl (2023) note that 11.6 billion pesos are spent on education transfers. For every peso transferred, 0.089 pesos of administrative costs are generated. From Parker and Vogl (2023), we can also back out the number of individuals exposed to the Progresa program: 726,212.[44] Thus, the per-person transfer is 15,973 pesos and the per-person

[44]On page 2801, Parker and Vogl (2023) calculate the average discounted lifetime earning gain for each women exposed to Progresa as children. They multiply this gain by the number of women exposed to Progresa to

administrative cost is 1,422 pesos.

Parker and Vogl (2023) then estimate that for those exposed to the program, annual earnings increased on average by 5,964 pesos for men and by 2,980 peso for women.[45] We use the same procedure as Parker and Vogl (2023) to translate this estimate into a discounted lifetime earnings gain except we use a discount rate of 0.11 reflecting the long-term government bond yield in Mexico at the time of program implementation whereas Parker and Vogl (2023) use a discount rate of 0.02.[46] We estimate per recipient, lifetime discounted earnings increase by 23,855 pesos for men and 11,919 pesos for women.[47] To estimate fiscal externalities that result from this earnings gain, we use the following tax rates: 16% consumption tax (we assume all earnings are spent in the period in which earned), a 1.92% income tax rate (this is the tax rate in the lowest bracket), and a 35% payroll tax.[48] Note, we assume that only 50% of consumption is actually subject to the consumption tax due to the large share of informal consumption in Mexico among poor households (Bachas et al., 2023), and we assume that only those working for a wage (as opposed to agriculture) are subject to payroll and income taxes. To estimate the share of individuals working for a wage, we use Table 1 of Parker and Vogl (2023) which shows that 58% of working men work for a wage and 67% of working women work for a wage. Thus, fiscal externalities are equal to: $23855 \times [0.58((1-0.0192) \times 0.16 \times 0.5 + 0.0192 + 0.35) + 0.42(0.16 \times 0.5)] = 6995$ pesos for men and $11919 \times [0.67((1-0.0192) \times 0.16 \times 0.5 + 0.0192 + 0.35) + 0.33(0.16 \times 0.5)] = 3890$ pesos leading to an average fiscal externality of 5442 pesos.[49]

Thus, to summarize costs for the CCT, we have a 15,973 peso transfer, a 1,422 peso administrative cost, and a 5,442 peso fiscal externality. We normalize these costs to be expressed per 1 peso of mechanical costs (i.e., the transfer plus the administrative cost). Thus, per 1 peso of mechanical cost, we have a 0.92 peso transfer, 0.08 peso administrative cost, and a 0.31 peso fiscal externality.

We now move onto estimating the WTP for a 0.92 peso CCT. For those households who are already sending their children to school, they value the transfer peso-for-peso. For those families induced to send their children to school in order to receive the grant, we assume they do not value

calculate an aggregate discounted lifetime earning gain. We can divide this aggregate gain by the individual gain to calculate the number of women exposed to the program. We then assume that the number of exposed men is the same as the number of exposed women to get the total number of individuals exposed.

[45]The average earnings increase for men and women come from averaging estimates across columns (1),(2),(3) and columns (5),(6),(7) of Table 3 in Parker and Vogl (2023), respectively; note these estimates are for monthly earnings. We multiply by 12 to convert to annual estimates.

[46]0.11 is the 10 year Mexican bond rate in July 2001 (https://fred.stlouisfed.org/series/IRLTLT01MXM156N).

[47]Following Parker and Vogl (2023), we assume that children start working 7 years after the start of the program and work for 45 years and that there is no earnings growth. Thus, we perform the following two calculations: $0.89^7(1-0.89^{45})/(1-0.89) \times 5964 = 23855$ and $0.89^7(1-0.89^{45})/(1-0.89) \times 2980 = 11919$. These calculations yield a present discounted value of lifetime earnings gain expressed in 2010 Mexican pesos.

[48]See https://mx.icalculator.com/income-tax-rates/2024.html.

