

Robustly Welfare Improving Tax Reforms

Paolo G. Piacquadio Kristen Vamsaeter

Second International Conference on Inequalities and Opportunities

Bari, 25th June 2026

- A central question in public finance is how to improve an existing income-tax schedule
- Answering this question requires taking a (normative) stand on social welfare
- Two main approaches
 - Choose a social welfare function and find welfare-improving reforms
 - Identify reforms that improve welfare for all individuals
- Each approach has shortcomings
 - Little consensus on which social welfare function to choose
 - Few reforms may survive the requirement that all individuals are better off
- Our question: is there a middle ground between these two approaches?

Our approach

- The appeal of the second approach is broad agreement across social welfare functions
- But the class of reforms is often too wide to be useful in practice
 - Few (if any) reforms qualify
- Our idea:
 - Narrow the class enough to find many more reforms
 - Preserve a broad support
- We aim at identifying *robustly welfare-improving reforms*

Our criterion

- We restrict attention to inequality-averse social welfare functions
- That is, those that weakly favour high-to-low-income transfers
- Intuition:
 - Sits between requiring everyone be made better off and committing to one welfare function
 - Empirically, most welfare functions used in public finance fit within this class
- Does it help?
 - The class is, of course, smaller than the unrestricted set of social welfare functions
 - But the payoff is that many more reforms turn out to count as improvements

- We develop an approach to identify and construct robustly welfare-improving reforms
- Central result: *necessary* and *sufficient* conditions for such a reform to exist
- The results are constructive
 - They identify which income brackets a reform should target
 - And, for each, whether the marginal tax rate should rise or fall
- Easy to apply with sufficient-statistics formulas for revenue effects of tax reforms
- Application: labor income taxes across U.S. states
 - Such reforms exist almost universally for standard behavioral responses

Outline

- 1 Framework for income-tax reform
- 2 Reforms that benefit everyone
- 3 Reforms robust across (most) welfare functions
- 4 Identifying and constructing robust reforms
- 5 Application to U.S. income tax schedules
- 6 Conclusion

Outline

- 1 Framework for income-tax reform
- 2 Reforms that benefit everyone
- 3 Reforms robust across (most) welfare functions
- 4 Identifying and constructing robust reforms
- 5 Application to U.S. income tax schedules
- 6 Conclusion

- A unit mass of workers i with preferences u over consumption c and earnings y
- A tax schedule T maps earnings to taxes, so $c = y - T(y)$
- Each worker chooses earnings

$$\arg \max_y u_i(y - T(y), y)$$

- For every schedule, the worker's problem has a unique interior solution $y_i(T)$
- This induces an earnings distribution F_Y

Local reforms and their effects

- A local reform perturbs the status quo T_0 in direction h with magnitude τ

$$T^{\tau,h}(y) = T_0(y) + \tau h(y)$$



- Revenue effect is the change in tax revenue R along h

$$D_R[h] = \left. \frac{d}{d\tau} R(T_0 + \tau h) \right|_{\tau=0}$$

- Combines a mechanical effect and a behavioural fiscal externality
- Reform is budget-neutral when $D_R[h] = 0$

Elementary tax reforms

- For identifying and constructing reforms, it helps to work with a simpler family
- We focus on *elementary reforms* b_z , built as 2 phase-in reforms



Elementary tax reforms effects

- Summarize each reform b_z by two effects
- Welfare effect: workers' total willingness to pay to avoid the increase b_z

$$B(z) = \int_0^{\infty} b_z(y) F_Y(dy)$$

- Revenue effect: revenue change from the same increase

$$\rho(z) = D_R[b_z]$$

- Their ratio: revenue $\rho(z)$ per unit of willingness to pay $B(z)$ for the reform

$$M(z) = \frac{\rho(z)}{B(z)}$$

Outline

- 1 Framework for income-tax reform
- 2 Reforms that benefit everyone**
- 3 Reforms robust across (most) welfare functions
- 4 Identifying and constructing robust reforms
- 5 Application to U.S. income tax schedules
- 6 Conclusion

Characterization

A local reform that makes everybody better off exists if and only if there is some income x with


$$M(x) < 0.$$

- Intuition: $M(x)$ turns negative only when a tax cut at x pays for itself
- Such a cut lowers some burdens *and* frees revenue for further cuts
- Typically happens when one the wrong side of the Laffer curve
- Direct test: check whether M ever goes below zero

Applying the test

- It can be shown that one might express M using sufficient statistics
- Necessary inputs are:
 - Earnings distribution f , F and tax rates T'
 - Behavioral parameters: compensated elasticity ε^c and income effects η
- Application: labor-income taxes across U.S. states
 - Earnings distribution and tax rates from IRS Statistics of Income tables
 - Behavioral parameters over a plausible range from the literature

Applying the test: Indiana



- Example application of the test to the income-tax schedule in Indiana
- Illustration with our 'favourite' specification of behavioural parameters:
 - Compensated elasticity $\varepsilon^c = 0.3$
 - Income effect $\eta = -0.5$
- $M(y) \geq 0$ for all y so no reform making everyone better off exists

Result: Few reforms make everyone better off

- Apply the test state by state, across behavioral parameters
- Such reforms are rare, and only appear when compensated elasticity ε^c is high
- Income effects make them rarer still; reach shrinks as η increases in magnitude

Comp. elasticity	Income effect	States with reform	Avg. population reach
0.2	0.0	0	—
0.2	-0.2	0	—
0.2	-0.5	0	—
0.5	0.0	1	0.6%
0.5	-0.2	1	0.0%
0.5	-0.5	1	0.0%
0.8	0.0	44	4.6%
0.8	-0.2	44	2.2%
0.8	-0.5	44	0.7%

