Delhi Air Quality 2025: Sources, Methods, and Findings

An Independent Analysis

Prepared by the office of Ajay Maken

13 June 2026

Table of Contents

Executive summary

In the last days of 2025 the Government of India announced, through the Press Information Bureau and the Annual Report of the Commission for Air Quality Management (CAQM), that Delhi's air had measurably improved. The centrepiece was a fall in fine-particle pollution (PM2.5) to 96 micrograms per cubic metre — described as the lowest reading other than the 2020 lockdown year.

This report is an independent examination of that announcement. We checked whether the government's published numbers are correct, and then asked the harder question: *do they mean what the announcement says they mean?* To do that we separated the part of any year-to-year change that is caused by the weather from the part caused by real changes in what Delhi emits into its air. The two are easy to confuse — a windy, rainy year disperses pollution and looks clean even if emissions never fell.

The government's published numbers are genuine. We reproduced every one of them from the raw monitoring data. But the headline improvement is, for the most part, a weather story rather than a policy story.

The single most important finding is this: once weather is removed, 2025's change in PM2.5 is not statistically distinguishable from zero — there is no demonstrable real improvement. The government's own observed fall on a fixed set of monitoring stations was 7.4 percent, but that figure is dominated by the weather of the comparison years. After de-weathering, two independent models place the central estimate of the emissions-driven change within a band of roughly plus-or-minus 2 to 3 points of zero — that is, indistinguishable from no change at all. The 7.4 percent the government published does not survive as a real improvement once the air is corrected for the weather it was measured in.

The supporting findings are consistent and, in several cases, uncomfortable:

- The government's headline annual figure rests on one freak-stagnant month: January 2024 — the year 2025 is measured against — was the stillest January in at least eight years (2018-2025).

- The July-to-December stretch that contains the entire Graded Response Action Plan (GRAP) season shows no improvement at all on any basis.
- December 2025 was the most polluted December in our records on both the raw and the weather-adjusted measure; favourable-versus-unfavourable winds explain only about half of that spike.
- Nitrogen dioxide and ground-level ozone both got genuinely worse.
- Across five post-monsoon seasons — the season GRAP exists for — there is no real improvement over 2021.
- Delhi remained the world's most polluted capital for the eighth year running, at roughly 19 to 20 times the World Health Organization's recommended limit.

Throughout, we distinguish clearly between what is directly measured and what is estimated by a model, and we attach honest uncertainty to every estimate.

1. Why this analysis, and what we examined

Two official documents form the basis of the government's improvement message, and they are the documents we examined.

The first is the PIB press release "2025 — Air Quality Perspective in Delhi" (Release ID 2210305, 31 December 2025). It carries the headline numbers, a month-by-month table of average air quality, and a specific explanation for a bad December.

The second is the CAQM Annual Report 2024-25. A full reading of all hundred pages confirms that its only quantitative air-quality content is a single comparative table on page 61 (printed page 50).

Stated neutrally, the specific claims we tested are:

- That "continued, concerted and persistent efforts… have further helped to improve the general air quality" (PIB, page 1).
- That the 2025 annual average Air Quality Index (AQI) of 201 was "the lowest in last 7 years… barring the COVID year 2020" (PIB, page 4).
- That PM2.5 of 96 µg/m³ and PM10 of 197 µg/m³ were the "lowest ever… barring the Covid affected year 2020" (PIB, page 3).
- That Delhi recorded 79 days of "Good" to "Satisfactory" air, the most other than 2020, and 8 "Severe" or worse days, the second-fewest since 2018 (PIB, pages 1 and 3).
- That December 2025's high monthly AQI of 351 was caused by "absolutely still / very low speed wind conditions and unfavourable meteorological conditions" (PIB, page 4).
- That CAQM's efforts "helped to increase the number of good to moderate days" (CAQM AR, page 61).
- A forward promise that air quality "will see further gradual but marked improvement, year on year" (PIB, page 7).

A point of method that matters before any finding is presented: we reproduced every number the government published directly from raw station data — the annual AQI of 201, PM2.5 of 96, PM10 of 197, 79 Good-to-Satisfactory days, and 8 Severe-plus days all check out. There is no dispute about the data. The disagreement is only about what the data means once weather is taken into account, and about whether the year fits the multi-year story the government tells.

The government's own annual series (its press-release table, reproduced below) makes the central point on its own: a single lower observed reading after four flat ones, against a promise of improvement "year on year."

Year	2018	2019	2020*	2021	2022	2023	2024	2025
Annual average AQI	225	215	185	209	209	204	209	**201**
PM2.5 (µg/m³)	113	107	94	105	98	100	104	**96**
PM10 (µg/m³)	241	217	180	211	211	205	212	**197**
Good–Satisfactory days	53	61	100	73	68	61	66	**79**
Severe+ days	20	24	15	24	6	15	17	**8**

* COVID lockdown year. The press release scopes 2025 PM2.5 as "till 29 December"; we verified this truncation is immaterial — worth about 1.6 AQI points in December and a negligible amount annually.

