Supplement

Respiration depth and the conditional diffusive oxygen supply to a Bay of Bengal oxygen minimum zone: a one-dimensional model

This Supplement documents the numerical implementation, the exact budget machinery, the convergence protocol, the forcing preparation, the parameter experiments, the particle-flux attribution method and the station-specific diagnostic limitations. Every number in the supplementary tables is generated by script from the frozen diagnostic archive; none is transcribed by hand.

Contents

S1 Numerical integration and operator sequence **2**

S2 Detrital sinking scheme and benchmark **2**

S3 Exact online oxygen and detritus budgets **3**

S4 Convergence criterion and diagnostic window **4**

S5 Forcing preparation and deep diffusivity **4**

S6 Parameter experiments and identifiability diagnostics **5**

S7 Source-tag particle-flux attribution **5**

S8 Denitrification diagnostic and structural robustness **7**

S9 Station-specific diagnostic limitations **7**

Supplementary tables **9**

Supplementary figures **22**

S1 Numerical integration and operator sequence

The model integrates five biogeochemical tracers on a uniform grid of $N = 200$ levels of thickness $\Delta z = 5$ m spanning 0–1000 m, with cell centres at 2.5–997.5 m and interfaces at 0–1000 m. The physical time step is $\Delta t = 3600$ s and the biological source–sink terms are advanced with 20 sub-steps per physical step. The calendar has 360 days in twelve 30-day months.

Within each physical step the operators are applied in the following order.

1. Forcing interpolation to the current model time.
2. Temperature and salinity: implicit vertical diffusion with relaxation to the prescribed climatology.
3. Nitrate restoring toward the reference climatology.
4. Oxygen restoring toward the reference climatology.
5. Implicit vertical diffusion of NO_3, P, Z and D.
6. Implicit vertical diffusion of O_2.
7. Biology, including detrital sinking and air–sea gas exchange, advanced in sub-steps.

Vertical diffusion is discretised implicitly and solved with the Thomas algorithm for the resulting tridiagonal system, which is unconditionally stable for the diffusive operator and imposes no time-step restriction from K_z.

The biological update uses a Patankar-type formulation in which each compartment's losses are placed in the denominator of an implicit update. In the implementation used here this guarantees positivity for the nitrogen compartments. For oxygen the update is not unconditionally positive by construction, because the respiration term enters the numerator; the model therefore applies a positivity clip and records every clipping event.

Positivity in practice. Instead of an abstract positivity guarantee, we report the measured behaviour: across every converged configuration analysed in this study the oxygen positivity clip was triggered *zero* times, and the minimum annual-mean oxygen concentration encountered was $3.2\,\mathrm{mmol\,m^{-3}}$. The clipping term is retained in the budget identity (Eq. (8) of the main text) and is diagnosed as exactly zero rather than assumed zero.

S2 Detrital sinking scheme and benchmark

Detritus sinks at a prescribed speed w_D, positive downward. Sinking is applied as a separate conservative flux operator, not as a source–sink term inside the biological update. The interface flux uses a total-variation-diminishing discretisation with the van Leer limiter,

$$F_{k+1/2} = w_D \left[D_k + \tfrac{1}{2}\,\phi(r_k)\,(1 - c_k)\,(D_{k+1} - D_k) \right], \tag{S1}$$

where $c_k = w_D \Delta t / \Delta z_k$ is the Courant number, $r_k = (D_k - D_{k-1})/(D_{k+1} - D_k)$ is the smoothness ratio and $\phi(r) = (r + |r|)/(1 + |r|)$ is the van Leer limiter. The surface interface carries no sinking flux and the bottom interface carries $w_D D_N$, which leaves the domain as export.

Consequence for this study. The earlier implementation of this model applied sinking as an in-place sweep, in which the flux entering cell k was computed from the already-updated value in cell $k - 1$. That is an implicit downstream Gauss–Seidel sweep rather than the explicit first-order upwind scheme it was documented as, and its numerical diffusion had not been measured.

An idealised benchmark quantifies the difference. A Gaussian detrital pulse of initial width $\sigma = 20.0$ m centred at 200 m is advected for 20 days at $w_D = 10\ \mathrm{m\,d^{-1}}$ with no biology and no physical diffusion, for which the analytic solution translates without spreading. The results are given in Table S6 and Fig. S5: the legacy sweep broadens the pulse to 37.5 m and explicit upwind to 37.4 m, corresponding to effective numerical diffusivities of 2.906×10^{-4} and 2.881×10^{-4} $\mathrm{m^2\,s^{-1}}$, whereas the van Leer scheme broadens it only to 20.9 m, corresponding to 1.055×10^{-5} $\mathrm{m^2\,s^{-1}}$. All three schemes conserve mass to 10^{-13} or better.

Physical thermocline diffusivity in this model is 1.0×10^{-5} $\mathrm{m^2\,s^{-1}}$ over the analysed depth range. A first-order scheme would therefore impose a spurious vertical spreading of the detrital field an order of magnitude larger than the physical mixing it competes with. Because the experiments in this paper deliberately alter the vertical distribution of particles, controlling that artefact is a prerequisite for interpreting the results, and the final model uses the van Leer scheme throughout.

S3 Exact online oxygen and detritus budgets

Method

Every oxygen source and sink is recorded *online* as the actual increment applied by its operator during the time step, accumulated into a depth-resolved array, and written at each output interval. No budget term is reconstructed after the fact from saved mean state. The distinction matters: an offline re-derivation must re-evaluate nonlinear terms on time-averaged fields, must approximate an implicit diffusion operator explicitly, and must assume a correspondence between the analysis calendar and the model calendar. Each of these introduces error that is invisible unless the budget is checked against the actual inventory change.

The identity verified here is

$$\frac{\mathrm{d}I}{\mathrm{d}t} = P_{\mathrm{photo}} + R_{\mathrm{resp}} + A_{\mathrm{remin}} + F_{\mathrm{AS}} + R_{\mathrm{restore}} + B_{\mathrm{diff}} + C_{\mathrm{clip}}, \tag{S2}$$

where I is the layer oxygen inventory. Both sides are evaluated over the *same* record set: the process terms are summed over records 1–12 of the diagnostic year and the inventory change is taken between records 0 and 12. Mismatching those intervals, for example by differencing annual means while integrating processes over the year, produces an apparent non-closure that is a sampling artefact rather than a physical imbalance.

Diffusive interface fluxes are reconstructed exactly from the recorded cell-wise increments. With G_k the net diffusive gain of cell k and no diffusive flux through the surface,

$$F_j = -\sum_{k<j} G_k, \tag{S3}$$

positive downward, so that $F_{\mathrm{top}} - F_{\mathrm{bot}}$ identically equals the layer integral of the increments. The reconstruction is therefore exact to roundoff by construction.

The detritus budget is assembled from direct online terms,

$$\Delta D = S_P + S_F + S_Z + L_{\mathrm{ox}} + L_{\mathrm{den}} + \nabla \cdot F_{\mathrm{sink}}, \tag{S4}$$

where S_P, S_F and S_Z are phytoplankton mortality, faecal pellet production and zooplankton mortality, L the two remineralisation losses and the final term the sinking divergence. Every term is the amount actually applied by the biological update.

Results

Closure results are given in Tables S4 and S5 and Fig. S4. Across every converged configuration the residual of Eq. (S2) is below $10^{-12}\,\mathrm{mmol\,m^{-2}\,d^{-1}}$. The detritus budget closes to 10^{-16}–10^{-14} relative to its largest individual term; the residual is normalised in this way, rather than by the inventory change, because the latter tends to zero at equilibrium and would give a meaningless ratio.

Why S1 requires the full identity. In the S2–S6 cores the photosynthesis, respiration and air–sea terms are exact zeros, and Eq. (S2) reduces to a three-term balance between demand, diffusion and restoring. At S1 they are not zero, because the oxygen minimum is some 75 m shallower and the core lies closer to the euphotic zone. The reduced form leaves a residual of order $10^{-5}\,\mathrm{mmol\,m^{-2}\,d^{-1}}$ there, against 10^{-13} for the full form: eight orders of magnitude worse (Fig. S4c). All S1 diagnostics in this study therefore use the complete seven-term identity. We note that the omitted terms are small in absolute terms, roughly 0.014 % of core remineralisation, so this is a requirement of exact closure rather than a statement that photosynthesis matters biologically in the S1 core. The biologically consequential S1 feature is in the overlying layer (Sect. S9).

S4 Convergence criterion and diagnostic window

The diagnostic window is set by a formal criterion rather than a fixed integration length. Simulations are integrated in five-year chunks. After each chunk, annual means are formed over complete model years only, and the criterion tests whether the year-on-year change of each monitored quantity remains below tolerance for the final consecutive years of the record. The monitored quantities are the layer-mean oxygen concentration in each budget layer, the depth-integrated total nitrogen inventory, and the depth-integrated inventories of P, Z and D. All must satisfy the criterion simultaneously.

Only complete years are used: a run of exactly N years produces one trailing record that would otherwise be treated as a one-sample year and generate a spurious drift. Second, the criterion requires that the run remains inside tolerance *to the end* of the record, not merely that some interval within it was quiet, so that a run which settles briefly and then drifts again is not declared converged.

Once the criterion is satisfied, the model is integrated for one further complete 360-day model year and all diagnostics are taken from that year. The correct description of this protocol is:

> Simulations were integrated until the convergence criterion was satisfied and diagnosed over the subsequent complete 360-day model year.

This is not a fixed spin-up discard. Convergence years for every configuration are given in Fig. S3a and Table S3; all configurations satisfied the criterion by model year four to six. The magnitude of the core oxygen inventory tendency over the diagnostic year is 10^{-7}–$10^{-5}\,\mathrm{mmol\,m^{-2}\,d^{-1}}$ (Fig. S3b), confirming that storage is negligible, though it is retained explicitly in Eq. (S2) rather than assumed zero.

S5 Forcing preparation and deep diffusivity

Physical forcing is prescribed from three sources, summarised in Table S9 and shown in Fig. S1.

Temperature and salinity are taken from the WOA23 monthly climatology and interpolated to the model grid; the model relaxes toward them on a 30-day timescale. Vertical diffusivity is *not* taken from a reanalysis mixing scheme. It is computed analytically from the mixed-layer depth by $K_z(z) = K_{\mathrm{bg}} + \frac{1}{2}(K_{\mathrm{ml}} - K_{\mathrm{bg}})[1 - \tanh((z - h_{\mathrm{ml}})/h_{\mathrm{tr}})]$ with $K_{\mathrm{ml}} = 10^{-3}\ \mathrm{m^2\,s^{-1}}$, $K_{\mathrm{bg}} = 10^{-5}\ \mathrm{m^2\,s^{-1}}$ and $h_{\mathrm{tr}} =$ 10 m (`scripts/preprocessing/build_model_inputs.py`, function `compute_kz_profile`). The only external input is the monthly mixed-layer depth h_{ml}, taken from the GLORYS12 variable `mlotst`. Surface shortwave radiation and 10 m wind speed are taken from ERA5 monthly climatologies.

Diffusivity over the analysed depth range. Because h_{ml} ranges only from 10.7 to 31.7 m across the six stations and the transition thickness is 10 m, K_z falls to within 5 % of K_{bg} by 60–70 m. Below 150 m it equals $1.000 \times 10^{-5}\ \mathrm{m^2\,s^{-1}}$ to machine precision at every station in every month. The analysed upper layer and core therefore lie entirely inside a single prescribed constant, and F_{top} is that constant multiplied by the modelled oxygen gradient at z_u. Two consequences follow and are stated rather than mitigated. First, the amplitude of every diffusive-flux response scales directly with a number that this framework does not constrain, which is why the main text reports a tested envelope rather than a single percentage. Second, the effective numerical diffusivity of the van Leer sinking scheme, 1.055×10^{-5} $\mathrm{m^2\,s^{-1}}$, is *comparable to* rather than negligible against the prescribed physical value over the analysed range; it is 29 times smaller than the first-order alternative, which is the comparison that motivates the scheme choice, but it is not small in absolute terms.

Reference climatology. WOA23 oxygen and nitrate are used for initialisation, as the restoring target, and as the reference profiles shown in Fig. S2. Because the model is relaxed toward these fields, agreement with them is partly imposed by construction. WOA23 is therefore a climatological reference and cannot serve as independent validation, and no skill statistic is computed from the comparison.

S6 Parameter experiments and identifiability diagnostics

The complete parameter set is listed in Table S1, extracted directly from the model parameter definitions. Only the detrital sinking speed w_D and the reference remineralisation rate r_d are perturbed.

Three families of experiment are used. The fixed-r_d sequences vary w_D alone at $r_d = 0.10$ $\mathrm{d^{-1}}$ and are the only configurations from which a w_D-specific response can be inferred. The sequential isolation path $(5, 0.15) \to (5, 0.10) \to (10, 0.10)$ separates the two steps of the headline weak-to-baseline comparison. The constant-ratio triplet $(5, 0.05)$, $(10, 0.10)$ and $(20, 0.20)$ at S3 tests whether the response is organised by w_D/r_d.

