

Yamuna at Delhi — BOD–DO modelling under monsoonal hydrology

A water-quality modelling case study — Wazirabad to Okhla, Delhi

Dr Venkat Nara · Research Fellow, University of Melbourne · venkat.nsn@gmail.com

Executive summary

The 22 km Wazirabad → Okhla reach of the Yamuna in Delhi is a difficult dry-season BOD–DO case. After upstream diversions the headwater flow drops to about 4 m³/s, while the Najafgarh, Shahdara and Sahibabad drains collectively inject roughly 52 m³/s of high-BOD wastewater along the reach. The question I set out to answer is which combination of source treatment, drain interception and managed environmental flow can move the reach toward CPCB Class C compliance, and which of the two Class C standards (DO ≥ 4 mg/L, BOD ≤ 3 mg/L) is actually binding. The analysis is screening-level: its purpose is to inform whether a higher-complexity tool (QUAL2K, WASP) is justified.

I built a 1D steady-state Streeter–Phelps model for the reach, calibrated it against 2018–2022 monthly medians at four CPCB stations, propagated parameter and forcing uncertainty through a 2000-run Latin Hypercube Monte Carlo, and translated the outputs into compliance probabilities under five management scenarios. The calibration is "Very Good" by Moriasi et al. (2007): pooled NSE = 0.78 (BOD) and 0.95 (DO), with PBIAS of +1.6% and –0.9%. Joint Class C compliance is effectively unreachable on this reach in the dry season under the tested scenarios — even the most aggressive scenario (S4: ~70% drain treatment plus a tripled Wazirabad release) leaves median Okhla BOD at about 7.7 mg/L. BOD is the binding standard, not DO. The DO standard, by contrast, is reachable: under the aggressive STP upgrade alone (S2), modelled P(DO ≥ 4) at Okhla rises from 4.8% to 81.5%, and S4 takes it to 86.9%. The Spearman ranking explains the asymmetry: drain BOD is the dominant driver of Okhla BOD ($\rho = +0.93$), while headwater Q is a weak driver ($|\rho| < 0.1$) because its dry-season magnitude is an order of magnitude smaller than the combined drain inflow.

Management insight
Binding constraint: BOD, not DO. Joint Class C is not reachable on this reach in the dry season.
Best near-term lever: drain treatment / STP upgrade. P(DO ≥ 4) rises from ~5% to ~82% at ~70% drain BOD removal alone.
Weak standalone lever: environmental-flow augmentation. ×3 Wazirabad release alone moves P(DO ≥ 4) only to ~8%.

1. Context and objectives

CPCB has classified the Yamuna at Delhi as Class E (fit only for irrigation and industrial cooling) at every station downstream of Wazirabad for two decades, despite about ₹3,000 crore of cumulative investment under the Yamuna Action Plan I–III (1993–present). The technical question is no longer whether the reach is non-compliant — that is settled — but which combination of source control, treatment and managed environmental flow can move it into Class C, and with what residual probability of failure.

The objectives are: (i) build a transparent, public-data 1D model of the reach that an engineering reviewer can audit on one page; (ii) calibrate it against the four CPCB stations along the reach; (iii) propagate plausible parameter and forcing uncertainty to a compliance probability; and (iv) compare five management scenarios on a common compliance metric, so an investment case can be argued in probability terms rather than as a

single median estimate. The model is screening-level, not a regulatory model: it is meant to test whether investment should prioritise treatment, flow release, interception or a combined pathway, before a higher-complexity tool (QUAL2K, WASP) is commissioned. Limitations are stated explicitly in §4.1.

2. Study area and data

The reach extends from Wazirabad barrage (km 0) to Okhla barrage (km 22) and falls within the National Capital Territory of Delhi. Three major drains discharge into the reach at km 1.5 (Najafgarh), km 13.5 (Shahdara), and km 18 (Sahibabad). Headwater flow, drain flows, and drain BOD/DO are taken from CWC Old Railway Bridge gauge records and CPCB Yamuna Action Plan reports. Calibration data are 2018–2022 monthly medians of BOD and DO at the four CPCB stations spanning the reach (Palla as upstream reference, Nizamuddin, ITO bridge, Okhla). Pre- and post-monsoon medians are used for calibration; monsoon medians are held back for evaluation.

