

Supplemental Material for Causal Drivers of Southern Ocean Greening Revealed by State-Estimate and Satellite Data

Gian Giacomo Navarra^{1*}, Curtis Deutsch² and Matt Mazloff³

^{1*}Atmosphere and Ocean Science, Princeton University, USA.

²Geoscience, Princeton University, USA.

³Scripps Institution of Oceanography, University of California, USA.

*Corresponding author(s). E-mail(s): gn5970@princeton.edu;

Contributing authors: cdeutsch@princeton.edu; mmazloff@ucsd.edu;

Abstract

This document provides supplemental figures and mathematical derivations supporting the main article. Section 1 presents the conceptual framework for light–iron co-limitation and its dependence on mixed-layer depth, together with the vertical structure of key biological parameters and seasonal trends in physical drivers. Section 2 presents the full seasonal PCMCI causal maps for all drivers. Section 3 provides bootstrap confidence intervals confirming the statistical robustness of the dominant causal pathways in the ACC. Section 4 extends the PCMCI analysis to the subtropical and high-latitude regions. Section 5 shows the sensitivity of causal results to predictor selection. Section 6 presents trends in net primary production from independent products. Section 7 provides partial correlation maps complementing the PCMCI analysis of mixed-layer depth. Section 8 shows the full spatial PCMCI maps for MLD variability. Section 9 derives the biomass-weighted limitation metrics. Sections 10–12 provide the full mathematical derivation of the GPDC–MCI test statistic and its statistical significance assessment.

1 Light and Iron Limitation in the Southern Ocean

Phytoplankton growth in the Southern Ocean is primarily constrained by two environmental factors: vertical iron supply and light availability. These form the conceptual framework for interpreting recent chlorophyll-*a* (Chl-*a*) changes. Iron is resupplied to the surface through deeper mixing and vertical entrainment of nutrient-rich waters, while light limitation is controlled by mixed-layer depth, seasonal stratification, and sea-ice cover. The vertical structure of these competing controls is illustrated in Supplemental Fig. S1. Light availability decreases approximately exponentially with depth due to radiative attenuation, whereas iron concentrations increase below the mixed layer through vertical resupply from deeper waters. We represent these processes using normalised resource functions

$$L(z) = e^{-z/z_L}, \quad F(z) = \frac{1}{1 + \exp(-(z - z_F)/\Delta z_F)}, \quad (1)$$

where $z_L = 35$ m is the optical decay scale and $z_F = 70$ m is the nutricline depth. Assuming co-limitation, potential growth rate is proportional to the product of both resources, normalised by its maximum:

$$G(z) = \frac{L(z) F(z)}{\max_z [L(z) F(z)]}. \quad (2)$$

Expressed in mixed-layer-depth space, deeper mixing simultaneously reduces light exposure and enhances iron entrainment,

$$L_{\text{MLD}} = e^{-\text{MLD}/z_L}, \quad F_{\text{MLD}} = 1 - e^{-\text{MLD}/z_F}, \quad (3)$$

producing maximum productivity at intermediate mixing depths. Shallow mixed layers are therefore iron-limited, deep mixed layers are light-limited, and atmospheric forcing influences chlorophyll variability by modulating mixed-layer depth, which jointly controls both resources.

1.1 Vertical structure of key biological parameters

The depth profiles of key biogeochemical and physical quantities are shown in Supplemental Fig. S3. These profiles demonstrate that phytoplankton growth rate, biomass, and Chl-*a* all peak within the upper 25 m of the water column, coinciding with the transition between the high-light surface layer and the iron-rich waters below. This vertical structure justifies the choice of 25 m as the integration depth for Chl-*a* and associated physical drivers throughout the analysis.

1.2 Seasonal dependence of physical driver trends

To characterise how the physical environment has changed over the analysis period, we computed seasonal trends for each physical driver by calculating the linear trend at each month of the year separately, retaining the full seasonal cycle rather than