Results are collected in Table S7 and Fig. S6.

First, the fixed-r_d response is non-monotonic at S3 and S4: the highest sampled upper-layer potential demand occurs at $w_D = 10$ $\mathrm{m\,d^{-1}}$, with four sampled speeds bracketing rather than locating a turning point. At S5 only the ascending, saturating limb fall within the tested range, so no turning point is claimed at those stations. The missing $w_D = 8$ $\mathrm{m\,d^{-1}}$ point at S1 and S5 is left as a gap and not interpolated.

Second, along the sequential path the r_d and w_D steps produce comparably sized changes in F_{top}. This is a decomposition along one tested path. Because interaction and order dependence were not evaluated, it does not establish a unique causal attribution, and we do not present one.

Third, the constant-ratio configurations agree to about 1 % in the oxygen minimum, realised core demand and F_{top}, across a fourfold change in each parameter. The natural interpretation is an effective attenuation length scale $L_R \sim w_D/r_{\mathrm{eff}}$. The agreement is close but not exact and varies systematically with w_D, as expected because r_{eff} is temperature dependent and a constant nominal ratio therefore does not give a constant attenuation length at every depth. We describe this as an approximate effective transport-to-remineralisation degeneracy and do not claim that only the ratio matters.

All plotted points are converged model runs. No regression, correlation coefficient or other inferential statistic is fitted to them, since they are deterministic evaluations of a model rather than samples from a population.

S7 Source-tag particle-flux attribution

The additivity problem

To determine where the detritus arriving at the core upper interface was produced, we carry passive tagged detrital tracers labelled by production depth: material present before tagging began, material produced above 50 m, between 50 m and z_u, and below z_u.

The van Leer limiter in Eq. (S1) is a nonlinear function of the local *total* detrital field. If each tagged component computed its own limiter, the tagged fluxes would not sum to the total flux, and any percentage derived from them would be meaningless. The flux is, however, linear in D once the limiter is frozen:

$$F_{k+1/2} = w_D \left[(1-a_k)D_k + a_k D_{k+1}\right], \qquad a_k = \tfrac{1}{2}\phi(r_k)(1-c_k), \tag{S5}$$

so a_k is diagnosed once per sub-step from the total field and reused unchanged for every tag. The remaining operators acting on detritus are already linear: implicit diffusion applies the same tridiagonal matrix to each tag, and the Patankar remineralisation update is a per-cell scalar multiplication.

Negative excursions. Applying the total-field limiter to a minority component can drive that component slightly negative, because the downwind correction term $a_k D_{k+1}$ can remove more of a locally depleted tag than the cell contains, even where the total field remains positive. Clipping the tags would inject unbudgeted tagged mass and destroy additivity. The tags are therefore left unclipped, and the excursion is quantified below. It amounts to 0.4–1.6 % of the band tag's own inventory, and it lies *entirely above 50 m*: it is exactly zero inside the attribution band and exactly zero below z_u in every case (Fig. S7d). It therefore cannot bias the flux at z_u. In addition, every accumulated tagged flux at z_u is non-negative, so the reported shares are a genuine non-negative partition of the annual flux.

The legacy tag falls to 10^{-19}–10^{-17} of the detrital inventory during a one-year tagging spin-up, so the diagnostic year carries no memory of the initial condition.

Additivity and subdomain-closure tests

Two tests were fixed in advance, and both had to be satisfied for the attribution to be reported.

1. **Tagged-flux additivity**: the sum of the tagged sinking fluxes at z_u must equal the total flux, using total-field limiter coefficients.
2. **Sub-domain budget closure**: the detritus budget for the interval between 50 m and z_u must close,

$$\Delta D_{50\to z_u} = S_D + [F_{\text{sink}}(50) - F_{\text{sink}}(z_u)] + [F_{\text{diff}}(50) - F_{\text{diff}}(z_u)] - L_D, \tag{S6}$$

   with the sinking fluxes taken from independently accumulated face fluxes rather than from the divergence, which would make the test circular.

Both were required to close to better than 10^{-10} relative. Both pass by four orders of magnitude in every case (Fig. S7a, b, Table S8).

Detrital diffusion does not vanish locally. The column integral of the detrital diffusive flux is zero to 5×10^{-12}, but the local flux at 50 m is 0.18–1.29 $\text{mmol N m}^{-2}\,\text{yr}^{-1}$ and is 10–100 times larger there than at z_u. The term is therefore retained explicitly in Eq. (S6) rather than neglected on the basis of the column integral.

Result and its interpretation

On the basis of both tests passing, 19–42 % of the modelled core-top sinking flux is attributed to detritus produced between 50 m and z_u under the frozen-total-D-limiter source-tag decomposition, the share rising as particle transfer weakens (Fig. S7c, Table S8).

This is a linear decomposition of a modelled flux, not a Lagrangian particle history. The tags do not follow individual particles and carry no trajectory information; they partition a flux among source regions under a frozen transport operator. The wording used throughout is deliberately *attributed to* rather than *originates from*.

Attenuation coordinate. The temperature-aware Damköhler coordinate $Da = \int r_{\text{eff}}/w_D \, dz$ is reported for two endpoints, from 50 m to the core upper interface ($Da_{50\to z_u}$) and to the core centre ($Da_{50\to z_c}$). These differ by a constant factor of 1.254 at S3 and 1.261 at S4 and must not share a symbol. The pure-decay expectation e^{-Da} underpredicts the measured flux survival by a factor of 1.3–1.8 (Table S8), because it describes only the through-flux and omits detritus produced within the interval. Correcting for the attributed in-band share recovers the measured ratio to 1–7 %, which is a consistency check between two independently computed diagnostics. Da is an ordering coordinate for attenuation, not a predictive law for absolute flux.

S8 Denitrification diagnostic and structural robustness

Denitrification: a formulation diagnostic. The denitrifying remineralisation pathway, Eq. (3) of the main text, is diagnosed online with the same machinery as the oxygen terms, and Table S11 reports the column- and core-integrated denitrifying loss and its share of total remineralisation. Denitrification consumes nitrate rather than oxygen and therefore does not appear in the oxygen identity, Eq. (S2), so none of the oxygen results in this study depends on it.

These numbers are a property of the model formulation and are not used for inference about real-ocean nitrogen loss. Three features of Eq. (3) constrain how far they can be read. The partition is a smooth $1 - f_{O_2}$ modulated by a nitrate limiter that stays between 0.967 and 0.984 over the analysed cores, so the denitrifying share equals $1 - \eta$ to within 0.006 and carries no information beyond the oxygen limitation factor already reported. The formulation has no low-oxygen threshold, so it assigns 16 % of core remineralisation to the denitrifying pathway at S5, where core oxygen is $20.3\,\text{mmol}\,\text{m}^{-3}$, and 22 % at S1 at $14.0\,\text{mmol}\,\text{m}^{-3}$, whereas canonical water-column denitrification is effectively absent at those concentrations. And nitrite is not resolved, so the model cannot represent the quantity that limits N_2 production in this basin. No comparison with observed nitrogen-loss rates is drawn anywhere in this study.

Structural robustness. Four prescribed choices set the core budget most directly and are not constrained within this framework: the vertical diffusivity K_z, the restoring timescale τ_{O_2}, the oxygen half-saturation constant K_{O_2}, and the width of the diagnostic layer. A fifth choice, the vertical extent over which oxygen restoring is applied, alters where a closure acts rather than how strong it is. Tables S10, S13 and S14 report these sensitivities for the weak and baseline configurations at S3 and S4. The K_z, τ_{O_2}, K_{O_2} and restoring-depth cases were re-integrated to convergence and diagnosed over the subsequent complete model year, following exactly the baseline protocol; the layer-width sensitivity was recomputed diagnostically from the frozen converged profiles without re-integration. No model code was changed, since every quantity is an existing model parameter.

The features required to survive were fixed in advance: the sign of the weak-to-baseline F_{top} response, the upper-boundary dominance of resolved diffusive supply, oxygen limitation with $\eta < 1$ declining with rising demand, and the reversal of the restoring term between configurations. They survive every perturbation of K_z, τ_{O_2}, K_{O_2} and layer width, over which the F_{top} response spans -14.3 to -47.3 %. Two results qualify that statement. First, η itself is not robust: varying K_{O_2} over a factor of sixteen moves it from 0.226 to 0.874, which is why the main text treats it as a diagnostic of a prescribed limitation rather than a measured property. Second, and more consequentially, the response does *not* survive excluding climatological oxygen restoring from the upper 100 m: it collapses to -1.2 % at S3 and reverses to $+1.1$ % at S4. Excluding only the upper 50 m leaves it intact at -20.0 and -16.8 %. The mechanism therefore requires that the water overlying the core not deoxygenate with it, and the headline percentages are specific to the baseline closure configuration.

S9 Station-specific diagnostic limitations

Three limitations affect how the standard diagnostics should be read at individual stations (Fig. S8).

S1: the 1000 m domain is not supported by the forcing. The S1 sub-threshold layer cannot be given a lower bound, because the forcing does not resolve the depths over which it would be measured. Two independent products place the seafloor at S1 (20.5° N, 89.5° E, on the Ganges–Brahmaputra shelf and upper slope) far above the model bottom: GLORYS12 temperature has no valid level below 77.9 m and WOA23 oxygen none below 125 m. The preprocessing step fills below the deepest valid level by constant downward extrapolation, so the S1 oxygen restoring target is a constant 16.18 mmol m^{-3} from 127.5 m to 1000 m and temperature is constant below 77.5 m. The apparent 885 m sub-threshold layer is that fill value lying below the 20 mmol m^{-3} threshold; it is a property of the extrapolation, not of the oxygen field.

S1 is therefore retained only as a forcing-limited contextual and numerical-diagnostic case. Its values appear below solely to demonstrate exact budget closure and to show what an unsupported, constant-filled column produces; none of them enters the regional inference of the main text. No S1 OMZ thickness is reported anywhere, no S1 quantity below 127.5 m is interpreted, and the common 1000 m domain is physically supported at S2–S6 only.

S1: euphotic upper layer. The geometric rule that defines the upper layer as the 50 m immediately above the core places it at 47.5–92.5 m at S1, inside the productive zone. These S1 values are forcing-limited numerical diagnostics only and are excluded from the quantitative synthesis. Its photosynthetic term, +2.74 mmol m^{-2} d^{-1} under the baseline configuration, is comparable in magnitude to its remineralisation term, −4.07 mmol m^{-2} d^{-1}, whereas the same term is an exact zero in the S2–S6 upper layers (Fig. S8c). The geometrically defined upper layer is therefore not the same biogeochemical object at every station, and the S1 upper-layer budget must not be read as a purely respiratory one. This is a limitation of the layer definition rather than of the mechanism, and it is the reason the upper-layer response at S1 is not interpreted as respiratory demand alone.

S5: threshold proximity. The S5 oxygen minimum spans only 19.2–21.2 mmol m^{-3} across the entire experimental sequence while straddling the 20 mmol m^{-3} level. Whether a threshold-based OMZ is diagnosed there therefore changes with configuration, but the underlying oxygen balance changes very little. S5 is best treated as a weak-OMZ end member, and the threshold crossing should not be read as a regime change.

Lower-boundary contribution. Across all configurations the lower core interface contributes between −1.1 and +25.9 % of the net resolved diffusive supply (Fig. S8d). The negative value occurs at S4 in the weak configuration, where the lower interface is a small net sink rather than a source; this is why the upper-boundary share exceeds 100 % in that case and why the share is not presented as a stacked fraction bounded by 100 %.

Supplementary tables

Table S1. Complete V2 parameter set, part 1: numerical and biological parameters, extracted directly from the model parameter definitions rather than transcribed. The basis column indicates the origin of each value using three tags. *Model lineage*: carried unchanged from the parent NPZD formulation. *Prescribed here*: set for this study, with the tested range indicated where one was explored. *Legacy, unattributed*: present in the legacy Fortran source with no attribution and not traceable to a specific publication. No literature source is claimed for any value in these tables. Particle, stoichiometric, oxygen and restoring parameters are in Table S2.

