Tales from Topological Oceans



Geoff Stanley
Linacre College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Michaelmas 2017

To my parents, Jan and Len, and my brother, Dave.

There is water at the bottom of the ocean.

-David Byrne

Acknowledgements

First and foremost, thanks go to my supervisor, David Marshall. I'm incredibly lucky to have the chance to work with you, and grateful for it. You gave me freedom to explore ideas even wilder than presented here, and guided me back when I'd gone into deep water. Thank you for your mentorship, and support, and insights over the years.

Thanks to Tim Dowling for getting me thinking outside the ocean, and teaching me a totally new perspective on geophysical fluid dynamics. Your enthusiasm is infectious. "I really like the ocean, I do not care that it is kicking my butt at the moment." I couldn't agree more.

To Francis Poulin and Marek Stastna, thank you for dunking the seed of fluid dynamics in my brain. I should have taken continuum mechanics before my final year of undergrad! To Oleg Saenko and Andrew Weaver, thank you for giving me the best start in oceanography that anyone could hope for. Thanks also to Neil Swart for teaching me how cool the Southern Ocean is. Thanks to Andrew Wells for all the eclectic scientific help over the years—you know how to solve a good problem, sir. Thanks also to Helen Johnson, for helping keep my eye on the broader side of physical oceanography, and generally being amazing.

To Paul Dellar and Chris Hughes, who examined this thesis, thank you for your insights and careful analysis. Also thanks to Laure Zanna and Tim Palmer for examining my transfer and confirmation reports, for your consideration and making sure I was on track. Looking back on the years, Laure, thank you for all those light-hearted coffee chats that morphed into serious check-ins. You might not have realized you were doing this, but I'm especially grateful for your earnest advice.

I'm very lucky to have spent a summer in Woods Hole on the GFD program. Thanks to everyone who made that possible, and who made it fun and stimulating. Thanks especially to Norm Lebovitz for providing an interesting research problem and guiding me through it with all the skill

and rapidity that summer required. I hope to return to those problems with more time.

Life in the Oxford Ocean-Climate group has been a terrific journey. In particular, Liam Brannigan, Ed Doddridge, Graeme MacGilchrist, Tomos David, and Laura Cimoli: I'm so glad to have shared it with you.

To my family, I am eternally indebted. To my parents, Jan and Len, thank you for teaching a young Geoffrey mathematics, and for inspiring me to pursue knowledge. To my brother, Dave, thank you for all the wisdom you've shared with me through the years, and for paving the way.

I'm truly grateful for the wonderful friends I've made in Oxford. I won't name names, but you know who you are. Thank you for all the positive transformations in life that we've shared, that were only possible under your wing, my dear Peletron. Leaving Oxford will be one of the hardest things I'll do.

I'm grateful for financial support from the Clarendon Fund, the Canadian Alumni Scholarship, Linacre College, and Atmospheric, Oceanic, and Planetary Physics.

Finally, thanks to the progressive rock band Yes for inspiring the title of this thesis. Your album, "Tales from Topographic Oceans" is widely regarded as one of your most bombastic moments, but there are some lovely segments (comparisons with this thesis withheld).

Abstract

Functional relationships between variables can offer great illumination to hard problems. In geophysical fluid dynamics these often arise via alignment of gradients, in which case the functional relationship is multivalued; its branches are determined by a topological analysis of the connectedness of level sets, accomplished by the Reeb graph.

The neutral tangent plane, along which mixing in the ocean predominantly occurs, is defined by $\nabla_n \rho = \rho_p \nabla_n p$, thus a multivalued function relates the *in-situ* density ρ and the pressure p. With topological analysis, a new approximately new surface is presented, called a topobaric surface. This surface possesses an exact geostrophic stream function, namely the exact geostrophic stream function for neutral surfaces, which McDougall (1989) proved must exist. A closed form expression for this stream function is presented.

For a steady, inviscid, buoyancy conserving fluid, a multivalued function relates the potential vorticity Q and the Bernoulli potential B on buoyancy surfaces. In the Antarctic Circumpolar Current (ACC), level sets of B are dominated by one circumpolar contour, so Q is studied as a single-valued function of B. A tight linear relationship is discovered. Its slope, appropriately non-dimensionalised, provides information about the degree of PV homogenisation on a buoyancy surface, and about the shear stability of the flow regime via Arnol'd's stability theorems.

The real ocean does not obey these assumptions, and Q is not exactly materially conserved. In the time-mean ACC, it is found that Q can vary along streamlines by up to 50%. The causes for these fluctuations are studied, revealing a balance between two large and opposite-signed forcings—mean vertical advection and eddy interfacial form stress—with friction significant near topography. Variation of Q along B contours adds complexity on top of a background state defined by constant Q along B contours.

Contents

1 Introduction **1**

1.1 The Southern Ocean 1
1.2 Neutral surfaces 7
1.3 Potential Vorticity 8
1.3.1 Potential vorticity conservation 9
1.3.2 Implications of PV conservation 11
1.3.3 Potential vorticity homogenization 11
1.3.4 Inhomogeneous potential vorticity 13
1.4 Topology 13
1.4.1 Contour alignment 13
1.4.2 The Reeb Graph 16
1.5 Data 18
1.6 Thesis outline 19
1.A Mathematical definitions 20
1.A.1 Functions 20
1.A.2 General topology 22
1.A.3 Graph Theory 24

2 The topology of neutral surfaces **26**

2.1 Introduction 26
2.2 Theory 28
2.2.1 Background and existence of neutral surfaces 28
2.2.2 Neutral surface characterisation 30
2.2.3 The exact geostrophic stream function on a neutral surface . . 31
2.2.4 The topology of neutral surfaces 34
2.2.5 The topology of neutral surface geostrophic stream functions . 40
2.2.6 Discussion 41
2.2.7 Boussinesq models 42

2.2.8 Summary of key equations . . . 42
2.3 Methods . . . 42
2.3.1 Initial guess . . . 43
2.3.2 Calculation of the Reeb graph . . . 44
2.3.2.1 Data preparation . . . 44
2.3.2.2 The Reeb graph from ReCon . . . 46
2.3.2.3 Reeb graph simplification . . . 47
2.3.3 Multisplines to fit $\hat{\pi}$ with matching conditions . . . 50
2.3.4 Integration to find $\hat{\rho}$ . . . 52
2.3.5 Root finding of $\tilde{z}$ . . . 53
2.3.6 Iteration . . . 53
2.4 Results . . . 53
2.4.1 Initial surface . . . 53
2.4.2 Comparison with other approximately neutral surfaces . . . 58
2.5 Conclusions . . . 62

3 Ertel potential vorticity versus Bernoulli potential on approximately neutral surfaces in the Southern Ocean 65
3.1 Introduction . . . 65
3.2 Theory . . . 69
3.2.1 Mathematical stage-setting . . . 69
3.2.2 Rossby Wave Phase Speed . . . 70
3.2.3 Mach number . . . 74
3.2.4 Triad relations between Q, B, and Ψ . . . 77
3.2.5 Orthobaric density . . . 77
3.3 Q versus B correlations spanning the Antarctic Circumpolar Current 80
3.3.1 Histograms of Q vs B . . . 80
3.3.2 Linear regression . . . 83
3.4 Local Q versus B correlations . . . 87
3.4.1 Q the dependent variable . . . 87
3.4.2 B the dependent variable . . . 90
3.5 Discussion and conclusion . . . 93
3.A Rossby wave vertical modes . . . 94
3.A.1 Derivation in depth coordinates . . . 95
3.A.2 Derivation in buoyancy coordinates . . . 96
3.A.3 Comparison of methods . . . 98

4 Non-conservation of potential vorticity in the Southern Ocean 101
4.1 Introduction . . . 101
4.2 Theoretical background . . . 106
4.2.1 Linear vorticity balance . . . 107
4.2.2 Approximate linear vorticity balance . . . 108
4.2.3 Potential vorticity balance . . . 109
4.3 The structure of PV along streamlines . . . 111
4.3.1 Thickness-weighted average PV . . . 112
4.3.1.1 Defining Buoyancy . . . 112
4.3.1.2 Calculating PV . . . 114
4.3.2 Hybrid Eulerian-mean PV . . . 114
4.3.3 Unweighted Eulerian-mean PV on potential density surfaces . 116
4.3.4 Results: PV along streamlines . . . 117
4.4 The forcing of PV along streamlines . . . 122
4.4.1 Time-mean model equations . . . 122
4.4.1.1 Momentum equation . . . 122
4.4.1.2 Buoyancy equation . . . 123
4.4.1.3 Time-mean PV budget . . . 124
4.4.2 Integration of PV forcing along time-mean geostrophic streamlines . . . 124
4.4.3 Results . . . 128
4.4.4 Discussion . . . 130
4.5 Conclusions . . . 134
4.A Where does potential vorticity live on a C grid? . . . 138
4.B Thickness-weighted average momentum, thickness, and PV budgets . 142
4.B.1 Theory . . . 142
4.B.1.1 Buoyancy-coordinate thickness equation . . . 144
4.B.1.2 TWA thickness equation . . . 144
4.B.1.3 Buoyancy coordinate momentum equation . . . 144
4.B.1.4 TWA momentum equation . . . 145
4.B.1.5 TWA Potential Vorticity equation . . . 146
4.B.2 Methods . . . 147
4.B.2.1 TWA numerics for Z level data . . . 147
4.B.2.2 Unaveraged thickness equation . . . 148
4.B.2.3 TWA momentum and thickness equations . . . 149
4.B.3 Results . . . 150

4.B.4 Helmholtz for the thickness equation 153
4.B.5 Discussion 154
4.C Alternative derivation of the eddy linear vorticity balance 156
4.C.1 Style 1: Integration between isopycnals 156
4.C.2 Style 2: Fast, separable derivation 159

5 Conclusion **161**
5.1 Key results 161
5.2 Future work 164
5.3 Concluding remarks 165

Bibliography **167**

List of Figures

1.1 Schematic of the Southern Ocean circulation. Westerly winds accelerate the Antarctic Circumpolar Current (yellow arrow) and cause northward Ekman transport; these latitudes are free of continental barriers, so the Ekman transport divergence must be balanced by upwelling, creating isopycnals rising towards Antarctica on the left. Curly arrows at the sea-surface indicate the buoyancy flux, and turbulent mixing is indicated by semi-circles around isopycnal interfaces. Two overturning cells are indicated, one involving Upper Circumpolar Deep Water (UCDW) upwelling and flowing north, another involving Lower Circumpolar Deep Water (LCDW) and North Atlantic Deep Water (NADW) upwelling and flowing south, forming Antarctic Bottom Water (AABW). On the right, variations to the sea-surface height and isopycnal depths are indicated in response to bottom topography. Figure reproduced from Olbers and Visbeck (2005). 4

1.2 Visualization of an island within a volcanic crater lake (left) and its contour tree (right). Contours of constant height are shown. Critical points are labelled with letters. A and B are maxima, C and D are saddle points, and E and F are minima. Figure reproduced from (Carr et al., 2010). 16

2.1 The pressure [dbar] of the the $\sigma_1 = 31.9$ isopycnal, where σ_1 is the potential density referenced to 1000 dbar, crudely constructed, for illustrative purposes only, using the 2004—2006 climatological fields from the OCCA of salinity and potential temperature. Grey denotes places where the isopycnal has outcropped or incropped, and black, land. Six contours of the pressure are shown: 700 dbar (thick black), 800 dbar (thin black), 900 dbar (thick grey), 1000 dbar (thin grey), 1100 dbar (thick white), and 1200 dbar (thin white). If this were a neutral surface, the specific volume on the 1100 dbar contour must be the same in both the southern Pacific and southern Indian oceans, whereas it can be different on the 1200 dbar contour in the same basins. 35

2.2 The depth, $\underset{\sim}{z}$, on the $\sigma_2 = 37.0$ isopycnal (now with $\underset{\sim}{z} < 0$). Two level sets of $\underset{\sim}{z}$ are shown: -2000 m (black) and -2250 m (white). The isopycnal has outcropped or incropped in grey regions. Note the saturated colourbar. Letters indicate some nodes of the simplified Reeb graph in Fig. 2.4, and are placed at the corresponding critical points: extrema of $\underset{\sim}{z}$ for leaf nodes (lower case), or saddle points of $\underset{\sim}{z}$ for internal nodes (upper case)). 54

2.3 The regions associated with each arc of the simplified Reeb graph of the isopycnal $\sigma_2 = 37.0$. Distinct colours are distinct regions, and grey indicates where the isopycnal has outcropped or incropped, or is disconnected from the main ocean. Letters are as in Fig. 2.2, and regions are coloured matching their arcs in Fig. 2.4. 55

2.4 A layout of the simplified Reeb graph of the depth $\underset{\sim}{z}$ of the isopycnal $\sigma_2 = 37.0$. Nodes are positioned according to the longitude and depth of their associated (critical) point in physical space, and labelled except where they overlap closely. The black and white level sets of Fig. 2.2 are indicated here as dashed lines. There are two cycles. The first is around Madagascar, given by nodes D, F, H, I, D. The second is around New Zealand, and given by nodes A, J, L, ..., H, G, A. 56

2.5 Two dimensional histograms of depth $\underset{\sim}{z}$ on the $\sigma_2 = 37.0$ isopycnal, vs (left) the *in-situ* density $\underset{\sim}{\rho}$, or (b) the partial derivative of *in-situ* density with respect to depth, $\underset{\sim}{\rho}_z$. Colour indicates the number of grid points in a 2D bin, going from black (1 data point per bin) through red to white (maximum data points per bin), and grey where 0 data points. In the left panel, a global quadratic fit of $\underset{\sim}{\rho}$ to $\underset{\sim}{z}$ was subtracted from $\underset{\sim}{\rho}$ (otherwise $\underset{\sim}{\rho}$ vs $\underset{\sim}{z}$ is hardly distinguishable from a straight line). A zoomed in region is also shown in the left panel. The branches of $\hat{\rho}$ are plotted in the left panel, offset by $-0.02\,\mathrm{kg\,m^{-3}}$. The branches for $\hat{\pi}$ are rather more messy and not shown. The topobaric surface algorithm finds new depths $\underset{\sim}{z}$ such that $\underset{\sim}{\rho} = \hat{\rho}(\underset{\sim}{z})$. Thus the left panel in a similar figure for the topobaric surface would simply show the white lines shown here. However, the right panel retains scatter, because of non-zero neutral helicity. See text for further discussion. 57

2.6 The change in depth between the topobaric surface (after four iterations, converged) and the $\sigma_2 = 37.0$ isopycnal from which the former was initiated. Red indicates that the topobaric surface is shallower, blue that it is deeper, and grey that either surface has outcropped or incropped or is not connected to the main ocean. 59

2.7 Comparison of zonal density gradient error on four approximately neutral surfaces: (a) potential density referenced to 2000 m, (b) neutral density, (c) ω-surface, and (d) topobaric surface. Colour shows $\log_{10}|\epsilon^{(x)}|$, red where $\epsilon^{(x)} > 0$ and blue where $\epsilon^{(x)} < 0$. The area-weighted L_1 and L_2 norms of $\epsilon^{(x)}$ are listed above each panel, where contributions are taken only from the largest connected region (for example, this excludes the Arctic in the case of potential density). Note neutral density is not defined north of 64°N (and the same goes for the ω-surface, since I used the neutral density surface as its initial guess). 60

3.1 The nine σ_ν surfaces examined in this study, plotted as grey curves in the meridional plane at 240°E longitude, for January 1–5, 2010. The corresponding values, from shallowest to deepest are 27.1, 27.2, 27.3, 27.4, 27.5, 27.6, 27.7, 27.75 and 27.8. The Ertel potential vorticity, Q [$\mathrm{m^{-1}\,s^{-1}}$] for this cross-section is shown in colour. A perturbation associated with a large eddy can be seen in the vicinity of -56° latitude. The ocean floor is indicated in black. Large PV is confined to a rather small area near the sea-surface. 81

3.2 Two dimensional histograms of potential vorticity vs Bernoulli potential in the ACC on orthobaric density surfaces on 22–26 March, 2006. Nine surfaces are shown, from $\sigma_\nu = 27.1$ to 27.8. The brightness maps linearly to the number of data points per 2D bin. Surfaces $\sigma_\nu = 27.1$, 27.2, and 27.3 are progressively offset by $-1 \times 10^{-11}\,\mathrm{m^{-1}\,s^{-3}}$; the relative zero is indicated with dotted lines at right. Left panel: all data within the region of interest based on time-mean sea-surface height. Right panel: also excluding data with $p < 300\,\mathrm{dbar}$; white lines indicate the line of best fit, with σ_ν values indicated. Note the ordinate is flipped to facilitate comparison with the northern hemisphere. 82

3.3 Time-series of the slope of Q vs B, i.e. q_1 in (3.19), for many orthobaric surfaces (labelled at left) over the whole ACC region of interest, excluding where the pressure is less than 300 dbar. 83

3.4 Time-series of the Rossby wave phase speed at 55°S associated with the frame shift α, for each orthobaric density surface. The orthobaric density for each surface is shown at right; next to this in parentheses is the offset value for that surface, in $\mathrm{cm\,s^{-1}}$. 85

3.5 Two dimensional histograms of potential vorticity vs (α-shifted) Bernoulli potential in the ACC region of interest, excluding $p \leq 300\,\mathrm{dbar}$, on $\sigma_\nu = 27.2$, 22–26 March 2006. Left panel: as in the right panel of Fig. 3.2. Right panel: Q vs $B^{(\alpha)}$, according to (3.14) and (3.20) and (3.21). The line of best fit is shown in white, and the R^2 coefficient of determination is indicated above. 86

3.6 Time-series of the inverse Mach number for each orthobaric density surface, offset by 0.5 per surface. The orthobaric density for each surface is indicated; next to this in parenthesis is the offset value for that surface. 87

3.7 The slope, $\mathrm{d}Q/\mathrm{d}B = q_1$, from simple linear regression of Q to B using data local to each water column, on the $\sigma_\nu = 27.3$ surface on 22–26 March, 2006. Note the nonlinear colour scale. 88
3.8 The Rossby wave zonal phase speed [m/s] on $\sigma_\nu = 27.3$, 22–26 March 2006: total (upper panel) or intrinsic (lower panel). Data local to each water column is used to fit Q to B and A as in (3.20), to obtain α. . 89
3.9 The inverse Mach number, "Ma"$^{-1}$ from (3.24), using the smoothed α field corresponding to the Rossby wave phase speed of Fig. 3.8, and simple linear regression of A to B to get A_1. Data is local to each water column, on the $\sigma_\nu = 27.3$ surface on 22–26 March, 2006. 90
3.10 The Rossby wave zonal phase speed [m/s] on $\sigma_\nu = 27.3$, 22–26 March 2006: total (upper panel) or intrinsic (lower panel). Data local to each water column is used to fit B to Q and A as in (3.29), to obtain α. . 91
3.11 The inverse Mach number, "Ma"$^{-1}$ from (3.24), using simple linear regression of B to Q and A to get α, and of A to B to get A_1. Data is local to each water column, on the $\sigma_\nu = 27.3$ surface on 22–26 March, 2006. 92
3.12 Comparison of the Rossby wave phase speed [m/s] around Drake Passage on 22–26 March 2005. The eigenvalue with the largest imaginary part (corresponding to the most unstable mode) is selected using depth (a,b) and orthobaric density (c,d) as the vertical coordinate. The bottom boundary condition is Neumann (b,c) or Dirichlet (a,d) corresponding to zero vertical flow (a,c) or zero horizontal flow (b,d). Each pixel is a water column, and grey pixels indicate all eigenvalues are real. Panels (e) and (f) use local estimates of the reference frame shift α on the orthobaric surface $\sigma_\nu = 27.3$, with Q as the dependent variable, (3.20), in (e), or B as the dependent variable, (3.29) in (f). 99

4.1 Figure 7 from Gille (1997). "Forcing terms responsible for changes in PV... The momentum advection and tracer advection terms refer to transient eddy fluxes, while the potential/in situ term results from potential and in situ density surfaces not coinciding, so that $[(\partial\bar{\boldsymbol{u}}/\partial z)\cdot\nabla]\bar{\sigma}$ is not zero. All terms are set to 0 at 180°W. Together tracer advection, tracer evolution, and momentum advection capture the leading variations in PV, while momentum and tracer diffusion play significant but smaller roles." Here σ is potential density referenced to 1000 dbar. PV is in $1 \times 10^{-8}\,\mathrm{kg\,s^{-1}m^{-4}}$. The streamline has an average depth near 1 km, and passes north of Kerguelen and deflects strongly northward downstream of Drake Passage. 105

4.2 Mean vertical Ekman velocity [metres / year, orange/purple] and geostrophic streamlines on the buoyancy surface $\overline{b} = b(133.25°\mathrm{E}, 54°\mathrm{S}, -1000\,\mathrm{m}, 18-22$ `Sept` $2008)$. The 13 streamlines shown here pass 133.25°E at a somewhat irregular spacing, crossing at 62, 60.5, 59, 58, 57, 56, 55, 54.5, 53.5, 52.4, 51.5, 50.7, 50°S (rounded). On streamlines, colour (blue/black/red) indicates the relative change of Eulerian mean PV from its mean on each streamline, and the colourbar is saturated. The 2000 m isobath is contoured black, and the 3000 m isobath grey. . . . 118

4.3 Comparison of different PV definitions and streamlines. **Top panel**: The unweighted Eulerian PV (4.1) on potential density surfaces (green), or the TWA PV, (4.3) on buoyancy surfaces (red) or the Eulerian PV (4.25) on buoyancy surfaces (blue). Five geostrophic streamlines are shown, passing 133.25°E at 60.5°S, 59°S, 57°S, 55°S, 52.4°S. The blue curves correspond to the 4th, 7th, 9th, 11th, and 12th geostrophic streamlines, from north to south, in Fig. 4.2. Alternate streamlines are darkened, for clarity. **Bottom panel**: The Eulerian PV along geostrophic streamlines (blue, as in the top panel), or along rotational streamlines (thick grey). 119

4.4 PV (thick black) along the (A) 2nd, (B) 4th, (C) 7th, or (D) 12th streamlines in Fig. 4.2 from north to south. They cross 133.25°E at (A) 50.7°S, (B) 52.4°S, (C) 55°S, or (D) 60.5°S. The PV is very well approximated by $(f + v_x - u_y)b_z$, which is decomposed as (red) only f varies, (blue) only b_z varies, or (grey) only $v_x - u_y$ varies (the other terms are set to their mean values along the streamline). The abscissa is distance along the streamline, and the first crossing of various longitudes are indicated on top. 120

4.5 PV (solid black) and its forcing integrated along the same four streamlines as in Fig. 4.4. Streamline A is the most northerly, D the most southerly. The PV forcings are (in order of the RHS of (4.44), from top to bottom then left to right) ageostrophic advection (solid magenta), buoyancy drift (dotted blue), eddy buoyancy fluxes (solid blue), buoyancy forcing (dashed blue), MVV advection (solid cyan), momentum drift (invisible—entirely negligible and not shown), eddy momentum fluxes (solid red), and momentum forcing (dashed red). Also shown is the integrated geostrophic advection of PV (light grey, RHS of (4.40)), which should match the PV itself, apart from numerical errors. . . . 127

4.6 Maps of selected PV forcing terms [s^{-4}]: the buoyancy tendency, the MVV advection, the eddy buoyancy flux, and the eddy momentum flux (top to bottom). Their signs are as in (4.44); thus $\breve{Q}$ increases following the geostrophic flow through red regions. Note the buoyancy tendency (top panel) has been multiplied by 10 to reveal the deep ocean. These fields are smoothed by 1/3° in both directions. The four streamlines in Fig. 4.4 and Fig. 4.5 are shown. 130

4.7 As in Fig. 4.6 but zoomed into Drake Passage and showing all terms: buoyancy tendency, buoyancy forcing, eddy buoyancy flux, MVV advection, momentum forcing, eddy momentum flux, ageostophic advection, and geostrophic advection (left to right then top to bottom). . . 131

4.8 Layout of a tracer cell on the C-grid. At the centre is the tracer grid, T. To the west, the south, and above are the zonal U, meridional V, and vertical W velocities. The components of the curl $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$ are on the upper-south, upper-west, and south-west edges of the tracer cell, and the divergence of the curl, $\nabla \cdot \boldsymbol{\omega}$ is on the upper south west corner of the tracer cell. This is where I define PV. The $\hat{i}$ (zonal), $\hat{j}$ (meridional), and $\hat{k}$ (upwards) unit vectors are indicated. 139

4.9 Caption on following page. 151
4.10 As in Fig. 4.9 but for $z_0 = -2000\,\mathrm{m}$ and only a selection of panels. . . 154
4.11 Just the shear forcing and advective error in the TWA PV budget, but here the z coordinate buoyancy advection ($\boldsymbol{v} \cdot \nabla b$) was computed using the offline velocity and buoyancy fields, on buoyancy surfaces $z_0 = -1000\,\mathrm{m}$ (top row, for comparison with Fig. 4.9), or $z_0 = -2000\,\mathrm{m}$ (bottom row, for comparison with Fig. 4.10). 155

Chapter 1

Introduction

1.1 The Southern Ocean

The opening of Drake Passage some 30 Myr ago (Meredith et al., 2011) drastically changed the world's climate, because it enabled the formation of an ocean current that could fully encircle Antarctica. This is the Antarctic Circumpolar Current (ACC), now by far the largest ocean current in the world, and one of the most turbulent, being populated at any time by hundreds of mesoscale eddies. The ACC extends to the sea-floor and must therefore navigate across complex bottom topography (Sverdrup et al., 1942, p. 607). By encircling Antarctica with a strong eastward flow, the ACC inhibits the meridional transport of tracers, including heat. With its formation, the polar Southern Hemisphere temperatures began to drop; then followed the growth of ice sheets and sea-ice on and around Antarctica, causing the south pole to grow colder still (Olbers et al., 2004).