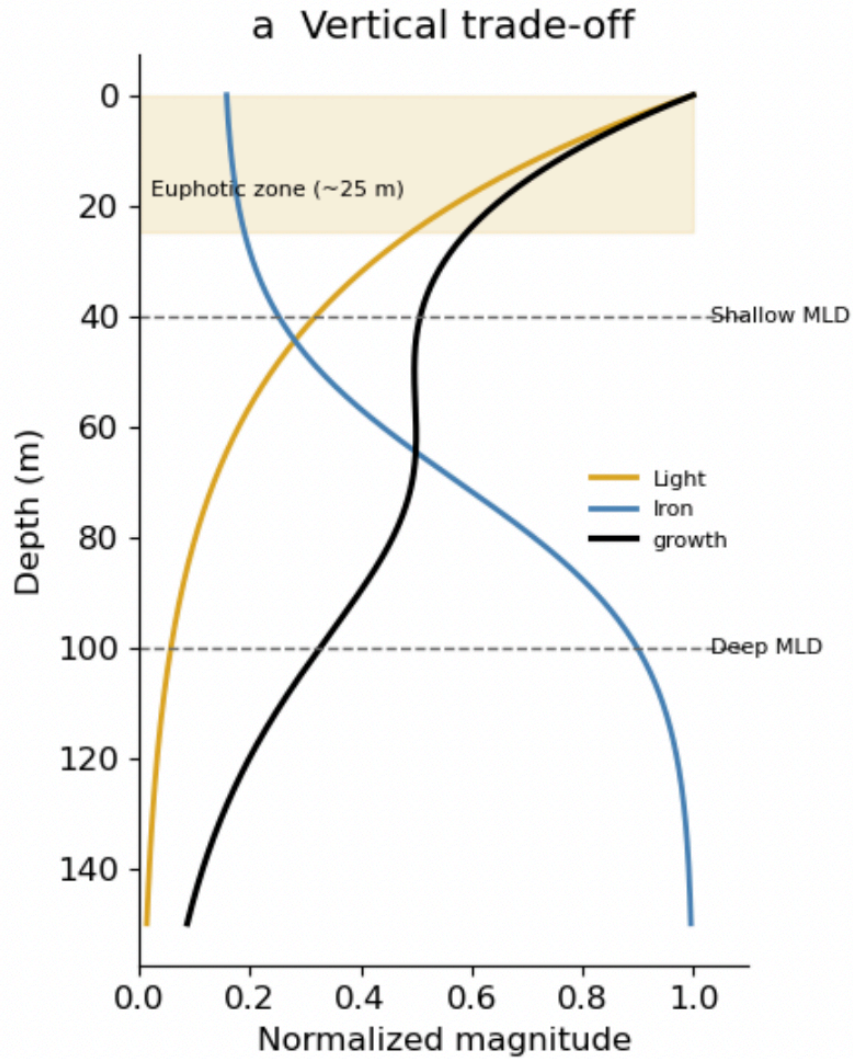


Fig. S1 Schematic of the vertical trade-off between light and iron limitation as a function of mixed-layer depth. Light availability (orange) decreases exponentially with depth, while iron availability (black) increases below the mixed layer through vertical entrainment. The product of the two resources (growth potential, blue) peaks at intermediate depths near the base of the euphotic zone (~25 m), providing the physical basis for the sensitivity of Southern Ocean productivity to mixed-layer depth variability.

removing it. The results are shown in Supplemental Fig. S4. This approach reveals the timing of changes in temperature, mixed-layer depth, salinity, surface heat flux, shortwave radiation, and wind amplitude, and allows the temporal ordering of physical driver anomalies to be assessed independently of the annual mean trend.

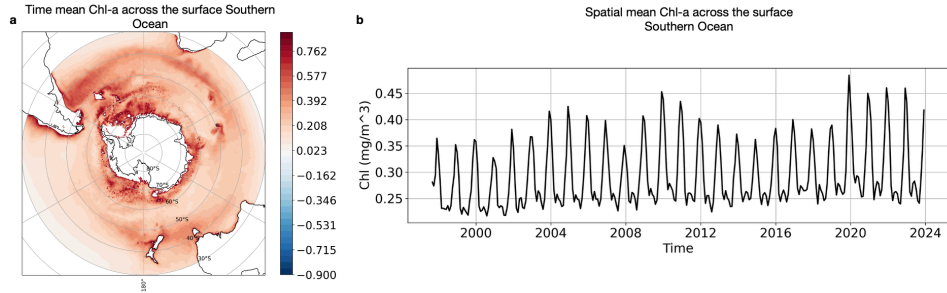


Fig. S2 Time-mean surface Chl-*a* and long-term trend across the Southern Ocean. Panel (a) shows the spatial distribution of the time-mean Chl-*a* concentration integrated over the upper 25 m, and panel (b) shows the corresponding basin-wide spatial mean time series. The positive trend in Chl-*a* over the observational period is statistically significant and consistent with independent satellite and BSOSE estimates.

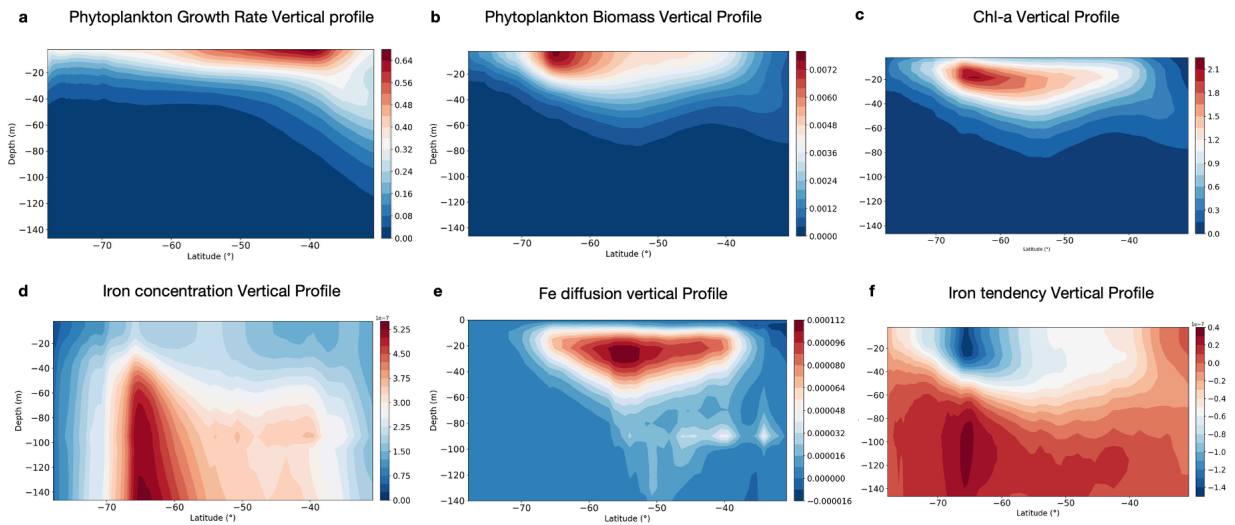


Fig. S3 Vertical structure of key biogeochemical and physical drivers in the Southern Ocean. Shown are summer Hovmöller sections of phytoplankton growth rate (a), phytoplankton biomass (b), chlorophyll-*a* (c), dissolved iron concentration (d), vertical iron diffusive flux (e), and the iron tendency (f). All fields exhibit a subsurface maximum in biological activity near 25 m depth, coincident with the transition between high light availability above and increased nutrient supply below. Iron concentrations and diffusive fluxes peak beneath the mixed layer, highlighting the role of vertical entrainment as the primary iron source to the euphotic zone.