Symbol	Description	Value	Unit	Basis	Perturbed
Δt	physical time step	3600	s	prescribed here; tested (Table S17)	no
n_{sub}	biological sub-steps per step	20	–	prescribed here	no
$\tau_{T,S}$	T/S nudging timescale	30	d	prescribed here	no
T_{ref}	reference temperature	20	°C	model lineage	no
$V_{\max}$	maximum phytoplankton growth rate	1.5	d^{-1}	legacy, unattributed	no
α	initial slope of P–I curve	0.03	$(W\ m^{-2})^{-1}\ d^{-1}$	legacy, unattributed	no
k_{NO_3}	half-saturation for NO_3 uptake	0.5	mmol m^{-3}	legacy, unattributed	no
k_w	seawater light attenuation	0.04	m^{-1}	legacy, unattributed	no
k_p	self-shading coefficient	0.03	m^2 (mmol N)$^{-1}$	legacy, unattributed	no
m_P	phytoplankton mortality	0.08	d^{-1}	legacy, unattributed	no
β_P	phytoplankton exudation	0.05	d^{-1}	legacy, unattributed	no
θ_P	temperature coefficient, phytoplankton	0.063	$°C^{-1}$	model lineage; $Q_{10} \approx 1.9$	no
$g_{\max}$	maximum grazing rate	1	d^{-1}	legacy, unattributed	no
k_P	half-saturation for grazing	0.5	mmol m^{-3}	legacy, unattributed	no
γ_Z	assimilation efficiency	0.7	–	legacy, unattributed	no
m_{Z1}	linear zooplankton mortality	0.05	d^{-1}	legacy, unattributed	no
m_{Z2}	quadratic zooplankton mortality	0.2	$(mmol\ m^{-3})^{-1}\ d^{-1}$	legacy, unattributed	no
β_Z	zooplankton excretion	0.05	d^{-1}	legacy, unattributed	no
θ_Z	temperature coefficient, zooplankton	0.063	$°C^{-1}$	model lineage; $Q_{10} \approx 1.9$	no

Table S2. Complete V2 parameter set, part 2: particle, stoichiometric, oxygen and restoring parameters. Tags as in Table S1. Only w_D and r_d are varied in the primary particle-transport experiment matrix, and the range over which they are varied is given in the basis column; K_z, τ_{O_2}, K_{O_2}, the restoring depth and the diagnostic layer width are altered only in the dedicated robustness experiments of Sect. S8. The K_z sensitivity multiplier is applied to the whole column in the experiments that use it (Table S10) and is unity in every other run, so no depth threshold for it is listed here. The value $r_{\mathrm{denit}} = 1.04$ is inherited from the legacy source and its stoichiometric derivation could not be reconstructed; it affects the diagnosed nitrate sink only and enters no oxygen result reported here.

Symbol	Description	Value	Unit	Basis	Perturbed
w_D	detrital sinking speed	10	m d^{-1}	prescribed here; perturbed 5–20	**yes**
r_d	reference remineralisation rate	0.1	d^{-1}	prescribed here; perturbed 0.05–0.20	**yes**
θ_D	temperature coefficient, remineralisation	0.063	°C^{-1}	model lineage; $Q_{10} \approx 1.9$	no
r_{op}	O_2 produced per N assimilated	8.625	mol O_2 (mol N)$^{-1}$	model lineage; Redfield 138/16	no
r_{on}	O_2 consumed per N remineralised	8.625	mol O_2 (mol N)$^{-1}$	model lineage; Redfield 138/16	no
K_{O_2}	half-saturation for oxic remineralisation	4	mmol m^{-3}	prescribed here; tested (Table S13)	no
r_{dn}	denitrification rate scaling	1	–	prescribed here; unity	no
r_{denit}	NO_3 consumed per N denitrified	1.04	mol N (mol N)$^{-1}$	legacy, unattributed	no
s_{O_2}	piston-velocity scale factor	1	–	prescribed here; unity, no tuning	no
$\tau_{\mathrm{NO_3}}$	NO_3 restoring timescale	60	d	prescribed here (closure)	no
τ_{O_2}	O_2 restoring timescale	365	d	prescribed here; tested (Table S10)	no
z_{rest}	shallowest restored depth	0	m	prescribed here; tested (Table S13)	no

Table S3. Complete station × case diagnostic dataset for every configuration integrated in this study, generated by script from the frozen diagnostic archive; no value is transcribed by hand. *conv* is the model year at which the automatic convergence criterion was first satisfied on all monitored variables; every case shown converged. *resid* is the residual of the full seven-term oxygen identity over the diagnostic year. S1 is included for completeness and is excluded from every quantitative result (Sect. 2 of the main text). Units mmol O_2 m^{-2} d^{-1} except O_2 (mmol m^{-3}), z_{min} (m) and the share (%).

Stn	Case	conv	O_2min	z_{min}	core O_2	A_{pot}	A_{act}	η	F_{top}	share	resid
S2	weak	4	3.90	222.5	4.38	0.0990	0.0533	0.538	0.08863	91.7	+1.8e-13
	baseline	4	2.73	207.5	3.00	0.9854	0.4215	0.428	0.05555	82.4	+1.8e-13
S3	weak	5	4.60	232.5	5.02	0.0583	0.0335	0.574	0.09952	97.7	+6.3e-15
	isolation	5	4.33	227.5	4.66	0.2275	0.1246	0.548	0.08472	95.6	+1.4e-14
	near-neutral	5	4.42	227.5	4.77	0.1744	0.0968	0.555	0.08807	96.2	+4.6e-15
	interm.-high	5	3.78	222.5	4.06	0.5688	0.2893	0.509	0.07356	92.5	+8.9e-15
	baseline	5	3.54	222.5	3.81	0.7291	0.3588	0.492	0.07216	91.9	-1.2e-14
	high	4	3.25	222.5	3.54	0.9417	0.4452	0.473	0.08002	93.7	+1.1e-14
	ratio-low	5	3.54	222.5	3.82	0.7244	0.3568	0.493	0.07235	92.0	+2.1e-14
	ratio-high	5	3.53	222.5	3.81	0.7329	0.3604	0.492	0.07197	91.9	+7.8e-15
S4	weak	5	5.68	362.5	6.63	0.0631	0.0402	0.638	0.12095	101.1	-6.5e-15
	isolation	5	5.41	222.5	6.16	0.2609	0.1598	0.613	0.10237	99.1	-3.4e-15
	near-neutral	5	5.54	222.5	6.30	0.1976	0.1224	0.620	0.10664	99.8	+2.6e-14
	interm.-high	4	4.63	222.5	5.32	0.6835	0.3913	0.573	0.08797	95.6	-1.2e-15
	baseline	4	4.29	222.5	4.96	0.8888	0.4931	0.555	0.08607	94.7	+5.6e-15
	high	5	3.84	222.5	4.52	1.1788	0.6274	0.532	0.09546	96.2	-4.9e-15
S5	weak	5	21.19	242.5	21.92	0.0438	0.0372	0.849	0.15585	91.8	-3.7e-14
	isolation	5	20.85	237.5	21.50	0.1749	0.1478	0.845	0.14471	90.1	+3.9e-14
	baseline	5	19.72	232.5	20.32	0.5778	0.4831	0.836	0.13645	87.2	-5.6e-15
	high	5	19.19	232.5	19.83	0.7607	0.6332	0.832	0.14497	88.3	+3.4e-14
S6	weak	6	15.74	222.5	18.67	0.0606	0.0500	0.825	0.15411	77.9	+7.0e-14
	baseline	6	14.37	222.5	17.24	0.5588	0.4527	0.810	0.13979	74.1	+4.4e-14
S1	weak	5	15.99	177.5	19.70	0.7810	0.6698	0.858	1.13267	99.8	-2.7e-12
	isolation	5	14.18	152.5	17.34	1.6827	1.3912	0.827	1.15720	98.5	+3.0e-13
	baseline	5	10.32	147.5	13.97	3.2482	2.5185	0.775	1.29595	97.1	+3.3e-13
	high	5	9.56	162.5	14.08	3.4093	2.6010	0.763	1.42302	98.2	+1.9e-13

Table S4. Exact online oxygen budget closure over the fixed core for every frozen configuration. All terms are annual means in mmol O_2 m^{-2} d^{-1}, integrated from the recorded per-operator increments over the same records as the inventory change. Terms are printed to six decimals so that the closure claim can be checked by summing the printed row; the residual column gives the machine-precision value. The identity is dI/dt = photo + resp + remin + air–sea + restore + diff + clip; the residual is the difference between the two sides. The air–sea and clipping columns are identically zero in every core listed, because no core reaches the surface and no positivity clipping occurred, and they are omitted from the display only. Photosynthesis and respiration are exact zeros in the S2–S6 cores, which all lie below the euphotic zone, but not at S1, whose 105 m core reaches into it.

Stn	Case	photo	resp	remin	restore	diff	dI/dt	residual
S1	weak	+0.000165	-0.000179	-0.669758	-0.465674	+1.135448	+1.582e-06	-2.7e-12
S1	isolation	+0.000168	-0.000182	-1.391229	+0.215885	+1.175354	-4.079e-06	+3.0e-13
S1	baseline	+0.000173	-0.000187	-2.518460	+1.183237	+1.335231	-6.134e-06	+3.3e-13
S1	high	+0.000175	-0.000189	-2.601018	+1.151376	+1.449649	-6.120e-06	+1.9e-13
S2	weak	+0.000000	-0.000000	-0.053302	-0.043337	+0.096651	+1.216e-05	+1.8e-13
S2	baseline	+0.000000	-0.000000	-0.421507	+0.354101	+0.067406	+5.103e-07	+1.8e-13
S3	weak	+0.000000	-0.000000	-0.033487	-0.068344	+0.101841	+1.027e-05	+6.4e-15
S3	isolation	+0.000000	-0.000000	-0.124597	+0.035973	+0.088628	+3.439e-06	+1.4e-14
S3	near-neutral	+0.000000	-0.000000	-0.096786	+0.005275	+0.091516	+4.953e-06	+4.6e-15
S3	intermediate-high	+0.000000	-0.000000	-0.289290	+0.209765	+0.079524	-7.181e-07	+8.9e-15
S3	baseline	+0.000000	-0.000000	-0.358839	+0.280337	+0.078502	-1.104e-06	-1.2e-14
S3	high	+0.000000	-0.000000	-0.445241	+0.359862	+0.085378	-4.829e-07	+1.1e-14
S4	weak	+0.000000	-0.000000	-0.040207	-0.079451	+0.119680	+2.189e-05	-6.5e-15
S4	isolation	+0.000000	-0.000000	-0.159794	+0.056552	+0.103253	+1.172e-05	-3.4e-15
S4	near-neutral	+0.000000	-0.000000	-0.122433	+0.015575	+0.106872	+1.417e-05	+2.6e-14
S4	intermediate-high	+0.000000	-0.000000	-0.391327	+0.299302	+0.092029	+3.269e-06	-1.4e-15
S4	baseline	+0.000000	-0.000000	-0.493057	+0.402173	+0.090886	+1.743e-06	+5.7e-15
S4	high	+0.000000	-0.000000	-0.627398	+0.528127	+0.099272	+1.770e-06	-4.8e-15
S5	weak	+0.000000	-0.000000	-0.037167	-0.132588	+0.169755	+4.364e-07	-3.7e-14
S5	isolation	+0.000000	-0.000000	-0.147791	-0.012848	+0.160634	-5.449e-06	+3.9e-14
S5	baseline	+0.000000	-0.000000	-0.483143	+0.326654	+0.156474	-1.534e-05	-5.6e-15
S5	high	+0.000000	-0.000000	-0.633207	+0.468923	+0.164266	-1.710e-05	+3.4e-14
S6	weak	+0.000000	-0.000000	-0.050003	-0.147942	+0.197900	-4.508e-05	+7.0e-14
S6	baseline	+0.000000	-0.000000	-0.452660	+0.264032	+0.188578	-5.050e-05	+4.4e-14

Table S5. Direct detritus budget closure over the water column for every frozen configuration, assembled from online source and loss terms. S2 and S6 were integrated in the weak and baseline configurations only. Units $\mathrm{mmol\,N\,m^{-2}\,yr^{-1}}$. The relative residual is normalised by the largest individual term because the inventory change tends to zero at equilibrium; that column is identically zero and is omitted from the display only.

