Welfare Analysis Meets Causal Inference: A Suggested Interpretation of Hendren

Amy Finkelstein*

March 2019

Abstract

In a pair of interconnected, important and impenetrable papers, Nathan Hendren has provided a framework for translating estimates of the causal effects of policies into welfare analyses of these policies. In this brief note, I describe the framework - which Hendren has named "The Marginal Value of Public Funds" (MVPF) - and how it can be used for empirical public finance welfare analysis. I also discuss how the MVPF relates to "traditional" public finance welfare analysis tools such as the marginal excess burden (MEB) and marginal cost of public funds (MCPF). Finally, I describe several recent empirical applications as a way of further illustrating and clarifying the approach.

Keywords: Welfare Analysis; Marginal Value of Public Funds; Fiscal Externality; Envelope Theorem; Causal Effects; Hendren

*I am grateful to Matt Notowidigdo, Ben Olken, and numerous seminar audiences for the confusion that prompted me to write this, to Alan Auerbach, Henrik Kleven and Sammy Young for helpful comments, and of course to Nathan Hendren for both the written work and the countless hours spent patiently answering (and re-answering) my many questions on it.

1 Introduction

Economists have made remarkable progress over the last several decades developing empirical techniques that can provide compelling evidence of causal effects - the so-called "credibility revolution" in empirical work (Angrist and Pishke 2010). For the most part, we have labor economists and applied econometricians to thank for this progress. Yet many of the applications have been in the domain of public finance: estimating the causal effects of *public policies.* To public finance economists, however, such descriptive estimates of the impacts of policies are only intermediate inputs. The ultimate question of interest is: what is the *welfare* impact of the public policy?

Until recently, there had been relatively little effort to harness the fruits of the "credibility revolution" for the public finance goal of welfare analysis. Instead, we in the empirical public finance community, guided by the marginal excess burden (a.k.a marginal deadweight loss) approach to welfare analysis, have found ourselves struggling to separate (non-distortionary) income effects of policies from their (distortionary) substitution effects. As Goolsbee (1999) has lamented "The theory largely relates to compensated elasticities, whereas the natural experiments provide information primarily on the uncompensated effects." Alternatively, guided by the "marginal cost of public funds" approach to welfare analysis, we have found ourselves working hard to estimate and monetize the non-monetary benefits of a policy change (e.g. through improved test scores, better health, or reduced pollution) only to compare these benefits of the policy to "the cost" of the policy - defined as expenditures on the policy multiplied by 1 plus "the" marginal cost of public funds - which everyone "in the know" knows to be 0.3. Our field may have been there for the revolution, but our core focus on welfare analysis seemed to be reaping little direct benefits from it.

Fortunately, there is light at the end of the empirical welfare tunnel. Hendren (2016) proposes a simple and transparent way to map from empirical estimates of causal effects of a public policy change to welfare analysis of that policy change. The key object for this mapping is what he calls the **marginal value of public funds** ($MVPF$).

The MVPF is defined simply as the ratio of the marginal benefit to the marginal cost of the policy. Equivalently, the MVPF is the ratio of the beneficiaries' willingness to pay for the increase in expenditure out of their own income to the cost to the government of the increase in expenditure per beneficiary:

$$MVPF = \frac{"Benefit"}{"Cost"}. \tag{1}$$

In this brief note we walk through the construction of the MVPF for a number of classic public finance welfare questions. We begin with a simple example: welfare analysis of a targeted cash transfer to a particular income group; we then show how the same approach can be naturally extended for welfare analysis of other policies such as spending on a in-kind transfer program, spending on a public good, or changes in corrective (Pigouvian) taxation.

We also discuss how to use the MVPF. It is most valuable for ranking two policies that target the same income group - whichever has a higher MVPF is preferable since it achieves a given outcome at lower cost. The MVPF is thus particularly useful for what is arguably the key comparative advantage

of economists in the policy sphere: faced with a given "societal" objective, which policy achieves it at the lowest resource cost? Of course, it is rare to have two policies that target exactly the same income group. Fortunately, Hendren (2017) extends the usefulness of the MVPF approach by showing how the MVPF of a policy that transfers resources to a specific income group can be usefully compared to the MVPF of a tax cut to this income group - i.e. the cost of transferring resources to that income group through the income tax code. He terms the MVPF of a tax cut to income group i the "efficient welfare weight" for that income group; policies that transfer to that income group with an MVPF that is higher than the efficient welfare weight represent potential Pareto improvements.

Another use of the MVPF is that one can compare the MVPF across policies with different distributional incidences and back out the implied relative social welfare weights such that one would be indifferent between those two policies. Importantly, the MVPF is not sufficient to determine *whether* either or both policies "should" exist - economics is not sufficient to resolve that question. In other words, the NBER working paper prohibition against "statements regarding which policies should (or should not) be adopted"[1], encodes a fundamental and important recognition of the limits of economic analysis. Still, the MVPF can provide important inputs into this type of analysis by quantifying the tradeoffs involved.

In Section 3 we endeavor to answer some common and natural questions about the approach - specifically how it relates to the "traditional" public finance welfare tools of marginal excess burden (MEB) and marginal cost of public funds (MCPF); we also discuss a key assumption behind the MVPF approach, namely the use of the envelope theorem, which poses challenges for analysis that involves non-marginal changes or non-optimizing agents. Finally, since a key value of the approach is the guidance it provides for how to map causal estimates of impacts of policy - where there has been enormous empirical progress and activity - into welfare analyses of those policies, in Section 4 we illustrate how the approach has been applied in some recent empirical papers. The last section briefly summarizes.

2 The Marginal Value of Public Funds (MVPF)

2.1 An initial application: the MVPF of an increase in public expenditure on a targeted cash transfer

We can (and will) apply the MVPF approach to analyze a number of different types of policies. For simplicity, we start with an analysis of the MVPF of an increase in public spending on a targeted cash transfer. This could be, for example, an increase in spending on a means tested cash welfare program like Temporary Assistance to Needy Families (TANF), or an increase in spending on a means tested tax credit such as the Earned Income Tax Credit (EITC). Our object of interest therefore is the MVPF of a \$1 increase in cash to a given income group.

Let's start with the numerator of the expression for the MVPF in equation (1) - the benefits from the cash transfer. This is, equivalently, the willingness to pay for \$1 more of the cast transfer. In considering willingness to pay, it is useful to distinguish between two classes of recipients. The

[1] http://papers.nber.org/wpsubmit/wp_submit.html

first is the infra-marginal recipients (whom we'll call I) who were already receiving the cash transfer policy; for them the \$1 transfer is valued at \$1. The second is the marginal recipients (whom we'll call M) who change their behavior in response to the change in policy and thus become newly eligible for the transfer; for example, they may reduce their labor supply in order to mean the means-tested eligibility threshold for TANF, or they may enter the labor force (increasing their labor supply) in order to qualify for the tax subsidy to wages for low-income workers through the EITC. For them, the **envelope theorem** plays a key role: it tells us that - if individuals are optimally choosing their behavior - the marginal individual who changes her behavior in response to a marginal change in the policy are, to first order, indifferent to their behavioral response. Thus the willingness to pay for the \$1 increase in the cash transfer is simply \$1 times the number of infra-marginal beneficiaries, or $\$I$. Note that, thus far we have not needed to estimate any causal effects.