2 Seasonal PCMCI Causal Maps

Supplemental Figs. S5 and S6 show the full seasonal spatial maps of PCMCI causal coefficients between Chl-*a* anomalies and all physical drivers considered in the analysis. The dominance of MLD and iron diffusive flux within the ACC is spatially coherent and evident across all seasons.

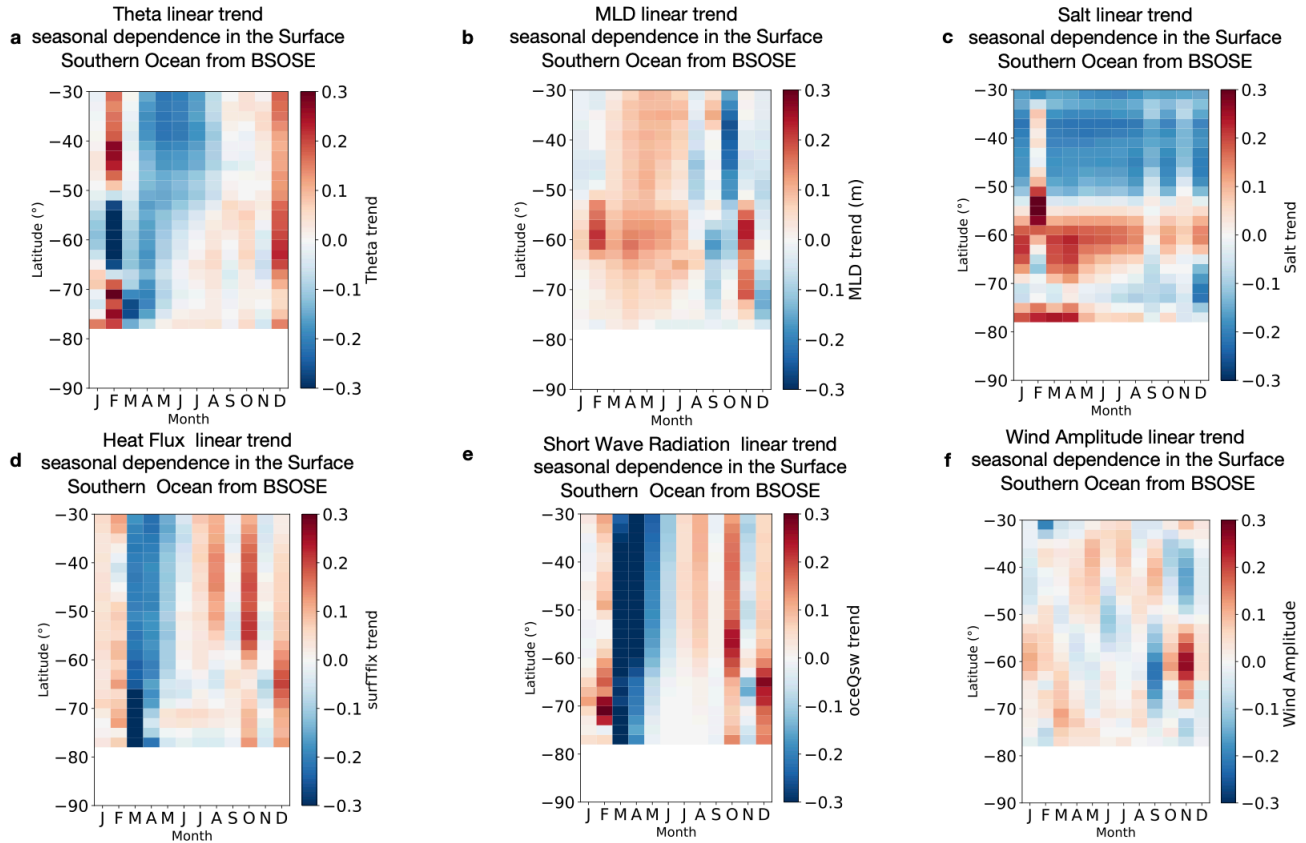


Fig. S4 Seasonal dependence of linear trends in physical drivers from BSOSE. Shown are zonally averaged, month-by-month trends in potential temperature (a), mixed-layer depth (b), salinity (c), surface heat flux (d), shortwave radiation (e), and wind amplitude (f). Reductions in shortwave radiation and surface heat flux precede mixed-layer deepening by one to two months, consistent with a thermodynamic control on mixed-layer variability.

3 Bootstrap Analysis of Causal Link Significance in the ACC

To confirm that the dominant causal pathways identified in the ACC are statistically robust, we applied a non-parametric spatial bootstrap over all grid points displaying a significant causal link with Chl-*a* (see Methods in the main text). For each region, we generated 2000 bootstrap realisations by resampling the significant grid-point values with replacement and computing the mean for each realisation, yielding an empirical distribution of the regional mean MCI coefficient. The 95% confidence interval is given by the 2.5th and 97.5th percentiles of this distribution.

