

Evaluation of Public Policies and Social Programs

Session 5: Modern Welfare Frameworks — MVPF & Sufficient Statistics

Andrés Barrios Fernández

Monday, August 17, 2026

MA in Economics, Universidad de los Andes, Chile — Second Semester 2026

Today

1. From Harberger's triangle to sufficient statistics
2. The MVPF framework
3. Guided calculation: MVPF for an education subsidy
4. Deep discussion and what we learned
5. Paper presentation (closes today's session)

This closes out Module A — the core toolkit.

Learning Objectives

By the end of today's session, you will be able to:

- Explain the sufficient-statistics logic: why local welfare effects can often be computed without a full structural model
- State the MVPF formula and its relationship to the traditional benefit-cost ratio
- Compute an MVPF from primitives (mechanical cost, WTP, behavioral response, tax rate, discount rate)
- Identify when the sufficient-statistics/MVPF approach is, and is not, well suited to evaluating a large, discrete project like the Astoreca school

From Harberger's Triangle to Sufficient Statistics

Recall: what do we mean by “welfare”?

- Session 3 (Hicks, 1943): an individual's welfare change from a policy is a *money-metric* object — compensating or equivalent variation, the dollar amount that maps a utility change onto a comparable scale
- Ordinary consumer surplus, WTP for a change: practical, estimable stand-ins for that same object (Willig, 1976, tells us when the approximation is safe)
- **Social welfare change:** sum these individual money-metric terms across everyone affected — Harberger's postulate 2, applied literally

Every “welfare” term today — Harberger's triangle, Chetty's sufficient statistics, MVPF's WTP numerator — is this same object, just computed a different way.

The problem: full welfare analysis is hard

Can we get a rigorous welfare number without estimating (or even specifying) the full structural model of behavior?

The hard way: a fully structural model

To compute welfare *exactly*, in principle you could:

1. Write down a complete utility function $U(c, l; \theta)$ with an unknown parameter vector θ
2. Estimate θ econometrically from data on behavior
3. Solve the individual's optimization problem under the status quo *and* under the counterfactual policy
4. Plug both solutions back into U to get the exact dollar-valued welfare change

This is what “a structural model” means: fully specified preferences (or technology), used to simulate any counterfactual you like.

Cost: strong, often untestable functional-form assumptions; heavy estimation burden; every downstream result is only as credible as the whole specified model.

Harberger (1964): the deadweight-loss triangle


Harberger, A.C. (1964), “The Measurement of Waste,” AER

For a small tax distortion t on a good with compensated elasticity ε , base price p and quantity q :

$$\text{DWL} \approx \frac{1}{2} \varepsilon \left(\frac{t}{p} \right)^2 p q$$

Notice what this formula does *not* require: the shape of the utility function, or any structural model of why people respond. **One elasticity** is enough.

The Harberger triangle, visually



$DWL \approx \frac{1}{2} \varepsilon (t/p)^2 pq$ — the shaded triangle, base Δq and height t . Following Harberger (1964).

From triangle area to Harberger's formula

The triangle is just $DWL = \frac{1}{2} \times \text{base} \times \text{height}$ — nothing more exotic than that.

- **Height** = t : the wedge between demand and supply price, at the new (post-tax) quantity
- **Base** = $\Delta q = \varepsilon(t/p)q$: the quantity response — literally the definition of the elasticity ε

$$DWL = \frac{1}{2} \times \underbrace{\varepsilon \left(\frac{t}{p} \right) q}_{\text{base}} \times \underbrace{t}_{\text{height}} = \frac{1}{2} \varepsilon \left(\frac{t}{p} \right)^2 pq$$

The formula *is* the triangle: one elasticity tells you how wide the base is, and the tax itself is the height.

Chetty, R. (2009), “Sufficient Statistics for Welfare Analysis,” AER

The welfare effect of a small policy change depends only on a small number of empirically estimable “sufficient statistics” (elasticities, pass-through rates, participation responses) — not on fully specifying preferences or technology.

This is Harberger’s triangle, generalized far beyond a single distortionary tax: to transfers, subsidies, insurance programs, and (as we’ll now see) essentially any government policy. Why does this work? The envelope theorem.


Start big: how would a large policy change reshape decisions?

Imagine each of these as a *large* change, not a marginal one:

- A payroll tax so large that it wipes out most of the return to working an extra hour — people would visibly work less
- A tuition subsidy so large that school becomes free — families who would never otherwise have paid would enroll
- A gas tax so large that driving triples in cost — people would relocate, switch cars, change jobs

Large changes obviously move behavior a lot, and the resulting welfare loss is correspondingly large — nobody would dispute that.