Off the coasts of Antarctica, surface waters grew cold and began to sink. At depth, though not at the surface (Nycander et al., 2015), this cold and fresh water, called Antarctic Bottom Water, is the densest in the world, displacing other water as it spreads globally along the sea-floor. To close the circulation, Antarctic Bottom Water is returned to its origin by both diabatic upwelling in the ocean interior (mixing) and by adiabatic upwelling along isopycnals in the Southern Ocean. Northward Ekman transport cannot be balanced by southward boundary currents in the continent-free latitudes of Drake Passage, and thus must be balanced by upwelling deep waters, thus inclining isopycnals in the Southern Ocean. This describes the abyssal Meridional Overturning Circulation (MOC), the Southern Ocean portion of which is schematised in Fig. 1.1. Some of the water that upwells in the Southern Ocean instead travels northward at the surface, eventually making its way to the North Atlantic where it sinks as North Atlantic Deep Water, to return at depth to the Southern Ocean,

though some, similarly, upwells diabatically in the low-latitudes. This is the North Atlantic MOC, which contributes to Europe's relatively warm climate. In fact, the North Atlantic MOC is largely driven by the work of winds over the Southern Ocean (Toggweiler and Samuels, 1995).

Moreover, the abyssal MOC and North Atlantic MOC are intertwined as a single MOC in the modern climate: Antarctic Bottom Water upwells diabatically in the Pacific and Indian oceans to a density less than that of North Atlantic Deep Water, leading to a figure-eight type picture of the MOC (Talley, 2013). Changes in Antarctic summer sea-ice extent can decouple this circulation to the conventional picture of independent, stacked MOC cells (Ferrari et al., 2014). The Southern Ocean, providing adiabatic upwelling, is thus crucial to the global MOC. Moreover, the MOC is intrinsically linked to the ACC through thermal wind balance. By transporting tracers such as heat, salt, oxygen, and the greenhouse-gas carbon dioxide, the MOC can be thought of as the lungs of the global ocean. More directly, the ACC dominates exchanges between the Pacific, Atlantic, and Indian oceans, so changes in the ACC can potentially change the large-scale properties of individual oceans. The Southern Ocean is clearly of tremendous climatic importance.

The significance of the ACC as a key player in climatic change began 30 Myr ago but continues under anthropogenic climate change. Today, an increasing trend in the Southern Annular Mode, and with it the strength of Southern Hemisphere surface westerlies, is generally accepted (Hartmann et al., 2013). How the ACC and the associated MOC will respond to changes in the westerly winds (their driving mechanism) is a key question for physical oceanographers. Coarse resolution general circulation models suggest these circulations will accelerate (Fyfe and Saenko, 2006). However, theoretical studies suggest that the ACC transport is largely insensitive to the wind-stress (Straub, 1993; Marshall et al., 2017). This phenomenon, termed "eddy saturation", has been verified in eddy-resolving simulations (Hallberg and Gnanadesikan, 2006; Munday et al., 2013). A similar phenomenon, termed "eddy compensation" (Viebahn and Eden, 2010) applies to the MOC (Hallberg and Gnanadesikan, 2006; Spence et al., 2009; Farneti and Delworth, 2010; Munday et al., 2013). The extent to which eddies are capable of shielding the ACC and MOC to changes in wind-stress remains an active area of research, an open question.

It is also a climatically imperative question. Currently, the Agulhas current down the east coast of southern Africa is largely retroflected by the ACC, but a southward shift of the ACC may allow more Indian ocean surface water to leak into the South Atlantic. This particularly salty water could accelerate the North Atlantic MOC

in contrast to the usual belief that the North Atlantic MOC will slow (Beal et al., 2011). Changes in the MOC impact the deep ocean reservoir of heat and carbon, playing into long-term climate feedbacks. Already the outgassing of carbon dioxide from the deep ocean has accelerated (Le Quere et al., 2007; Lovenduski et al., 2008); the future of this trend is debated. As Antarctic Bottom Water holds more carbon than North Atlantic Deep Water, one way to change the ocean's total carbon content is by changing the ratio by volume of the two water masses (Marinov et al., 2006). A major determinant of this ratio is the slope of isopycnals in the Southern Ocean. This slope determines, through thermal wind balance, the baroclinic transport of the ACC, and maps the meridional surface density gradient onto the deep interior stratification of the global ocean (Nikurashin and Vallis, 2011). With steeper isopycnals, more potential energy is available, which mesoscale eddies extract through baroclinic instability.

Given this climatic importance, accurate projections of future climate change depend on accurate representation of the Southern Ocean in models. The ACC transport is commonly used as a benchmark to judge climate models—or to tune them. The modelled ACC transport amongst 18 climate models, actively used for 21'st century IPCC-class projections, ranged from 90 Sv to almost 250 Sv (Meijers et al., 2012). From observations, the canonical estimate of the ACC baroclinic transport through Drake Passage is roughly 130 Sv to 140 Sv ($1\,\mathrm{Sv} \equiv 10^6\,\mathrm{m^3\,s^{-1}}$) (Whitworth and Peterson, 1985; Cunningham et al., 2003), though recent estimates include a more substantial barotropic transport, bringing the total to 173 Sv (Donohue et al., 2016). As an understatement, there is room for improvement here; a good scientific understanding of the Southern Ocean helps.

Thus many have asked: what sets the transport of the ACC? Answers have evolved over the years. A flat-bottom ocean model required unrealistically high values for frictional parameters in order to obtain a realistic ACC transport (Hidaka and Tsuchiya, 1953). Including topography greatly alleviated this problem: the early GFDL model of a homogeneous ocean (equivalently called a barotropic ocean, for a uniform density ocean has no vertical shear of horizontal velocities, by thermal wind) reached a steady-state ACC transport of over 600 Sv with a flat bottom ocean and realistic bottom-friction, reduced to under 30 Sv with realistic topography (Bryan and Cox, 1972). This drastic reduction is readily understood in terms of the potential vorticity (PV), which in a barotropic ocean is approximately f/H, where $H > 0$ is the ocean depth, $f = 2\Omega\sin(\phi)$ is the Coriolis parameter, with $\Omega = 2\pi/(86164\text{ s})$ Earth's angular velocity and ϕ the latitude. (The relative vorticity has been neglected given the

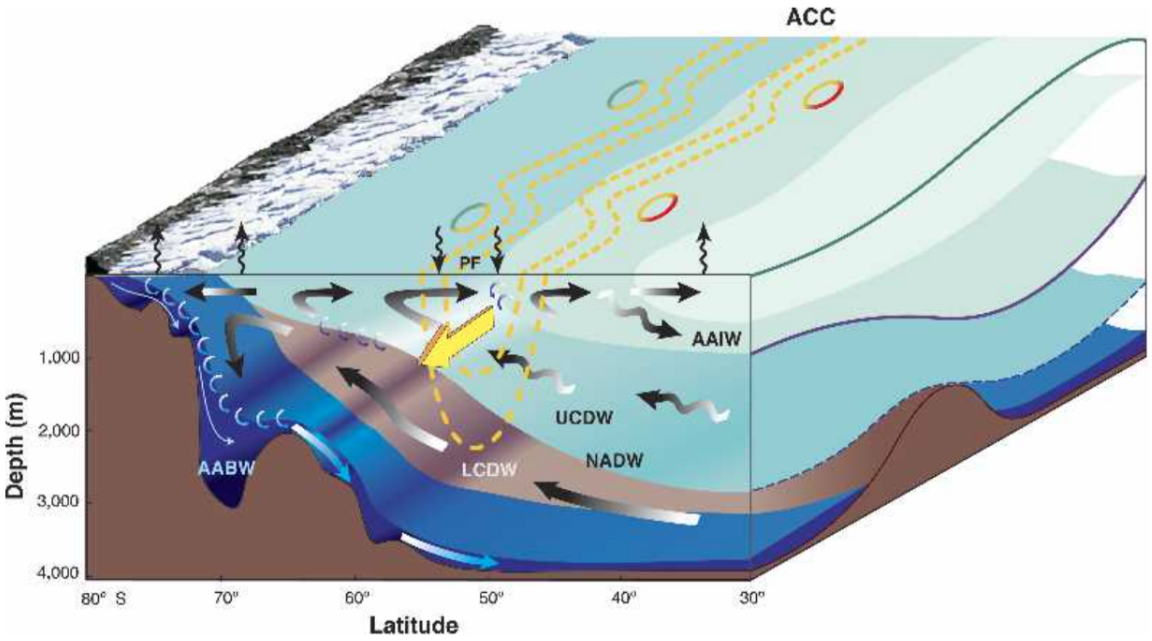


Figure 1.1: Schematic of the Southern Ocean circulation. Westerly winds accelerate the Antarctic Circumpolar Current (yellow arrow) and cause northward Ekman transport; these latitudes are free of continental barriers, so the Ekman transport divergence must be balanced by upwelling, creating isopycnals rising towards Antarctica on the left. Curly arrows at the sea-surface indicate the buoyancy flux, and turbulent mixing is indicated by semi-circles around isopycnal interfaces. Two overturning cells are indicated, one involving Upper Circumpolar Deep Water (UCDW) upwelling and flowing north, another involving Lower Circumpolar Deep Water (LCDW) and North Atlantic Deep Water (NADW) upwelling and flowing south, forming Antarctic Bottom Water (AABW). On the right, variations to the sea-surface height and isopycnal depths are indicated in response to bottom topography. Figure reproduced from Olbers and Visbeck (2005).

large-scale assumption.) The PV should be quasi-materially conserved, meaning the PV of a fluid parcel following the flow is roughly constant, and would be constant if not for mechanical forcing and dissipation (and the effect of neglecting the relative vorticity in the PV). These terms are usually small in the momentum equation, so a homogeneous ocean must nearly follow contours of f/H, the "geostrophic contours". In the Southern Ocean, circumpolar geostrophic contours north of the Antarctic shelf are rare, largely due to Drake Passage and the shallow topography of Scotia Ridge. This explains the mere 30 Sv solution above: the barotropic transport of the ACC must be small.

A stratified ocean can support baroclinic flow in which a strong surface current is somewhat decoupled from a weaker bottom flow; the stratification effectively shields

the surface currents from the steering influence of topography (Veronis, 1981).

Munk and Palmén (1951) made the early suggestion that the eastward momentum imparted to the ACC by the westerly winds should be predominantly balanced by bottom form drag—pressure gradients across bathymetry. This is schematised on the right of Fig. 1.1. The yellow arrow indicates the ACC flow direction, and an elevated sea-surface height is established upstream of elevated bathymetry. Deep isopycnals heave downwards to compensate, but the net effect is of higher pressure on the upstream side of the topographic bump, thereby transferring eastward momentum from the ACC to the solid earth, and decelerating the ACC.

Taking the idea of Munk and Palmén (1951) to its logical conclusion, Johnson and Bryden (1989) developed an analytic model in which 100% of the eastward momentum imparted by the wind is balanced by bottom form drag across topography, with eddy interfacial form stress acting as the communication mechanism for the surface wind stress to reach the bottom. Eddy interfacial form stress is part and parcel of the process of baroclinic instability in which available potential energy is released by a poleward eddy flux of buoyancy. Marshall et al. (1993) illuminated this balance between wind-stress and eddy interfacial form stress as the dominant terms of a more fundamental balance between the (wind-driven) acceleration round a streamline encircling Antarctica generating vorticity, and the lateral eddy PV fluxes (now speaking of the full baroclinic PV); the latter are accomplished by both eddy momentum fluxes (Reynolds stresses) that flux relative vorticity, and, more importantly, by eddy buoyancy fluxes that induce the eddy interfacial form stress. Johnson and Bryden (1989) used the dominant term, the poleward eddy heat flux, parameterised as proportional to both the mean zonal shear and the mean meridional temperature gradient. But thermal wind relates the mean meridional temperature gradient to the mean zonal shear, closing the model: the (given) surface wind-stress dictates the eddy interfacial form stress, which dictates the mean zonal shear in the ACC, which can be integrated up to give the total baroclinic transport. They obtained a best estimate of 96 Sv.

This discussion has, it is hoped, motivated the need for theoretical models of the Southern Ocean. By refining the questions we ask, they help us probe the nature of the real Southern Ocean, and of climate models, both of which are too complicated for complete understanding alone. But theoretical models must also be probed: do their simplifying assumptions lead us to an inaccurate view of the full system? Observations and climate models can help with this: it is a two-way street. This thesis investigates three facets underlying many theoretical models.

Many theoretical models of the Southern Ocean are based on some common assumptions. From the Navier-Stokes equations, the "primitive equations" have already combined the centrifugal and gravitational body forces, neglected horizontal accelerations by the Coriolis force arising from vertical motions and neglected vertical accelerations caused by horizontal motions (traditional approximation), assumed fluid properties vary much faster in the vertical than geometric properties (shallow-fluid approximation), and assumed the pressure is simply the weight of overlying fluid (hydrostatic balance) (Vallis, 2006). From here, the "planetary geostrophic" equations add the assumptions of geostrophic balance and incompressibility to the primitive equations. They effectively represent the large-scale oceanic flow, which is on much larger scales than the deformation radius.

Another common assumption is that buoyancy is conserved following the flow. In other words, the flow is constrained to lie in surfaces of constant buoyancy; two dimensional flow being far simpler than three dimensional flow, this is a great boon. However, the non-linear equation of state of sea-water makes a precise definition of buoyancy impossible. Surfaces of constant buoyancy are called neutral surfaces (McDougall, 1987a), but such surfaces are commonly believed to be ill-defined (McDougall and Jackett, 1988). Nonetheless, such a boon they would be, oceanographers instead use well-defined surfaces that are only approximately neutral. Isopycnals (isosurfaces of potential density) are one such approximately neutral surface, but there are many, each with their own advantages and drawbacks.

Buoyancy conservation and planetary geostrophy are a powerful combination, because the whole system boils down to a single equation for material conservation of the PV. Thus, many theoretical models of the Southern Ocean take this combination as a starting point. An even more powerful assumption to add to the mix is homogenization of PV (Rhines and Young, 1982) on approximately neutral surfaces in the Southern Ocean (Marshall et al., 1993; Marshall, 1995b). The main work of this thesis addresses these assumptions:

- Chapter 2: What are neutral surfaces, and what is the geostrophic stream function on a neutral surface?
- Chapter 3: To what degree is potential vorticity homogenized on approximately neutral surfaces in the Southern Ocean?
- Chapter 4: To what degree is potential vorticity conserved along geostrophic streamlines in the Southern Ocean?

These three main questions are given further background in this introduction. Neutral surfaces are discussed in Section 1.2, and PV in Section 1.3. From both discussions arise a common topological question, which is discussed in Section 1.4. The data used in this thesis is briefly described in Section 1.5. Appendix 1.A provides some background on mathematical topics.

1.2 Neutral surfaces

For background, neutral surfaces are discussed here, descriptively. Precise and mathematical discussions are reserved until Section 2.2.

Fluid parcels exhibit certain behavioural characteristics: they dislike moving in a particular direction, and instead they prefer to move in a plane normal to this direction. This direction is called the dianeutral vector, and the plane called the neutral tangent plane (McDougall, 1987a). The neutral tangent plane is defined everywhere in the ocean by the local gradients of salinity and temperature (used in this section as shorthands for Absolute Salinity and Conservative Temperature), such that fluid parcel, wrapped in an impermeable and perfectly insulating bag, can move in the neutral tangent plane while always having the same density as its surroundings. In other words, there is no buoyancy restoring force, and gravity has no effect upon it. Thus by "prefer", I mean they can move in this plane with no energy cost, and by "dislike", I mean that there is an energy cost associated with this motion.

Salt and temperature may vary on the neutral tangent plane, but in perfectly compensating ways: the original fluid parcel may have a different salinity and temperature than its new environment, but its *in-situ* density is now the same. As the *in-situ* density is a function of salinity, temperature, and pressure, and as the pressure is the same between the fluid parcel and its new environment, there is an implicit equation involving the salinity and temperature of the environment, and some constant parameters (the original parcel's salinity and temperature, and the pressure at the new location). Moving the parcel elsewhere on the neutral tangent plane picks up more such equations, and thus the salinity and temperature on the neutral tangent plane are functionally related.

Now, it would be convenient if these neutral tangent planes, which are defined everywhere in the ocean from local properties, could be somehow "stitched together" to form a simple two dimensional surface, extending globally throughout the ocean, called a neutral surface (McDougall, 1987a). Unfortunately, complications due to the non-linear equation of state prevent this (McDougall and Jackett, 1988). Instead, a

fluid parcel travelling a closed loop in latitude and longitude and always following the local neutral tangent plane will return to its original position at a different depth than it began; this is called the "neutral helix" (McDougall and Jackett, 1988). The change in depth is proportional to a measurable quantity called the helicity, that depends only on the local values and gradients of salinity, temperature, and pressure; only for an ocean with everywhere-zero helicity would neutral surfaces exist as simple two dimensional surfaces (McDougall and Jackett, 1988). Implicit in this description is that the fluid parcel is continually mixing with its surroundings during its journey (which can be confusing since the neutral tangent plane is defined based on insulated fluid motions).

Despite these complications, neutral surfaces remain incredibly important for oceanographers. The smallness of mechanical dissipation measured in the ocean supports the view that, locally, the flow is very nearly aligned with the neutral tangent plane (McDougall et al., 2014). A consequence of this is that both active and passive tracers spread along neutral surfaces. Layered ocean models strive to take advantage of this near two-dimensional nature of the flow, and thus require an accurate representation of neutral surfaces. Thus, for practical purpose, oceanographers resort to approximately neutral surfaces, which are well-defined, usually an iso-surface of some three-dimensional quasi-conservative density—today these are most commonly potential density (Wüst, 1935; Reid and Lynn, 1971) and neutral density (Jackett and McDougall, 1997), but other choices are available.

For a theoretical neutral surface, McDougall (1989) proved that an exact geostrophic stream function must exist, but a closed form expression for it has never been given. I derive one in Chapter 2, and this answer leads to new, fundamental ideas about neutral surfaces based on topological thinking. These new ideas are applicable to the construction of a new approximately neutral surface with an exact geostrophic stream function.

1.3 Potential Vorticity

Potential vorticity is at the core of many great theories of geophysical fluid dynamics. Planetary geostrophy and quasi-geostrophy both boil the primitive equations down to a single equation for a single variable, the PV, from which all other variables may be determined; it has proved a treasure-trove of results, including the following: the dispersion relation for Rossby waves, the nature of baroclinic instability (Eady,

1949; Charney, 1947) and shear instability more broadly (Arnol'd, 1966), PV homogenization (Rhines and Young, 1982), ventilation theory (Luyten et al., 1983), non-acceleration of the mean flow by eddies (Eliassen and Palm, 1961; Andrews, 1983), and jet focussing and mixing barriers (Dritschel and McIntyre, 2008).

Potential vorticity owes its importance to two properties: (1) it is quasi-conserved by the flow, and (2) it can be inverted to obtain the large-scale velocity and density, consistent with prescribed boundary conditions (Hoskins et al., 1985). I will not directly deal with (2), though Chapter 3 will discuss an indirect effect of the PV distribution on flow stability.

The derivation of the first property is given next, then its implications are discussed. By quasi-conservation, we mean that the material derivative of PV would be zero if not for forcing and dissipation, both of the mechanical and the thermodynamic varieties, and if not for non-zero helicity associated with the non-linear equation of state (see Chapter 2).

1.3.1 Potential vorticity conservation

The classical derivation of the potential vorticity equation by Ertel (1942) (see also Schubert et al., 2004) is achieved in three steps: first the curl of the momentum equation is taken, then the dot product with a conservative variable is taken, and finally mass conservation is employed. One could instead begin from the tracer equation, or even from the mass equation, but these derivations feel less natural: there's a certain pleasure in annihilating gradients in the momentum equation by taking the curl. With hindsight, I present a slightly cleaner and more systematic derivation that actually unites these three derivations.

The momentum equation, tracer equation, and mass equation, for a compressible fluid may be written as follows

$$\frac{\partial \boldsymbol{v}}{\partial t} + \boldsymbol{\omega}_a \times \boldsymbol{v} + \nabla K + \nabla \Phi = -\nu \nabla p + \boldsymbol{F}, \tag{1.1a}$$

$$\frac{\mathrm{D}\gamma}{\mathrm{D}t} = \dot{\gamma}, \tag{1.1b}$$

$$\frac{\mathrm{D}\nu}{\mathrm{D}t} = \nu(\nabla \cdot \boldsymbol{v}), \tag{1.1c}$$

where $\boldsymbol{v}$ is the 3D velocity, $\boldsymbol{\omega}_a = 2\boldsymbol{\Omega} + \nabla \times \boldsymbol{v}$ is the absolute vorticity, $\boldsymbol{\Omega} = \Omega(\cos\phi \hat{\boldsymbol{j}} + \sin\phi \hat{\boldsymbol{k}})$ Earth's angular velocity vector, $K = \boldsymbol{v} \cdot \boldsymbol{v}/2$ is the kinetic energy per unit mass, Φ is the geopotential, F represents external mechanical forcing and internal

viscous stresses, γ is a tracer, $\dot{\gamma}$ represents sources and diffusion of γ, $\nu = \rho^{-1}$ is the specific volume, and $D/Dt = \partial/\partial t + \boldsymbol{v} \cdot \nabla$ is the material derivative.

The potential vorticity is defined as

$$Q = (\boldsymbol{\omega}_a \cdot \nabla\gamma)\ \nu. \tag{1.2}$$

Considering the form of $\partial_t Q$, we are led to combine the equations of motion as follows:

$$\nu\nabla\gamma \cdot \nabla \times (1.1\text{a}) + \nu\boldsymbol{\omega}_a \cdot \nabla(1.1\text{b}) + (\boldsymbol{\omega}_a \cdot \nabla\gamma)(1.1\text{c}). \tag{1.3}$$

Taking the curl of (1.1a)—this pleasure remains—and adding $0 = \partial_t(2\boldsymbol{\Omega})$ yields the vorticity equation:

$$\frac{\partial \boldsymbol{\omega}_a}{\partial t} + \nabla \times (\boldsymbol{\omega}_a \times \boldsymbol{v}) = \nabla p \times \nabla\nu + \nabla \times \boldsymbol{F}. \tag{1.4}$$

Thus, the procedure (1.3) produces the following:

$$\begin{aligned}
&\nu\nabla\gamma \cdot \partial_t\boldsymbol{\omega}_a && + \nu\boldsymbol{\omega}_a \cdot \partial_t\gamma && + (\boldsymbol{\omega}_a \cdot \nabla\gamma)\partial_t\nu \\
+&\nu\nabla\gamma \cdot \nabla(\boldsymbol{\omega}_a \times \boldsymbol{v}) && + \nu\boldsymbol{\omega}_a \cdot \nabla(\boldsymbol{v} \cdot \nabla\gamma) && + (\boldsymbol{\omega}_a \cdot \nabla\gamma)\boldsymbol{v} \cdot \nabla\nu \\
=&\nu\nabla\gamma \cdot (\nabla p \times \nabla\nu) + \nu\nabla\gamma \cdot \nabla \times \boldsymbol{F} && + \nu\boldsymbol{\omega}_a \cdot \nabla\dot{\gamma} && + (\boldsymbol{\omega}_a \cdot \nabla\gamma)\nu\nabla \cdot \boldsymbol{v}.
\end{aligned} \tag{1.5}$$

The PV tendency term appears immediately in the first line of (1.5). On the second line, the vector triple product unites two terms into one:

$$\nu\boldsymbol{\omega}_a \cdot \nabla(\boldsymbol{v} \cdot \nabla\gamma) + \nu\nabla\gamma \cdot \nabla \times (\boldsymbol{\omega}_a \times \boldsymbol{v}) = \nu\nabla \cdot ((\boldsymbol{\omega}_a \cdot \nabla\gamma)\boldsymbol{v}). \tag{1.6}$$

Thus the second line in (1.5) is $\nabla \cdot (Q\boldsymbol{v})$, and the last term on the RHS modifies this into $\boldsymbol{v} \cdot \nabla Q$, whence

$$\frac{\mathrm{D}Q}{\mathrm{D}t} = \nu\nabla\gamma \cdot (\nabla p \times \nabla\nu) + \nu\nabla\gamma \cdot \nabla \times \boldsymbol{F} + \nu\boldsymbol{\omega}_a \cdot \nabla\dot{\gamma}. \tag{1.7}$$

If the fluid is unforced and non-dissipative ($\boldsymbol{F} = 0$), if the tracer γ is materially conserved ($\dot{\gamma} = 0$), and if γ is a function of p and ν alone, then the RHS of (1.7) vanishes and Q is a conserved quantity. Mechanical forcing or dissipation modifies PV by its curl in the plane of γ surfaces. Sources of γ modify PV when they vary in the direction of the vorticity vector.