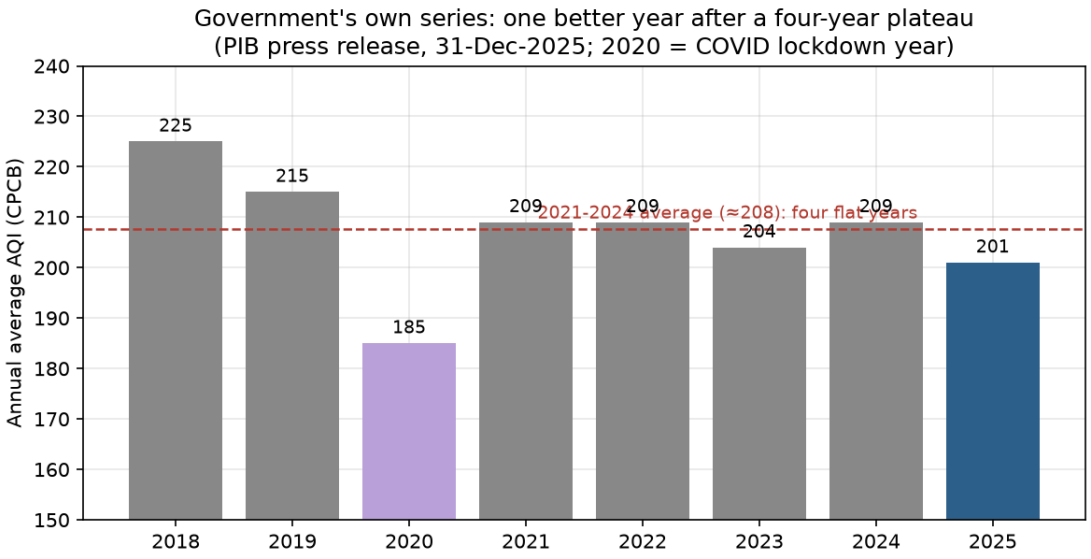

Delhi's annual average AQI, 2018-2025. The AQI sat at 209, 209, 204, 209 from 2021 to 2024 before falling to 201 in 2025.

Read across the AQI row: 209, 209, 204, 209, then 201. Read the PM2.5 row from 2022: 98, 100, 104 — rising for three straight years. The promise of "marked improvement, year on year" is contradicted by the government's own table, which shows four statistically flat years followed by a single lower observed reading — and even that reading, as the rest of this report shows, does not survive once weather is removed.

2. Our data sources

Everything in this report is built from public, reproducible data. Nothing relies on a private feed or a proprietary dataset.

Source	What it provides	Provenance	Period
CPCB monitoring stations (via CREA)	Hourly ground readings of seven pollutants from Delhi's official network	Central Pollution Control Board stations, downloaded through the Centre for Research on Energy and Clean Air	2021 – 12 June 2026
ERA5 / Open-Meteo weather	Hourly wind speed and direction, rainfall, boundary-layer height (how high the air mixes), temperature	Copernicus ERA5 reanalysis, via Open-Meteo and Google Earth Engine	2021 – 12 June 2026
Sentinel-5P (TROPOMI) satellite	Column measurements of nitrogen dioxide and formaldehyde over Delhi from space	European Space Agency Copernicus programme, via	2019 – 2025

Source	What it provides	Provenance	Period
		Google Earth Engine	
MODIS active-fire detections	Daily counts of stubble-fire pixels over Punjab and Haryana	NASA MODIS (Terra + Aqua), via Google Earth Engine	October–November, 2024 and 2025
Comparator-city PM2.5	Daily city averages for Lucknow, Kanpur, Jaipur, Patna	Same CPCB-via-CREA public process as Delhi	2024 – 2025
Government documents	The claims under examination	PIB Release 2210305; CAQM Annual Report 2024-25 (page 61)	2025

The scale of the ground data is substantial: roughly 1.7 gigabytes of hourly readings across six years, and about 47,700 hours of matched weather data. Section 5 sets out the full count.

3. How we separate real change from weather ("de-weathering")

The hardest problem in judging an air-quality year is that the weather changes the numbers as much as emissions do. A windy week scatters pollution and the monitors read low; a still, cold week traps it and they read high — even if the city emitted exactly the same amount of smoke and dust on both. To know whether Delhi's air genuinely got cleaner, you have to strip the weather out. The standard scientific method for doing this is called "de-weathering."

> **For the general reader.** Imagine two photographs of the same street taken a year apart. In one, a strong wind is blowing the haze away; in the other, the air is dead still. The street looks cleaner in the first photo, but that tells you about the wind, not about how much pollution was produced. De-weathering is a way of asking the data a fairer question: *what would each day have looked like under ordinary, average weather?* We feed a computer model years of records — every day's pollution alongside that day's wind, rain, temperature, and how high the air was mixing — until it learns how weather pushes pollution up and down. Then we ask it to replay every day under typical weather. What is left over is the part of the change that the weather cannot explain — the part that reflects real changes in what the city emits.
>
> **For the technical reader.** We trained two independent machine-learning engines per pollutant and per station: a random forest (the R package `ranger`) and a boosted-regression-tree model (the R `gbm` / `deweather` stack). Both were fitted on a fixed-station panel using *only* meteorological and meteorological-lag features — wind speed and direction, precipitation, boundary-layer height, temperature, and their lags — with no pollutant lags included, so that nothing about the pollution history can leak into the prediction and contaminate the attribution. The weather-adjusted series is produced by climatological normalisation: the meteorology is repeatedly resampled and the model

predictions averaged, which holds weather at its climatological distribution and isolates the emissions-driven signal. For particulate matter we anchor on the random-forest and legacy models; the gas figures use the adjudicated all-pollutant series. The city AQI is computed from per-pollutant sub-indices on the standard Central Pollution Control Board averaging windows — 24-hour means for PM2.5, PM10, sulphur dioxide and nitrogen dioxide; rolling 8-hour maxima for carbon monoxide and ozone — with the city index taken as the worst sub-index, and a particulate-only fallback index carried alongside.

Two principles govern how we report every result in Section 6:

- **Observed versus weather-adjusted are different quantities, and we label them as such every time.** An observed figure is direct measurement. A weather-adjusted ("real") figure is a model's estimate of the emissions-driven part. We never let one pass for the other.
- **Weather-adjusted figures carry honest uncertainty.** Any single annual weather-adjusted change is reported with a band of roughly plus-or-minus 2 to 3 percentage points, and any single-month weather-adjusted value with a band of roughly 5 to 16 AQI points. For that reason single-month weather-adjusted numbers are not published as headline claims, and every weather-adjusted point estimate is labelled "estimate" or given as a range.