Now let's consider the denominator of the MVPF in equation (1) - the cost of the cash transfer. What is the cost of this increase in program generosity by \$1? Again, it is useful to think of two different classes of cost. The first is what Hendren calls the "**mechanical cost**" of the policy - i.e. increased government expenditures due to the policy in the absence of any behavioral response. In this example the mechanical cost is simply \$1 for each infra-marginal recipients, or $\$I$ total. The second is what Hendren calls the "**fiscal externality**" from the policy; this is the impact of any behavioral response to the policy on the government's budget outlays. For example, if individuals reduce their labor supply in order to become eligible for TANF, this will in turn reduce income tax revenue collected by the government on earnings; this reduction in income tax revenue (along with the increased expenditure on TANF for the marginal individuals who reduce their labor supply to become eligible for TANF in response to the policy change) needs to be included in the fiscal externality. Similarly, if individuals enter the labor force in order to become eligible for the EITC, this will increase government expenditures on the EITC above what it would have been in the absence of this behavioral response. More generally, the fiscal externality must account for the full impact of any behavioral response on tax revenue - this could include, for example, changes in eligibility for other public programs and hence government expenditures on that program, changes in consumption patterns and hence sales taxes, changes in public health care spending through Medicare and Medicaid (if the program affects health) etc. etc. These behavioral responses to the program may include those of both marginal and infra-marginal recipients; an infra-marginal recipient who was already eligible for the program might - due to the income effect - change some other behavior - such as consumption decisions - that affects government revenue.

Crucially, we cannot invoke the envelope theorem and 'ignore' these behavioral responses because the agents making (by assumption, privately optimal) behavioral changes in response to the policy change do not internalize the *external effects* of the policy on the government budget. In other words, the causal effect on the government's budget matters because of a fiscal externality. The envelope theorem guarantees that behavioral responses by individuals do not affect their utility directly; however, when prices do not reflect their resource costs, behavioral responses impose a cost on those bearing the difference between the prices faced by the individual and their resource costs.[2]

[2]By the same token, if the behavioral responses to the policy have external effects on other individuals besides

For comparing across policies, it is useful to normalize the MVPF by the mechanical cost of the policy ($\$I$) so that the MVPF is reported per \$1 increase in mechanical expenditure per infra-marginal beneficiary of the policy. In this case the MVPF of the \$1 increase in cash benefits is:

$$MVPF^{\$} = \frac{1}{1 + FE} \tag{2}$$

where FE denotes the impact on the budget (i.e. fiscal externality) of the policy per dollar increase in the mechanical expenditure per infra-marginal beneficiary of the policy. Note that the fiscal externality may be positive or negative; policies that have a positive net effect on the government budget (e.g. by improving health and reducing public spending on health care) reduce the effective cost of the policy below the mechanical cost; conversely, policies that have a negative net effect on the government budget increase the effective cost of the policy above the mechanical cost.

It is here, with the concept of a policy's fiscal externality, that the applied econometrics literature on causal inference connects with the welfare analysis that is at the core of public finance: welfare analysis requires estimates of the *causal effects* of the policy on the government's budget. This is familiar territory. An enormous empirical literature studies the impact of various programs on taxable income, and participation in (and expenditures from) other public programs, all of which are potentially part of the fiscal externality. It is incumbent upon the researcher to try to estimate any a priori plausible channels for this fiscal externality.

A note on priority

It is apparently a natural tendency to attribute an idea to the source from which one learned it,[3] and I plead guilty on this account. I have written this note as a clarification of Hendren (2016)[4] because it was through that paper - and related conversations with Hendren - that I first became aware of the empirical attraction of this approach for the type of empirical welfare analysis problems I have often grappled with in my research and teaching.

However neither I - nor Hendren (2016) - mean to claim that the core ideas here are new. Indeed, Hendren (2016) contains an extensive discussion of the intellectual history, noting (see footnote 6, p.83) "to the best of my knowledge, this measure of the marginal value of public funds (MVPF) was initially proposed by Mayshar (1990)" and highlighting (see page. 54) closely related work by Slemrod and Yitzhaki (1996, 2001), Kleven and Kreiner (2006), and others. Likewise - the two key concepts underlying the MVPF - the fiscal externality and the envelope theorem are also not new. The fiscal externality is a property that has been present in our intellectual tradition since at least Ramsey (1927), although its crystallization and importance has become apparent more recently (see e.g. Feldstein 1999, Saez 2004, Kleven and Kreiner 2005, Chetty and Saez 2010). Likewise the key insight behind the envelope theorem (Envelope, 1845), has been used in almost all previous empirical

recipients of the policy, these would also have to be taken into account; we cover this below when we consider cases of "multiple beneficiaries"; for now, for simplicity, we assume that government policy is the only pre-existing distortion, and hence the only source of potential "external effects".

[3] I learned this from Scott Stern, so I cite him here

[4] with full appreciation of the inherent irony of this endeavor: In the working paper version of this paper, Hendren (2013) notes "Relative to [the existing] literature, the primarily contribution of this paper is a clarification."

welfare analyses, including Harberger (1964). My eagerness to clarify and illustrate the approach lies not in its novelty but in its utility.

2.2 Efficient Welfare Weights and the Use of the MVPF

Now that we can estimate the MVPF of a targeted cash transfer what do we do with it? Is an MVPF of a policy of 0.8 or 1.3 "good" or "bad"? How do we actually "use" the MVPF for welfare analysis? This is fundamentally a question of the relevant benchmarks against which to evaluate the MVPF.

The most straightforward use of the MVPF is for comparing policies that have the same distributional incidence. If two different policies can be used to transfer resources to group i, the policies with the higher MVPF is preferred - it represents a lower cost way of enacting this transfer. This represents a key attraction of the MVPF framework: it allows one to do comparative welfare analyses across policies without having to take a stand on how exactly the program is funded (which is rarely known or clear) and without having to take a stand on the curvature of either the social welfare function or the individual utility function (which is also usually unceratin). It is always more comfortable as an economist to answer the question: given the goal of redistributing X amount to population Y, can this be more efficiently accomplished through instrument Z_1 or Z_2 than to ask whether a specific policy Z_1 is "worth its costs". The latter requires a series of normative / ideological assumptions / preferences in the form of the social welfare function and individual utility function. Imagine however two policies that attempt to transfer to the same population - for example, expanding the earned income tax credit or expanding cash TANF benefits for the same population. A comparison of estimates of the MVPF of these two policies would allow us to answer the question: if we want to transfer to this population, is it better to do it through the earned income tax credit or TANF? The answer would be: whichever of the two policies has a higher MVPF.

Of course, it is rare that two policies target exactly the same populations. Eligibility requirements or thresholds may differ as may the take-up decision among the eligible population. Hendren (2017) develops a useful tool for comparing the MVPF of policies that affect different populations: each can be compared to the benchmark of the MVPF of the tax schedule for the relevant affected population. The tax schedule provides a useful benchmark because it spans what is arguably the most important form of heterogeneity in the population (differences in earnings potentials). One natural benchmark is therefore to compare the MVPF of a policy that affects a population of a given income level y_i to the MVPF of a targeted, \$1 income tax cut for individuals at income y_i. Hendren (2017) defines the MVPF of a \$1 income tax cut to those of income group i as the "**efficient welfare weight**" ν_i. The efficient welfare weight ν_i is thus the cost of transferring \$1 to individuals of income i through an income tax cut to them. The formula for ν_i is given by equation (2); with a tax cut, the FE in the denominator of equation (2) is captured by the impact of the tax cut on tax revenue via any behavioral response to the tax cut.