The results are shown in Supplemental Fig. S7. In panel (a), the confidence intervals for MLD do not overlap with those of wind amplitude, salinity, or buoyancy frequency across all seasons, confirming that MLD is the statistically dominant causal driver of

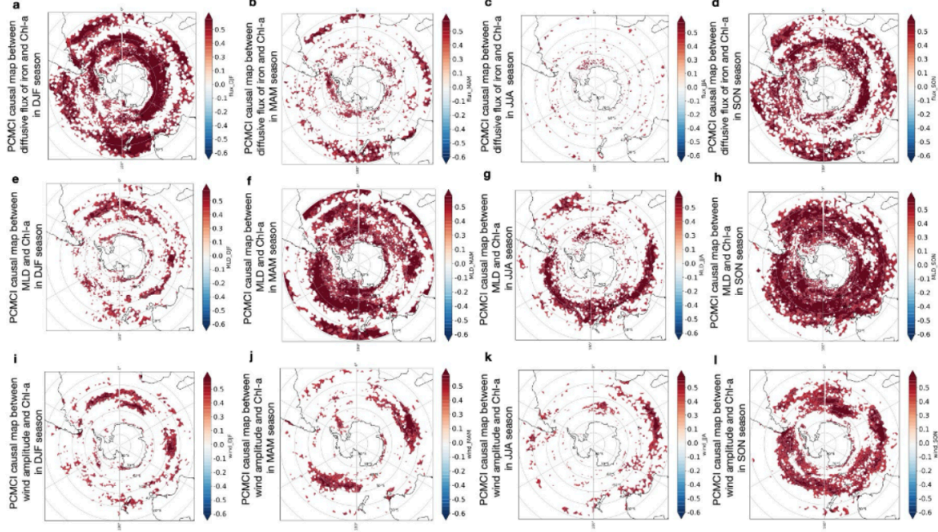


Fig. S5 Seasonal spatial maps of PCMCI causal links between Chl-*a* anomalies and selected physical drivers. Shown are the MCI causal coefficients, computed separately for each season (DJF, MAM, JJA, SON), for the diffusive flux of iron (a–d), mixed-layer depth (e–h), and wind amplitude (i–l). Significant links ($p < 0.05$, FDR-corrected) are shown in colour; grey indicates non-significant grid points. The dominance of MLD and iron diffusive flux within the ACC is evident across all seasons.

Chl-*a* variability in the ACC. In panel (b), the confidence intervals for surface heat flux and shortwave radiation do not overlap with those of wind amplitude in most seasons, confirming that thermodynamic forcing exerts a statistically distinguishable stronger control on MLD variability than mechanical wind-driven mixing. These results provide quantitative support for the claims in the main text and are not an artefact of spatial averaging.

4 PCMCI Analysis for the Subtropics and High-Latitude Southern Ocean

To assess whether the causal structure identified within the ACC extends to other dynamical regimes, we repeated the PCMCI analysis separately for the subtropical and high-latitude regions of the Southern Ocean. The results, including bootstrap confidence intervals, are shown in Supplemental Fig. S8. In contrast to the ACC, no single driver achieves consistently dominant causal strength in either region across all seasons, indicating that chlorophyll variability outside the ACC reflects spatially heterogeneous and comparable contributions from multiple environmental factors.

5 Sensitivity of Causal Results to Predictor Selection

To verify that the causal pathways identified in the main analysis are not artefacts of the specific set of predictors included, we repeated the PCMCI analysis while systematically excluding one variable at a time. The resulting causal networks, shown in




Fig. S6 Seasonal spatial maps of PCMCI causal links between Chl-*a* anomalies and additional physical drivers. Shown are the MCI causal coefficients for buoyancy frequency (a–d), potential temperature (e–h), and salinity (i–l), computed separately for each season (DJF, MAM, JJA, SON). The stronger and more spatially coherent causal influence of buoyancy frequency during summer, relative to the other drivers, is consistent with the transition to a stratification-controlled regime in the ACC during the austral growing season.

Supplemental Fig. S9, are qualitatively unchanged regardless of which variable is withheld. The dominant causal influence of MLD on Chl-*a* within the ACC is preserved in all cases, confirming that the detected relationships reflect intrinsic dependencies within the coupled physical–biogeochemical system rather than sensitivity to predictor selection.

6 Changes in Net Primary Production

To assess whether the Chl-*a* trends are accompanied by changes in carbon fixation, we computed linear trends in net primary production (NPP) from three independent products: BSOSE, the Eppley-VGPM satellite-based estimate, and the standard VGPM product. Results are shown in Supplemental Fig. S10. All three products show positive NPP trends over the analysis period, consistent with the Chl-*a* increase representing a real enhancement of phytoplankton productivity rather than solely a change in cellular pigment concentration.

