The Advantages of the MVPF: A Comment on García and Heckman (2022)*

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

April 12, 2022

Evaluating the desirability of government policies requires the consistent application of a welfare metric designed to compare across policies. Hendren and Sprung-Keyser [2020] popularize the Marginal Value of Public Funds (MVPF) as a unified metric for evaluating the welfare consequences of government policy. In a recent working paper, "On Criteria for Evaluating Social Programs", García and Heckman [2022] question the validity of the MVPF and advocate for alternative metrics such as the Net Social Benefits criterion. In this comment, we address their criticisms of the MVPF and argue they are without merit.

We proceed in three 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 other methods of cost-benefit analysis and highlight its unique advantages. Third, we clarify several incorrect statements about the MVPF approach made in García and Heckman [2022]. We welcome any further feedback on the MVPF framework and encourage readers to reach out to us with any questions.

1 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. Using the notation from García and Heckman [2022], let ΔW denote the benefits that the policy provides to individuals in the population. Let ΔE denote the government's upfront expenditure on the policy. And let ΔC denote the reduction in government costs due to the the causal effect of the policy.

*We thank Jamie Emery for excellent research assistance, and we are grateful to Bob Hahn and Robert Metcalfe for helpful comments.

(This captures any long-run government savings, such as increases in taxes collected if a policy increases earnings.) Hendren and Sprung-Keyser [2020] define the MVPF as[1]:

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

For a given policy with $MVPF = A$, the policy delivers \$$A$ of benefits per dollar of net government spending.

Given a set of MVPF estimates, one can then conduct welfare comparisons across policies. In particular, take any two policies where the MVPF of the first policy is $MVPF = A$ and the MVPF of the second policy is $MVPF_2 = B$. Now, consider increasing 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 $\bar{\eta}_1 A > \bar{\eta}_2 B$, where $\bar{\eta}_j$ is the average social marginal utility of the beneficiaries of policy $j = 1, 2$. The MVPF framework formalizes the tradeoff of moving a dollar between different groups of beneficiaries and is transparent about the distributional incidence of doing so.[2]

2 Comparison to Traditional Approaches

García and Heckman [2022] argue that the MVPF is inferior to traditional metrics for conducting welfare analysis. In this section, we compare the MVPF to the two more traditional metrics discussed in García and Heckman [2022]: the benefit-cost ratio (BCR) and the net social benefits (NSB) of a policy. We argue that the MVPF has unique advantages relative to both approaches.

2.1 Comparison to BCR

García and Heckman [2022] compare the MVPF to a traditional benefit-cost ratio (BCR). In their note, they define the benefit-cost ratio as[3]:

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

[1] García and Heckman [2022] provide an incorrect definition of the MVPF. They use the following expression:

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

Their definition adjusts the denominator 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. This adjustment is motivated by the incorrect belief that metrics for welfare anlaysis must close the budget constraint. By contrast, the correct denominator in the MVPF captures net government revenue and does not envision any specific method of closing the budget constraint. As discussed in Section 2.2.3, 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.

[2] In 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 2.3.2.

[3] It is worth noting that this definition in García and Heckman [2022] differs from the definition of the BCR presented in much of previous literature (E.g. Heckman et al. [2010]). In many cases, the BCR is given by

$$BCR = \frac{\Delta W + \Delta C}{\Delta E\,(1+\phi)} \tag{2}$$

This prior definition does not adjust the government savings term (ΔC) by the deadweight loss of taxation $1+\phi$. While we argue that welfare analysis should avoid budget constraint assumption embedded in these $1+\phi$ adjustments, we believe the current definition represents an improvement over the previous approach. If increases in government revenue are adjusted based on the cost of raising revenue, then it follows that government savings should be adjusted to account for savings due to reduced distortions in the future.

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 assumed 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.

Despite these differences, García and Heckman [2022] argue that there is "no special advantage" associated with the MVPF.[4] In truth, the differences between the MVPF and the BCR are quite consequential. We discuss each in turn.

2.1.1 Capturing Government Savings

The MVPF measures the benefits of a policy per dollar in 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 $\Delta E = 1$ and a value to participants 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 $\Delta W = 3$ 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.[5] 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 BCR is $\frac{3}{1+\phi}$. Thus, the traditional cost-benefit 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 which paid for itself. 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.

2.1.2 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 2.2.3, this

[4] In their note, García and Heckman also argue that both the MVPF and the BCR are inferior to net social benefits (NSB), the metric discussed in Section 2.2. We still discuss the differences between BCR and MVPF because the use of BCR is widespread in the literature, including in previous and recent work by García and Heckman [García et al., 2020, 2021, Heckman et al., 2010].

[5] This is true regardless of the value of ϕ.

