WEEKLY DEEP ANALYSIS · DELHI

Delhi Air: the week in review

Week ending 18 July 2026 · prepared by Ajay Maken MP and Team

Across the seven days to 18 July 2026, Delhi's air averaged **196** on the national index. The worst day was **Monday 13 July**, at **310**. Not one of the week's 311 station readings was rated Good.

196
The week's average air-quality index, as published each morning

99
The level the city's own pollution justifies, once the weather is removed

+122
Points against the same days last year, matched stations

102
Of 311 station readings this week were rated Poor or worse

1 · The seven days at a glance

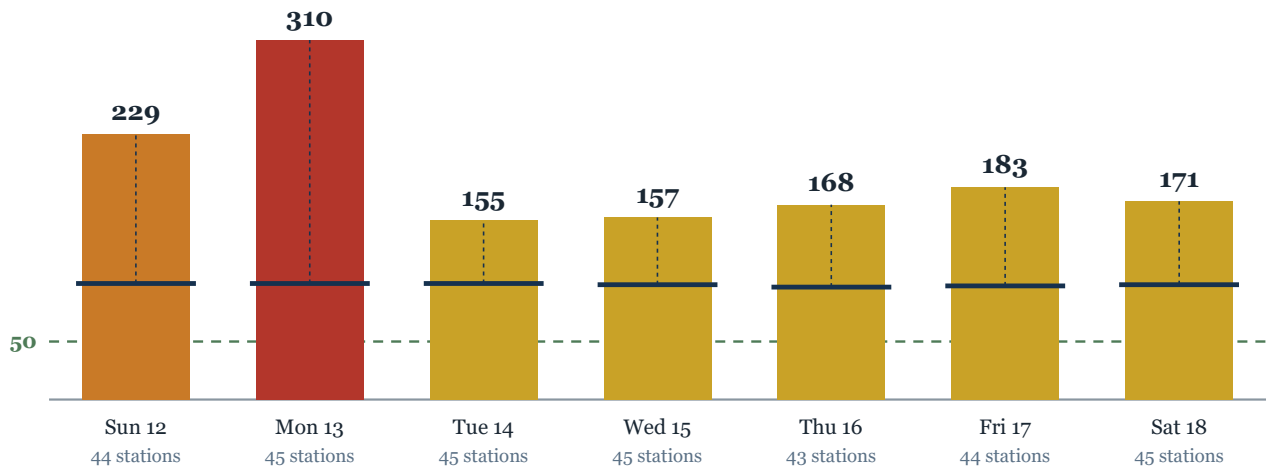

Each column is one day, coloured by the category the index fell in. The navy rule on each column marks the level the city's own pollution justifies once the weather is removed, so the gap between that rule and the top of the column is the weather's share of the day. A day with no bulletin is drawn as an empty slot rather than dropped, because a shortened week would misrepresent it. The green line at 50 is the top of the clean-air band: below it the air is rated Good, and Delhi did not reach it on any day this week.

2 · The week, day by day

Day	Published that morning	Stations then	Recomputed for this report	Stations now	Category	Weather-corrected
Sun 12 Jul	229	44	234	45	Poor	100
Mon 13 Jul	310	45	309	45	Very Poor	100
Tue 14 Jul	155	45	155	45	Moderate	100
Wed 15 Jul	157	45	157	45	Moderate	99
Thu 16 Jul	168	43	168	44	Moderate	97
Fri 17 Jul	183	44	185	45	Moderate	98
Sat 18 Jul	171	45	171	45	Moderate	99

Why two columns. The figure in the second column is the one published on the morning of that day. The figure in the fourth column is the same day recalculated for this report. They can differ for one ordinary reason: some monitoring stations report late, so a day's city average is built on a few more stations a week later than it was on the morning itself. The station count beside each figure shows exactly that. This week 3 days have since moved: 12 Jul from 229 to 234; 13 Jul from 310 to 309; 17 Jul from 183 to 185. Where the two differ, the figure published that morning is the one that was published, and this report does not overwrite it.