These efficient welfare weights provide a way of comparing the MVPF of a policy that transfers resources to income group $y_i (MVPF_i)$ to the $MVPF$of a policy that transfers resources to income group y_j $(MVPF_j)$, because they give us a way of "projecting" both policies onto the same population. Given that we can make transfers from income group $y_i$to income group $y_j$through the tax code, we

can "turn" the policy targeted at income group i into a comparable MVPF for a policy targeted to income group j by multiplying the $MVPF_i$by the ratio of efficient welfare weights ($\nu_{j/}\nu_i$). Thus we prefer policy i to policy j if $(MVPF_i)\frac{\nu_j}{\nu_i} > MVPF_j$. Policies that target the same income group naturally have the same efficient welfare weights, which is why the comparison between them becomes a direct comparison of which has a higher MVPF.

The efficient welfare weights also provide a direct way of evaluating a single policy. If it targets income group i, the relevant benchmark is to compare to MVPF of \$1 of increased spending on this policy to the MVPF of a \$1 tax cut to income group i - in other words, to the efficient welfare weight ν_i, which provides a benchmark for Pareto improvements: If the $MVPF$ exceeds ν_i, there is a potential Pareto improvement: cut income taxes on group i by \$1 (with MVPF of ν_i) and transfer it back to the group through the policy with $MVPF> \nu_i$. If there were no distortions from taxation or the social welfare weights for all individuals affected by the policy were 1, the natural benchmark would be to compare the MVPF to 1. However, when these conditions fail (as they almost always will), the natural benchmark is to compare the MVPF of the policy to the MVPF of the "next best alternative"; Hendren (2017) argues that the MVPF of targeted tax cuts through the income distribution provides a useful natural alternative for comparisons: policies whose MVPF is above the efficient social welfare weight for the income group affected should be expanded, those that are below should be contracted. In other words, the efficient welfare weights allow the analyst to search for potential feasible Pareto improvements.[5] Of course, a key assumption underlying this exercise is that conditional on income, society attaches the same social welfare weights to all individuals, so that we do not have to consider distributional consequences across individuals (of say different health) conditional on income.

Hendren (2017) provides empirical estimates of the distribution of ν_i given the existing tax schedule, tax data on the distribution of income, and existing empirical estimates of taxable income elasticities at different income levels. For example, he estimates that the MVPF of a \$1 tax cut at the high end of the income distribution is about \$1.5 - a dollar tax cut for the rich generates \$1 in benefits and a cost of \$0.65. However, at the bottom of the income distribution, he estimates that the MVPF of a \$1 tax cut to the poor (through the EITC) has an MVPF of about \$0.88 - again it generates \$1 in benefits but the cost is \$1.15.[6] Put differently, transferring (through the income tax) \$1 from the top of the income distribution can generate about \$0.65/1.15 = \$0.57 of welfare to someone at the bottom of

[5]Hendren (2017) notes that these efficient welfare weights provide a way of implementing the Kaldor-Hicks test for efficiency that accounts for the distortionary cost of taxation. Kaldor and Hicks proposed a notion of "economic efficiency" based on the sum of individuals' willingness to pay for a policy; if this were negative, a set of individual-specific lump-sum transfers could reach an allocation that would be Pareto superior to the alternative policy (Hicks 1940). In other words, everyone would prefer the transfers to the policy. Moreover, if this sum is positive, Kaldor (1939) argued that the winners could likewise compensate the losers using individual-specific lump-sum transfers. Of course, in practice such compensating, individual-specific lump-sum transfers are not feasible, and any attempt to enact a set of transfers to achieve a Pareto improvement involves distortionary costs, a point that Kaldor and Hicks recognized. Hendren's approach provides the set of "efficient" social welfare weights that modifies the Kaldor-Hicks experiments to incorporate the distortionary cost of redistribution through the tax and transfer system. The efficient social welfare weight at an income level \$y is equal to the marginal cost of providing \$1 of welfare (cash) to individuals earning \$y.

[6]Note that both tax cuts - at the top or bottom of the income distribution - generate increased labor supply. However, this increased labor supply produces a negative fiscal externality at the high end of the income distribution and a positive fiscal externality at the bottom of the income distribution. This is because at the high end, marginal tax rates are positive, while at the bottom end individuals face negative marginal tax rates in the phase-in range of the EITC. We discuss this in more detail in Section 4 below.

the distribution. Conversely, transferring $1 from the bottom of the income distribution can generate around $1.15/0.65=$1.77 of welfare to those at the top of the income distribution. Note that a lower MVPF for a tax cut to the poor than to the rich is consistent with an optimal tax system in the presence of a social welfare function that places greater weight on the marginal value of resources for the poor than the rich (e.g. a utilitarian social welfare function).

Thus far we have talked about using the MVPF to make comparative evaluations *across* policies. That of course is a different question for whether a policy the targets a given income group and has a given MVPF "should" be enacted. Economics alone cannot provide an answer to this question - it depends on social preferences. Therefore it requires us to take a stand on a social welfare function. Economists have never had a comparative advantage at specifying societal preferences. Given a set of preferences however, the MVPF of a policy targeted at a given population can be used to say whether it "should be done." For example, consider a hypothetical policy of a policy that takes a $1 from group B and spends it on a policy directed at group A. The ratio of the MVPF of the policy aimed at group A to the policy aimed at group B reveals the implicit welfare weights at which society would be indifferent to this policy. If one is willing to take a stand on the social welfare function, we can compare those implicit welfare weights to actual social welfare weights. For example, suppose the ratio of the MVPF of a policy that transfers to group A relative to taking away from group B is $0.8 and group A has a consumption level that is 50 percent of the consumption of group B. Suppose further that we have a utilitarian social welfare function - so that social welfare is the (unweighted) sum of the individual utilities of everyone in society, and the individual utility functions exhibit constant relative risk aversion over consumption $u(c_i) = \frac{c_i^{1-\sigma}}{1-\sigma}$ with a coefficient of relative risk aversion (σ) of 3 (i.e. marginal utility of consumption is $\frac{1}{c^3}$); these assumptions deliver a set of social welfare weights. Specifically, these assumptions, - which yield that the marginal utility of consumption of a recipient in group A is 8 times that of group B - imply that society would desire a transfer policy from group B to group A with an MVPF of at least 0.125. Alternatively, we can compare the MVPFs of two policies with different distributional implications and calculate the revealed social welfare weights at which society would be indifferent between the two policies.

Put differently, redistribution across individuals is costly: attempts to redistribute the pie invariably shrink the pie. This equity-efficiency trade-off lies at the heart of economics and of public policy. It was memorably characterized by Arthur Okun through the metaphor of a leaky bucket: transferring resources from the rich to the poor, Okun suggests, is like trying to transfer water in a leaky bucket - some of the water leaks out along the way (Okun 1975). The MVPF quantifies the leaks in Okun's bucket. An estimate of the MVPF of a policy of $0.8 is akin to saying that for every dollar expenditure on the policy, it generates $0.8 in benefits, and $0.2 in "leaks". With an assumption of a social welfare function we can say whether society is willing to incur those leaks.

