MIT Open Access Articles

Welfare analysis meets causal inference

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Finkelstein, Amy and Nathaniel Hendren. "Welfare analysis meets causal inference." Journal of Economic Perspectives, 34, 4 (September 2020): 146–167 © 2020 The Author(s)

As Published: 10.1257/JEP.34.4.146

Publisher: American Economic Association

Persistent URL: <https://hdl.handle.net/1721.1/130347>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.

Welfare Analysis Meets Causal Inference

Amy Finkelstein and Nathaniel Hendren

Economists have made remarkable progress over the last several decades in developing empirical techniques that provide compelling evidence of causal effects—the so-called “credibility revolution” in empirical work (as discussed in this journal by Angrist and Pischke 2010). But while it is interesting and important to know what the effects of a policy are, we are often also interested in a normative question as well: Is the policy a good idea or a bad idea? Or in the more careful language of economics: What is the welfare impact of the policy?

Until recently, there had been relatively little effort to harness the gains of the “credibility revolution” for the goal of welfare analysis. Instead, we in the empirical public finance community have struggled with other approaches. One venerable tradition is to take an estimate of the benefits of an expenditure policy and compare it to the “cost” to the government. This cost is usually defined as the expenditures on the program, multiplied by 1 plus “the marginal cost of public funds,” which is designed to take account of the distortionary effects of the taxation needed to finance the policy, which everyone “in the know” knows to be 0.3, or maybe 0.5 if you’re feeling pessimistic. Thus, buried in the last section of an empirical paper that painstakingly estimates the impact of a policy, is an ad hoc analysis that compares the benefits to the cost, multiplied by a smudge factor of 1.3 (for example, Finkelstein and McKnight 2008; Olken 2007). The other common method is the “marginal

■ *Amy Finkelstein is Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. Nathaniel Hendren is Professor of Economics, Harvard University, Cambridge, Massachusetts. Their email addresses are afink@mit.edu and nhendren@fas.harvard.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at <https://doi.org/10.1257/jep.34.4.146>.

excess burden” or “deadweight loss” approach, which requires valiant attempts to separate non-distortionary income effects of policies (which are transfers rather than welfare losses) from their distortionary substitution effects (which lead to deadweight losses). As Goolsbee (1999) has lamented, “The theory largely relates to compensated elasticities, whereas the natural experiments provide information primarily on the uncompensated effects.”

Fortunately, glimpses of light have appeared at the end of the empirical-welfare tunnel. In this essay, we describe a transparent framework for mapping empirical estimates of causal effects of a public expenditure (or tax) change to welfare analysis of that policy change. Following Hendren and Sprung-Keyser (2020), we refer to it as the “marginal value of public funds” (MVPF). The MVPF is the ratio of the marginal benefit of the policy to the net marginal cost to the government of the policy; crucially, this net marginal cost is inclusive of the impact of any behavioral responses to the policy on the government budget.

Our goal is not to break new theoretical ground with the MVPF framework—as we will discuss, its mathematical formulation has been around for decades (for example, Mayshar 1990). Instead, we proselytize for the underrecognized empirical usefulness of this approach in the wake of the “credibility revolution.” Its key advantage is that it relies on the causal effects of policy changes to conduct empirical welfare analysis. We provide guidance on how to implement and interpret this approach, with the hope that it will facilitate empirical welfare analysis across a variety of fields.

To do so, we start with a benchmark case of a small increase in a cash transfer that only affects its recipients, whose response to the policy is privately optimal. Under these assumptions, we show that estimates of causal effects of the policy are needed only for estimating the policy’s costs, not its benefits. We discuss how this logic is adapted as we relax each simplifying assumption in the benchmark case.

Once we estimate the MVPF, how can we use it? An MVPF of, say, 1.5 means that every \$1 of net government spending provides \$1.50 of benefits to the beneficiaries of the policy—or in other words, the beneficiaries would be willing to pay up to \$1.50 for that \$1 policy. One can use the MVPF to compare this “bang for the buck” of policies that affect the same group of individuals (the policy with the higher MVPF is preferable) or for comparing policies that affect different groups. In the latter case, the MVPF quantifies the implicit tradeoffs involved: given policies A and B, policy A is preferred to policy B if and only if one prefers giving $MVPF_A$ to policy A beneficiaries over giving $MVPF_B$ to policy B beneficiaries. Of course, economists have no special powers that allow them to declare such tradeoffs are appropriate.¹ But economics can clarify the tradeoffs embodied in the policy decisions society faces.

¹In other words, the National Bureau of Economic Research prohibition for its working papers against “statements regarding which policies should (or should not) be adopted” (at http://papers.nber.org/wpsubmit/wp_submit.html) encodes a fundamental and important recognition of the limits of economic analysis.

We then endeavor to answer some common and natural questions about the MVPF approach, including how it relates to “traditional” public finance welfare tools like marginal excess burden and marginal cost of public funds. Finally, we offer some examples of how the MVPF approach has or can be applied to some recent empirical applications across a variety of fields, including public finance, labor economics, development economics, trade, and industrial organization.

How to Construct the Marginal Value of Public Funds: An Initial Illustration

The MVPF is defined as the ratio of the marginal benefit of the policy to the marginal cost of the policy. Equivalently (and more usefully for operationalizing it), it is the ratio of the beneficiaries’ willingness to pay for the increase in expenditure out of their own income to the net cost to the government of the increase in expenditure per beneficiary:

$$\text{MVPF} = \frac{\text{Beneficiaries' Willingness to Pay}}{\text{Net Cost to Government}}.$$

Let’s consider how to calculate the MVPF for a \$1 change in cash benefits in a public program. This could be, for example, a means-tested cash welfare program like Temporary Assistance to Needy Families (TANF) or a means-tested tax credit such as the Earned Income Tax Credit (EITC). For concreteness, we’ll talk about a cash increase but we could just as easily consider a cash decrease; the MVPF would be the same number because both willingness to pay and cost would be negative. For this initial example, we make several assumptions that we’ll relax later in the discussion: a cash (not in-kind) transfer; the policy change is small; individuals exhibit privately optimal behavioral responses to the policy change; and no impacts of the policy on the people who were not the policy’s direct recipients.

Costs

Consider first the denominator of the expression for the MVPF: What is the cost of increasing the program’s cash benefits by \$1? It is useful to think of two different classes of cost: the mechanical cost and the fiscal externality. The “mechanical cost” of the policy is the increase in government expenditures due to the policy in the absence of any behavioral response. If the number of infra-marginal individuals who were already receiving the cash transfer policy is I , then the mechanical cost of increasing payments by \$1 for each infra-marginal recipient will be $\$I$.

The “fiscal externality” from the policy captures the effect of any behavioral response to the policy on the government’s net budget outlays. For example, if individuals reduce their labor supply to become eligible for additional welfare benefits, this will reduce tax revenue collected by the government on earnings. Conversely, if individuals enter the labor force in order to become eligible for an expanded Earned Income Tax Credit and this decreases their use of other government

transfer programs such as food stamps, this will increase net government revenue (but also raise government costs from increased EITC payments). The fiscal externality must account for the full impact of any behavioral response by both marginal and infra-marginal recipients on government spending and tax revenue: for example, it must consider changes in government net revenue arising from changes in eligibility for (and hence spending on) other public programs, changes in sales taxes from modified consumption patterns, changes in public health care spending through Medicare and Medicaid (if the program affects health), and so on.

