Journal of Sustainable Development and Green Technology

♦ ISSN: 2836-5399 Vol : 06, No. 02, 2026 ♦ PP. 53–61

♦ Article DOI: https://doi.org/10.54216/JSDGT.060206



FULL LENGTH ARTICLE

Sustainability Policy for Urban Air Quality through Equity Constrained Abatement Allocation

Adam Marks[1,*]

[1] Higher Education Leadership Academy, Dubai, United Arab Emirates

Email: adam.marks@ihele.org

Received: February 05, 2026 Revised: March 10, 2026 Accepted: April 24, 2026 ⋆ Corresponding author

ABSTRACT

Two allocations can give the same mean reduction in fine particulate matter ($PM_{2.5}$) and still leave different cities above the reference concentration. We examine this choice with 2021 records for 40 Spanish cities from the World Health Organization's January 2024 database. The model minimises population-weighted squared excess above 5 $\mu g/m^3$ subject to a resource budget, local reduction caps and an optional minimum share for each city, and numerical optimisation confirms its analytical solution. At half of assumed total capacity, optimal allocation and proportional allocation both bring the weighted mean concentration down from 9.882 to 8.404 $\mu g/m^3$; optimal allocation cuts squared excess loss by 9.44% and between-city excess variance by 83.39% relative to proportional allocation. A minimum share of one-quarter of each local reduction cap extends reductions to every city above the reference concentration and raises squared excess loss by 3.01%. A separate feasibility test puts a common concentration ceiling of 10 $\mu g/m^3$ within reach at 17.12% of assumed total capacity. A ceiling of 9 $\mu g/m^3$ would push one city past its local reduction cap. These conditional planning results set targets for later assessment of specific interventions.

Keywords: Sustainability policy ▪ Urban air quality ▪ Distributional equity ▪ Convex optimisation ▪ Particulate matter

1. INTRODUCTION

A programme can bring down the mean concentration across a portfolio of cities while the cities with the highest concentrations stay well above the target. Equal percentage cuts have a different weakness: giving every municipality the same percentage reduction is easy to defend on grounds of participation, but residents of different cities would still face markedly different concentrations once the cuts had been made. An authority that sets the target is therefore also deciding, explicitly or not, which disparities the programme will address and which local constraints it will accept.

Goforth and Nock [1] model how alternative energy transitions distribute air pollution, and Picciano et al. [2] find disparities that persist after average exposure has fallen; both work forward from emissions through an atmospheric model. Our starting point comes later. An authority has already chosen an aggregate target and now has to decide which cities should deliver the improvement needed to reach it, and that allocation decision is what we model.

City means hide a good deal. Rentschler and Leonova [3] link pollution exposure to poverty, while Behrer and Heft-Neal [4] report that in most low- and middle-income countries it is the wealthier districts that are more polluted, so the direction of the association depends on the setting. Scale enters as well. With nationwide individual-level data from the Netherlands, van den Brekel et al. [5] find exposure differences that follow the intersection of ethnic and socioeconomic characteristics, and no city mean can resolve those. We use equity in a narrower sense, to mean disparities in concentration between

© 2026 The Author(s). Published by ASPG. This article is licensed under CC BY 4.0.

cities and the minimum share of local reduction capacity each municipality is guaranteed.

We ask, first, how a fixed resource budget should be allocated when the objective weights high residual excess more heavily than low residual excess. How much does squared excess loss rise when every eligible city (every city above the reference concentration) receives a minimum share? And when does a common concentration ceiling become infeasible despite adequate aggregate resources? Inputs and decisions are kept apart: we take observed concentrations C_i and published population weights w_i as given, solve for conditional reductions x_i under assumed local reduction caps u_i and resource coefficients a_i, derive the allocation for each city and then test how the result changes as those assumptions vary.

2. POLICY SETTING AND CITY SELECTION

2.1 Local air quality measures and concentration targets

Pisoni et al. [6] find that the response to local air quality policy varies across European cities. Even within one city the pollutant matters: Gómez-Losada and Pires [7] follow the phases of Madrid's low emission zone and separate the particulate response from the responses of other pollutants, which makes it hard to justify carrying any single intervention effect over to every city. Concentration reductions are solved for first. Evidence on transport, heating and industrial measures would then establish which of those reductions a city could deliver.

Gentry et al. [8] connect emissions to exposure inequality with a reduced-complexity atmospheric model, and Wang et al. [9] examine emissions strategies that vary by location. Both need source–receptor information. The workbook supplies concentrations without the relationships needed to trace pollution from emissions to each city, so each city's concentration reduction is allowed to vary independently when targets are compared, and transport across administrative boundaries is left to the implementation study that would follow.

2.2 Monitoring coverage and population weights

We use the January 2024 edition V6.1 of the World Health Organization (WHO) Ambient Air Quality Database [10]. Of its 333 Spanish records for 2021, 197 report $PM_{2.5}$. Applying the 75% annual monitoring coverage requirement leaves 121 of them, and once a positive population figure is also required, 40 unique city records remain (coverage among them runs from 77% to 100%, median 97%, and their reported populations add up to 16,965,345).

These 40 cities are in the analysis because their records are complete. We use the published population field as a scenario weight, because the workbook has no harmonised 2021 census and does not match population boundaries to monitoring boundaries, so the weighted means that result cannot be read as estimates of personal exposure; Shairsingh et al. [11] set out the measurement differences behind WHO's caution about ranking cities or comparing them across countries. The cities are listed alphabetically and allocations are compared only within this portfolio, with no inference drawn about national prevalence, within-city inequality or the causal effect of any particular policy.

