

The Case for Using the MVPF in Empirical Welfare Analysis*

Nathaniel Hendren (Harvard University and Policy Impacts)
Ben Sprung-Keyser (Harvard University and Policy Impacts)

May, 2022

Abstract

This paper outlines the case for using the Marginal Value of Public Funds (MVPF) in empirical welfare analysis. It compares the MVPF approach with more traditional welfare metrics such as the Cost-Benefit Ratio and the Net Social Benefits criterion. It outlines the advantages of the MVPF approach relative to these metrics. In building the case for the MVPF, this paper also addresses several misconceptions about the MVPF that appear in recent literature.

1 Introduction

Evaluating the desirability of government policies requires the consistent application of a welfare metric designed to compare across policies. Hendren and Sprung-Keyser [2020] popularized the Marginal Value of Public Funds (MVPF) as a unified metric for evaluating the welfare consequences of government policy. Some recent work, most notably García and Heckman [2022a, 2022b], has argued in favor of alternate criteria for evaluating social programs such as a Net Social Benefits criterion. In this paper we discuss the value of the MVPF relative to alternate metrics.

We proceed in four steps. First, we provide the definition of the MVPF and outline the conceptual policy experiment that motivates its use. Second, we compare the MVPF to the traditional Cost-Benefit Ratio and highlight the MVPFs unique advantages. Third, we conduct a similar comparison between the MVPF and the Net Social Benefits criterion. Finally, we address several misconceptions about the MVPF approach that are present in recent criticisms.

*We thank Jamie Emery for excellent research assistance, and we are grateful to Robert Hahn and Robert Metcalfe for helpful comments. This research was supported by Policy Impacts.

2 The Definition of the MVPF

In order to debate the relative merits of various metrics for welfare analysis, it is important to clearly define those metrics and outline the underlying conceptual experiment justifying their use. For any policy change, the MVPF is the ratio of the benefits that the policy provides to its recipients, divided by the policy’s net cost to the government. Hendren and Sprung-Keyser [2020] define the MVPF is defined as:¹

$$MVPF = \frac{Benefits}{Net\ Govt\ Cost} = \frac{\Delta W}{\Delta E - \Delta C} \quad (1)$$

where ΔW denotes the benefits that the policy provides to individuals in the population, ΔE denotes the government’s upfront expenditure on the policy, and ΔC denotes the long-run reduction in government costs due to the the causal effect of the policy.² For example, if a policy increases earnings and, consequently, increases taxes, those tax savings are incorporated in ΔC . Similarly, if a new spending policy results in a reduction in government transfers, that reduction in transfers is incorporated in ΔC . The MVPF measures a policy’s “bang for the buck”: For a given policy with $MVPF = A$, the policy delivers $\$A$ of benefits per dollar of net government spending.

As is highlighted in the discussion that follows, the novelty of the MVPF approach centers on its treatment of government costs and the way in which it closes the government budget constraint. The MVPF captures the long-run return on net government spending, rather than measuring the return on initial budgetary outlays. Moreover, the MVPF approach avoids arbitrary assumptions about how the budget constraint is closed (ie, how a policy is paid for.) Traditional welfare metrics adjust expenditures for the deadweight loss of taxation.³ By contrast, the MVPF does not impose a specific means of closing the budget constraint. This allows for an intuitive comparison of non-budget neutral policies. One can close the budget constraint by simply comparing the MVPF of an expenditure policy with the MVPF of a revenue raiser.

For example, take any two policies where the MVPF of the first policy is $MVPF_1 = A$ and the MVPF of the second policy is $MVPF_2 = B$. The MVPF allows us to assess the welfare consequences of spending on Policy 1 financed by raising revenue from Policy 2. This combined policy increases social welfare if and only if one prefers to give $\$A$ to Policy 1 beneficiaries relative to $\$B$ to Policy 2 beneficiaries. Formally, the combined policy increases social welfare iff

$$\bar{\eta}_1 MVPF_1 > \bar{\eta}_2 MVPF_2 \quad (2)$$

where $\bar{\eta}_j$ is the average social marginal utility of the beneficiaries of policies $j = 1, 2$. Given MVPFs A and B , the combined policy increases welfare iff $\bar{\eta}_1 A > \bar{\eta}_2 B$. The MVPF framework formalizes the tradeoff of