The same station, first on every day of the week

On each of the 7 days of this week the dirtiest air in Delhi was measured at **Anand Vihar**, and on each of those days the pollutant responsible was PM10. The station that came second changed during the week; the station that came first never did. Its lead over the day's runner-up was never smaller than 49 index points and reached 136. A single day's worst station can be an accident of one faulty afternoon. 7 days out of 7, with that margin, is not an accident.

The index stops at 500. That is where the scale ends, not where the air stopped getting worse, so this station's weekly average of 384 is a floor rather than a middle.

Anand Vihar is one site and it is not the whole city. The city-wide figure makes the same point: across the week's 311 station readings, 102 were Poor or worse and not one was rated Good.

3 · The weather, or the city's own pollution?

Measured against the same seven days last year, the week ran **+122 points** on the meter. After the weather is removed, **+21 points** of that survive as the part the city's own pollution produced. The difference between those two numbers is the weather's contribution, in whichever direction it fell.

> **Why this distinction carries the whole report.** A calm, still week makes the air look worse and a windy, wet week makes it look better, and neither is anyone's policy. Raw numbers therefore flatter or punish a government for the sky's luck. The weather-corrected figure is the accountability signal, because it answers whether the city's own pollution actually rose or fell.

Which winds bring Delhi its dirtiest air

A pollution rose is a compass built from the last 90 days. Each spoke points to a direction the wind blew from. **The length of a spoke shows how often the wind came from that direction, so a long spoke is a frequent wind and not a dirty one.** The colour bands along each spoke show how dirty the air was on those days. Read length and colour together: a long pale spoke is a common but clean wind, while a shorter deep-red one is less frequent but dirtier.

> **The dirtiest and the cleanest wind.** Over this period days with a dominant **NW** wind (24 such days) carried the highest typical fine-particle level, while **SE** winds (16 days) brought the cleanest air. Only sectors with at least 5 days are ranked, so a one-day or two-day wind cannot be crowned. These two are the highest and the lowest of the 4 sectors ranked, and no significance test stands behind that ordering, so read them as the ends of this period's range rather than as a measured gap.