This concept of a policy's fiscal externality is where the applied econometrics literature on causal inference connects directly with welfare analysis. The fiscal externality logic clarifies which causal effects are and are not necessary for estimating program costs for purposes of welfare analysis. Specifically, it is sufficient to estimate the net impact of the increase in benefit levels on net government revenue, without decomposing the impact into these various channels (Kleven and Kreiner 2005). For example, a large literature has analyzed a wide variety of potential effects of an increase in the level of unemployment insurance benefits on a range of behaviors including unemployment duration (for a recent review, see Shmieder and von Wachter 2016), exit rates into unemployment (Jäger, Schoefer, and Zweimüller 2019) and re-employment wages (Nekoei and Weber 2017). For welfare analysis, however, one needs the net impact of these behavioral changes on the government budget; the individual channels of response are neither necessary nor sufficient.

Benefits

Now consider the numerator of the expression for the MVPF: the benefits from the \$1 increase in cash transfer—that is, the willingness to pay by recipients out of their own income for \$1 more of the cash transfer. In many cases, this is harder to estimate than the costs of the program (although, deliberately, it is not hard in our first example). But a key insight of the MVPF framework is that they need not depend directly on behavioral responses to—or causal effects of—the policy.

It is useful to distinguish between two classes of recipients of the transfer. For the infra-marginal recipients who were already receiving the cash transfer, the \$1 transfer is valued at \$1. How much would you be willing to pay for an extra dollar? One dollar. But for the marginal recipients who change their behavior in response to the change in policy and thus become newly eligible for the transfer, how much do they value their new benefit? For example, suppose they decrease the amount they work to become income-eligible for the cash policy. Well, if they are making privately optimal decisions, they must be indifferent to changing their behavior. Why? Because they had already chosen their behavior (in this example, the hours they work) at the optimal level—balancing private costs and benefits of another hour of work—under the old policy, and we've made only a very small (\$1) change in the policy. More generally, if the policy change is small and individuals are making privately optimal decisions, the private cost of whatever behavioral change

the marginal recipient undertakes to become eligible for the benefit is equal to the private benefit from becoming eligible.²

Given these assumptions, the \$1 increase in the cash transfer has no welfare effect on marginal beneficiaries (in other words, those who change their behavior in response to the policy change in order to become newly eligible). The willingness to pay is just \$1 times the number of infra-marginal beneficiaries (I). Note that we have not needed to estimate any causal effects to determine the benefits of the policy for either inframarginal or marginal recipients. In other words, despite a large empirical literature (to which we plead guilty of making contributions) on the potential benefits of public policies—the impacts of unemployment insurance on eventual re-employment wages, the impacts of health insurance on health and consumption smoothing, and so on—in this benchmark case, these studies do not directly inform recipient willingness to pay. Causal effects are needed only for the fiscal externality cost term in the denominator because in that setting, an agent who is making (by assumption) privately optimal behavioral changes in response to the policy change will not internalize the external effects of the policy on the government budget.³

Putting it Together

It is convenient to normalize the willingness to pay (the numerator of the MVPF) by the mechanical cost to the government. Recall that this mechanical cost was \$ I , which was also the willingness to pay for the policy. Thus, the MVPF of a \$1 increase in cash benefits is:

$$\text{MVPF}^{\$1} = \frac{1}{(1 + FE)}$$

where FE denotes the 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 may have a positive net effect on the government budget (say, by improving health and reducing public spending on health care) or a negative net effect on the government budget (say, by discouraging work effort).

The MVPF measures a policy's bang-for-the-buck. Every \$1 of net spending on a tax cut delivers \$MVPF of benefits to the recipients of that tax cut. Conversely, every \$1 of net revenue raised through a tax increase imposes a cost equivalent to \$MVPF

²This is known as the envelope theorem. The envelope theorem guarantees that behavioral responses to marginal policy changes by utility-maximizing 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. Behavioral responses to policies therefore have first-order effects on policy *costs*—because of the fiscal externality—but only second-order effects on recipient welfare—because of the envelope theorem.

³By the same token, if the behavioral responses to the policy have external effects on other individuals besides recipients of the policy, these effects would also have to be taken into account. We cover this possibility 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.”

to the beneficiaries. Spending more resources on policies with higher MVPFs delivers greater welfare to those beneficiaries per dollar spent; raising revenue from policies with lower MVPFs does so with lower welfare loss to those paying for the revenue.

How to Use the Marginal Value of Public Funds for Welfare Analysis

Now that we know how to construct the MVPF (at least for one specific example; we'll discuss more applications in a moment), what do we do with it? For example, an MVPF of 1.3 means that every \$1 of net government spending provides \$1.30 of benefits to the beneficiaries of the policy, or in other words, the beneficiaries would be willing to pay up to \$1.30 for that \$1 policy. But is an MVPF of 1.30 "good"? What about an MVPF of 0.8? In other words, what do we do with the MVPF once we've estimated it?

We start with a special case where welfare analysis with the MVPF is comparatively easy: when a policy's net cost to the government (the denominator) is negative and the willingness to pay (the numerator) is positive. In this case, the government spending "pays for itself." Hendren and Sprung-Keyser (2020) define such policy to have an infinite MVPF.⁴ A classic example would be if cutting tax rates increases net tax revenue, perhaps because of an increase in labor supply in response to the lower marginal tax rates. This is often referred to as being on the "wrong side of the Laffer curve" because the government can simultaneously *cut* taxes and *increase* revenue.

However, most government expenditures have net positive costs to the government. In this case, the most straightforward use of the MVPF framework is to compare two policies that seek to transfer benefits to the same group of people. For example, imagine a comparison of two policies designed to transfer resources to lower income individuals: expanding the earned income tax credit (a wage subsidy for low income workers) or expanding cash welfare benefits (a direct cash transfer to low-income individuals). If these policies have the same distributional incidence, then spending more money on the one with the higher MVPF is preferred. For the same cost, it creates more transfers to the targeted group. The higher MVPF policy gets more "bang for the buck." This means that one can construct a budget-neutral policy that increases individuals' welfare by spending more on the policy with the high MVPF, with such policy financed by reduced spending on the policy with the low MVPF.

Of course it is rare that two policies target exactly the same population. We even fudged a bit in our preceding example because potential recipients of cash welfare and of the earned income tax credit are overlapping but not identical groups. If two policies target two different groups, how can researchers use the marginal cost of public funds that they have calculated? One option is to compare the MVPF of the policy to a calibrated MVPF for a modification of the tax schedule

⁴At least one of us has wondered why they define a term that is negative as being infinite. The authors explain that they define it as an infinite MVPF to make clear it's "even better" than any finite MVPF.

at the same region of the income distribution (Hendren 2020). Another option is to use the MVPFs to quantify the tradeoff involved in making a budget-neutral change between the policies; in other words, taking a dollar from one policy and adding it to the other. Given two policies, A and B, spending more money on policy A financed by reduced spending on B generates $MVPF_A$ dollars of welfare gain for policy A beneficiaries and $MVPF_B$ dollars of welfare loss for policy B beneficiaries. So, if $MVPF_A = 1$ and $MVPF_B = 2$, this means one can take \$2 from policy B beneficiaries and generate \$1 for policy A beneficiaries.