¹Both García and Heckman 2022a and 2022b provide an incorrect definition of the MVPF. They use the following expression:

$$"MVPF" = \frac{\Delta W}{(\Delta E - \Delta C)(1 + \phi)}$$

This definition adjusts the denominator by $1 + \phi$, where ϕ is marginal deadweight cost of taxation. They argue that ϕ is set to 0 in Hendren and Sprung-Keyser, 2020 because that work ignores the distortionary cost of taxation. In fact, there should be no ϕ in the MVPF calculation (even if it is costly to raise revenue). This is because the MVPF approach fundamentally differs from traditional cost-benefit analyses by allowing users to choose how to close the budget constraint by comparing MVPFs. As discussed in Section 3.2, this is one of the primary advantages of the MVPF approach. The welfare consequences of raising revenue can be evaluated by using MVPF estimates of specific policy instruments, rather than making a catch-all assumption about the deadweight loss of taxation.

²The equations above follow notation from García and Heckman [2022a]. This notation differs from García and Heckman [2022b]. In the latter case, ΔW is expressed as B , ΔC is expressed as D , and ΔE is expressed as $(1 - \omega)D$.

³The rarely stated implicit assumption motivating this adjustment is that the policy is paid for using a change in the linear income tax. The deadweight loss of taxation is then determined by the elasticity of taxable income associated with this hypothetical income tax change.

moving a dollar between different groups of beneficiaries and is transparent about the distributional incidence of doing so.⁴ An extensive set of MVPF estimates produced by researchers can be found in the Policy Impacts Library at www.policyimpacts.org.

3 Comparison to Benefit Cost Ratios

Benefit-cost ratios are one of the most commonly used metrics for evaluating the welfare consequences of government policy. They pervade discussions of policy effectiveness, such as research on early childhood education (see, e.g., Barnett and Masse [2007], Heckman et al. [2010], García et al. [2020]).⁵ In this section, we compare the MVPF to this traditional metric and argue the MVPF has unique advantages.

Using the notation above, the benefit-cost ratio is defined as:

$$BCR = \frac{\Delta W + \Delta C(1 + \phi)}{\Delta E(1 + \phi)} \quad (3)$$

The BCR and the MVPF are created using many of the same components, but this definition differs from the MVPF in two important ways. First, in the benefit-cost ratio, the benefits that flow back to the government (ΔC) are added in the numerator rather than subtracted in the denominator. Second, changes in government revenue are multiplied by $1 + \phi$, where ϕ is marginal deadweight cost of taxation. This implicitly assumes that funds are raised via an increase in a linear income tax rate that imposes a welfare cost of $1 + \phi$ per dollar of revenue that is raised. These two differences between the MVPF and the BCR are quite consequential. We discuss each in turn.

3.1 Capturing Government Savings

The MVPF measures the benefits of a policy per dollar of net government expenditure. All savings to the government are counted in the denominator as offsetting initial costs. The BCR measures benefits per dollar of upfront government expenditure. It treats government savings as a benefit with value $1 + \phi$. This second approach can cause the government to miss potential Pareto improvements, leaving money on the table.

Consider, for example, a hypothetical comparison between two policies. Suppose the first policy is a college scholarship program with an upfront cost of $\Delta E = 1$ per participant and a value to each participant of $\Delta W = 1$. This policy increases earnings and consequently increases government tax revenue. Those tax revenue gains offset the cost of the policy so that $\Delta C = 1$. Suppose the second policy is an expenditure policy that delivers $\Delta W = 3$ per beneficiary in benefits and costs the government $\Delta E = 1$, but the behavioral response to this policy has no impact on the government budget so $\Delta C = 0$.

The first policy provides benefits to individuals at no government cost. Any expenditure on that policy represents a Pareto improvement.⁶ In this case, the MVPF is infinite. By contrast, the BCR of this Pareto-improving policy is $\frac{2+\phi}{1+\phi}$ since the \$1 in increased tax revenue is considered a social benefit and counted in the numerator (rather than subtracted from the denominator). For the second policy, the MVPF is 3 and the

⁴In the case where the two policies have similar beneficiaries so that $\bar{\eta}_1 = \bar{\eta}_2$, then the combined experiment increases welfare if and only if the MVPF of policy 1 is higher than policy 2: $A > B$. For more discussion of distributional incidence, see Section 4.4.

⁵In recent work, García and Heckman have argued that benefit cost-ratios are not “an appropriate criterion for determining social optimality.” That said, given the prominence of this metric in their own work, we believe it is appropriate to compare the MVPF to the BCR.