3. MATHEMATICAL FRAMEWORK

3.1 Objective and feasible allocations

The reference concentration is $\tau = 5\ \mu g/m^3$, the annual guideline value discussed by Rentschler and Leonova [3], and excess d_i is measured above it. Equation (1) also sets the local reduction cap u_i with $\rho = 0.30$: however much of the resource budget remains, no city can be assigned a reduction larger than 30% of its baseline concentration, or larger than its excess where that is smaller. We do not treat τ as a zero-risk threshold or a legal limit. ρ is an assumption, and an engineering estimate for a particular intervention would need its own evidence.

$$\begin{aligned} d_i &= \max(C_i - \tau, 0), \\ u_i &= \min(d_i, \rho C_i), \\ w_i &= \frac{P_i}{\sum_j P_j}, \qquad \tau = 5, \quad \rho = 0.30 \end{aligned} \tag{1}$$

For each city the decision x_i is a concentration reduction in $\mu g/m^3$. Each city also carries an objective weight v_i, a normalised population weight w_i and a relative resource coefficient a_i; the central specification, used to compare allocation rules at the same resource budget, sets $v_i = w_i$ and $a_i = 1$. The minimum share η enters as a lower bound, $x_i \geq \eta u_i$, so what it guarantees is a fraction of the local reduction capacity assigned to a city. Equation (2) gives the objective and both constraints.

$$\begin{aligned} \min_{\mathbf{x}} \quad & J(\mathbf{x}) = \sum_i v_i (d_i - x_i)^2 \\ \text{subject to} \quad & \sum_i w_i a_i x_i \leq B, \\ & \eta u_i \leq x_i \leq u_i. \end{aligned} \tag{2}$$

When $a_i = 1$, one unit of the resource budget buys one unit of population-weighted concentration reduction in the model, and the budget is measured in concentration units. Define $B_{\max} = \sum_i w_i u_i = 2.954836\ \mu g/m^3$, the assumed total capacity, and $B = bB_{\max}$, with b running from zero to one in the central specification ($b = 0.50$ is called half capacity below). A city below τ has $u_i = 0$. Only cities with positive capacity can receive a reduction, and with equal resource coefficients the minimum share is affordable when $\eta \leq b \leq 1$. Monetary expenditure and emissions tonnage lie outside this accounting.

Figure 1 traces the allocation between A Coruna and Alcala De Henares at $b = 0.50$. Holding the other allocations at the 40-city optimum leaves these two cities a resource constraint, drawn as a straight line across contours of their joint contribution to squared excess loss, and the marked point is the solution to the full problem seen in these two coordinates.

3.2 Derivation and uniqueness

Squared excess loss is strictly convex because every objective weight v_i is positive. Once the resource budget covers the lower bounds on city reductions, the feasible allocations form a nonempty, closed and bounded convex set, and the optimum is unique. Equations (3) and (4) use 2λ as the multiplier on the resource constraint and nonnegative multipliers α_i and β_i

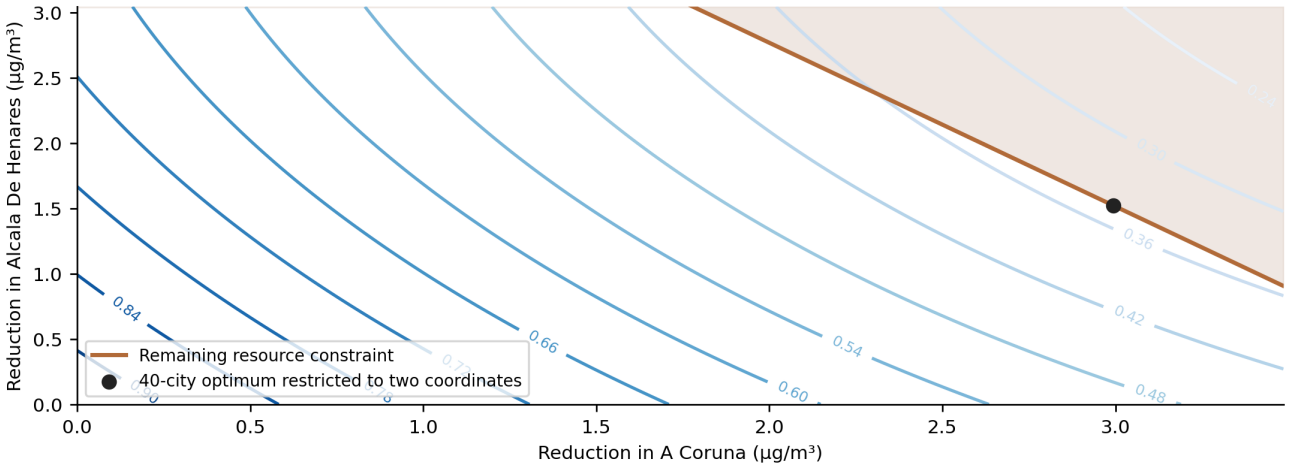


Figure 1. Geometry of the constrained allocation problem. Author calculations for A Coruna and Alcala De Henares, with $b = 0.50$, $\eta = 0$ and the allocations to the other cities fixed at the full optimum. Shading marks allocations that exceed the remaining resource budget; contours give the two cities' contribution to squared excess loss.

on the lower and upper bounds, respectively.

$$\begin{aligned}\mathscr{L} &= J(\mathbf{x}) \\ &\quad + 2\lambda\left(\sum_i w_i a_i x_i - B\right) \\ &\quad + \sum_i \alpha_i(\eta u_i - x_i) + \sum_i \beta_i(x_i - u_i)\end{aligned} \tag{3}$$

$$-2v_i(d_i - x_i) + 2\lambda w_i a_i - \alpha_i + \beta_i = 0 \tag{4}$$

A slack constraint has a zero multiplier by complementary slackness. At an interior solution the marginal reduction in squared excess loss per unit of weighted resource is equal across cities, and clipping that interior value to the lower or upper bound, wherever it would fall outside the feasible interval, gives equation (5).

$$x_i^* = \min\left\{u_i, \max\left[\eta u_i,\ d_i - \frac{\lambda w_i a_i}{v_i}\right]\right\} \tag{5}$$

For a budget strictly between the lower-bound requirement and saturation, λ is chosen to satisfy $\sum_i w_i a_i x_i = B$; the left-hand side is continuous and non-increasing in λ, so bisection or any other bracketed root solver will find it. At the endpoints the allocation may still be unique even though more than one multiplier satisfies the conditions, and we report a shadow price only where λ is identified (Table 3 shows a dash at $\eta = b = 0.50$ for this reason); city reductions are labelled zero, interior or capped using a tolerance of 10^{-9}.