Now shrink the change: what happens near the optimum?



The curve is *flat* right at the peak — that's what "optimal" means. A small horizontal nudge (green) costs almost nothing vertically; only a large one (orange) does.

Why behavioral responses (mostly) wash out for the individual

- At someone's privately optimal choice, they are — by definition — indifferent to small deviations from it: that's the zero-slope, flat-top picture you just saw, and it's exactly what a first-order condition means
- Starting from an undistorted baseline, a small policy-induced nudge in their behavior costs *them* only a **second-order** (quadratic-in-the-wedge) welfare loss — the envelope theorem
- You already saw this: $DWL \approx \frac{1}{2}\varepsilon(t/p)^2 pq$ — that squared term *is* the second-order envelope-theorem loss, sitting in the Harberger formula from three slides ago

But first-order for anyone who doesn't get to reoptimize

- The *same* behavioral response hits someone else's budget a full order larger: government tax revenue, $R = t \cdot q(t)$, moves as $dR/dt = q + t \cdot dq/dt$
- That second term doesn't vanish the way the DWL triangle does — it shrinks only as fast as t itself (first order), not as t^2 (second order), because the government just collects tax on whatever quantity results; it never reoptimizes around your choice the way you did
- This is Chetty's sufficient statistic in one sentence: not the whole behavioral model, just the one number (an elasticity) that captures this first-order impact on a constraint nobody optimizes away

Fine print: the DWL triangle strictly uses the *compensated* elasticity, the fiscal externality the *actual* behavioral response — they coincide when income effects are small, which we'll assume throughout.

This is the direct ancestor of MVPF

Harberger's triangle $\longrightarrow$ Chetty's sufficient statistics $\longrightarrow$ Hendren &
Sprung-Keyser's MVPF


Same underlying idea, each generation broader: use a handful of credibly estimated behavioral responses to answer a welfare question, without needing the whole model.

The MVPF Framework

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


$$\text{Net Cost} = \text{Mechanical Cost} - \text{Fiscal Externality}$$

How Hendren & Sprung-Keyser actually build these numbers



Net cost to government of admitting a marginal student to Florida International University: the mechanical cost (red) minus every downstream tax/spending effect of their higher earnings (blue), ending in a net cost far below the sticker price. Hendren & Sprung-Keyser (2020), Figure I, Panel A, based on Zimmerman (2014).

Same logic, the willingness-to-pay side



WTP is built the same way: what the student gives up (tuition, red) against what they gain (higher lifetime after-tax earnings, blue), summing to the baseline WTP (green). **This is exactly today's guided calculation** — mechanical cost, behavioral/fiscal response, and WTP, all built from the same three ingredients. Hendren & Sprung-Keyser (2020), Figure I, Panel C.

Why $WTP \approx \text{mechanical cost}$, in the frictionless benchmark

- Recall the hilltop picture: at someone's own optimal choice, a small nudge costs them almost nothing — the envelope theorem, again
- Now flip it around: if the *government* hands someone a dollar spent exactly the way they would have spent it themselves, that dollar is worth — to them — almost exactly one dollar. There is no “translation loss” between a dollar of mechanical spending and a dollar of value received
- That is the entire content of “ $WTP \approx \text{mechanical cost}$ ”: not an assumption about generosity or altruism, just the same flat-top logic applied to a transfer instead of a tax

Which elasticity belongs in the denominator?

Whatever governs the **fiscal externality** — not a generic demand elasticity.

- Specifically: the response of taxable earnings, program take-up, or transfer receipt to the policy — the one behavioral margin that feeds back onto the government's budget
- This is a direct callback to Chetty's lesson from Section 1: it's the elasticity of whatever constraint the individual *doesn't* optimize away that matters, not the individual's own private margin

Why the benchmark can fail

- **Liquidity constraints:** a cash-strapped family can't act on their own long-run optimum today, so what they'd pay for a dollar *now* exceeds a dollar of future value — WTP and mechanical cost decouple
- **Internalities (Session 7):** if people don't optimize even in their own interest (self-control problems, present bias), their revealed WTP is not a reliable stand-in for their true welfare
- **Pre-existing distortions elsewhere:** a second wedge already in the economy (a tax, a market failure) breaks the clean envelope-theorem cancellation — “second best” territory

In every case, the fix is the same: estimate WTP directly instead of assuming it equals mechanical cost.

$$\text{MVPF} = \frac{\text{Willingness to Pay}}{\text{Net Cost}}$$

$$\text{Net Cost} = \text{Mechanical Cost} - \text{Fiscal Externality}$$

The fiscal externality term is exactly Harberger's and Chetty's behavioral-response logic, applied to the government budget.