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	∞

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 chooses to construct a budget neutral policy, that can be done by examining the MVPF of a particular policy that raises the necessary revenue, rather than assuming a one-size-fits-all distortion applies to all revenue increases.

2.2 Comparison to Net Social Benefit

García and Heckman [2022] define the net social benefit (NSB) of the policy as

$$NSB = \Delta W - (1+\phi)(\Delta E - \Delta C) \tag{4}$$

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 order to argue for the superiority of the NSB approach, García and Heckman [2022] use the example presented in Table I.

In this example, 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 ($\Delta C = 100$).[6] 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 [2022] argue that the MVPF can provide poor guidance for policy. We believe this example does not prove the superiority of the NSB over the MVPF. The criticism in García and Heckman [2022] falls short for four reasons.

2.2.1 Upfront versus Net Spending

The conceptual experiment outlined in García and Heckman [2022] relies heavily on the assumption that the government has a fixed quantity of funds for upfront investment, ΔE. It measures programs by the size of their upfront cost rather than the size of their net cost. In order to see the trouble with this approach, imagine

[6] García and Heckman [2022] 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.

a policy 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 has a much lower NSB. That said, this policy only costs the government \$1. The government could implement the 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. This trouble arises because the example above categorizes policies based on their upfront cost, rather than their net cost.[7]

2.2.2 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.

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 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 mechancially true that the NSB of Perry Preschool is low, but only because the program itself was small.[8] 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]

2.2.3 Closing the Budget Constraint

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

[7]It 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.

[8]Taken to its logical extreme, this argument implies that all policies evaluted with regression discontinuity designs have a NSB of approximately 0. The policy experiment in those papers is only defined in infinitessimally small regions around the threshold, and so the size of the policy is approximately zero.

[9]When 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 benfits, one need multiply by the size of net government expenditures.

.

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

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

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

We can see from these expressions that the NSB equation in García and Heckman [2022] is 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.

2.2.4 Distributional Considerations

García and Heckman [2022] suggest that a social planner should implement all policies with positive net social benefit, $SB > 0$. That said, their discussion sets aside any distributional concerns. In practice, one may wish to implement a policy that has $SB < 0$ if the benefits disportionately 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.[10]

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 1 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 is the average social marginal utility of the beneficiaries of policy j.[11] 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. By constrast, 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 if it imposes a cost of \$1.85 on a top earner? The MVPF approach helps quantify the equity tradeoffs faced in policy decisions.

[10] In order to account for this and to restore the $SB > 0$ condition, one must modify the NSB equation given in García and Heckman [2022] 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:

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

[11] Written another way, one prefers the policy if and only if $\frac{\bar{\eta_1}}{\bar{\eta_2}} > \frac{MVPF_2}{MVPF_1}$.

3 Misconceptions about the MVPF

In this final section, we discuss a couple of misconceptions about the MVPF approach that are expressed in García and Heckman [2022]. In the opening paragraph of their article, García and Heckman levy two criticisms of the MVPF. First they write:

> [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.

Next, they 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.

Both of these statements are inaccurate. We discuss each in turn.

3.1 Fixed Government Budget

García and Heckman [2022] argue that the MVPF takes as given a fixed govenrment budget. In particular, they write *"The MVPF takes $\bar{E}$ as given and asks for programs with the highest value added per unit expenditure for a fixed budget."*

This is false as the MVPF approach does not take as given any amount of spending on programs. To the contrary, the MVPF approach can be used to determine whether expanding government is desirable. As noted above, this simply requires comparing the MVPFs of tax policies alongside the MVPFs of expenditure policies. 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 another policy that has $MVPF = B$. For every \$1 of revenue raised, this policy imposes a welfare cost of \$$B$. This could be a linear income tax, but it could also be any other method of raising revenue.[12]

In this set-up, the government should expand the size of the budget if and only if $A > B$.[13] In fact, the government should continue to increase the size of its budget until $A = B$.

3.2 Assuming Optimality of Existing Programs

García and Heckman [2022] claim that the MVPF approach assumes that existing government programs are set optimally. This again is false.

The MVPF approach makes zero assumptions about optimality of 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, at no point does the MVPF approach 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

[12]For 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.

[13]As 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.

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.

4 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 measure 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. We encourage researchers to incorporate MVPF estimates into their work.[14]

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.

Jorge Luis García and James J Heckman. On criteria for evaluating social programs. Technical report, Institute of Labor Economics (IZA), 2022.

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.

Jorge Luis García, Frederik H Bennhoff, Duncan Ermini Leaf, and James J Heckman. The dynastic benefits of early childhood education. Technical report, National Bureau of Economic Research, 2021.

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.

[14]We 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.