1.3.2 Implications of PV conservation

Consider an unforced, non-dissipative geophysical fluid in steady state, and with no sources/sinks of buoyancy, and with Boussinesq reference specific volume ν_0. A parcel of such a fluid flows along the intersection of isosurfaces of b (generically called isopycnals) and isosurfaces of the Bernoulli potential $B = \nu_0 p + \Phi + K$. Assuming ∇b and ∇B are not exactly collinear, this defines a one dimensional streamline for the flow. The potential vorticity, since it too is conserved by the flow, must be constant on this streamline. Giving values for b and B selects a family of streamlines, each of which has a potentially different value for Q. Hence, Q is a multivalued function of b and B.

1.3.3 Potential vorticity homogenization

This multivalued function $Q(b, B)$ can be used in a hierarchy of ways. At one extreme, $Q(b, B)$ can take the form of a constant Q; with this extreme simplicity came great illumination of the nature of baroclinic instability (Eady, 1949).

At the next level, one takes $Q = Q(b)$: PV is homogenized on isopycnals. This is a decent approximation for interior ocean dynamics (away from the surface and away from topography) where there are no sources of potential vorticity. Indeed, in the absence of any injection of potential enstrophy at large-scales, a forward potential enstrophy (Q^2) cascade implies there is low PV variance on large-scales, i.e. that PV is fairly homogenized on large-scales. This argument was formally developed by Rhines and Young (1982). By assuming (i) that the eddy-flux divergence of PV has an intermediate strength, larger than dissipation and forcing but weaker than the background advection of PV by the mean flow, and (ii) that the eddy flux of PV is expressible as a diffusion of large-scale PV and that this is (in an integrated sense, around a closed streamline) down-gradient, they proved that, within a closed streamline, gradients of PV across streamlines are small.

What is meant by small? Rhines and Young (1982) compared PV data with the PV of a layer of uniform thickness $\overline{h}$, taken as the areal mean layer thickness, at rest on the β-plane, namely $\beta y/\overline{h}$, where $\beta = \mathrm{d}f/\mathrm{d}y$ is the planetary vorticity gradient. Plotting PV versus latitude revealed regions where PV roughly followed this background gradient, and other regions where PV was essentially uniform compared to this slope. Essentially, this compares the PV gradient with $\beta/\overline{h}$. This is physically questionable since "a baroclinic fluid is a moving fluid, even in the presence of viscosity" (Vallis, 2006, pp. 166), so unless these layers are flat, the fluid will not be

at rest. There are other background states one could compare to, but this still begs a physical justification. An alternative approach exploits the fact the PV gradient supports Rossby waves, whose phase speed can be compared to the local flow speed (Dowling, 2014). Indeed, this view opens up a new perspective on shear instability (Dowling, 2014) that I take up in Chapter 3.

By assuming the forcing and dissipation of PV is small, the theory of Rhines and Young (1982) applies to relatively deep, "unventilated", layers, or more specifically to closed streamlines on such layers which do not outcrop. Another theory was developed by Luyten et al. (1983) that applies to shallower, "ventilated" layers; here PV is set at the outcrop and moves into the interior under Ekman-downwelling, with homogenization achievable if the boundary conditions along the outcrop give the same PV.

Keffer (1985) tested these theories with observational maps of PV on several isopycnals in all the ocean basins. He found homogeneous PV in layers of the South Atlantic, North and South Pacific, and Indian oceans, all of which are dynamically isolated from near-surface forcing and dissipation; here the mechanism of Rhines and Young (1982) was evidently at work. In the North Atlantic, however, isopycnals tend to outcrop south of the line of zero wind-stress curl, whence Luyten et al. (1983) theory may apply, though also subpolar mode water intrudes from the north-east. Despite this atmospheric forcing, near PV homogenization was still found in the North Atlantic. In a further detailed analysis of PV in the North Pacific, Talley (1988) found homogenized PV in the subtropical gyre on isopycnals that outcrop in the subpolar gyre, and found homogenized PV across both subtropical and subpolar gyres on somewhat deeper (but not too deep) isopycnals. O'Dwyer and Williams (1997) found homogenized PV in the deep North Pacific and eastern Atlantic.

Should we expect PV to be homogenized in the Southern Ocean? There, the steeply inclined isopycnals are indeed connected to the surface, so the Rhines and Young (1982) theory might not apply. However, it is also a region of highly intense geostrophic turbulence, so perhaps the scale separation that Rhines and Young required may yet apply. Marshall et al. (1993) examined the isopycnal potential vorticity of the Southern Ocean and found gradients of PV on a given isopycnal to be small when deeper than about 400 m, with strong gradients above this as one would expect, owing to atmospheric forcing.

Models built from the principle of PV homogenization can be quite revealing, with quasi-analytic solutions for the background flow accessible. For instance, Marshall (1995b) was able to construct the entire 3D buoyancy and velocity from knowledge of

the surface buoyancy b_s along a single longitude transect of the ACC, as follows. The large-scale PV $Q = f\partial_z b$, paired with the functional form $Q(b)$, defines an ordinary differential equation; b_s thus determines the full water column buoyancy, including the bottom buoyancy b_b. Next, the Montgomery potential M at the sea-floor, M_b, is assumed to be a function of b_b. Then the hydrostatic relation, which in buoyancy coordinates is $\partial_b M = -z$, is integrated up to obtain the surface Montgomery potential, M_s. Thus in any water column c, the surface Montgomery potential is a function of the surface buoyancy, $M_s^c(b_s)$; this function implicitly depends on geographic parameters (f and the fluid depth H) and the functions $Q(b)$ and $M_b(b_b)$. Along the transect where b_s is measured, M_s is computed as above, giving a functional relation $M_s^t(b_s)$ ("t" for transect). At other fluid columns in the ACC then, an implicit equation is solved for b_s, namely $M_s^c(b_s) = M_s^t(b_s)$. By prescribing $Q = Q(b)$, the interior dynamics are made uninteresting and the entire 3D circulation problem has been reduced to 2D problems at the surface and bottom boundaries, rather reminiscent of the Eady problem. Another application of $Q = Q(b)$ is a non-linear Rossby wave equation by Marshall (2016) in the presence of a sheared background flow and variable topography.

1.3.4 Inhomogeneous potential vorticity

A level higher, one retains the B dependence: $Q = Q(b, B)$. Marshall and Stephens (2001) extended the above Marshall (1995b) model to include Q variations on a b surface. In Chapter 3 I study the variation of Q with B on b surfaces in the ACC, including the time-evolution of this Q vs B relationship.

Finally, one accepts that $Q = Q(b, B)$ is not a solution for a forced, dissipative, and time-dependent fluid, and one can investigate the structure of Q when these idealizations are relaxed. In Chapter 4 I study the structure of Q along time-mean streamlines in the ACC, where Q is forced (predominately by eddies) and no longer conserved.

1.4 Topology

1.4.1 Contour alignment

The preceding discussions of neutral surfaces (Section 1.2) and potential vorticity (Section 1.3.2) described functional relationships between variables: salt and temperature in the former case, potential vorticity and Bernoulli potential in the latter

case. These were argued descriptively, but they can also be derived mathematically. Looking ahead, (2.7) defines the neutral tangent plane, while (3.11) follows from the assumptions of Section 1.3.2 but working in buoyancy coordinates. Both these equations are of the form

$$\nabla\phi = \alpha\nabla\psi, \tag{1.8}$$

for three functions ϕ, ψ, and α, with α nowhere zero. Since $\nabla\phi$ is always orthogonal to a contour of constant ϕ, and similarly for ψ, (1.8) implies that the contour of ϕ through any point is aligned with the contour of ψ through the same point. Since this is true at all points, a contour[1] of constant ϕ must be exactly a contour of constant ψ: the two contours have the same set of points. In the simplest scenario, where there is one contour for any given value of ψ, then ϕ must be a function of ψ: specifying a value of ψ specifies a contour of ψ, which is identical to a contour of ϕ, and ϕ is constant on these points. Note, ϕ and ψ are interchangeable, so without loss of generality only contours of ψ determine whether this simple scenario applies.

In general, there are multiple disjoint contours for a given value of ψ. More precisely, the level set of ψ at some $p \in \mathbb{R}$ has multiple connected components. Now, ϕ can take different values on each contour, so ϕ is a *multivalued* function of ψ.

Seeing that connectedness of contours is important, one is led to study the topology of the system. Specifically, the Reeb graph is the mathematical tool that tracks changes in the connected components of level sets, and so is the correct way to analyse the relationship between ϕ and ψ. Before defining it (Section 1.4.2), I discuss some of its previous uses.

Though defined over 70 years ago, the Reeb graph (Reeb, 1946) has not, to the best of my knowledge, previously been used for studying multivalued functions arising from (1.8). Pure mathematics has found great use for it, for example used by Arnol'd (1957) in solving Hilbert's superposition problem (see also Arnold, 2006). With the Reeb graph, Agarwal et al. (2004) generalized the concept of elevation on Earth to more complex shapes lacking a canonical reference point (on Earth, the centre of mass), with applications to protein docking. Gorban (2013) developed the geometric side of chemical kinetics using the Reeb graph. A transition from one chemical state to another may be suggested by a decrease in the Gibbs energy, but is not realizable if any path between these states is non-monotonic in the Gibbs energy. If this were a one dimensional dynamical system, the state cannot move from one side of an

[1] When the domain is two dimensional, these contours are one dimensional in the usual sense. When the domain is an n-dimensional manifold, contours are $(n-1)$-dimensional manifolds. Nonetheless, the same argument applies.

equilibrium point to the other, even if the Gibbs energy is lower on the other side. Gorban (2013) extended these ideas to highly multidimensional systems.

As an applied tool, the Reeb graph has primarily been used in computer graphics (see Biasotti et al., 2008, for a review), initiated by Shinagawa et al. (1991). Graphical *representation* of an object using a triangular mesh, or using 2D pixels or 3D voxels, was well advanced, but these methods lacked any sort of high-level *description* of the object's shape. The Reeb graph describes an object's identity in these high-level terms. For instance, Takahashi et al. (1995) advanced the field of GIS by using the Reeb graph[2] to correctly identify peaks, pits, passes, ridges, and ravines from discretised topographic data. Representing topographic information as a function of latitude and longitude is one thing, but describing it by peaks and passes is another. This sort of high-level information gave rise to new methods to segment objects and visually render them. For example, scalar data returned by magnetic resonance imaging can be simplified using the Reeb graph, for interactive exploration and segmentation of various regions of the brain (Carr et al., 2010). Some other examples include fuel combustion, CT scans, and the probability distribution of a nucleon (Weber et al., 2007; Bremer et al., 2011), skeletal animation (Kanongchaiyos and Shinagawa, 2000), and human body segmentation (Werghi et al., 2006). The list goes on.

While better visualization certainly aids scientific discovery, topological tools can also be a direct part the scientific analysis (Bethel et al., 2009). Thus far, a related topological tool called the Morse-Smale complex has primarily been used towards this goal. For instance, in a simulation of Rayleigh-Taylor instability, Laney et al. (2006) use the Morse-Smale complex to segment and count bubbles, which gives new information on the spatial scale and evolution of the turbulence. Of oceanographic relevance, Xue and Gu (2009) perform automated detection of ocean water masses using the Morse-Smale complex.

I propose to use the Reeb graph towards this goal: it is not just for visualization. Mathematics is the language of the universe. Some branches of mathematics, however, do not describe the physical universe, or at least we have not yet realized how they might do so. I intend to help move the Reeb graph into the camp of mathematical objects that does, in a real and meaningful way, describe the physical world.

[2] Since the domain here is simply connected, the Reeb graph becomes a tree, called the contour tree.

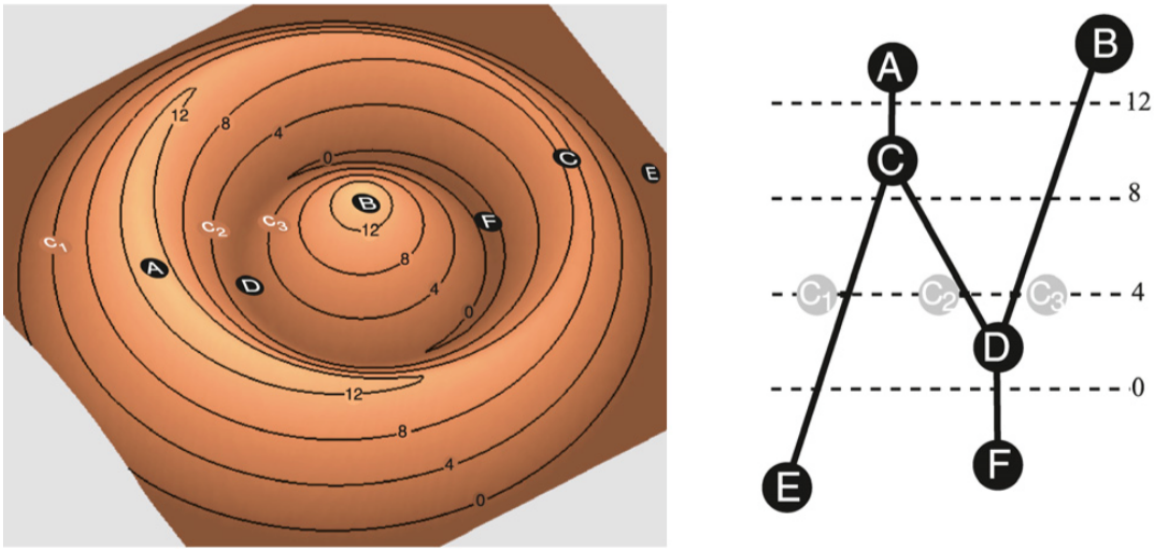


Figure 1.2: Visualization of an island within a volcanic crater lake (left) and its contour tree (right). Contours of constant height are shown. Critical points are labelled with letters. A and B are maxima, C and D are saddle points, and E and F are minima. Figure reproduced from (Carr et al., 2010).

1.4.2 The Reeb Graph

Two definitions of the Reeb graph are given, first using general topology, then using graph theory. Mathematical background for both fields is given in Appendix 1.A.

Given a continuous function $f : X \to \mathbb{R}$, where X together with some definition of its open sets is a topological space, the **Reeb graph** (Reeb, 1946) is formally defined as the quotient space $X/\sim$, where the equivalence relation $\sim$ is such that $x \sim y$ if x and y are elements of the same connected component of the level set of f at z, for some $z \in \mathbb{R}$.

Loosely speaking, the function f assigns a value to each point in space X, and each connected component of each level set of f is contracted to a single point in the new space. In particular, note $x \sim y \implies f(x) = f(y)$, but the converse is not necessarily true: the level set of f at $z \in \mathbb{R}$, which is $f^{-1}(z) = \{x \in X : f(x) = z\}$, can have multiple connected components, so x and y can be in different components.

This formal definition, rooted in general topology, is not required to make practical use of the Reeb graph. I will give a more informal definition, using graph theory (which the formal definition does not use whatsoever).

A good example to start with is a topographic one, such as shown in Fig. 1.2, and suppose we extend the domain to infinity in all directions. A function $f : \mathbb{R}^2 \to \mathbb{R}$ assigns, to each point in X, the height of topography. The level set of f at 4, say, is the

set of all points (x, y) such that $f(x, y) = 4$. This set has three connected components; each of these is called a contour. One such contour is on the outside of the volcano, one on its inside, and one on the volcanic island in the middle. Now, consider how the level set of f at z changes as the isovalue z is swept from high to low values. The level set at $z > f(B)$ is the empty set (there are 0 connected components). When $z = f(B)$, the level set has 1 contour, which contains a single point in space, B. A new contour has been created at $z = f(B)$. As z decreases, this contour grows, but nothing really changes until $z = f(A)$, when a second contour is created. Now the level set of z has 2 connected components. Continuing, again nothing really changes, until $z = f(C)$, when the contour that grew from A has reached around the entire rim of the volcano, and meets itself at C; then it splits in two (like DNA during mitosis) as z is decreased further. Now there are 3 connected components of the level set; each of the 3 contours contains one of the points C_1, C_2, or C_3. Descending further, the contour born at B that encircles the island coalesces with one of the contours born at C when $z = f(D)$: a contour has been destroyed. One more contour is destroyed when $z = f(F)$, and the last is destroyed when $z = f(E)$.

The points A through F are *critical points* of the function f: they are maxima, minima, or saddle points. At the *critical values* (the value of f at the critical points), important things happened to the topology of the level set. Its connectedness changed as z passes a critical value. At regular points, the level set changes but its connectedness does not.

The nodes of the Reeb graph represent the critical points of f. An extremum is a node of degree one in the Reeb graph, and a saddle is a node of degree three or more[3]. Consider two critical points n and m. If there is a path between n and m such that f is monotone along this path, and no point on this path is on a contour of any other critical value, then the Reeb graph has an arc between nodes n and m. In other words, nothing (topologically) interesting happens on this path between two critical points. This arc represents all points in space (which are *regular points*) that are on contours that intersect this path. These points are called "associated vertices" for this arc. (Each data point in space will be defined as a vertex on a simplical mesh, hence the choice of "vertices".) The Reeb graph discards these regular points, but the **augmented Reeb graph** retains them. They can be thought of as degree two nodes in the graph, or more usefully as a set of vertices associated to each arc.

[3] Under certain conditions, namely that f is Morse, all saddles are degree three nodes. For f to be Morse requires, among other things, that all critical points of f are in the interior of the manifold. This will not generally be true for the examples in this thesis.

The above example is a special case of a Reeb graph, known as a contour tree. The graph, containing no cycles, is a tree. When the manifold is simply connected (any path on the manifold can be continuously contracted to a point), the Reeb graph reduces to this simpler case, a tree. Since the ocean contains islands (including bigger islands known as continents), many manifolds embeddable in the ocean are not simply connected. Thus, we must deal with Reeb graphs containing cycles, rather than the simpler contour trees. Indeed, if $\mathbb{M}$ is an orientable 2-manifold without boundary and with genus g, then the Reeb graph of a Morse function on $\mathbb{M}$ has a cycle basis with g elements (Cole-McLaughlin et al., 2003; Agarwal et al., 2004).

Reeb graphs are part of a broader mathematical field called Morse theory, which studies the topology of a differentiable manifold $\mathbb{M}$ indirectly, by studying the properties of a smooth real-valued function on the manifold. Most functions suffice, with some restrictions on their critical points. Specifically, a **Morse function** on a manifold $\mathbb{M}$ is a smooth function $f : \mathbb{M} \to \mathbb{R}$ satisfying the following three properties (Cole-McLaughlin et al., 2003): (1) all critical points of f are non-degenerate and lie in the interior of $\mathbb{M}$; (2) all critical points of the restriction of f to $\partial\mathbb{M}$, the boundary of $\mathbb{M}$, are non-degenerate; (3) all critical values of f and of the restriction of f to $\partial\mathbb{M}$ are unique.

If $\mathbb{M}$ is embeddable in $\mathbb{R}^n$, it is often useful to choose f as the height of $\mathbb{M}$ given a particular embedding. In such a case, property (3) can often be satisfied by an infinitesimal solid body rotation of the manifold (Carr et al., 2003). The failure of properties (1) and (2) result in nodes in the Reeb graph having degree greater than 3, but this can be handled (Doraiswamy and Natarajan, 2013).

Practical use of a Reeb graph requires translating the ideas of Morse theory, involving smooth functions, to discretised data, and pairing that with algorithms that efficiently operate on large data-sets. This will be described further in Section 2.3.

1.5 Data

I use data from two ocean reanalysis products in this thesis, both stemming from the Estimating the Circulation and Climate of the Ocean (ECCO) consortium. Chapter 2 uses ECCO2 (Menemenlis et al., 2005), while Chapters 3 and 4 use the Southern Ocean State Estimate (SOSE, Mazloff et al., 2010). These both use the MITgcm to minimize, in a least-squares sense, a misfit function, which compares the model state with oceanographic and meteorological observational data. To minimize this misfit function, the initial conditions and boundary conditions are both modified, but no

artificial sources are added to the ocean interior. Thus, output from such models can be thought of as a close estimate of the real ocean (if not exactly converged), but dynamically self-consistent.

ECCO2 is a global model run on a cube-sphere grid, with a mean horizontal grid spacing of 18 km and 50 vertical levels varying from 10 m at the surface to 445 m at the bottom. The results are interpolated onto a latitude-longitude grid (an equirectangular projection) at 1/4° resolution, and archived by Eulerian averaging on three day intervals. I use data averaged between June 20 and June 22, 2002, prior to model drift in the Southern Ocean (Azaneu et al., 2014). For full details, see Menemenlis et al. (2005).

SOSE is a Southern Ocean-only model, with a sponge region at 25°S. It is run on a latitude-longitude grid of 1/6° resolution with 42 vertical layers, ranging from 10 m at the surface to 250 m at the bottom. For full details, see Mazloff et al. (2010).

1.6 Thesis outline

The primary chapters of this thesis follow, after a review of some mathematics. Chapter 2 studies neutral surfaces and the geostrophic stream function on neutral surfaces. Chapters 3 and 4 study potential vorticity in the Southern Ocean; the former studies PV variations across streamlines, while the latter studies PV variations along individual streamlines. Final remarks are offered in Chapter 5.

A word on the title of this thesis is warranted. The major insights of Chapter 2 are topological in nature, and it is only in this chapter that explicit topological calculations are performed. Chapters 3 and 4 both study PV in the Southern Ocean, so arguably the title should be along the lines of "Potential vorticity dynamics of the Antarctic Circumpolar Current". However, the same topological considerations of Chapter 2 apply to PV, and shape some of the thinking in Chapter 3. Explicit topological calculations are not performed there because there is plenty to learn first without them, and the circumpolar streamlines in the Southern Ocean actually greatly simplify the topology. That said, these topological considerations only arise for PV if it is conserved following the flow, so Chapter 4 can be viewed as investigating the degree to which the real ocean exhibits such topological properties. Finally, there are more geophysical applications of this topological analysis, some of which are discussed while concluding in Chapter 5.

Appendix

1.A Mathematical definitions

The mathematical foundations of the Reeb graph lie in topology and graph theory, so both are reviewed below, mostly as a sequence of definitions. Only elementary aspects of these fields are needed. In particular, only general topology (also called point-set topology) is needed, and naive set theory suffices. I begin with some general definitions focusing on functions, then follow with topology and graph theory. This review is not intended to be exhaustive or completely self-contained. Crucial definitions for the Reeb graph are critical points, level sets, connected components, and graphs. Much of this material is in preparation for Chapter 2.

1.A.1 Functions

A **set** is a collection of distinct objects. An object x that belongs to the set X is called an element of X, denoted $x \in X$. A set X is a **subset** of a set Y, denoted $X \subseteq Y$, if $x \in Y \ \forall x \in X$. Two sets X and Y are equal, denoted $X = Y$, if $X \subseteq Y$ and $Y \subseteq X$.

A **multiset** is a collection of objects, not necessarily distinct. The **multiplicity** of an element is the number of instances of that element in the multiset.

A **tuple** is a finite ordered list of objects.

The **Cartesian product** of two sets X and Y, denoted $X \times Y$, is the set of tuples $\{(x, y) : x \in X, y \in Y\}$. When $Y = X$, this is denoted X^2. This definition extends recursively to define X^n.

A **binary relation** between two sets X and Y is a subset $R \subseteq X \times Y$. When $(x, y) \in R$, we write $x \sim y$ and say that x is R-related to y. A binary relation on a single set X is a binary relation between X and X.

An **equivalence relation** is a binary relation on X that is reflexive, $x \sim x$, symmetric, $x \sim y \iff y \sim x$, and transitive, $x \sim y \ \& \ y \sim z \implies x \sim z$. These properties must hold for all $x, y, z \in X$.

A **function** f from a set X to a set Y, denoted $f : X \to Y$, is a subset of $X \times Y$ such that $\forall x \in X$, there is exactly one $y \in Y$ such that $(x, y) \in f$, which is equivalently written $f(x) = y$.

The **image** of $S \subseteq X$ under a function $f : X \to Y$ is the set $\{f(s) : s \in S\}$. It is denoted $f(S)$, and is a subset of Y.

The **preimage** of $T \subseteq Y$ under a function $f : X \to Y$ is the set $\{x \in X : f(x) \in T\}$. It is denoted $f^{-1}(T)$, and is a subset of X.

A **level set** of f at z is the preimage of the singleton set $\{z\}$ under f. Specifically, it is $f^{-1}(z) = \{x \in X : f(x) = z\}$.