4. The scale of the computation

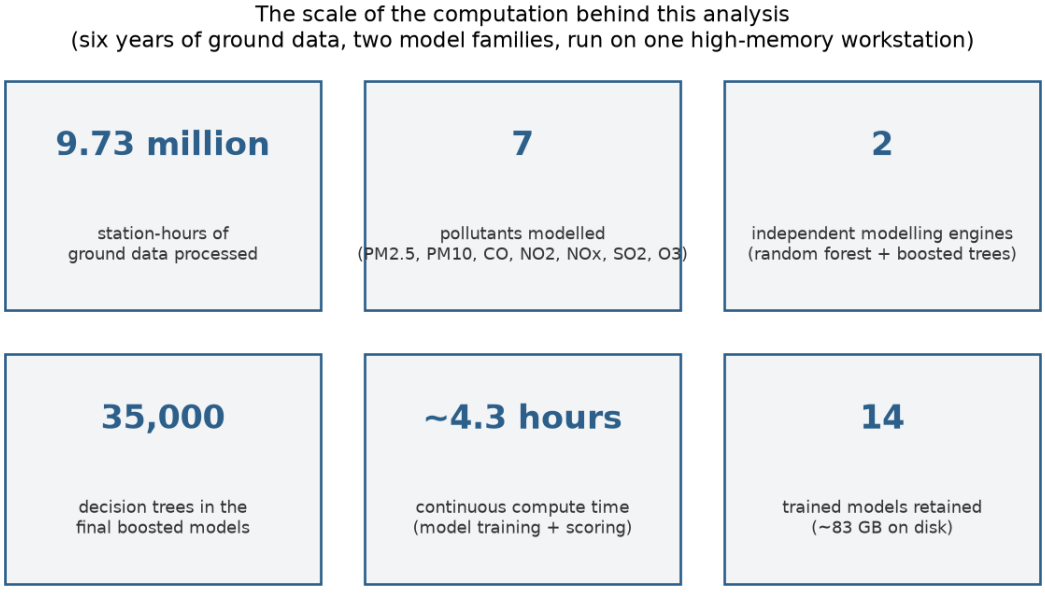

The scale of the computation: nearly ten million station-hours of ground data, seven pollutants, two independent model families totalling 35,000 decision trees, run over about 4.3 hours on one high-memory workstation.

The Boss asked a fair question of any analysis that contradicts an official one: how rigorous is it? The honest answer is in the size of the work. The de-weathering was run on a single high-memory workstation (an Intel Core i9-Ultra 285K with 256 GB of RAM). The work is computation on the processor — the random-forest and boosted-tree fits run across the machine's cores, not on a graphics card — and it took several hours of continuous running.

Quantity	Figure
Pollutants modelled	7 (PM2.5, PM10, CO, NO2, NOx, SO2, O3)
Station-hours of ground data processed	~9.73 million (9,734,329)
Monitoring stations (per pollutant)	39 to 55
Years of ground data	6 (2021 through 12 June 2026)
Hours of matched weather data	47,736
Modelling engines	2 (random forest + boosted regression trees)
Decision trees in the final boosted models	35,000 (7 pollutants × 5,000)
Final-model decision trees, both engines	40,600
Model-fits performed	182 (14 final models + 168 cross-validation fits)
De-weathering predictions per engine pass	~9.7 million
Predictions across training, validation and scoring	tens of millions (~20–40 million)
Continuous compute time	~4.3 hours of model training and scoring
Trained models retained on disk	14 (≈83 GB random forest + ≈106 MB boosted)
Satellite composites (ozone-regime analysis)	96 (24 months × 2 products × 2 regions)

In plain terms: nearly ten million hours of station readings, every one of them matched to that hour's weather, were used to teach two separate families of models — 35,000 decision trees among them — what weather does to each pollutant. The models then re-scored that entire record under typical weather. This is not a back-of-envelope comparison of two annual averages; it is the kind of analysis that lets us say, with a stated margin of error, how much of a change the weather can account for.

5. Findings

(a) The headline: the PM2.5 "improvement" does not survive de-weathering

PM2.5 — the fine particles small enough to enter the bloodstream and the pollutant the government led with — is where the gap between appearance and reality is widest.

On a fixed panel of 38 monitoring stations the government's observed annual average fell from 106.1 to 98.3 µg/m³, a decline of 7.4 percent. That figure is a direct measurement, and we do not dispute it. But once weather is removed, the 7.4 percent does not survive: the two independent engines agree closely, and the weather-adjusted change — about −2.1 percent (boosted trees) to −2.8 percent (random forest) — sits inside the model's own uncertainty band. Set against the plus-or-minus 2-to-3-point uncertainty described in Section 3, the central estimate is statistically indistinguishable from zero. **There is no demonstrable real improvement in PM2.5: once weather is removed, the change cannot be told apart from no change at all.** The government's 7.4 percent is a measurement of the weather it was taken in, not of a cleaner city.

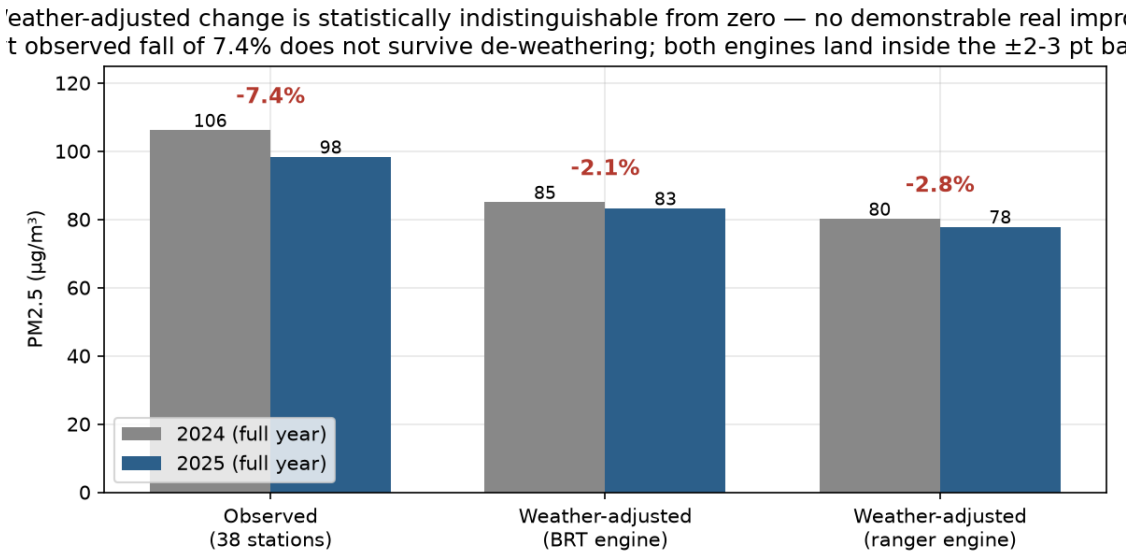

Delhi PM2.5 across the full year: the government's observed decline (7.4%) does not survive weather adjustment — the weather-adjusted change is statistically indistinguishable from zero.