⁶This is true regardless of the value of ϕ .

BCR is $\frac{3}{1+\phi}$. Thus, the traditional benefit-cost framework would prioritize the second policy, even though the first policy represents a Pareto improvement.

Hendren and Sprung-Keyser [2020] show that this distinction can matter in practice. For example, they consider the Medicaid expansion to children born after September 30, 1983. Research shows this policy reduced the prevalence of chronic health conditions and healthcare utilization, such as hospitalization and emergency room visits, fell accordingly. The savings in future healthcare expenditure (ΔC) more than offset the upfront spending on the policy ΔE . This policy effectively paid for itself, which is defined to be an infinite MVPF. Hendren and Sprung-Keyser also consider the BCR of this policy and find it barely exceeds 1.

3.2 Closing the budget constraint

A second important distinction between the MVPF and the BCR is that the MVPF does not assume a particular method of closing the budget constraint. In the BCR equation, the costs ΔC and ΔE are all adjusted by $1 + \phi$, where ϕ is marginal deadweight cost of taxation. Implicitly, this definition assumes that the policy is financed by a change in the linear income tax. As discussed in more detail in Section 4.3, this adjustment is unnecessary, and potentially misleading. There is no need to close the budget constraint when comparing the welfare consequences of various policies. If one seeks to compare the relative effectiveness of two different spending policies, that can be done without making an assumption about how the budget constraint is closed. For example, when evaluating the relative effectiveness of similar expenditures on early childhood spending or on college scholarships, incorporating the deadweight loss of taxation is not informative for conducting that comparison. In fact, such an adjustment makes it more difficult to compare results across bodies of research as different researchers might make different (and arbitrary) assumptions regarding the DWL of taxation.

That said, the MVPF framework does enable one to conduct budget neutral policy comparisons. As explained in Section 2, one can take the MVPF of an expenditure policy and compare it with the MVPF of the the associated revenue raiser. The key here is that there is no need to assume a one-size-fits-all distortion applies to all revenue increases. Imagine, for example, that the government is considering a large fiscal transfer to families with children and intends to finance that policy by increasing capital gains taxes. Those policies each have an associated MVPF. There is no reason to measure the welfare consequences of the transfer policy by arbitrarily assuming it is paid for using a linear income tax. Instead, one can construct the MVPF of the fiscal transfers, $MVPF_{transfers}$, and the MVPF of the capital gains taxes, $MVPF_{capitalgains}$. One can directly conclude from the MVPFs that the policy increases welfare iff $\bar{\eta}_{transfers} MVPF_{transfers} > \bar{\eta}_{capitalgains} MVPF_{capitalgains}$, where $\bar{\eta}_{transfers}$ and $\bar{\eta}_{capitalgains}$ are the social welfare weights for those affected by the transfers and capital gains taxes, respectively.

4 Comparison to Net Social Benefit

Some have argued that the welfare consequences of government policy should be evaluated using the Net Social Benefit (NSB) criteria (García and Heckman, 2022a, 2022b). In this section, we compare the MVPF to NSB and highlight the advantages of the MVPF. The NSB of a policy is defined as:

$$NSB = \Delta W - (1 + \phi)(\Delta E - \Delta C) \quad (4)$$

The NSB focuses on differences rather than ratios, calculating the welfare gain of a policy and subtracting its net cost. The net cost of a policy is adjusted by $1 + \phi$, the distortionary cost of raising the linear income tax.

In evaluating the tradeoffs between the two metrics, it is helpful to examine the case for the NSB, as presented in García and Heckman (2022a, 2022b). They argue for the superiority of the NSB using the example presented in Table 1 below.

Table 1: MVPF vs. Net Social Benefit

Term	Formula	Policy 1	Policy 2
Benefits to Beneficiaries	ΔW	500	100
Upfront Cost	ΔE	100	100
Benefits to Taxpayers ($\phi = 1/3$)	ΔC	25	100
Net Government Cost ($\phi = 1/3$)	$\Delta E - \Delta C$	75	0
Welfare Cost of Raising Revenue	$(1 + \phi)(\Delta E - \Delta C)$	100	0
Net Social Benefit (eq. 4)	$NSB = \Delta W - (1 + \phi)(\Delta E - \Delta C)$	400	100
MVPF	$MVPF = \Delta W / (\Delta E - \Delta C)$	6.67	∞

