DELHI AIR QUALITY · RETROSPECTIVE

Weather's share vs the city's own pollution

2025 compared with 2024 — the full year, weather removed

This retrospective applies the same weather-vs-own-pollution analysis as the daily report, but to the whole of **2025** compared with **2024** — separating how much of the year-on-year change in Delhi's air was the **weather** from how much was the **city's own doing**.

What the two lines are. Each line is 2025's level minus the average of the same ±7 days in 2024, in µg/m³. The **grey line** is the **raw observed** change — what the monitors measured, weather and all. The **red line** is the same change **after our model removes the weather**, leaving what the city's own emissions actually did.

Why it matters. A calm year looks worse and a windy, wet year looks better — neither is policy. The red line is the accountability signal: did the city's own pollution run higher or lower in 2025 than in 2024? The **vertical gap between the lines is the weather's estimated share** of the change.

What de-weathering points to here. Where the red line sits above zero, 2025's own pollution ran higher than 2024 regardless of the weather; below zero, it genuinely ran cleaner.

PM2.5 — 2025 vs 2024 (µg/m³, 7-day average)

■ Observed change (weather included) ■ Weather-corrected change — the city's own pollution (model estimate)

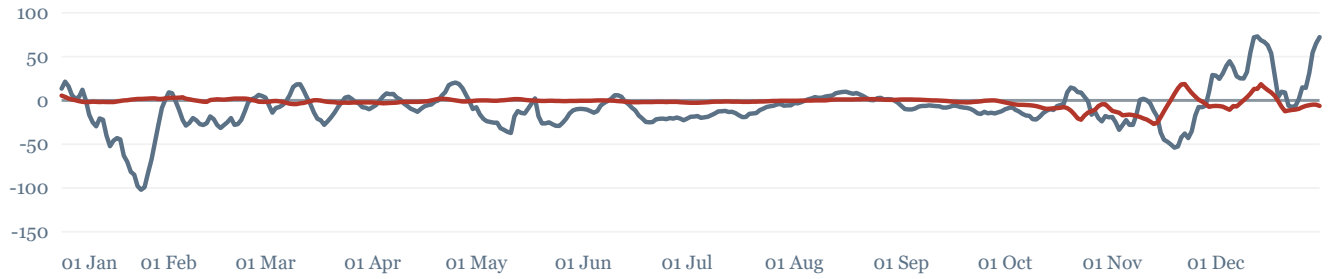

PM10 — 2025 vs 2024 (µg/m³, 7-day average)

■ Observed change (weather included) ■ Weather-corrected change — the city's own pollution (model estimate)