Fine particles, PM2.5

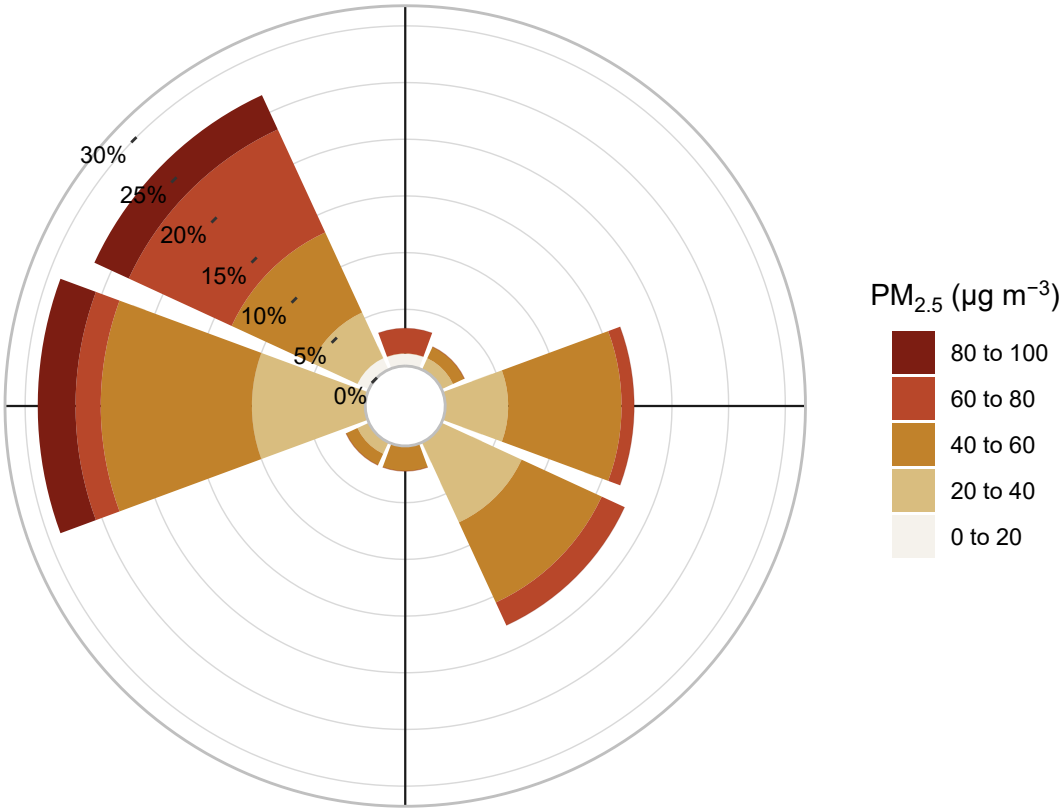

Which winds bring Delhi its dirtiest air, continued

Coarse dust, PM10

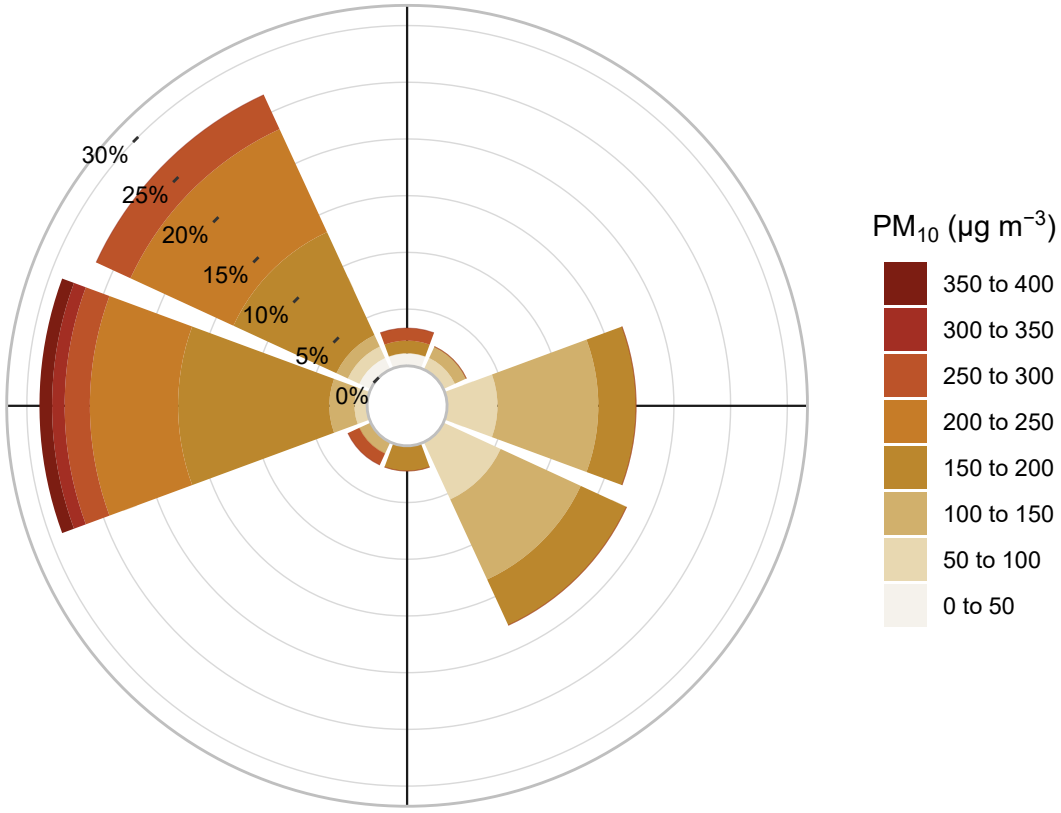

Read length and colour together. A long pale spoke is a common but clean wind; a shorter deep-red spoke is a less frequent but dirtier one. The dirtiest wind named earlier is the sector with the highest typical reading, which need not be the longest spoke on the rose.