Is such a transfer from group B to group A desirable? That depends on how one feels about these two different groups (sometimes referred to as their “social welfare weights”). If one places equal value on \$1 in the hands of A beneficiaries and \$1 in the hands of B beneficiaries, then the transfer from group B to group A would not be desirable—instead it would be desirable to increase spending on policy B and reduce spending on policy A. But if one values giving \$2 to policy-B beneficiaries less than giving \$1 to policy-A beneficiaries, one would prefer spending more on policy A financed by less spending on policy B.

How should one decide whether \$1 to group A is preferable to \$2 to group B? Perhaps by introspection. Or on philosophical grounds (Saez and Stantcheva 2016). We don’t have “satisfying” answers because economics don’t generally have a comparative advantage at specifying societal preferences. People disagree. But the marginal cost of public funds quantifies the tradeoff, which is a crucial first step in deciding whether one “likes” it or not. And it’s where economists can most directly contribute to these interesting and difficult questions.

Relaxing Assumptions

Now that we have some idea of how to construct and use the MVPF, we’d like to walk through a bunch of real-world applications. But before we can do so, we promised that we would discuss how to relax a bunch of the heroic assumptions we made for the sake of our “benchmark” example. Here we go.

What if the Policy Changes Are Large?

We considered a \$1 change to a policy. That was one of two key assumptions needed for the argument that for marginal recipients (that is, recipients who change their behavior in response to the policy change to become newly eligible), we could assume their willingness to pay was zero.

In practice, of course, many policy changes are large, and the approximation that marginal individuals who react to a policy change experience no net benefits may fail, perhaps spectacularly. For large policies, the marginal cost of public funds remains a useful guide, but measuring willingness to pay can be less straightforward because it now requires incorporating some value of benefits to marginal recipients. Kleven (forthcoming) provides a recent discussion of this point and some possible approaches for analyzing large reforms.

Fundamentally, we need an estimate of the marginal recipients' demand curve for the increase in public expenditure: that is, willingness to pay is the area under that demand curve. For a large increase in a public cash transfer, preexisting recipients still value the transfer at its dollar value (a dollar is still worth a dollar), but for recipients who change their behavior in order to access the larger public cash transfer, we need to know their willingness to pay for that cash transfer, net of the utility cost of their behavioral change.

Estimating demand is a bread-and-butter task of empirical economics, so we are in familiar—if sometimes empirically challenging—territory. It is all the more challenging when the good is not typically traded in a well-functioning market, so that demand cannot be directly estimated. In the example above, one standard approach (really just a short cut) is to count 50 percent of the increased transfer payments to marginal recipients; this 50 percent approximation follows from an assumption of linearity in the response function and the geometry of triangles (Plimpton 1800BC and Pythagoras 500BC). This approach is popular for its ease of implementation, if not necessarily, its realism; Finkelstein, Hendren, and Luttmer (2019) and Hendren and Sprung-Keyser (2020) are recent examples. More ambitiously, one can specify and estimate an economic model of behavior and use that to derive the demand system. Below, we discuss an application to the MVPF of an increase in import tariffs, using Fajgelbaum et al.'s (2020) constant-elasticity-of-demand system to estimate the welfare impacts on marginal actors.

What if Behavior Isn't Privately Optimal?

Our assumption that individuals make privately optimal decisions was the second key to being able to ignore welfare consequences for marginal recipients. However, a large literature in behavioral economics suggests that individuals commonly make mistakes. In this case, we can no longer assume that the welfare impact of the policy change for marginal recipients is zero, even if the policy change itself is small. For example, a \$1 increase in the cigarette tax may induce people to smoke less; if individuals smoke more than they would like to, their reduction in smoking may provide first-order benefits to them.

Here we find ourselves in the world of behavioral welfare analysis. It's no longer enough to estimate the marginal recipients' demand curve because their choices (demand) may not reveal their preferences. Either the researcher must assert that she "knows" the individual's utility function (for example, Bound et al. 2004; Finkelstein, Hendren, and Luttmer 2019) or try to elicit their true valuations, perhaps by experimentally eliminating bias and then eliciting demand (for example, see Allcott and Taubinsky 2015).

What if the Policy Provides an In-Kind Transfer?

Relaxing the first two assumptions led us to the observation that we needed to consider benefits for marginal recipients. To this point, we have also assumed that the transfers are in cash. This made life easy (or at least, easier) because it seems reasonable to assume that *inframarginal* recipients place a value of \$1 on receiving a

cash transfer of \$1. However, it is not obvious how inframarginal recipients value a \$1 increase of spending on in-kind assistance. In-kind transfers in the form of health care, education, housing, job training, and food are a substantial share of government expenditures in the United States and in other high-income countries (Currie and Gahvari 2008).

Consider an increase in \$1 of government spending on an in-kind benefit, such as additional government spending per pupil at public universities. Because we can no longer assume that the mechanical cost is valued by infra-marginal recipients dollar for dollar, we need to estimate the willingness to pay by infra-marginal recipients out of their own income per dollar of the in-kind benefit, W . The more general formula for the MVPF is then:

$$\text{MVPF}^{\text{inkind}} = \frac{W}{(1 + FE)}.$$

Here, W denotes the willingness to pay per infra-marginal recipient for the increased spending on their education. In the cash case, we knew that $W = 1$; \$1 of expenditures in the form of a dollar transfer is valued at \$1 by those who didn't change their behavior to receive it. But an in-kind transfer might be valued at less than the expenditure on it (in other words, $W < 1$), if it causes the infra-marginal recipients to consume more of the in-kind good than they would if given cash. Alternatively, an in-kind transfer might be valued at more than the government expenditure on it (in other words, $W > 1$) if the government can provide the good at lower cost than is available on the private market.

Estimating W can be relatively straightforward if the transferred good is also traded in the market at observed prices. In that case, estimating the demand curve for the good among the infra-marginal recipients gives us W . But inferring W becomes considerably more challenging when the expenditure is on an in-kind good or service that is not traded in a market—for example, increases in spending in public school, spending used to reduce pollution, or expanded public health insurance.

Of course, the empirical challenge of estimating willingness to pay when demand is not directly observed is not specific to the MVPF framework. Any form of welfare analysis must grapple with how to estimate the monetized value of specific goods and services. Fortunately, a range of techniques have been productively employed. One is to infer willingness to pay from other market transactions—such as wages if the good is bundled into workplace amenities, or house prices if the good is concentrated locally (for example, Rosen 1974; Greenstone 2017). Another is to calibrate (a fancy word for “make up”) a utility function for the goods delivered. This approach has been used, for example, in the literature valuing increased generosity of public unemployment insurance benefits (for example, Gruber 1997) or expansions of public health insurance eligibility (Finkelstein, Hendren, and Luttmer 2019). Another option is for the researcher to ask hypothetical questions to elicit the willingness of individuals to pay for a private good, such as health insurance (for example, Krueger and Kuziemko 2013). Yet another approach is to offer the good at

randomized prices and thus estimate willingness to pay directly, as Fischer, Frölich, and Landmann (2018) did for eliciting the value of health insurance in rural Pakistan. Finally, researchers can estimate the benefits of the policy itself and then attempt to monetize these benefits. For example, improvements in test scores are frequently mapped to monetary values through the relationship between test scores and earnings (as in Kline and Walters 2016) and researchers monetize estimated health benefits by relating their estimates to the value of a statistical life or a quality-adjusted life year (as in Currie and Gruber 1996; Goodman-Bacon forthcoming).