Stn	Case	srcP	srcF	srcZ	lossR	lossN	sinkdiv	residual	rel.
S1	weak	+122.716	+201.810	+305.639	-612.155	-18.010	+0.000	+2.1e-12	+3.5e-15
S1	isolation	+121.256	+182.537	+272.799	-551.183	-25.409	-0.000	-2.2e-12	-4.0e-15
S1	baseline	+118.975	+155.673	+227.959	-454.648	-47.958	-0.001	+4.5e-12	+9.9e-15
S1	high	+117.516	+140.530	+203.174	-392.082	-68.631	-0.507	+5.2e-12	+1.3e-14
S2	weak	+117.613	+166.383	+243.153	-510.076	-17.073	-0.000	-1.2e-12	-2.3e-15
S2	baseline	+113.990	+123.285	+173.146	-344.281	-65.220	-0.919	-6.7e-14	-1.9e-16
S3	weak	+111.003	+145.331	+209.556	-453.600	-12.289	-0.000	-3.7e-12	-8.1e-15
S3	isolation	+109.794	+128.270	+182.517	-400.655	-19.923	-0.002	+1.2e-13	+3.0e-16
S3	near-neutral	+110.075	+132.113	+188.561	-413.167	-17.581	-0.001	-9.4e-13	-2.3e-15
S3	intermediate-high	+108.430	+110.812	+155.428	-338.348	-36.146	-0.175	-1.6e-12	-4.8e-15
S3	baseline	+107.877	+104.155	+145.225	-311.304	-45.183	-0.770	+2.2e-13	+7.1e-16
S3	high	+106.621	+89.912	+123.645	-241.719	-63.617	-14.843	+2.9e-12	+1.2e-14
S4	weak	+114.842	+186.628	+276.200	-562.398	-15.273	-0.000	+1.7e-12	+3.0e-15
S4	isolation	+113.558	+167.188	+243.774	-500.936	-23.582	-0.002	+6.4e-12	+1.3e-14
S4	near-neutral	+113.859	+171.627	+251.123	-515.586	-21.021	-0.001	+9.0e-13	+1.7e-15
S4	intermediate-high	+112.103	+146.608	+210.152	-427.012	-41.653	-0.197	+4.0e-13	+9.4e-16
S4	baseline	+111.515	+138.630	+197.306	-394.639	-51.926	-0.886	+1.7e-12	+4.2e-15
S4	high	+110.204	+121.391	+169.944	-310.250	-73.406	-17.883	+3.7e-12	+1.2e-14
S5	weak	+109.757	+121.817	+170.189	-394.048	-7.714	-0.000	+1.2e-13	+3.0e-16
S5	isolation	+108.720	+105.695	+145.227	-349.932	-9.707	-0.002	+1.0e-12	+3.0e-15
S5	baseline	+106.994	+83.450	+111.796	-285.555	-16.022	-0.663	+2.7e-12	+9.6e-15
S5	high	+105.722	+70.982	+93.602	-236.762	-20.907	-12.636	+1.7e-12	+7.3e-15
S6	weak	+104.515	+86.571	+115.873	-299.928	-7.031	-0.000	+8.0e-13	+2.7e-15
S6	baseline	+101.386	+61.004	+79.210	-225.817	-15.144	-0.639	+2.0e-12	+8.8e-15

Table S6. Gaussian-pulse sinking benchmark: pure advection at $w_D = 10$ m d^{-1} for 20 days with no biology or physical diffusion, initial $\sigma = 20.0$ m. K_{num} is the effective numerical diffusivity inferred from variance growth. The prescribed physical diffusivity over the analysed depth range is 1.0×10^{-5} m^2 s^{-1}.

Scheme	final σ (m)	spreading (m)	K_{num} (m^2 s^{-1})	mass error
legacy	37.472	+17.472	2.906e-04	1.45e-13
upwind	37.361	+17.361	2.881e-04	-1.13e-15
TVD	20.892	+0.892	1.055e-05	-5.67e-16

Table S7. Parameter-identifiability diagnostics. (a) Fixed-r_d sequences ($r_d = 0.10$ d^{-1}). Normalised values are relative to each station's own $w_D = 5$ m d^{-1} case. S1 and S5 were not run at $w_D = 8$ m d^{-1}.

Station	w_D (m d^{-1})	$A_{\mathrm{up,pot}}$	normalised	F_{top}	normalised
S1	5	5.1442	1.0000	1.15720	1.0000
S1	10	4.2552	0.8272	1.29595	1.1199
S1	20	2.7653	0.5375	1.42302	1.2297
S3	5	0.4493	1.0000	0.08472	1.0000
S3	8	0.7454	1.6589	0.07356	0.8683
S3	10	0.8307	1.8488	0.07216	0.8518
S3	20	0.8076	1.7974	0.08002	0.9446
S4	5	0.5317	1.0000	0.10237	1.0000
S4	8	0.9135	1.7181	0.08797	0.8593
S4	10	1.0290	1.9355	0.08607	0.8407
S4	20	1.0202	1.9189	0.09546	0.9325
S5	5	0.3092	1.0000	0.14471	1.0000
S5	10	0.6088	1.9687	0.13645	0.9429
S5	20	0.6118	1.9785	0.14497	1.0018

(b) Sequential parameter-isolation path. These are responses along one tested path and do not constitute a unique causal decomposition, since interaction and order dependence were not evaluated.

Station	weak (5, 0.15)	(5, 0.10)	baseline (10, 0.10)	r_d step (%)	w_D step (%)
S3	0.09952	0.08472	0.07216	-14.9	-14.8
S4	0.12095	0.10237	0.08607	-15.4	-15.9

(c) Constant $w_D/r_d = 100$ configurations at S3, normalised to the baseline. The spread quantifies an approximate effective transport-to-remineralisation degeneracy; it is not exact.

Quantity	(5, 0.05)	(10, 0.10)	(20, 0.20)	norm.	norm.	norm.	spread (%)
F_{top}	0.07235	0.07216	0.07197	1.0027	1.0000	0.9974	0.53
$A_{\mathrm{core,act}}$	0.35684	0.35884	0.36044	0.9944	1.0000	1.0045	1.01
O_2min	3.54479	3.53795	3.53246	1.0019	1.0000	0.9984	0.35

Table S8. Particle attenuation and core-top flux attribution. Da is the temperature-aware attenuation coordinate $\int r_{\mathrm{eff}}/w_D\,\mathrm{d}z$ integrated from 50 m to the core upper interface (z_u) or to the core centre (z_c); the two endpoints give different values and must not share a symbol. The attributed share is the fraction of the annual core-top sinking flux assigned to detritus produced between 50 m and z_u under the frozen-total-D-limiter source-tag decomposition; it is a linear flux decomposition and not a Lagrangian particle history. Both additivity and subdomain-closure tests were fixed in advance and required residuals below 10^{-10}.

Stn	Case	$Da_{50\to z_u}$	$Da_{50\to z_c}$	$F(z_u)/F(50)$	e^{-Da}	attributed (%)	additivity	sub-domain
S3	weak	4.3007	5.3926	0.0249	0.0136	42.37	2.3e-14	7.6e-15
S3	near-neutral	3.1538	3.9546	0.0671	0.0427	34.00	2.2e-14	1.3e-14
S3	intermediate-high	1.7919	2.2469	0.2274	0.1666	25.54	3.1e-14	1.5e-14
S3	baseline	1.4336	1.7975	0.3161	0.2385	23.63	2.9e-14	2.7e-15
S4	weak	4.3282	5.4587	0.0225	0.0132	36.95	3.4e-14	1.1e-14
S4	near-neutral	3.1740	4.0030	0.0617	0.0418	28.79	3.3e-14	1.4e-16
S4	intermediate-high	1.8034	2.2744	0.2139	0.1647	20.86	2.3e-14	6.9e-15
S4	baseline	1.4427	1.8196	0.2987	0.2363	19.15	2.9e-14	4.7e-15

Table S9. Forcing and reference data. WOA23 is used for initialisation and as the restoring target and is therefore a climatological reference, *not* an independent validation dataset.

Source	Variables	Model use	Interpolation	Known limitation
WOA23	T, S, O_2, NO_3 climatology	initial conditions, restoring target, reference profiles in Fig. S2	monthly climatology linearly interpolated to the model calendar and grid	not independent of the model solution; cannot serve as validation
GLORYS12	mixed-layer depth `mlotst`, the sole input to the analytic K_z profile of Eq. (6) of the main text; no reanalysis diffusivity field is used	sets the depth of the analytic mixing transition	monthly climatology, 10.7–31.7 m across the six stations	K_z equals the background value 10^{-5} m^2 s^{-1} at every depth below 60–70 m, so the analysed layers sit inside a single prescribed constant
ERA5	10 m wind speed, shortwave radiation	air–sea gas exchange and light forcing	monthly climatology	climatological forcing only; no interannual variability is represented

Table S10. Structural robustness of the headline mechanism to the two prescribed physical choices that most directly set the core budget: the vertical diffusivity and the restoring timescale. K_z is scaled over the whole column; τ_{O_2} is varied about its baseline of 365 d. Every configuration was re-integrated to convergence and diagnosed over the subsequent complete model year, exactly as in the baseline experiments. The test is not whether the numbers reproduce but whether the four structural features survive: the sign of the weak-to-baseline F_{top} response, the upper-boundary dominance of resolved diffusive supply, oxygen limitation with $\eta < 1$, and the reversal of the restoring term between the weak and baseline configurations. All four survive in all eight perturbed cases. F_{top} and R_{restore} in $\text{mmol}\,O_2\,m^{-2}\,d^{-1}$.

Stn	Physics	F_{top} weak	F_{top} baseline	ΔF_{top} (%)	η baseline	share (%)	R_{restore} baseline
S3	baseline	0.09952	0.07216	-27.5	0.492	91.9	+0.2803
S3	$K_z \times 0.5$	0.03920	0.02821	-28.0	0.491	90.7	+0.2783
S3	$K_z \times 2$	0.30889	0.24209	-21.6	0.514	95.3	+0.1990
S3	$\tau_{O_2} = 180$ d	0.08439	0.07034	-16.7	0.521	94.2	+0.3055
S3	$\tau_{O_2} = 730$ d	0.13473	0.08566	-36.4	0.447	91.0	+0.2317
S4	baseline	0.12095	0.08607	-28.8	0.555	94.7	+0.4022
S4	$K_z \times 0.5$	0.05697	0.04146	-27.2	0.554	96.4	+0.3954
S4	$K_z \times 2$	0.32011	0.24181	-24.5	0.570	96.0	+0.3338
S4	$\tau_{O_2} = 180$ d	0.12166	0.10293	-15.4	0.587	99.4	+0.4179
S4	$\tau_{O_2} = 730$ d	0.13632	0.07646	-43.9	0.493	89.4	+0.3532

Table S11. Denitrification response to particle transport. The denitrifying loss is the detritus remineralised along the denitrifying pathway, diagnosed online, and the fraction is that loss as a percentage of total (oxic plus denitrifying) remineralisation. Column values integrate the full 0–1000 m domain; core values integrate the fixed baseline-defined core layer. The core-layer partition requires the depth-resolved detrital loss terms, which were not archived for the S2 and S6 integrations, so only full-column values are given there and the core columns are left blank rather than estimated. Denitrifying loss in $\mathrm{mmol\,N\,m^{-2}\,yr^{-1}}$, oxygen in $\mathrm{mmol\,m^{-3}}$.

				Full column		Core layer	
Station	Case	O_2min	Core O_2	denitrifying	fraction (%)	denitrifying	fraction (%)
S1	weak	15.99	19.70	18.01	2.9	4.49	13.8
S1	baseline	10.32	13.97	47.96	9.5	29.46	21.9
S1	high	9.56	14.08	68.63	14.9	32.61	23.1
S2	weak	3.90	4.38	17.07	3.2	–	–
S2	baseline	2.73	3.00	65.22	15.9	–	–
S3	weak	4.60	5.02	12.29	2.6	1.02	42.2
S3	baseline	3.54	3.81	45.18	12.7	15.18	50.3
S3	high	3.25	3.54	63.62	20.8	20.36	52.3
S4	weak	5.68	6.63	15.27	2.6	0.93	35.8
S4	baseline	4.29	4.96	51.93	11.6	16.18	44.0
S4	high	3.84	4.52	73.41	19.1	22.55	46.3
S5	weak	21.19	21.92	7.71	1.9	0.27	14.9
S5	baseline	19.72	20.32	16.02	5.3	3.89	16.2
S5	high	19.19	19.83	20.91	8.1	5.23	16.5
S6	weak	15.74	18.67	7.03	2.3	–	–
S6	baseline	14.37	17.24	15.14	6.3	–	–

Table S12. Annual-mean oxygen offset of each configuration from the WOA23 restoring target, integrated over the baseline-defined core layer. The offset is derived two ways: $R_{\text{restore}}\,\tau_{O_2}/h_{\text{core}}$ from the recorded restoring increment, and direct integration of target minus model over the same layer. The two agree to three significant figures throughout, which is an internal consistency check on the restoring implementation. Negative bias means the model core is poorer in oxygen than the target. The baseline configuration is not the closest to the climatology at any station; it was not selected by fit. Dashes mark configurations not run at that station.