Does wind direction really matter?

The chart below groups the daily city readings by the day's dominant wind direction, so you can see which sectors tend to sit high and which sit low. But a raw difference between sectors can mislead: in Delhi the cleaner-looking winds are often simply the stronger winds, and strong winds clear pollution no matter where they come from. The table that follows runs the honest test. It asks whether direction still explains pollution after wind speed and the depth of the atmosphere have already been accounted for.

Delhi PM2.5 and PM10 by dominant wind direction (last 90 days)

Daily city-averages mapped to the day's predominant wind (24h to 10AM IST). The day count is printed under each sector, and the 4 sectors below 5 days are drawn unfilled because they are too thin to rank.

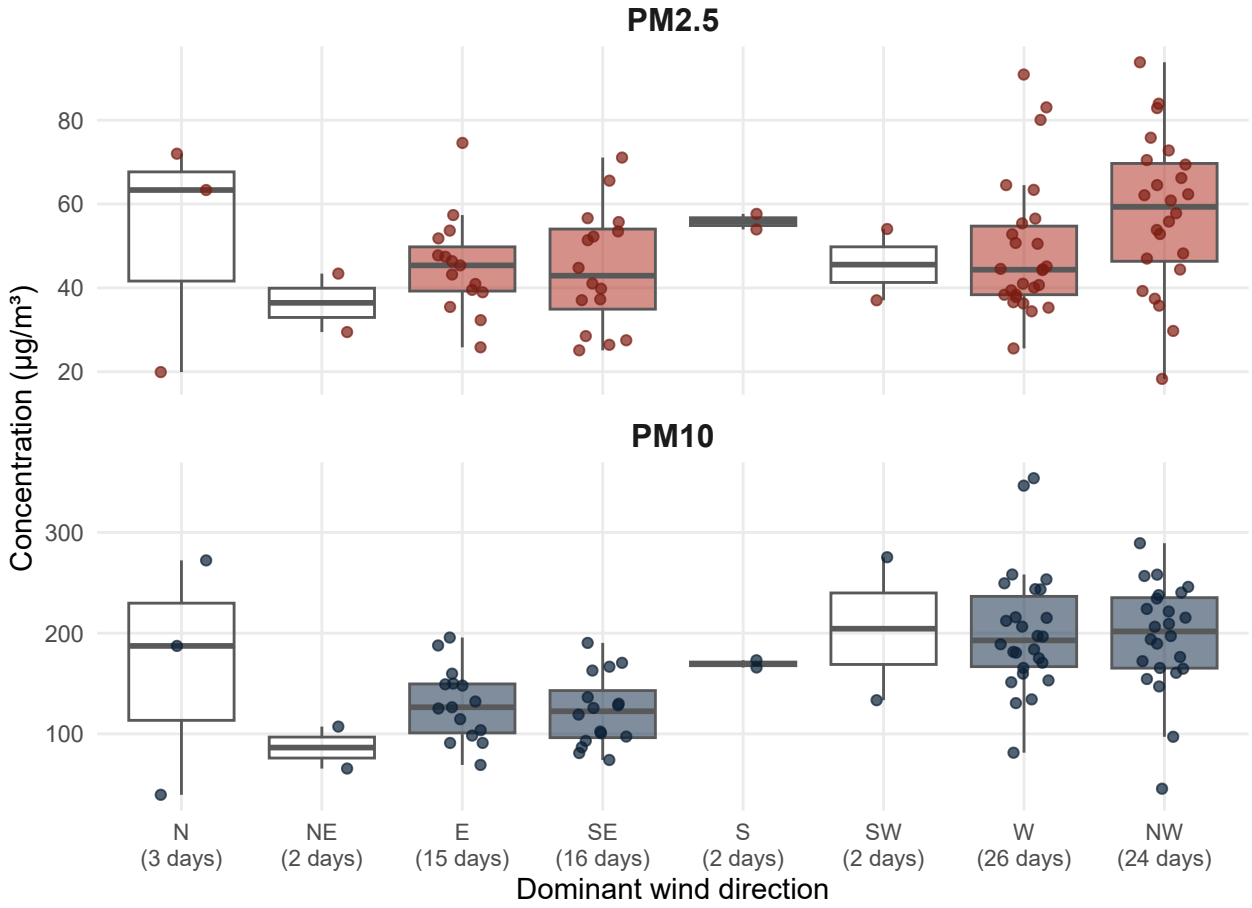

Does wind direction survive the honest test?