External Effects of Policies

So far we have (implicitly) focused on policies that have effects only on their intended recipients. However, many policies have indirect effects beyond the obvious beneficiaries. For example, health insurance subsidies to low income individuals may reduce uncompensated care costs to hospitals and therefore provide benefits to hospital shareholders as well (Garthwaite, Gross, and Notowidigdo 2018). A tax on carbon may affect not only those who use fossil fuels, but also those who benefit from reduced global warming. Vaccine subsidies may provide benefits not just to those receiving the vaccine, but also to those who do not obtain the vaccine and yet benefit from the reduced spread of a virus.

The MVPF framework readily captures these effects. The key extension is to measure the willingness to pay of everyone in the population affected by the policy, including those indirectly affected by the change in the policy. The formulas remain the same as in the above examples, except that the estimation of willingness to pay for inframarginal individuals now includes people who are not direct recipients of the policy.

For example, consider the MVPF of a \$1 subsidy to the price of vaccines, which presumably generates positive (health) externalities on the population. As before, 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 changes in labor market participation and productivity which may affect income tax revenue.

What about the benefits of this \$1 subsidy to vaccines? The group of infra-marginal recipients who were already getting the vaccine value the \$1 decrease in its price: \$1. Again, the group of marginal recipients who choose to get the vaccine because of the price reduction have no net welfare change because they are indifferent between not receiving the additional subsidy and not changing behavior, or receiving the additional subsidy and changing behavior (under the assumption that they were already behaving in a privately optimal manner). However, the fact that these individuals become vaccinated may generate external effects on the health of the rest of the population. The magnitude of these welfare effects depends on the magnitude and the sum total of any benefits (positive or negative) for the externally affected population from the increased vaccination of marginal recipients,

as measured by the externally affected population's willingness to pay; this is not equal to 1 but rather captures their willingness to pay for the marginal beneficiaries to be subsidized to obtain the vaccine. The more others benefit from the vaccine, the higher the MVPF. One would need to estimate the willingness to pay for non-recipients and calculate the numerator as the average willingness to pay across the infra-marginal recipients (those who value the subsidy at \$1) and the externally affected population (those who value the increase in the number of people receiving the vaccine by some amount W that would have to be estimated).

In contrast, a \$1 subsidy on carbon emissions could have negative externalities. This means the MVPF of a carbon subsidy will be lower than it would be in the absence of these externalities. In this sense, carbon taxes will impose less welfare loss on individuals per dollar of government revenue raised—it will be a more desirable tax than in the absence of the externalities.

Frequently Asked Questions

Why Doesn't Welfare Analysis Have to Think about How Policies Are Financed and the Distortionary Costs of Taxation (as in the Marginal-Cost-of-Public-Funds Approach)?

As we mentioned near the start of this paper, a common approach to welfare analysis is to try to measure the benefits of a policy change and then compare this to “the cost” of raising revenue to pay for the policy, which in turn is commonly defined as expenditures on the policy multiplied by 1 plus “the” marginal cost of public funds. Conventional wisdom usually places this cost somewhere between 0.3 (Poterba 1996) and a more conservative 0.5 (Heckman et al. 2010). In short, the marginal cost of public funds approach seeks to account for the distortionary cost of raising the tax revenue to finance that expenditure. Presto: welfare analysis.⁵

Most commonly, the marginal cost of public-funds approach imagines that the revenue for an expenditure is raised through a linear tax on income that leads to distortions in behavior. However, it has been recognized that this is not the only way to raise revenue, and as a result, there is no single marginal cost of public funds (Kleven and Kreiner 2006). The marginal cost of funds will vary depending on whether we increase taxes on the rich or reduce transfers to the poor. For some taxes, such as carbon taxes, the marginal cost of public funds is potentially negative because taxing carbon can have large benefits in the long run that offset its costs today.

By contrast, an attraction of the MVPF approach is that it severs spending analysis from revenue-raising analysis. We can then think separately about the MVPF both of the spending policy and of various policies to finance it—including reduced spending on other policies, increases in specific taxes, or deficit financing. Thus, the MVPF approach “closes the budget constraint” by comparing two MVPFs to

⁵Of 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.

form (hypothetical) budget-neutral policies rather than assuming a specific form of (hypothetical) financing for the policy, as the MVPF does.

Why Don't Researchers Need to Estimate Income and Substitution Effects of the Policy Separately (as in the Marginal-Excess-Burden Approach)?

The other common approach in public finance to welfare analysis is the concept of the deadweight loss of a policy (due to Harberger 1964) and its extension to marginal deadweight loss—also known as marginal excess burden (due to Auerbach 1985; Auerbach and Hines 2002).

The marginal excess burden of a tax change is commonly defined as the welfare impact of conducting the policy and simultaneously requiring that the beneficiaries pay for it through individual-specific, lump-sum transfers (Auerbach and Hines 2002). Because the conceptual experiment involves not only the policy envisioned but also these compensatory transfers, calculating marginal excess burden requires measuring the “compensated” response to the policy that excludes the income effect.

It is well-known that estimation of the marginal deadweight loss 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, this literature has been steeped in the view that it is essential to separate out income effects from substitute effects of the policy, which is challenging to estimate. Moreover, this approach is based on an unrealistic thought experiment in which individual-specific, lump-sum taxation (a policy instrument that doesn't exist) is used to finance the policy. Once again, the distinction between the MVPF and the marginal excess burden approach comes down to how the budget constraint is closed; here, the marginal excess burden approach imagines hypothetical lump-sum taxes, whereas, as discussed earlier, in the MVPF approach, one compares MVPFs of two policies to form hypothetical budget-neutral policies.

How Does the Marginal Value of Public Funds Framework Account for Policies that Affect a Diverse Group of Beneficiaries?

Policies rarely affect a homogenous group of people. Once there are different kinds of beneficiaries to a policy (either because the direct recipients are a heterogeneous group or because of external effects), welfare analysis needs to take account of the fact that societal preferences over transferring resources to different groups may differ. In terms of the example we discussed above, the beneficiaries of a subsidy for health insurance to low-income consumers may include not only the low-income recipients, but also hospital shareholders.

When a policy affects diverse groups, the MVPF is still constructed as previously described. However, it becomes more difficult to think about whether an MVPF of 0.8 or 1.3 is “good” or “bad.” To do so, one wants to take account of societal preferences toward the various recipients *within* the group of beneficiaries of a policy. If (for the sake of concreteness) one places lower social value on providing benefits to hospital shareholders than to low income individuals, then for a given MVPF, the

policy will be less desirable if more of the benefits accrue to shareholders than if they accrue to low-income recipients.⁶

Isn't This an Old Idea That's Been around for a Long Time?

Yes. The core ideas of the MVPF are explored in impenetrable detail in Hendren (2016), which itself notes that the mathematical definition of the MVPF is not new.⁷ It was initially proposed by Mayshar (1990), where it was referred to (incorrectly) as the “marginal excess burden.” In related work by Slemrod and Yitzhaki (1996, 2001) and Kleven and Kreiner (2006), it is referred to as the “marginal cost of funds” (or “marginal benefit of projects” in the case of expenditures).