7 Partial Correlation Analysis of Mixed-Layer Depth Anomalies

To complement the PCMCI analysis, we computed partial correlations between MLD anomalies and key oceanic and atmospheric drivers, controlling for the influence of

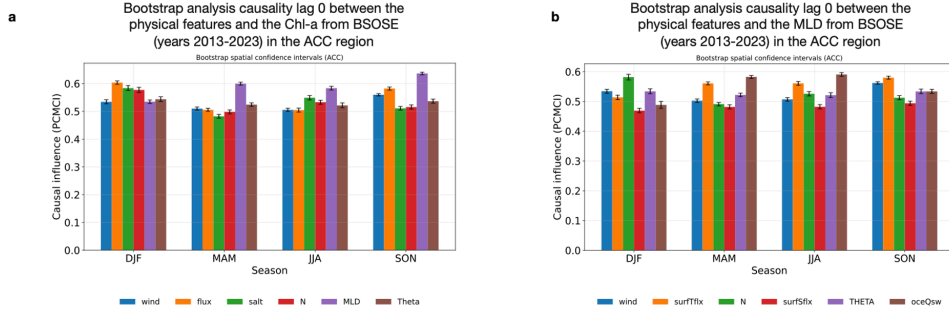


Fig. S7 Bootstrap confidence intervals for the seasonal mean PCMCI causal coefficients (MCI) in the ACC region. Panel (a) shows causal influences on Chl-*a* anomalies and panel (b) shows causal influences on MLD anomalies, for each driver (wind amplitude, surface heat flux, salinity, buoyancy frequency N , mixed-layer depth MLD, and potential temperature Θ) and season (DJF, MAM, JJA, SON). Error bars represent the 95% confidence interval obtained from 2000 spatial bootstrap realizations. The non-overlapping confidence intervals for MLD relative to other drivers in panel (a), and for surface heat flux and shortwave radiation relative to wind amplitude in panel (b), confirm the statistical robustness of the dominant causal pathways identified in the main text.

all other variables. Results are shown in Supplemental Fig. S11. The partial correlation maps confirm the dominance of thermodynamic forcing — surface heat flux and buoyancy flux — over wind-driven mechanical mixing in controlling MLD variability across the ACC, consistent with the causal analysis in the main text.

8 PCMCI Causal Analysis of Mixed-Layer Depth Variability

Supplemental Fig. S12 shows the full spatial maps of PCMCI causal coefficients between MLD anomalies and its physical drivers across the Southern Ocean. These results complement the seasonal mean analysis shown in Fig. 5 of the main text and demonstrate the spatial coherence of the thermodynamic control on MLD variability throughout the ACC.

9 Biomass-Weighted Integration of Light and Iron Limitation Factors

The depth-integrated light and iron limitation terms used in the analysis are computed as biomass-weighted vertical means, emphasising the depths where phytoplankton are most concentrated:

$$\text{Light}_{\text{weighted}} = \frac{\int \text{Light}(z) \text{biomass}(z) dV}{\int \text{Chl}_a(z) dV}, \quad (4)$$

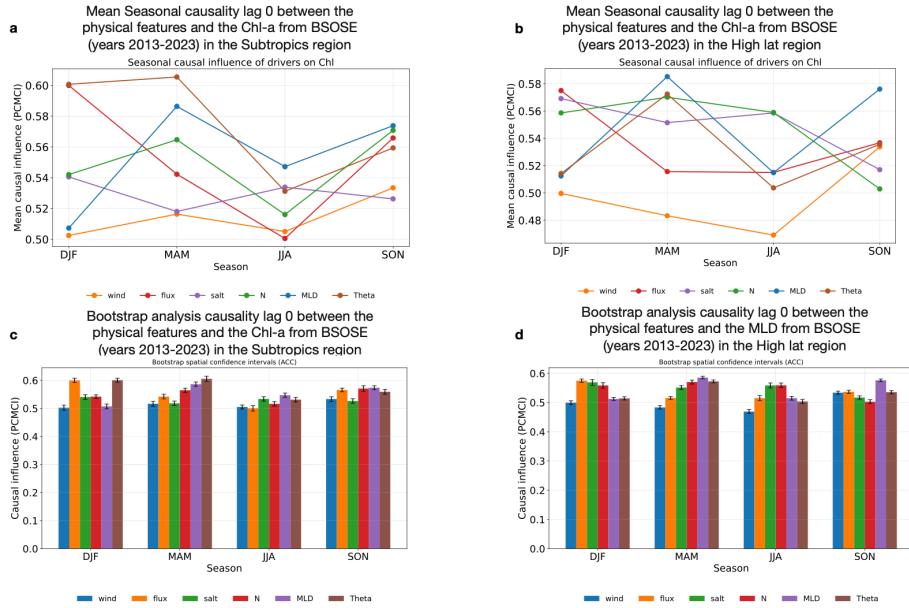


Fig. S8 Seasonal PCMCI causal analysis for the subtropical (a, c) and high-latitude (b, d) regions of the Southern Ocean. Panels (a) and (b) show the mean seasonal causal influence of individual physical drivers (wind amplitude, surface heat flux, salinity, buoyancy frequency N , mixed-layer depth MLD, and potential temperature Θ) on Chl- a anomalies; panels (c) and (d) show the corresponding 95% bootstrap confidence intervals obtained from 2000 spatial bootstrap realisations. In contrast to the ACC (Supplemental Fig. S7), no single driver achieves consistently dominant causal strength in either region across all seasons, indicating that chlorophyll variability outside the ACC reflects spatially heterogeneous contributions from multiple environmental factors.

$$\text{Fe}_{\text{weighted}} = \frac{\int \text{Fe}(z) \text{biomass}(z) dV}{\int \text{Chl}(z) dV}. \quad (5)$$

Both quantities are dimensionless scores on $[0, 1]$, where values closer to 1 indicate weaker limitation, as defined in the BSOSE biogeochemical module.