- “Bang for the buck” from the beneficiary’s perspective
- $MVPF = \infty$ for a fully self-financing policy
- The fiscal externality in the denominator is exactly the behavioral-response term Harberger and Chetty taught us to isolate

Comparing frameworks: BCR vs. MVPF, precisely

$$\text{BCR} = \frac{\text{WTP} + \text{FE}}{(1 + \varphi) C}, \quad \text{MVPF} = \frac{\text{WTP}}{C - \text{FE}}$$

Same three ingredients (WTP, mechanical cost C , fiscal externality FE) — combined two very different ways:

- **BCR** puts the fiscal externality in the *numerator*, as an extra benefit, and inflates cost by a *fixed, policy-independent* φ (the assumed deadweight loss of raising \$1 through distortionary taxation, typically 0.3–0.5)
- **MVPF** nets the fiscal externality directly against cost in the *denominator*, using *that policy's own*, actual, estimated fiscal externality — no external assumption about how the government would raise the money

Same inputs, different plumbing — and the plumbing matters exactly when the fiscal externality is large.

Where the two formulas pull apart


A policy that exactly pays for itself: $FE = C$

$$BCR = \frac{WTP + C}{(1 + \varphi)C} \approx \frac{2C}{1.5C} \approx 1.33 \quad (\varphi = 0.5, WTP \approx C)$$

$$MVPF = \frac{WTP}{C - C} = \infty$$

BCR reports an unremarkable 1.33; MVPF correctly flags a **Pareto improvement** — the government gets its money back and beneficiaries still gain.

This split shows up across 133 real policies



BCR and MVPF track each other closely up to about 1 — then visibly pull apart as MVPF grows large or infinite (FIU GPA, top-tax and Medicaid-expansion policies at far right, where BCR stays bounded but MVPF diverges). Hendren & Sprung-Keyser (2020), Figure IX, Panel B.

A case where the ranking flips: EITC vs. the tax that funds it


Hendren & Sprung-Keyser (2020), Section VI.B

The 1993 EITC expansion: $BCR = 0.74$ (fails a strict cost-benefit test, at $\varphi = 50\%$) but $MVPF = 1.12$.

- A strict BCR test would reject the EITC expansion outright: the assumed 50% deadweight cost of raising the funds is “too large to justify” it
- But that verdict bakes in *one specific, uniform* financing assumption. MVPF instead asks: compare the EITC’s MVPF (1.12) to the MVPF of the 1993 top-tax-rate increase that roughly funded it (1.85, in their estimates)
- If society would rather give \$1.12 to a low-income worker than \$1.85 to a top earner, the trade is welfare-improving — *despite* the low BCR

This is Okun’s leaky bucket (Session 8) in miniature: BCR hides the redistributive trade-off inside a fixed φ ; MVPF forces you to compare it explicitly.

Calibrating your intuition: what's a big MVPF?



133 US policies, MVPF against the average age of beneficiaries. Child education/health cluster at the top (some infinite); adult cash transfers and UI cluster near 1; some tax cuts fall below

1. Hendren & Sprung-Keyser (2020), Figure III.

Today's guided calculation gives an MVPF of 50.

That is unusually high — looking at the plot again, only a handful of child-education and child-health policies reach that far up. It is a useful benchmark for how sensitive MVPF is to hitting the marginal-recipient margin hard: most of the government's outlay comes back as tax revenue, precisely because the earnings response is large relative to the mechanical cost.

What MVPF doesn't tell you

Two policies: $MVPF = \infty$ and $MVPF = 2$. Is the first one obviously better?

- Not if the first program reaches 10 people and the second reaches 100,000 — MVPF is silent on *scale* (total dollars, people affected)
- It's also silent on whether the recipients are the people you most want to reach — that's Session 8's distributional weights, not MVPF itself
- MVPF answers “is this an efficient use of the marginal government dollar,” not “should we fund this instead of that,” and not “how much should we spend on it”

Guided Calculation: MVPF for an Education Subsidy

Government subsidizes \$1000 of tuition

- Mechanical cost: \$1000
- $WTP = \$1000$
- Future earnings: +\$200/yr for 30 years
- Tax rate: 25%
- Discount rate: 3%
- 10% of recipients are marginal (would not have attended otherwise)

Before you compute: which term dominates?

Mechanical cost is \$1000. Will Net Cost end up close to \$1000, close to \$0, or negative?