A function $f : \mathbb{R}^n \to \mathbb{R}^m$ is **differentiable at** $\boldsymbol{x} \in \mathbb{R}^n$ if $\lim_{\boldsymbol{y} \to \boldsymbol{x}} ||f(\boldsymbol{y}) - f(\boldsymbol{x}) - J_f^{\boldsymbol{x}}(\boldsymbol{y} - \boldsymbol{x})|| \ / \ ||\boldsymbol{y} - \boldsymbol{x}|| = 0$ for some linear map $J_f^{\boldsymbol{x}} : \mathbb{R}^n \to \mathbb{R}^m$. Here, vertical bars represent the Euclidean norm in $\mathbb{R}^m$ or $\mathbb{R}^n$ for the numerator or denominator, respectively. When the limit exists and equals zero, the linear map $J_f^{\boldsymbol{x}}$ is the **derivative** of f at $\boldsymbol{x}$. As a linear map, $J_f^{\boldsymbol{x}}$ can be expressed by a matrix (m rows, n columns), called the **Jacobian matrix**, in which case its columns are the partial derivatives of f in the chosen basis, and these partial derivatives are guaranteed to exist. The set $\{f(\boldsymbol{x}) + J_f^{\boldsymbol{x}}(\boldsymbol{y} - \boldsymbol{x}) : \boldsymbol{y} \in \mathbb{R}^n\}$ defines a **tangent plane** to f at $\boldsymbol{x}$, which is a good approximation of $f(\boldsymbol{y})$ in a sufficiently small region around $\boldsymbol{x}$, since the above limit is zero. Intuitively, $J_f^{\boldsymbol{x}}(\boldsymbol{h})$ gives the rate of change of f at $\boldsymbol{x}$ in the direction $\boldsymbol{h}$. A function is **differentiable** if it is differentiable at every point in its domain.

The gradient is a special case of the Jacobian when $m = 1$ above. Specifically, the **gradient** of a differentiable scalar function $f : \mathbb{R}^n \to \mathbb{R}$ at $\boldsymbol{x} \in \mathbb{R}^n$ is the linear map $J_f^{\boldsymbol{x}} : \mathbb{R}^n \to \mathbb{R}$. Being a linear map, it can, as before, be expressed as a matrix, but now with only one row. Thus the gradient is usually defined as a vector field ∇f such that $(\nabla f) \cdot \boldsymbol{h} = J_f^{\boldsymbol{x}}(\boldsymbol{h})$, where $\boldsymbol{x}$ exists implicitly on the left hand side as the point at which ∇f is evaluated.

A **regular point** of a differentiable scalar function $f : \mathbb{R}^n \to \mathbb{R}$ is a point where ∇f is non-zero. In this case, $J_f^{\boldsymbol{x}}(\nabla f / ||\nabla f||) = ||\nabla f|| \geq (\nabla f) \cdot \boldsymbol{h} = J_f^{\boldsymbol{x}}(\boldsymbol{h})$ for any unit vector $\boldsymbol{h}$, so ∇f is parallel to the unit vector $\boldsymbol{h}$ that maximizes $J_f^{\boldsymbol{x}}(\boldsymbol{h})$. Most points of a non-pathological function are regular.

A **critical point** of a differentiable scalar function $f : \mathbb{R}^n \to \mathbb{R}$ is a point where the gradient ∇f is zero. A critical point $\boldsymbol{x}$ has a corresponding **critical value**, $f(\boldsymbol{x})$.

The **Hessian** of a twice continuously differentiable scalar function $f : \mathbb{R}^n \to \mathbb{R}$ at $\boldsymbol{x} \in \mathbb{R}^n$ is the linear map $H_f^{\boldsymbol{x}} : \mathbb{R}^n \to \mathbb{R}^n$ defined by $H_f^{\boldsymbol{x}} = J_{\nabla f}^{\boldsymbol{x}}$. That is, the Hessian is the Jacobian of the gradient. As a linear map, the Hessian can be expressed as the **Hessian matrix**, given a basis $\{x_i : i = 1..n\}$, having (i, j) element $\partial_{x_i} \partial_{x_j} f$.

A critical point $\boldsymbol{x}$ is further classified by the Hessian of f at $\boldsymbol{x}$. Since the Hessian matrix is symmetric, all its eigenvalues are real. If the Hessian matrix is non-singular (it is invertible, its determinant is non-zero, and none of its eigenvalues are zero), then the critical point is **non-degenerate**. The number of negative eigenvalues I of $H_f^{\boldsymbol{x}}$ is the **index** of the critical point $\boldsymbol{x}$. If $I = 0$, the function is concave-up here and

$\boldsymbol{x}$ is a **local minimum**. Similarly, if $I = n$, the function is concave-down here and $\boldsymbol{x}$ is a **local maximum**. For intermediate values of I, $\boldsymbol{x}$ is a **saddle** point, with f increasing in some directions and decreasing in others.

1.A.2 General topology

A **topological space** is a pair (X, τ), where X is a set and τ is a set of subsets of X. Elements of τ are called **open sets**, and τ is called a **topology** on X. The elements of τ must satisfy three axioms. First, the empty set and X both belong to τ. Second, any union, whether finite or infinite, of elements of τ belongs to τ. Third, any finite intersection of elements of τ belongs to τ.

A function $f : X \to Y$ between two topological spaces is **continuous** if, for any open $V \subseteq Y$, the preimage of V, namely $U = \{u \in X : f(u) \in V\}$, is an open subset of X. Thus, for any $x \in X$ and any open $V \subseteq Y$ containing $f(x)$, there exists an open set in X, namely U, that contains x and satisfies $f(U) \subseteq V$. This connects with the more familiar definition of a continuous function f between metric spaces X and Y: f is continuous at $x \in X$ if, for any $\delta > 0$, one can always find an $\epsilon > 0$ such that f maps all points a distance strictly less than ϵ from x (an open set, U) to a distance strictly less than δ from $f(x)$ (an open set, V). Thus, open sets give rise to notions commonly related to distance, without actually imposing a metric (which quantifies distance) on the topological space.

A **homeomorphism** is a continuous function from one topological space to another that has a continuous inverse function. Intuitively, a homeomorphism represents a continuous stretching and twisting of one geometric object to another. Two topological spaces are said to be **homeomorphic** to one another if there exists a homeomorphism between them. For instance, a mug with a handle is homeomorphic to a doughnut, but not to a cake: with continuous deformation, the handle of the mug becomes the hole at the centre of the doughnut, but without any such hole the cake is topologically different.

A **neighbourhood** N of $x \in X$ is a subset of X (it need not be open) such that $x \in U \subseteq N$ for some open set U.

An n-**manifold** is a topological space (X, τ) such that for each point $x \in X$, there exists a homeomorphism, called a **chart**, from a neighbourhood of x to an open subset of $\mathbb{R}^n$. Thus, an n-manifold locally resembles n-dimensional Euclidean space, but globally these spaces need not be topologically equivalent. A circle is a 1-manifold: the function that maps a point on the circle to its angle is a homeomorphism from a neighbourhood of that point to the real line. But the circle and the real line are

not homeomorphic: no homeomorphism exists that maps the whole circle to an open subset of the real line. For example, the angle function maps the circle to $[0, 2\pi)$, not an open set. A plus sign, +, is not a manifold: the centre where two lines intersect does not resemble Euclidean space of any dimension. An isohaline, i.e. a level set of the function that maps every point in the ocean to its salinity (a continuous 3D scalar field), is a 2-manifold. This surface may have holes, for example surrounding islands. The depth of an isohaline need not be a function of geographic position—a water column may have multiple depths with the same salinity, so that isohaline surface can bend over itself—but it is still a manifold.

An **atlas** is a collection of charts such that the union of the domain of all charts is the entire manifold. Two charts with domains having a non-empty intersection can, through the composition of one with the inverse of the other, define a homeomorphism of an open subset of $\mathbb{R}^n$ to another open subset of $\mathbb{R}^n$. This composition, called a **transition map**, tells how the manifold is "patched together" from various bits of Euclidean space.

A **differentiable n-manifold** is an n-manifold for which an atlas exists whose transition maps are all differentiable. Intuitively, the differential structure of Euclidean space can be imported to a manifold by its charts; if these charts are all mutually compatible, then the differential structure of Euclidean space is brought to the manifold in a self-consistent way: there are no non-differentiable transitions on the manifold when moving from one chart to another. Every n-manifold for $n \leq 3$ is a differentiable manifold.

The definition of **critical points** of functions whose domain is $\mathbb{R}^n$ extends naturally to a function f defined on a differentiable n-manifold $\mathbb{M}$. Given an atlas for $\mathbb{M}$ and a point $\boldsymbol{x} \in \mathbb{M}$, one locates the chart $U : \mathbb{M} \to \mathbb{R}^n$ in the atlas containing $\boldsymbol{x}$, then classifies the point $U(\boldsymbol{x})$ as regular or critical (and further subclassification) using the composed function $f \circ U^{-1}$.

There are two elementary ways of creating new topological spaces from old ones: by taking a subset, and by relating elements via an equivalence relation. From the perspective of the set only, the latter is a special case of the former. However, these operations imbue the new topological space with different topologies.

Given a topological space (X, τ_X), a **subspace** is a topological space (S, τ_S) where $S \subseteq X$ and $\tau_S = \{S \cap U : U \in \tau\}$.

Given a topological space (X, τ_X), an equivalence relation $\sim$ defines the **quotient space** (Y, τ_Y) as follows. First define the **canonical map** $q : X \to X/\sim$ by $q(x) = \{a \in X : a \sim x\}$. For each $x \in X$, the subset $q(x) \subseteq X$ is called the **equivalence**

class of X given $\sim$, and is sometimes denoted $[x]$. For the quotient space, the set Y, often denoted $X/\sim$, is the set of such equivalence classes, namely $Y = \{q(x) : x \in X\}$. Intuitively, one thinks of "gluing" equivalent (under $\sim$) points in X together to a single point, and placing this "representative point" in Y. The **quotient topology** τ_Y contains all subsets of Y whose preimage under q is an open set in X: $\tau_Y = \{U \subseteq Y : q^{-1}(U) \in \tau_X\}$. Intuitively, one thinks of an open subset $U \subseteq Y$ being "unglued", returning to all points in X that were glued into U, and finding that this subset in X is open.

A **disconnected** space is a topological space for which there exist two disjoint, nonempty, open sets whose union is the full topological space. If a topological space is not a disconnected space, it is a **connected** space.

A **connected component** of a topological space (X, τ) is a connected subspace that is not a subspace of any other connected subspace. If X has exactly one connected component, then X is **connected**. The union of all connected components of X is X itself. Any two connected components of X are disjoint.

An n-**manifold with boundary** is defined similarly to an n-manifold, except some points (the union of which is called the boundary) have no chart as defined for an n-manifold, but do have a **boundary chart** that is a homeomorphism from a neighbourhood of the point to an open set in $\mathbb{R}^n_+ = \{(x_1, ..., x_n) : x_n \geq 0\}$ imbued with the subspace topology from $\mathbb{R}^n$. The boundary of an n-manifold with boundary is a $(n-1)$-manifold, and the interior (those points not in the boundary) of an n-manifold with boundary is an n-manifold.

1.A.3 Graph Theory

An **undirected graph** is a tuple $G = (N, A)$ where N is a set of nodes (or vertices) and A a set of arcs (or edges), and elements of A are subsets of N having precisely two elements. Intuitively, a graph is a collection of points (nodes), and lines connecting any two points (arcs).

An arc a is said to be **incident upon**, or **adjacent to**, a node n if $n \in a$.

The **degree** of a node is the number of distinct arcs that are incident upon it.

A node n is a **neighbour** of a node m if there is an arc incident upon both n and m.

When A is a multiset (its elements need not be unique), G is an undirected **multigraph**. A multigraph can have multiple arcs incident upon the same two nodes. The Reeb graph is actually a multigraph, but the word graph is used throughout this thesis as shorthand.

A **walk** is an ordered sequence that starts and ends with nodes, and alternates between nodes and arcs, such that every arc in this sequence is incident upon both nodes next to it in the sequence.

A **simple cycle** (or a **cycle** for short) is a walk whose first node is also its last node, and contains no repeated arcs. An n**-cycle** is a cycle with n arcs in its walk.

A **connected** graph is a graph such that there is a walk between any two of its nodes.

A **tree** is a connected graph with no cycles.

A **forest** is a graph with no cycles—and thus is a collection of trees.

A **spanning subgraph** of a graph $G = (N, A)$ is a graph $H = (N, F)$ having $F \subseteq A$. A subgraph has the same set of nodes, but possibly fewer arcs.

A **minimum spanning tree** is a subgraph that is also a tree.

An **Eulerian graph** is a graph whose nodes all have even degree.

The **cycle space** of a graph G consists of all Eulerian subgraphs of G. The cycle space is a vector space over the two-element finite field $\mathbb{Z}_2$, containing only two elements 0 and 1. In $\mathbb{Z}_2$, operations are like in the integers $\mathbb{Z}$ but modulo 2. For example, 1+1=0 in $\mathbb{Z}_2$. Two elements of the cycle space are added (as in $\mathbb{Z}_2$) on their arcs; multiplication by 1 is the identity operation, and by 0 gives the empty subgraph of G.

A **cycle basis** is a basis for the cycle space.

A **directed** graph is a graph whose arcs are ordered 2-tuples rather than unordered sets. That is, the arcs of a directed graph can only be walked in the specified direction.

An **arc-weighted** graph is a graph where each arc is associated with a weight, usually a real number.

An **node-weighted** graph is a graph where each node is associated with a weight, usually a real number.

Chapter 2

The topology of neutral surfaces

2.1 Introduction

Strong stratification throughout most of the ocean inhibits vertical motion, largely confining the oceanic flow to a two dimensional surface. This is a great simplification, but only if we know the orientation of this surface; it is not a geopotential surface. Crafting such 2D surfaces has been a perennial goal for physical oceanographers. It is perennial because it is both highly desirable and impossible. At any point in the ocean, the neutral tangent plane is defined as the direction in which fluid parcels that are displaced while conserving their heat and salt are not subject to a buoyant restoring force (McDougall, 1987a). This is well-defined, and locally the flow is largely aligned with this plane. Globally stitching together these neutral tangent planes could form a 2D surface, called a neutral surface. Unfortunately, non-linearity in seawater's equation of state leads to a path-dependence underlying the definition of neutral surfaces; therefore, neutral surfaces are commonly said to be ill-defined (McDougall and Jackett, 1988). Nonetheless, approximately neutral surfaces can be constructed.

Neutrality, as the above property is called, is one of three desirable properties of variable like potential temperature in a dry atmosphere. First, potential temperature is materially conserved (excluding the effects of irreversible processes like molecular diffusion and radiation, in light of which we call this property quasi-conservative). Second, transforming the equations of motion to use potential temperature as the vertical coordinate introduces no additional accelerations. Third, potential temperature surfaces are identical to neutral surfaces. In fact, these three properties are closely linked. The second property guarantees the existence of an exact geostrophic stream function in potential temperature surfaces. The first and second properties imply the Ertel potential vorticity is quasi-conserved (de Szoeke and Springer, 2009).

However, it is known that no quasi-conservative potential vorticity exists for a multi-component fluid like the ocean (Müller, 1995, and references therein); thus the first and second property cannot both be satisfied.

The goal, therefore, is to define a variable that satisfies three properties as well as possible. Of course there is a trade-off between properties. The first such variables were potential density (Wüst, 1935) and specific volume anomaly (Montgomery, 1937). Lynn and Reid (1968) revealed the highly undesirable property that potential density can exhibit unphysical overturns far away from its reference pressure (i.e. property three can fail disastrously). Similar issues affect the specific volume anomaly when far from the reference potential temperature and salinity. Nowadays, one can use patched potential density (Reid and Lynn, 1971), neutral density (Jackett and McDougall, 1997), orthobaric density (de Szoeke et al., 2000), and generalized patched potential density and thermodynamic neutral density (Tailleux, 2016). One purpose-built 2D surface exists: the ω-surface (Klocker et al., 2009). Each have their own advantages and disadvantages.

Having the flow confined to a 2D surface is much more useful if we can describe the flow on that surface by a stream function. The geostrophic stream function serves this purpose, and though it approximates the flow, it is useful because it is calculable using only hydrographic information. Geostrophic stream functions are ubiquitous in physical oceanography, for three reasons. First, they allow us to study maps of the large-scale flow based purely on hydrographic data (Montgomery, 1938; Zhang and Hogg, 1992) Second, they provide a simple way to analyse quantities following the (geostrophic) flow (e.g. Gille, 1997; Shuckburgh et al., 2009) Third, as a divergence-free representation of a simplified flow field, they serve as a building block for many theoretical models (Stommel, 1948; Stommel and Arons, 1959; Welander, 1971; Rhines and Young, 1982; Luyten et al., 1983; Marshall and Radko, 2003).

It is an advantage if the geostrophic stream function exists for a particular 2D surface, but this is not always the case. Of the aforementioned surfaces, only specific volume anomaly and orthobaric density surfaces possess an exact geostrophic stream function. In fact, any pycnotropic variable, i.e. a function of *in-situ* density and pressure alone, possesses an exact geostrophic stream function (McDougall, 1989; de Szoeke, 2000). Hence isosurfaces of p and isosurfaces of ρ both possess an exact geostrophic stream function, but these surfaces are not particularly useful: they are close to geopotential surfaces, and far from neutral. Isosurfaces of potential density (isopycnals) do not have an exact geostrophic stream function (McDougall, 1989), nor

do neutral density surfaces (de Szoeke and Springer, 2009; McDougall and Klocker, 2010), nor do the ω-surfaces defined by Klocker et al. (2009).

I present a fresh perspective to this problem, based on topology and the Reeb graph, yielding 2D nearly neutral surfaces possessing an exact geostrophic stream function. With the second property satisfied, attention rests on the third property: as a 2D surface not a 3D variable, material conservation cannot be evaluated, so the first property is ignored for now. These surfaces derive from the topological connectedness of pressure contours, and so I call them *topobaric* surfaces.

This chapter is laid out as follows. Section 2.2 reviews the theory of neutral surfaces, derives their exact geostrophic stream function, and develops a topological theory of neutral surfaces. Section 2.3 details the method used to construct topobaric surfaces. Results of this method are described in Section 2.4. Conclusions are given in Section 2.5.

2.2 Theory

2.2.1 Background and existence of neutral surfaces

It is best to work with specific volume, $\nu = \rho^{-1}$, where ρ is the *in-situ* density. The gradient of specific volume is

$$\nabla\nu = \nu_S \nabla S + \nu_\theta \nabla\theta + \nu_p \nabla p, \tag{2.1}$$

where subscripts represent partial derivatives.[1] Consider a fluid parcel of specific volume ν_0 being adiabatically infinitesimally displaced by $\mathrm{d}\boldsymbol{r}$. The specific volume of its new surroundings is $\nu_0 + \nabla\nu \cdot \mathrm{d}\boldsymbol{r}$, which has contributions from salt, potential temperature[2], and pressure changes (substitute (2.1)). However, its specific volume after displacement is only $\nu_0 + \nu_p \nabla p \cdot \mathrm{d}\boldsymbol{r}$. The difference, $(\nu_S \nabla S + \nu_\theta \nabla\theta) \cdot \mathrm{d}\boldsymbol{r}$, creates a buoyant restoring force. If the displacement $\mathrm{d}\boldsymbol{r}$ is perpendicular to the dianeutral vector

$$\boldsymbol{A} = \nu_S \nabla S + \nu_\theta \nabla\theta, \tag{2.2}$$

[1] Often this equation is divided by ν, so the coefficients are the thermal expansion coefficient $\nu^{-1}\nu_\theta$, the haline contraction coefficient $-\nu^{-1}\nu_S$, and the adiabatic compressibility $\nu^{-1}\nu_p$. This complicates further differentiation and so is not used here.

[2] The distinctions between potential temperature and Conservative Temperature, and between practical salinity and Absolute Salinity, are not important here. One should use Conservative Temperature and Absolute Salinity when dealing with real ocean data, but use potential temperature and practical salinity when dealing with a model that uses these in its equation of state.

then the buoyant restoring force is zero. The neutral tangent plane, in which a fluid parcel can move with no buoyant restoring force, is the plane normal to the dianeutral vector. Except in rare places where $\boldsymbol{A} = \boldsymbol{0}$ identically, such as truly mixed regions, this is well-defined everywhere in the ocean.

A neutral surface is everywhere tangent to the neutral tangent plane. Suppose a neutral surface exists, and consider a closed path $\partial\Omega$ on the neutral surface which bounds a subset Ω. Take the line integral of $\boldsymbol{A}$ along $\partial\Omega$, and use Stokes' theorem:

$$\int_{\partial\Omega} \boldsymbol{A} \cdot \mathrm{d}\boldsymbol{r} = \int_{\Omega} \nabla \times \boldsymbol{A} \cdot \mathrm{d}\boldsymbol{\Omega}. \tag{2.3}$$

The LHS is zero, because $\mathrm{d}\boldsymbol{r}$ is tangent to the neutral surface, which by definition means it is in the neutral tangent plane, hence orthogonal to $\boldsymbol{A}$. For the RHS, since the surface of integration is a neutral surface, $\mathrm{d}\boldsymbol{\Omega} = \mathrm{d}\Omega\boldsymbol{A}/|\boldsymbol{A}|$, hence

$$0 = \int_{\Omega} \frac{H}{|\boldsymbol{A}|} \mathrm{d}\Omega, \tag{2.4}$$

where $H = \boldsymbol{A} \cdot \nabla \times \boldsymbol{A}$ is the neutral helicity. Since Ω was arbitrary, we require $H = 0$ everywhere. Expanding $\boldsymbol{A}$ and using the chain rule on $\nabla\nu_S$ and $\nabla\nu_\theta$ akin to (2.1), we find

$$H = (\nu_\theta \nu_{Sp} - \nu_S \nu_{\theta p}) \nabla p \cdot (\nabla S \times \nabla\theta). \tag{2.5}$$

The term in parentheses is non-zero; it is $(\nu_S \nu)$ times the thermobaricity[3]. Thus $H = 0$ if and only if the three vectors ∇p, ∇S, and $\nabla\theta$ are coplanar, i.e. if one of the vectors is a weighted sum of the other two. This is not generally true (see McDougall and Jackett, 2007, for analysis of oceanic data), and $H \neq 0$ in most of the ocean. Thus, (2.4) is false, so a contradiction is reached. The false assumption was to assume the integration path was closed. It is usually concluded that neutral surfaces are not well-defined (McDougall and Jackett, 1988).

Still, oceanographers need well-defined surfaces, i.e. with a unique depth for any geographic location, on which to analyse ocean properties. A useful goal is to craft a well-defined surface that is everywhere nearly tangent to the neutral tangent plane: this is a so-called "approximately neutral surface". In Section 2.2.4, a different kind of topological consideration of (exact) neutral surfaces is considered, leading to a new way to craft an approximately neutral surface. However, first we consider the geostrophic stream function on neutral surfaces in Section 2.2.3, as these will naturally lead into the subsequent material. Before either can be done, a different characterisation of neutral surfaces is required.

[3] This equivalence is straightforward but takes a few lines to derive. The thermobaricity is defined (McDougall, 1987b) as $a_p - ab^{-1}b_p$, where $a = -\rho^{-1}\rho_\theta$ is the thermal expansion coefficient and $b = \rho^{-1}\rho_S$ is the haline contraction coefficient.

2.2.2 Neutral surface characterisation

First, we require another characterisation of neutral surfaces, as follows. The 3D gradient of some tracer C is decomposed into a component parallel and a component normal to the neutral tangent plane, as such: $\nabla C = \nabla_N C + (\nabla C \cdot \hat{\boldsymbol{n}})\hat{\boldsymbol{n}}$ where $\hat{\boldsymbol{n}} = \boldsymbol{A}/|\boldsymbol{A}|$ (McDougall et al., 2014). Thus, displacements $\mathrm{d}\boldsymbol{r}$ in the neutral tangent plane satisfy

$$\begin{aligned} 0 = \mathrm{d}\boldsymbol{r} \cdot \boldsymbol{A} &= \mathrm{d}\boldsymbol{r} \cdot (\nu_S \nabla_N S + \nu_S (\nabla S \cdot \hat{\boldsymbol{n}})\hat{\boldsymbol{n}} + \nu_\theta \nabla_N \theta + \nu_\theta (\nabla \theta \cdot \hat{\boldsymbol{n}})\hat{\boldsymbol{n}}) \\ &= \mathrm{d}\boldsymbol{r} \cdot (\nu_S \nabla_N S + \nu_\theta \nabla_N \theta). \end{aligned} \tag{2.6}$$

Since $(\nu_S \nabla_N S + \nu_\theta \nabla_N \theta)$ is in the neutral tangent plane and $\mathrm{d}\boldsymbol{r}$ is arbitrary in that plane, then $\hat{\boldsymbol{n}}$ must be such that

$$\nu_S \nabla_N S + \nu_\theta \nabla_N \theta = \mathbf{0}, \tag{2.7}$$

which provides another, sometimes more useful, definition for the neutral tangent plane. McDougall et al. (2014) show that this is equivalent to[4]

$$\nu_{\underset{\sim}{S}} \nabla \underset{\sim}{S} + \nu_{\underset{\sim}{\theta}} \nabla \underset{\sim}{\theta} = \mathbf{0}, \tag{2.8}$$

where the under-tilde indicates restriction to the neutral surface.

Specifically, we have a function $\underset{\sim}{S} : \mathbb{M} \to \mathbb{R}$ that assigns a salinity to each point in space, here a 2-manifold $\mathbb{M}$, which is simply the surface of a sphere (of constant radius), with some holes. The holes are islands, seamounts, continents, and other places where the neutral surface has grounded or outcropped (in what follows, only islands are listed, for poetics). The same can be done for other scalar fields. Whenever one sees a variable with an under-tilde, it can be thought of simply as a function of geographic location, i.e. latitude and longitude (though undefined for some locations).