3.3 Distributional accounting and feasibility tests

With $a_i = 1$, any allocation that spends the whole resource budget lowers the population-weighted mean concentration by exactly B, so the weighted mean cannot tell such allocations apart. Residual excess $e_i = d_i - x_i$ is more informative. Its squared weighted mean is the same for all of them; what differs is the weighted variance of residual excess across cities, which we call between-city excess variance, and equation (6) splits the central objective into these two parts.

$$\begin{aligned}\bar{C}_{\text{post}} &= \bar{C}_0 - B, \\ \mu_e &= \sum_i w_i e_i, \\ J &= \mu_e^2 + \mathrm{Var}_w(e), \\ \mathrm{Var}_w(e) &= \sum_i w_i (e_i - \mu_e)^2\end{aligned} \tag{6}$$

At a fixed binding budget the first term in equation (6) is a constant, and minimising J is then the same as minimising between-city excess variance. That equivalence is for comparisons at one value of B, with population weights and equal resource coefficients. Across budgets it can fail, and Section 4.2 shows it doing so: once local reduction caps bind, extra resources can go only to cities whose concentrations are already relatively low, and between-city excess variance may then rise.

For every city to reach a common concentration ceiling $T \geq \tau$, each must receive at least $r_i(T) = \max(C_i - T, 0)$. Equation (7), written for $\eta = 0$, adds those reductions up into a resource requirement and states the separate bound that each city has to satisfy.

$$\begin{aligned}B_{\text{req}}(T) &= \sum_i w_i a_i \max(C_i - T, 0), \\ r_i(T) &= \max(C_i - T, 0), \quad r_i(T) \leq u_i \text{ for every } i, \\ T &\geq \tau\end{aligned} \tag{7}$$

The aggregate requirement $B_{\text{req}}(T)$ can fit within the resource budget even when the reduction required in one city exceeds u_i, and a planner could then enlarge the aggregate resource budget and still miss the ceiling. In Section 4.4 that city is Murcia. At an interior budget where the optimal value function is differentiable, equation (8) gives the marginal change in squared excess loss as the budget grows.

$$\frac{dJ^*}{dB} = -2\lambda \tag{8}$$

The derivative in equation (8) is measured in squared excess loss per additional resource unit. It has no monetary meaning.

Table 1. City observations and conditional reductions at half capacity

ID	City	Pop. M	C	Cover %	Cap. u	x opt.	x floor	C opt.
01	A Coruna	0.248	11.617	98	3.485	2.993	2.769	8.624
02	Alcala De Henares	0.198	10.147	98	3.044	1.523	1.299	8.624
03	Alicante	0.370	8.944	94	2.683	0.320	0.671	8.624
04	Almeria	0.201	6.953	97	1.953	0.000	0.488	6.953
05	Arrecife	0.065	6.961	85	1.961	0.000	0.490	6.961
06	Aviles	0.078	8.363	99	2.509	0.000	0.627	8.363
07	Bilbao	0.350	10.051	90	3.015	1.427	1.203	8.624
08	Castellon de la plana	0.174	10.422	93	3.127	1.798	1.574	8.624
09	Ceuta	0.084	7.736	94	2.321	0.000	0.580	7.736
10	Collado Villalba	0.064	10.850	99	3.255	2.226	2.002	8.624
11	Cordoba	0.344	12.250	98	3.675	3.626	3.402	8.624
12	Elda	0.087	8.887	91	2.666	0.263	0.667	8.624
13	Ferrol	0.066	7.766	98	2.330	0.000	0.582	7.766
14	Gijon	0.272	10.259	98	3.078	1.635	1.411	8.624
15	Huelva	0.144	7.010	89	2.010	0.000	0.502	7.010
16	Jaen	0.113	9.869	97	2.961	1.245	1.021	8.624
17	Las Palmas	0.381	7.688	83	2.306	0.000	0.577	7.688
18	Linea De La Concepcion	0.064	9.996	78	2.999	1.372	1.148	8.624
19	Lugo	0.099	10.805	97	3.241	2.181	1.957	8.624
20	Madrid	6.618	10.471	77	3.141	1.847	1.623	8.624
21	Malaga	0.590	7.968	97	2.390	0.000	0.598	7.968
22	Murcia	0.500	12.951	100	3.885	3.885	3.885	9.066
23	Ourense	0.106	9.924	97	2.977	1.300	1.076	8.624
24	Oviedo	0.220	10.430	98	3.129	1.806	1.582	8.624
25	Pontevedra	0.083	9.424	94	2.827	0.800	0.707	8.624
26	Puerto De La Cruz	0.067	8.446	79	2.534	0.000	0.633	8.446
27	Sagunto	0.067	10.209	98	3.063	1.585	1.361	8.624
28	Salamanca	1.796	7.651	95	2.295	0.000	0.574	7.651
29	San Sebastian Donostia	0.188	8.337	96	2.501	0.000	0.625	8.337
30	Santa Cruz De Tenerife	0.368	8.931	89	2.679	0.307	0.670	8.624
31	Santiago De Compostela	0.098	8.793	98	2.638	0.169	0.659	8.624
32	Sevilla	0.704	10.844	81	3.253	2.220	1.996	8.624
33	Telde	0.214	7.600	97	2.280	0.000	0.570	7.600
34	Toledo	0.086	11.713	93	3.514	3.089	2.865	8.624
35	Torrevieja	0.085	4.573	100	0.000	0.000	0.000	4.573
36	Valdemoro	0.077	11.446	99	3.434	2.822	2.598	8.624
37	Valencia	0.834	10.622	83	3.187	1.998	1.774	8.624
38	Valladolid	0.305	11.084	93	3.325	2.460	2.236	8.624
39	Vigo	0.307	12.219	98	3.666	3.595	3.371	8.624
40	Vitoria Gasteiz	0.254	9.517	98	2.855	0.893	0.714	8.624

Source data are the 2021 Spanish records from WHO [10]. Pop. M gives reported population in millions and Cover gives annual $PM_{2.5}$ monitoring coverage. Concentrations and reductions are in $\mu g/m^3$. C is observed. Cap. u is the assumed local reduction cap, x opt. and C opt. are the optimal reduction and resulting concentration at $\eta = 0$, and x floor is the optimal reduction at $\eta = 0.25$. City IDs follow the alphabetical source names. The population weights come from the workbook and may differ from harmonised 2021 counts.