2.3 Additional applications

Once we realize the power of comparing the MVPF across different policies - with the help of the "efficient welfare weights" if they affect different populations - we suddenly realize there are many policies we can compare with this approach. It is conceptually straightforward to extend the core

MVPF framework we laid out in Section 2.1 in the context of a $1 targeted cash transfer to other types of policies, although doing so may raise additional empirical challenges. We briefly walk through the extension to a few canonical cases here.

2.3.1 Welfare analysis of an increase of $1 of subsidy

Consider for example welfare analysis of an increase in the subsidy of a good. To be concrete, consider for example the analysis in Finkelstein, Hendren and Shepard (2017) of changes in health insurance subsidies for low income adults. The MVPF expression here is identical to that in equation (2) for cash transfers: a $1 increase in subsidy is valued at $1 for infra-marginal recipients and at $0 for marginal recipients, and has a mechanical cost of $1 per infra-marginal recipient. What are the fiscal externalities in this case? One natural one is the fiscal externality imposed on the government budget by the fact that more people will enroll when the premiums are lowered; the amount of this fiscal externality depends on the magnitude of the enrollment response as well as the cost to the government per marginal enrollee. The cost to the government per marginal enrollee in turn depends on the difference between the premiums paid by the marginal enrollees and the average expenditures (e.g. health care costs and potentially administrative costs) by the government for the marginal enrollees. Another potential fiscal externality would arise if the increase in subsidies affected labor supply and hence government income tax revenue. Thus the specific causal effects that have to be estimated would differ, but the framework is the same.

2.3.2 Non-cash policies: Welfare analysis of an increase in $1 of government spending on an in-kind transfer

Welfare analysis of in-kind transfers is conceptually straightforward under the MVPF approach, although (as always) considerably more challenging empirically.

Consider for example an increase in $1 of an in-kind transfer such as health insurance, public housing, or food stamps. The envelope theorem still tells us we can ignore any benefits to marginal enrollees. And the costs of the program remain the mechanical cost (the $1 increase to all infra-marginal enrollees) plus the fiscal externality of the policy on the government budget. The key difference with an in-kind transfer is that we can no longer assume that it is valued by infra-marginal recipients dollar for dollar. On the one hand, an in-kind transfer might be valued at less than cost if it causes the infra-marginal recipients to consume more of the in-kind good than they would if given cash. On the other hand, an in-kind transfer might be valued at more than cost if it is not available at cost in the private market; this would apply, for example, to an insurance market suffering from adverse selection or a good that is provided in an imperfectly competitive market. We denote by W the average willingness to pay (out of their own income) by infra-marginal recipients per dollar of in-kind benefit. The MVPF is then:

$$MVPF^{inkind} = \frac{W}{1+FE}. \tag{3}$$

The (considerable) added empirical challenge is now to estimate W. This is relatively straightfor-

ward if the transferred good is also traded in the market at observed prices; the demand curve for the good among the infra-marginal recipients gives us W. Estimating demand is the bread and butter of empirical economics, so we are in familiar - if sometimes empirically demanding - territory. The exercise of inferring W becomes considerably more challenging when the expenditure is on an in-kind good that we do not observe offered in a market. How much do recipients value improvements in a public school's test scores, reduced pollution, or reduced mortality risk? One approach is to infer value from other market transactions - such as wages if the good is bundled into workplace amenities, or house prices if the good is concentrated locally (Rosen 1974; Greenstone 2017). Another approach is to specify a particular, calibrated utility function for the good; Finkelstein, Hendren and Luttmer (2015) offer several variants of this approach in estimating W in the context of Medicaid, the publicly provided health insurance program for low income individuals. Of course, these empirical challenges are not specific to the MVPF framework; any form of welfare analysis must grapple with how to estimate the monetized value of specific goods.

2.3.3 Multiple affected policies

Thus far we have considered policies that target a given income group i and where all costs are born by the government. Some policies have multiple beneficiaries of different income group; one natural example is a public good that can be used by people of different income levels. Other policies have impacts on non-recipients. For example, Finkelstein, Hendren and Shepard (2017) suggest that part of the benefits from expanding public subsidies for health insurance to low income adults accrue to the members of society who were previously financing uncompensated care for low income uninsured adults. Likewise, Chetty and Saez (2010) point out that when social insurance expansions affect the private insurance market (for example because of adverse selection in the private market), the impact on the private market needs to be accounted for in welfare analysis. In all of these cases, the MVPF approach requires simply that we modify it to account for all affected parties, weighted by appropriate social welfare weights. Put differently, we now have to consider both potential benefits to non-recipients in the numerator, and potential fiscal externalities on non-government parties in the denominator.

We illustrate this general point by way of two specific examples.

Welfare analysis of an increase in $1 of government spending on a public good One of the first problems we teach in public finance courses is the valuation of spending on a public good like a lighthouse. The well-known result (dating back to Samuelson (1954)) is that one needs to compare the sum of benefits across all people in society from the public goods expenditure to the cost of the expenditure. Suppose the lighthouse affects N people, and let i index these people, and W_i denote individual i's willingness to pay for a dollar more spending on the lighthouse out of her own income. Then the MVPF can be written:

$$MVPF^{publicgood} = \frac{\frac{1}{N}\sum_i \eta_i W_i}{1 + FE}. \tag{4}$$

The numerator is the average willingness to pay for the policy weighted by each individual's social welfare weight η_i, where the social welfare weight measures how much society values the marginal consumption of individual i. If all affected individuals have the same social welfare weights (for example they are all of the same income level), we can simply sum their willingness to pay. However, if society puts different weights on the utility of different affected individuals, we need to account for the fact that different affected individuals may have different willingness to pay $W_{i.}$

Where do the social welfare weights $\eta_i$come from? Typically the researcher assumes and imposes them through the choice of social welfare function. Reasonable people may disagree about how society's indifference point between say \$1 in the hands of "the rich" and \$0.X in the hands of the "poor". An alternative is to use Hendren's (2017) "efficient welfare weights" which measure the cost of moving money to each income level i through the tax schedule, and provide a natural benchmark for searching for potential Pareto improvements.

Welfare analysis of policies with externalities A classic public policy is a corrective, Pigouvian tax or subsidy to correct a negative (or positive) externality by forcing the externality generator to "internalize" the externality. Consider for example, increasing by \$1 a subsidy to the price of vaccines, which presumably generate positive (health) externalities on the population. The mechanical cost of the subsidy is simply \$1 times the number of infra-marginal recipients (I) who were already receiving the vaccine. The fiscal externality (FE) cost includes any impact of the subsidy on the government budget, for example through changes in health which may affect other publicly-financed health care expenditures or labor market participation and productivity and hence income tax revenue.

What about the benefits? Again, the group I of infra-marginal recipients who were getting the vaccine already value the \$1 decrease in its price at \$1. And a group of marginal recipients M who choose to get the vaccine because of the price reduction and, by the envelope theorem, are indifferent to their behavioral response (i.e. receive no net private welfare change). However the benefits include external effects from the behavior of the marginal recipients on the health of the rest of the population; the magnitude of these welfare effects depends on the magnitude and value of any external benefits E_i, as well as the social welfare weights for externally affected individuals (which we denote by η_i), who may differ from the infra-marginal recipients (whose social welfare weight we normalize to 1). We can thus write:

$$MVPF^{Pigou} = \frac{1 + \frac{\sum \eta_i E_i}{I}}{1 + FE}. \tag{5}$$

Even a policy not specifically designed to correct an externality, could have an effect on external parties (i.e. non recipients), and this would need to be accounted for using the same formula.