Outline

- 1 Framework for income-tax reform
- 2 Reforms that benefit everyone
- 3 Reforms robust across (most) welfare functions**
- 4 Identifying and constructing robust reforms
- 5 Application to U.S. income tax schedules
- 6 Conclusion

- To assess implications of tax reforms, we need a common scale to compare workers
- For this, we use consumption-equivalent well-being E_i
- Specifically, the consumption at baseline earnings y_i^0 that leaves i indifferent

$$E_i(c_i, y_i) = k \iff u_i(c_i, y_i) = u_i(k, y_i^0)$$

- In other words, a worker's willingness to pay for a reform is the change in E_i

The robustness criterion

- We evaluate reforms with inequality-averse planners
- Each values progressive transfers in consumption-equivalent well-being E_i
- These are exactly the welfare functions

$$V^\Psi(c, y) = \int \Psi(E_i) di, \quad \Psi \text{ increasing and concave}$$

- Ψ increasing ensures Pareto; Ψ concave ensures favouring equality in E_i
- A reform is *robustly welfare-improving* if it raises welfare for *every* such Ψ
 - Stronger than committing to one welfare function
 - Weaker than requiring everyone be made better off

What robustness requires

Robustness, transparently

A reform h is robustly welfare-improving if and only if

$$\int_0^t h(y) F_Y(dy) \leq 0 \quad \text{for every income cutoff } t.$$

- An inequality-averse planner values progressive transfers
- But we fix *no* degree of aversion, so a reform must look redistributive to all of them at once
- That is exactly the condition: net gains to everyone below *every* cutoff
 - Enough to satisfy even the planner most focused on the bottom
 - Fail it anywhere, and a planner averse enough to weight that cutoff objects

Outline

- 1 Framework for income-tax reform
- 2 Reforms that benefit everyone
- 3 Reforms robust across (most) welfare functions
- 4 Identifying and constructing robust reforms**
- 5 Application to U.S. income tax schedules
- 6 Conclusion

A sufficient condition

Sufficiency

If there are incomes $x_1 < x_2$ with

$$M(x_1) < M(x_2),$$

then the reform


$$h = -\rho(x_2) b_{x_1} + \rho(x_1) b_{x_2}$$

is budget-neutral and robustly welfare-improving.

- Constructive: cut taxes at the lower income x_1 , raise them at the higher x_2
- Same direct test as before: look for a pair where M does not decrease

- Budget-neutral?
 - By construction
 - Each reform scaled by the other's revenue effect, so the two cancel
- Robustly welfare-improving?
 - Tax cuts fall on low incomes, increases on high
 - $M(x_1) < M(x_2)$ ensures that bottom cut costs less revenue per unit of welfare than top rise
 - So one can give the bottom more in consumption-equivalents than the top loses
 - Every inequality-averse planner ranks such a transfer an improvement

Applying the test: Indiana



- Application of the test to Indiana's income-tax schedule
- Baseline behavioural parameters:
 - Compensated elasticity $\varepsilon^c = 0.3$
 - Income effect $\eta = -0.5$
- $M(y)$ is non-decreasing on $[50, 120]$ and $[180, 230]$ (in \$1,000s)
- Robust budget-neutral welfare-improving reforms exist in these ranges

The condition is also necessary

Characterization

A nontrivial robust budget-neutral reform exists if and only if there are incomes $x_1 < x_2$ with

$$M(x_1) < M(x_2).$$

- As for the everyone-better-off case, the condition is also necessary
- So the test is complete: if a robust reform exists, it is sufficient to look at M 's shape

Intuition: necessity

- Robustness has a sharp meaning: from the bottom up, the cumulative gain must never turn negative
- Suppose, for contradiction, M fell everywhere in income
 - Revenue raised per unit of welfare is then lowest at the top, highest at the bottom
- A reform favouring the bottom must cut low and raise high
 - But cutting low forgoes a lot of revenue, raising high yields little
- Revenue can't balance, so one violates budget-neutrality
- So a robust budget-neutral reform exists only if M rises somewhere

Outline

- 1 Framework for income-tax reform
- 2 Reforms that benefit everyone
- 3 Reforms robust across (most) welfare functions
- 4 Identifying and constructing robust reforms
- 5 Application to U.S. income tax schedules**
- 6 Conclusion

Robust reforms exist almost everywhere

- Apply the test to each state's combined income-tax schedule
- Robust reforms exist in nearly every state, across almost all parameters
 - Most states for (almost) all behavioural responses
 - Far more prevalent than reforms making everyone better off
- What varies most is reach: the population share a reform covers

ε^c	η	Everyone-better-off states	Robust states	Robust reach
0.2	0.0	0	20	5.5%
0.2	-0.2	0	50	38.4%
0.2	-0.5	0	51	70.3%
0.5	0.0	1	49	21.1%
0.5	-0.2	1	50	31.9%
0.5	-0.5	1	51	43.7%
0.8	0.0	44	49	24.7%
0.8	-0.2	44	50	30.6%
0.8	-0.5	44	51	37.6%

Outline

- 1 Framework for income-tax reform
- 2 Reforms that benefit everyone
- 3 Reforms robust across (most) welfare functions
- 4 Identifying and constructing robust reforms
- 5 Application to U.S. income tax schedules
- 6 Conclusion

- A middle ground between making everyone better off and using single welfare function
 - Reforms that improve welfare for every inequality-averse planner
- Can identify and construct reforms using revenue per welfare M of elementary reforms
 - A robust welfare improving reform exists iff M is non-decreasing
- Easy to apply: M reduces to standard sufficient statistics
- Across U.S. states, such reforms exist almost everywhere
 - Far more prevalent than reforms making everyone better off