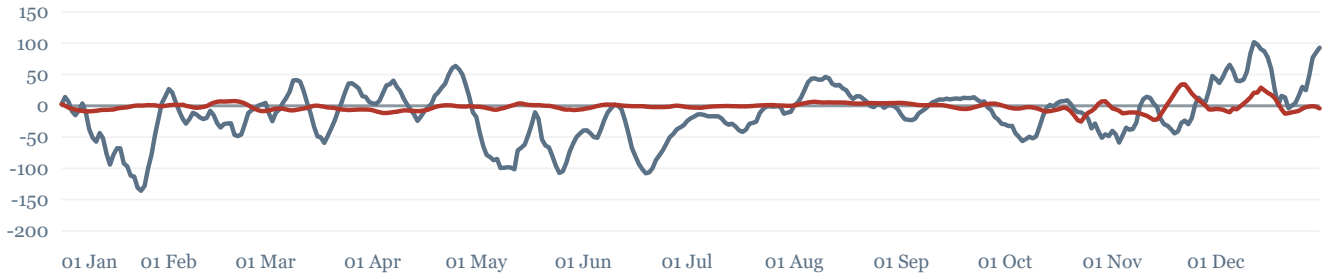

'2024' is the average of the same ±7 days in 2024 (a 15-day average, not a single day). This is a **retrospective re-estimate**: every 2025 and 2024 day was de-weathered with our current weather-correction model (trained through 2026-07-11), so the two years are compared on one consistent basis. Because that model was fit on data through 2026, both years are in-sample re-estimates — a descriptive lens for comparing the years, not what was forecastable at the time. The red line is a model estimate; the ±model margin shown per panel and on the summary page is the model's single-day error over each month's matched stations (deliberately conservative — a month's average must clear it before we state a magnitude), so the gap between the lines is an estimate too. Additionally smoothed with a 7-day centered average (needs ≥3 days; 363 of 365 days reconstructed, 363 usable). Same-station matched comparison; ±model-margin shown per panel overleaf.

2 · Month by month, at a glance

The one-glance summary of the pages that follow: after removing the weather, in which months did 2025's own pollution genuinely run above or below 2024?

Each bar is a month's **average** weather-corrected change (2025 minus 2024), µg/m³ per day — right of centre = 2025's own pollution ran higher, left = lower. **Full colour** = clears the model margin (a firm magnitude); half-tone = within margin but the same direction on ≥80% of days; grey = no clean signal. The small grey "obs" number is the observed change (weather included) — the gap to the bar is the weather's share.

PM2.5 — monthly average change, 2025 vs 2024 (µg/m³ per day, weather-corrected)

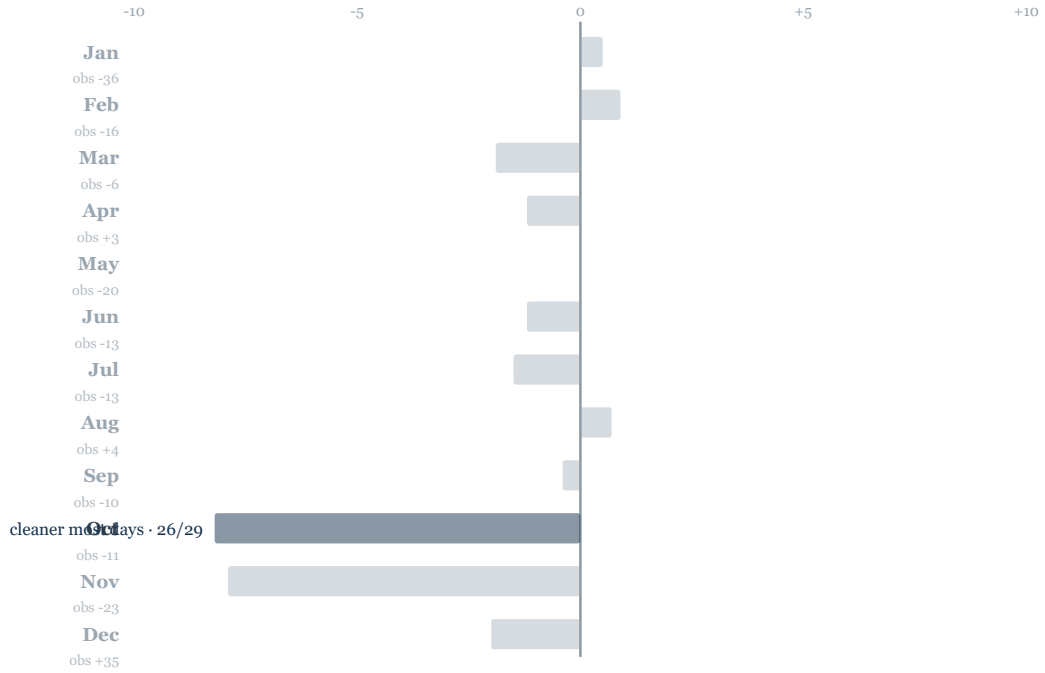

PM10 — monthly average change, 2025 vs 2024 (µg/m³ per day, weather-corrected)

Statistic: the **mean** of each month's daily weather-corrected residuals — not a sum (a sum would just reflect how many days a month has). Margin = the model's single-day error over the month's matched stations (deliberately conservative); a month is credited with a magnitude only when its mean clears it. Feb = 28 matched days (29 Feb 2025 has no 2024 counterpart).

