

Evaluation of Public Policies and Social Programs

Session 8: Distributional Weights & Social Welfare Functions

Andrés Barrios Fernández

Monday, August 31, 2026

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

Today

1. Social welfare functions and distributional weights
2. Exercise: a new school vs. a broad-based subsidy
3. When to use distributional weights, and project application
4. What we learned

Today we relax an assumption every session so far has quietly made. No paper presentation today.

Learning Objectives

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

- State Harberger's postulate 2 precisely, and explain what relaxing it requires
- Derive isoelastic distributional weights from a social welfare function, and compute them for $\eta = 0, 1, 2$
- Apply weights to a policy comparison and show that they can flip the ranking of two options
- Take a defensible position on when a social evaluation should use distributional weights at all

Social Welfare Functions


Session 1: “sum the gains and losses to all parties affected, with no regard to who gains and who loses.”

Every dollar of benefit has counted the same, regardless of whose pocket it's in. Today we ask: **should it?**

Atkinson, A. (1970), “On the Measurement of Inequality,” JET

If utility $u(c)$ is concave in consumption c , an extra dollar is worth more, in utility terms, to a poor household than to a rich one. Atkinson formalized how to aggregate this into a single inequality-sensitive social welfare measure.

Diminishing marginal utility, visually



The slope of $u(c)$ — the marginal utility of income — is exactly what a distributional weight is built from. Standard concave-utility diagram, following Atkinson (1970).

Three social welfare functions

- **Utilitarian** ($\eta = 0$): $W = \sum_i c_i$ — a dollar counts equally regardless of income
- **Isoelastic / log-utility** ($\eta = 1$): $W = \sum_i \ln c_i$
- **Rawlsian** ($\eta \rightarrow \infty$): $W = \min_i u(c_i)$ — only the worst-off household matters

General isoelastic form: $W = \sum_i c_i^{1-\eta} / (1 - \eta)$

The distributional weight

$$\lambda_i = \left(\frac{\bar{y}}{y_i} \right)^\eta$$

- η : your inequality-aversion parameter — how much more a dollar to the poor counts
- $\eta = 0$: $\lambda_i = 1$ for everyone (Harberger's postulate 2, unweighted)
- Higher η : weights tilt sharply toward low-income beneficiaries
- MVPF's implicit take: a government's revealed program choices imply a set of weights, whether or not anyone wrote them down (Hendren & Sprung-Keyser, 2020, Section II.C)

Exercise: A New School vs. a Broad-Based Subsidy

Two ways to spend \$8M/year

- A** Broad, income-blind school-choice voucher: 8,000 families, mean income \$55,000, benefit \$300/family
- B** Astoreca-style school targeted at a low-income neighborhood: 1,200 families, mean income \$18,000, benefit \$1,200/family

Task (pairs, 20 minutes)

For $\eta = 0, 1, 2$:

1. Compute the distributional weight for each group
2. Compute weighted and unweighted total benefits for each option
3. Does distributional weighting change which option looks better?

Assume a regional mean income of $\bar{y} = \$40,000$ for this calculation.


Solution: unweighted benefits

Option A: $8,000 \times \$300 = \$2,400,000$

Option B: $1,200 \times \$1,200 = \$1,440,000$

Unweighted ($\eta = 0$): Option A wins.

Solution: weighted benefits, $\eta = 0, 1, 2$



As η rises, Option B overtakes Option A — targeting the same \$8M at a poorer population turns out to generate more social welfare once you stop treating every dollar as identical.

Deep discussion: which η is right?

There is no consensus value of η in economics — estimates in the optimal-taxation literature range roughly from 1 to 3.

- Your choice of η is a value judgment, not a data-generating fact — own it explicitly in your report
- A donor with an explicit mission to serve low-income families is implicitly choosing a high η — does the Astoreca project's funder have such a mission?

When to Use Distributional Weights

The case for and against

For

Declining marginal utility of income is empirically well-grounded; ignoring it is itself a value judgment ($\eta = 0$)

Against

Kaplow (2004): redistribution is better handled through the tax system — don't distort individual project choices to do equity's job twice

- US OMB practice: unweighted CBA is standard; distributional analysis reported *separately*
- UK Green Book: some contexts use explicit weights (transport appraisal)

Saez, E. & Stantcheva, S. (2016), “Generalized Social Marginal Welfare Weights for Optimal Tax Theory,” AER

Rather than committing to one social welfare function, treat the welfare weight on each group as a parameter you can vary and report sensitivity over — the same spirit as this course’s overall approach to uncertainty.

How should your team present distributional analysis for Astoreca?

Minimum bar: report the unweighted headline number *and* a distributional table (who gains, by income group) — let the reader apply their own η if they disagree with yours. With no presentation today, use the rest of class for team project work — instructor circulates for feedback.

What We Learned Today

- Distributional weights formalize exactly the assumption Harberger's postulate 2 asks us to make — and relaxing it is a value judgment, not a technical correction
- Isoelastic weights $\lambda_i = (\bar{y}/y_i)^\eta$ let you show how a policy ranking changes continuously as inequality aversion rises
- Targeting the same budget at a poorer population can flip a policy comparison, even when the unweighted numbers say otherwise
- Best practice: report the unweighted number and a distributional table, and be explicit about which η (if any) you used

- Problem Set 2 (self-study checkpoint: by Session 12)