2.4 Summary

The MVPF of a policy is the ratio of its benefits (willingness to pay) to it costs. Costs require estimates of the *causal effects of the policy on the government budget*; therefore any behavioral response that might potentially affect the government revenue or expenditures needs to be estimated; any behavioral

responses with no public sector budgetary implications does not. Armed with the advances in applied micro in estimating such casual effects, the researcher now has guidance as to which causal effects are relevant to estimate for welfare analysis of a given policy.

What about benefits? For cash transfers or tax changes nothing needs to be done - cash is valued dollar for dollar by infra-marginal recipients and has no net utility impact on marginal recipients by the envelope theorem . For in kind benefits, the envelope theorem still applies for marginal recipients, but the value of the in-kind transfer to the infra-marginal recipients needs to estimated; it may be less than or greater than its cash cost, and estimating it poses its own empirical challenges. Finally, when there are multiple beneficiaries - such as in the case of a public good or an externality - a set of social welfare weights for all affected individuals needs to be assumed or estimated, again with its own empirical challenges. The purpose of the framework is not to "solve" all these empirical challenges but to clarify what objects are needed.

The derivation of this MVPF framework required **two key assumptions:** marginal changes and individuals making (privately) optimal choices. Together, these allow us to exploit the power of the **envelope theorem** to infer the welfare benefits to marginal recipients.

3 Frequently Asked Questions

Why don't we need to separately estimate income and substitution effects of the policy (as with Marginal Excess Burden)?

Those of us brought up in the traditional school of public finance welfare analysis have been steeped in the view that welfare analysis requires us to estimate compensated elasticities not uncompensated ones; the distortionary effects of government policy arise from its substitution effects, not its income effects; estimation of the deadweight loss of a tax policy - the excess of lost consumer welfare over the surplus over tax revenue collected - can be badly biased if the uncompensated (Marshallian) demand curve is used to measure consumer welfare rather than the compensated (Hicksian) demand curve (Hausman 1981). As a result, calculation of deadweight loss requires the researcher to calculate the behavioral response to a hypothetical policy that holds individual utilities constant through individual-specific lump-sum transfers, so that the substitution effect then governs the behavior of interest.

Since Harberger (1964), the concept of deadweight loss and the extension to marginal deadweight loss - also known as marginal excess burden (due to Auerbach 1985 and Auerbach and Hines 2002) - is arguably the most common welfare framework used to evaluate the welfare impact of a change in tax or transfer policies. The marginal excess burden of a tax change is defined as the difference between the Hicksian equivalent (or compensating, depending on researcher preferences) variation associated with the tax change and the change in tax revenue. Marginal excess burden is thus based on an arguably arbitrary and unrealistic thought experiment in which individual lump sum taxation (a policy instrument that doesn't exist in our second best world) is used to hold individual utility constant in the face of a given policy change. This necessitates measures of compensated elasticities - i.e. behavioral responses holding utility constant.

Of course,government policy changes are not compensated through individual-specific lump sum

taxation, so identifying compensated elasticities is challenging. It is easier for the empirical researcher to produce credible estimates of uncompensated behavioral responses rather than compensated ones. The MVPF takes advantage of this by using uncompensated behavioral responses. Importantly therefore, as we emphasized in Section 1.2, a given MVPF estimate has no independent meaning. Rather it is useful primarily as metric to compare to other policies. If two policies target the same population, the one with the higher MVPF is preferable. If they target different populations, they can still be compared by using estimates of the MVPF of the tax system (i.e. "efficient welfare weights") to determine the resource cost of moving the resources targeted by one policy to the policy targeted by the other policy (Hendren, 2017). Importantly, the MVPF is most directly useful in making comparisons across policies. By computing and comparing the MVPF across a range of policies, economists can provide guidance on how policy reforms can improve social welfare by moving resources from policies with low MVPF to high MVPF, regardless of whether they are tax policies or expenditure policies or combinations therefore. If the policies target different parts of the income distribution, one can use estimates of the MVPF of marginal income tax changes - which tell us the cost of moving resources to a given population through the tax code - to put them "on the same" footing (Hendren, 2017).

How does this all relate to the marginal cost of public funds (MCPF)?

Another venerable traditional in public finance welfare analysis is to take an estimate of the benefits of an expenditure policy and compare it to the cost, defined as the expenditures on the program multiplied by 1 plus "the marginal cost of public funds (MCPF)". Those of us in the empirical know have regularly invoked the conventional wisdom that the MCPF (or λ, as we insiders call it) is about 0.3. In other words, the MCPF is intended to provide an adjustment to the estimate of the welfare consequences of public spending that accounts for the distortionary cost of raising the tax revenue to finance that expenditure. Many the last section of an empirical paper on the impact of a policy will compare the painstakingly estimated benefits from that policy to the cost of the policy multiplied by 1.3 (see e.g. Finkelstein and McKnight 2008; Olken 2007). Presto: welfare analysis.[7]

Hendren (2016) however emphasizes the key limitation to the MCPF approach: it does not account for *who* the money is raised from. The idea that there is a single "cost of public funds" - as we typically assume in empirical welfare analysis - is absurd. Beginning with Mirlees (1971), the optimal tax literature has emphasized that the cost of raising revenue differs across the income distribution. Empirically, the literature had tended to find that it is cheaper to raise revenue from the poor than the rich; this is the mirror image of Okun's (1975) "leaky bucket" (i.e. that it is costly to redistribute to the poor) - as a result, it is cheap to raise revenue from them by reducing redistribution. For example raising money from the poor can be done through a reduction in the EITC, which has an MVPF of around 0.9, while raising revenue from the rich requires increasing top marginal income tax rates which has an MVPF of about 1.5 or higher (Hendren 2016, 2017). A lower MVPF for raising revenue from the poor than the rich is consistent with an optimal tax system in the presence of a social welfare function that places greater weight on the marginal value of resources for the poor than the rich (e.g.

[7] Of course, usually some of the costs are just transfers, and those only should be multiplied by 0.3, not 1.3. Presto: "insightful" public finance seminar comment.

a utilitarian social welfare function).

More broadly, the MCPF concept is based on the premise that we need to compare the welfare benefits from any given expenditure to the costs of raising the revenue to pay for the policy is "worth it". This suffers from two related issues. First, as already noted "the cost" of raising the revenue to pay for the policy depends on how that revenue is raised - i.e. what taxes on whom are increased to "pay for" the policy, since the cost of raising revenue differs across the income distribution (as well as across tax instruments).[8] Second, and relatedly, it will shock no modern reader of this piece to hear that, in practice, most government policy changes are not budget neutral, at least in the short run. Hendren's (2016) MVPF framework provides a way of comparing across (non budget neutral) policies[9]; augmented by his notion of "efficient welfare weights" (Hendren 2017), these comparisons can be made across policies that affect different parts of the income distribution, without having to assume and impose a particular social welfare function. This allows us to determine which policy instrument among potential options can achieve a given goal at the lowest resource cost, as well as what that resource cost is.