3. Methodology

3.1 Process model

A 1D steady-state Streeter–Phelps representation is solved along the reach with explicit mixing nodes at each drain confluence. BOD decays first-order at rate k_d ; DO recovers via reaeration k_a (O'Connor–Dobbins, depth/velocity-driven) and is consumed by BOD oxidation. Both rate constants are temperature-corrected ($\theta_d = 1.047$, $\theta_a = 1.024$) using a season-typical water temperature (28 °C pre-monsoon, 27 °C monsoon, 19 °C post-monsoon). DO saturation $O_s(T)$ follows Benson–Krause; the Delhi elevation correction is < 3% and neglected. Hydraulic depth and velocity scale with Q (Manning-type exponents) so the same kinetic model applies across seasons.

Streeter–Phelps is appropriate at this level of analysis because it captures the first-order BOD–DO dynamics with parameters that are physically interpretable and easy to audit. The assumptions break in three known places — diurnal DO swings, sediment oxygen demand, and the slack-water algal cycles above each barrage — and a full regulatory model would replace the kinetic core with QUAL2K or WASP. §4.1 lists the specific limitations.

3.2 Calibration

Two parameters — $k_{d,20}$ and a multiplier on the O'Connor–Dobbins k_a — are calibrated by a 25×25 grid search on a sum-of-squared-residual objective covering pre- and post-monsoon BOD and DO at the four stations (DO weighted twice). The full grid is reported so a reviewer can see the calibration valley rather than only its minimum. The fit gives $k_{d,20} = 0.46$ 1/d and k_a multiplier = 1.10. Pooled NSE is 0.78 for BOD and 0.95 for DO; pooled PBIAS is +1.6% (BOD) and –0.9% (DO), classified as "Very Good" under Moriasi et al. (2007). On the held-back monsoon medians, the model gives a mean absolute percentage error (MAPE = mean of $|obs - mod|/obs$ across stations, excluding Palla DO because near-saturation values inflate the percentage error) of about 9% on DO and 44% on BOD. The monsoon BOD residual is large because the steady-state 1D structure does not represent in-channel storage above each barrage or the lagged drain loading during the wet season. The dry-season case — which is the design case for compliance planning — is unaffected by this residual, and the DO field is reproduced to within about a tenth of the observed value at all four stations.

3.3 Uncertainty

Five inputs are propagated through a 2000-run Latin Hypercube: headwater Q multiplier (Lognormal, $\sigma = 0.35$); $k_{\{d,20\}}$ (Uniform 0.20–0.55 1/d, widened from the Chapra (1997) urban-stream range to bracket the calibrated optimum at ~ 0.46 1/d); k_a multiplier (Uniform 0.7–1.4); drain BOD multiplier (Lognormal, $\sigma = 0.25$); and drain Q multiplier (Triangular 0.6–1.0–1.3). Priors are kept bounded rather than over-specified, and DO is floored at zero in the Monte Carlo output because negative DO is unphysical. Spearman rank correlations identify the drain BOD multiplier as the dominant driver of Okhla outputs ($\rho = +0.93$ against BOD, -0.56 against DO). The k_a multiplier and $k_{\{d,20\}}$ are the next most leveraged inputs on DO ($\rho = +0.53$ and -0.54). Headwater Q is essentially decoupled from both outputs ($|\rho| < 0.1$), because dry-season Q_{head} (~ 4 m³/s) is an order of magnitude smaller than the combined drain Q (~ 52 m³/s); even doubling Q_{head} barely shifts the mixing ratio. The near-one-for-one sensitivity of Okhla BOD to drain BOD means source-strength uncertainty dominates over kinetic uncertainty in this regime: better data on drain loading is more valuable to the planning question than tighter bounds on k_d or k_a .