The table below asks whether wind direction still explains pollution after wind speed and the depth of the atmosphere have already been accounted for. Only if direction survives that adjustment can we say a sector is genuinely bringing pollution in, rather than simply being a calm or a windy day with a directional label.

Month	Pollutant	Days tested	Wind sectors	How likely this is chance	After the multiple-test screen	Does direction matter?
April	PM2.5	30	7	0.467	0.533	No
April	PM10	30	7	0.575	0.575	No
May	PM2.5	31	5	0.081	0.216	No
May	PM10	31	5	0.001	0.005	Yes
June	PM2.5	30	7	0.001	0.005	Yes
June	PM10	30	7	0.193	0.309	No
July	PM2.5	18	7	0.262	0.350	too few days
July	PM10	18	7	0.161	0.309	too few days

A lower figure means wind direction is doing more work, and the screen makes the test harder to pass. A version of this test corrected for day-to-day persistence in the weather was built and tested; at the number of days available in a single month it proved unreliable in both directions, so this table reports the ordinary test with the sample size shown and the corrected version is held back until it can be made dependable.

Weather against PM2.5, month by month

Each square answers one question: in this month, when the weather ran unusual for the time of year, did the pollution run unusual in the same way? A coloured square is a link the report is willing to stand behind. A white square with a dotted edge did not pass the reliability screen, and its number is printed for completeness only.

Against the same days last year

This grid asks whether, on the days when Delhi's weather differed from the same days a year earlier, its pollution differed in step. That is the question of a whole season pushed off course.

	January	February	March	April	May	June	July
Wind speed *cleaner air when higher*	−0.08 (−0.06)	−0.23 (−0.28)	−0.26 (−0.28)	−0.18 (−0.15)	−0.18 (−0.18)	−0.21 (−0.18)	−0.07 (−0.05)
Temperature *no fixed direction*	+0.43 (+0.24)	**+0.45** (+0.47)	**+0.70** (+0.75)	**+0.65** (+0.45)	**+0.66** (+0.62)	**+0.70** (+0.68)	+0.63 (+0.68)
Depth of the mixing layer *cleaner air when higher*	−0.31 (−0.22)	+0.40 (+0.34)	+0.39 (+0.38)	+0.12 (+0.14)	+0.30 (+0.25)	**+0.43** (+0.45)	+0.18 (+0.16)
Rainfall *cleaner air when higher*	+0.08 (+0.18)	−0.27 (−0.28)	−0.15 (−0.31)	**−0.45** (−0.32)	−0.16 (−0.29)	−0.27 (−0.31)	−0.49 (−0.53)
Morning inversion strength *cleaner air when lower*	−0.37 (−0.19)	+0.07 (+0.03)	−0.03 (−0.00)	+0.22 (+0.37)	+0.06 (+0.18)	−0.09 (+0.00)	*not measured*
Morning inversion height *cleaner air when higher*	+0.08 (+0.07)	**−0.54** (−0.16)	−0.00 (−0.02)	**−0.45** (−0.38)	−0.36 (−0.48)	−0.33 (−0.40)	*not measured*

Against the day before

This grid asks a sharper question. On the days when the weather changed from the day before, did the pollution change with it? Comparing a day only with the day before removes everything that drifts slowly, which means the season, the calendar and any change in the city itself, and leaves only sudden movements.