As a check on the clipped solution, we solved the same problem by constrained sequential quadratic optimisation in three configurations; objective values agreed to within 1.6×10^{-11}, and no city's allocation differed by as much as 3.5×10^{-5} $\mu g/m^3$. Both the resource budget and the local bounds were verified on the unrounded values.

4. RESULTS

4.1 City records and the central allocation

Table 1 lists the 40 observed concentrations, from 4.573 $\mu g/m^3$ in Torrevieja to 12.951 $\mu g/m^3$ in Murcia, together with each city's conditional reduction at $b = 0.50$. Torrevieja is the only city below the reference concentration. The baseline population-weighted mean is 9.882 $\mu g/m^3$. Madrid alone, with a reported 6.618 million residents, carries about 39% of the weight, so a reduction there moves the weighted mean far more than the same reduction in Elda or Ceuta.

Figure 2 follows the alphabetical order of Table 1. At the optimum, 13 cities receive zero reduction (one of them the city already below τ), 26 receive an interior allocation and end at $\tau + \lambda = 8.624011$ $\mu g/m^3$, and one reaches its local reduction cap. The 13 cities without a reduction all start below 8.624 $\mu g/m^3$ and keep their observed concentrations; Salamanca, at 7.651 $\mu g/m^3$ and with the second-largest weight, is one of them. Murcia is the capped city. Its full reduction of 3.885 $\mu g/m^3$ takes it only to 9.066 $\mu g/m^3$, and we read the zero allocations elsewhere as the model's priority at this budget, which says nothing about local environmental need.

4.2 Mean concentration and between-city excess variance

At $b = 0.50$ the resource budget is $B = 1.477418$ $\mu g/m^3$, and optimal allocation and proportional allocation ($x_i = bu_i$) both leave a weighted mean concentration of 8.404358 $\mu g/m^3$. The residual excess ends up in different cities. Under optimal allocation squared excess loss J is 11.850138, against 13.085168 under proportional allocation (9.44% lower), and between-city excess variance is 0.245969 against 1.480998 (83.39% lower). Proportional allocation leaves Murcia, Cordoba and Vigo above 10 $\mu g/m^3$; optimal allocation leaves no city there.

Equation (6) accounts for the whole gap between the two rules. Both leave a weighted mean excess of 3.406489 $\mu g/m^3$, the squared-mean terms cancel, and the 1.235030 $(\mu g/m^3)^2$ by which their squared excess losses differ is exactly the difference in between-city excess variance, so a programme scored on mean concentration alone would rate them as equal. Table 2 also contains a case that runs the other way. At $b = 0.10$ optimal allocation leaves 18 cities above 10 $\mu g/m^3$ and proportional allocation 14, because with so little to spend the loss falls fastest when the largest excesses are cut, and the number of cities above 10 $\mu g/m^3$ is not what the objective minimises.

From $b = 0.75$ to assumed total capacity, squared excess loss keeps falling (7.261824 to 4.698305) while between-city excess variance climbs from 0.144772 to 0.976988 (Figure 3). The local reduction caps are responsible: Murcia is already capped at $b = 0.50$ and Cordoba and Vigo are within 0.1 $\mu g/m^3$ of their caps, so the remaining resources can go only to cities with capacity left, and the weighted mean keeps dropping as the leftover excess becomes more uneven across the portfolio. The calculation concerns concentrations; no health effect is estimated from it.

Table 2. Resource budgets and residual concentration outcomes

b (%)	Rule	B	Mean C	J	Variance	$N > 10$
10	Optimal	0.295	9.586	22.385	1.331	18
10	Prop.	0.295	9.586	23.020	1.966	14
25	Optimal	0.739	9.143	17.964	0.781	0
25	Floor 25%	0.739	9.143	18.958	1.776	9
25	Prop.	0.739	9.143	18.958	1.776	9
50	Optimal	1.477	8.404	11.850	0.246	0
50	Floor 25%	1.477	8.404	12.206	0.602	0
50	Prop.	1.477	8.404	13.085	1.481	3
75	Optimal	2.216	7.666	7.262	0.145	0
75	Floor 25%	2.216	7.666	7.339	0.222	0
75	Prop.	2.216	7.666	8.332	1.215	1
100	Optimal	2.955	6.927	4.698	0.977	0

Note. B and Mean C in $\mu g/m^3$; squared excess loss J and between-city excess variance in $(\mu g/m^3)^2$. $N > 10$ is the number of cities above 10 $\mu g/m^3$. Proportional allocation sets $x = bu$. Floor 25% is optimal allocation with a minimum share $\eta = 0.25$ and is infeasible for $b < 25\%$. At assumed total capacity ($b = 100\%$) the three rules give the same allocation. All values are model scenarios.

4.3 Minimum shares and municipal participation

The minimum share decides who receives a reduction (Table 3). Guaranteeing 25% of each local reduction cap at half capacity raises squared excess loss from 11.850138 to 12.206235, which is 3.01% above the optimum without a minimum share, and cuts the number of cities with zero reduction from 13 to one; that one is Torrevieja, below the reference concentration and with no capacity. Every eligible city therefore gets a reduction, and the weighted mean concentration does not change.

Table 3. Minimum shares at the half-capacity budget

η (%)	Mean C	J	ΔJ (%)	Var.	No act.	λ
0	8.404	11.850	0.00	0.246	13	3.624
5	8.404	11.902	0.44	0.298	1	3.660
10	8.404	11.962	0.94	0.358	1	3.699
20	8.404	12.113	2.22	0.509	1	3.797
25	8.404	12.206	3.01	0.602	1	3.848
30	8.404	12.311	3.89	0.707	1	3.903
40	8.404	12.562	6.01	0.958	1	4.027
50	8.404	13.085	10.42	1.481	1	—

Note. η is the minimum share, expressed as a fraction of each local reduction cap. ΔJ is the percentage change in squared excess loss relative to $\eta = 0$. No action counts Torrevieja, which lies below τ. At $\eta = b = 0.50$ the lower bounds use up the whole resource budget, λ is not unique and is shown as a dash; units are as in Table 2.