The reason is visible in the raw weather record, with no model needed. **January 2024 — the month 2025 is measured against — was the stillest January in at least eight years (2018-2025).** Its average wind was 6.58 km/h, the lowest of the eight years, against 8 to 9 km/h in every other January (the others span 8.24 to 9.39 km/h); it had seventeen near-calm days — days whose mean wind fell below 6 km/h — against just three to nine in other years, and a single day of rain. January 2025, by contrast, was an ordinary month. The government's headline annual figure rests on this one freak-stagnant month in its base year: an annual average is dragged by a single trapped, windless January, exactly the kind of distortion an annual summary hides and de-weathering is built to catch. (Tellingly, the annual *median* wind for the two years is nearly identical; the January extreme is a tail of the distribution that the yearly summary conceals.)

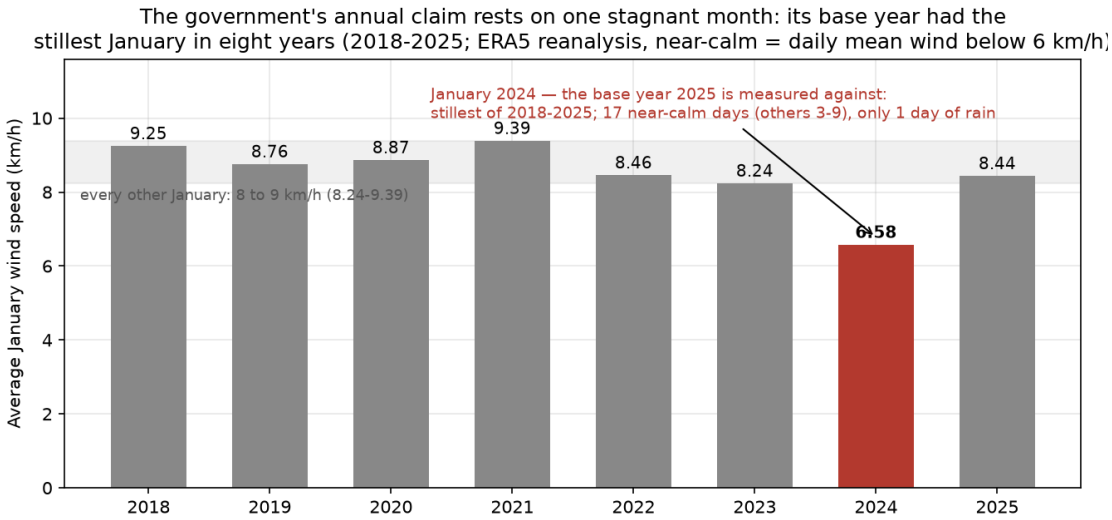

Average January wind speed by year, 2018-2025. January 2024 — the base year 2025 is measured against — was the stillest January of the eight years, with far more near-calm days and almost no rain.

(b) The season that matters: post-monsoon and GRAP

If policy were driving the improvement, it should show up most clearly in the second half of the year, when Delhi's pollution peaks and when the Graded Response Action Plan and CAQM's directions are in force. It does not.

July to December — the half of the year that contains the entire GRAP season — shows **no improvement on any basis**: observed +1.2 percent, weather-adjusted +0.8 to +1.0 percent. October to December was in fact slightly worse than 2024 (+2.4 percent observed, +2.3 percent weather-adjusted). The whole of the government's observed annual decline is concentrated in January to June 2025 — the half of the year that, for data-coverage reasons, cannot be fully de-weathered, and the half dominated by the freak-stagnant January-2024 base month.

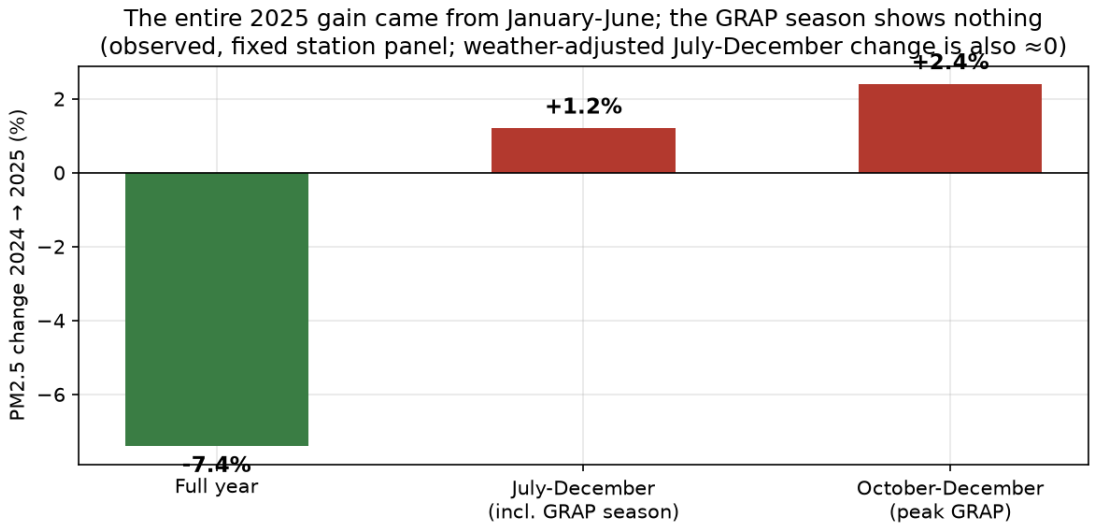

Where the government's observed decline sits: January-June versus July-December. It is entirely in the first half — the season GRAP exists for shows nothing.