Likewise, the idea of a fiscal externality is not new: It traces back at least to Ramsey (1927), although its crystallization and importance has become more salient recently (for example, Feldstein 1999; Saez 2004; Kleven and Kreiner 2005; Chetty and Saez 2010). Finally, the key insight that when small shifts in incentives lead to behavioral shifts, the net welfare effect on individuals is zero has been used extensively in previous empirical welfare analyses, including Harberger (1964). Our desire to clarify and illustrate the approach lies not in its novelty but in its usefulness: the fundamental novelty of the MVPF approach is not its mathematics, but its empirics: it relies on the causal effects of the policy and therefore provides a path to welfare analysis that leverages the tools generated in the credibility revolution.

Applications

In this section, we aim to reinforce the ideas behind the MVPF approach, as well as the usefulness of this approach, by giving some examples of how it has or can be applied in a variety of fields.

Income Tax Rates

A classic question in public finance concerns analysis of changes in marginal income tax rates. The MVPF of a tax cut that targets a particular income group tells you the welfare gain to those beneficiaries per dollar of net cost to the government. The benefits (numerator) of a tax cut are straightforward: cutting taxes by a dollar increases welfare by a \$1 (that is, \$1 is valued at \$1 by individuals who would be in that income group even without the tax cut). This \$1 valuation requires us

⁶Formally, Hendren and Sprung-Keyser (2020) show that one needs to use the incidence-weighted average social welfare weight when comparing MVPFs across policies.

⁷With apparently little sense of irony, Hendren (2013) notes in the working paper version: “Relative to [the existing] literature, the primary contribution of this paper is a clarification.” In turn, the current article is a revised version of a teaching note (Finkelstein 2019) in which the mathematical derivations of the MVPF is attributed to Hendren (2016), because it is apparently a natural tendency to attribute an idea to the source from which one learned it. Finkelstein learned this tendency from Scott Stern—we therefore wish to cite Scott Stern—appropriately here.

to assume that individuals are making privately optimal decisions so that we can ignore any benefits to marginal recipients who change their behavior in response to the tax cut.

The cost of the tax cut is the sum of the mechanical cost and the fiscal externality. The mechanical cost of the tax cut—that is, the cost per inframarginal recipient, holding behavior constant—is a dollar. The fiscal externality of the tax cut is how the tax cut affects the government budget. Possible behavioral responses may include changes in labor supply and changes in the use of tax sheltering strategies, among others. The key is the elasticity of taxable income (and hence tax revenue) with respect to the tax rate (Feldstein 1999).

This causal object has been the subject of a vast empirical literature in labor economics and public finance. For example, drawing on existing causal estimates of various tax reforms, Hendren and Sprung-Keyser (2020) estimate that for every \$1 of revenue raised from the 1993 tax increase on top earners, the government lost \$0.46 in revenue from behavioral responses that reduced top earners' taxable income. Therefore the net "cost" of the tax increase on the government budget (mechanical cost of plus fiscal externality of $-\$0.46$) is \$0.54, for an MVPF of 1.85 ($= 1/0.54$). The parameter for a tax increase can be used in reverse to think about a tax cut: that is, a dollar of tax cuts on high earnings costs less than its mechanical cost of a dollar because increases in labor supply (or decreases in tax shelters), increase taxable earnings and tax revenue. The MVPF of a tax cut on top earners is greater than 1 because \$1 in tax cuts generates \$1 in benefits but costs less than \$1 due to the negative fiscal externality.

Indeed, if the fiscal externality of a tax cut at the top is less than -1 , then we are on the "wrong side" of the Laffer curve. Cutting tax rates raises revenue, and the tax cut "pays for itself." Hendren and Sprung-Keyser (2020) calculate that the fiscal externality from the Reagan tax cut of 1981 was -1.51 , so that the tax cut "paid for itself"—although they caution that there is a wide degree of statistical uncertainty in the estimates of the behavioral response (for example, they can't statistically reject an MVPF of 1).

By contrast, a tax cut at the bottom of the income distribution—say, in the form of an increase in the Earned Income Tax Credit—has a different fiscal externality. When individuals at the bottom of the income distribution enter the labor market, they impose a negative externality on the government budget by taking EITC benefits (which increase government spending), but a positive benefit by taking less in transfers that would go to those with lower incomes (in the form of welfare, food stamps, and other benefits). On net, the calculations in Hendren and Sprung-Keyser (2020) suggest the reduction in transfer payments slightly outweighs the increased EITC costs, so that a \$1 mechanical increase in the EITC leads to a fiscal externality that reduces net government costs by .08. This implies an MVPF of \$1.12 ($= 1/(1 - .08)$).

It is perhaps not surprising that the MVPF appears to be lower for a tax cut to the poor than to the rich: this outcome is what would be expected in an optimal tax system set by a planner that places greater social welfare weight on the

marginal value of resources for the poor than the rich. The “bang for the buck” is higher for tax cuts at the top than the bottom, but tax cuts at the top may not be desirable given the greater social value of resources at lower incomes. It is cheaper to raise revenue from the poor, but this of course has adverse distributional implications.

Education

The government is a large provider and funder of education, especially primary schooling. How do we calculate the MVPF for an increase in school funding? To illustrate this, consider the work of Jackson, Johnson, and Persico (2016), who study the effect of K–12 school spending on children’s long-run outcomes. They use variation from school finance equalizations to show that increased spending led to an increase in children’s earnings trajectory over their life cycle.

To calculate the MVPF in this context, first consider the costs to the government of the policy. There is the upfront cost from increased school spending. This is offset, however, by any increases in future tax revenue paid by the children as a result of their increased earnings. Hendren and Sprung-Keyser (2020) translate the estimates from Jackson, Johnson, and Persico (2016) into a projection of lifetime tax revenue paid. They find that this increase in tax revenue is actually sufficient to cover the initial spending on education (accounting for real government interest rates of 3 percent—which is of course an assumption), so that the net cost of the policy is negative. This implies an infinite MVPF, regardless of the size of the willingness to pay for the policy; as long as willingness to pay is positive (in other words, the children are personally better off from the additional spending), the policy increases welfare without costing the government any money. As a result, we can skirt the more conceptually and empirically challenging task of estimating the willingness to pay for this increase in education spending; we discussed potential approaches to estimating willingness to pay for in-kind transfers in the extensions section earlier, but are glad not to have to actually implement them here.

More generally, Hendren and Sprung-Keyser (2020) have provided a “library” of estimates of the MVPF for over 100 US expenditure policies, including changes in spending on education, job and vocational training, housing subsidies, food stamps, health insurance, and many more. It would be a useful task to develop a comparable MVPF “library” for public expenditure programs in other countries.

De-Worming

The educational example above assumes the only beneficiaries from the expenditures are the individual students themselves. However, many government programs can have externalities onto others. In the education example, expanded education may increase the earnings of the rest of the population through complementarities in production (leading to a higher MVPF). Conversely, some of the estimated wage gains may come from sorting/signaling and therefore impose negative externalities on others (leading to a lower MVPF).