With $\eta = 0.50$ at $b = 0.50$ the minimum shares alone spend

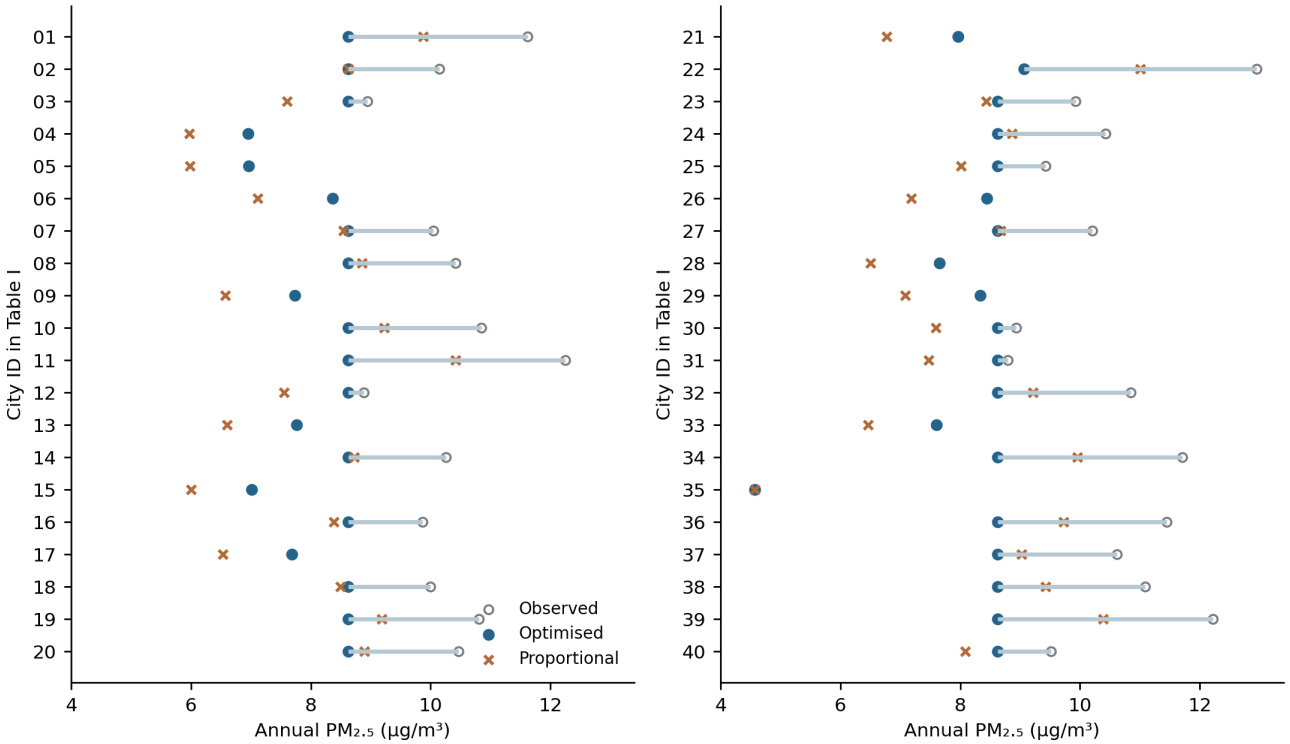


Figure 2. Observed and conditional city concentrations. Observed concentrations from WHO [10], with city IDs as in Table 1. For each city an open circle marks the observation and a line runs to the filled circle for its concentration under optimal allocation; crosses give proportional allocation, all at $b = 0.50$.

the entire resource budget, every city gets $x_i = 0.50u_i$, and the result is proportional allocation exactly (the 50% row of Table 3 repeats the proportional row of Table 2). If a planner wants cities with high residual excess to come first, the minimum share has to stay below the budget fraction b, because at $\eta = b$ the lower bounds settle the allocation and squared excess loss has no say in it.

4.4 Common ceilings and local bottlenecks

Table 4 checks common concentration ceilings city by city. A ceiling of 10 $\mu g/m^3$ needs $B_{req} = 0.505840$ $\mu g/m^3$, or 17.12% of assumed total capacity, with every city's required reduction inside its local reduction cap, and 9.5 $\mu g/m^3$ (28.80%) passes the same check. The 9 $\mu g/m^3$ ceiling fails. The arithmetic requirement is only 40.84% of assumed total capacity, yet one city would have to exceed its local reduction cap, and enlarging the aggregate resource budget cannot make that ceiling feasible.

The lowest feasible common concentration ceiling is $\max_i(C_i - u_i) = 9.0657$ $\mu g/m^3$, set by Murcia (12.951 − 3.885 $\mu g/m^3$), whereas at assumed total capacity the weighted mean concentration drops to 6.926941 $\mu g/m^3$ (Figure 4). A mean target is met through a weighted sum, and a large reduction in Madrid can stand in for small ones elsewhere. A common concentration ceiling has to be met in every city, and here Murcia's local reduction cap decides how low it can go.

Table 4. Common concentration ceilings and local feasibility

T	B req.	Cap. %	Feasible	Viol. caps
12.0	0.037	1.25	Yes	0
11.0	0.121	4.10	Yes	0
10.0	0.506	17.12	Yes	0
9.5	0.851	28.80	Yes	0
9.0	1.207	40.84	No	1
8.5	1.586	53.66	No	3
8.0	1.976	66.87	No	6
7.5	2.409	81.53	No	10
7.0	2.895	97.96	No	18
6.0	3.889	131.61	No	27
5.0	4.884	165.29	No	36

Note. T and Required B in $\mu g/m^3$. Required B gives the minimum resource requirement before local reduction caps are applied, with $a = 1$ and $\eta = 0$. Capacity % expresses that requirement as a percentage of assumed total capacity (2.954836). A "No" means at least one city would need to exceed the assumed 30% local reduction cap to reach T. The table tests analytical ceilings.