The MVPF offers an alternative way to form a budget neutral policy experiment: suppose we have estimated the MVPF of policy 1 and policy 2; then a budget neutral policy that raises revenue from policy 2 (e.g. by raising taxes or cutting spending on the policy) and spending money on policy 1 (e.g. by lowering some tax or increasing some expenditure) will increase welfare if and only if the MVPF of policy 1 exceeds the MVPF of policy 2.

How does the MVPF incorporate higher levels of government spending?

Many policy proposals are critiqued on the grounds of "they cost too much money." Think about "Medicare for All" or a Universal Basic Income program. Another example is the trade-off between direct public provision of a good and vouchers that subsidize consumers' purchase of the good on the private market. This is typically described by the trade-off between allocative inefficiency of direct public provision - some people would have preferred to purchase more of the good but cannot pay only for the incremental component - and the cost in terms of increased public expenditures on vouchers to those consumers who would have purchased a good on the private market even with the option of a free public good (Peltzman 1973, Cutler and Gruber 1996, Einav, Finkelstein and Williams 2016).

For those of us used to thinking about the marginal cost of public funds, this increased public expenditure seems like a cost. However, in the MVPF framework, this is just part of the mechanical part of the cost for inframarginals ($\$I$), since their behavior is not affected (except potentially via an income effect which I ignore). But isn't the increase in government spending via transfers to these inframarginals relevant for something? Yes. The relevant question becomes: what is the MVPF of

[8] The underlying literature of course recognizes this point - even if it is sometimes lost on us practitioners. For example, Kleven and Kreiner (2006) calculate how the MCPF evolves through the income distribution. Still, this leaves open the question of what tax instrument to assume is being used when financing a particular policy. Indeed, for some taxes - such as carbon taxes - the MCPF is arguably negative.

[9] Indeed the MVPF is not well-defined for budget neutral policies since by definition of budget neutrality the cost of the policy - i.e. the denominator - is 0. If one wants to think about a specific expenditure policy that is "paired with" a specific tax policy to "pay for it", the way to do it would be to consider the two policies separately, compute the MVPF of each, and then compare them.

this form of transfer as opposed to some other policy with the same distributional incidence. In the case of vouchers instead of direct public education (Peltzman 1973), we might think the transfers are distributed to relatively high income individuals (who use private schools), so the MVPF (of approximately $1) might compare poorly with the higher MVPF from a tax cut on high income individuals.

What if we don't want to invoke the envelope theorem? non marginal changes and non-optimizing agents

The envelope theorem is key to the MVPF formula - it allows us to not have to estimate or value the impacts of any behavioral responses to the policy except as they impact the government budget, since impacts on the government budget represent a "fiscal externality" that the optimizing agent does not internalize. Thus the envelope theorem buys us a lot. But it is not an innocuous assumption. There are (at least) two natural reasons why it may be violated in the context of a given policy change: non-marginal changes and non-optimizing agents.

For non-marginal policy changes, the approximation of the envelope theorem may fail, perhaps spectacularly. Kleven (2018) provides a formal analysis of this point and some possible approaches for analyzing large reforms. There are no silver bullets: Additional work - in the form of assumptions and/or estimates - will be needed. Fortunately, the MVPF framework is still a useful starting point - including the emphasis on constructing a ratio of benefits to costs which can then be compared across policies, and the fact that the costs can be usefully decomposed into the mechanical cost and the fiscal externality. However, the benefits to marginal recipients can no longer be ignored - they must be estimated and incorporated. Finkelstein, Hendren and Luttmer (2015) provide an empirical example in the context of (non-marginal) expansions in Medicaid eligibility. Finkelstein and Notowidigdo (2018) provide an empirical example in the context of (potentially non-marginal) reductions in hassle costs of signing up for a means-tested benefit program. Relatedly, if agents are not behaving optimally, there may be first order welfare effects for marginal recipients. Again, these must be estimated and included in the numerator of the MVPF formula; Finkelstein and Notowidigdo (2018) provide an empirical example.

4 Empirical Examples

In this last section we aim to further reinforce the ideas behind the MVPF by giving some specific, recent empirical examples.

The Earned Income Tax Credit (EITC).

The Earned Income Tax Credit (EITC) is one of the primary forms of tax subsidies for low-income individuals in the United States. Consider the thought experiment of an EITC expansion: a marginal shift outward of the entire schedule. Equation (2) tells us that to estimate the MVPF of an EITC expansion requires us to estimate one object: the fiscal externality of this expansion or, in other words, the causal impact of behavioral responses to the expansion on the government budget. Why is this?

Well what are the benefits of the policy per dollar of "mechanical cost" (i.e. direct expenditure on it)? Infra-marginal beneficiaries - those already receiving the EITC - value the $1 increase in benefits at $1; marginal beneficiaries - those who change their labor supply behavior to become newly eligible for the EITC in response to its expansion (e.g. those who enter the labor force) - are, by the envelope theorem, indifferent between their old and new allocation; their net utility gain is a 0. Thus the benefit or willingness to pay for the expansion per dollar of direct expenditure is 1. The costs of the expansion per dollar of direct expenditure are the mechanical cost ($1) plus the FE.

The empirical work comes in estimating the FE. In principle, it could include a potentially limitless set of ways that the EITC could affect the government budget, including effects on taxable labor income, effects on taxable capital income, public health expenditures (if increases in the EITC have health effects for recipients who are also on Medicare or Medicaid), criminal justice expenditures (by the same token), take up of other social programs such as SSDI or SNAP etc etc. Note that these fiscal externalities from increasing EITC benefits can be created by both marginal enrollees and infra-marginal enrollees whose behavior changes through the receipt of higher income, if that behavioral change impacts the government budget. Estimating the FE could well prove a full employment program for generations of empirical applied micro researchers.

For simplicity, Hendren (2016) focuses his illustration of the FE of the EITC solely on the impact of an expansion on taxable labor income. He draws on an extensive empirical literature summarized by Hotz and Scholz (2003) documenting how EITC expansions cause an extensive marginal labor force response: individuals enter the labor market on the extensive margin in response to the expansion of the EITC. This increase in labor force participation creates a negative fiscal externality on the government through its impact on reduced income tax revenue (since for these marginal entrants, income tax rates are negative); Hendren (2016) draws on the existing empirical estimates to conclude that this FE on the extensive margin increases the cost of the EITC expansion by an additional $0.09 for every dollar in mechanical cost, reflecting the wage subsidies paid to the marginal enrollee (if that marginal enrollee lost other public benefits such as Section 8 housing vouchers, that should be accounted for as well). In addition, the EITC expansion can have an effect on labor supply (And hence income tax revenue) on the intensive margin as individuals move their labor earnings toward the revenue-maximizing kink point Drawing on estimates from Chetty et al. (2013), Hendren (2016) estimates that this creates an additional fiscal externality cost of about $0.05 for every dollar of mechanical cost. Combining these two margins, Hendren (2016) concludes that the FE due to the externality is about $0.14 per dollar of mechanical cost; thus the MVPF of the EITC expansion is 1/(1+0.14) or $0.88. In other words,, for every dollar of expenditure on the EITC, about $0.88 ends up as welfare benefits to recipients; $0.12 "leaks" out or is otherwise the "resource cost" of transferring this money to the low income recipients. Note that, conversely, raising revenue by increasing taxes on the poor has a negative resource cost, since we can do it by cutting the EITC which has an MVPF of less than 1. Conversely, as we will now see, raising revenue by increasing taxes on the rich has an MVPF of less than 1 - i.e. a positive resource cost.