In this set-up, there are two different policies with the same upfront cost of \$100, $\Delta E = 100$. The first policy delivers \$500 in benefits, $\Delta W = 500$, while the second policy delivers \$100 in benefits, $\Delta W = 100$. The first policy recoups 25% of its initial costs ($\Delta C = 25$) while the second policy pays for itself ($\Delta C = 100$).⁷ García and Heckman argue that the first policy is strictly preferable because the net social benefits are 400, rather than $NSB = 100$ for the second policy. (The MVPF of the first policy is 6.67 and the MVPF of the second policy is infinite.) Observing this discord between NSB and the $MVPF$, García and Heckman [2022a] argue that the MVPF can provide poor guidance for policy. This argument falls short for 4 reasons.

4.1 Upfront versus Net Spending

The case for the NSB relies very heavily on the assumption that the government has a fixed quantity of funds for upfront expenditure. The example provided in García and Heckman [2022a] measures programs by the size of their upfront cost, ΔE , rather than the size of their net cost. In order to see the trouble with this approach, imagine a new policy 3 with $\Delta W = 100$, $\Delta E = 100$, and $\Delta C = 99$. This is a policy that produces \$100 in benefits but only costs the government \$1 because the government recoups \$99 of its initial expenditure. When comparing this to Policy 1 above, this Policy 3 has a much lower NSB. That said, this policy only costs the government \$1. The government could implement this new policy 75 times over before its total costs would equal that of Policy 1 above. If the government implemented the policy 75 times, the benefits would be more than 7400, far in excess of the benefits from Policy 1.⁸

⁷García and Heckman [2022a] provide an example where $(\Delta E - \Delta C)(1 + \phi) = 100$ but they do not specify ϕ . They incorrectly define the MVPF to include $(\Delta E - \Delta C)(1 + \phi)$ in the denominator, as opposed to $(\Delta E - \Delta C)$. As a result, we must make an assumption about ϕ to correctly infer the MVPF that would have been estimated in their example. We assume $\phi = 1/3$ for the calculations in Table 1.

⁸It is worth noting that the NSB approach makes a peculiar assumption about the value of government savings. It assumes that all government savings are rebated to taxpayers through reductions in a linear income tax, even if the policy at hand has a higher MVPF than the linear income tax. For example, if a policy produces \$5 for each dollar in upfront government expenditures, government savings are still valued at $1 + \phi$ rather than 5.

4.2 From Research to Policy – Ratios versus Levels

The net social benefit, NSB , measures the benefits of the policy in levels ($Benefits - Cost$), rather than in ratios, $\frac{Benefits}{Cost}$. While there is nothing inherently wrong with evaluating policies in this manner, we believe it is also valuable to measure the dollar-for-dollar return on government spending. It provides an additional piece of crucial information when evaluating potential government policies.

The advantage of ratios is particularly clear when thinking about the translation of economic research into economic policy recommendations. It is common practice amongst researchers to conduct empirical evaluations of policies on a small set of individuals. Those small scale studies are then used to make policy recommendations for a larger set of individuals. For example, those aiming to understand the value of pre-school have studied small scale programs such as Perry Preschool or Abecedarian [Heckman et al., 2010, Barnett and Masse, 2007].

If one were to rely exclusively on the NSB, these studies would be overlooked. It is mechanically true that the NSB of Perry Preschool is low, but only because the program itself was small. It would, however, be silly to dismiss these programs out of hand because they represent small scale policy experiments. This is part of the reason that ratios are quite common in economic research. It is valuable to measure the return on a dollar of additional spending.⁹

García and Heckman [2022b] seek to address this concern in the appendix to their work by advocating for modified NSB metrics that measures NSB per recipient or per dollar of upfront expenditure.¹⁰ They argue these alternate metrics should be used to rank policies. The problem with these solutions is that they reintroduces an arbitrary normalization. When evaluating policies per policy recipient, the definition of a policy recipient and the quantity spent on each recipient will have a large impact on the ranking of policies. It is not clear why a \$5000 scholarship given to 100 students should be regarded as twice as good as a \$2500 scholarship given to 200 students. Similarly, if one normalizes by upfront expenditures, the resulting expression is simply a modified benefit-cost ratio, which we discussed above in Section 3.¹¹