Looking specifically at the post-monsoon window across five seasons — the season the entire GRAP apparatus exists for — there is no real improvement over 2021. Our best estimate points to a deterioration of roughly 20 AQI points (an estimate, not a statistically firm result).

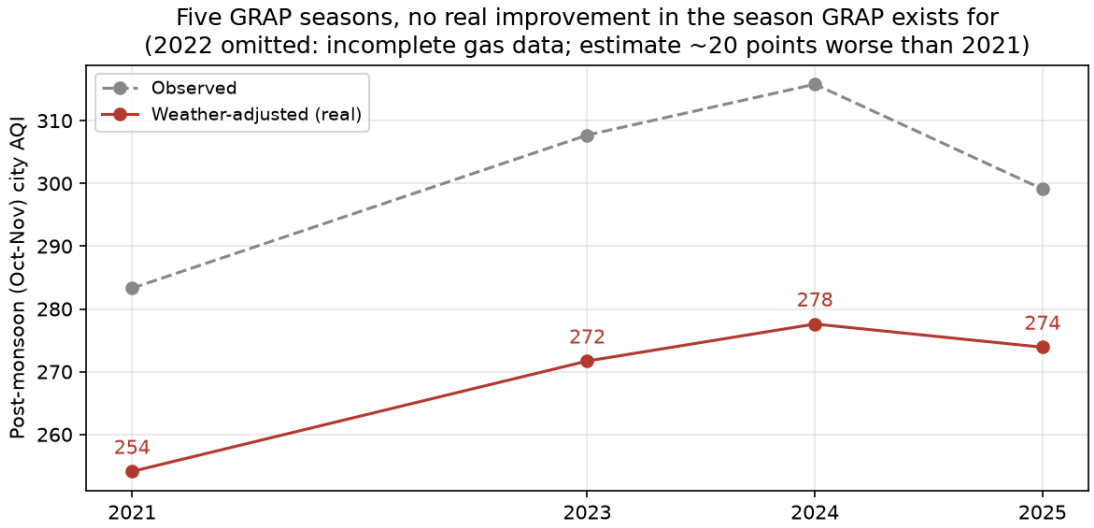

The post-monsoon/GRAP season across five years: no real improvement over 2021.

(c) December 2025: the claim that does not hold

The press release attributes December 2025's monthly AQI of 351 to "absolutely still / very low speed wind conditions." This is the one explanation that does not survive examination — though it fails by half rather than completely.

	Dec 2022	Dec 2023	Dec 2024	Dec 2025
Observed city AQI	329.0	359.2	306.5	**362.1**
Weather-adjusted	334.6	332.0	324.2	**348.2**

December 2025 was, together with December 2023, the most polluted December since at least 2021 on the observed basis, and the highest on the weather-adjusted basis outright, by a clear margin (348 against a next-highest of 335). Of the roughly 56-point observed jump over December 2024, about **24 points survive weather adjustment**, with high statistical confidence on both model generations. In plain terms: still winds explain only about half of the December the government blamed entirely on the weather. The other half was a real worsening — inside the very year the announcement celebrates.

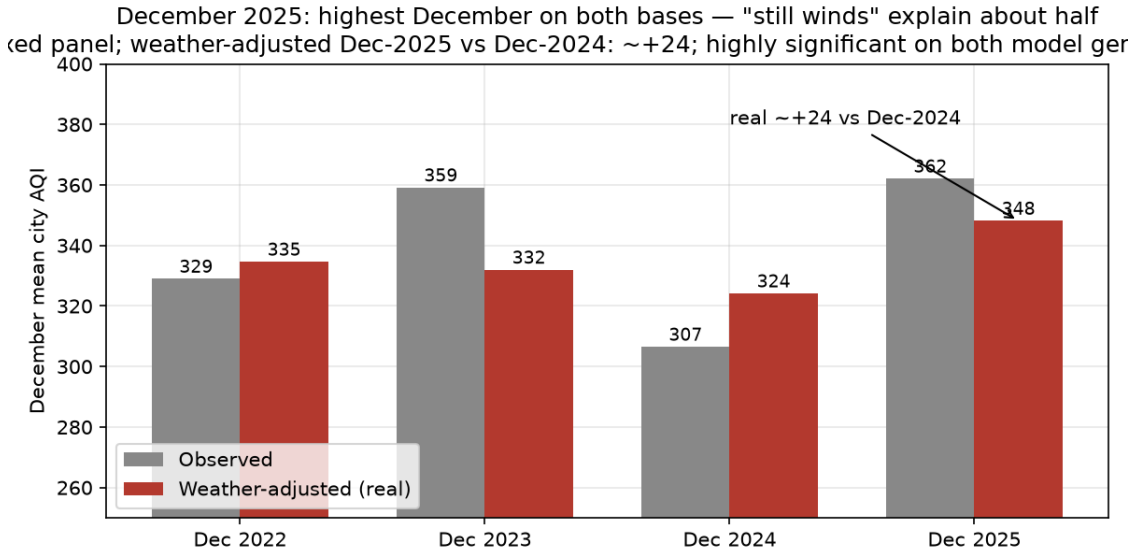

December across four years, on both the observed and weather-adjusted bases.

(d) The gases: nitrogen dioxide and ozone got worse

The pollutants the announcement did not lead with tell their own story, and it is not a story of improvement.