10 Markov Discovery Algorithms

Markov discovery algorithms reconstruct the causal structure of a system from observational data by exploiting conditional independence relationships implied by graphical models. In a causal graph, each variable is assumed to be independent of all non-descendant variables given its direct parents — the Markov property. Algorithms in this class identify causal connections by systematically testing whether statistical dependencies between variables vanish once appropriate conditioning variables are included. Classical constraint-based approaches such as the PC algorithm iteratively remove edges that satisfy conditional independence criteria and orient the remaining links according to logical consistency rules. However, their direct application to time-series data is limited because geophysical variables exhibit strong autocorrelation

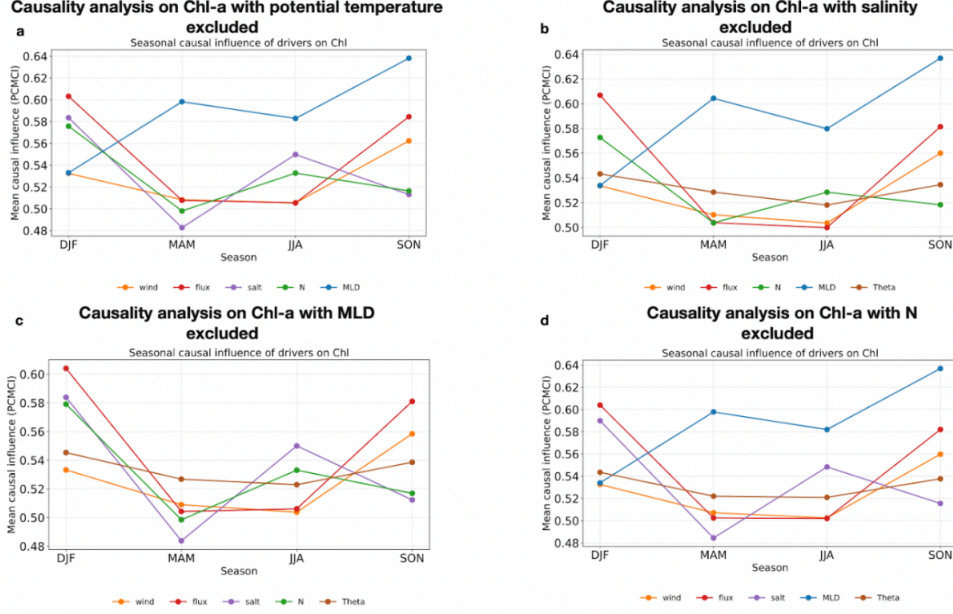


Fig. S9 Sensitivity of PCMCI causal results to predictor selection. Each panel shows the seasonal mean causal influence of the remaining drivers on Chl-*a* anomalies when one variable is excluded from the analysis: potential temperature excluded (a), salinity excluded (b), MLD excluded (c), and buoyancy frequency excluded (d). The dominant causal structure is preserved in all cases, demonstrating robustness to variable selection.

and lagged interactions across multiple timescales. Extensions to temporal settings, including PCMCI, adapt the Markov discovery framework by explicitly incorporating time-lagged variables and restricting causal directions to follow temporal ordering. This allows the reconstruction of directed dynamical networks in which links represent information transfer across time rather than contemporaneous statistical association.

11 Computation of Causal Strength in the PCMCI–GPDC Framework

For a tested causal link $X_i^{t-\tau} \rightarrow X_j^t$, the PCMCI algorithm quantifies causal strength using the Momentary Conditional Independence (MCI) statistic computed via the Gaussian Process Distance Correlation (GPDC) test. The conditioning set \mathbf{Z} is the union of the parents of the response variable $\hat{\mathcal{P}}(X_j^t)$ and the parents of the driver variable $\hat{\mathcal{P}}(X_i^{t-\tau})$ identified during the PC stage, excluding $X_i^{t-\tau}$ itself:

$$\mathbf{Z} = \hat{\mathcal{P}}(X_j^t) \cup \hat{\mathcal{P}}(X_i^{t-\tau}) \setminus \{X_i^{t-\tau}\}.$$

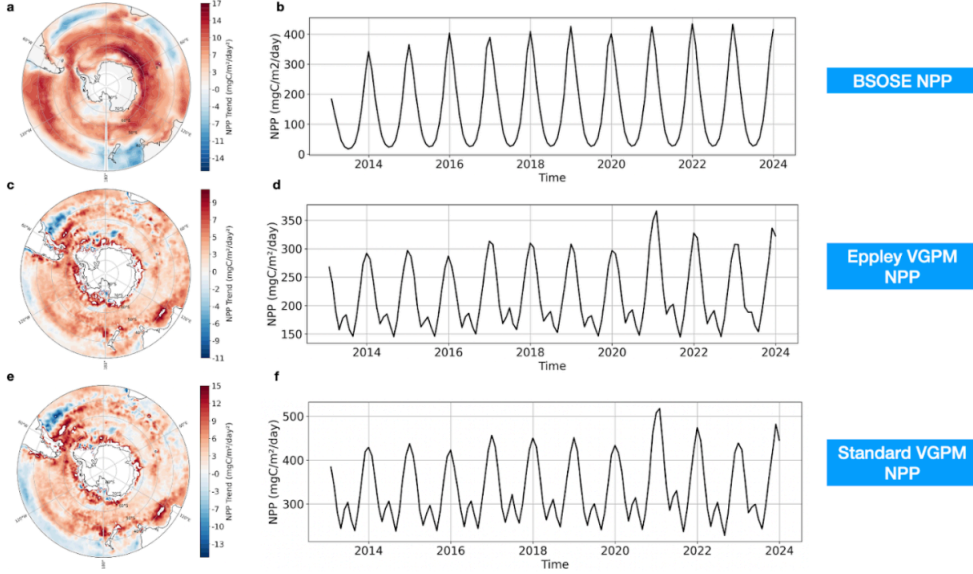


Fig. S10 Trends in net primary production (NPP) from BSOSE (a, b), the Eppley-VGPM satellite product (c, d), and the standard VGPM product (e, f). Left panels show the spatial distribution of the linear trend; right panels show the corresponding basin-wide spatial mean time series. The agreement in sign and seasonal structure across independent products confirms that Southern Ocean greening is accompanied by enhanced biological carbon fixation.