	January	February	March	April	May	June	July
Wind speed *cleaner air when higher*	**−0.63** (−0.67)	−0.46 (−0.41)	**−0.60** (−0.57)	−0.33 (−0.34)	−0.46 (−0.40)	**−0.56** (−0.53)	+0.16 (+0.18)
Temperature *no fixed direction*	+0.18 (+0.15)	+0.00 (+0.00)	+0.18 (+0.09)	+0.38 (+0.46)	+0.32 (+0.32)	+0.21 (+0.22)	+0.56 (+0.63)
Depth of the mixing layer *cleaner air when higher*	**−0.53** (−0.56)	−0.35 (−0.24)	−0.36 (−0.38)	−0.05 (−0.03)	−0.18 (−0.19)	−0.08 (−0.10)	+0.17 (+0.17)
Rainfall *cleaner air when higher*	−0.01 (+0.06)	+0.35 (−0.09)	−0.15 (−0.44)	−0.04 (−0.21)	−0.32 (−0.28)	−0.39 (−0.26)	−0.20 (−0.63)
Morning inversion strength *cleaner air when lower*	−0.11 (+0.08)	+0.40 (+0.04)	+0.25 (+0.30)	+0.22 (+0.17)	−0.01 (+0.02)	−0.16 (−0.09)	*not measured*
Morning inversion height *cleaner air when higher*	−0.09 (−0.19)	−0.26 (+0.10)	−0.29 (−0.10)	−0.15 (−0.16)	−0.27 (−0.18)	−0.14 (+0.29)	*not measured*

The second line under each weather name says which direction usually helps Delhi's air. A number in brackets is a second measure of the same link, one that is not thrown by a handful of extreme days; where the two differ materially the bracketed figure is the more trustworthy of the two.

Weather against PM10, month by month

Each square answers one question: in this month, when the weather ran unusual for the time of year, did the pollution run unusual in the same way? A coloured square is a link the report is willing to stand behind. A white square with a dotted edge did not pass the reliability screen, and its number is printed for completeness only.

Against the same days last year

This grid asks whether, on the days when Delhi's weather differed from the same days a year earlier, its pollution differed in step. That is the question of a whole season pushed off course.

	January	February	March	April	May	June	July
Wind speed *cleaner air when higher*	−0.08 (−0.04)	−0.24 (−0.22)	−0.01 (−0.06)	−0.06 (−0.04)	−0.21 (−0.19)	−0.15 (−0.02)	+0.31 (+0.14)
Temperature *no fixed direction*	+0.39 (+0.18)	**+0.58** (+0.62)	**+0.67** (+0.70)	**+0.74** (+0.47)	**+0.78** (+0.71)	**+0.76** (+0.76)	+0.80 (+0.86)
Depth of the mixing layer *cleaner air when higher*	−0.35 (−0.30)	+0.34 (+0.28)	**+0.49** (+0.47)	+0.03 (+0.01)	+0.34 (+0.28)	**+0.51** (+0.52)	+0.51 (+0.46)
Rainfall *cleaner air when higher*	−0.03 (−0.02)	−0.27 (−0.25)	−0.26 (−0.39)	**−0.40** (−0.20)	−0.06 (−0.23)	−0.32 (−0.44)	−0.50 (−0.70)
Morning inversion strength *cleaner air when lower*	−0.28 (−0.04)	+0.10 (+0.08)	−0.05 (−0.00)	+0.12 (+0.29)	+0.08 (+0.06)	−0.01 (+0.10)	*not measured*
Morning inversion height *cleaner air when higher*	+0.04 (−0.01)	**−0.48** (−0.08)	+0.02 (+0.04)	−0.27 (−0.16)	−0.25 (−0.41)	−0.29 (−0.22)	*not measured*

Against the day before

This grid asks a sharper question. On the days when the weather changed from the day before, did the pollution change with it? Comparing a day only with the day before removes everything that drifts slowly, which means the season, the calendar and any change in the city itself, and leaves only sudden movements.

