The Compensation Gap: A Unit-Invariant Measure for Distributional Policy Evaluation

Audric De Bevere* Gilles Grandjean†

Draft – July 24, 2026

Abstract

Winner and loser headcounts are widely used to evaluate the distributional effects of public policies, but their policy ranking can change when the statistical unit switches from households to individuals. We introduce the *compensation gap*: the average monetary shortfall that must be filled to keep losses below a chosen threshold. For any policy, population subgroup, and common loss threshold, reporting this shortfall per household or per person changes its scale but not the policy ranking. The measure combines the incidence and intensity of losses and has a direct interpretation as a residual compensation budget. We illustrate the measure using a Belgian carbon-pricing microsimulation. Headcounts rank equal per-household and equal per-capita dividends differently depending on the statistical unit. The compensation gap yields a unique ranking: eliminating the remaining losses requires €6.03 per household per month under the per-household dividend and €6.48 under the per-capita dividend.

Keywords: policy evaluation; distributional incidence; compensation; unit of analysis; carbon pricing; revenue recycling.
JEL codes: H22; H23; D31; Q52; Q58.

1 Introduction

Distributional evaluations often report the shares of winners and losers created by a policy. Such headcounts are intuitive, but they do not necessarily produce a unique policy ranking. Household-level statistics give every household the same weight, whereas individual-level statistics give greater weight to outcomes experienced by members of larger households. Whenever policy effects are related to household size, changing the statistical unit can therefore change the ranking of policies.

The problem can arise when comparing taxes, benefit reforms, subsidies, administered price changes, or any other policy that affects resources at the household level. Both household and individual headcounts remain meaningful: they simply answer different questions. The

*UCLouvain Saint-Louis Bruxelles; FRESH grantee of the Fonds de la recherche scientifique–FNRS.
†UCLouvain Saint-Louis Bruxelles.

difficulty arises when a policy analyst seeks an overall comparison and the conclusion depends on which population units are counted.

We propose the *compensation gap* to address this measurement problem. For a given policy and tolerated loss, it sums the monetary shortfalls of units whose post-policy resources fall below their policy-specific benchmark. The resulting total is the ex-post budget that would be required, under perfect targeting and abstracting from administrative costs, to bring every affected household up to that benchmark. Normalising this budget by the number of households or by the number of people produces different monetary scales but the same policy ranking.

The construction is related to the poverty-gap measure (Foster et al. 1984). The important difference is the reference line. A conventional poverty gap measures shortfalls from a common poverty threshold. The compensation gap measures shortfalls from household-specific benchmarks based on pre-policy resources and a common tolerated loss. It is therefore designed to compare how effectively alternative policies limit adverse changes, rather than to measure the level of poverty.

We first define household and individual headcounts and show why their rankings may differ. We then establish the rank invariance of the compensation gap. As an illustration, we compare equal per-household and equal per-capita recycling of carbon-pricing revenue in Belgium. Carbon-pricing studies commonly examine both the incidence and heterogeneity of household losses (Pizer and Sexton 2019; Cronin et al. 2019; Douenne 2020), while uniform household or per-capita dividends provide prominent recycling benchmarks (Berry 2019; Fremstad and Paul 2019; Budolfson et al. 2021). This application makes the unit-of-analysis problem especially transparent, without limiting the indicator's broader use.

2 Measurement

2.1 Headcounts and the unit of analysis

Let h index households, n_h denote household size, and $w_h > 0$ be a survey weight. Consider a fixed subgroup S and a policy r. Household h has benchmark disposable income Y_h and post-policy disposable income Y_{hr}. For a tolerated proportional loss $x \in [0, 1]$, define its monetary shortfall as

$$g_{hr}(x) = [(1 - x)Y_h - Y_{hr}]_+ , \qquad [a]_+ \equiv \max(a, 0). \tag{1}$$

Thus, $g_{hr}(x) > 0$ if household h loses more than x of benchmark income under policy r.