Station	Configuration	R_{restore}	offset (restoring)	offset (direct)	model	bias
		$\text{mmol}\,\text{m}^{-2}\,\text{d}^{-1}$	$\text{mmol}\,\text{m}^{-3}$	$\text{mmol}\,\text{m}^{-3}$	$\text{mmol}\,\text{m}^{-3}$	%
S1	isolation (5, 0.10)	+0.2159	+0.750	+0.751	17.335	-4.2
	weak (5, 0.15)	-0.4657	-1.619	-1.619	19.705	+8.9
	baseline (10, 0.10)	+1.1832	+4.113	+4.113	13.973	-22.7
	high (20, 0.10)	+1.1514	+4.002	+4.002	14.084	-22.1
WOA23 core-layer target at S1: 18.086 $\text{mmol}\,\text{m}^{-3}$						
S3	isolation (5, 0.10)	+0.0360	+0.125	+0.125	4.662	-2.6
	weak (5, 0.15)	-0.0683	-0.238	-0.238	5.024	+5.0
	near-neutral (5, 0.11)	+0.0053	+0.018	+0.018	4.768	-0.4
	intermediate-high (8, 0.10)	+0.2098	+0.729	+0.729	4.057	-15.2
	baseline (10, 0.10)	+0.2803	+0.975	+0.974	3.812	-20.4
	high (20, 0.10)	+0.3599	+1.251	+1.251	3.536	-26.1
WOA23 core-layer target at S3: 4.787 $\text{mmol}\,\text{m}^{-3}$						
S4	isolation (5, 0.10)	+0.0566	+0.197	+0.197	6.159	-3.1
	weak (5, 0.15)	-0.0795	-0.276	-0.276	6.632	+4.3
	near-neutral (5, 0.11)	+0.0156	+0.054	+0.054	6.302	-0.9
	intermediate-high (8, 0.10)	+0.2993	+1.040	+1.040	5.316	-16.4
	baseline (10, 0.10)	+0.4022	+1.398	+1.398	4.958	-22.0
	high (20, 0.10)	+0.5281	+1.836	+1.836	4.520	-28.9
WOA23 core-layer target at S4: 6.356 $\text{mmol}\,\text{m}^{-3}$						
S5	isolation (5, 0.10)	-0.0128	-0.045	-0.045	21.502	+0.2
	weak (5, 0.15)	-0.1326	-0.461	-0.461	21.918	+2.1
	baseline (10, 0.10)	+0.3267	+1.136	+1.136	20.322	-5.3
	high (20, 0.10)	+0.4689	+1.630	+1.630	19.828	-7.6
WOA23 core-layer target at S5: 21.457 $\text{mmol}\,\text{m}^{-3}$						
S2	weak (5, 0.15)	-0.0433	-0.151	-0.151	4.385	+3.6
	baseline (10, 0.10)	+0.3541	+1.231	+1.231	3.003	-29.1
WOA23 core-layer target at S2: 4.234 $\text{mmol}\,\text{m}^{-3}$						
S6	weak (5, 0.15)	-0.1479	-0.514	-0.514	18.674	+2.8
	baseline (10, 0.10)	+0.2640	+0.918	+0.918	17.242	-5.1
WOA23 core-layer target at S6: 18.159 $\text{mmol}\,\text{m}^{-3}$						

Table S13. Sensitivity of the weak-to-baseline response to the oxygen half-saturation constant K_{O_2} and to the vertical extent over which oxygen restoring is applied (z_{rest}; restoring acts below this depth). All runs use $\tau_{O_2} = 365$ d, were integrated to the same automatic convergence criterion and diagnosed over the subsequent complete model year. The $K_{O_2} = 4$ / $z_{\text{rest}} = 0$ setting is the baseline and reproduces the frozen archive exactly. ΔF_{top} is the fractional change from the weak to the baseline particle configuration. Budget residuals are 5×10^{-16} to 2.7×10^{-14} $\text{mmol}\,\text{m}^{-2}\,\text{d}^{-1}$ throughout.

Stn	Perturbation	η weak	η base	F_{top} weak	F_{top} base	ΔF_{top}	share	R_{restore}
				$\text{mmol}\,\text{m}^{-2}\,\text{d}^{-1}$		%	%	sign reversal
S3	$K_{O_2} = 1$	0.841	0.759	0.09715	0.06706	-31.0	88.0–97.4	yes
	$K_{O_2} = 4$ (baseline)	0.574	0.492	0.09952	0.07216	-27.5	91.9–97.7	yes
	$K_{O_2} = 16$	0.257	0.226	0.10518	0.08653	-17.7	95.8–98.1	yes
	$z_{\text{rest}} = 50$ m	0.573	0.504	0.09809	0.07846	-20.0	93.1–97.7	yes
	$z_{\text{rest}} = 100$ m	0.571	0.519	0.09740	0.09623	-1.2	95.0–97.5	yes
S4	$K_{O_2} = 1$	0.874	0.805	0.11839	0.08002	-32.4	90.8–100.9	yes
	$K_{O_2} = 4$ (baseline)	0.638	0.555	0.12095	0.08607	-28.8	94.7–101.1	yes
	$K_{O_2} = 16$	0.311	0.273	0.12727	0.10285	-19.2	98.7–101.3	yes
	$z_{\text{rest}} = 50$ m	0.638	0.578	0.12151	0.10108	-16.8	96.9–101.1	yes
	$z_{\text{rest}} = 100$ m	0.639	0.603	0.13009	0.13153	+1.1	99.0–101.0	yes

Table S14. Sensitivity of the weak-to-baseline response to the width of the fixed baseline-centred diagnostic layer, recomputed from the frozen converged profiles without re-integration. The layer is centred on the baseline oxygen-minimum depth (222.5 m at both stations) and the half-width is varied. The 50 m half-width is the configuration used throughout the main text. Layer-mean budget imbalance dI/dt is 10^{-7}–10^{-5} $\mathrm{mmol\,m^{-2}\,d^{-1}}$ in every column.

Stn	half-width m	interval m	thickness m	η weak	η base	F_{top} weak $\mathrm{mmol\,m^{-2}\,d^{-1}}$	F_{top} base	ΔF_{top} %
S3	25	195.0–250.0	55	0.547	0.476	0.03385	0.01784	-47.3
	50	170.0–275.0	105	0.574	0.492	0.09952	0.07216	-27.5
	75	145.0–300.0	155	0.633	0.530	0.23890	0.20466	-14.3
S4	25	195.0–250.0	55	0.609	0.534	0.04728	0.02529	-46.5
	50	170.0–275.0	105	0.638	0.555	0.12095	0.08607	-28.8
	75	145.0–300.0	155	0.689	0.589	0.20187	0.16064	-20.4

Table S15. Independent BGC-Argo evaluation with its processing uncertainty. Each observational metric is computed along three defensible paths differing only in depth-bin labelling and in whether individual profiles or monthly means are pooled; spread is (max−min)/mean across those paths. Model values are the annual mean of the diagnostic year of the baseline configuration on the same 25 m grid. No parameter was adjusted in response to these data and WOA23 is not used. The gradient at z_u is not constrained by these observations and no agreement is claimed for it.

Metric	Stn	path 1	path 2	path 3	spread (%)	model
core-layer O_2 ($\mathrm{mmol\,m^{-3}}$)	S2	1.736	1.986	2.419	33	2.964
	S3	2.924	2.938	3.044	4	3.894
	S4	2.799	2.990	3.080	9	5.095
	S5	14.337	14.331	18.188	25	20.151
O_2 minimum ($\mathrm{mmol\,m^{-3}}$)	S2	1.561	1.546	1.421	9	2.747
	S3	2.790	2.770	2.744	2	3.545
	S4	2.490	2.499	2.157	14	4.324
	S5	13.521	13.698	7.313	55	19.745
depth of O_2 min (m)	S2	225	225	250	11	200
	S3	200	225	225	12	225
	S4	250	300	325	26	225
	S5	200	225	425	79	225
$\partial O_2/\partial z$ at z_u ($\mathrm{mmol\,m^{-4}}$)	S2	-0.138	-0.213	-0.275	66	-0.112
	S3	-0.025	-0.052	-0.126	149	-0.108
	S4	-0.056	-0.091	-0.141	88	-0.106
	S5	-0.059	-0.087	-0.184	113	-0.173

Matchup radius 200 km. Profiles / unique floats: S2 298/8; S3 361/16; S4 343/7; S5 170/4.
Platform-block bootstrap, 1000 replicates over whole floats,
2.5–97.5 % interval on O_2 at z_u: S2 1.9–8.7; S3 2.3–4.1; S4 3.0–6.3; S5 12.4–21.1.
Modelled and observed oxyclines fall in the same 25 m bin at all four stations.

Table S16. Matchup-radius sensitivity and platform-level uncertainty of the BGC-Argo evaluation. Each station is re-evaluated at three horizontal matchup radii; *floats* is the number of distinct WMO platforms contributing, which is the effective sample size because profiles from one drifting float are strongly autocorrelated. The final column gives the 2.5–97.5 % interval on oxygen at z_u from 1000 bootstrap replicates resampling whole floats with replacement at the 200 km radius. S1 is not evaluated (Sect. 2 of the main text) and S6 has only three profiles within 200 km. No model parameter was adjusted in response to these data.

Stn	R (km)	profiles	floats	median dist (km)	core O_2	O_2 at z_u ($\mathrm{mmol\,m^{-3}}$)	$z(O_2\mathrm{min})$ (m)	O_2 at z_u bootstrap 95 %
S2	100	44	4	72	2.51	8.15	238	1.9–8.7
	150	112	5	116	2.59	7.97	238	
	200	298	8	164	2.47	6.24	238	
S3	100	81	7	82	3.02	3.13	212	2.3–4.1
	150	183	11	105	3.13	3.47	212	
	200	361	16	149	2.87	3.38	212	
S4	100	95	5	70	2.11	3.88	262	3.0–6.3
	150	209	6	106	2.71	4.23	262	
	200	343	7	133	2.87	4.56	262	
S5	100	49	3	78	15.40	16.74	212	12.4–21.1
	150	96	3	100	15.10	16.36	212	
	200	170	4	135	14.27	15.34	212	

Core-layer oxygen and the depth of the oxygen minimum are stable across radii at S3 and S4. The gradient at z_u is not tabulated here because it is not constrained by these data: it varies by 66–150 % of its own mean between processing paths (Table S15).

Table S17. Physical time-step sensitivity at S3. Only the physical time step is varied; the number of biological sub-steps per physical step is held at 20, so the biological sub-step duration scales with Δt. The 5 m vertical grid is fixed in every station forcing file and could not be varied without regenerating all model inputs from the source products, so vertical-resolution dependence is reported as a limitation (Sect. S2) rather than tested. Every run was integrated to the same convergence criterion and diagnosed over the subsequent complete model year.

Δt (h)	F_{top} weak	F_{top} baseline	ΔF_{top} (%)	core O_2 baseline	η baseline
0.5	0.099528	0.072167	-27.491	3.8124	0.4922
1.0	0.099519	0.072158	-27.493	3.8121	0.4922
2.0	0.099501	0.072142	-27.497	3.8112	0.4922

Spread across a fourfold change in Δt: 0.005 percentage points in ΔF_{top}, 0.031 % in core oxygen.