4.3 Closing the Budget Constraint

The traditional net social benefit criterion assumes that tax revenue is valued at $1 + \phi$, where ϕ is the marginal deadweight cost of taxation. This valuation comes from the embedded assumption that the budget constraint is closed using a linear income tax. By contrast, the MVPF does not require the budget constraint to be closed in any particular manner. In fact, the Net Social Benefits criterion with a deadweight loss of ϕ is simply a special case of the MVPF condition outlined in Equation 2 above. We can see that in the following manner:

Suppose we envision a spending policy with $MVPF^{spend} = \frac{\Delta W}{\Delta E - \Delta C}$ and we imagine closing the budget constraint with an alternative revenue-raising tax policy, call it $MVPF^{tax}$. As explained above, this policy only increases social welfare iff $\bar{\eta}_1 MVPF^{spend} > \bar{\eta}_2 MVPF^{tax}$. Assume for the moment that $\bar{\eta}_1 = \bar{\eta}_2$. In the

⁹When comparing ratios, it is important to compare policies of similar sizes. As Hendren and Sprung-Keyser [2020] write: “Consider the case where Policy 1 was a \$1M government expenditure and Policy 2 was a \$2M government expenditure. Comparing Policy 1 and Policy 2 would require the MVPF for a version of Policy 1 that is scaled up to cost \$2M. This same logic would also apply if considering a large-scale expenditure on a policy that had previously been analyzed with a narrower RCT – one would have to make the additional assumption that the average treatment effect of this expanded policy is given by the effect identified in the RCT.” In order to translate these ratios into total benefits, one need multiply by the size of net government expenditures.

¹⁰They refer to this second metric as the “marginal social value of government expenditures” and argue the appropriate acronym for this object should be MVPE.

¹¹Formally, $\frac{NSB}{\Delta E} = \frac{\Delta W - (1 + \phi)(\Delta E - \Delta C)}{\Delta E} = (1 + \phi)[BCR - 1]$. For a fixed value of ϕ this will produce an identical ranking of policies as the BCR.

net social benefit framework, the welfare consequences of the combined policy are as follows:

$$SB = \Delta W - MVPF^{tax} (\Delta E - \Delta C)$$

Here, this budget neutral policy would increase welfare if and only if:

$$MVPF^{spend} = \frac{\Delta W}{\Delta E - \Delta C} > MVPF^{tax}$$

These expressions demonstrate that the NSB equation is simply a special case where the budget constraint is closed using a linear income tax, ie $MVPF^{tax} = 1 + \phi$. We see this as a reason to prefer the MVPF framework. It treats the return on spending and revenue raising as separable, but it allows them to be combined to search for welfare improvements.

4.4 Distributional Considerations

The Net Social Benefits criterion is typically used to search for policies where $NSB > 0$. As explained in García and Heckman 2022b, any such policy “enhances the social utility possibility frontier.” The trouble with such a condition, $NSB > 0$, is that it sets aside any distributional concerns. In practice, one may wish to implement a policy that has $NSB < 0$ if the benefits disproportionately flow to a disadvantaged population.

Suppose, for example, that we are evaluating a policy that expands preschool programs to low-income children. The traditional NSB approach imagines that we finance this policy with an increase in the linear income tax. This means that the costs of the policy disproportionately fall on the affluent, whereas the benefits of the policy accrue to low-income children. If society has a preference for redistributing from rich to poor, it may be desirable to pursue this policy even if the net social benefits are negative.¹²

The MVPF approach helps to clarify why $SB > 0$ may not always be the right test. It enables researchers to consider distributional incidence in a straightforward manner. Recall from section 2 that a combined budget-neutral policy increases social welfare iff $\bar{\eta}_1 MVPF_1 > \bar{\eta}_2 MVPF_2$, where $\bar{\eta}_j$ is the average social marginal utility of the beneficiaries of policy j .¹³ This expression provides intuition for distributional consequences in the MVPF framework. For example, imagine two policies with $MVPF_1 = 1$ and $MVPF_2 = 2$. Would it increase welfare to spend on policy 1 financed by policy 2? Only if providing \$1 to beneficiaries of policy 1 is valued more than providing \$2 to beneficiaries of policy 2.