PM2.5 — month by month (2025 vs 2024)

■ Observed change (weather included) ■ Weather-corrected change — the city's own pollution (model estimate)

January · scale ±110

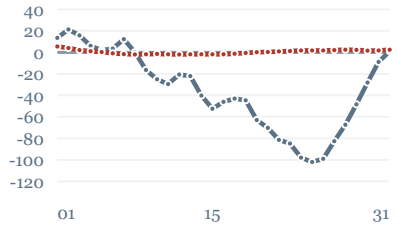

about level with 2024 — within the ±11 µg/m³ model margin (stations 39–40)

February · scale ±40

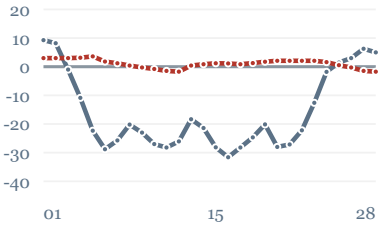

about level with 2024 — within the ±11 µg/m³ model margin (stations 39–40)

March · scale ±30

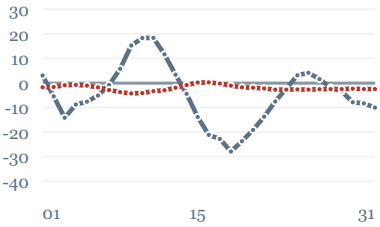

about level with 2024 — within the ±11 µg/m³ model margin (stations 40)

April · scale ±30

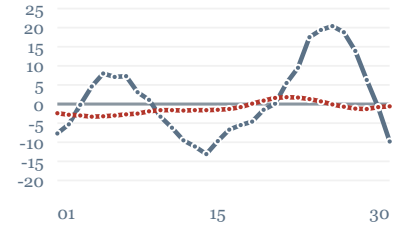

about level with 2024 — within the ±11 µg/m³ model margin (stations 40)

May · scale ±40

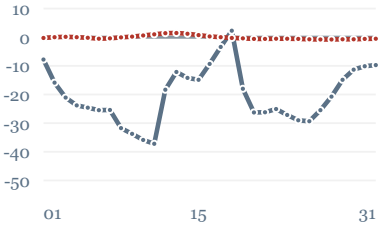

about level with 2024 — within the ±11 µg/m³ model margin (stations 40)

June · scale ±30

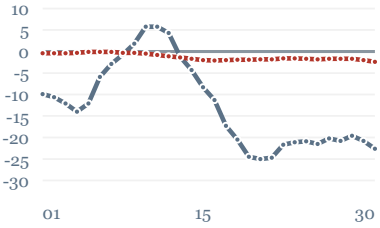

about level with 2024 — within the ±11 µg/m³ model margin (stations 40)

July · scale ±30

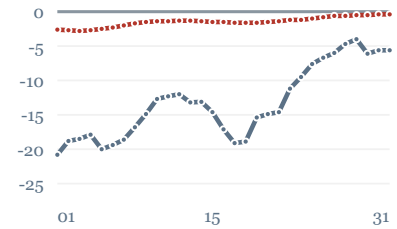

about level with 2024 — within the ±11 µg/m³ model margin (stations 40)

August · scale ±10

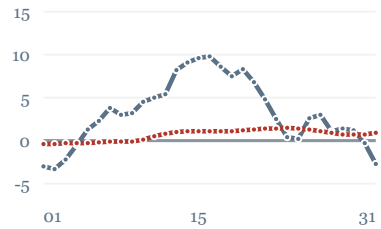

about level with 2024 — within the ±11 µg/m³ model margin (stations 40)