Supplementary figures

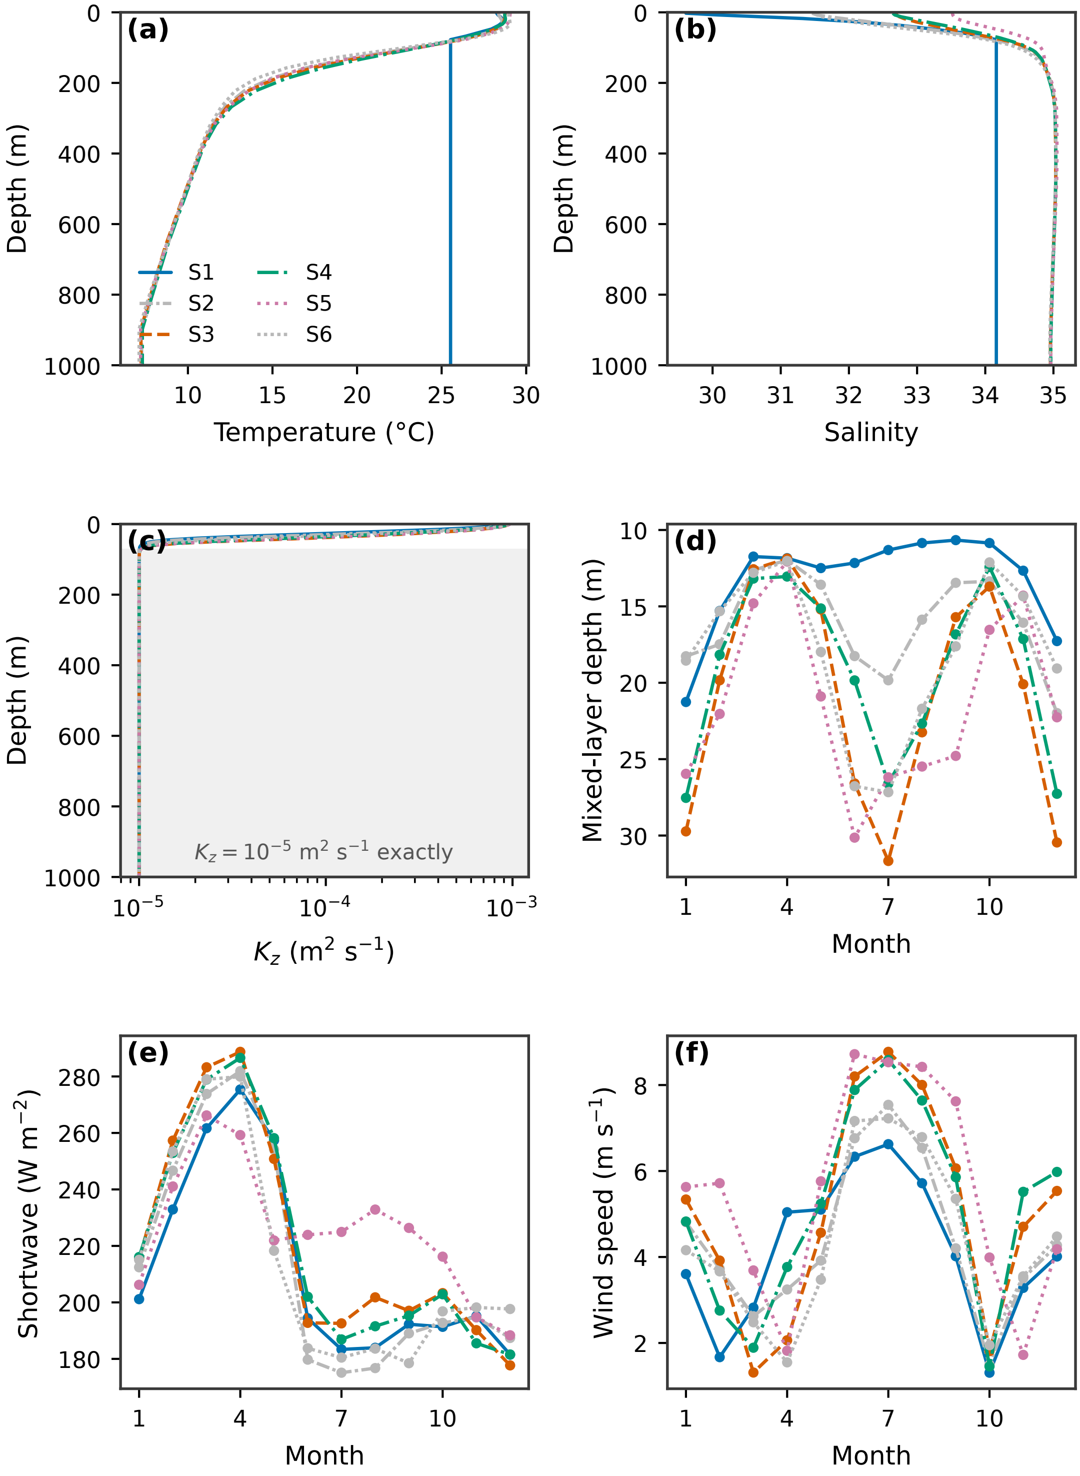

Figure S1. Prescribed physical forcing. Annual-mean (a) temperature, (b) salinity and (c) vertical diffusivity over the full 1000 m model domain, and monthly climatological (d) mixed-layer depth, (e) shortwave radiation and (f) wind speed. The five stations analysed in the main text are coloured; S1 is grey. Vertical diffusivity is not a reanalysis field: it is computed analytically from the GLORYS12 mixed-layer depth, and because that depth is 10.7–31.7 m at these stations the profile reaches its background value of 1.0×10^{-5} m^2 s^{-1} by 60–70 m and equals it below 150 m at every station in every month. The shading in (c) marks that constant, which spans the whole of the analysed upper layer and core. All quantities are prescribed model input, not model output.

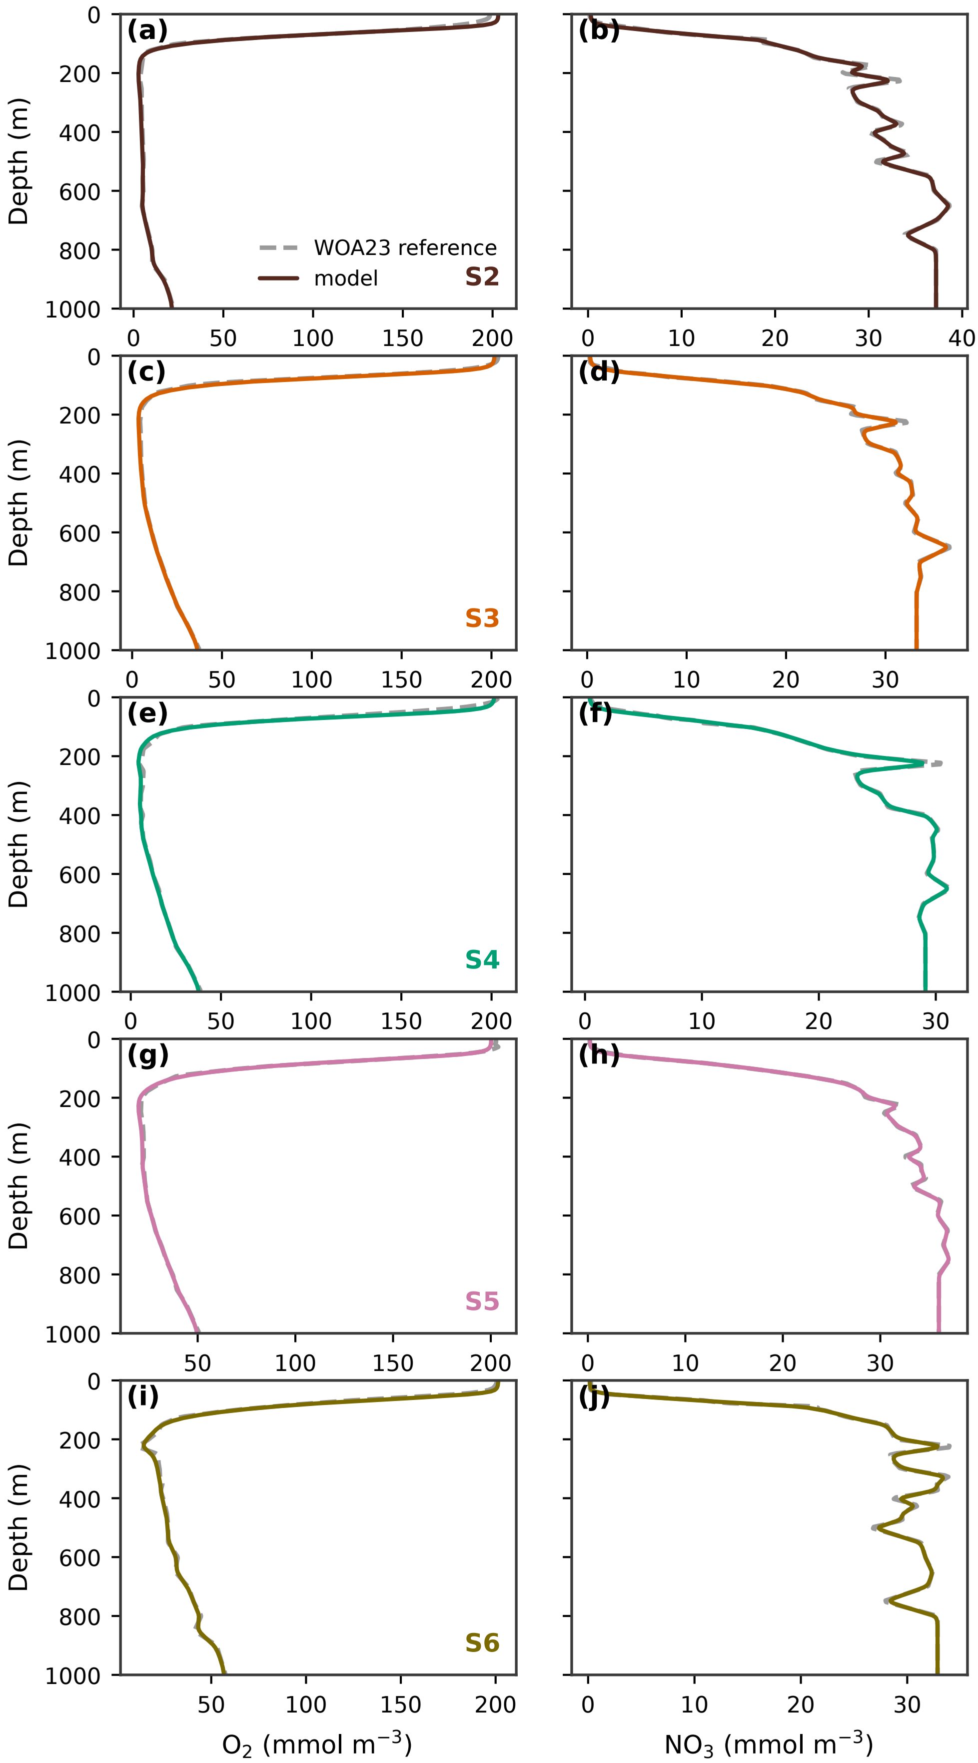

Figure S2. Simulated annual-mean oxygen (left) and nitrate (right) profiles at the five analysed stations, shown against the WOA23 climatology. WOA23 is used for model initialisation and as the restoring target, so it is a climatological reference rather than an independent validation dataset and the comparison should be read as a model–reference difference. No skill statistics are computed. Simulations were integrated until the convergence criterion was satisfied and diagnosed over the subsequent complete 360-day model year.

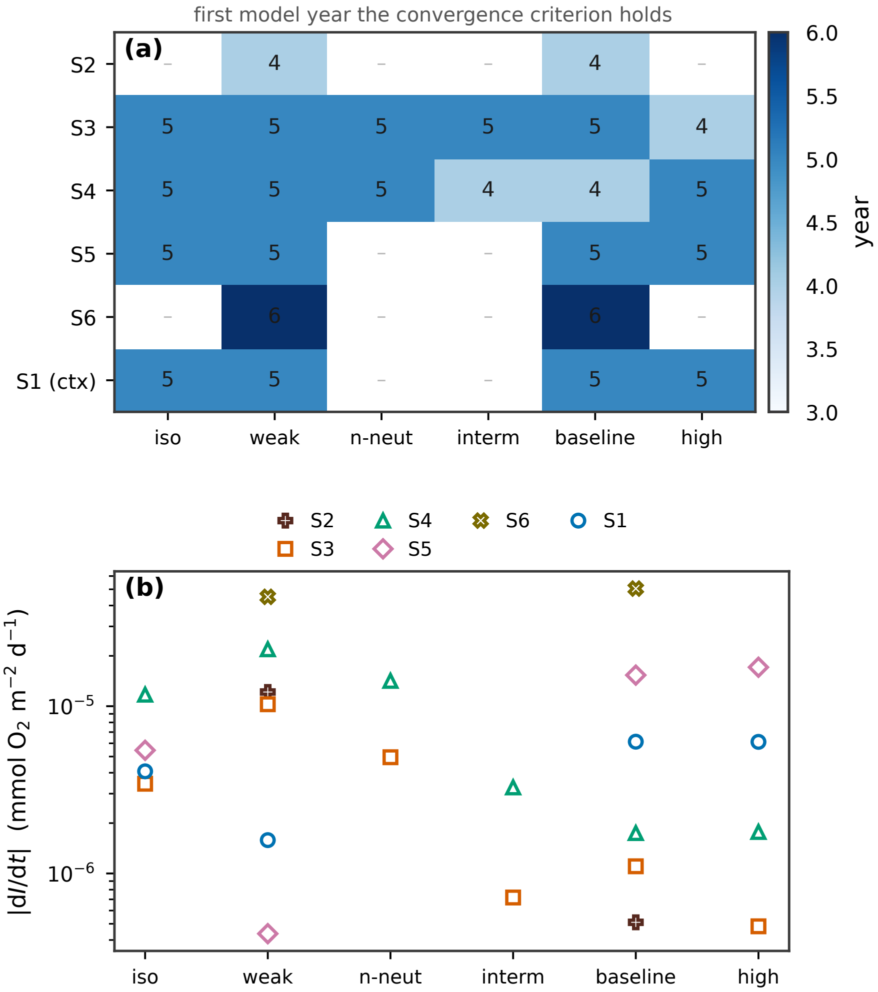

Figure S3. Convergence of the retained experiments. (a) First model year from which the automatic convergence criterion holds to the end of the run, for every station and particle-transport configuration; dashes mark combinations not run. Convergence is assessed simultaneously on layer-mean oxygen, total nitrogen and each biological compartment. All diagnostics are taken from the complete 360-day model year following that point. (b) Magnitude of the core oxygen inventory tendency over that diagnostic year, confirming that storage is negligible in every retained case.