Hendren and Sprung-Keyser [2020] illustrate the value of this approach using the example of the Omnibus Budget Reconciliation Act of 1993. This policy change led to an increase in the top marginal income tax rate and an expansion of the Earned Income Tax Credit (EITC). They find that the MVPF of the top tax rate increase is 1.85; whereas the expansion of the EITC is 1.12. This suggests that the policy took \$1.85 from affluent individuals for each \$1.12 that it provided to EITC beneficiaries. A simple net social benefit test would suggest the 1993 reform was not desirable – its net social benefits were negative. By contrast, the MVPF approach frames the debate as a question of the social preference for redistribution. Whether society prefers this policy depends on one’s social preferences: does one wish to provide \$1.12 to an EITC beneficiary

¹²In order to account for this and to restore the $SB > 0$ condition, one must modify the NSB equation given in García and Heckman [2022a] to incorporate distributional weights. In particular, let $\bar{\eta}$ denote the welfare weight society places on the low-income children in the preschool program relative to those paying the income tax. The weighted social benefit is then given by:

$$WSB = \bar{\eta}\Delta W + (1 + \phi) (\Delta E - \Delta C)$$

¹³Written another way, one prefers the policy if and only if $\frac{\bar{\eta}_1}{\bar{\eta}_2} > \frac{MVPF_2}{MVPF_1}$.

if it imposes a cost of \$1.85 on a top earner? The MVPF approach helps quantify the equity tradeoffs faced in policy decisions. Focusing on a single condition, $NSB > 0$, can obfuscate these concerns.

5 Misconceptions about the MVPF

In this final section, we discuss five misconceptions about the MVPF approach that are expressed in previous literature. These include statements that the MVPF assumes a fixed government budget, that the MVPF assumes optimality of existing government programs, that the MVPF does not account for the social cost of raising revenue, that the MVPF misses potential Pareto improvements, and that the MVPF - and other ratio statistics - rely on arbitrary decisions about what goes into the numerator and the denominator. Each of these concerns are without merit. We discuss them in turn.

5.1 Fixed Government Budget

Some have suggested that the MVPF approach cannot be used to shed light on the optimal size of the government budget. This argument is misguided.¹⁴ The MVPF certainly can be used to examine whether expanding the government is desirable and to characterize the optimal size of the government budget. As noted above, this simply requires comparing the MVPFs of tax policies alongside the MVPFs of expenditure policies. The MVPF of a tax change tells us how much individuals are willing to pay to avoid the tax increase per dollar of net government revenue that is raised.¹⁵ So, suppose we have a proposal for a new expenditure with an $MVPF = A$. This policy delivers \$A benefits per dollar of net government cost. Suppose we also have a proposal to finance this spending with a tax increase that has an $MVPF = B$. For every \$1 of revenue raised, the tax increase imposes a welfare cost of \$B. This could be a linear income tax, but it could also be any other method of raising revenue.¹⁶ In this set-up, the government should expand the size of the budget if and only if $A > B$.¹⁷ In fact, the government should continue to increase the size of its budget until $A = B$.

As noted above, this logic is akin to the test of whether $NSB > 0$. Policies may have diminishing marginal returns as they scale, but each marginal expansion of Policy A is welfare enhancing as long as $A > B$.¹⁸

5.2 Arbitrariness of Ratios

Some have argued that ratios serve as ineffective welfare criteria because the placement of terms in the numerator and denominator can be arbitrary. For example, Boardman et al. [2017] argue in a prominent

¹⁴For example, García and Heckman [2022a] write that the MVPF “evaluates programs for a fixed government budget, whereas traditional cost-benefit analyses also consider the benefits of expanding the size of the budget to finance new programs.” and that “MVPF is silent on the optimality of the government budget and does not directly address what its size should be. It does not provide a proper criterion for evaluating programs that expand or contract the size of the government”.

¹⁵For example, Hendren and Sprung-Keyser [2020] show that the 1993 increase in the top marginal income tax rate had an MVPF of 1.85, suggesting it delivered a welfare loss of 1.85 for every \$1 of government revenue collected.

¹⁶For example, one could consider the MVPF of a consumption tax, such as a gas tax, or increased enforcement of existing taxation via expanded audits. The MVPF framework offers flexibility in how the government chooses to close or relax its budget constraint.

¹⁷As noted above, if the tax and spending policies have different distributional incidence, then one needs to ask whether one wishes to provide \$A to the spending beneficiaries at a cost of \$B to those paying the taxes.