September · scale ±20

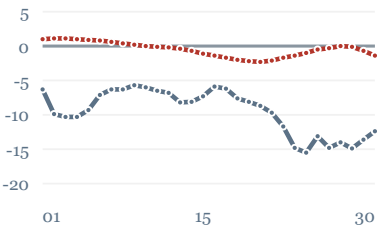

about level with 2024 — within the ±11 µg/m³ model margin (stations 40)

October · scale ±30

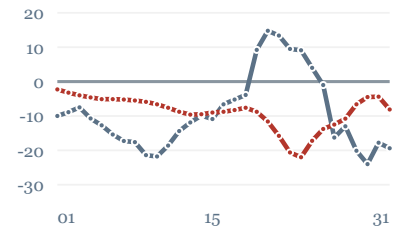

within the ±11 µg/m³ margin, but ran lower than 2024 on 26 of 29 days (stations 39–40)

November · scale ±60

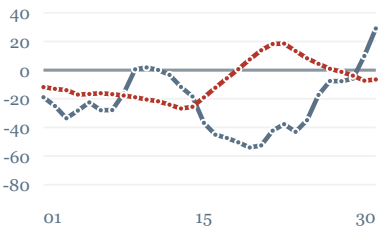

about level with 2024 — within the ±11 µg/m³ model margin (stations 40)

December · scale ±80

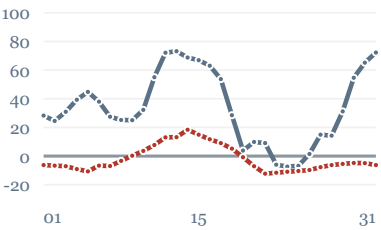

about level with 2024 — within the ±11 µg/m³ model margin (stations 40–41)

Each small chart is one calendar month of 2025 vs the same ±7 days in 2024; the 7-day average is centered and continuous, so a month's first and last few days blend into the neighbouring months. Grey = observed change, red = the city's own (weather-corrected) change.

PM10 — month by month (2025 vs 2024)

■ Observed change (weather included) ■ Weather-corrected change — the city's own pollution (model estimate)

January · scale ±140

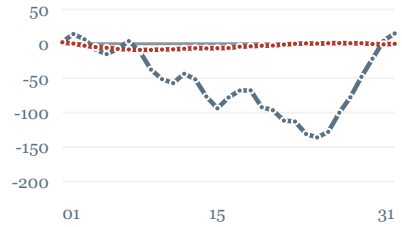

about level with 2024 — within the ±18 µg/m³ model margin (stations 38–39)