Pollutant	2024	2025	Observed (govt basis)	Weather-adjusted	Reading
PM2.5 (µg/m³)	106.1	98.3	−7.4%	within ±2-3 pts of zero (estimate)	No demonstrable real change; weather-driven

Pollutant	2024	2025	Observed (govt basis)	Weather-adjusted	Reading
PM10 (µg/m³)	213.3	199.7	−6.4%	within ±1-2 pts of zero (estimate)	No demonstrable real change; weather-driven (observed-only govt claim)
NO2 (µg/m³)	42.0	44.9	**+7.1%**	**+7.7%**	Worsened; none of it is weather
Ozone, 8h max (µg/m³)	60.8	67.7	**+11.2%**	**+13.6%**	Worsened; weather masked part of it
SO2 (µg/m³)	15.1	14.8	−2.1%	within ±2-3 pts of zero (estimate)	No demonstrable real change
CO, 8h max (µg/m³)	1808	1813	+0.3%	within ±3 pts of zero (estimate)	Essentially flat

Nitrogen dioxide rose 7.1 percent on the observed basis and 7.7 percent once weather is removed — meaning the rise is entirely real, with none of it explained by favourable or unfavourable weather. Ground-level ozone rose 11.2 percent observed and an estimated 13.6 percent in real terms — the largest worsening of any pollutant, and one the weather actually masked. Sulphur dioxide and carbon monoxide both sat within their uncertainty bands of zero — no demonstrable real change either way (carbon monoxide has a small monitoring panel and is never headlined here).

(e) PM10: an observed government number that does not survive de-weathering

Coarse particles (PM10) deserve a clear statement because the government leans on them. The observed annual average fell 6.4 percent (213.3 to 199.7 µg/m³), and on the observed fixed panel 2025 is the lowest PM10 reading of 2021-2025 — so the government's "lowest ever, barring 2020" PM10 claim is, as a raw measurement, reproduced in our own data. But that is an observed-only number: once weather is removed, the change collapses to within one or two points of zero, statistically indistinguishable from no change. The low 2025 reading reflects the weather PM10 was measured in, not a real reduction in what Delhi puts into the air. The government's record is a record of the weather, not of the city getting cleaner.

(f) The ozone question: what the satellite shows

Ozone's 11 percent rise invites a chemistry defence. In some cities, cutting vehicle emissions can mechanically push ozone *up* for a period, so a rising ozone reading can in principle be a side-effect of a successful clean-up rather than a failure. Whether that defence applies to Delhi depends on the chemical "regime" the city's air is in — and that regime can be read from space.

The European Sentinel-5P satellite measures the ratio of formaldehyde to nitrogen dioxide in the air column (the "FNR"). The value of that ratio indicates the regime: a low ratio means ozone formation is limited by available nitrogen oxides (the regime where cutting NOx can raise ozone); a high ratio means it is limited by volatile organic compounds (where the titration defence does not apply). We computed this ratio over the Delhi core for every month of 2024 and 2025.

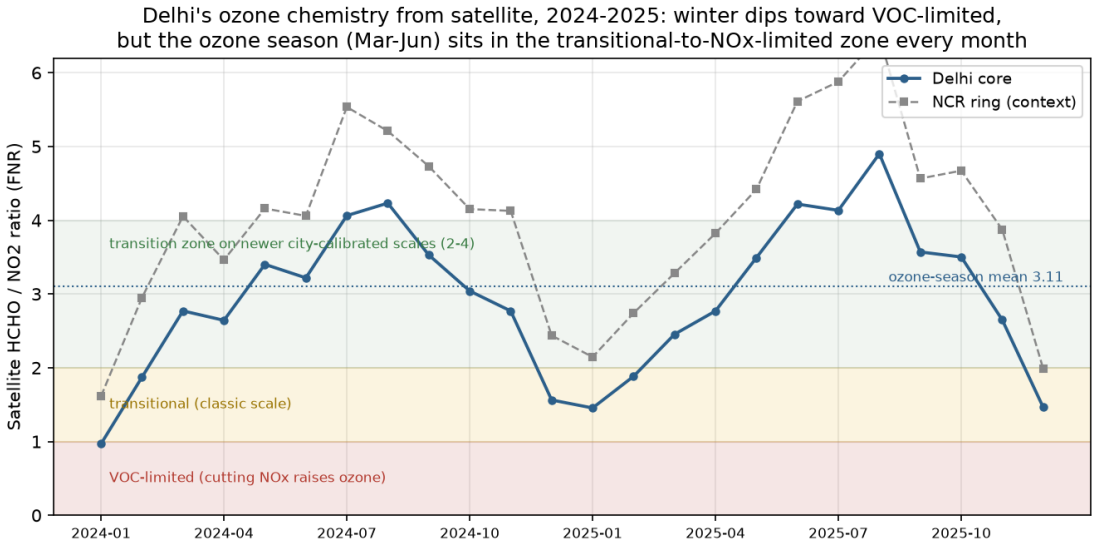

The satellite formaldehyde-to-NO2 ratio over Delhi, by month, against the regime thresholds.

The finding is two-sided. In winter, Delhi's ratio dips low (a seasonal value of 1.46, with January 2024 the only month below 1.0), so the "cutting emissions raises ozone" mechanism is plausible for the winter months. But in the ozone season itself — March to June, when ozone actually crosses health limits — the ratio averages 3.11 and never falls below 2.45 in any month, which reads transitional-to-NOx-limited on every published scale, including one calibrated for India. In plain terms: the chemistry defence cannot excuse the ozone rise in the season where ozone actually matters. The data also does not prove the opposite regime, so the formal classification is "transitional" and no one may claim the question is fully settled.

(g) Did the policies work? Three independent attribution tests

The deepest question is whether any specific policy can be credited for a real change in Delhi's air. We ran three pre-registered tests — designs locked in before the results were seen — and none of them credits a policy with a real reduction in the pollution Delhiites actually breathe.

Test 1 — a break at the policy dates. We looked for a sudden improvement in the de-weathered daily series at each major GRAP date (a regression-discontinuity-in-time design). No improvement break is detectable at any policy date. The one signal the test flagged (Stage-III, 11 November 2025, +39 AQI points) has the wrong sign for policy credit and is an artefact of how GRAP works: the stage was invoked *because* AQI had crossed 425 that morning, so the rule chases the pollution rather than the other way round. Across 150 placebo dates in earlier years, 21.3 percent produced a larger swing — so the signal sits inside ordinary noise. This is a low-powered test: it can rule out detectable breaks, not all effects.