	January	February	March	April	May	June	July
Wind speed *cleaner air when higher*	**−0.73** (−0.77)	**−0.55** (−0.56)	**−0.65** (−0.64)	−0.11 (−0.20)	−0.24 (−0.16)	−0.27 (−0.16)	+0.33 (+0.22)
Temperature *no fixed direction*	+0.20 (+0.11)	+0.19 (+0.10)	+0.31 (+0.16)	+0.45 (+0.57)	+0.39 (+0.41)	**+0.51** (+0.36)	+0.58 (+0.72)
Depth of the mixing layer *cleaner air when higher*	**−0.60** (−0.59)	−0.39 (−0.31)	−0.45 (−0.42)	−0.01 (−0.04)	−0.31 (−0.22)	+0.23 (+0.03)	+0.39 (+0.27)
Rainfall *cleaner air when higher*	−0.22 (−0.12)	+0.17 (−0.12)	−0.23 (−0.45)	−0.08 (+0.01)	−0.29 (−0.22)	−0.32 (−0.17)	−0.16 (−0.67)
Morning inversion strength *cleaner air when lower*	−0.05 (+0.20)	+0.30 (+0.01)	+0.28 (+0.29)	+0.11 (+0.01)	−0.03 (−0.01)	−0.16 (−0.24)	*not measured*
Morning inversion height *cleaner air when higher*	−0.07 (−0.08)	−0.16 (+0.18)	**−0.45** (−0.21)	+0.02 (+0.07)	−0.06 (+0.07)	−0.15 (+0.29)	*not measured*

The second line under each weather name says which direction usually helps Delhi's air. A number in brackets is a second measure of the same link, one that is not thrown by a handful of extreme days; where the two differ materially the bracketed figure is the more trustworthy of the two.

How to read these grids

What the grids measure

Neither grid is about how dirty the air was. Both are about departures: how far a day sat above or below what would be normal for it. A month can be filthy throughout and still show no link here, if its filth did not rise and fall with the weather.

The two grids can disagree, and when they do it is usually because they are answering different questions rather than because one of them is wrong. Read each for the question it asks.

One square, read in full

Take the square where wind speed meets fine particles in June. Suppose it reads about minus one half. Here is what produces a number like that.

On one day the wind drops well below normal and the pollution climbs above normal. That is one day, and one day settles nothing. On another the wind picks up and the pollution falls. Now take every consecutive pair of days in the month. On most of them the wind and the pollution moved in opposite directions, which is the direction Delhi would want; on the rest they moved the same way.

The single number summarises all of those days at once. The minus sign says that when the wind picked up, the pollution usually fell. The size says how firmly. Squared, it becomes the plainest statement of strength: a correlation of about minus one half means the two shared roughly thirty per cent of their day-to-day movement, and the other seventy per cent moved independently.

That seventy per cent is not a rounding error. There are days in every month when the wind rises and the pollution rises anyway. A correlation is a tendency across a month. It is never a rule that must hold on a given morning.

Three things a square does not say

It does not say the weather caused the change. In Delhi the wind, the rain, the temperature and the depth of the mixing layer all rise and fall together, so any one of them will appear to move with pollution partly because the others did.

It does not say the weather is why Delhi's air is bad. Wind decides how much of a day's pollution drifts away and how much settles over the city. It does not decide how much the city puts into the air to begin with. In a month when the wind had been generous on every single day, Delhi's vehicles, construction sites, industry and waste fires would have emitted precisely what they emitted.

And it does not carry to another month. The same pairing can be strong in one month and absent in the next, which is exactly why this report prints every month rather than a single figure for the year.

The two tests, and why there are two

Why one significance test is not enough

Delhi's weather does not start afresh each morning. A still spell lasts three or four days, and so does a windy one, and the pollution beneath it behaves the same way. Thirty days of weather are therefore not thirty separate pieces of evidence. The ordinary significance test assumes that they are, and so it can credit a month with more confidence than the month has actually earned.

This report therefore runs a second test alongside the ordinary one, a correction published by Whitney Newey and Kenneth West. Rather than treating each day as fresh, it measures how much each day still resembles the day before it, the day two before and the day three before, and rebuilds the margin around the correlation to match. The number of days it looks back is fixed by a rule and not chosen by hand, and for a month of about thirty days that rule gives three.