February · scale ±50

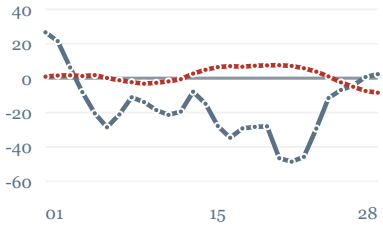

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

March · scale ±60

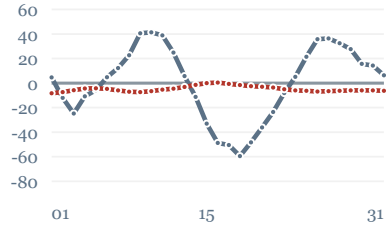

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

April · scale ±70

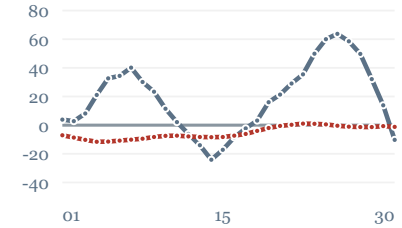

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

May · scale ±110

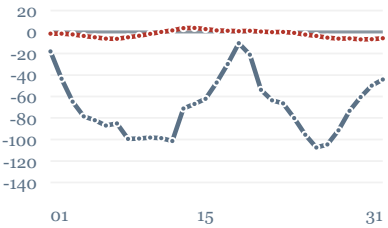

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

June · scale ±110

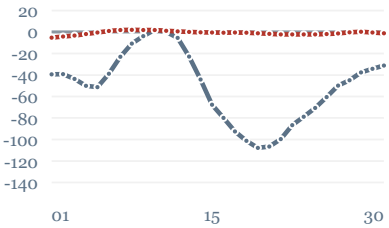

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

July · scale ±50

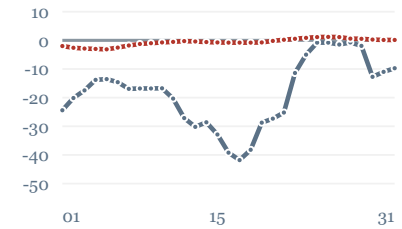

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

August · scale ±50

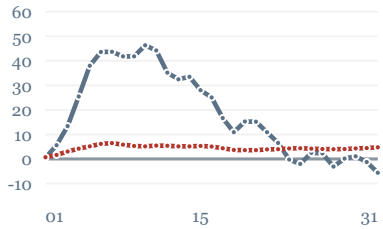

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

September · scale ±30

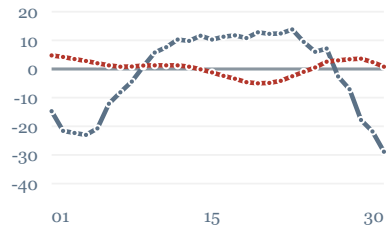

about level with 2024 — within the ±18 µg/m³ model margin (stations 38–39)

October · scale ±60

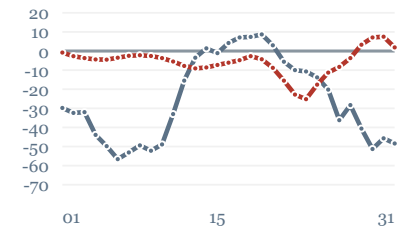

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

November · scale ±60

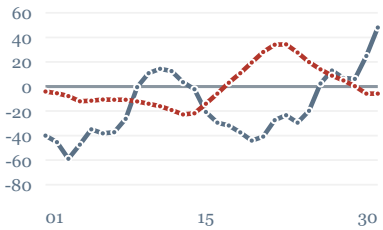

about level with 2024 — within the ±18 µg/m³ model margin (stations 39)

December · scale ±110

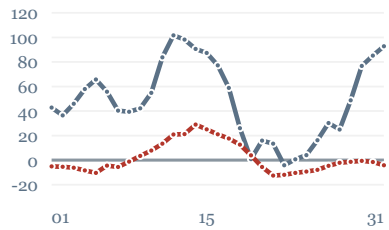

about level with 2024 — within the ±17 µg/m³ model margin (stations 39–40)

Each small chart is one calendar month of 2025 vs the same ±7 days in 2024; the 7-day average is centered and continuous, so a month's first and last few days blend into the neighbouring months. Grey = observed change, red = the city's own (weather-corrected) change.

5 · What Delhi actually breathes — population-weighted

Every chart so far gives each monitoring station equal weight. But stations sit in areas of very different population. This page re-draws the weather-corrected comparison **weighting each station by the population it represents** — closer to **what Delhi's people actually breathe** than a plain average of monitors.

Where these lines track the equal-weight charts, the plain average was already representative; where they part, exposure is concentrated where the pollution is.

PM2.5 — population-weighted, 2025 vs 2024 (µg/m³, 7-day centered average)

■ Observed change (weather included) ■ Weather-corrected change — the city's own pollution (model estimate)