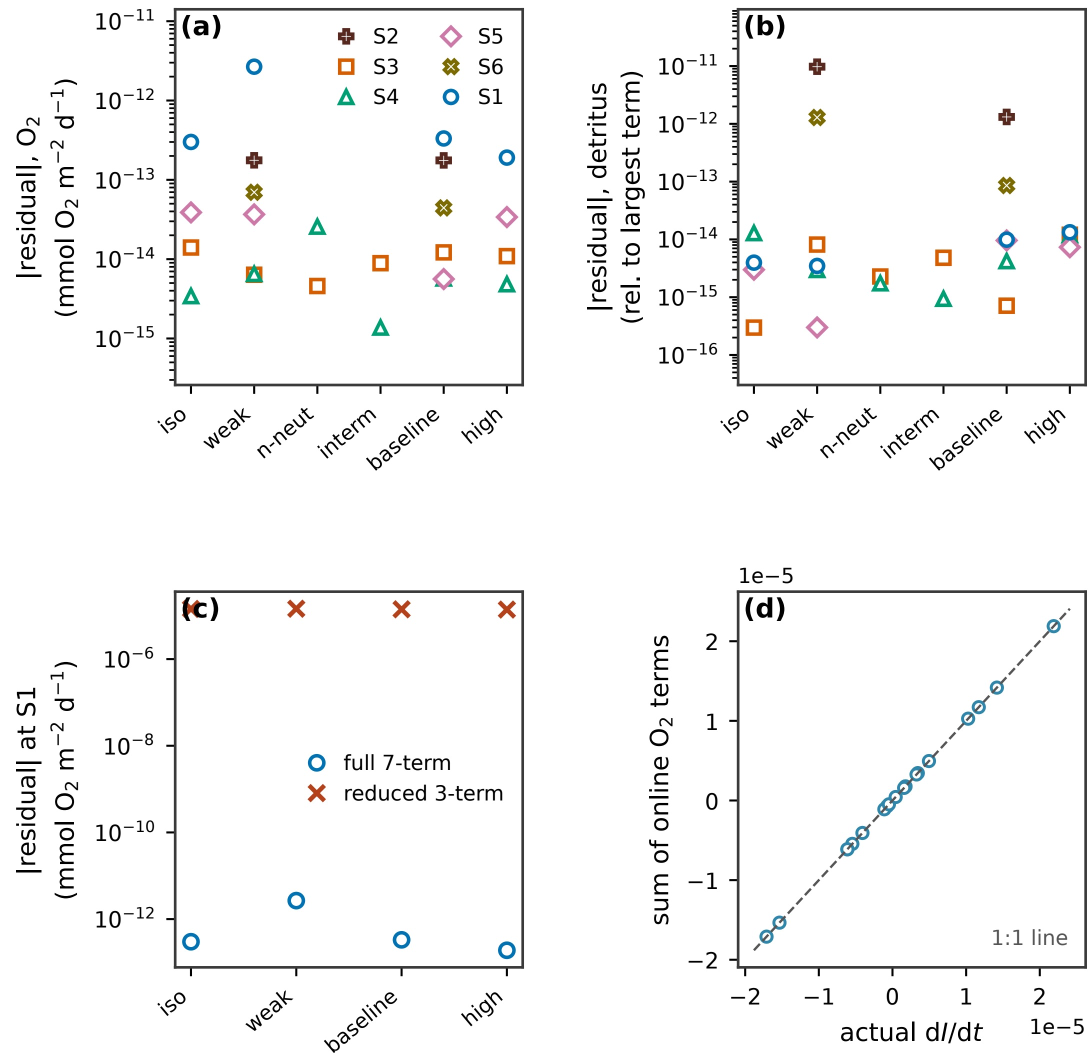

Figure S4. Closure of the online budgets over the diagnostic year. (a) Absolute residual of the full seven-term core oxygen identity and (b) residual of the discrete detritus budget assembled from direct online source and loss terms, normalised by the largest individual term because the inventory change itself tends to zero at equilibrium. (c) At S1 the reduced three-term oxygen balance leaves a residual many orders of magnitude larger than the full seven-term identity, because the S1 core carries small but non-zero photosynthesis and respiration; all S1 diagnostics therefore use the full form. (d) Sum of the online oxygen terms against the actual core inventory tendency for every station and configuration, with the 1:1 line. Simulations were integrated until the convergence criterion was satisfied and diagnosed over the subsequent complete 360-day model year.

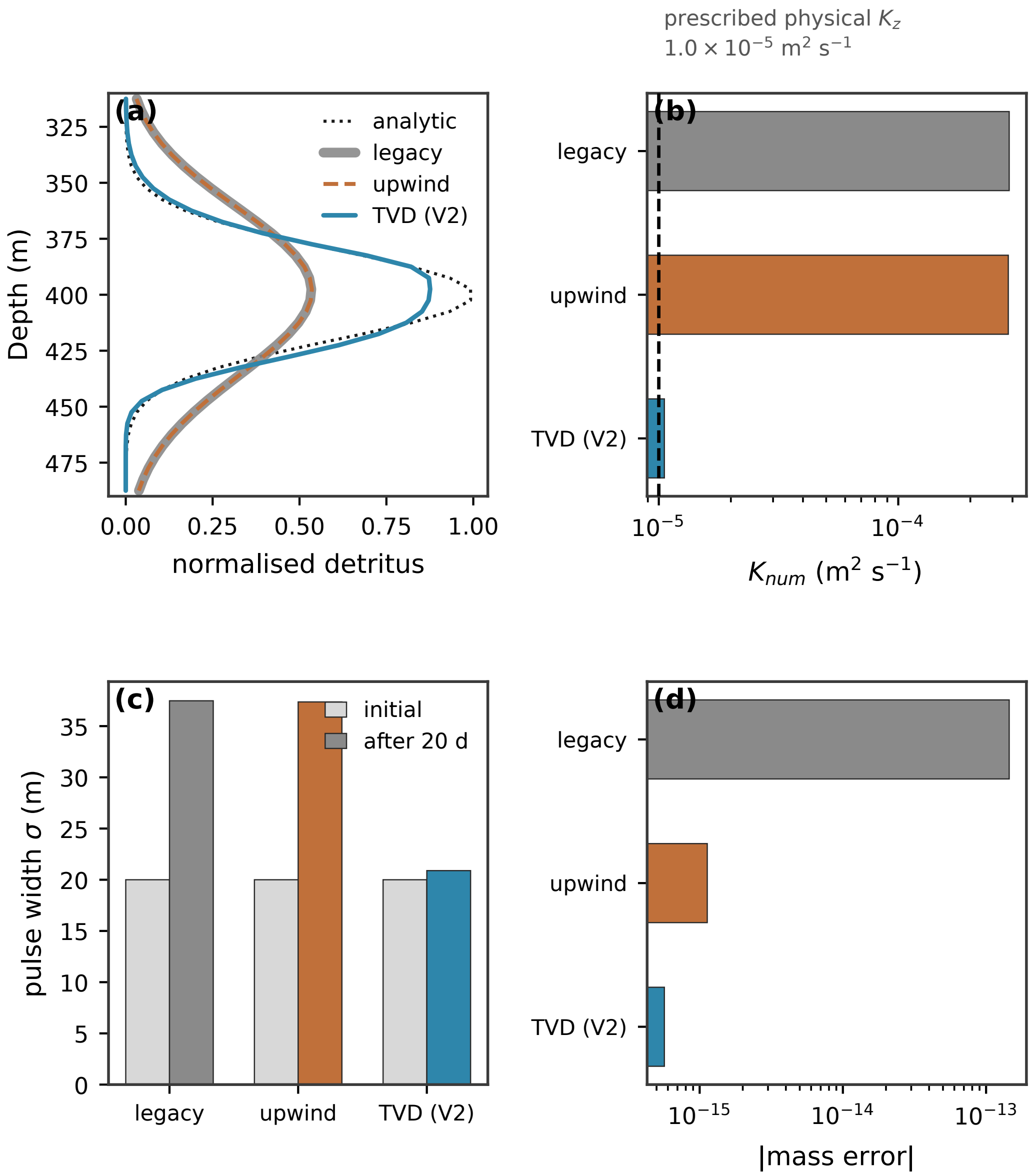

Figure S5. Numerical behaviour of the detrital sinking schemes, from a Gaussian pulse advected for 20 days at 10 m d^{-1} with no biology or physical diffusion. (a) Actual final numerical profiles after 20 days of pure advection, with the analytic solution retaining its initial width. (b) Effective numerical diffusivity inferred from variance growth, against the prescribed physical diffusivity over the analysed depth range. (c) Initial and final pulse width: the legacy in-place sweep and explicit upwind both nearly double it, and are so nearly identical to each other (37.47 against 37.36 m) that their profiles in (a) coincide, which is why legacy is drawn as the wide solid line beneath the dashed upwind one; the van Leer TVD scheme used here broadens it by under 5 %. (d) Absolute mass-conservation error; all three schemes conserve mass to 10^{-13} or better.

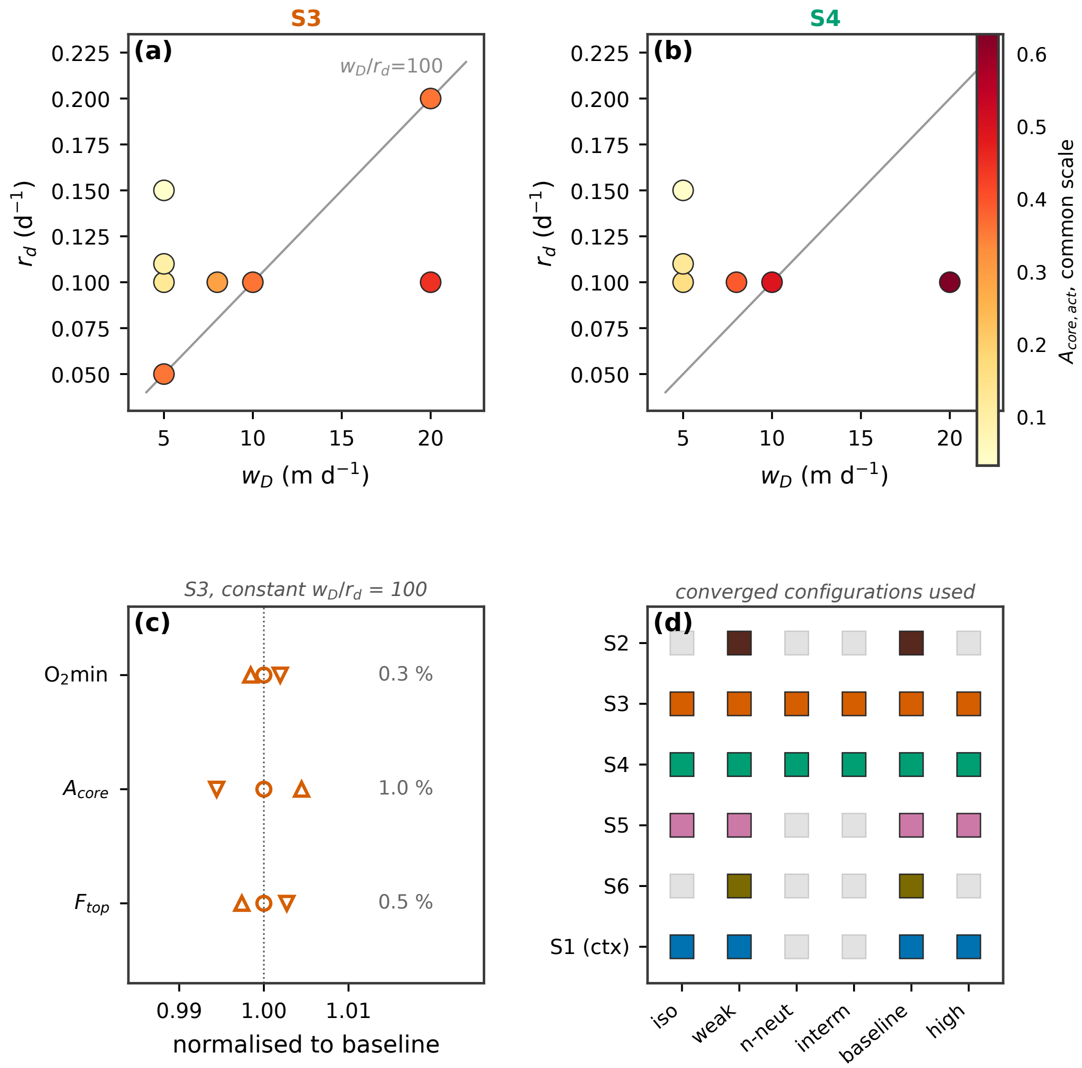

Figure S6. Particle-transport parameter structure. (a, b) Realised core respiratory oxygen demand across the converged configurations at S3 and S4 in the (w_D, r_d) plane; the grey line marks constant $w_D/r_d = 100$. Both panels use one common colour normalisation, so colours are directly comparable between them. (c) The three S3 configurations sharing that ratio, each response normalised to the baseline with the total spread given: within the tested range the oxygen response is much better constrained by the effective transport-to-remineralisation ratio than by w_D or r_d individually, though the degeneracy is approximate rather than exact. (d) Configurations available at each station; S1 and S5 were not run at $w_D = 8\ \mathrm{m\,d^{-1}}$. All plotted points are converged runs and no statistical model is fitted to them.