With this notation, the gradients in (2.8) are naturally gradients on the sphere of constant radius: distance is measured without regard to the radial (vertical) variations of the surface. This is the "projected non-orthogonal gradient" introduced by Starr (1945) and used throughout studies of neutral surfaces. Indeed, (2.8) is identical to one of two equivalent definitions of neutral surfaces from McDougall (1987a), the other being identical to the following equation, obtained simply by using (2.1) in (2.8):

$$\nabla \underset{\sim}{\nu} = \nu_{\underset{\sim}{p}} \nabla \underset{\sim}{p}. \tag{2.9}$$

I henceforth use (2.9) because having p a monotonic function of depth will be useful.

[4] McDougall et al. (2014) use the notation ∇_n, which is the projected nonorthogonal gradient, defined as $\nabla_n \equiv \frac{\partial}{\partial x}|_n \boldsymbol{i} + \frac{\partial}{\partial y}|_n \boldsymbol{j}$, where $\boldsymbol{i}$ and $\boldsymbol{j}$ are the eastward and northward unit vectors.

2.2.3 The exact geostrophic stream function on a neutral surface

McDougall (1989) first proved that an exact geostrophic stream function exists for a neutral surface but could not provide a closed expression for it. Since then, many approximate geostrophic stream functions have been crafted. Zhang and Hogg (1992) upgraded the Montgomery (1937) potential for use in potential density surfaces. Cunningham (2000) combined the Bernoulli potential and the potential enthalpy which, as McDougall and Klocker (2010) show, can also be seen as a modification of the dynamic height. McDougall and Klocker (2010) derive a further extension from Zhang and Hogg (1992) by assuming a quadratic dependence between Conservative Temperature and pressure along a neutral surface.

A closed form expression for an exact geostrophic stream function on a neutral surface has never been found. Even if neutral surfaces do not embed in the real ocean, an important step for theoretical oceanography would be to find an expression for the geostrophic stream function on a neutral surface. For practical purposes, then, such an expression could be applied to approximately neutral surfaces with some, hopefully small, error.

The basis of all geostrophic stream functions is the geostrophic approximation of the horizontal momentum equation:

$$-\boldsymbol{k} \times f\boldsymbol{u} = \nu \nabla_z p, \tag{2.10}$$

where ∇_z is the 2D gradient operator along geopotentials, i.e. holding the depth (below the surface geoid) z constant. Thus a geostrophic stream function $\underset{\sim}{\Psi}$ on a particular 2D surface should satisfy[5] [6] [7]

$$\nabla \underset{\sim}{\Psi} = \underset{\sim}{\nu} \nabla_z p. \tag{2.11}$$

The first step is usually to transform the gradient of p from one at constant z to one in the surface. The general coordinate transformation relating 2D gradients of A (a 3D scalar field) from one surface (r) to another (s) in the projected nonorthogonal

[5] For some surfaces, it is useful to instead require $\nabla \underset{\sim}{\Psi} = \nabla_z p$, which gives a stream function for $\rho f \boldsymbol{u}$ rather than $f\boldsymbol{u}$ (McDougall, 1989).

[6] McDougall and Klocker (2010) include the geopotential of the sea surface in their definition, $\nabla \underset{\sim}{\Psi} = \underset{\sim}{\nu} \nabla_z p - \nabla(g\eta)$, where η is the sea-surface height. This distinction is not important: $g\eta$ can be absorbed or not easily into $\underset{\sim}{\Psi}$. (Further note their assumption that $p = 0$ at the sea-surface renders $\nabla(g\eta) = \nabla_p \Phi_0$ with $\Phi_0 = gz$ the surface geopotential, which was their actual expression.

[7] One might put an under-tilde beneath $\nabla_z p$ in (2.11), but that would seem gratuitous.

coordinate system is $\nabla_r A = \nabla_s A + \partial_s A \, \nabla_r s$ (McDougall and Jackett, 1988). This, together with the hydrostatic relation $\partial_z p = -g\rho$, transforms (2.11) into[8]

$$\nabla \underset{\sim}{\Psi} = \underset{\sim}{\nu} \nabla \underset{\sim}{p} + g \nabla \underset{\sim}{z}. \tag{2.12}$$

Further progress requires use of the neutral relation, (2.9), which implies there is a functional relationship between $\underset{\sim}{\nu}$ and $\underset{\sim}{p}$. Specifically,

$$\underset{\sim}{\nu} = \hat{\nu}(\underset{\sim}{p}) \tag{2.13}$$

for some function $\hat{\nu}$. One might notice that (2.9) is of the form (1.8), so $\hat{\nu}$ should be a multivalued function of p. This will be pursued in the next section, but presently let us consider a sufficiently small neighbourhood of a point on the neutral surface that $\hat{\nu}$ is single-valued. Taking the gradient of (2.13) yields

$$\nabla \underset{\sim}{\nu} = \frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p} \nabla \underset{\sim}{p}. \tag{2.14}$$

Now, use the product rule in (2.12), then use (2.14), and finally use the product rule again, as follows:

$$\begin{aligned}
\nabla \underset{\sim}{\Psi}^{(n)} &= \nabla \Big[g \underset{\sim}{z} + \underset{\sim}{\nu} \underset{\sim}{p} \Big] - \underset{\sim}{p} \nabla \underset{\sim}{\nu} && (2.15)\\
&= \nabla \Big[g \underset{\sim}{z} + \underset{\sim}{\nu} \underset{\sim}{p} \Big] - \underset{\sim}{p} \frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p}(\underset{\sim}{p}) \nabla \underset{\sim}{p} && (2.16)\\
&= \nabla \Big[g \underset{\sim}{z} + \underset{\sim}{\nu} \underset{\sim}{p} \Big] - \frac{1}{2} \frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p}(\underset{\sim}{p}) \nabla \underset{\sim}{p}^2. && (2.17)
\end{aligned}$$

This, the first step of an iterative procedure, is done. A superscript (n) has been added to $\underset{\sim}{\Psi}$ to denote that this is the geostrophic stream function in a neutral surface. The second step proceeds exactly as the first:

$$\begin{aligned}
\nabla \underset{\sim}{\Psi}^{(n)} &= \nabla \Big[g \underset{\sim}{z} + \hat{\nu}(\underset{\sim}{p}) \, \underset{\sim}{p} - \frac{1}{2} \frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p}(\underset{\sim}{p}) \, \underset{\sim}{p}^2 \Big] + \frac{1}{2} \underset{\sim}{p}^2 \nabla \left[\frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p}(\underset{\sim}{p}) \right] && (2.18)\\
&= \nabla \Big[g \underset{\sim}{z} + \hat{\nu}(\underset{\sim}{p}) \, \underset{\sim}{p} - \frac{1}{2} \frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p}(\underset{\sim}{p}) \, \underset{\sim}{p}^2 \Big] + \frac{1}{2} \underset{\sim}{p}^2 \frac{\mathrm{d}^2\hat{\nu}}{\mathrm{d}p^2} \nabla \underset{\sim}{p} && (2.19)\\
&= \nabla \Big[g \underset{\sim}{z} + \hat{\nu}(\underset{\sim}{p}) \, \underset{\sim}{p} - \frac{1}{2} \frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p}(\underset{\sim}{p}) \, \underset{\sim}{p}^2 \Big] + \frac{1}{3!} \frac{\mathrm{d}^2\hat{\nu}}{\mathrm{d}p^2} \nabla \underset{\sim}{p}^3, && (2.20)
\end{aligned}$$

having also used (2.13). This procedure can iterate indefinitely, yielding a power series that, if it converges, gives

$$\underset{\sim}{\Psi}^{(n)} = g \underset{\sim}{z} - \sum_{m=1}^{\infty} \frac{(-1)^m}{m!} \frac{\mathrm{d}^{m-1}\hat{\nu}}{\mathrm{d}p^{m-1}}(\underset{\sim}{p}) \, \underset{\sim}{p}^m \tag{2.21}$$

[8] More generally, the gravitational potential Φ can be used as the vertical coordinate, in which case hydrostatic balance is $\mathrm{d}p/\mathrm{d}\Phi = -\rho$, and (2.12) becomes $\nabla \underset{\sim}{\Psi} = \underset{\sim}{\nu} \nabla \underset{\sim}{p} + \nabla \underset{\sim}{\Phi}$.

plus an arbitrary constant, here chosen as zero. The series above resembles a Taylor series, but with two defects: the derivative on $\hat{\nu}$ is off by one, the series starts at $m = 1$. Now, observe that for any constant p_r, $0 = \underset{\sim}{\nu}\nabla p_r$ can be added to the RHS of (2.12), and this merely replaces $\underset{\sim}{p}$ with $(\underset{\sim}{p} - p_r)$ in the above manipulations, except where $\underset{\sim}{p}$ is an argument to $\hat{\nu}$ or any of its derivatives. Define $V(p) = \int_{p_r}^{p} \hat{\nu}(p')\mathrm{d}p'$, such that $\mathrm{d}V(p)/\mathrm{d}p = \hat{\nu}(p)$. With this, (2.21) becomes

$$\underset{\sim}{\Psi}^{(n)} = g\underset{\sim}{z} - \sum_{m=1}^{\infty} \frac{1}{m!} \frac{\mathrm{d}^m V}{\mathrm{d}p^m}(\underset{\sim}{p})\ (p_r - \underset{\sim}{p})^m. \tag{2.22}$$

Note the alternating sign was absorbed into $(\underset{\sim}{p} - p_r)^m$. Finally, adding and subtracting $V(\underset{\sim}{p})$ enables the series to be written starting from $m = 0$:

$$\underset{\sim}{\Psi}^{(n)} = g\underset{\sim}{z} + V(\underset{\sim}{p}) - \sum_{m=0}^{\infty} \frac{1}{m!} \frac{\mathrm{d}^m V}{\mathrm{d}p^m}(\underset{\sim}{p})\ (p_r - \underset{\sim}{p})^m. \tag{2.23}$$

The series above is simply the Taylor series for $V(p_r)$ about $\underset{\sim}{p}$. However, $V(p_r) = 0$. Thus, the exact geostrophic stream function for a neutral surface is

$$\boxed{\underset{\sim}{\Psi}^{(n)} = g\underset{\sim}{z} + \int_{p_r}^{\underset{\sim}{p}} \hat{\nu}(p')\ \mathrm{d}p'.} \tag{2.24}$$

This stream function is easily confirmed to be exact by taking its gradient, using Leibniz integral rule, which yields (2.12) upon using (2.13).

In fact, (2.24) is the exact geostrophic stream function on any surface for which $\underset{\sim}{\nu}$ is a function of $\underset{\sim}{p}$. A not very useful example is a surface of constant ν. More useful are isosurfaces of orthobaric density (de Szoeke et al., 2000). Orthobaric density is a pycnotropic variable (a function only of ν and p) so its isosurfaces have this property. (Note, dividing (2.9) by ν_p, the entire argument can be reversed, with $\underset{\sim}{p}$ a function of $\underset{\sim}{\nu}$; any surface having that property also possesses an exact geostrophic stream function—for instance, an isobaric surface, though again not very useful.)

In all the above expressions, $\underset{\sim}{z}$ is obtained by inverting then integrating hydrostatic balance, $p_z = -\rho g$, in the local water column:

$$\underset{\sim}{z} = \eta - \frac{1}{g} \int_{p_\eta}^{\underset{\sim}{p}} \nu\big(S(p'), \theta(p'), p'\big)\ \mathrm{d}p', \tag{2.25}$$

where η is the sea-surface height and p_η the pressure at the sea-surface (atmospheric loading), and the horizontal dependence of S and θ has been suppressed. Note p_r need not equal p_η. This transformation is standard when studying geostrophic stream functions. For example, multiplying (2.25) by g and taking the gradient in an isobaric

surface yields $\nabla(g\underset{\sim}{z})$ on the LHS, which is the entire RHS of (2.12) when the under-tilde represents an isobaric surface, and thus

$$\underset{\sim}{\Psi}^{(p)} = g\eta - \int_{p_\eta}^{p} \nu\big(S(p'), \theta(p'), p'\big)\ \mathrm{d}p' \tag{2.26}$$

is the exact geostrophic stream function in an isobaric surface: this is the well known "dynamic height".[9] Though $\underset{\sim}{\Psi}^{(p)}$ is really just $g\underset{\sim}{z}$, the form (2.26) is essential because the integral is computable from a hydrographic cast, measuring salinity and potential temperature as functions of pressure. With model data, the actual depth of an isobaric or other such surface may be known, but this is not so for real ocean measurements.

One might argue that neither (2.21) nor (2.24) is a closed form expression for $\underset{\sim}{\Psi}^{(n)}$, since $\hat{\nu}(p)$ is unspecified. Or is it? We are used to thinking there is ambiguity about neutral surfaces, given their embedding issues, and this ambiguity transfers to $\hat{\nu}$. But if we actually, hypothetically, do have a neutral surface, then we do have the function $\hat{\nu}$. If the surface is truly neutral, a scatter plot of $\underset{\sim}{\nu}$ vs $\underset{\sim}{p}$ must have precisely zero scatter about a curve through the data: the data determines $\hat{\nu}$ precisely. Consider that the exact geostrophic stream function on, for example, an isobaric surface, (2.26), has an analogous definition: it too is defined using a function of specific volume, and this function is determined from data. The difference is that this data is taken vertically above or below the location in question. For a neutral surface, $\underset{\sim}{\Psi}^{(n)}$, the problem lies in determining the local neutral surface. Assuming it is known, $\underset{\sim}{\Psi}^{(n)}$ from (2.24) is the exact geostrophic stream function on that surface.

2.2.4 The topology of neutral surfaces

In studying the geostrophic stream function, we took a local view, such that $\hat{\nu}$ was a single-valued function of p. To obtain more and more accurate estimates of $\hat{\nu}$ requires getting more and more data from the local neutral surface. Equivalently, it enables higher and higher derivatives of $\hat{\nu}$ to be estimated at the local point $\underset{\sim}{p}$, so hopefully the power series for $\hat{\nu}(\underset{\sim}{p} + \Delta p)$ converges for all $|\Delta p|$ within some radius of convergence. However, the radius of convergence may be larger than the actual domain of $\hat{\nu}$ as a single-valued function. Moreover, high order polynomial interpolants notoriously overshoot and undershoot; attempting to evaluate a power series like in

[9] Numerical calculations of the dynamic height can be made more accurate by instead using the specific volume anomaly, namely $\delta = \nu(S(p'), \theta(p'), p') - \nu(S_0, \theta_0, p')$ in the integrand of (2.26); this is then called the "dynamic height anomaly". With 64 bit floating point arithmetic, this is no longer necessary (McDougall and Klocker, 2010).

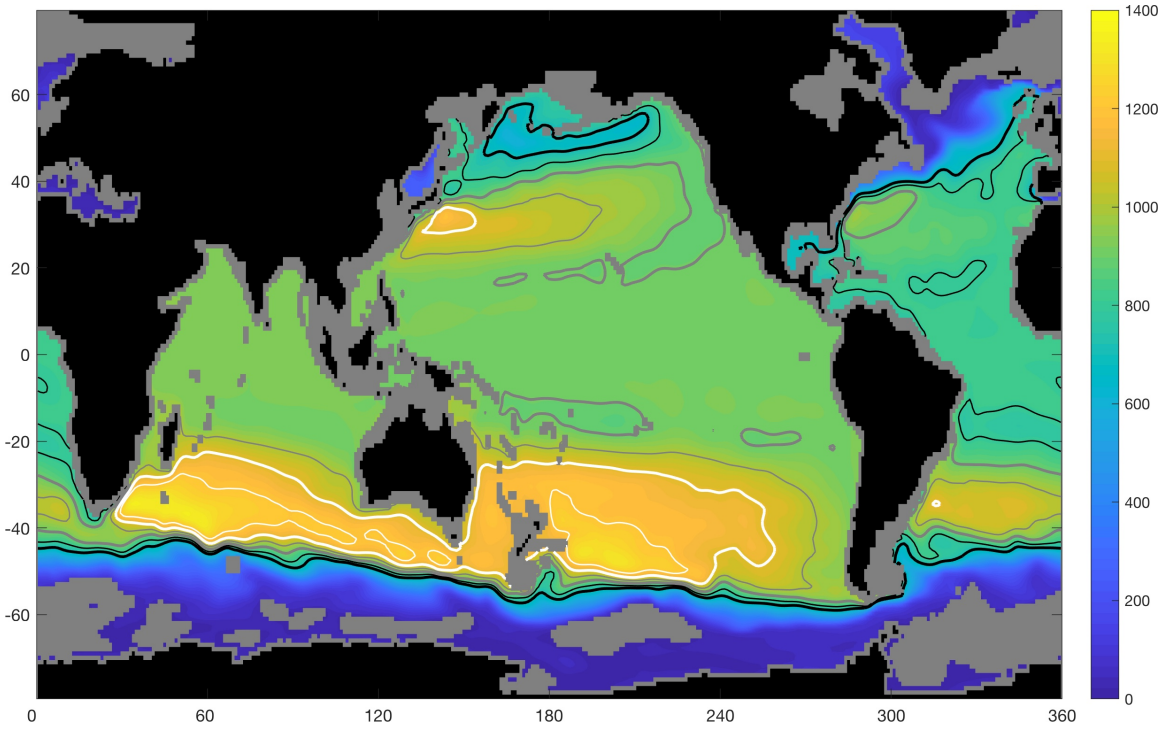


Figure 2.1: The pressure [dbar] of the the $\sigma_1 = 31.9$ isopycnal, where σ_1 is the potential density referenced to 1000 dbar, crudely constructed, for illustrative purposes only, using the 2004—2006 climatological fields from the OCCA of salinity and potential temperature. Grey denotes places where the isopycnal has outcropped or incropped, and black, land. Six contours of the pressure are shown: 700 dbar (thick black), 800 dbar (thin black), 900 dbar (thick grey), 1000 dbar (thin grey), 1100 dbar (thick white), and 1200 dbar (thin white). If this were a neutral surface, the specific volume on the 1100 dbar contour must be the same in both the southern Pacific and southern Indian oceans, whereas it can be different on the 1200 dbar contour in the same basins.

(2.21) is fraught with problems. Instead, (2.24) is much more stable, but it requires knowledge of $\hat{\nu}$ away from the local region.

Thus, we now consider global topological properties of neutral surfaces. Suppose the ocean has zero neutral helicity, so neutral surfaces are embeddable in the ocean. (Our further topological considerations have nothing to do with embeddings.) Because the real ocean has non-zero neutral helicity, what follows is not strictly true, but the effort will yield approximately neutral surfaces.

Where does assuming zero neutral helicity lead us? It means (2.9) is exact, which, being of the form (1.8), has profound topological implications. It implies that $\varkappa$ must be a multivalued function of pressure p. The different branches of this multivalued function are defined on separate domains: in one geographic region, the specific vol-

ume is one value, but in a separate geographic region, it is a different value *at the same pressure*.

Consider, for example, Fig. 2.1 which shows the pressure $\underset{\sim}{p}$ on an isopycnal using the OCCA climatology (Forget, 2010). Six level sets of $\underset{\sim}{p}$ are shown. A level set is composed of any number of connected components, called contours. For instance, the level set of $\underset{\sim}{p} = 1100\,\text{dbar}$ has three contours (the thick white contours). If this isopycnal were truly a neutral surface, the specific volume on the 1100 dbar contour in the subtropical gyre in the north Pacific would be a single value, but different from the single value of specific volume on the 1100 dbar contour in the southern Pacific's subtropical gyre. The specific volume is the same on the southern Pacific 1100 dbar contour and the Indian 1100 dbar contour, because these are evidently the same contour, but the specific volume takes different values between the Pacific and the Indian oceans on the 1200 dbar (thin white) contour. Similar distinctions can arise not with closed $\underset{\sim}{p}$ contours, but with contours that intersect the boundary. For instance, the 800 dbar (thin black) contour in the North Atlantic, which grounds at Europe and North America, supports a distinct specific volume to that on the 800 dbar contour in the South Atlantic, which grounds into Africa. Our task is to determine geographic regions such that any level set of $\underset{\sim}{p}$ has no more than one connected component in each region. Then, we can generate functions $\hat{\nu}$ for each region satisfying (2.13).

The (augmented) Reeb graph of $\underset{\sim}{p} : \mathbb{M} \to \mathbb{R}$ determines these regions, and also how they meet and nest into a global structure. To make this concrete, suppose the Reeb graph has A arcs and N nodes. Each node $n \in \{1, ..., N\}$ is associated with (1) a geographic location $\boldsymbol{x}_n \in \mathbb{M}$ that is a critical point of $\underset{\sim}{p}$, and (2) its critical value $p_n = \underset{\sim}{p}(\boldsymbol{x}_n)$. Each leaf node is a maximum or minimum of $\underset{\sim}{p}$, and each internal node is a saddle point. Each arc $a \in \{1, ..., A\}$ in the Reeb graph is incident upon two nodes l_a and h_a, with $p_{l_a} < p_{h_a}$. (The choice l and h denotes "low" and "high".) In the augmented Reeb graph, each arc a is associated with a geographic region which is a set denoted $X_a \subseteq \mathbb{M}$. From of the data $\{\underset{\sim}{p}(\boldsymbol{x}) : \boldsymbol{x} \in X_a\}$ and $\{\underset{\sim}{\nu}(\boldsymbol{x}) : \boldsymbol{x} \in X_a\}$, we will empirically build a function $\hat{\nu}_a : [p_{l_a}, p_{h_a}] \to \mathbb{R}$. We call $\hat{\nu}_a$ a *branch*[10] of the multivalued function $\hat{\nu}$.

To recap, a topologically important region is associated with an arc of the Reeb graph. Anywhere in this region, the pressure is between the pressures associated with the two nodes adjacent to this arc. We will define a function that returns a specific

[10] This is not meant to connotate that there are discontinuities or singularities of $\hat{\nu}$ at branch points, as in complex analysis. There are not.

volume for any pressure value between these two endpoints. This function is a branch of the multivalued function $\hat{\nu}$.

How should we choose $\hat{\nu}$? We could choose some approximately neutral surface, collect $\underset{\sim}{p}$ and $\underset{\sim}{\nu}$ data on it, then empirically fit $\hat{\nu}$. With $\hat{\nu}$ decided, (2.13) constitutes a root-finding problem for the pressure of the neutral surface, $\underset{\sim}{p}$, so we can update our estimate of the neutral surface. But will the new surface be more neutral?

If the surface is neutral, then (2.9) and (2.14) together require

$$\frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p}(\underset{\sim}{p}) = \underset{\sim}{\nu}_{\underset{\sim}{p}}. \tag{2.27}$$

Taking the curl of (2.9) yields $0 = \nabla\underset{\sim}{\nu}_{\underset{\sim}{p}} \times \nabla\underset{\sim}{p}$, so $\underset{\sim}{\nu}_{\underset{\sim}{p}}$ must *also* be a multivalued function of $\underset{\sim}{p}$:

$$\underset{\sim}{\nu}_{\underset{\sim}{p}} = \hat{\mu}(\underset{\sim}{p}). \tag{2.28}$$

Using (2.28), (2.27) may be integrated from some starting pressure P to get

$$\hat{\nu}(p) = \int_P^p \hat{\mu}(p')\,\mathrm{d}p' + C \tag{2.29}$$

for some arbitrary constant C, bearing in mind the multivalued nature of $\hat{\mu}(p)$ when evaluating this integral. This merely shifts the burden to determining a different multivalued function, $\hat{\mu}$. The upside is that, having $\hat{\mu}$, the neutrality condition (2.9) is guaranteed.

Loosely written, (2.29) hides some complexity. Let us talk through it before defining it more precisely. You are a fluid parcel who starts with pressure P at some geographic location, which is inside one of our regions. This region has a local branch of $\hat{\mu}$. The spirit of adventure strikes, and you set off to explore the deep ocean, barometer in hand. You begin moving towards the boundary of this region, always keeping to the neutral tangent plane, and all the while integrating the local branch of $\hat{\mu}$ using your barometer. When you reach the the end of the domain of this branch of $\hat{\mu}$, you have arrived at the boundary of this region. Assuming you didn't go the wrong way (towards an extremum of $\underset{\sim}{p}$), you may now choose a new region to move through. Continue moving in this way, always integrating the local $\hat{\mu}$ branch as you go, until you have made it to your final destination, with pressure p.

Let us now precisely define (2.29). Select a starting location $\boldsymbol{s} \in \mathbb{M}$, set $P = \underset{\sim}{p}(\boldsymbol{s})$, the pressure at this location, and find the arc a_0 such that $\boldsymbol{s} \in X_{a_0}$. Now choose any continuous path through $\mathbb{M}$, starting at $\boldsymbol{s}$ and finishing at some point $\boldsymbol{f} \in \mathbb{M}$. Let $p = \underset{\sim}{p}(\boldsymbol{f})$, and find the arc $a_{\mathcal{I}}$ such that $\boldsymbol{f} \in X_{a_{\mathcal{I}}}$. The path through $\mathbb{M}$ corresponds

to a walk in the Reeb graph, alternately passing through nodes and along arcs, in the order $n_1, a_1, n_2, a_2, ..., a_{\mathcal{I}-1}, n_{\mathcal{I}}$. The integral (2.29) is, more precisely,

$$\hat{\nu}_{a_{\mathcal{I}}}(p) = \int_P^{p_{n_1}} \hat{\mu}_{a_0}(p') \, \mathrm{d}p' + \sum_{i=1}^{\mathcal{I}-1} \int_{p_{n_i}}^{p_{n_{i+1}}} \hat{\mu}_{a_i}(p') \, \mathrm{d}p' + \int_{p_{n_{\mathcal{I}}}}^{p} \hat{\mu}_{a_{\mathcal{I}}}(p') \, \mathrm{d}p' + C. \tag{2.30}$$

Note, the pressure of the final destination *and* its region (and hence its branch) matter.