¹⁸The logic here reinforces the notion that, when comparing ratios, it is helpful to compare policies of similar sizes. As Hendren and Sprung-Keyser [2020] write: “Consider the case where Policy 1 was a \$1M government expenditure and Policy 2 was a \$2M government expenditure. Comparing Policy 1 and Policy 2 would require the MVPF for a version of Policy 1 that is scaled up to cost \$2M. This same logic would also apply if considering a large-scale expenditure on a policy that had previously been analyzed with a narrower RCT – one would have to make the additional assumption that the average treatment effect of this expanded policy is given by the effect identified in the RCT.” The construction of budget neutral policy comparisons requires that MVPF estimates are valid for the relevant-sized policy changes.

textbook on cost-benefit analysis that “ratios are subject to manipulation.”¹⁹ While this argument may apply to the traditional cost-benefit ratios, the structure of the MVPF avoids this concern.

Consider, for example, a policy of expanding health insurance subsidies to low-income adults. The cost of these subsidies may be partially offset by a reduction in Medicaid reimbursements to hospitals. This is because the policy reduces Medicaid’s obligation to compensate hospitals for healthcare provided to the uninsured.²⁰ Should the “upfront” cost of the policy be equal to total spending on the subsidies or the subsidy cost net of the reduced Medicaid outlays? The traditional approach is silent on this – practitioners must exercise their discretion to construct a benefit-cost ratio.²¹ Arbitrary decisions made by the practitioner limit the ability to use these statistics to compare across policies.

By contrast, when constructing the MVPF, there is no judgement call necessary regarding the placement of terms in the numerator or denominator. All government costs are placed in the denominator and all benefits to individuals are placed in the numerator. So, in the case of the subsidy example above, the reduction in Medicaid spending goes in the denominator because it is a reduction in government cost. This set-up ensures logical consistency, avoids arbitrary deduction making, and enables intuitive comparisons across MVPF estimates. Indeed, the determinacy of the numerator and denominator in the MVPF is critical for ensuring that the MVPF provides a unified method of comparing policies.

5.3 Assuming Optimality of Existing Programs

When evaluating the impact of a policy on net government costs, the MVPF approach incorporates “fiscal externalities” that treat the status quo tax and transfer system as given. This has led some to argue that the MVPF approach assumes that existing government programs are set optimally.²² This, however, is not the case. The MVPF approach does not assume the optimality of any government policy. As with any welfare analysis, MVPF estimates can lead authors to invoke assumptions of individual optimization (e.g. the envelope theorem) when measuring ΔW . The MVPF approach examines the impact of the policy on the government’s budget (i.e. ΔC) in the context of existing tax and transfer programs. However, the MVPF approach does not assume these policies are set optimally.

To the contrary, the MVPF approach enables researchers to identify the sub-optimality of government policies. Indeed, a core conclusion of Hendren and Sprung-Keyser [2020] is that many policies targeting low-income children have infinite MVPFs. These are policies that pay for themselves in the long-run, suggesting that status quo policy has clearly been inefficiently underinvesting in low-income children. The government has failed to take advantage of potential Pareto improvements, and the MVPF approach helps identify those opportunities.

5.4 Ignoring the Social Cost of Raising Revenue

The MVPF approach also differs from traditional benefit-cost analysis because individual MVPF estimates do not include an adjustment for the deadweight loss from tax revenue (i.e. there is no “ ϕ ” in the MVPF

¹⁹García and Heckman [2022b] make a similar argument with regard to the MVPF, arguing that “practioners may differ in their judgement regarding whether a component is part of [the benefits, the net cost, or the upfront cost of a policy. The [net social benefit] is a determinant sum that ranks programs without any dependence on this arbitrary decision.”

²⁰Formally, offsets the costs of uncompensated care through Disproportionate Share Hospital (DSH) Payments.

²¹This concern applies not only to traditional cost-benefit ratios but also to the “net social benefits per dollar of upfront expenditure” ratio as proposed by García and Heckman [2022b] in their appendix.

²²For example, García and Heckman [2022a] write “The traditional approach does not necessarily assume that the basket of current government programs is optimally determined, whereas the Hendren and Sprung-Keyser approach does.”

expressions above). This has led some to conclude that the MVPF ignores the social cost of raising revenue.²³ This, however, is a misunderstanding of how the MVPF approach works. The MVPF approach does not ignore the distortions generated from raising taxes to pay for social programs with different budget requirements. As shown in Hendren and Sprung-Keyser [2020], the MVPF of tax changes measures the welfare cost of raising revenue. This can be directly compared to the MVPF of the proposed expenditure policy to assess whether the combined policy increases social welfare. As we note above in Section 5.1, increasing spending on a policy with $MVPF = A$ financed by a tax policy with $MVPF = B$ will increase welfare if and only if $A > B$ (abstracting from distributional concerns). Using this approach, the MVPF allows researchers to properly account for the social cost of raising revenue. In fact, researchers can draw upon the set of previously constructed MVPF estimates to incorporate a distortion associated with wide range of revenue raisers.