Conditioning on the parents of the driver, rather than on all lagged variables, is what distinguishes PCMCI from standard Granger causality: it removes the autocorrelation of X_i itself before testing whether $X_i^{t-\tau}$ carries additional information about X_j^t , thereby avoiding inflated false-positive rates. The resulting MCI statistic measures the strength of direct conditional dependence and should not be interpreted as a mechanistic causal effect, but as evidence of directed statistical dependency in a dynamical system. The computation proceeds in three steps.

Step 1: Removal of conditioning effects.

The shared predictability explained by \mathbf{Z} is removed from both the driver and the response by non-parametric Gaussian process (GP) regression. Let

$$X = X_i^{t-\tau}, \quad Y = X_j^t,$$

and let $\mathbf{Z} = (\mathbf{z}_1, \dots, \mathbf{z}_T)^\top$ denote the $T \times m$ matrix of conditioning observations. GP regression is fitted separately to predict X and Y from \mathbf{Z} at the training points:

$$\hat{\mathbf{x}} = K_{ZZ}(K_{ZZ} + \sigma_X^2 I)^{-1} \mathbf{x}, \quad \hat{\mathbf{y}} = K_{ZZ}(K_{ZZ} + \sigma_Y^2 I)^{-1} \mathbf{y},$$

where K_{ZZ} is the $T \times T$ kernel matrix with entries $[K_{ZZ}]_{pq} = k(\mathbf{z}_p, \mathbf{z}_q)$ and $k(\cdot, \cdot)$ is a radial basis function (RBF) kernel. GP regression is used because it makes no

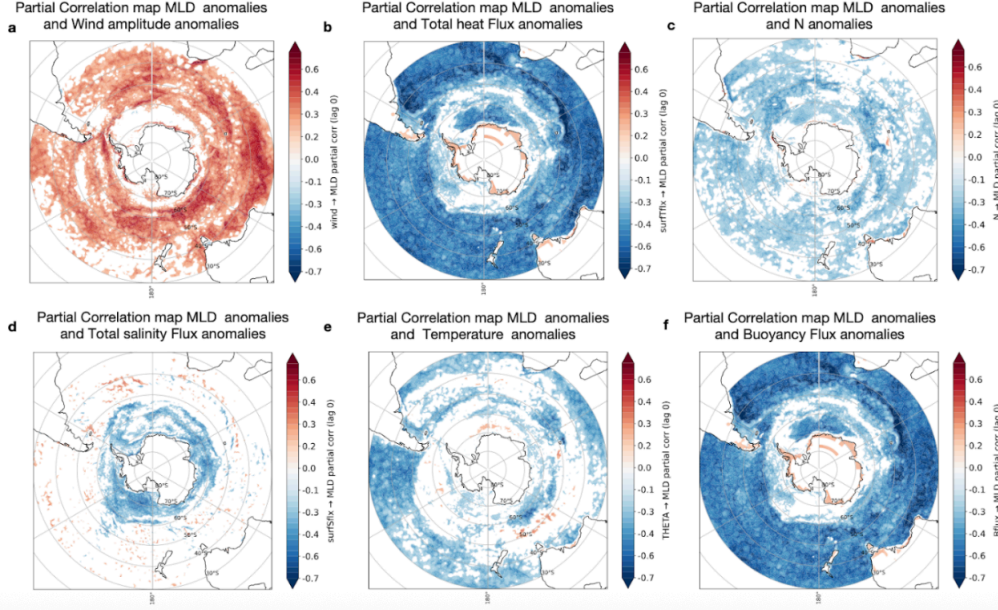


Fig. S11 Partial correlation maps between MLD anomalies and physical drivers. Shown are partial correlations with wind amplitude (a), total surface heat flux (b), buoyancy flux (c), salinity (d), potential temperature (e), and buoyancy frequency (f). Negative partial correlations with heat flux and buoyancy flux across the ACC are consistent with thermodynamic forcing as the dominant control on mixed-layer depth, while the weaker and less spatially coherent wind correlation supports a secondary role for mechanical mixing.

assumption about the functional form of the dependence, allowing it to capture non-linear and regime-dependent relationships between the conditioning variables and the target. The residuals

$$\varepsilon_X = \mathbf{x} - \hat{\mathbf{x}}, \quad \varepsilon_Y = \mathbf{y} - \hat{\mathbf{y}}$$

represent the portions of X and Y that remain unexplained by the conditioning set. If $X_i^{t-\tau}$ truly causes X_j^t , a residual dependence between ε_X and ε_Y will persist after this removal.

Step 2: Transformation to uniform marginals.

To obtain a distribution-free test statistic whose null distribution depends only on sample size T , the residuals are mapped to uniform marginals via the empirical cumulative distribution function:

$$u_p = \frac{\text{rank}(\varepsilon_{X,p})}{T+1}, \quad v_p = \frac{\text{rank}(\varepsilon_{Y,p})}{T+1}.$$

This rank transformation removes any dependence of the subsequent test on the shape of the residual distributions, making the significance threshold universal across variables with different marginal behaviour.