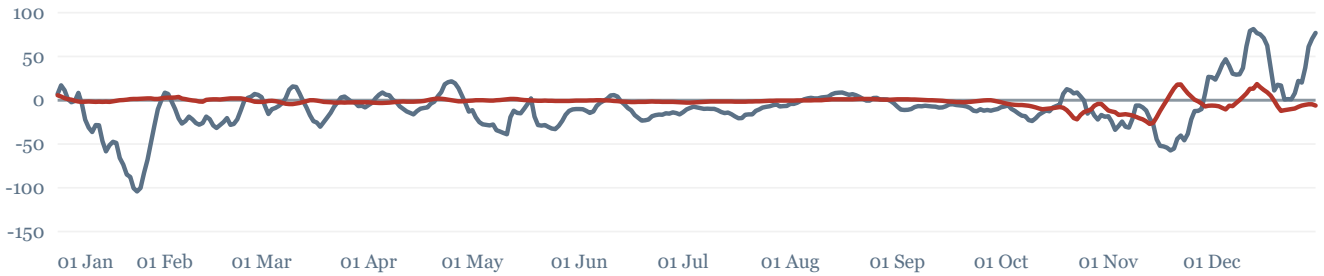

Over the year the population-weighted weather-corrected change averages -1.9 µg/m³, against -1.8 µg/m³ on the plain station average for the same days — a difference of 0.1 µg/m³.

PM10 — population-weighted, 2025 vs 2024 (µg/m³, 7-day centered average)

■ Observed change (weather included) ■ Weather-corrected change — the city's own pollution (model estimate)

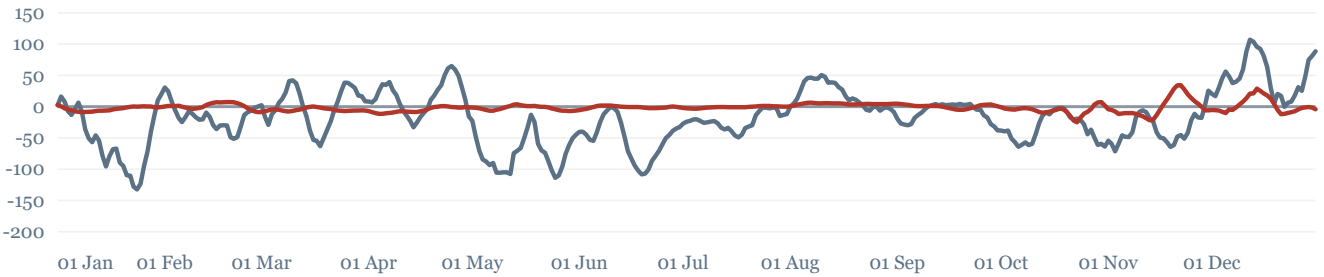

Over the year the population-weighted weather-corrected change averages -1.3 µg/m³, against -1.2 µg/m³ on the plain station average for the same days — a difference of 0.1 µg/m³.

Population weights cover the stations with a mapped population catchment (about 44); days without enough weighted coverage are left blank. Same de-weathering and the same ±7-day matched baseline as the rest of the report — only the city average is population-weighted here instead of a plain station mean.

6 · Robustness & data integrity

Why a ±7-day window. The baseline is a 15-day average — the same calendar date ±7 days — not a single day, for a firm reason. Delhi's day-to-day air is driven by weather systems (passing western disturbances, stagnant high-pressure spells and the mixing they control) that come and go every three to seven days; a 15-day window spans two to three of those cycles, so no single storm, calm night or shower can define the "normal". Because neighbouring days are correlated, the 15 days act like about seven independent readings — enough to cut the baseline's one-day noise roughly two-and-a-half-fold — while staying narrow enough to remain inside the same pollution season (a month-long window would blur the sharp changes at the monsoon's arrival and the late-autumn crop-burning surge, and multiply the turning-point bias). A single-date comparison would be *more* sensitive to one day's weather, not less, so the window is deliberately the harder-to-game choice — the standard "anomaly from normal" method (as a temperature "departure from normal" compares today to the long-run normal for the date). The same ±7-day window is used identically in the daily bulletin, the daily report and this retrospective. One honest limit: festival and crop-burning emissions are real pollution and are kept in the numbers by design, so where a moving festival such as Diwali falls at different calendar positions across the two years, its spike can sit inside one year's window and not the other's — October–November months carry that caveat.