Here, we provide a specific example of how such externalities are incorporated into the MVPF framework in the context of a health policy implemented in a developing country. This example is due to Baird et al. (2016), who study the impact of school-based de-worming treatments in Kenya. They document that these treatments led to improvements in health and long-run earnings for the children in these schools. In addition, the treatments also provided benefits for students in neighboring primary schools—who did not receive the de-worming treatment—through reductions in transmissions of infection.

Computing the MVPF in this case would therefore involve measuring the willingness to pay for the treatments not only for the children who were directly treated (and their families/communities) but also people in the neighboring areas who also saw improvements in their health (and school attendance rates). We discussed this earlier when we talked about the possibility that policies may have external effects beyond the direct recipients. However, as with the Jackson, Johnson, and Persico (2016) estimates of spending on K–12 education in the United States, Baird et al. (2016) estimate that the net-cost to the government of de-worming is negative: the long-run tax revenue from increased earnings in adulthood is sufficient to cover the government cost of the de-worming efforts. Once again, we are spared having to calculate the willingness to pay for in-kind transfers. Given the estimated effects, de-worming policy has an infinite MVPF and is a win-win for the government and its citizens. Of course, in other settings where net costs are positive, one would have to estimate the affected individuals' willingness to pay for the de-worming using the methods for estimating willingness to pay for in-kind transfers that we discussed earlier.

Import Tariffs

A classic question in international trade concerns the welfare consequences of import tariffs, such as the 2018 tariffs imposed by the United States on goods from China (Fajgelbaum et al. 2020; Cavallo et al. 2019; Amiti, Redding, and Weinstein 2019a). We consider the MVPF of an increase in import tariffs from the perspective of the home country. We therefore ignore any costs or benefits for other countries—this could, of course, be incorporated.

To begin, suppose that an increase in tariffs does not lead to a domestic price change, and there are no retaliatory responses by foreign governments to their tariffs. In this case, a \$1 increase in tariffs leads foreigners to pay \$1 more in taxes and imposes no costs on domestic citizens. With no change in domestic prices, the willingness to pay by residents in the home country will be zero, resulting in an MVPF of zero. From the home country's perspective, the tariff would be an effective way of raising revenue—or, equivalently, an import subsidy would be a poor use of government revenue.

More commonly though, tariffs increase domestic prices. Indeed, Fajgelbaum et al. (2020), Cavallo et al. (2019), and Amiti, Redding, and Weinstein (2019a) all find that the 2018 tariffs were passed-through in full to domestic prices; in other words, domestic prices went up by the amount of the tariff. In terms of the “benefits”

to US consumers (that is, their willingness to pay to avoid a price increase), individuals would be willing to pay \$1 to avoid an increase in prices of \$1. If demand for other goods is not affected by the tariff on imports, the “benefits” of the tariff are \$1. (Actually, it’s $-\$1$, but the denominator will also be negative so they will cancel. Taxes and subsidies on the same good(s) have the same MVPF.)

Turning now to costs, one can think of the tariff as a tax on imported goods, so there is a mechanical cost proportional to the current expenditure on imported goods. But, there is also a potential negative fiscal externality if the tariff reduces consumption of imported goods; the fiscal externality is the impact of this behavioral response to the tariff on tariff revenue. In this case, the net revenue raised by the policy will be less than the mechanical cost. To calculate the fiscal externality, we need a causal estimate of the elasticity of imports with respect to the tariff.

Amiti, Redding, and Weinstein (in this journal, 2019a) estimate that the total government revenue raised by the 2018 tariffs is \$15.6 billion; this includes the sum of a mechanical cost of \$32 billion and the negative fiscal externality of $-\$16.4$ billion. If the tariff were thought of as “small,” the benefits would simply be equal to the mechanical cost of \$32 billion; domestic consumers’ willingness to pay for \$32 billion in revenue is just the increase in revenue, so the MVPF would be simply $\$32/\$15.6 = 2.05$. But \$32 billions is not small. We are now in the world we discussed above: “what if the policy changes are large?” and must try to estimate willingness to pay for non-marginal policy changes.

The approach that Amiti, Redding, and Weinstein (2019a) take is in the spirit of the famous Harberger (1964) triangle: while the first dollar of the tariff raises revenue proportional to \$32 billion, as one raises the tariff further, consumers who choose to consume fewer imported goods are less affected by further increases in the tariff. As a result, the last dollar of the tariff imposes a welfare cost of \$15.6 billion in contrast to the initial \$32 billion. Assuming consumers substitute away from imported goods in a linear fashion, this implies that half of the reduction in tax revenue due to behavioral responses of \$16.4 billion is “valued” by consumers. This implies a willingness to pay to avoid the tariffs of $\$32 \text{ billion} - \$8.2 \text{ billion} = \$23.8 \text{ billion}$. Putting this together implies an MVPF of $23.8/15.6 = 1.5$. Every \$1 raised by the government imposes a \$1.50 (that is, $\$23.8 \text{ billion}/\15.6 billion) negative benefit on US consumers.

A further concern from a domestic perspective is that raising tariffs leads to a change in prices of exported goods through terms-of-trade effects. Fajgelbaum et al. (2019) use a trade model to capture the spillover effects of price changes onto substitutes and complements for each product and conclude that US individuals would have been willing to pay a total of \$41.6 billion to avoid the increase in tariffs. This differs from the Amiti, Redding, and Weinstein (2019a) number both because it incorporates spillover effects of price changes and because of various implementation choices. Likewise, they estimate a different impact of the tariff on net revenue of \$34.3 billion. This implies an MVPF of 1.2 (that is, $\$41.6 \text{ billion}/\34.3 billion), so that every \$1 of government revenue raised imposes a \$1.20 welfare loss on the domestic population because of increased prices and reduced export demand.

Interestingly, these implied MVPFs of 1.2–1.5 for tariffs are in a range similar to that of raising revenue through the income tax.⁸ Of course, remember that we have not considered potential policy responses by other countries in the form of a trade war, which would negatively affect domestic consumers in a way we have not captured (but in principle could).

Government Procurement Policy

A classic question in industrial organization considers the optimal design of government procurement contracts (Laffont and Tirole 1993). Empirical researchers have studied public procurement contracts for highways (Lewis and Bajari 2014), defense (Carril and Duggan 2018), health insurance (Decarlois 2015; Cabral, Geruso, and Mahoney 2018), durable medical equipment (Ji 2019), and other goods.

In this case, the MVPF measures the monetary benefit of a change in procurement contract per dollar increase in public costs. To be concrete, consider an increase in the government payment to private insurers to provide insurance coverage to elderly individuals through the Medicare Advantage program. In the United States, individuals eligible for Medicare—the public health insurance program for elderly and disabled individuals—can choose between the publicly-provided, fee-for-service Traditional Medicare program and obtaining subsidized coverage through their choice of a privately-provided Medicare Advantage insurance plan. About 30 percent of the 44 million Medicare enrollees choose Medicare Advantage. One key design question for the government is how much to subsidize purchases of these private plans. Cabral, Geruso, and Mahoney (2018) have analyzed the impact of these subsidies empirically. We would like to analyze the MVPF of a \$1 increase in the subsidy per enrollee.