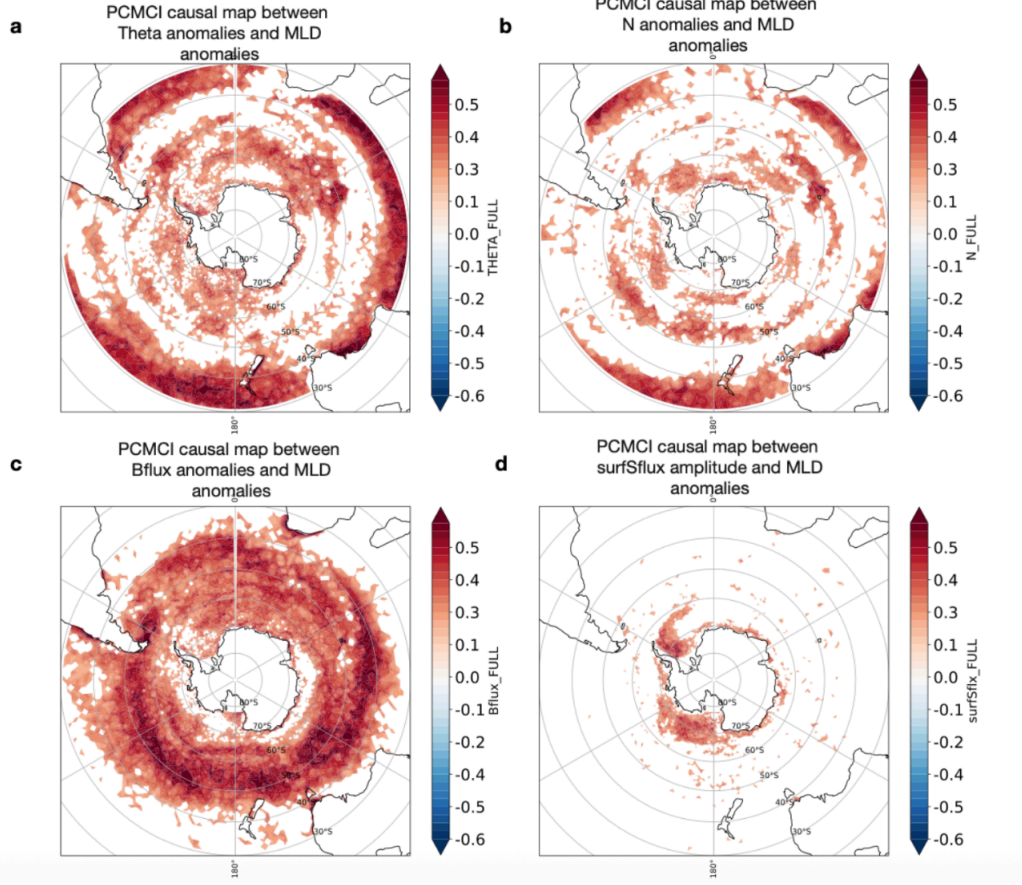


Fig. S12 Spatial maps of PCMCI causal coefficients (MCI) between MLD anomalies and physical drivers: potential temperature (a), buoyancy frequency (b), buoyancy flux (c), and shortwave radiation (d). Significant positive causal links between surface heat flux, shortwave radiation, and MLD are concentrated within the ACC, consistent with thermodynamic forcing as the primary control on mixed-layer depth in this region.

Step 3: Distance correlation and the MCI statistic.

Residual dependence between \mathbf{u} and \mathbf{v} is quantified by distance correlation. Define the $T \times T$ pairwise distance matrices

$$a_{pq} = |u_p - u_q|, \quad b_{pq} = |v_p - v_q|,$$

and their double-centred forms

$$A_{pq} = a_{pq} - \bar{a}_p - \bar{a}_q + \bar{a}.., \quad B_{pq} = b_{pq} - \bar{b}_p - \bar{b}_q + \bar{b}..,$$

where \bar{a}_p , \bar{a}_q , and $\bar{a}.$ are the row, column, and grand means of a_{pq} , respectively. The sample distance covariance and variance are

$$\widehat{\text{dCov}}^2(U, V) = \frac{1}{T^2} \sum_{p,q} A_{pq} B_{pq}, \quad \widehat{\text{dVar}}^2(U) = \widehat{\text{dCov}}^2(U, U),$$

and the MCI statistic is the normalised distance correlation

$$\text{MCI} = \widehat{\text{dCor}}(U, V) = \frac{\widehat{\text{dCov}}(U, V)}{\sqrt{\widehat{\text{dVar}}(U) \widehat{\text{dVar}}(V)}}.$$

The MCI value lies in $[0, 1]$, where zero indicates conditional independence and larger values indicate stronger direct causal coupling between the driver at lag τ and the response at time t , after removing the influence of all conditioning variables. Three properties make it well suited to this application. First, unlike partial correlation, distance correlation equals zero if and only if U and V are truly statistically independent, detecting both linear and nonlinear dependencies — important because phytoplankton responses to physical forcing in the Southern Ocean are often nonlinear. Second, the sequential conditioning in Steps 1–3 means MCI measures *direct* causal influence: it is not inflated by autocorrelation of the driver nor by indirect pathways mediated through other variables in \mathbf{Z} . Third, because the MCI is computed from rank-transformed residuals, its null distribution depends only on T and can be precomputed, avoiding computationally expensive permutation tests.