5. SENSITIVITY ANALYSIS AND POLICY INTERPRETATION

5.1 Objective weights and resource coefficients

In the central specification reported population sets both the objective weights and the population weights in the resource constraint, and the two sensitivity tests vary these separately. The first raises P_i to an exponent θ and normalises the result-

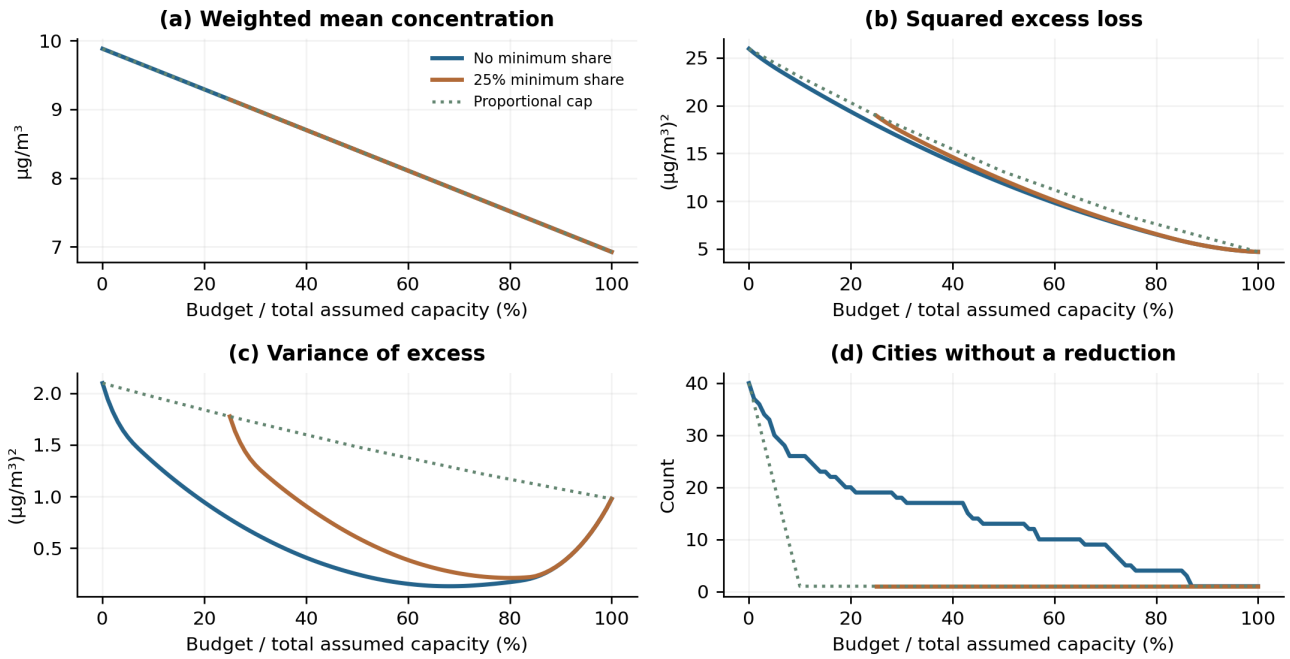


Figure 3. Mean improvement and distributional outcomes across budgets. Author calculations with $a = 1$ and $\rho = 0.30$. Panels show weighted mean concentration, squared excess loss, between-city excess variance and the number of cities receiving zero reduction. The 25% minimum-share curve starts at the smallest budget at which it is feasible, and the dotted segments for proportional allocation interpolate between the reported budget scenarios.

ing objective weights, with the resource constraint and B held fixed. The second keeps the objective weights at their central values and lets the resource coefficient rise with baseline excess. Equation (9) specifies both tests.

$$v_i(\theta) = \frac{P_i^{\theta}}{\sum_j P_j^{\theta}}, \qquad 0 \le \theta \le 1,$$
$$a_i(\gamma) = 1 + \gamma \frac{d_i}{\max_j d_j}, \qquad \gamma \ge 0 \tag{9}$$

At $\theta = 0$ every city carries the same objective weight. Evaluated with the original population weights, so that the two losses are comparable, the resulting allocation has a squared excess loss of 14.956753 against 11.850138 at $\theta = 1$; both allocations give a weighted mean concentration of 8.404358, since they deliver the same weighted physical reduction. The difference in squared excess loss is the consequence of choosing one objective. Equal city weights or population weights is a choice we leave to the authority, and the 26% higher loss is its price when the result is scored with population weights.

Raising γ to 1 (with $\eta = 0$) cuts the weighted physical reduction that the same resource budget buys from 1.477418 to 0.844191 μg/m^3, about 43% less, so the weighted mean concentration only reaches 9.037585 and three cities stay above 10 μg/m^3 (Table 5); in this test the mean and the dispersion both move. The result follows from assuming that resource coefficients rise with baseline excess, and real intervention costs would have to be measured before it could be taken further. Figure 5 puts this test next to the change in objective weights.

5.2 Implications for sustainability policy design

Under equal resource coefficients and a binding budget, a weighted mean target cannot distinguish the allocations that

Table 5. Sensitivity to assumed resource heterogeneity

γ	η (%)	B	Physical Δ	Mean C	J	$N > 10$
0.00	0	1.477	1.477	8.404	11.850	0
0.00	10	1.477	1.477	8.404	11.962	0
0.25	0	1.477	1.245	8.636	13.611	0
0.25	10	1.477	1.250	8.632	13.725	0
0.50	0	1.477	1.075	8.806	15.001	0
0.50	10	1.477	1.084	8.798	15.114	0
1.00	0	1.477	0.844	9.038	17.034	3
1.00	10	1.477	0.858	9.024	17.140	3

Note. Every row has a resource budget of 1.477418 in baseline-normalised units, with resource coefficients $a = 1 + \gamma d / \max(d)$. Physical Δ is $\sum_i w_i x_i$ in μg/m^3, the units of Mean C, and J is evaluated with population weights. $\eta = 0.10$ is feasible for every set of coefficients in the table.