What is the beneficiaries' willingness to pay for a \$1 increase in the subsidy? By now, this should be old hat: inframarginal beneficiaries value the \$1 transfer at a \$1.⁹ Marginal beneficiaries are those who switch from Traditional Medicare to Medicare Advantage in response to the increase in subsidy. Cabral, Geruso, and Mahoney (2018) estimate that every \$1 of subsidy increases Medicare Advantage enrollment by about 0.09 percentage points. We employ the logic that the marginal actors were indifferent between not receiving the additional subsidy and

⁸It is important to mention one caveat about this result: Further increases in the tariff rate may not actually increase government revenue. In Amiti, Redding, and Weinstein (2019b), the authors note that increasing the tariff from 10 to 25 percent on \$200 billion of Chinese imports would lower government revenue, implying an infinite MVPF, so that lowering the tariff would raise welfare.

⁹As Cabral, Geruso, and Mahoney (2018) emphasize, there are two potential types of infra-marginal beneficiaries: consumers who were in Medicare Advantage and insurers who were selling Medicare Advantage. The extra \$1 of subsidy from the government to Medicare Advantage may be split between increases in consumer surplus (in the form of lower prices or higher quality) or higher profits to the firms. They estimate the “pass through” rate is about 54 percent to consumers (virtually all of which comes in the form of lower prices) and 46 percent to firms. How does this distributional analysis affect the MVPF analysis? As discussed in the multiple beneficiaries section above, it does not affect the calculation of the MVPF per se, but rather the interpretation of the result.

not switching, or receiving the additional subsidy and switching, so the net welfare change for these switchers is zero, and the size of the enrollment effect does not directly enter the MVPF estimate. Thus, the benefit of the \$1 subsidy per existing enrollee is simply \$1.

What are the costs of the dollar increase in the subsidy? In the absence of any behavioral response, the mechanical cost of the policy per existing enrollee would simply be \$1 as well. Whether the fiscal externality is negative or positive depends on whether Medicare Advantage saves money so that the 0.09 percentage point increase in enrollment leads to an increase or decrease in costs. Existing estimates suggest that the government ends up paying 3–6 percent more for individuals enrolled in Medicare Advantage than it would have if they had enrolled in Traditional Medicare (Medicare Payment Commission 2018; Curto et al. 2014). Even using the 6 percent number would imply that the MVPF of the increase in the Medicare Advantage subsidy is roughly equal to 1 ($= 1/(1 + .0009 \times .06)$).

Conclusion

The MVPF framework offers a powerful approach to empirical welfare analysis of a change in public expenditures or taxes. The approach focuses on the ratio of affected individuals' own willingness to pay for the policy change to the causal effect of the policy on government's net costs.

A key attraction of this approach is that it allows researchers to incorporate causal estimates of policy changes directly into a welfare analysis. In addition, the MVPF provides an important guide for future empirical work on which behavioral responses matter for welfare. Specifically, empirical economists interested in translating the benefits of the “credibility revolution” into progress on applied welfare analysis should focus their efforts on estimating behavioral responses that have fiscal externalities on the government budget, not on behavioral responses whose costs are (approximately) fully internalized by the responding individuals. The approach seems both more robust and easier to interpret than the traditional methods of welfare analysis, which may require estimating effects of hypothetical policies in which those affected are “compensated” for the change through lump-sum transfers.

Of course, the MVPF approach is no panacea. As we emphasized, estimating the willingness to pay for the policy change can be challenging, especially if the policy involves in-kind transfers (such as subsidized education) or effects on individuals not directly targeted by the policy change. Here, we have described how a variety of arrows in the empirical economists' quiver—including structural modeling, calibration exercises, and quasi-experimental or experimental techniques—may usefully be brought to bear. The core value of the MVPF is that it provides clarity on what objects are needed for welfare analysis. In doing so, it can potentially remove the silos across different fields and place welfare analyses on the same playing field.

Only in rare cases will welfare analysis of real-world public policy be clear-cut and straightforward. But the MVPF framework has the flexibility to be applied in a wide range of situations.

■ *We are grateful to Alan Auerbach, Dave Donaldson, Xavier Jaravel, Amy Kim, Henrik Kleven, Enrico Moretti, Matthew Notowidigdo, Ben Olken, Ben Sprung-Keyser, and Sammy Young for helpful comments; to the JEP Editors (and especially Timothy Taylor) for extensive and helpful comments and edits; and to the numerous students and seminar audiences whose (understandable) questions and confusions prompted us to write this essay.*

References

- Allcott, Hunt, and Dmitry Taubinsky. 2015. "Evaluating Behaviorally Motivated Policy: Experimental Evidence from the Lightbulb Market." *American Economic Review* 105 (8): 2501–38.
- Amiti, Mary, Stephen J. Redding, and David E. Weinstein. 2019a. "The Impact of the 2018 Tariffs on Prices and Welfare." *Journal of Economic Perspectives* 33 (4): 187–210.
- Amiti, Mary, Stephen J. Redding, and David E. Weinstein. 2019b. "New China Tariffs Increase Costs to U.S. Households." *Liberty Street Economics*, May 23. <https://libertystreeteconomics.newyorkfed.org/2019/05/new-china-tariffs-increase-costs-to-us-households.html>.
- Angrist, Joshua D., and Jörn-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 J. 1985. "The Theory of Excess Burden and Optimal Taxation." In *Handbook of Public Economics*, Vol. 1, edited by Alan J. Auerbach and Martin Feldstein, 61–127. Amsterdam: Elsevier.
- Auerbach, Alan J., and James R. Hines Jr. 2002. "Taxation and Economic Efficiency." In *Handbook of Public Economics*, Vol. 3, edited by Alan J. Auerbach and Martin Feldstein, 1347–1421. Amsterdam: Elsevier.
- Baird, Sarah, Joan Hamory Hicks, Michael Kremer, and Edward Miguel. 2016. "Worms at work: Long-Run Impacts of a Child Health Investment." *The Quarterly Journal of Economics* 131 (4): 1637–80.
- Bound, John, Julie Berry Cullen, Austin Nichols, and Lucie Schmidt. 2004. "The Welfare Implications of Increasing Disability Insurance Benefit Generosity." *Journal of Public Economics* 88 (12): 2487–2514.
- Cabral, Marika, Michael Geruso, and Neale Mahoney. 2018. "Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage." *American Economic Review* 108 (8): 2048–87.
- Carril, Rodrigo, and Mark Duggan. 2018. "The Impact of Industry Consolidation on Government Procurement: Evidence from Department of Defense Contracting." NBER Working Paper 25160.
- Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang. 2019. "Tariff Passthrough at the Border and at the Store: Evidence from US Trade Policy." NBER Working Paper 26396.
- Chetty, Raj, and Emmanuel Saez. 2010. "Optimal Taxation and Social Insurance with Endogenous Private Insurance." *American Economic Journal: Economic Policy* 2 (2): 85–114.
- Currie, Janet, and Firouz Gahvari. 2008. "Transfers in Cash and In-Kind: Theory Meets the Data." *Journal of Economic Literature* 46 (2): 333–83.
- Currie, Janet, and Jonathan Gruber. 1996. "Saving Babies: The Efficacy and Cost of Recent Changes in the Medicaid Eligibility of Pregnant Women." *Journal of Political Economy* 104 (6): 1263–96.
- Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya. 2014. "Can Health Insurance Competition Work? Evidence from Medicare Advantage." NBER Working Paper 20818.