The correction does not always make this report stricter, and we will not pretend that it does. On the day-on-day grid each reading is already a change from the day before, and that subtraction has usually over-corrected for the carry-over. When that happens the rebuilt margin comes out narrower than the ordinary one, and the correction makes a finding easier to reach rather than harder. Measured across all 160 tests in this edition, the correction moved the answer towards caution in only 49 of them.

A correction that can loosen the standard cannot be the only guard. So this report computes both answers and keeps whichever of the two is the more cautious. A square has to satisfy both tests before it is allowed to carry colour. That is a stricter standard than either test on its own, and it cannot be embarrassed by either one misbehaving.

Why a grid needs a second screen

One grid on these pages is not one test. For most months it is twelve, being six weather measurements set against two pollutants. Across the months and both grids, this edition runs 160 of these tests in all.

That is the difficulty. Toss a coin often enough and some run of heads will look remarkable. Compare enough pairs of unrelated numbers and some pair will look connected. Testing twelve things at once and then pointing at the most striking of them is how a report manufactures a finding that was never there.

The screen applied here was published by Yoav Benjamini and Yosef Hochberg. It ranks a grid's results from the most striking to the least and raises the bar for each in turn, so a result has to be considerably more striking to be believed when it is one of twelve than when it stands alone. The bar is set at ten per cent, and the plain meaning of that setting is this: among all the squares this report allows to carry colour, we expect about one in ten to be a coincidence rather than a real link, and we cannot tell you which one. Each month and each grid is screened on its own, so a weak month cannot borrow credibility from a strong one.

A square computed on fewer than 20 usable pairs of days is not tested at all, because a margin that corrects for carry-over cannot be trusted on a handful of days. Its number still prints, in grey. Of the 160 squares in this edition, 28 carry colour.

One honest limit. This screen works cleanly when the tests on a grid do not conspire with one another. Delhi's weather does not fully oblige, because wind, rain, temperature and mixing depth rise and fall together, so the twelve tests on a grid are related to each other. Read the ten per cent as a sound working guide rather than as a guarantee.

Annexure · Method, in plain English

The daily building block

Every pollution number is a 24-hour average ending at 10 in the morning, the same window Delhi's pollution control board uses, so a day in this report always means yesterday morning to this morning. A monitoring station counts on a given day only if it reported for at least sixteen of those twenty-four hours, which keeps a half-broken sensor from skewing the city figure.

Why only two pollutants

Delhi's monitors also measure four gases, but that feed froze on 23 June 2026 and its numbers cannot be trusted, so this edition reports fine particles and coarse dust only. Their measurements are independent of that fault. This is also why the daily report shows blank days in late June where this one does not: a national index is a composite that needs the gas readings, so it could not be formed on days when the particle monitors were reporting perfectly well.

The week's figures

The week is the average of the figures published each morning, never a fresh calculation of the whole week at once. Those daily figures were produced by one version of the weather-correction model, and recomputing the week now would use a model none of those mornings used.

Prepared by Ajay Maken MP and Team.

What the correction can and cannot say

The weather-corrected figure is an estimate from a model, not a measurement. It is reported without a margin in this edition, because the honest way to measure that margin is still being built and quoting a made-up one would be worse than quoting none. No claim on these pages depends on it: the week's verdict rests on a gap of many points, the widespread finding is a count of stations, and the persistence finding is a count of days.

The screen, in one paragraph

Each square is tested twice, by an ordinary significance test and by one that tolerates the way Delhi's weather persists for days at a time, and the more cautious of the two answers is the one that counts. Every grid is then screened again for the fact that it runs twelve tests at once, at a ten per cent false-discovery rate. A square computed on fewer than 20 usable pairs of days is not tested at all. Squares that fail print their number in grey.

Provenance

Pollution figures come from the CREA compilation of the pollution control board's station data. Weather comes from the ERA5 and Open-Meteo hourly record for Delhi, plus the Safdarjung weather balloon. Every figure in this report was recomputed from those files at the moment it was built, and the exact input file is named in the footer of every page.