Does the window choice change the answer? (robustness) The ±7-day window is fixed in the code and applied to every day, but the fair test is whether the verdict survives a *different* window. Below is the full-year mean weather-corrected change (2025 minus 2024, µg/m³ per day) recomputed at four window widths:

	±3	±7	±14	±30
PM2.5	-2.6	-2.6	-2.5	-2.3
PM10	-2.0	-1.9	-1.9	-1.8

Across every window from ±3 to ±30 days the de-weathered year gap moves by less than 0.4 µg/m³ and never changes sign — the ±7-day choice does not drive the result. (City-average; the matched-station ±7 headline used elsewhere in this report differs by under 1 µg/m³ from station composition. The window only re-averages the baseline — each day's weather-corrected value is unchanged — so no re-modelling is involved.)

Diwali 2025 — why 20–21 October are excluded. Delhi's monitoring network went dark precisely when Diwali pollution peaked. On the night of 20–21 October 2025, PM2.5 station coverage fell from about 39 to as few as **12 of ~40 stations** between midnight and 3 AM — the firecracker peak — recovering to full coverage by late morning. Because the daily average uses a "16 of 24 hours" rule, stations that dropped only during the peak still counted, but with their highest hours missing — so the official Diwali AQI is **biased low**. The Supreme Court was told only about 9 of 37 stations reported continuously through the period and sought a report from the CAQM and CPCB; the matter is also before the National Green Tribunal, alongside allegations of water-sprinkling near monitors. The data proves an extensive, peak-timed reporting failure; whether it was deliberate is not established. Either way the readings understate reality, so we exclude these two days from the weather-corrected comparison rather than report an "improvement" the data cannot support. Full detail in the companion monitoring-forensic brief.

How this report is made

The comparison. Every pollution number is a 24-hour average ending at 10 AM IST (the CPCB window). A station counts on a day only if it reported ≥16 of 24 hours. The "city value" is the plain average across stations that passed, and each 2025 day is compared with a **matched-station** average of the same ±7 calendar days in 2024. Why ±7 days — and how much the window choice matters — is set out on the previous page ("Robustness & data integrity").

De-weathering. A random-forest model (trained on 2021–2026 hourly data) predicts what each station-hour's pollution would have been under that day-of-year and hour's *normal* weather, holding the actual traffic/day-of-week/seasonal trend. The gap between the observed and the weather-normalised value is the weather's contribution. One consistent current model is used for both years so the red line's movement reflects the city's own pollution, not a change of model.

Honest margins. The weather-corrected line is a model estimate; each month's panel shows the ±model margin (from the model's cross-validated error and the month's station count), and a month is only called "above" or "below" 2024 when it clears that margin. Grey/settled months are read on their shape, not any single day.

Coverage. 363 of 365 days of 2025 were reconstructed (363 passed the full-network floor). PM2.5 and PM10 only; the gas channels are excluded.

Baseline coverage. Each station's last-year figure averages the days it reported within the 15-day window — not always the full 15. Across the 363 compared days, 116 had at least one station whose baseline rested on fewer than 8 days (thinnest: 1). Those are the "composition-sensitive" days already flagged on the panels, so the comparison never silently treats a thin average as a full 15-day one.

Was the baseline year a freak-weather year? It does not distort the red line. Because both years are put through the same weather-correction model, the weather-corrected (red) comparison is **weather-neutral by construction** — an unusually calm or windy 2024 cannot bias it. Weather shows only in the observed (grey) line, so the gap between grey and red at any point *is* the weather's share; a freak-weather baseline year would appear there as a wide grey–red gap, not as a shift in the red line. (A per-date check of whether 2024's weather was itself unusual, matching the daily bulletin's climatology flag, is a separate enhancement.)

Pollution: CREA compilation of CPCB/DPCC station data (raw hourly). De-weathering: the AQI_Analysis Z2 random-forest model. All figures recomputed from source at build time.