Now, we must consider how to match the branches at the saddle points. The following two matching conditions are not actually true of neutral surfaces in a zero neutral helicity ocean, but are needed to construct a well-defined approximately neutral surface, the topobaric surface.

First, we must ensure the branches of $\hat{\nu}$ are well-defined: (2.30) must be independent of the chosen path. Backtracking does not pose a problem, but alternative paths around an island can pose a problem. These create holes in the surface, and thus holes in $\mathbb{M}$, which can create cycles in the Reeb graph. Indeed, even in a zero neutral helicity ocean, islands can cause neutral surfaces to become ill-defined, in the sense that they exist at multiple depths for a given single geographic location. A neutral trajectory around an island returns to its starting geographic location at a different depth than it began, and thus the neutral surface resembles a "spiral staircase", or complex logarithm, around islands. One can imagine replacing the island with seawater such that the neutral helicity at the former island's boundary remains zero, yet is non-zero in the former island's interior; thus a neutral trajectory around the former island creates a height change, as per the discussion around (2.3) to (2.5). By requiring that the integral of $\hat{\mu}$ around any cycle in the Reeb graph is zero, we gain path-independence for $\hat{\nu}$, and the topobaric surface that will be constructed from $\hat{\nu}$ will be well-defined. Thus we require

$$\sum_{i=1}^{\mathcal{I}-1} \int_{p_{n_i}}^{p_{n_{i+1}}} \hat{\mu}_{a_i}(p) \, \mathrm{d}p = 0 \tag{2.31}$$

for any cycle in the Reeb graph, which is a walk $n_1, a_1, n_2, a_2, ..., a_{\mathcal{I}-1}, n_{\mathcal{I}}$ that starts and ends at the same node ($n_1 = n_{\mathcal{I}}$) and has no repeated arcs.

Second, we must join branches at saddle points in a way that maintains neutrality, but in fact this happens automatically. Suppose n is an internal node of the Reeb graph, with degree d. The incident arcs to n are $a_1, ..., a_d$. Node n corresponds to a saddle point $\boldsymbol{s} \in \mathbb{M}$, with pressure $\underset{\sim}{p}(\boldsymbol{s})$. Satisfying neutrality (2.9) here requires, for all $i \in \{1, ..., d\}$, that $\underset{\sim}{\nu_p} \nabla \underset{\sim}{p} = \nabla \hat{\nu}_{a_i}(\underset{\sim}{p}) = \hat{\mu}_{a_i}(\underset{\sim}{p}) \nabla \underset{\sim}{p}$, evaluated at $\boldsymbol{s}$. But $\nabla \underset{\sim}{p} = \mathbf{0}$

here, so $\hat{\mu}_{a_i}$ can take any value at $\underset{\sim}{p}(\boldsymbol{s})$, suggesting that $\hat{\mu}$ might be discontinuous at the saddle pressures. If this is the case, then $\underset{\sim}{\nu_p}$ is discontinuous at $\boldsymbol{s}$, which implies the salinity and/or potential temperature are discontinuous at $\boldsymbol{s}$. Though there may be sharp gradients in some places, we ultimately regard the ocean state as varying continuously in space (after making the continuum assumption, and aside from shocks, which we have no reason to expect are occurring at pressure saddles). Thus, the branches of $\underset{\sim}{\nu_p}$ should meet continuously at the pressure saddles, and we will require

$$\hat{\mu}_{a_1}(\underset{\sim}{p}(\boldsymbol{s})) = \hat{\mu}_{a_i}(\underset{\sim}{p}(\boldsymbol{s})) \qquad \forall i \in \{2, ..., d\}, \tag{2.32}$$

for all internal nodes n of the Reeb graph (which determine d and a_i as above).

Now, purely for insight, we go one small step further. Taking the gradient of (2.28), and also using the chain rule on ν_p as a function of S, θ, and p, we get

$$\frac{\mathrm{d}\hat{\mu}}{\mathrm{d}p}(\underset{\sim}{p})\,\nabla\underset{\sim}{p} = \nabla\underset{\sim}{\nu_p} = \underset{\sim}{\nu_{pS}}\nabla\underset{\sim}{S} + \underset{\sim}{\nu_{p\theta}}\nabla\underset{\sim}{\theta} + \underset{\sim}{\nu_{pp}}\nabla\underset{\sim}{p}. \tag{2.33}$$

Thus, the quantity $\nu_{pS}\nabla\underset{\sim}{S} + \nu_{p\theta}\nabla\underset{\sim}{\theta}$, which differs from $\underset{\sim}{\nu_S}\nabla\underset{\sim}{S} + \underset{\sim}{\nu_\theta}\nabla\underset{\sim}{\theta}$ because ν_S and ν_θ are functions of pressure, must be parallel to $\nabla\underset{\sim}{p}$. This will be satisfied naturally from the relation between $\underset{\sim}{\nu}$ and $\underset{\sim}{p}$, (2.13), but it is *not* satisfied by algorithms that seek $\underset{\sim}{p}$ by minimizing some kind of error from neutrality, i.e. minimizing some norm of $\underset{\sim}{\nu_S}\nabla\underset{\sim}{S} + \underset{\sim}{\nu_\theta}\nabla\underset{\sim}{\theta}$, as Klocker et al. (2009) do. Such an algorithm minimizes the error from neutrality at the expense of this condition on the higher derivatives.

Finally, note from the left equation of (2.33) that $\mathrm{d}^2\hat{\nu}/\mathrm{d}p^2$—which equals $\mathrm{d}\hat{\mu}/\mathrm{d}p$ by combining (2.27) and (2.28) then differentiating— is a multivalued function of p. One can continue this to find $\mathrm{d}^n\hat{\nu}/\mathrm{d}p^n$ is a multivalued function of p for any $n \in \mathbb{N}$. However, $\mathrm{d}^2\hat{\nu}/\mathrm{d}p^2 \neq \underset{\sim}{\nu_{pp}}$. Indeed, using (2.7) to re-write $\nabla\underset{\sim}{S} = -\underset{\sim}{\nu_\theta}\underset{\sim}{\nu_S}^{-1}\nabla\underset{\sim}{\theta}$, (2.33) becomes

$$\left(\underset{\sim}{\nu_{p\theta}} - \underset{\sim}{\nu_{pS}}\frac{\underset{\sim}{\nu_\theta}}{\underset{\sim}{\nu_S}}\right)\nabla\underset{\sim}{\theta} = \left(\frac{\mathrm{d}\hat{\nu}}{\mathrm{d}p}(\underset{\sim}{p}) - \underset{\sim}{\nu_{pp}}\right)\nabla\underset{\sim}{p}. \tag{2.34}$$

We know the LHS term in parentheses is non-zero. In a region where $\nabla\underset{\sim}{p} \neq \mathbf{0}$, $\nabla\underset{\sim}{\theta}$ is parallel to $\nabla\underset{\sim}{p}$ or $\nabla\underset{\sim}{\theta} = \mathbf{0}$, and either way $\underset{\sim}{\theta}$ is a function of $\underset{\sim}{p}$ in the region. Thus $\underset{\sim}{\theta} = \hat{\theta}(\underset{\sim}{p})$ for some function $\hat{\underset{\sim}{\theta}}$. Then (2.33) becomes

$$\frac{\mathrm{d}^2\hat{\nu}}{\mathrm{d}p^2}(\underset{\sim}{p})\,\nabla\underset{\sim}{p} = \left[\left(\underset{\sim}{\nu_{p\theta}} - \underset{\sim}{\nu_{pS}}\frac{\underset{\sim}{\nu_\theta}}{\underset{\sim}{\nu_S}}\right)\frac{\mathrm{d}\hat{\underset{\sim}{\theta}}}{\mathrm{d}p}(\underset{\sim}{p}) + \underset{\sim}{\nu_{pp}}\right]\nabla\underset{\sim}{p}. \tag{2.35}$$

This gives, in regions where $\nabla\underset{\sim}{p} \neq \mathbf{0}$, an expression for $\mathrm{d}^2\hat{\nu}/\mathrm{d}p^2$. Note it differs from $\underset{\sim}{\nu_{pp}}$, because $\underset{\sim}{\theta}$ and $\underset{\sim}{S}$ also vary along the neutral surface, thereby affecting the

specific volume. Were an insulated fluid parcel moved horizontally (holding its S and θ constant) but always allowed to find its level of neutral buoyancy, it would drift off the neutral surface. (Of course, a similar expression could be obtained focusing on S instead of θ.)

2.2.5 The topology of neutral surface geostrophic stream functions

The definition (2.24) for the geostrophic stream function was local, in the sense that $\hat{\nu}$ was considered a single-valued function. Extending this to a global definition requires considering $\hat{\nu}$ a multivalued function. As $\hat{\nu}$ was turned from a local definition, loosely written in (2.29), to a global definition, carefully written in (2.30), so too can $\underset{\sim}{\Psi}^{(n)}$ be made global. After all, its local definition (2.24) is analogous to that for $\hat{\nu}$, (2.29).

Now considering $\underset{\sim}{\Psi}^{(n)}$ globally, we are faced with the same path-dependence problem as for $\hat{\nu}$, which was overcome by adding the constraints (2.31). Should we add further, similar constraints for $\underset{\sim}{\Psi}^{(n)}$? Having found a perfectly neutral surface (say), it would seem strange for necessities of the geostrophic stream function to modify that surface; such a modification may destroy its neutrality.

The values of $\underset{\sim}{\Psi}^{(n)}$ are indeed path-dependent, but only the gradient $\nabla\underset{\sim}{\Psi}^{(n)}$ really matters. One can happily define the path of integration in (2.24) to loop arbitrarily many times around islands, without affecting $\nabla\underset{\sim}{\Psi}^{(n)}$ in the *interior* of any region (on which a branch of $\hat{\nu}$ is defined). One can then imagine islands like spiral staircases for $\underset{\sim}{\Psi}^{(n)}$. They can be followed to higher or lower values of $\underset{\sim}{\Psi}^{(n)}$, and exited at any time, say to move back to an original geographic position. Thus, on a neutral surface (supposing zero neutral helicity), the geostrophic stream function $\underset{\sim}{\Psi}^{(n)}$ is a multivalued function of geographic location! (Indeed, the proof by McDougall (1989) for the existence of $\underset{\sim}{\Psi}^{(n)}$ was local in scope, so this issue could not have been foreseen.)

This does hamper the utility of mapping $\underset{\sim}{\Psi}^{(n)}$, i.e. selecting a single value of $\underset{\sim}{\Psi}^{(n)}$ for each geographic location. In this view, curves along which will $\underset{\sim}{\Psi}^{(n)}$ is discontinuous will emanate from islands, yet this has no relevance whatsoever to the flow field, as at any location, even the boundary between regions, there are branch choices available such that $\nabla\underset{\sim}{\Psi}^{(n)}$ is unaffected. Nonetheless, given these complications, this chapter will not pursue practical calculation of $\underset{\sim}{\Psi}^{(n)}$, but rather focuses on creating topobaric surfaces—a considerable challenge itself.

2.2.6 Discussion

It has been noticed before in the literature that $\underset{\sim}{\nu}$ and $\underset{\sim}{p}$ are functionally related, or equivalently that $\underset{\sim}{S}$ and $\underset{\sim}{\theta}$ are functionally related. In their critique of orthobaric density, McDougall and Jackett (2005a) highlighted differences in the $\underset{\sim}{\theta}$ vs $\underset{\sim}{S}$ relationship between the North and South Atlantic as a cause of poor performance of orthobaric density. They correctly state that there are "two parcels on a neutral trajectory that share the same pressure but different densities", and that this is possible "when a pressure extremum exists on a neutral trajectory". On a one-dimensional neutral trajectory, this is correct, but it is incorrect for a two-dimensional (approximately) neutral surface. Rather, the possibility of multiple branches of the $\underset{\sim}{\nu}$ vs $\underset{\sim}{p}$ relationship opens when there are multiple disconnected contours of $\underset{\sim}{p}$. Considering a 1D trajectory and observing a pressure minimum, say, is not sufficient: a single contour of pressure slightly above that minimum can, in 2D space, wrap around the pressure minimum, in which case the density is the same on either side of the pressure minimum. A new branch need not have opened.

Nonetheless, the Atlantic equatorial region is a good place to observe distinct branches of $\underset{\sim}{\nu}$ vs $\underset{\sim}{p}$, because the trade winds cause Ekman upwelling here. This lifts up deep "isopycnals" in a zonal band, so the situation is quasi-1D. de Szoeke and Springer (2005) extended the original orthobaric density to use a different virtual compressibility for the North and the South Atlantic. The resulting orthobaric densities were aligned at the equator where both orthobaric density surfaces were assumed to be level, but the orthobaric density was discontinuous here. A fictitious force must be applied to a water parcel crossing this discontinuity to keep it on the same orthobaric density surface (de Szoeke and Springer, 2005). This technique was taken further by de Szoeke and Springer (2009), who crafted "extended orthobaric density" by further dividing the $\underset{\sim}{\nu}$ vs $\underset{\sim}{p}$ relationship into arbitrarily many latitude bins. But latitude circles are not the topologically correct boundaries between $\underset{\sim}{\nu}$ and $\underset{\sim}{p}$ relations, so this still suffers the same discontinuities as did the single split at the equator, albeit spread out (de Szoeke and Springer, 2009).

With the full topological analysis, such discontinuities are not necessary. The Reeb graph enables different functional relations between $\underset{\sim}{\nu}$ and $\underset{\sim}{p}$ in different regions, and tells us where to match them continuously.

2.2.7 Boussinesq models

Boussinesq ocean models swap the *in-situ* density ρ for a constant reference density ρ_c in the horizontal momentum equations, thus the geostrophic relation becomes

$$-\boldsymbol{k} \times f\boldsymbol{u} = \rho_c^{-1}\nabla_z p = \rho_c^{-1}\nabla\underset{\sim}{p} + g\rho_c^{-1}\underset{\sim}{\rho}\nabla\underset{\sim}{z}. \tag{2.36}$$

The task of a geostrophic stream function is to bring $\underset{\sim}{\rho}$ inside the gradient with $\underset{\sim}{z}$. The logic of Section 2.2.3 can be again followed, but it is vastly easier to work with ρ rather than ν.

A second, perhaps less appreciated, change is that the *in-situ* density, which is still used in the vertical momentum equation (possibly reduced to hydrostatic balance), is calculated from the equation of state not with the pressure, but with the depth, as $\rho(S, \theta, -g\rho_c z)$ (Young, 2010). So, the gradient of ρ is

$$\nabla\rho = \rho_S\nabla S + \rho_\theta\nabla\theta + \rho_z\nabla z \tag{2.37}$$

where $\rho_z = -g\rho_c\rho_p$. The neutral surface relation (2.9) becomes

$$\nabla\underset{\sim}{\rho} = \rho_{\underset{\sim}{z}}\nabla\underset{\sim}{z}. \tag{2.38}$$

Thus for a Boussinesq model, we work with ρ rather than ν, and the roles of pressure and depth get flipped. Thus, a function $\hat{\rho}(z)$ takes the place of $\hat{\nu}(p)$, and $\hat{\pi}(z) = \mathrm{d}\hat{\rho}(z)/\mathrm{d}z$ takes the place of $\hat{\mu}(p)$.

2.2.8 Summary of key equations

To recap, the Reeb graph of $\underset{\sim}{p}$ determines the regions in which different branches of $\hat{\mu}$ apply. These branches are fit to (2.28), subject to conditions (2.31) and (2.32) determined with the help of the Reeb graph. Then, path integrals of $\hat{\mu}$ give $\hat{\nu}$, as encoded by walks on the Reeb graph, (2.30). Finally, the pressure and specific volume on the approximately neutral surface are compared using $\hat{\nu}$ in (2.13), and updates to $\underset{\sim}{p}$ could be made as necessary in an iterative scheme (Section 2.3).

2.3 Methods

This section describes practical application of the theory of Section 2.2 to generate topobaric surfaces. Everything will be described as for a Boussinesq ocean, but is easily translatable to the non-Boussinesq ocean. This is made easy by treating $z > 0$, and likewise the code does this. The basic procedure is as follows.

1. Begin with an approximately neutral surface, with depth $\underset{\sim}{z}$.
2. Calculate the Reeb graph of $\underset{\sim}{z}$, involving the following sub-steps:
   (a) pre-process $\underset{\sim}{z}$;
   (b) calculate the Reeb graph;
   (c) post-process the Reeb graph;
   (d) simplify the Reeb graph.
3. Fit a branch $\hat{\pi}(z)$ for each region, i.e. each arc on the Reeb graph, while satisfying the matching conditions.
4. Integrate $\hat{\pi}(z)$ to get $\hat{\rho}(z)$, using the constant of integration to maintain the general depth of the surface.
5. Solve for the new depth $\underset{\sim}{z}$ satisfying $\rho(S(\underset{\sim}{z}), \theta(\underset{\sim}{z}), \underset{\sim}{z}) = \hat{\rho}(\underset{\sim}{z})$ in each water column.
6. Return to Step 2, unless some convergence test is passed.

Each of these steps will be discussed in turn.

I construct topobaric surfaces using this algorithm from ECCO2 (Menemenlis et al., 2005) data on 22–24 December 2002. A single archived time-step is chosen to make the task as hard as possible: the Reeb graph of a smoother climatology is considerably simpler. Nonetheless, the simplification algorithm makes it possible to work with this eddy-rich data.

2.3.1 Initial guess

Any approximately neutral surface can be used as an initial guess. I use potential density, to ensure that the algorithm works even when the initial guess is not very good, globally. Specifically, I use potential density referenced to 2000 m (i.e. to $10^{-4} \cdot 2000\,\mathrm{m} \cdot g\rho_c$, in dbar) with isovalue $1037.0\,\mathrm{kg\,m^{-3}}$. The depth is around 2000 m in the main ocean basins; the Arctic is disconnected, and the Indonesian throughflow is closed.

2.3.2 Calculation of the Reeb graph

Fast, robust, general computation of Reeb graphs from discretised data has only been achieved recently. Carr et al. (2003) developed the first fast algorithm to compute the contour tree in any dimension. This was leveraged by Doraiswamy and Natarajan (2013) to quickly compute the Reeb graph in any dimension. I use the freely available Java software of Doraiswamy and Natarajan (2013), called ReCon, available through `http://vgl.csa.iisc.ac.in/`.

Using it requires some pre-processing, some understanding of the algorithm, and some post-processing, each of which is discussed in turn, next. As a quick overview, selected islands (and seamounts, outcroppings, etc, essentially any hole in $\underset{\sim}{z}$) are filled with extremely large values such that all contours go around them, then $\underset{\sim}{z}$ is turned into a simplical decomposition. The Reeb graph is calculated using ReCon. Post-processing of the Reeb graph removes nodes and arcs corresponding to those filled holes. Finally, the Reeb graph is simplified by leaf pruning and vertex reduction (Carr et al., 2010), to remove insignificant maxima and minima of $\underset{\sim}{z}$.

Where this processing involves graph theory, I use the MATLAB toolbox `GAIMC` by David Gleich, available at `www.mathworks.com/matlabcentral/fileexchange/24134`.

2.3.2.1 Data preparation

Some pre-processing is needed and/or helpful before calculating the Reeb graph. First, as it greatly helps with the simplification algorithm, most missing data of $\underset{\sim}{z}$ are filled with extremely large (and unique) values, larger than the maximum valid value of $\underset{\sim}{z}$. This removes holes in the manifold, thereby removing cycles from the Reeb graph. Contours of $\underset{\sim}{z}$ will pass entirely around these filled regions, so in physical space the effect is to connect previously disconnected contours, which joins arcs in the graph and closes cycles.

I filled all islands other than Antarctica, Madagascar, and New Zealand (which, at this depth, has united). Many other islands have already joined one mainland, including Australia, the Indonesian archipelago, Greenland, Iceland, the Americas, and Afro-Eurasia.

As in our case here, discretised data is commonly measured or calculated on a rectilinear grid. In two dimensions, rectilinear data is often treated with bilinear interpolation. This simply performs linear interpolation in one dimension, then the next, and as such is a quadratic interpolant. Non-linear interpolants introduce new

critical points, thereby changing the topology in a way that is dependent upon choice of interpolant.

This problem is avoided by using a simplical mesh and a piecewise linear interpolant. In two dimensions, a simplical mesh is a collection of triangles; in three dimensions it is a collection of tetrahedrons, and so on to simplices of arbitrary degree. If measurements of some real scalar field are made at points $\{x_1, ...x_n\}$, obtaining data $\{y_1, ..., y_n\}$, then a simplical mesh can be defined having vertices $\{x_1, ...x_n\}$, and a piecewise-linear function f having $f(x_i) = y_i$ for all $i = 1, ..., n$. This f is not smooth and so is not technically Morse, but the choice of simplical mesh with linear interpolants guarantees that f has no critical points inside the mesh, and thus the Reeb graph of f will be identical to the Reeb graph of a smooth function g that also has $g(x_i) = y_i$ for all $i = 1, ..., n$ (Carr et al., 2003).

The simplical mesh decomposition of arbitrary data points is open to choice. Assume our data arrives on a rectilinear grid. One simplical mesh is obtained by adding a new data point at the centre of each rectangle, four-way averaging the data onto that new data point, then splitting each rectangle into four triangles. This is closest to how contours are typically computed from rectilinear data using the marching cubes algorithm. Another option is simply to split each rectangle into two triangles. The choice of simplical decomposition can matter (Carr et al., 2006), but I discerned no significant difference using these methods on global oceanographic data. Therefore, the simpler, two triangle method is used. Where a rectangle has one data point missing, only one triangle is produced.

For simplicity, only one connected component of the approximately neutral surface is handled at a time, guaranteeing the Reeb graph is a connected graph. A graph is constructed from the simplical mesh, having a node at every vertex of the mesh, and an arc between two nodes when their corresponding vertices share a face of the mesh. The connected components of this graph are calculated (Tarjan, 1972), and only the largest one retained. (Faster, image manipulation methods used on the original rectilinear grid are not sufficient, because the simplical decomposition can remove the odd grid point, such as those with ground for 7 of 8 neighbours.)

Finally, any duplicate values of $\underset{\sim}{z}$ are found and perturbed by a machine-precision amount until all vertices have unique values, to satisfy the third property of Morse functions, discussed in Section 1.4.2.

2.3.2.2 The Reeb graph from ReCon

Both the ReCon algorithm (Doraiswamy and Natarajan, 2013) and the Carr et al. (2003) contour tree algorithm underlying ReCon are described briefly here. ReCon computes the Reeb graph in $\mathcal{O}(v \log v + sn)$ time, where v is the number of vertices, n is the number of triangles, and s is the number of saddles of the simplical mesh.

It is useful to further classify saddles as an *up-forking saddle* if a new level set is created just above the critical value, and and *down-forking saddle* if a new level set is created just below the critical value.

The Carr et al. (2003) algorithm is to sort the vertices (hence the $\mathcal{O}(v \log v)$ dependence) by their function value, then compute the *join tree* in a single sweep of the data, followed by a second sweep to compute the *split tree*. First define a *join component* at p in the image of f to be a connected component of the set $\{x \in \mathbb{M} : f(x) \leq p\}$. Then, the join tree is like a Reeb graph except the arcs correspond to join components, rather than components of the level sets. When $p = \max f$ the join component is the whole manifold; here a new node is created, and this is the root of the join tree. Sweeping p downwards, new join components are created at down-forking saddles of f, and these are internal nodes of the join tree. Finally, join components can be destroyed when p passes a local minimum of f, and these are the other leaf nodes of the join tree. The opposite procedure produces the split tree. The two trees are then merged using an iterative algorithm that, essentially, ensures all internal nodes of both the join and the split tree are copied to the contour tree. For full details, see Carr et al. (2003).

Doraiswamy and Natarajan (2013) first identify all possible loop saddles, which are nodes in the Reeb graph that uniquely identify a cycle in the Reeb graph. An easy and convenient identification is the lowest up-forking saddle in the loop. From the perspective a downward sweep through the data, the loop is destroyed at this saddle. The function value at these points is then used to split the data into loop-free regions, upon which the contour tree is computed using the Carr et al. (2003) algorithm. Finally, the Reeb graph is constructed by merging maxima-minima pairs in the contour trees that originated from a common loop saddle.

ReCon reads its input and gives its outputs using external files. I added new methods to the ReCon's source code to enable it to communicate directly with MATLAB, or many other languages including Python. The input to ReCon is the vertices defining the simplical mesh, and the faces of the simplices. For a d-dimensional manifold, the former is an v by $(d + 1)$ array: each of the v rows defines a vertex by its

first d columns, and the function value on that vertex by the $(d+1)$ column. The latter is a n by 3 array: each of the n rows defines a face by indexing its 3 vertices.