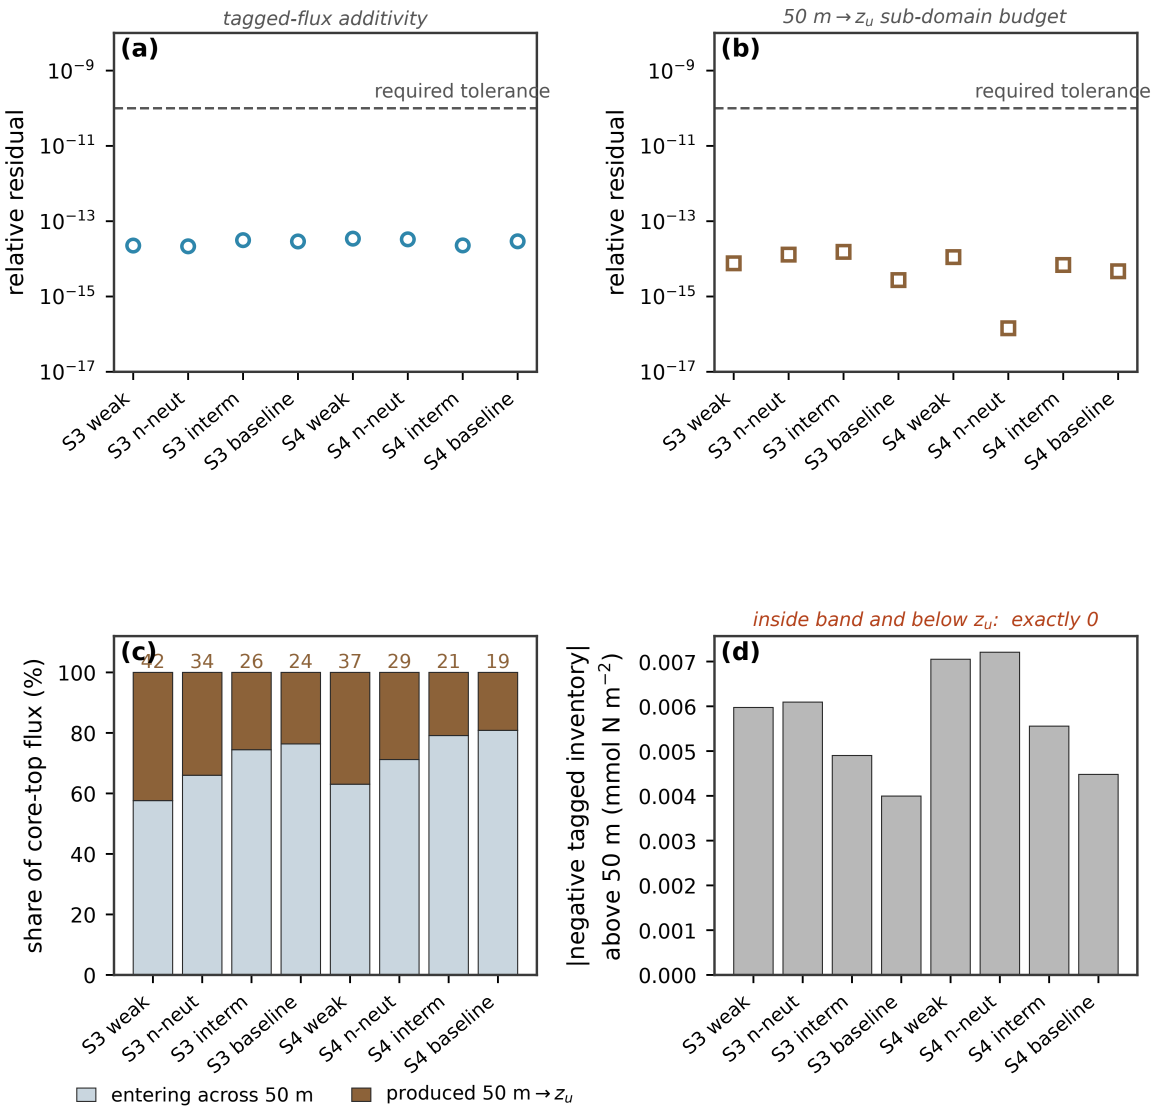

Figure S7. Quality control of the core-top particle-flux attribution. (a, b) The two additivity and subdomain-closure tests fixed in advance: additivity of the tagged sinking fluxes at the core upper interface, and closure of the detritus budget for the 50 m to z_u sub-domain. Both pass by four orders of magnitude against the 10^{-10} relative threshold. (c) Decomposition of the annual core-top sinking flux into material entering across 50 m and material attributed to biological production between 50 m and z_u, under the frozen-total-D-limiter source-tag decomposition; these are shares of a linear flux decomposition and not Lagrangian particle histories. (d) Negative tagged inventory produced by applying the total-field flux limiter to a minority component. It lies entirely above 50 m and is exactly zero both inside the attribution band and below z_u, so it cannot bias the flux at z_u.

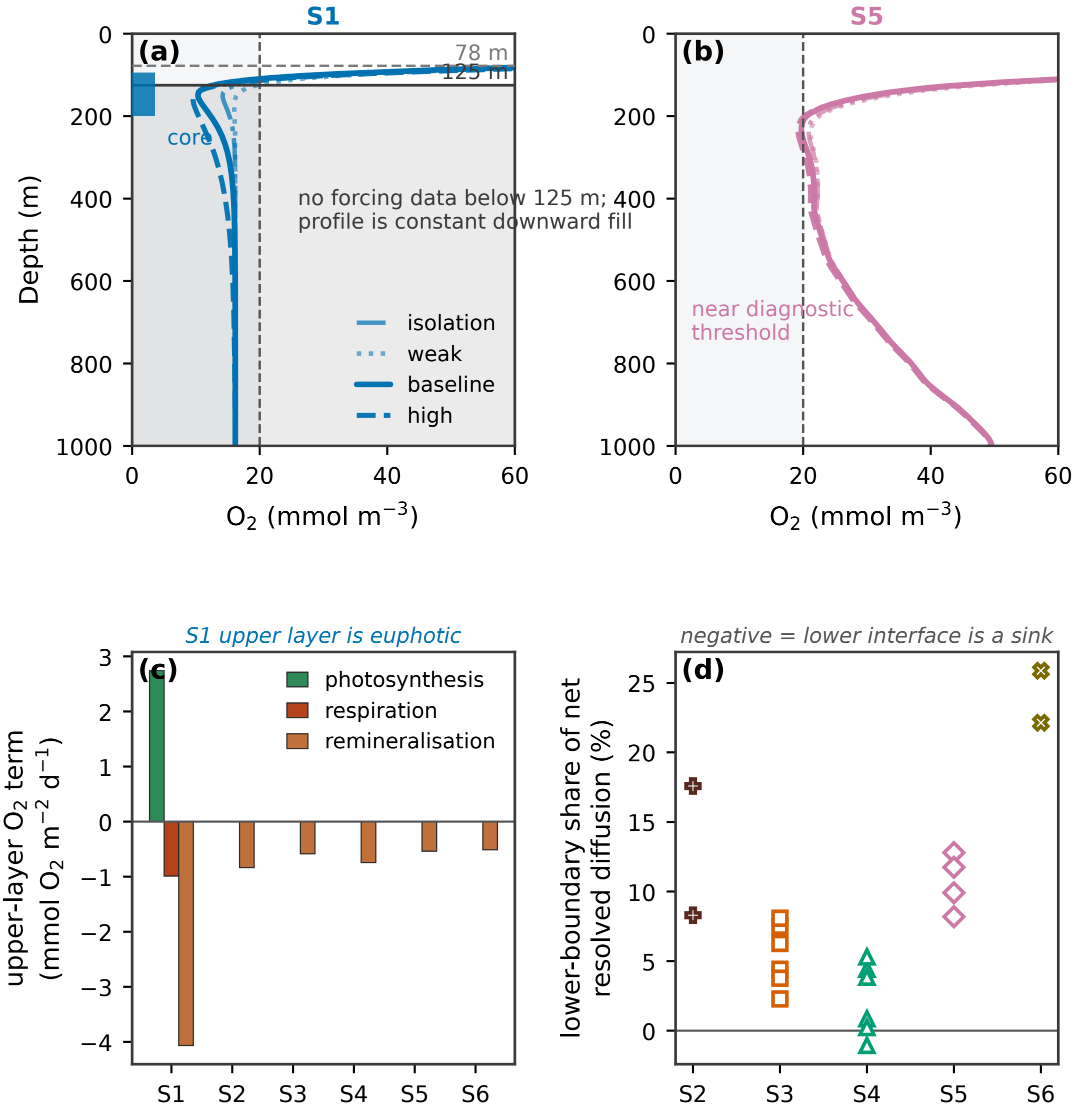

Figure S8. Station-dependent limits of common OMZ diagnostics. Annual-mean oxygen profiles across all tested configurations at (a) S1 and (b) S5; shading and the dashed line mark 20 mmol m^{-3}. At S1 no OMZ thickness is reported: the oxygen target below 127.5 m is a constant downward extrapolation of the deepest WOA23 level, so the apparent sub-threshold layer there is an artefact of that fill and is not interpreted. At S5 the oxygen minimum spans only 19.2 to 21.2 mmol m^{-3} across the whole sequence, so threshold exceedance reflects proximity to the threshold and is not a tipping point. (c) Baseline upper-layer oxygen terms: at S1 photosynthesis is comparable in magnitude to remineralisation, whereas it is an exact zero at S2–S6, so the geometrically defined upper layer is not the same biogeochemical object at every station. (d) Contribution of the lower core interface to net resolved diffusive supply; negative values indicate a net sink there.

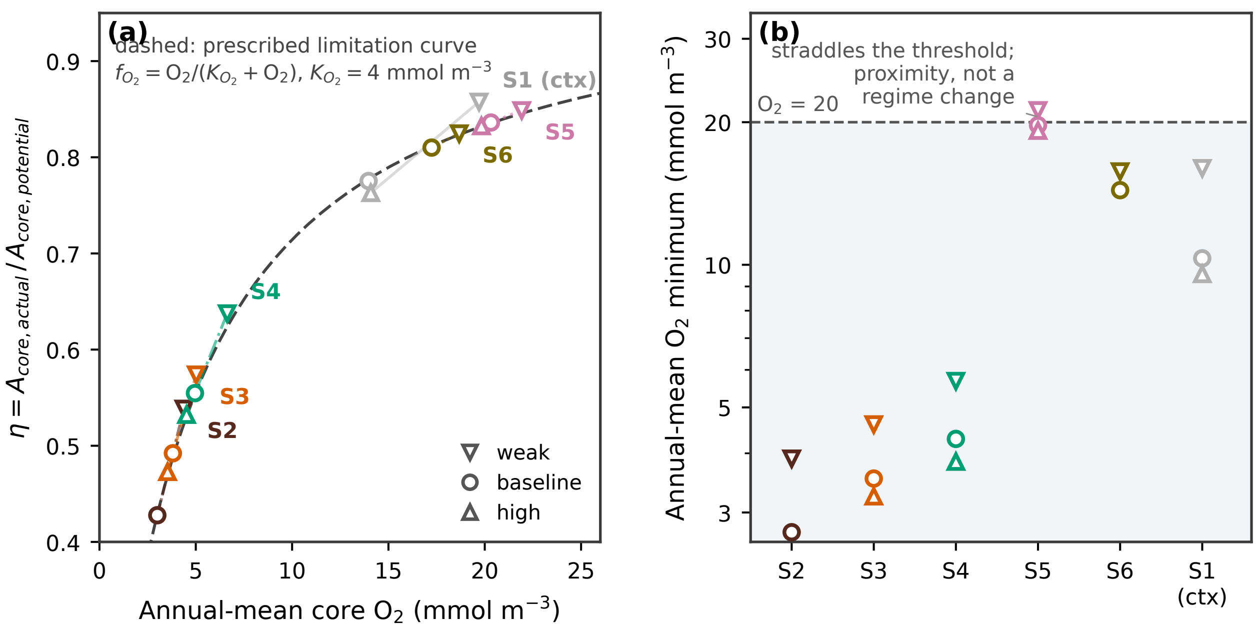

Figure S9. Oxygen-limitation diagnostic and its state dependence. (a) The limitation factor η, the ratio of realised to conditional potential core remineralisation demand, against annual-mean core oxygen, with the prescribed Michaelis–Menten function $f_{O_2} = O_2/(K_{O_2} + O_2)$, $K_{O_2} = 4\,\mathrm{mmol\,m^{-3}}$, drawn as the dashed curve. η is a demand-weighted mean of that function over the core, so the diagnosed points lie on the curve by construction; the panel shows where each column sits on a prescribed relationship and is not evidence of an empirical law. Because that identity is algebraic (Sect. 3.2 of the main text), the panel is presented as supporting material rather than as a result. (b) Annual-mean oxygen minimum by station and configuration; the shaded region and dashed line mark the $20\,\mathrm{mmol\,m^{-3}}$ diagnostic OMZ threshold, which is a descriptive boundary and not a denitrification threshold. At S5 the minimum spans only 19.2 to $21.2\,\mathrm{mmol\,m^{-3}}$ across the whole sequence, so straddling that line reflects proximity to it and is not a regime change. S1 is shown for context only; 71 % of its diagnostic core is constant downward extrapolation and it is not used quantitatively.