Top marginal tax rates.

What is the MVPF of a tax cut through a reduction on the top marginal income tax rate, which is applied to all incomes above some threshold. What is the MVPF of this policy? Hopefully by now the benefit part is clear: infra-marginal recipients of the tax cut value the $1 of tax cut at a $1, marginal recipients derive no net utility due to the envelope theorem. The cost should also be clear: it is the mechanical cost (which we have normalized to 1) plus the FE - the impact of behavioral responses to the tax cut on the government budget.

Again, there are many potential channels by which behavioral responses to the tax cut can impact the government budget; Hendren (2017) concentrates on what is arguably the most direct: the impact of behavioral responses to the tax cut on tax revenue. He notes that there is a large empirical literature estimating how behavioral responses to changes in the top marginal tax rate affects taxable income. He shows how these estimates - - together with estimates of the shape of the income distribution - can be used to estimate the impact of behavioral responses to changes in the top marginal tax rate on tax revenue. Implementing this approach using data on from tax records on the shape of the income distribution and existing estimates of behavioral responses and their effects on taxable income, Hendren (2017) estimates that the FE from reducing the top marginal tax rate is about $-\$0.3$ for every dollar of mechanical reduction in the tax liability of top earners. Here the fiscal externality is negative - so that the total cost of the policy is less than its mechanical cost through lost tax revenue, because the reduction in marginal tax rates causes behavioral changes that increase total tax revenue; these may include, for example, both increases in labor supply and reduced uses of tax sheltering strategies. As a result, the MVPF of reducing the top marginal tax rate is $1/(1\text{-}0.3) = 1.4$. Importantly, this exercise suggests that it would be beneficial for subsequent empirical work to directly estimate the impact of changes in tax rates on government revenue, since this is the object of interest for the MVPF.

We could of course easily do the thought exercise in reverse: what is the MVPF of a tax increase through an increase in the top marginal tax rate. Once again, the MVPF is 1.4 - every dollar the government collects, takes $1.40 in welfare away from the people whose taxes were increased.

Health Insurance Subsidies for Low Income Adults

The U.S. government spends an enormous amount of money on health insurance for low-income adults. Since the enactment of the Affordable Care Act in 2010, a key way these expenditures are made is through subsidizing health insurance premiums for low income adults on the state or federal health insurance exchanges. What is the MVPF of a $ increase in subsidies for health insurance for low income adults. Finkelstein, Hendren and Shepard (2017) provide estimates. Once again, the empirical work is in estimating the FE. The benefit of the subsidy increase is $1 to infra-marginal recipients and 0 (by the envelope theorem) for the marginal recipients who take up health insurance in response to the increase in subsidy. The costs are the $1 in mechanical costs - increased subsidy expenditures - plus the FE of the increase in the policy.

Once again, the set of behavioral responses to the increase in the subsidy which could affect the government budget are numerous. The authors ignore potential behavioral responses for infra-marginal

enrollees (due to increased income which might affect e.g. their receipt of other benefits) and focus on the fiscal externalities from behavioral changes from marginal enrollees. Here too, they focus on just one way that marginal enrollees' behavioral response could impact the government budget, namely through changes in government expenditures due to increased enrollment (i.e. the marginal recipients who take up health insurance in response to the increase in subsidy).[10] The fiscal externality in this setting is the increase in enrollment times the public subsidy per marginal enrollee (i.e the total cost of the health insurance for the marginal enrollee minus the part of that cost born by the marginal enrollee (rather than the government) through the enrollee premium). Note that the public subsidy per marginal enrollee could be negative - and thus the fiscal externality could be negative. This would be the case if due to a market failure (such as adverse selection) the willingness to pay of the marginal enrollee (i.e. the enrollee premium) was actually above the cost of the marginal enrollee. In other words, it is possible that the MVPF of these subsidies could be above 1.

Finkelstein, Hendren and Shepard (2017) use this approach, together with estimates of how demand for insurance and the costs of enrollees are affected by changes in the subsidy to estimate the MVPF of the subsidy. Naturally, the MVPF of an additional $ of subsidy for public health insurance - will vary depending on the initial condition - i.e. what percent of individuals already have insurance; this affects the MVPF because it affects the fiscal externality through two channels: the slope of the demand curve (and hence the enrollment response) can vary, and the difference between the cost for marginal enrollees and their willingness to pay. In their context, the MVPF measurement also needs to account for the fact that there are multiple beneficiaries: when the marginal uninsured low income adults takes up formal health insurance in response to an increase in the subsidy, the parties who were previously providing some uncompensated care (charity care) to the previously "uninsured" also benefit. These need to be accounted for by adding the reduced uncompensated care expenses per $1 increase in subsidy times the number of marginal enrollees and social welfare weight of the benefiting party; if the benefiting party is the government, one could alternatively subtract it from costs. Assuming the incidence is on the government, the authors estimate that when about 30% of the market is insured, the MVPF of an additional subsidy is 1.3, but when 90% is insured, it falls to $0.8.

Provision of in-kind goods.

If the government increases its expenditure on an in-kind good - such as public housing, public health insurance (Medicaid) or food stamps (SNAP) - by $1 the key modification to the MVPF is, as discussed above, that we can no longer assume that the infra-marginal enrollees value this increase in spending at a $. They might value it by less than a dollar - if the provision of the in kind good is more than they would choose to purchase if given the equivalent amount in cash - or at more than a dollar - if the provision of the in kind good is one that cannot be bought at marginal cost on the private market (due for example to market failures such as imperfect competition or adverse selection).

[10] Another natural way that marginal enrollees' behavior could affect the government budget is if they changed their labor market activity in order to newly-qualify for the expanded means tested subsidy; if individuals reduce their earnings in response to higher subsidies, this would introduce an added source of fiscal externality.

Figure 1: Estimates of Policy MVPFs and Efficient Welfare Weights from Hendren (2017)

Figure 8: Using Efficient Welfare Weights for Policy Evaluation

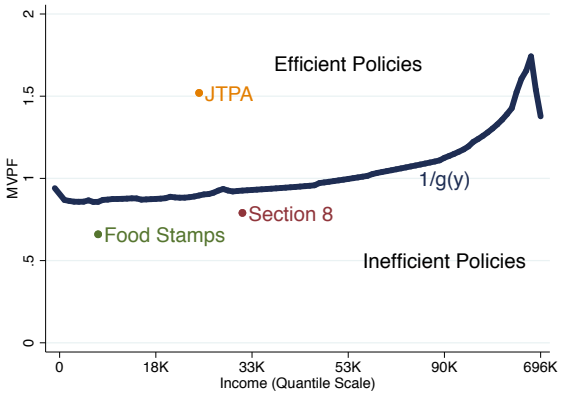

Notes: This figure illustrates the use of the efficient welfare weights for assessing the efficiency of government policy changes. The line presents the value of $\frac{1}{g(y)}$, which represents the amount of welfare that can be delivered to each portion of the income distribution per dollar of government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

Assessing the value of in-kind transfers can be empirically challenging. If they are traded in a well-functioning market, one can estimate demand (willingness to pay) for the in-kind good and thus value it; for example, this has been done for food stamps. If, however, one must find other creative ways to estimate the value of the in kind good to the infra-marginal recipients. Finkelstein, Hendren and Luttmer (2015) show how one can use causal estimates of the impact of Medicaid on outcomes that are arguments of the utility function to try to value this. They end up with an MVPF of about 0.2 to 0.4.