- Decarolis, Francesco.** 2015. "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?" *American Economic Review* 105 (4): 1547–80.
- Fajgelbaum, Pablo D., Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal.** 2020. "The Return to Protectionism." *Quarterly Journal of Economics* 135 (1): 1–55.
- Feldstein, Martin.** 1999. "Tax Avoidance and the Deadweight Loss of the Income Tax." *Review of Economics and Statistics* 81 (4): 674–80.
- Finkelstein, Amy.** 2019. "Welfare Analysis Meets Causal Inference: A Suggested Interpretation of Hendren." <https://economics.mit.edu/files/16272>. (accessed August 1, 2020).
- 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 Economics* 92 (7): 1644–68.
- Finkelstein, Amy, Nathaniel Hendren, and Erzo F.P. Luttmer.** 2019. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy* 127 (6): 2836–74.
- Fischer, Torben, Markus Frölich, and Andreas Landmann.** 2018. "Adverse Selection in Low-Income Health Insurance Markets: Evidence from a RCT in Pakistan." IZA Discussion Paper 11751.
- Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigdo.** 2018. "Hospitals as Insurers of Last Resort." *American Economic Journal: Applied Economics* 10 (1): 1–39.
- Goodman-Bacon, Andrew.** Forthcoming. "The Long-Run Effects of Childhood Insurance Coverage: Medicaid Implementation, Adult Health and Labor Market Outcomes." *American Economic Review*.
- Goolsbee, Austan.** 1999. "Evidence on the High-Income Laffer Curve from Six Decades of Tax Reform." *Brookings Papers on Economic Activity* 1999 (2): 1–64.
- Greenstone, Michael.** 2017. "The Continuing Impact of Sherwin Rosen's 'Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition.'" *Journal of Political Economy* 125 (6): 1891–1902.
- Gruber, Jonathan.** 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review* 87 (1): 192–205.
- Harberger, Arnold C.** 1964. "The Measurement of Waste." *American Economic Review* 54 (3): 58–76.
- Hausman, Jerry A.** 1981. "Exact Consumer's Surplus and Deadweight Loss." *American Economic Review* 71 (4): 662–76.
- Heckman, James, J., Seong Hyeok Moon, Rodrigo Pinto, Peter A. Savelyev, and Adam Yavitz.** 2010. "The Rate of Return to the HighScope Perry Preschool Program." *Journal of Public Economics* 94 (1–2): 114–28.
- Hendren, Nathaniel.** 2013. "The Policy Elasticity." NBER Working Paper 19177.
- Hendren, Nathaniel.** 2016. "The Policy Elasticity." *Tax Policy and the Economy* 30 (1): 51–89.
- Hendren, Nathaniel.** 2020. "Measuring Economic Efficiency Using Inverse Optimum Weights." *Journal of Public Economics* 187.
- Hendren, Nathan, and Ben Sprung-Keyser.** 2020. "A Unified Welfare Analysis of Government Policies." *Quarterly Journal of Economics* 135 (3): 1209–1318.
- Hicks, J.R.** 1940. "The valuation of social income." *Economica* 7 (26): 105–124.
- Jackson, C. Kirabo, Rucker C. Johnson, and Claudia Persico.** 2016. "The Effects of School Spending on Educational and Economic Outcomes: Evidence from School Finance Reforms." *Quarterly Journal of Economics* 131 (1): 157–218.
- Jäger, Simon, Benjamin Schoefer, and Josef Zweimüller.** 2019. "Marginal Jobs and Job Surplus: A Test of the Efficiency of Separations." NBER Working Paper 25492.
- Ji, Yunan.** 2019. "The Impact of Competitive Bidding in Health Care: The Case of Medicare Durable Medical Equipment." <https://scholar.harvard.edu/files/yunan/files/dme.pdf>. (accessed July 27, 2020).
- Kleven, Henrik Jacobsen.** Forthcoming. "Sufficient Statistics Revisited." *Annual Review of Economics*.
- Kleven, Henrik Jacobsen, and Claus Thustrup Kreiner.** 2005. "Labor Supply Behavior and the Design of Tax and Transfer Policy." *Danish Journal of Economics* 143: 321–58.
- Kleven, Henrik Jacobsen, and Claus Thustrup Kreiner.** 2006. "The Marginal Cost of Public Funds: Hours of Work versus Labor Force Participation." *Journal of Public Economics* 90 (10–11): 1955–73.
- Kline, Patrick, and Christopher R. Walters.** 2016. "Evaluating Public Programs with Close Substitutes: The Case of Head Start." *Quarterly Journal of Economics* 131 (4): 1795–1848.
- Krueger, Alan B., and Ilyana Kuziemko.** 2013. "The Demand for Health Insurance among Uninsured Americans: Results of a Survey Experiment and Implications for Policy." *Journal of Health Economics* 32 (5): 780–93.

- Laffont, Jean-Jacques, and Jean Tirole.** 1993. *A Theory of Incentives in Procurement and Regulation*. Cambridge: MIT Press.
- Lewis, Gregory, and Patrick Bajari.** 2014. "Moral Hazard, Incentive Contracts, and Risk: Evidence from Procurement." *The Review of Economic Studies* (81) 3: 1201–28.
- Mayshar, Joram.** 1990. "On Measures of Excess Burden and Their Applications." *Journal of Public Economics* 43 (3): 263–89.
- Medicare Payment Commission (MedPAC).** 2018. *Report to Congress: Medicare Payment Policy*. Washington D.C.: MedPAC
- Nekoei, Arash, and Andrea Weber.** 2017. "Does Extending Unemployment Benefits Improve Job Quality?" *American Economic Review* 107 (2): 527–61.
- Olken, Benjamin A.** 2007. "Monitoring Corruption: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 115 (2): 200–49.
- Plimpton.** 1800 BC. https://en.wikipedia.org/wiki/Pythagorean_theorem#History.
- Poterba, James.** 1996. "Government Intervention in the Markets for Education and Health care: How and Why?" In *Individual and Social Responsibility: Child Care, Education, Medical Care, and Long-Term Care in America*, edited by Victor R. Fuchs, 277–308. Chicago: University of Chicago Press.
- Pythagoras** 500 BC. https://en.wikipedia.org/wiki/Pythagorean_theorem#History.
- Ramsey, Frank.** 1927. "A Contribution to the Theory of Taxation." *The Economic Journal* 37 (145): 47–61.
- Rosen, Sherwin.** 1974. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82 (1): 34–55.
- Saez, Emmanuel.** 2004. "Reported Incomes and Marginal Tax Rates, 1960–2000: Evidence and Policy Implications." *Tax Policy and the Economy* 18: 117–72.
- Saez, Emmanuel, and Stefanie Stantcheva.** 2016. "Generalized Social Welfare Weights for Optimal Tax Theory." *American Economic Review* 106 (1): 24–45.
- Schmieder, Johannes F., and Till von Wachter.** 2016. "The Effects of Unemployment Insurance Benefits: New Evidence and Interpretation." *Annual Review of Economics* 8: 547–81.
- Slemrod, Joel, and Shlomo Yitzhaki.** 1996. "The Costs of Taxation and the Marginal Efficiency Cost of Funds." *International Monetary Fund Staff Papers* 43 (1): 172–98.
- Slemrod, Joel, and Shlomo 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.