The output of ReCon is several arrays that describe the Reeb graph, having N nodes and A arcs. There are arrays of length N storing the node vertices (the index of the vertex in the simplical mesh), node values (the value of the function at the node's vertex), and node types (indicating correspondence to regular points, minima, saddles, or maxima). There are three arrays of length A: the lower node attached to each arc, the higher node attached to each arc, and the arc identifier (usually just the numbers 1 through A). There are two ragged arrays[11] of length N, storing the up-arcs attached to each node, and the down-arcs attached to each node. The augmented Reeb graph also has one ragged array of length A, storing identifiers for the associated vertices for each arc (usually the numbers 1 through V where V is the number of vertices in the simplical mesh). This data structure, rather than a sparse adjacency matrix, can represent a graph with multiple arcs between the same two nodes.

2.3.2.3 Reeb graph simplification

The Reeb graph of ς, the depth of the aforementioned isopycnal from a 3-day average of ECCO2 data (1/4° resolution), has over 40,000 arcs. The ocean is full of turbulence, which creates many small extrema in the pressure field on an isopycnal. In principle, these could all have different density vs pressure relations, but in practice it is better to join together these small-scale fluctuations into larger regions. This is achieved by simplifying the Reeb graph, reducing the number of nodes and arcs. Exactly how the joining is done and what is meant by "small-scale" is a matter of choice, and corresponds to how one simplifies the Reeb graph.

The method I use broadly follows the leaf-pruning and vertex-reduction method of Carr et al. (2010), which is described first. However, this method is designed for Reeb graphs without cycles (contour trees), and so I extend it to handle cycles in a manner similar to Pascucci et al. (2007). Graph simplification is a complicated task; the following description is intended to be complete for purposes of reproducing my simplification algorithm, but as a result is highly technical.

A Reeb graph is an undirected graph with values associated to each node: these are the function value at the extrema or saddles. Given a node n, an arc a adjacent to n is called an up-arc from n when the node at the other end of a has a larger value. Similarly, a is a down-arc from n when the node at the other end has smaller value.

[11] A ragged array of length n is an array of length n whose elements are arrays of variable length.

An arc a connects two nodes, and is thus an up-arc from one node and a down-arc from the other. Where unambiguous, the "from n" is often dropped. Furthermore, Reeb graph is used throughout as shorthand for the augmented Reeb graph.

Leaf pruning removes a leaf[12] from the graph and merges the leaf arc onto another arc from its parent in the same direction. Pruning a leaf whose only arc is a down-arc (the leaf is a local maxima) merges the leaf arc onto another down-arc from the parent. Pruning a leaf whose only arc is an up-arc (the leaf is a local minima) merges the leaf arc onto another up-arc from the parent. To merge an arc onto another, all vertices associated to the first arc are associated to the second, and the first arc is deleted.[13] In physical space, merging two arcs corresponds to connecting their previously disconnected level sets. In the language of topology, the topological or physical space is, in the region of the two arcs, glued together at common values of the function. One can also think of leaf pruning as flattening the region around an extremum to the level of its nearest saddle, though this is not quite correct: the function is unchanged, whereas the *space* is changed. With this in mind, a leaf is "protected" if it is the last up-arc or the last down-arc from its parent: there is no space with which to glue this region.

The algorithm begins by building a priority queue[14] of all leaves. Their measure of priority is typically the "persistence" (Edelsbrunner et al., 2000), which is the difference in function value between leaf and parent, always taken to be positive. Carr et al. (2010) describes how other local geometric measures of an arc can be used. I use the vertex count (thinking backward through the simplical mesh, this is the number of grid points) on each augmented arc as a local geometric measure of area. Finally, the priority of an arc is a weighted sum of the persistence and the vertex count, both divided by their maximum values over all arcs; I take an equal weighting.

The simplification algorithm proceeds iteratively. On each iteration, an element is popped from the queue and processed. If it is leaf (with the extension, it can be a

[12] A leaf is a degree one node. Therefore it has exactly one arc adjacent to it, so we can unambiguously also refer to that arc as a leaf arc, and unambiguously identify the node at the other end of the leaf arc as the parent node.

[13] The associated vertices are handled efficiently during the iterative procedure by maintaining a partition of the arcs: a set of disjoint sets whose union is all the arcs. Only after completing the iteration are the associated vertices merged onto the remaining arcs. Also, an array of deleted vertices is maintained.

[14] A priority queue is a data structure, holding arbitrary objects together with a "priority" for each object. One can "push" a new object into the queue with a given priority, "pop" the most urgent object off the queue, and check when the queue is empty. Because the priority is maintained as a sorted vector, the time complexity of these operations is constant.

cycle) and is not protected, it is pruned immediately. Protected leaves are discarded. When the parent has more than two arcs in the leaf's direction, a choice must be made. I choose to merge the leaf arc onto the arc with the largest persistence. This minimizes an undesirable but not really problematic effect of leaf pruning, which is that it destroys the property of the Reeb graph that all vertices on an arc have function values intermediate to the function values at the arc's two nodes. If the parent was a degree three node, after leaf pruning it is a degree two node. In most cases it has one up-arc and one down-arc, and thus is a regular node: the topology no longer changes at this node, and thus it can be deleted and its down-arc merged onto its up-arc. This is called "vertex reduction" (Carr et al., 2010). Moreover, vertex reduction can make a protected leaf node become prunable. Thus, any neighbour of the reduced node that is a leaf is pushed into the priority queue, with priority given by the function value difference between the leaf and its new parent. The iteration is stopped when a desired number of arcs remain.

Carr et al. (2010) proved that, for an acyclic Reeb graph (a contour tree), this algorithm will not stall until only two nodes remain. If the graph has one cycle, for example, the iteration will stall when there are only four remaining nodes, connected consecutively but with a double arc connecting the middle two. Indeed, this algorithm will never remove a cycle from the graph, hence why select islands were filled in the pre-processing step.

It must be said that there are many other ways of simplifying the Reeb graph, and the results do differ between methods. Graph theoretic techniques such as centrality measures can be used to choose cut vertices (a node whose removal increases the number of connected components of the graph). Or, considering the shortest path between pairs of leaves (a subset of the all-pairs shortest path problem), extrema of the function value on such paths can be used to choose cut vertices. Or, cycles can be closed in a similar iterative simplification and no holes filled in pre-processing.[15]

[15] The importance of cycles can be classified similarly to leaf nodes, based on their maximum and minimum nodes (Agarwal et al., 2004). Closing a cycle is done by deleting all arcs in the cycle, ordering the cycle's nodes by function value, then adding new arcs between adjacent nodes (Pascucci et al., 2007, see Fig. 5b of). This removes one arc from the graph. It also adjusts the associated vertices of all arcs in the cycle. Each new arc is associated with all vertices having function value between the range of its incident nodes function values, which are unchanged. If the maximum or minimum nodes from the cycle become regular as a result, they are immediately reduced. This method becomes complicated for cycles that, walking from the minimum node to maximum node, have a non-monotonic sequence of associated function values, but can nonetheless still be handled. Closing cycles can modify the nodes and arcs associated with other cycles, so after closing any cycle, the next smallest cycle must be recomputed, ideally using a least common ancestors method (Agarwal et al., 2004).

Ultimately, for this oceanographic application, the simplification method must not be too sensitive to small changes in the graph, otherwise iterations to find the topobaric surface might diverge. The leaf pruning method performs well in this regard, and precisely choosing which islands are filled helps with this.

2.3.3 Multisplines to fit $\hat{\pi}$ with matching conditions

All branches of $\hat{\pi}$ are fit empirically, while satisfying the matching conditions (2.31) and (2.32), using a new tool called `MULTISPLINEFIT`, which extends the `SPLINEFIT` MATLAB function by Jonas Lundgren, available at `www.mathworks.com/matlabcentral/fileexchange/13812`, both of which are described here. This description begins simply with polynomials, then extends to splines, which are piecewise polynomials.

To fit a polynomial of order n, $f(x) = \sum_{i=0}^{n-1} a_i x^i$, to some data $\boldsymbol{x} = \{x_i\}_{i=1}^m$ and $\boldsymbol{y} = \{y_i\}_{i=1}^m$, both $(m \times 1)$ column vectors, one can find a least-squares solution to $X\boldsymbol{a} = \boldsymbol{y}$, where X is the $m \times n$ matrix with $X_{i,j} = x_i^{j-1}$, i.e. the Vandermonde matrix of $\boldsymbol{x}$, $V(\boldsymbol{x})$. The solution vector, $\boldsymbol{a}$, gives the coefficients $\{a_i\}_{i=0}^{n-1}$.

Now suppose we want the above fit subject to a single linear constraint $f(x_0) = y_0$, which can be written as $B\boldsymbol{a} = y_0$ where $B = V(x_0) = [1, x_0, ..., x_0^{n-1}]$. Linear constraints on a derivative, say $f'(x_0) = y_0$, is written with $B = [0, 1, 2x_0, ..., (n-1)x_0^{n-2}]$, and this is easily extended to higher derivatives. A set of c such linear constraints is thus represented as $B\boldsymbol{a} = \boldsymbol{y_0}$, where B is a $c \times n$ matrix and $\boldsymbol{y_0}$ an $c \times 1$ column-vector.

To fit the model subject to these linear constraints, first find a particular solution $\boldsymbol{a_0}$ satisfying $B\boldsymbol{a_0} = \boldsymbol{y_0}$, and the null-space Z (an $n \times z$ matrix, with $z \leq n$) of B satisfying $BZ = 0$. Then, find the least-squares solution to $XZ\boldsymbol{b} = \boldsymbol{y} - X\boldsymbol{a_0}$, and obtain the final solution $\boldsymbol{a} = \boldsymbol{a_0} + Z\boldsymbol{b}$. One can immediately see that $B\boldsymbol{a} = B\boldsymbol{a_0} + BZ\boldsymbol{b} = \boldsymbol{y_0}$ is guaranteed, and $||X\boldsymbol{a} - \boldsymbol{y}||^2 = ||X\boldsymbol{a_0} + XZ\boldsymbol{b} - \boldsymbol{y}||^2$ is indeed minimized.

Now, suppose the model f is instead a piecewise polynomial of P pieces, with knots $\{k_i\}_{i=1}^{P+1}$, usually taking $k_1 = \min(\boldsymbol{x})$ and $k_{P+1} = \max(\boldsymbol{x})$:

$$f(x) = \sum_{i=0}^{n-1} a_{i,j} x^i, \qquad k_j \leq x < k_{j+1}. \tag{2.39}$$

By collapsing $a_{i,j}$ to an $(nP \times 1)$ column vector, this can easily be fit with the above machinery: linear constraints are used to ensure the pieces match at the knots with the desired class of continuity. Note that X becomes a block diagonal matrix; the j'th

block is $V(\boldsymbol{x}_j)$ where $\boldsymbol{x_j} = \{x_i \in \boldsymbol{x} : k_j \leq x_i < k_{j+1}\}$. For large P, X is represented as a sparse matrix.

So far, [1, x, ..., x^{n-1}] has been used as a basis for the space of polynomials of order n. This is badly conditioned, and so for numerical purposes a different basis will be used, namely B-splines, but the above ideas are essentially unchanged.

A spline is a piecewise polynomial of order n that is continuous at all knots, and, assuming all knots are distinct, its first $n-2$ derivatives are also continuous at all knots; e.g. when $n = 4$ we have a continuous piecewise cubic polynomial with continuous first and second derivatives at the knots, but discontinuous third derivatives. At the first knot k_1, there are n degrees of freedom, which completely determine the polynomial from k_1 to k_2. At each internal knot k_j, $2 \leq j \leq P$, an additional degree of freedom is gained, for the choice of the $(n-1)$'th derivative of the next piece. In total, there are $n + P - 1$ degrees of freedom. B-splines are a basis for the function space of splines of order n over given knots $\{k_i\}_{i=1}^{P+1}$, and there are $n + P - 1$ of them, for all distinct knots. Then, when solving $X\boldsymbol{a} = \boldsymbol{y}$ with X appropriate to B-splines, the coefficients $\boldsymbol{a}$ give the spline f by linear combination of B-splines (De Boor, 1978). Essentially, this builds many linear constraints into the problem. Not only does this make finding the null space Z easier, it helps keep XZ sparse.

`SPLINEFIT`, as just described, is now extended to a multivalued spline fitting algorithm, `MULTISPLINEFIT`, that simultaneously fits multiple splines over arbitrary domains, each to its own data, but subject to matching conditions between certain splines at certain locations, and subject to integral constraints involving multiple splines. Linear integral constraints such as $\int_{x_1}^{x_2} f(x)\ \mathrm{d}x = y_0$ are added with $B = [(x_2 - x_1), ..., (x_2^n - x_1^n)/n]$, and this can be easily extended to repeated integrals.

Thus a spline $\hat{\pi}_a(z)$ is fit for each arc $a \in \{1, ...A\}$ in the Reeb graph, from the data $\underset{\sim}{z}(X_a)$ and $\underset{\sim}{\rho}(X_a)$. The matching conditions (2.32) are easily added as the above linear constraints (simply move one term to the other side so the constant $y_0 = 0$). If desired, continuity of higher order derivatives can be imposed between adjacent splines.

For the matching conditions (2.32), we must first determine the cycles in the graph. In fact, not all cycles are needed. Adding, modulo 2, the arcs of two (different) cycles produces a third cycle. But if (2.32) holds for the first two cycles, it is guaranteed to hold for the third as well. Thus, we need a cycle basis, from which all cycles of the Reeb graph can be constructed by such linear combinations. The cycle basis is constructed by first finding a minimum spanning tree (Prim, 1957), then for each arc

a in the Reeb graph that is missing from the minimum spanning tree, a breadth-first search in the minimum spanning tree yields the shortest path between the adjacent nodes of a, and this path together with a is a cycle, and in particular an element of the cycle basis. Then, (2.32) is added as a constraint for each cycle in the cycle basis.

I use piecewise linear functions ($n = 2$) to fit each branch $\hat{\pi}_a$, and choose the number of pieces P for each spline as $\operatorname{ceil}((\max \precsim - \min \precsim)/300\,\mathrm{m})$, i.e. one piece when the range of $\precsim$ is less than 300 m, two when it is between 300 m and 600 m, and so forth. This yields piecewise quadratic C^1 branches $\hat{\rho}_a$. However, there is considerable flexibility in how to fit the branches $\hat{\pi}_a$.

2.3.4 Integration to find $\hat{\rho}$

Having chosen the branches $\hat{\pi}_a$, they are integrated via (2.30) to obtain the branches $\hat{\rho}_a$. First, the point $\boldsymbol{s} \in \mathbb{M}$ is chosen. For simplicity, I choose the global minimum of $\precsim$ (i.e. shallowest—recall we are being $z > 0$ people), which is the lowest node on the Reeb graph. Suppose that node is n_1, having value z_1; being a leaf, it has only one adjacent arc, a_1. The branch $\hat{\pi}_{a_1}$ is integrated (analytically—this is trivial for splines) from z_1 to arbitrary z, giving the branch $\hat{\rho}_{a_1}$. Now, let $n_2 = h_{a_1}$ be the other node connected to a_1, with value z_2. A similar integration gives the branches of $\hat{\rho}$ on all other arcs adjacent n_2, and the constant of integration C is taken as $\hat{\rho}_{a_1}(z_2)$, to ensure continuity.

This procedure is continued until all branches have been integrated. Specifically, a breadth-first search is started from the lowest node on the Reeb graph. Each step of a breadth-first search "discovers" a node adjacent to a previously discovered node. The arc connecting these nodes, at each step, is integrated as above. Any arc not integrated by the end of the breadth-first search is then integrated.[16] The matching condition (2.31) ensures the resulting branches are independent of the processing order.

Finally, one constant of integration remains to be chosen, to be added to all branches $\hat{\rho}_a$. (It begins as the constant of integration for the first branch, $\hat{\rho}_{a_1}$, but the above procedure adds this to all other branches.) A deep cast is chosen in the central Pacific at $(180°\mathrm{E}, 0°\mathrm{N})$, and the constant chosen so that the updated surface lies at the same depth as the old surface there, i.e. (2.13) is forced to hold at this cast. This particular cast was chosen because it is deep, and neutrally connected to a large volume of the global ocean; other casts can, of course, be used.

[16] Essentially, this is a topological ordering for processing the arcs.

2.3.5 Root finding of $\underset{\sim}{z}$

With all branches of $\hat{\rho}$ chosen, (2.13) now constitutes a root-finding problem for $\underset{\sim}{z}$, the depth of the updated surface. This is done by bisection on each cast[17]. Multiple roots are possible, so an initial guess is provided, based on $\underset{\sim}{z}$ from the previous iteration. A small interval around the initial guess is tested for a sign change at its limits. If a sign change is found, bisection proceeds inside this interval. Otherwise, the interval is (geometrically) expanded until a sign change is found and bisection can proceed, or the shallowest and deepest grid cells are reached. If the latter, no root is found, meaning the updated surface has outcropped or incropped. (There is currently no capacity for "wetting", whereby subsequent iterations retest casts that previously outcropped or incropped. This would require a way to attach a cast to a particular branch of $\hat{\rho}$, which could be ambiguous at the boundary between regions, though the impact of this is likely insignificant.)

2.3.6 Iteration

After one iteration of the above procedure, we have an updated estimate of the depth $\underset{\sim}{z}$ an approximately neutral surface. As such, the Reeb graph of the new $\underset{\sim}{z}$ will be different, so this procedure is iterated. The convergence is quite rapid. For the chosen isopycnal, the root-mean-square change of $\underset{\sim}{z}$ after the first iteration is 30 m, whereas by the fourth iteration it is 0.7 m. I stop when the root-mean-square change of $\underset{\sim}{z}$ is less than 1 m—four iterations in this case.

2.4 Results

Results from the first step (for descriptive purposes) of the iteration of the above algorithm are described now, followed by a comparison of the converged topobaric surface with other approximately neutral surfaces.

2.4.1 Initial surface

The depth of the initial isopycnal (see Section 2.3.1) is shown in Fig. 2.2, here reverting to $z < 0$ for display. Many distinct regions are evident. For example, just north of the ACC there are five major regions: the Atlantic, the western Indian, south of Australia, the western Pacific, and the eastern Pacific. The split of the southern Pacific subtropical gyre is particularly prominent. The level sets for two depths

[17]MATLAB's code generation is used to turn this into efficient C MATLAB executables (MEX).

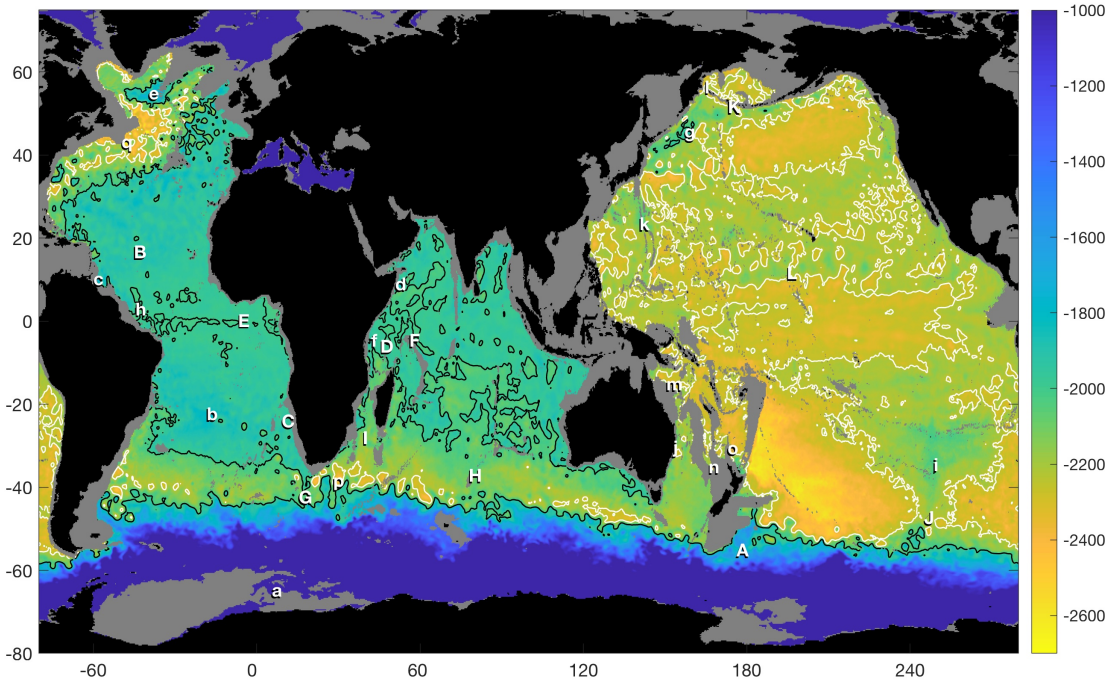


Figure 2.2: The depth, $\underset{\sim}{z}$, on the $\sigma_2 = 37.0$ isopycnal (now with $\underset{\sim}{z} < 0$). Two level sets of $\underset{\sim}{z}$ are shown: $-2000\,\mathrm{m}$ (black) and $-2250\,\mathrm{m}$ (white). The isopycnal has outcropped or incropped in grey regions. Note the saturated colourbar. Letters indicate some nodes of the simplified Reeb graph in Fig. 2.4, and are placed at the corresponding critical points: extrema of $\underset{\sim}{z}$ for leaf nodes (lower case), or saddle points of $\underset{\sim}{z}$ for internal nodes (upper case)).

are shown; as there are many connected components (contours) for each level set, the unsimplified Reeb graph will be quite complex. As was mentioned, it contains over 40,000 arcs, each corresponding to different regions in the ocean. This reveals a practical difficulty that a region corresponding to an arc of the Reeb graph can contain grid points scattered far apart in geographic space, when contours are filamentary.

As such, the initial Reeb graph is simplified. The regions associated to each arc are mapped in Fig. 2.3, and the graph itself is shown in Fig. 2.4.[18] Examining the two maps together helps one understand the Reeb graph. Starting at a in shallow near Antarctica, one moves deeper and northward until A, when a choice must be made to go west or east of New Zealand. This is the beginning of a cycle in the Reeb graph. Going east into the Pacific, another choice must be made at J whether to move up to the central south Pacific (to i), or down towards the western Pacific. If down, one can

[18] Drawing graphs—deciding the placement of their nodes, and the shape of their arcs—is non-trivial. The drawing shown is a simple choice that conveys useful information about depth and longitude, but at the cost of some arcs overlapping. See Heine et al. (2011) for further discussion of drawing Reeb graphs.

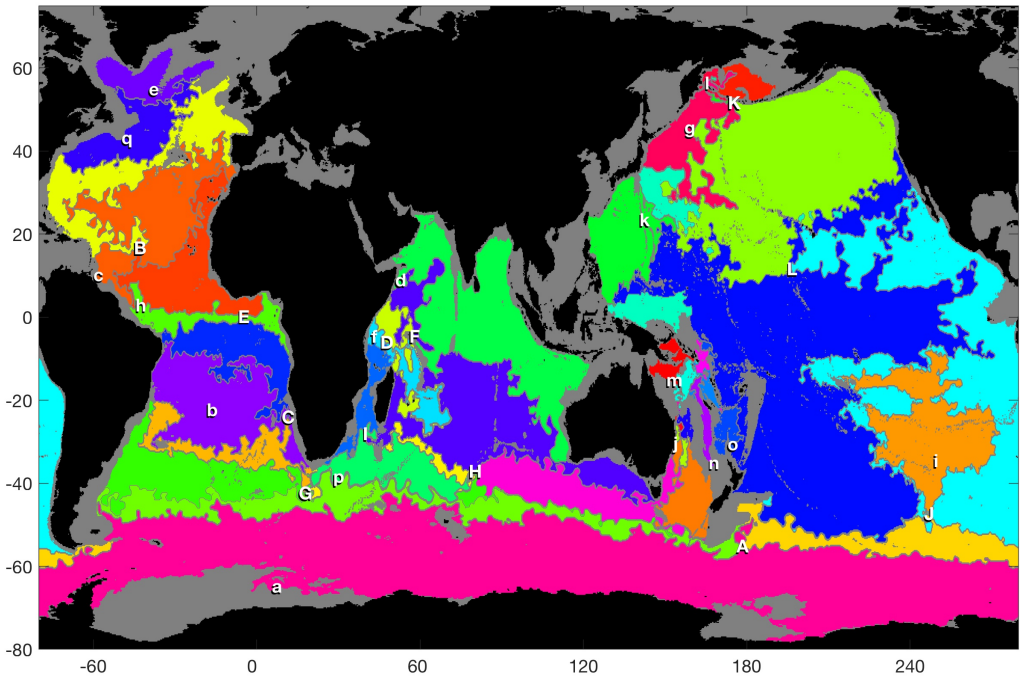


Figure 2.3: The regions associated with each arc of the simplified Reeb graph of the isopycnal $\sigma_2 = 37.0$. Distinct colours are distinct regions, and grey indicates where the isopycnal has outcropped or incropped, or is disconnected from the main ocean. Letters are as in Fig. 2.2, and regions are coloured matching their arcs in Fig. 2.4.

go far down, as the isopycnal is around its deepest in the western South Pacific. Note, there is no node on the simplified Reeb graph corresponding to this extremum of $\underset{\sim}{z}$. The original Reeb graph had this node, but the leaf pruning simplification removed it; as discussed, leaf pruning destroys the property that $\underset{\sim}{z}(X_a) \in [\underset{\sim}{z}(\boldsymbol{x}_{l_a}), \underset{\sim}{z}(\boldsymbol{x}_{h_a})]$, i.e. that the function values at all points associated to an arc lie between the function values at the two nodes that arc is incident upon. However, this is not a problem for forming topobaric surfaces.