Combining the applications

A key attraction of the MVPF framework is the ability to compare across policies to determine the most efficient way to do accomplish a given goal. When the policies have the sample distributional incidence, this is simple: the one with the highest MVPF is preferred. When they have different distributional incidence, Hendren (2017) suggests that that the MVPF of a policy that affects a given income level be compared MVPF of a tax cut to that income level (i.e. to the efficient welfare weight). Expansions of policies whose MVPF is above the relevant efficient welfare weight represent potentials for Pareto improvements; conversely cuts in policies whose MVPF is below the efficient welfare weight represent potential improvements. To illustrate this, Figure 1 reproduces the empirical summary given in Figure 8 of Hendren (2017).

5 Conclusion

The MVPF framework offers a powerful framework for welfare analysis of a policy change based on affected individual's own willingness to pay for the policy change to the causal effect of the policy on government's costs. When the policy involves cash to one group, willingness to pay is trivially constructed; when it involves in-kind transfers or cash to multiple groups with potentially different social welfare weights, additional empirical challenges emerge but the conceptual framework remains. A key attraction of this approach for empirical work is that it allows one to directly incorporate causal estimates of behavioral responses to policies - of the type that is now the bread and butter of applied micro economics - into welfare analysis: these behavioral responses matter through their impact on the government budget. The interplay between the framework and empirical work runs both ways: at the same time, the MVPF provides an important guide for future empirical work on *which* behavioral responses matter for welfare. Empirical public finance economists interested in translating the benefits of the "credibility revolution" into progress on applied welfare analysis would be well advised to focus their empirical efforts on estimating behavioral responses that have fiscal externalities on the government budget, and not behavioral responses whose costs are likely fully (or approximately fully) internalized by the responding individuals.

References

Angrist, Joshua D., and Jorn-Steffen Pischke (2010). The Credibility Revolution in Empirical Economics: How Better Research Design is Taking the Con out of Econometrics. Journal of Economic Perspectives, 24(2), 3-30.

Auerbach, Alan. 1985. The theory of excess burden and optimal taxation. Handbook of public economics, Elsevier.

Auerbach, Alan and Jim Hines. 2002. Taxation and economic efficiency. Handbook of public economics, Elsevier.

Chetty, Ray and Emmanuel Saez. 2010. Optimal Taxation and Social Insurance with Endogeneous Private Insurance. American Economic Journal: Economic Policy 2: 85-114.

Chetty, Raj, John Friedman and Emmanuel Saez. 2013. Using differences in knowledge across neighborhoods to uncover the impacts of EITC on earnings. American Economic Review.

Cutler, David M., and Jonathan Gruber. 1996. "Does Public Insurance Crowd Out Private Insurance?" Quarterly Journal of Economics 111 (2): 391–430

Einav, Liran, Amy Finkelstein and Heidi Williams. 2016. Paying on the margin for medical care: evidence from breast cancer treatments." American Economic Journal: Economic Policy. 8(1):52-79.

Feldstein, M. 1999. "Tax Avoidance and the Deadweight Loss of the Income Tax." Review of Economics and Statistics 81 (4): 674–80

Finkelstein, Amy, Nathan Hendren and Erzo F.P. Luttmer. 2015. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." NBER Working Paper 21308.

Finkelstein, Amy, Nathan Hendren and Mark Shepard. 2017. "Subsidizing health insurance for low-income adults: Evidence from Massachusetts." NBER Working Paper 23668.

Finkelstein, Amy and Robin McKnight. 2008. "What Did Medicare Do? The Initial Impact of Medicare on Mortality and Out of Pocket Spending." Journal of Public Ecopnomics 92: 1644-1669.

Finkelstein, Amy and Matthew Notowidigdo. 2017. "The Effects of Information and Application Assistance: Evidence from a SNAP Take-up Experiment." https://economics.mit.edu/files/14789

Goolsbee, Austan. 1999. "Evidence on the High-Income laffer curve from six decades of tax reform." Brookings Papers on Economic Activity 2:1-64.

Greenstone, Michael. 2017. "The Continuing Impact of Sherwin Rosen's 'Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition". in John List and Harald Uhlig (eds) "The Past, Present and Future of Economics: a Celebration of the 125 Year Anniversary of the JPE and of Chicago Economics." 125 (6): 1891-1902.

Hausman, Jerry. 1981. "Exact Consumer Surplus and Deadweight Loss." American Economic Review.

Harberger, Arnold. 1964. The measurement of waste. American Economic Review 84: 58-76.

Hendren,. Nathan. 2013. "The policy elasticity." National Bureau of Economic Research Working Paper 19177.

Hendren, Nathan. 2016. "The policy elasticity". Tax Policy and the Economy 30 (1)

Hendren, Nathan, 2017. "Efficient welfare weights." NBER Working Paper 20351.

Hicks, J.R. 1940. "The valuation of social income." Economica 7 (26): 105-124.

Hotz, J and K Scholz. 2003. The earned income tax credit. In Moffit (ed) Means-Tested Transfer Programs in the United States. University of Chicago Press.

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

Kleven, H. J., and C. T. Kreiner. 2005. "Labor Supply Behavior and the Design of Tax and Transfer Policy." Danish Journal of Economics 143:321–58.

Kleven, H.J., and C.T. Kreiner. 2006. "The Marginal Cost of Public Funds: Hours of Work versus Labor Force Participation." Journal of Public Economics 90:1955–73.

Kleven, H.J. 2018. "Sufficient Statistics Revisited". Working Paper. https://www.henrikkleven.com/uploads/3/7

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

Mirlees, James. 1971. "An exploration into the theory of optimal income taxation." Review of Economic Studies 38 (2): 175-208.

Peltzman, Sam. 1973. "The Effect of Government Subsidies-in-Kind on Private Expenditures: The Case of Higher Education." Journal of Political Economy 81 (1): 1–27.

Ramsey, Frank. 1927. "A contribution to the theory of taxation." The Economic Journal 37 (

Saez, Emmanuel. 2004. "Reportred Incomes and Marginal Tax Rates, 1960-2000: Evidence nad Policy Implications." Tax Policy and the Economy 18: 117-172.

Slemrod, J., and S. Yitzhaki. 1996. "The Social Cost of Taxation and the Marginal Cost of Funds." International Monetary Fund Staff Papers 43 (1):172–98.

Slemrod, J. and S. Yitzhaki. 2001. "Integrating Expenditure and Tax Decisions: The Marginal Cost of Funds and the Marginal Benefit of Projects." National Tax Journal 54 (2): 189–202

Okun, Arthur. 1975. Equality and Efficiency. Brookings Insitution Press

Olken, Benjamin. 2007. "Monitoring Corruption: Evidence from a Field Experiment in Indonesia." Journal of Political Economy 115(2): 200-249.

Rosen, Sherwin. 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82 (1): 34-55.

Samuelson, Paul. 1954. The Pure Theory of Public Expenditure. The Review of Economics and Statistics 36(4): 387-389..