Table 2 compares at any one budget. We would report three quantities beside it. Squared excess loss J combines the level of residual excess with its spread; between-city excess variance isolates the spread, and Section 4.2 shows it rising once local reduction caps bind even while J falls. A count of cities above 10 μg/m^3 is easy to communicate but says nothing about how far above they are, and at $b = 0.10$ it moves in the opposite direction to J.

In this model the instrument comes after the allocation. A planner fixes the objective, the minimum share and a feasible common concentration ceiling, and only then asks whether specific measures (a low emission zone, heating standards, controls on industrial sources) can deliver the assigned reductions at the assumed resource coefficients, which needs evidence on costs, emission sources, meteorology and atmospheric transport. Spillovers between cities, or cost differ-

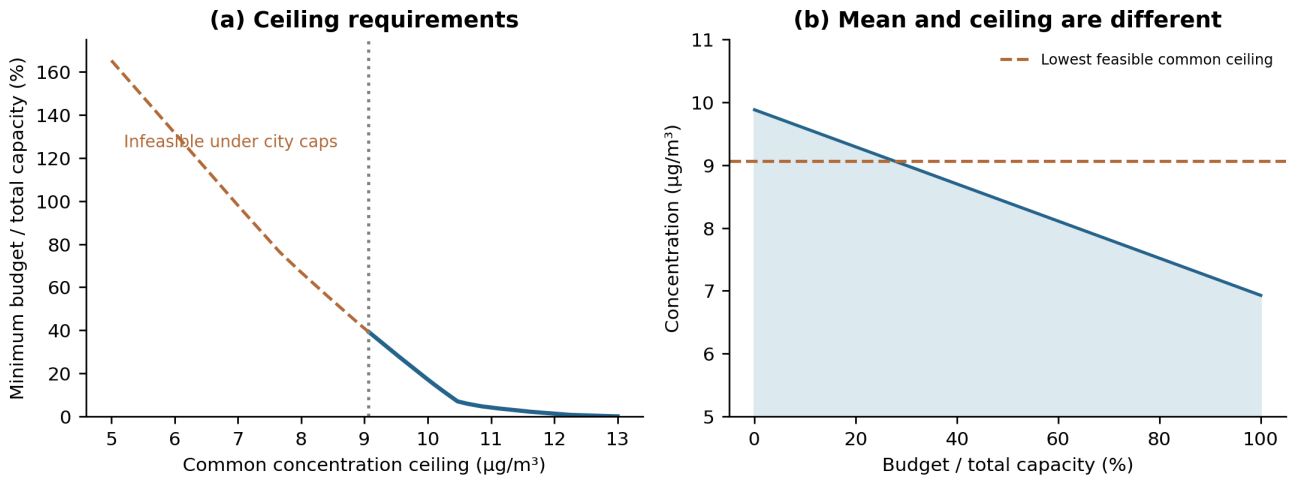


Figure 4. Aggregate resources and local ceiling constraints. Author calculations. In panel (a), the dashed segment gives arithmetic resource requirements where at least one local reduction cap prevents attainment, and the vertical guide marks the lowest feasible common concentration ceiling of 9.0657 $\mu g/m^3$. Panel (b) plots the weighted mean concentration under optimal allocation against this ceiling.

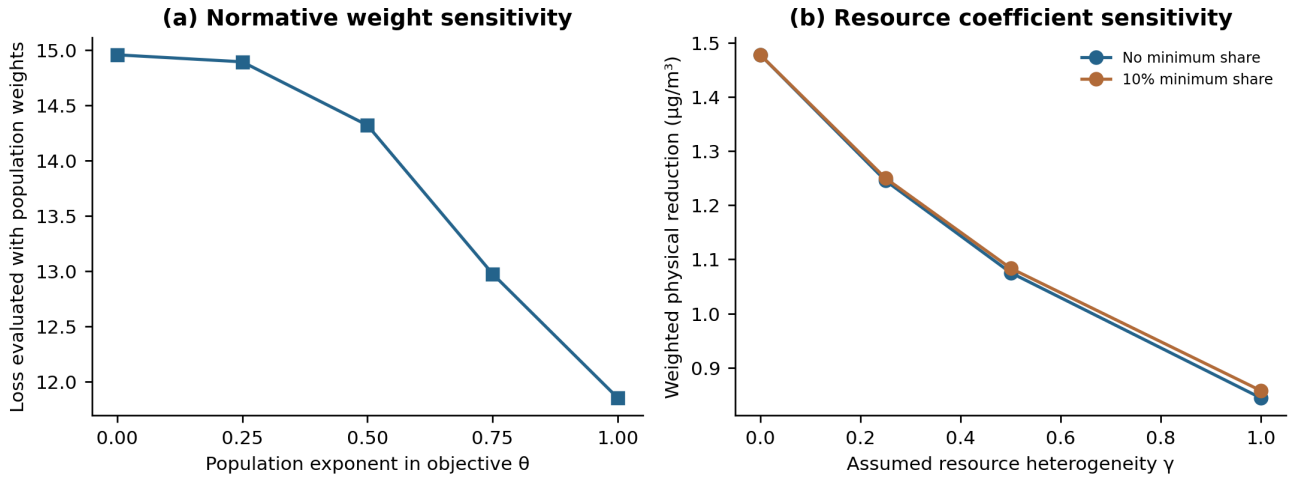


Figure 5. Sensitivity to objective weights and resource coefficients. Author calculations, $B = 1.477418$. Panel (a) evaluates each allocation with the reported population weights after varying the objective weights. In panel (b) the weighted physical reduction is plotted against γ, and θ and γ are both scenario parameters.

ences large enough to change the feasible set, would mean recalculating the allocation.

Monitoring is one place to start. Kelp et al. [12] place sensors to target environmental injustice, Considine et al. [13] treat access to monitoring and the accuracy of low-cost sensors as separate problems, and Wei et al. [14] estimate daily $PM_{2.5}$ on a 1 km grid, which would show variation within and between cities that a single monitor misses. None of this verifies the reductions themselves. That would take observations made before and after an intervention.