Test 2 — cities without Delhi's policies. We measured four Indo-Gangetic-plain cities with none of Delhi's machinery — no GRAP, no CAQM directions, no electric-bus fleet, no end-of-life-vehicle ban — on the identical public process. Their observed 2025 readings moved the same way Delhi's did: Lucknow −4.9 percent (58.8 to 55.9), Kanpur −5.5 percent (46.5 to 43.9), Jaipur −6.8 percent (56.3 to 52.5). Delhi, run through the same process, came to −7.8 percent (104.5 to 96.4) — barely 1 to 3 points off a regional pattern shared by cities that did none of what Delhi did. Once that shared regional weather movement is accounted for, nothing Delhi-specific is left to credit: the gap is the same size as Delhi's own weather-adjusted change, which is itself indistinguishable from zero. Two unrelated methods land on the same answer — no demonstrable real improvement unique to Delhi. (Patna fell 26.8 percent, but it is an outlier on a multi-year slide where a monitoring-network change cannot be excluded; no conclusion here rests on it.)

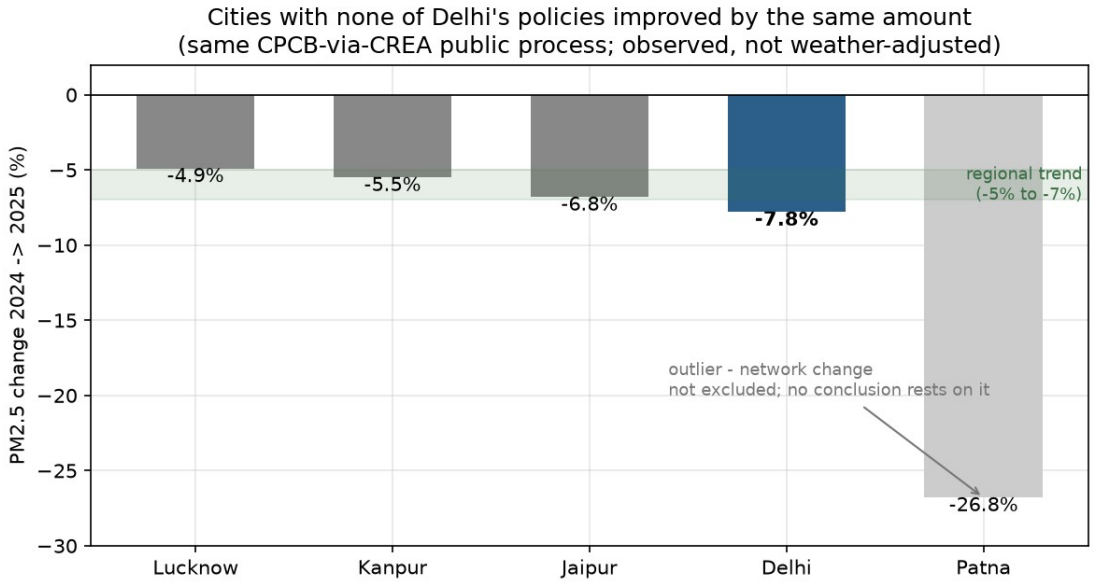

The 2024-to-2025 PM2.5 change in four cities with none of Delhi's policies, against Delhi run through the identical process. Delhi sits barely beyond the regional cluster; Patna is flagged as an outlier.

Test 3 — stubble fires. Counted stubble-fire detections over Punjab and Haryana fell 17.4 percent (30,976 to 25,577 MODIS pixels) and the 2025 season ran about 6 to 9 days earlier. But Delhi's de-weathered October-November air was essentially flat (277.6 to 273.9), and the

statistical link between fires and Delhi's air did not weaken by the pre-registered margin. Fire *counts* are a detection index, not a measure of smoke emitted; the satellite-measured burnt *area* was flat (19.17 versus 19.14 lakh hectares). The MODIS −17.4 percent must never be confused with the much larger figures quoted in the news (a VIIRS-sensor "−70 percent" or CAQM's "−85 percent") — different sensors, different baselines, not comparable.

The only Delhi-specific satellite signal — and it sits inside a worsening family. Satellite NO2 column readings show the Delhi core declining faster than the surrounding NCR ring from 2024 to 2025 (core −7.7 percent versus ring −2.6 percent), and the core has fallen every year since the 2021 peak. This is *consistent with* — not proof of — Delhi-specific traffic, electrification and NOx measures. But it concerns the nitrogen-oxide column only, the very family in which the surface pollution Delhiites actually breathe got worse: ground-level NO2 rose 7.7 percent and ozone rose an estimated 13.6 percent in the same period. A faster-falling NOx column above a city whose surface NO2 and ozone both worsened is not an air-quality win — it is a trade-off inside a family that deteriorated where it counts.