3.4 Compliance metric

Three compliance probabilities are reported side by side: $P(\text{DO} \geq 4 \text{ mg/L})$, $P(\text{BOD} \leq 3 \text{ mg/L})$, and the joint $P(\text{both})$. All three are evaluated at the Okhla outflow in the pre-monsoon — the dry-season design state (low Q, warm water) — which is the appropriate reference for a planning-grade compliance probability. Reporting the marginals separately keeps the binding standard visible scenario by scenario, rather than burying it inside a single joint number.

4. Key findings

Finding 1 — $\text{BOD} \leq 3 \text{ mg/L}$ is effectively unreachable on this reach in the dry season under the tested scenarios. Even the combined scenario (S4: $\sim 70\%$ drain treatment plus a tripled Wazirabad release) leaves median Okhla BOD at about 7.7 mg/L, and the BOD-marginal compliance probability is below 1% across all five scenarios. Reaching 3 mg/L would likely require above 90% drain BOD removal, which is not credible at the combined loading scale of Najafgarh, Shahdara and Sahibabad without effectively diverting the entire dry-season municipal load.

Finding 2 — $\text{DO} \geq 4 \text{ mg/L}$ is reachable, and the route is drain treatment rather than flow augmentation. $P(\text{DO} \geq 4)$ at Okhla rises from 4.8% under the baseline to 81.5% under S2 (aggressive STP upgrade alone) and to 86.9% under S4 (S2 plus a tripled Wazirabad release). The flow-augmentation scenario alone (S3) only moves the probability to 7.8%, because dry-season Q_{head} is an order of magnitude smaller than the combined drain inflow and tripling it does not materially change the mixing ratio. The reason BOD remains binding while DO improves is mechanical: DO has a second recovery route through reaeration, whereas BOD is purely load-driven and only responds to source removal.

Finding 3 — The dominant driver of Okhla BOD is the drain BOD multiplier (Spearman $\rho = +0.93$). Headwater Q is decoupled from the output under current dry-season operations ($|\rho| < 0.1$). The kinetic parameters k_d and k_a are required for the model to fit but are second-order from a policy-leverage standpoint. Investment that reduces drain BOD will move Okhla BOD almost one-for-one in this regime; investment in any other lever will not.

S	Description	Med. BOD	Med. DO	P(DO≥4)	P(BOD≤3)	P(joint)
S0	Baseline (current operations)	27.4	1.82	4.8%	0%	0%
S1	Drain interception + STP retrofit ($\eta \approx 50\%$)	13.8	3.91	46.4%	0%	0%
S2	Aggressive STP upgrade ($\eta \approx 70\%$)	8.3	4.77	81.5%	0%	0%
S3	E-flow release at Wazirabad ($\times 3$ head Q)	24.3	2.29	7.8%	0%	0%
S4	S2 + S3 combined	7.7	4.91	86.9%	0%	0%

Table 1. Five-scenario summary at Okhla (pre-monsoon). Median values are run-medians of the 2000-run Latin Hypercube Monte Carlo. $P(DO \geq 4)$ and $P(BOD \leq 3)$ are the marginal probabilities; $P(joint)$ is the probability of meeting both standards simultaneously. The marginals are shown separately so the binding standard (BOD) is visible scenario by scenario.