- The tax rate is 25%, and the earnings gain lasts 30 years — a long time for those taxes to accumulate
- Think back to the hilltop picture: none of that earnings response is a *loss* to net anyone out against — it's pure fiscal externality, working in the government's favor
- *Discuss in pairs for a minute before the next slide.*

Where does “19.6” come from?

The present value of \$1/year for 30 years at $r = 3\%$ is an annuity factor:

$$\text{Annuity factor} = \frac{1 - (1 + r)^{-n}}{r} = \frac{1 - (1.03)^{-30}}{0.03} \approx 19.6$$

Always know where your PV factors come from — you will build many of these by hand this semester.

Solution

$$\text{Fiscal Externality} = 0.25 \times 200 \times 19.6 = 980$$

$$\text{Net Cost} = 1000 - 980 = 20$$

$$\text{MVPF} = 1000/20 = 50$$

Every \$1 of net government spending generates \$50 of value to recipients.

The fiscal externality (\$980) very nearly offsets the entire mechanical cost — Net Cost collapses to \$20, exactly the mechanism from the previous slide.

The same policy, run through the BCR formula

Using the same three ingredients, at $\varphi = 0.5$:

$$\text{BCR} = \frac{\text{WTP} + \text{FE}}{(1 + \varphi) C} = \frac{1000 + 980}{1.5 \times 1000} = \frac{1980}{1500} \approx 1.32$$

$\text{BCR} \approx 1.3$: a solid, respectable policy. $\text{MVPF} = 50$: an exceptional one.

Same numbers, same policy — the gap between 1.3 and 50 is entirely the plumbing difference from two slides ago, not a disagreement about the underlying facts.

1. Recompute both the MVPF *and* the BCR if 20% of recipients are marginal instead of 10% (does the marginal share even enter either formula — why or why not?)
2. Recompute both if the tax rate were 10% instead of 25%. Which ratio moves more, and why?
3. In your own words: what has to be true about a policy for BCR and MVPF to tell substantially different stories?

Deep Discussion

Which framework for the Astoreca project?

Is the Astoreca school evaluation more naturally...

An MVPF problem?

A government/donor spending choice:
does it have a clean mechanical cost, and
a fiscal/budget response to net against it?

A traditional social CBA problem?

Capital costs, community and sibling
externalities, no single “funder’s budget”
that all costs and benefits net against

Small groups: argue both sides — then be ready to name *one concrete number* your side’s framework would need that the other side’s wouldn’t.

A sharper caveat: is Astoreca a marginal change?

Sufficient statistics are, by construction, a *first-order (local)* approximation.

- Harberger's triangle and Chetty's envelope-theorem logic are derived for *small* policy changes
- Opening one new school for 600 students is a large, discrete, lumpy project — not a marginal tweak to an existing program
- Does that make MVPF/sufficient-statistics reasoning *wrong* for this project, or just something to apply with more caution? What would you want to check?

Instructor's guidance: use both, for different jobs

- **Lead with social CBA** for the headline recommendation: Astoreca is a discrete capital project, not a marginal tweak to an existing budget line, and CBA naturally accommodates the community/sibling/parental-labor-supply externalities that never touch any single funder's budget — MVPF's denominator is silent on all of them
- **Nest MVPF inside it** wherever the project contains an actual marginal government spending choice: e.g., a means-tested tuition subsidy for a subset of families, or an incremental teacher hire. Those sub-decisions are exactly what MVPF was built for, and its “bang for the buck” framing is more informative there than a project-wide BCR would be
- The two frameworks are not rivals for the same job — **the envelope decision is CBA-shaped; the financing decisions inside it are MVPF-shaped**

What We Learned Today

- Harberger's DWL triangle, Chetty's sufficient statistics, and the MVPF are one intellectual lineage: use credible, estimable behavioral responses instead of a full structural model
- MVPF and the traditional BCR can disagree because they answer different questions and make different implicit assumptions about weighting
- Sufficient-statistics reasoning is a *local* approximation — a caveat that matters directly for evaluating a large, discrete project like a new school
- Module A is complete: you now have the private appraisal baseline, the shadow-pricing/social-cost toolkit, and the modern MVPF framework

- Problem Set 1 (self-study checkpoint: by Session 7)
- Teams should have a one-paragraph project framing ready informally by Session 6

Next: Module B — Valuing Non-Market Benefits

Paper Presentation (75–90 min)

Pedro Vergara & Felix Gimeno

Greenstone & Gallagher (2008), “Does Hazardous Waste Matter? Evidence from the Housing Market and the Superfund Program,” *QJE*