[49]Note, we assume that the consumption tax is paid on after-income-tax income.

the transfer at all (i.e., they are all indifferent between receiving the grant and sending their child to school and not receiving the grant and not sending their child to school). Bergstrom and Dodds (2021) estimate that 11% of the Progresa recipients are marginal (i.e., would not send their child to school if the grants were removed).[50] Thus, for 0.92 pesos transferred, the average WTP across Progresa recipients is $0.82 = 0.92 \times 0.89$. We assume the households do not value the extra lifetime earnings by the envelope theorem. Thus, the average willingness to pay for a 0.92 peso CCT is 0.82 pesos. Finally, we normalize the average welfare weight on Progresa households to be 1. Thus, the MVPF and WMVPF for the CCT are both equal to $MVPF_{CCT} = WMVPF_{CCT} = \frac{0.82}{0.92+0.08-0.31} = 1.19$.

C.2 UCT Estimates

We consider a counter-factual policy in the Progresa setting: a cash transfer offered to the same set of eligible households without the conditions that children must attend school. We assume administrative costs of a UCT are half of those of the CCT (CCTs will incur greater administrative costs due to the monitoring of children in school). Thus, for every peso transferred via a UCT, we assume 0.045 pesos of administrative costs are generated. Normalizing the mechanical cost to be 1 peso, we get 0.96 pesos in UCT and 0.04 pesos in administrative costs.

Next, we assume a 1 peso UCT in this setting generates a fiscal externality that is 44% as large as a 1 peso CCT (Bergstrom and Dodds (2021) estimate that a UCT in this setting would generate a school enrollment effect that is 44% as a CCT in this setting). Given that a 0.92 peso CCT generates a 0.31 fiscal externality, we estimate that a 0.96 peso UCT generates a $0.31 \times 0.44 = 0.14$ peso fiscal externality.

For WTP, we assume households value the UCT peso-for-peso. Thus, a 0.96 peso UCT generates a WTP of 0.96 pesos (we again invoke the envelope theorem such that any changes in behavioral induced by the UCT have no first-order impacts on household utility). Finally, because we assume the set of households who are offered the UCT are the same as the CCT, the average welfare weight on UCT households is also 1. Thus, the MVPF and WMVPF for the UCT are both equal to $MVPF_{UCT} = WMVPF_{UCT} = \frac{0.96}{0.96+0.04-0.14} = 1.12$.

C.3 WMVPF of a Consumption Tax

Consider a consumption tax of τ and two types of consumption: formal sector and informal sector consumption, c^f, c^i. Only formal sector consumption is taxed. Let the price (exclusive of taxes) of formal seector consumption be given by p^f and the price of informal sector consumption be given by 1. Assuming 100% pass-through of taxes onto prices, a household's WTP for a change in the consumption tax of $\Delta\tau$ is equal to $-\Delta\tau p^f c^f$.

[50] Table 1 of Bergstrom and Dodds (2021) shows that enrollment increases from 68.3% to 76.7% with the grants. Thus, of the 76.7% enrolled, 11% are marginal.

Now consider the change in government revenue from a $\Delta\tau$ change in the consumption tax. The change in revenue is equal to $\Delta R = \Delta\tau\left(p^f c^f + \tau\frac{d(p^f c^f)}{d\tau}\right)$. The first term represents the mechanical increase in revenue and the second term represents the change in revenue resulting from fiscal externalities. Let $\xi = \frac{d(p^f c^f)}{d(1+\tau)}\frac{1+\tau}{p^f c^f} = \frac{d(p^f c^f)}{d\tau}\frac{1+\tau}{p^f c^f}$. We get the change in revenue is equal to $\Delta R = \Delta\tau\left(p^f c^f + \xi\frac{\tau}{1+\tau}p^f c^f\right)$. The MVPF of changing the consumption tax by $\Delta\tau$ is therefore given by $\frac{-\Delta\tau p^f c^f}{-\Delta\tau\left(p^f c^f + \xi\frac{\tau}{1+\tau}p^f c^f\right)} = \frac{1}{1+\xi\frac{\tau}{1+\tau}}$ where ξ measures the percentage change in formal sector spending (net of taxes) when $1+\tau$ increases by 1%. We set $\xi = -1$ consistent with Bachas et al. (2023). The consumption tax in Mexico is 16%. Thus, the MVPF of a consumption tax is $MVFP_\tau = 1/(1 - 0.16/0.84) = 1.24$.[51]