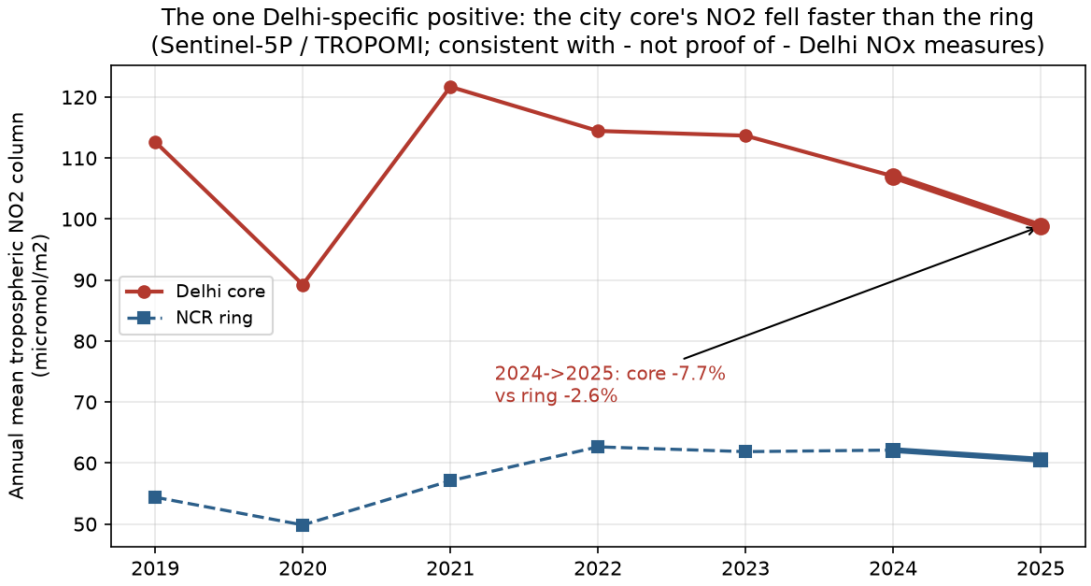

Annual satellite NO2 column over the Delhi core versus the surrounding NCR ring. From 2024 to 2025 the core fell faster than the ring (−7.7 versus −2.6 percent).

The overall attribution verdict: on these designs, no specific policy can be credited with a real reduction in the particulate pollution Delhiites breathe; the one Delhi-specific satellite signal sits in the nitrogen-oxide column — the very family where surface NO2 and ozone got worse — so it is a trade-off, not a cleaner-air result.

(h) The human cost and the global context

The numbers above are abstract until set against a body and a map.

At 96 µg/m³, Delhi's 2025 PM2.5 is about **2.4 times India's own annual standard** (40 µg/m³) and about **19 to 20 times the World Health Organization guideline** (5

µg/m³). The Air Quality Life Index, a standard public-health metric, implies that sustained exposure at these levels costs a Delhi resident roughly **10 years of life expectancy**. Delhi remained the **world's most polluted capital for the eighth consecutive year** in 2025. For scale, winter PM2.5 in Beijing runs around 29 µg/m³; in Delhi it runs around 163. With no demonstrable real change once weather is removed, the city is sitting still on a health emergency — nineteen to twenty times above the level the WHO considers safe to breathe.

6. What our analysis can and cannot say

Honest science states its own limits as plainly as its findings.

- **Observed and weather-adjusted are distinct quantities.** Observed findings stand on direct measurement. Weather-adjusted findings are model estimates of the emissions-driven part, and we label them as estimates throughout.
- **Weather-adjusted figures carry a model-uncertainty band.** Any single annual weather-adjusted change carries roughly plus-or-minus 2 to 3 percentage points; any single-month weather-adjusted value roughly 5 to 16 AQI points. For that reason single-month weather-adjusted claims are not published, and point estimates always carry the word "estimate" or a range.
- **The comparator-city analysis is observed-only, not weather-adjusted.** City monitoring networks differ, and the underlying city-average aggregation has not been independently audited. The Patna figure carries a network-change flag and supports no conclusion here.
- **Fire counts are a detection index, not a measure of emissions.** A change in counted fires is not a change in smoke; MODIS counts are not comparable with VIIRS or with the government's own figures.
- **The policy-break test is low-powered.** Delhi's de-weathered daily series is highly persistent, which widens the true uncertainty beyond what a naive calculation shows. The test can flag candidates and rule out detectable breaks; it cannot settle attribution by itself.
- **Attribution is framed as "consistent with, not proof of."** No specific policy can yet be credited or blamed for a real change in Delhi's particulate air on these designs.
- **The ozone regime is formally "transitional."** Neither the NOx-limited nor the titration interpretation may be cited as settled; the titration defence applies to winter only.
- **There are data-coverage limits inside the period.** The 2021 winter window is partial, and 2022 lacks the full gas record needed for an all-pollutant AQI; we make no all-pollutant AQI claims for 2022's seasons, and carbon monoxide is never headlined.

What we *can* say, with confidence, is the through-line: the published numbers are genuine; the government's headline PM2.5 fall does not survive weather adjustment — once weather is removed there is no demonstrable real improvement, the change being statistically indistinguishable from zero; the GRAP season shows no improvement; December was the most polluted on record on both bases; nitrogen dioxide and ozone worsened; and across five post-monsoon seasons there is no real gain. Each of these survived independent recomputation.

7. Reproducibility statement

This analysis is built entirely on public data and standard methods, and we are happy to share specifics on request.

- **Ground data** are CPCB station readings downloaded through CREA, the public source any researcher can access.
- **Weather data** are the Copernicus ERA5 reanalysis, obtained through Open-Meteo and Google Earth Engine.
- **Satellite data** are European Space Agency Sentinel-5P (TROPOMI) and NASA MODIS products, both freely available through Google Earth Engine.
- **Methods** are standard and widely published: random-forest and boosted-regression-tree de-weathering with climatological normalisation; CPCB sub-index AQI computation on the official averaging windows; regression-discontinuity-in-time and comparator-city designs for attribution; the formaldehyde-to-NO2 satellite ratio for ozone-regime classification.
- **Quality control** follows the conventional rule that a station-day needs at least 16 valid hourly readings out of 24 to count toward any average.
- **Verification.** Every finding in this report was independently recomputed from source data and adjudicated before publication; the attribution claims cleared a four-of-four adversarial check, and the satellite ozone-regime values were reproduced to eight significant figures from the archive. The figures presented are the verified set.

The conclusion stands on the data, not on trust: anyone with the same public sources and the standard methods can reproduce it.

Prepared by the office of Ajay Maken — 13 June 2026. Sources: PIB Release 2210305 (31 December 2025); CAQM Annual Report 2024-25; CPCB station data 2021-2026 (via CREA); ERA5 (Copernicus, via Open-Meteo and Google Earth Engine); Sentinel-5P and MODIS (via Google Earth Engine).