Let $W_S = \sum_{h \in S} w_h$ be the weighted number of households and $N_S = \sum_{h \in S} w_h n_h$ the corresponding population. The household and individual loser headcounts are

$$L_r^H(S, x) = \frac{1}{W_S} \sum_{h \in S} w_h \mathbf{1}\{g_{hr}(x) > 0\} , \tag{2}$$

$$L_r^I(S, x) = \frac{1}{N_S} \sum_{h \in S} w_h n_h \mathbf{1}\{g_{hr}(x) > 0\} . \tag{3}$$

Here, $\mathbf{1}\{\cdot\}$ is the indicator function, equal to one when the condition holds and zero otherwise. Equation (3) assigns the outcome of a household to each of its members. It consequently gives household h a weight proportional to n_h, whereas equation (2) gives it weight w_h. If the households protected by one policy differ systematically in size from those protected by another, the two headcounts can rank the policies differently.

2.2 The compensation gap

Define the total residual compensation budget as

$$B_r(S, x) = \sum_{h \in S} w_h g_{hr}(x). \tag{4}$$

The compensation gaps per household and per person are, respectively,

$$\mathrm{CG}_r^H(S, x) = \frac{B_r(S, x)}{W_S}, \qquad \mathrm{CG}_r^I(S, x) = \frac{B_r(S, x)}{N_S}. \tag{5}$$

The first is the average amount per household required to prevent losses exceeding x; the second reports the same aggregate requirement per person. At $x = 0$, the benchmark is pre-policy income and the gap is the amount needed to eliminate all monetary losses. A positive x focuses the measure on losses exceeding a tolerated proportion of income.

Proposition 1 (Rank invariance)**.** *For any two policies r and s, evaluated for the same subgroup S and tolerance x,*

$$\mathrm{CG}_r^H(S, x) \leq \mathrm{CG}_s^H(S, x) \quad \Longleftrightarrow \quad \mathrm{CG}_r^I(S, x) \leq \mathrm{CG}_s^I(S, x).$$

Proof. Both normalising factors, W_S and N_S, are positive and common to the policies. Each inequality is therefore equivalent to $B_r(S, x) \leq B_s(S, x)$. □

Rank invariance is specific to the compensation gap: headcounts and other distributional indicators may still depend on the unit of analysis. The gap should therefore be used alongside headcounts and subgroup results, since two policies can require the same aggregate compensation while imposing losses on different households. Nor is the gap a welfare measure. Its interpretation requires a cardinal monetary outcome, a common benchmark, a common subgroup, and a common tolerance. Within those conditions, it combines information discarded by a headcount: a policy performs worse if it leaves either more uncompensated units or larger shortfalls among them.

3 Application: two carbon-dividend designs

3.1 Microsimulation and policies

We illustrate the measure using the 2016 Belgian Household Budget Survey, covering approximately 4,500 households representative of the Belgian population. We simulate the immediate direct effect of a €45/tCO_2 price on residential natural gas and heating oil and on gasoline and diesel. Fuel quantities are held fixed, the carbon price is fully passed through to consumers, and only carbon payments made by households are recycled. Survey weights are used throughout.

The same revenue is returned under two budget-neutral rules. Under the *Household* design, every household receives €21.1 per month. Under the *Capita* design, every person receives €9.2 per month. The first provides relatively greater protection to smaller households, while the second directs larger transfers to households with more members. The application is deliberately narrow: its role is to expose the measurement issue, not to suggest that the compensation gap applies only to these two policies.

3.2 Headcount reversal

Table 1 reports winners and large losers. A winner has post-policy income above benchmark income; a large loser experiences a loss exceeding 0.5% of disposable income. With households as the unit, Household produces more winners and fewer large losers than Capita. With individuals as the unit, both rankings reverse.

Table 1: Headcount indicators by statistical unit

	Household dividend	Per-capita dividend
Households as the unit of analysis		
Winners (%)	56.8	51.2
Large losers (%)	14.8	20.0
Individuals as the unit of analysis		
Winners (%)	51.5	61.0
Large losers (%)	15.1	13.3

Notes: A large loser has a net carbon burden exceeding 0.5% of disposable income. Individual-level rates assign every person the outcome of their household. All estimates use survey weights.

Neither set of headcounts is incorrect. Household headcounts describe the proportion of family units affected, while individual headcounts describe the proportion of people living in those units. Their disagreement nevertheless prevents a unit-independent policy ranking.