To get the average welfare weight for the consumption tax, we follow Section 5.1 and set the average welfare weight $\bar{\eta} = \bar{z}^{-\gamma}$ where $\bar{z}$ represents the average income of policy beneficiaries and $\gamma = 1$. Average per-capita income for Mexico for 2010 was \$9,823 (current USD; World Bank (2024)). Converting this to 2010 Mexican pesos, we get a per-capita income of 91,730 pesos.[52] The average income of Progresa recipients is 10,020 pesos per-capita.[53] Because we normalize the average welfare weight on UCT/CCT households to be 1, the average welfare weight for those impacted by the consumption tax is equal to $91730^{-1}/10020^{-1} = 0.11$. Hence, the WMVPF for a consumption tax equals 0.11×1.24.

C.4 Calculating the WNSB for the CCT and UCT

Finally, we calculate the WNSB for the CCT and UCT as follows:

$$WNSB_{CCT} = \frac{1 \times 0.82 + 0.11 \times 1.24 \times 0.31}{1} = 0.86$$

$$WNSB_{UCT} = \frac{1 \times 0.96 + 0.11 \times 1.24 \times 0.14}{1} = 0.98$$

Note, that the government gets to spend the fiscal externalities generated by each policy on reducing the consumption tax next generation (and does not have to finance the mechanical spending on the CCT or UCT as we are considering an example where the government is given a grant of \$1,000,000 and we assume that this is a small grant so that we can just scale both the numerators and denominators of the WMVPFs and WNSBs by 1,000,000 without having to worry about curvature in the welfare or cost functions).

[51] The consumption tax change occurs for the next generation; hence, we need to discount the WTP and the net cost back to when Progresa was implemented. However, because this discounting occurs on both the numerator and denominator, it will not affect $MVPF_\tau$.

[52] We use an average exchange rate of \$1 USD = 17.13 pesos for 2024 and use https://www.inflationtool.com/mexican-peso to deflate 2024 values to 2010 values.

[53] See Table 1 of Parker and Vogl (2023): average monthly earnings per capita is 835 pesos per capita (average across men and women). Multiplying by 12 gives 10,020 pesos per-capita.

D Appendix: Tables

Table 7: Augmented Internal WMVPFs for Selected Federal Policies

Program	Average Income	Net Cost	WTP	$\bar{\eta}$	Internal Cost	External Cost	Augmented Internal WMVPF
EITC Expansion to Adults without Dependents (Paycheck Plus)	10534	1074	1070	0.95	1080	-5.85	0.94
Disability Insurance Benefit Generosity	26445	1.04	1.00	0.38	1.04	0.01	0.36
Top Income Tax Reductions	737931	0.86	1.00	0.01	0.88	-0.02	0.02
Higher Education Tax Deductions	28411	-460	485	0.35	-187	-273	-1.06
Medicaid Expansions to Young Children	48096	-491	3681	0.21	-409	-83	-1.89
Housing Vouchers (Moving to Opportunity Experiment)	16331	-9215	69601	0.61	-8302	-913	-5.15

Note: This Table displays the "augmented internal WMVPF" from Proposition 3 for the six policies discussed in Section 5.1. Data on average income, net cost, and WTP are taken from Hendren and Sprung-Keyser (2020). Data on costs accruing to the federal government (internal costs) and costs accruing to state governments (external costs) are calculated by the authors using data in Hendren and Sprung-Keyser (2020). For all of of the programs other than the higher education tax deductions, the only fiscal externality we assume accrues to states is income tax revenue, which, following Hendren and Sprung-Keyser (2020), we assume is 2.6% on average. For the higher education tax deductions, we also assume that 53% of marginal government higher-education costs are borne by states (taken from the fraction of state higher education expenditures in 2015 from the Delta Cost Project). We assume $\bar{\eta} = 10000 \times \text{Average Income}^{-1}$ (where the factor of 10000 is for readability). We assume the numeraire policy, p_{N+1}, for states is a lump-sum transfer which generates no behavioral responses. The WMVPF for this lump-sum transfer equals $\bar{\eta}_{N+1} MVPF_{N+1}$ where $MVPF_{N+1} = 1$. We assume $\bar{\eta}_{N+1} = 10000 \times 97360^{-1}$ where 97360 is the average U.S. household income in 2016.