5.5 Identifying Pareto Improvements

The MVPF is generally defined to be the ratio of a policy’s benefits, ΔW , to the policy’s net cost, $\Delta E - \Delta C$. But, in cases where a policy provides positive benefits without imposing any net government costs, the MVPF is defined as infinite. These are policies that provide a Pareto improvement – the government can provide benefits to society without actually needing to spend any money. Some have mistakenly suggested that Pareto improving policies have a negative MVPF, and, therefore, the MVPF approach suggests these policies are inadvisable.²⁴ This is a definitional misunderstanding that leads to the incorrect conclusion. Indeed, a core finding of the MVPF analysis in Hendren and Sprung-Keyser [2020] is that many historical policies that invested directly in low-income children had infinite MVPFs. The paper argues that these are Pareto improving policies.

As noted previously, traditional cost-benefit analysis is not able to identify Pareto improving policies. To see this, consider a tax cut of $\Delta E = 1$. This provides $\Delta W = 1$ of benefits to the taxpayers receiving the tax cut. Suppose however that current tax rates are at the top of the Laffer curve. Taxpayers adjust their income in response to the tax cut and so tax revenue is increased by $\Delta C = 1$, fully offsetting the upfront cost. The MVPF of the tax cut is infinite, suggesting that there is a Pareto improvement to be achieved by lowering taxes. By contrast, the benefit-cost ratio of this policy is less than or equal to 2 (the BCR is given by $BCR = \frac{1+1+\phi}{1+\phi}$ where ϕ is the DWL associated with the method used to raise revenue to finance the tax cut). As illustrated in Hendren and Sprung-Keyser [2020], there are many policies that have benefit-cost ratios exceeding 2 that do not pay for themselves. This means that the traditional approach does not respect the Pareto principle, but the MVPF does.

6 Conclusion

The MVPF approach provides a flexible and transparent method for measuring the welfare impacts of government policies. For a given policy, it measures the policy’s benefits per dollar of net government spending. We believe this is a valuable way to measure a policy’s welfare effects. By focusing on net government expenditures rather than the return to initial expenditures, it generates an ordering of policies that respects the Pareto principle. (The same cannot be said of benefit-cost ratios.) This approach enables researchers to mea-

²³For example, García and Heckman [2022b] write: “Unlike the [net social benefit], the MVPF ignores ϕ and thus does not account for the distortions generated when raising taxes to pay for social programs with different budget requirements.”

²⁴For example, García and Heckman [2022b] write: “If the government faces a type of Laffer Curve, it could be that [the net government revenue is negative, $\Delta E - \Delta C < 0$]. The program is a money pump. While having a negative MVPF, it would nonetheless expand social welfare gauged by NSB.”

sure the welfare impacts of policies without imposing ad-hoc assumptions about how the budget constraint is closed. Moreover, the MVPF approach enables researchers to assess the role of distributional considerations in a straightforward way. We believe these are unique advantages of the MVPF over traditional metrics used in welfare analysis.²⁵

References

- W Steven Barnett and Leonard N Masse. Comparative benefit–cost analysis of the abecedarian program and its policy implications. *Economics of Education Review*, 26(1):113–125, 2007.
- Anthony E Boardman, David H Greenberg, Aidan R Vining, and David L Weimer. *Cost-benefit analysis: concepts and practice*. Cambridge University Press, 2017.
- Jorge Luis García and James J Heckman. On criteria for evaluating social programs. Technical report, Institute of Labor Economics (IZA), March 2022a.
- Jorge Luis García and James J Heckman. On criteria for evaluating social programs. Technical report, NBER Working Paper No. 30005, April 2022b.
- 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, 2020.
- 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, 2010.
- Nathaniel Hendren and Ben Sprung-Keyser. A unified welfare analysis of government policies. *The Quarterly Journal of Economics*, 135(3):1209–1318, 2020.

²⁵We encourage readers who are interested in the MVPF approach to visit www.policyimpacts.org. Further descriptions of the MVPF approach and its comparison to the traditional approach are provided at <https://policyimpacts.org/mvpf-explained/what-is-the-mvpf>.