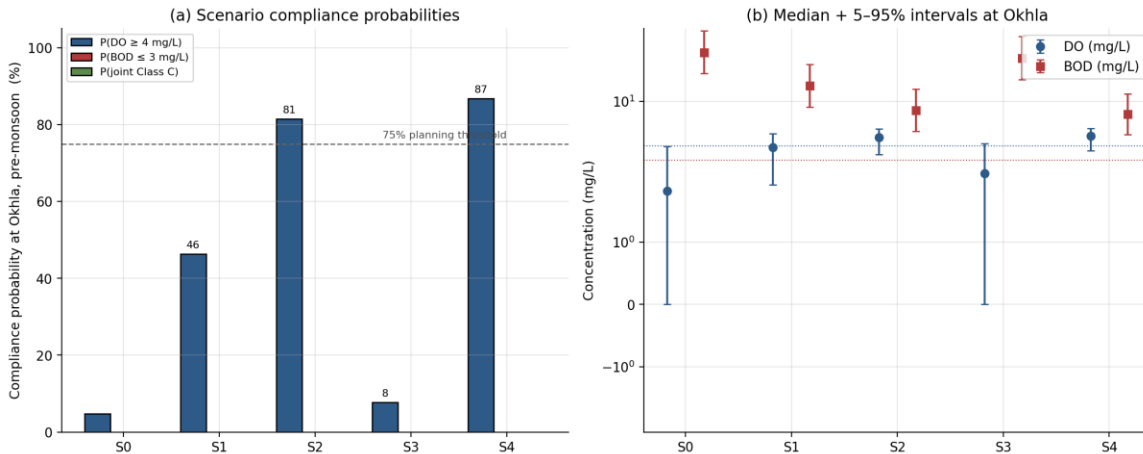


Figure 1. Compliance probability and BOD/DO spread by scenario, generated from the 2000-run Latin Hypercube Monte Carlo around the calibrated Streeter–Phelps model.

4.1 Limitations and assumptions

The model is steady-state 1D and built from public data. It does not resolve diurnal DO swings, sediment oxygen demand, nitrification kinetics, or the slack-water reservoirs above each barrage, where eutrophic algal cycles dominate the DO budget. Drain BOD and Q are treated as time-invariant within each season, although both carry a diurnal and a stormwater signature in practice. The dry-season Wazirabad release is the most contestable single input — CWC and CPCB report different values — but the sensitivity analysis shows it is not a high-leverage input on this reach, because dry-season Q_{head} is an order of magnitude smaller than the combined drain Q. Where the inputs are ambiguous, the conservative value has been used. The model is intended for dry-season planning and should not be used to infer monsoon BOD compliance quantitatively without a storage-aware dynamic model.

5. Relevance and transferability

The same workflow — calibration against public CPCB data, uncertainty propagation by Latin Hypercube, and translation to a compliance probability — applies to other regulated river-reach problems in South Asia: BOD/DO compliance on the Hindon, Sabarmati or Musi, or any reach where municipal wastewater discharge, an upstream barrage and a CPCB Class designation interact. The Streeter–Phelps core can be replaced by QUAL2K or WASP without changing the calibration or Monte Carlo structure, and the same is true if the receptor of interest changes from BOD/DO to ammonia, faecal coliform or instream chlorophyll-a.

Reporting outputs as compliance probabilities rather than as median point estimates also gives the model some use as a staged decision-support tool. Where joint Class C is unreachable in the near term but DO recovery is, the probability framing makes the interim case explicit rather than burying it inside a single pass/fail number.

References

1. Chapra, S.C. (1997). *Surface Water-Quality Modeling*. McGraw-Hill, New York.
2. CPCB (2021). *Yamuna Water Quality Status Report 2018–2020*. Central Pollution Control Board, New Delhi.
3. Emerson, K., Russo, R.C., Lund, R.E., Thurston, R.V. (1975). Aqueous ammonia equilibrium calculations: effect of pH and temperature. *J. Fish. Res. Bd. Canada* 32: 2379–2383.
4. Moriasi, D.N., Arnold, J.G., Van Liew, M.W., et al. (2007). Model evaluation guidelines for systematic quantification of accuracy in watershed simulations. *Trans. ASABE* 50(3): 885–900.
5. O'Connor, D.J. and Dobbins, W.E. (1958). Mechanism of reaeration in natural streams. *Trans. ASCE* 123: 641–684.
6. Streeter, H.W. and Phelps, E.B. (1925). *A Study of the Pollution and Natural Purification of the Ohio River*. US Public Health Service Bulletin 146.
7. USEPA (2013). *Aquatic Life Ambient Water Quality Criteria for Ammonia — Freshwater 2013*. EPA-822-R-13-001.