3.3 Compensation-gap ranking

Table 2 reports the compensation gap at zero tolerance. Without recycling, eliminating all losses would require an average payment of €21.10 per household per month. The requirement falls to €6.03 under Household and €6.48 under Capita. Household therefore leaves a residual shortfall that is €0.45 smaller per household per month. Reporting the same budgets per

person would rescale both values by the same average household size and leave this ranking unchanged.

Table 2: Compensation gap at zero tolerance, monthly euros per household

	No redistribution	Household dividend	Per-capita dividend
All households	21.10	6.03	6.48
Energy-poor households	—	8.62	13.18

Notes: Energy poverty is measured before the policy using the Belgian energy-poverty indicator (Fondation Roi Baudouin 2024). The no-redistribution value for this subgroup is not reported.

The difference is larger among households classified as energy poor before the policy: the gap is €8.62 under Household and €13.18 under Capita. This subgroup calculation illustrates an important safeguard. The aggregate gap supplies a unique ranking of the remaining monetary requirement, while subgroup gaps reveal where that requirement is concentrated.

The ranking also holds when the tolerated loss is set to 0.5% of disposable income: the gaps fall to €1.64 under Household and €2.33 under Capita. Unlike the corresponding large-loser headcount, the compensation gap records how far households lie below the tolerated benchmark. The result is substantively consistent across the two thresholds in this application, although policy rankings need not be invariant to the choice of tolerance.

4 Conclusion

Winner and loser headcounts may rank policies differently when analysts switch between households and individuals. The compensation gap addresses this ambiguity by measuring the residual monetary budget required to keep policy losses below a common threshold. Because its per-household and per-person versions divide the same aggregate shortfall by constants common to all policies, their policy ranking is identical.

The measure is general: it can be applied to any policy represented by comparable monetary resources, to any fixed population subgroup, and to any common tolerated loss. In the Belgian carbon-dividend illustration, headcount rankings reverse with the statistical unit, whereas the compensation gap consistently favours the equal per-household dividend.

The compensation gap is intended to complement, not replace, existing distributional statistics. Headcounts show how many households or people lose, subgroup results identify who is exposed, and the gap quantifies the remaining intensity of losses in directly interpretable budgetary units. Reporting these objects together separates a normative choice about whom to count from the practical question of how much compensation remains necessary.

References

Berry, Audrey (2019). "The Distributional Effects of a Carbon Tax and Its Impact on Fuel Poverty: A Microsimulation Study in the French Context". In: *Energy Policy* 124, pp. 81–94. DOI: 10.1016/j.enpol.2018.09.021.

Budolfson, Mark et al. (2021). "Protecting the Poor with a Carbon Tax and Equal Per Capita Dividend". In: *Nature Climate Change* 11.12, pp. 1025–1026. DOI: 10.1038/s41558-021-01228-x.

Cronin, Julie Anne, Don Fullerton, and Steven Sexton (2019). "Vertical and Horizontal Redistributions from a Carbon Tax and Rebate". In: *Journal of the Association of Environmental and Resource Economists* 6.S1, S169–S208. DOI: 10.1086/701191.

Douenne, Thomas (2020). "The Vertical and Horizontal Distributive Effects of Energy Taxes: A Case Study of a French Policy". In: *The Energy Journal* 41.3. DOI: 10.5547/01956574.41.3.tdou.

Fondation Roi Baudouin (2024). *Baromètre de la précarité énergétique 2024*.

Foster, James, Joel Greer, and Erik Thorbecke (1984). "A Class of Decomposable Poverty Measures". In: *Econometrica* 52.3, pp. 761–766. DOI: 10.2307/1913475.

Fremstad, Anders and Mark Paul (2019). "The Impact of a Carbon Tax on Inequality". In: *Ecological Economics* 163, pp. 88–97. DOI: 10.1016/j.ecolecon.2019.04.016.

Pizer, William and Steven Sexton (2019). "The Distributional Impacts of Energy Taxes". In: *Review of Environmental Economics and Policy* 13.1, pp. 104–123. DOI: 10.1093/reep/rey021.