The minimum share in this model is municipal, and it answers a question about which cities take part. Within a city, the mean can hide exposure near industrial sites or major roads, and the published population field gives no weight to income, ethnicity, age or pre-existing disadvantage. An operational programme would need neighbourhood vulnerability data, a chance for affected communities to review the objective and its weights, and independent scrutiny of the local reduction caps, since the feasibility calculation rests on them.

5.3 Limitations

Of the 333 Spanish records for 2021, only 40 pass every selection criterion, most of the losses being records without a population figure or with too little monitoring coverage, and even among the 40 the coverage and the area a single monitor represents vary. We kept Salamanca's unusually large entry of 1.796 million so that the calculation can be reproduced from the published workbook, but its municipal population count still needs verification. Before any operational use, populations would have to be checked independently and population and monitoring boundaries matched; the results apply to these 40 records only and support no statement about national levels, rankings of cities by pollution or differences in personal exposure.

The model is static and deterministic. Four inputs set its feasible allocations (the 30% local reduction cap, the reference concentration, the resource coefficients and the minimum shares), and to know whether a city could actually deliver its assigned reduction we would need a monetary budget, a timescale for intervention and a causal estimate of the concentration response, plus evidence on political bargaining,

rebound effects and atmospheric spillovers; any statement about health would need an epidemiological dose–response relationship as well. Within these limits, what the allocation rule does is keep mean progress, the distribution of residual excess and the feasibility of a common concentration ceiling apart.

6. CONCLUSION

Fixing the aggregate reduction does not fix who is left with the excess. For the 40-city Spanish portfolio the clipped solution is unique; at half of assumed total capacity it leaves the same weighted mean concentration as proportional allocation (8.404 $\mu g/m^3$), lowers squared excess loss by 9.44% and brings Murcia, Cordoba and Vigo below 10 $\mu g/m^3$, which proportional allocation does not. A 25% minimum share for every eligible city costs 3.01% in squared excess loss. A common concentration ceiling of 10 $\mu g/m^3$ is affordable at 17.12% of assumed total capacity, whereas 9 $\mu g/m^3$ is out of reach at any budget because of Murcia's cap. We suggest that authorities report between-city excess variance alongside mean concentration and check local feasibility before committing to a common ceiling, although independent estimates of intervention costs, atmospheric responses and neighbourhood vulnerability are still needed before an allocation of this kind is put into operation.

REFERENCES

[1] T. Goforth and D. Nock, "Air pollution disparities and equality assessments of US national decarbonization strategies," *Nature Communications*, vol. 13, no. 1, p. 7488, 2022.

[2] P. Picciano, M. Qiu, S. D. Eastham, M. Yuan, J. Reilly, and N. E. Selin, "Air quality related equity implications of U.S. decarbonization policy," *Nature Communications*, vol. 14, no. 1, p. 5543, 2023.

[3] J. Rentschler and N. Leonova, "Global air pollution exposure and poverty," *Nature Communications*, vol. 14, no. 1, p. 4432, 2023.

[4] A. P. Behrer and S. Heft-Neal, "Higher air pollution in wealthy districts of most low- and middle-income countries," *Nature Sustainability*, vol. 7, no. 2, pp. 203–212, 2024.

[5] L. van den Brekel, V. Lenters, J. D. Mackenbach, G. Hoek, A. Wagtendonk, J. Lakerveld, D. E. Grobbee, and I. Vaartjes, "Ethnic and socioeconomic inequalities in air pollution exposure: A cross-sectional analysis of nationwide individual-level data from the netherlands," *The Lancet Planetary Health*, vol. 8, no. 1, pp. e18–e29, 2024.

[6] E. Pisoni, P. Thunis, A. De Meij, and B. Bessagnet, "Assessing the impact of local policies on PM2.5 concentration levels: Application to 10 european cities," *Sustainability*, vol. 14, no. 11, p. 6384, 2022.

[7] Á. Gómez-Losada and J. C. M. Pires, "Air quality assessment during the low emission zone implementation in madrid (spain)," *Urban Climate*, vol. 55, p. 101995, 2024.

[8] B. M. Gentry, A. L. Robinson, and P. J. Adams, "EASIUR-HR: A model to evaluate exposure inequality caused by ground-level sources of primary fine particulate matter," *Environmental Science & Technology*, vol. 57, no. 9, pp. 3817–3824, 2023.

[9] Y. Wang, J. S. Apte, J. D. Hill, C. E. Ivey, R. F. Patterson, A. L. Robinson, C. W. Tessum, and J. D. Marshall, "Location-specific strategies for eliminating US national racial-ethnic PM2.5 exposure inequality," *Proceedings of the National Academy of Sciences*, vol. 119, no. 44, p. e2205548119, 2022.

[10] World Health Organization, "WHO ambient air quality database january 2024 edition v6.1," 2024.

[11] K. Shairsingh, G. Ruggeri, M. Krzyzanowski, P. Mudu, M. Malkawi, J. Castillo, A. Soares da Silva, M. Saluja, K. Cervantes-Martínez, J. Mothe, and S. Gumy, "WHO air quality database: Relevance, history and future developments," *Bulletin of the World Health Organization*, vol. 101, no. 12, pp. 800–807, 2023.

[12] M. M. Kelp, T. C. Fargiano, S. Lin, T. Liu, J. R. Turner, J. N. Kutz, and L. J. Mickley, "Data-driven placement of PM2.5 air quality sensors in the united states: An approach to target urban environmental injustice," *GeoHealth*, vol. 7, no. 9, p. e2023GH000834, 2023.

[13] E. M. Considine, D. Braun, L. Kamareddine, R. C. Nethery, and P. deSouza, "Investigating use of low-cost sensors to increase accuracy and equity of real-time air quality information," *Environmental Science & Technology*, vol. 57, no. 3, pp. 1391–1402, 2023.

[14] J. Wei, Z. Li, A. Lyapustin, J. Wang, O. Dubovik, J. Schwartz, L. Sun, C. Li, S. Liu, and T. Zhu, "First close insight into global daily gapless 1-km PM2.5 pollution, variability, and health impact," *Nature Communications*, vol. 14, no. 1, p. 8349, 2023.