Loss of this property also can make for stranger effects, notably the cyan arc between L and an unlabelled saddle leading to k. This cyan arc's associated region is south-east of Japan, *and* north of Papua New Guinea (and even has few small patches just north of Hawaii, but the large cyan region in the eastern Pacific is a different region, and a slightly different colour). As the Reeb graph is supposed to collect connected components of level sets, this is unfortunate. Examining Fig. 2.2, the part south-east of Japan is deeper, the part north of Papua New Guinea shallower. These were (hundreds of) separate arcs in the original Reeb graph. But even this is not a problem for forming topobaric surfaces: though a single branch of $\hat{\pi}$ is fit for the cyan

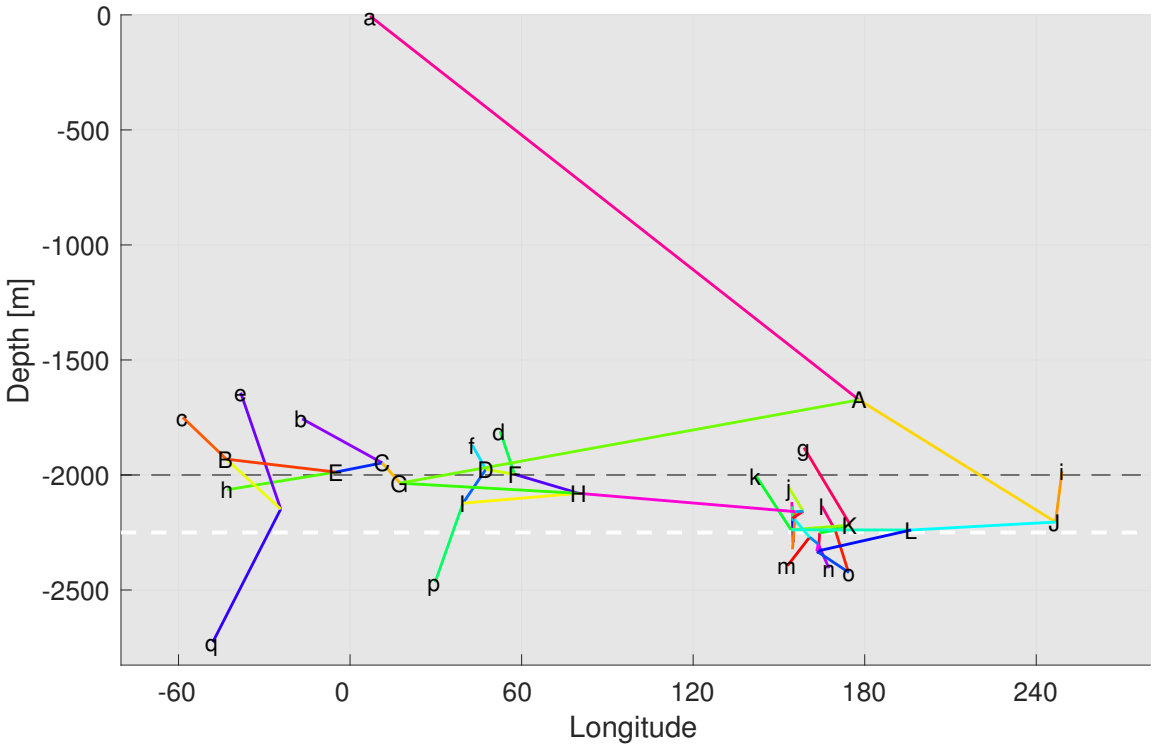


Figure 2.4: A layout of the simplified Reeb graph of the depth $\underset{\sim}{z}$ of the isopycnal $\sigma_2 = 37.0$. Nodes are positioned according to the longitude and depth of their associated (critical) point in physical space, and labelled except where they overlap closely. The black and white level sets of Fig. 2.2 are indicated here as dashed lines. There are two cycles. The first is around Madagascar, given by nodes D, F, H, I, D. The second is around New Zealand, and given by nodes A, J, L, ..., H, G, A.

arc, its deeper parts and shallower parts are not in conflict.

Continuing the tour, from L one could instead travel north of New Zealand, along many small arcs. These small regions remain in the simplified Reeb graph because the leaf pruning algorithm does not prune a leaf from the Reeb graph when it is the last up or the last down arc from its parent saddle node. Passing these regions, one arrives at H, where an arc with a rather filamentary region can be followed to G, and then back to A, closing the cycle around New Zealand. Or, from G, one can go to C and then to E.

Node E is a saddle point near the equator, separating the North Atlantic from the South Atlantic. The equatorial Atlantic, focused upon by McDougall and Jackett (2005a) and de Szoeke and Springer (2005), does appear distinctive in the Reeb graph. From Fig. 2.2 alone, one might glean that there is a saddle of $\underset{\sim}{z}$ somewhere near the equatorial Atlantic, but it is not obviously more important than other saddles of $\underset{\sim}{z}$.

For that, one should study scatter plots of $\underset{\sim}{\rho}$ vs $\underset{\sim}{z}$, as shown in Figure 2.5. The

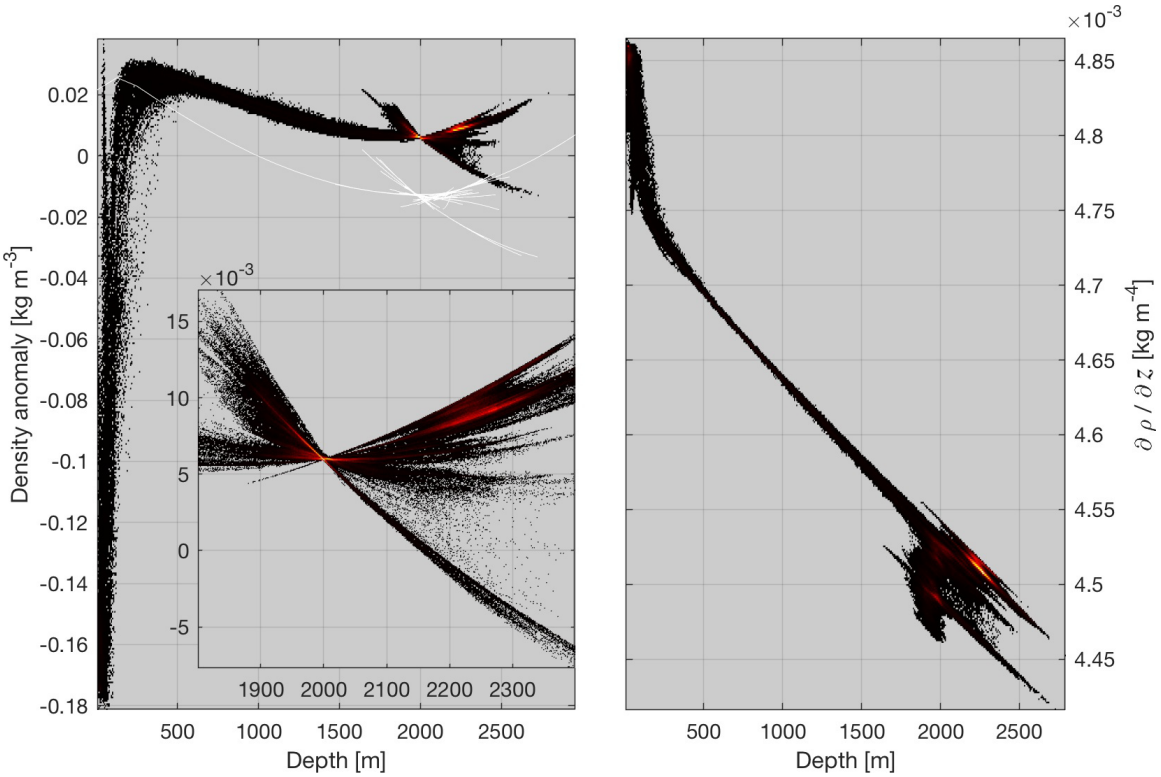


Figure 2.5: Two dimensional histograms of depth $\underset{\sim}{z}$ on the $\sigma_2 = 37.0$ isopycnal, vs (left) the *in-situ* density $\underset{\sim}{\rho}$, or (b) the partial derivative of *in-situ* density with respect to depth, $\underset{\sim}{\rho}_z$. Colour indicates the number of grid points in a 2D bin, going from black (1 data point per bin) through red to white (maximum data points per bin), and grey where 0 data points. In the left panel, a global quadratic fit of $\underset{\sim}{\rho}$ to $\underset{\sim}{z}$ was subtracted from $\underset{\sim}{\rho}$ (otherwise $\underset{\sim}{\rho}$ vs $\underset{\sim}{z}$ is hardly distinguishable from a straight line). A zoomed in region is also shown in the left panel. The branches of $\hat{\rho}$ are plotted in the left panel, offset by $-0.02\,\mathrm{kg\,m^{-3}}$. The branches for $\hat{\pi}$ are rather more messy and not shown. The topobaric surface algorithm finds new depths $\underset{\sim}{z}$ such that $\underset{\sim}{\rho} = \hat{\rho}(\underset{\sim}{z})$. Thus the left panel in a similar figure for the topobaric surface would simply show the white lines shown here. However, the right panel retains scatter, because of non-zero neutral helicity. See text for further discussion.

multivalued nature of $\underset{\sim}{\rho}$ vs $\underset{\sim}{z}$ is evident in the left panel. The population of data centred at 2000 m and angled at 10 o'clock is mostly the Atlantic, whereas the population with $\underset{\sim}{z} > 2000\,\mathrm{m}$ and angled at 2 o'clock is mostly the Pacific; the Indian is somewhat mixed between the two. The source of North Atlantic Deep Water on this isopycnal is seen around 1700 m and $\underset{\sim}{\rho}$ anomaly around $0.02\,\mathrm{kg\,m^{-3}}$, clearly distinct from Southern Ocean waters with lower $\underset{\sim}{\rho}$ anomaly. The entire ocean coalesces in $\underset{\sim}{\rho}$–$\underset{\sim}{z}$ space at 2000 m, simply due to the use of potential density referenced to 2000 m. A fluid parcel on this isopycnal has $\rho(S, \theta, 2000\,\mathrm{m}) = 1037.0\,\mathrm{kg\,m^{-3}}$, regardless of its location. When

it actually is at 2000 m depth, its *in-situ* density is this constant value, globally. For other depths, its *in-situ* density can take other values.

To complete the first iteration, new depths are computed such that (2.13) holds everywhere. Thus, the $\underset{\sim}{\rho}$ vs $\underset{\sim}{z}$ histogram at the start of the second iteration looks exactly like the white curves shown in Fig. 2.5 for this iteration. However, the $\underset{\sim}{\rho}_z$ vs $\underset{\sim}{z}$ histogram does not coalesce in this way, though (2.28) suggests it should for a truly neutral surface. To see why, consider the gradient of density in an approximately neutral surface: it will be akin to (2.9) but with additional gradients of salinity and potential temperature, i.e. add the LHS of (2.7) to the RHS of (2.9)). Technically, these extra gradients prevent the $\underset{\sim}{\rho}$ and $\underset{\sim}{z}$ contours from aligning, so $\underset{\sim}{\rho}$ is not a perfect multivalued function of $\underset{\sim}{z}$. Nonetheless, we are free to force it to be so at every water column. Doing so uses up the one free constraint we have per water column, namely the choice of $\underset{\sim}{z}$. Only in a zero-neutral helicity ocean will this choice of $\underset{\sim}{z}$ also yield $\underset{\sim}{\rho}_z$ a perfect multivalued function of $\underset{\sim}{z}$.

There are two cycles in the Reeb graph for this $\underset{\sim}{z}$: one around Madagascar and one around New Zealand. (This isopycnal has a very narrow passage open between New Zealand and Australia, though hard to see in Fig. 2.2 or Fig. 2.3.) All other islands were filled, except Antarctica. However, Antarctica does not produce a cycle in the Reeb graph, because there exists a circumpolar contour of $\underset{\sim}{z}$. Since the Reeb graph makes equivalent all points in physical space on the same contour (a connected component of a level set), to travel from the Indian, Pacific, or Atlantic oceans to Antarctica one must cross this circumpolar contour of $\underset{\sim}{z}$, but crossing at one place is equivalent to crossing at any other place.

2.4.2 Comparison with other approximately neutral surfaces

The depth of the topobaric surface, after four iterations of the above method, is shown in Fig. 2.6 by way of its difference to the initial depth (Fig. 2.2). Recall that the topobaric surface is pinned to its initial depth at $(180°\text{E}, 0°\text{N})$, and thus the equatorial Pacific doesn't change much. The topobaric surface is raised in the Atlantic and Indian Oceans, and lowered in the Southern Ocean, relative to the original isopycnal. The isopycnal slope in the Southern Ocean is steeper than the slope of the topobaric surface.

How close to neutral is a topobaric surface, and how does it compare with other approximately neutral surfaces? To answer this, we must first discretise the neutral surface relationship (2.9). Suppose we have a neutral surface in a Boussinesq ocean model with a rectilinear grid, indexed by (i, j) for the zonal and meridional

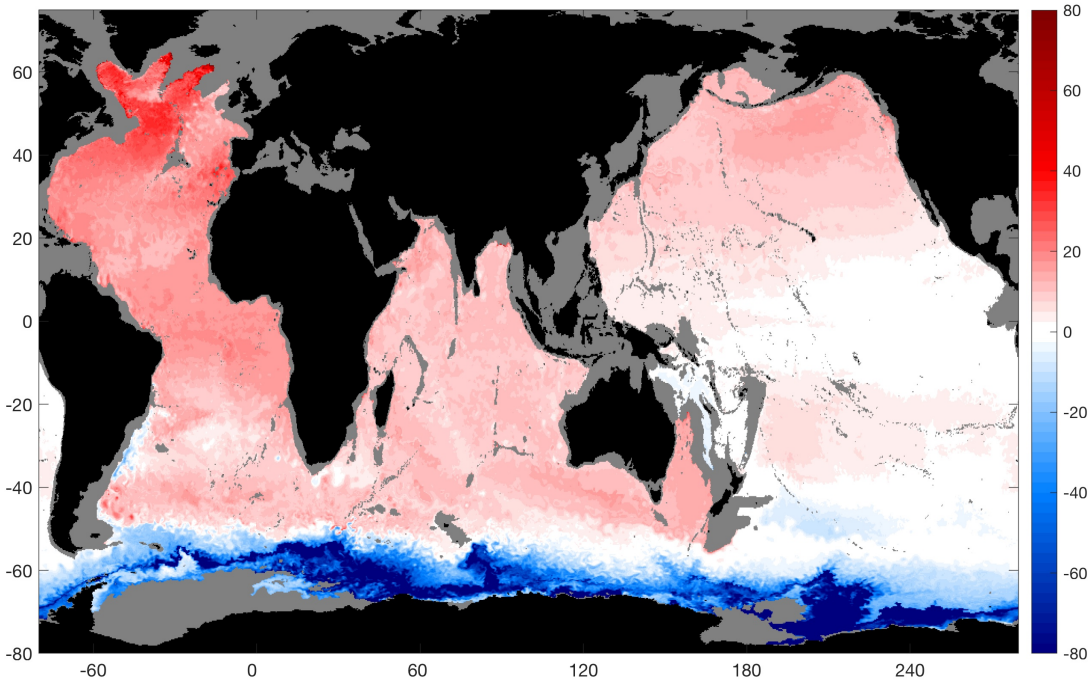


Figure 2.6: The change in depth between the topobaric surface (after four iterations, converged) and the $\sigma_2 = 37.0$ isopycnal from which the former was initiated. Red indicates that the topobaric surface is shallower, blue that it is deeper, and grey that either surface has outcropped or incropped or is not connected to the main ocean.

directions respectively. The salinity, potential temperature, and depth on an approximately neutral surface are $S_{i,j}$, $\theta_{i,j}$, and $z_{i,j}$ at water column i, j (under-tildes are dropped for clarity). The zonal, in-surface density gradient is discretised simply as $(\rho(S_{i,j}, \theta_{i,j}, z_{i,j}) - \rho(S_{i-1,j}, \theta_{i-1,j}, z_{i-1,j})/\Delta x_{i,j}$. Expanding both of these terms by Taylor series about the common midpoint $(\overline{S}^i, \overline{\theta}^i, \overline{z}^i)$, where $\overline{S}^i = (S_{i,j} + S_{i-1,j})/2$, and similarly for θ and z, we find that

$$\epsilon^{(x)} = \frac{\rho(S_{i,j}, \theta_{i,j}, z_{i,j}) - \rho(S_{i-1,j}, \theta_{i-1,j}, z_{i-1,j})}{\Delta x_{i,j}} - \rho_z(\overline{S}^i, \overline{\theta}^i, \overline{z}^i) \, \frac{z_{i,j} - z_{i-1,j}}{\Delta x_{i,j}}, \qquad (2.40)$$

should be almost zero for a neutral surface, for which the terms involving ρ_S and ρ_θ cancel. Specifically, the quadratic terms in the Taylor series cancel identically, so a neutral surface has $\epsilon^{(x)} \sim \mathcal{O}([(z_{i,j} - z_{i-1,j})/\Delta x_{i,j}]^3)$. A similar construction yields the meridional component of the density gradient error for an approximately neutral surface, $\boldsymbol{\epsilon}$.

Figure 2.7 shows $\epsilon^{(x)}$ for four approximately neutral surfaces, as well as its area-weighted l_1 and l_2 norms[19] over the largest connected region. Surface (a) is the

[19] The l_1 norm (mean absolute error) of a vector $\boldsymbol{x}$ with weighting $\boldsymbol{w}$ is $||\boldsymbol{x}||_1 = \sum_i(w_i|x_i|)/\sum_i(w_i)$. Similarly, the l_2 norm (root mean square) is $||\boldsymbol{x}||_2 = \sqrt{\sum_i(w_i x_i^2)/\sum_i(w_i)}$.

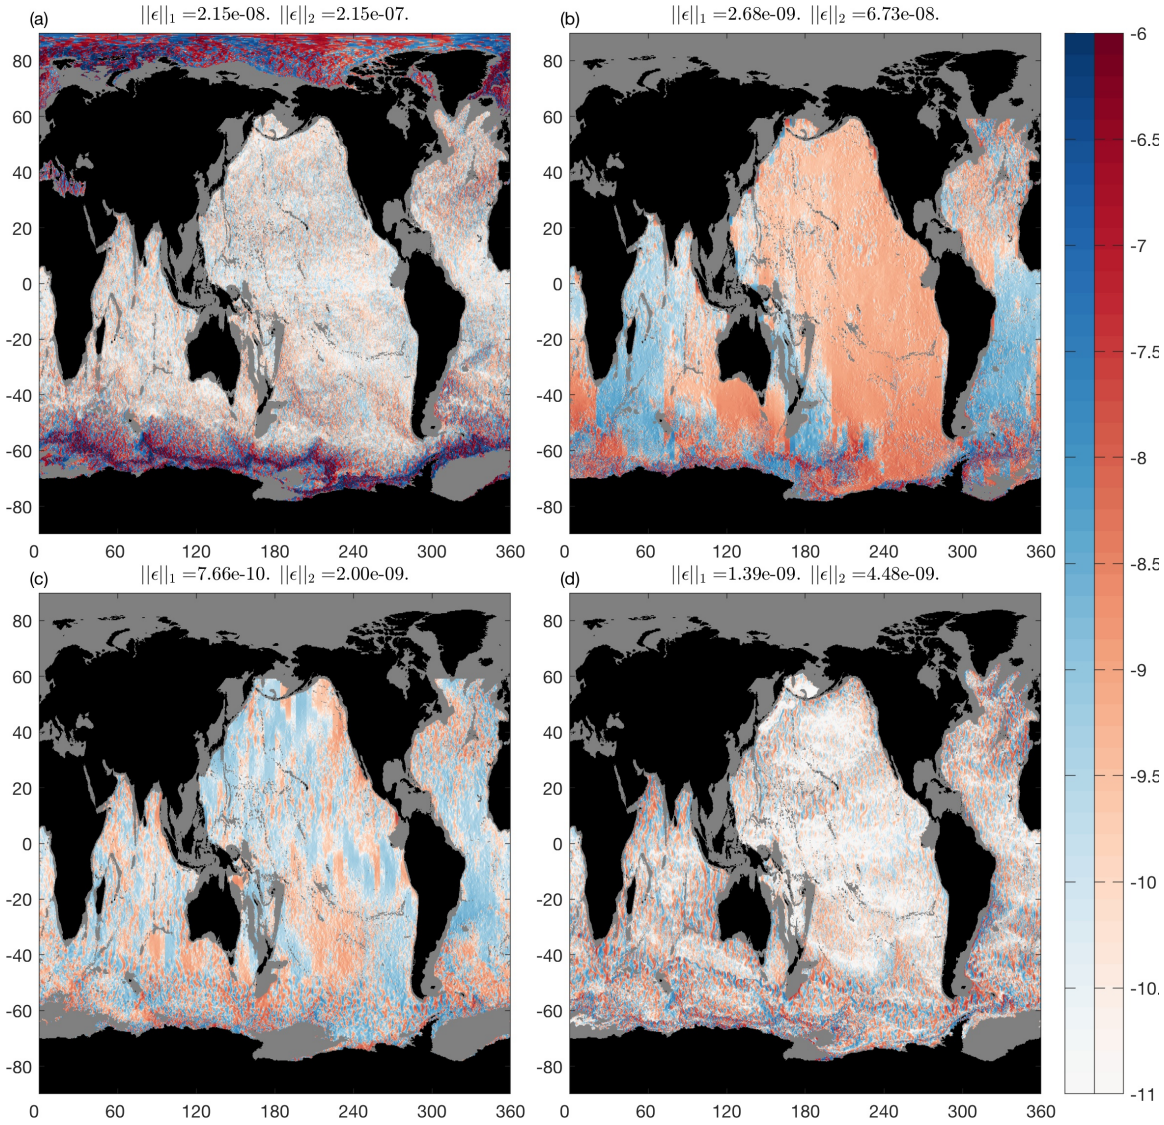


Figure 2.7: Comparison of zonal density gradient error on four approximately neutral surfaces: (a) potential density referenced to 2000 m, (b) neutral density, (c) ω-surface, and (d) topobaric surface. Colour shows $\log_{10}|\epsilon^{(x)}|$, red where $\epsilon^{(x)} > 0$ and blue where $\epsilon^{(x)} < 0$. The area-weighted L_1 and L_2 norms of $\epsilon^{(x)}$ are listed above each panel, where contributions are taken only from the largest connected region (for example, this excludes the Arctic in the case of potential density). Note neutral density is not defined north of 64°N (and the same goes for the ω-surface, since I used the neutral density surface as its initial guess).

original potential density surface from Section 2.3.1, and surface (d) is the topobaric surface generated from it. Surface (b) is the neutral density surface $\gamma^n = 27.9645$, and surface (c) is an ω-surface initiated from a neutral density surface very close to surface (b). All four surfaces intersect at $(180°\text{E}, 0°\text{N}, 2276\,\text{m})$.

For the potential density surface, $|\epsilon^{(x)}|$ is small where the surface is near the reference pressure, but becomes very large elsewhere, particularly in the Southern Ocean where the isopycnal rises to the sea-surface, so on average $|\epsilon^{(x)}|$ is relatively large. For the neutral density surface, $|\epsilon^{(x)}|$ is smaller in the Southern Ocean, but can be uniformly large over vast regions, leading to an average $|\epsilon^{(x)}|$ that is about 1/8 or 1/3 that of the potential density surface, for the l_1 or l_2 norms, respectively. Sharp divisions between $\epsilon^{(x)} > 0$ and $\epsilon^{(x)} < 0$ are found, particularly in the Southern Ocean. These are aligned with the latitude-longitude grid, and are due to the underlying 4° by 4° grid of the WOCE atlas from which neutral density interpolates (Jackett and McDougall, 1997). The ω-surface (Klocker et al., 2009) has $|\epsilon^{(x)}|$ uniformly very small, with l_1 norm about 30% that of the neutral density surface, whereas for the l_2 norm this is only 3%.

This distinction, between 30% and 3%, is because ω-surfaces are constructed to minimize $\boldsymbol{\epsilon}$ in a least-squares sense (they used the equivalent formulation in terms of salinity and potential temperature, not density and depth/pressure). More precisely, Klocker et al. (2009) minimize $||(\boldsymbol{\epsilon})||_2$ without changing its two-dimensional curl: they take an initial approximately neutral surface with error $\boldsymbol{\epsilon}$ and find a scalar potential field Φ' that minimizes $||\boldsymbol{\epsilon} + \nabla\Phi'||_2$, then convert this Φ' into a depth change for the neutral surface, iterating this procedure until it converges. They claim that the two-dimensional curl of $\boldsymbol{\epsilon}$ is approximately proportional to the neutral helicity and inversely proportional to the squared buoyancy frequency, and thus completely determined by the hydrography, hence why they attempt not to change it. To the extent that this approximation is true, there can be no hope to produce an approximately neutral surface with smaller $||\boldsymbol{\epsilon}||_2$.

Nonetheless, the topobaric surface performs admirably, with $||\boldsymbol{\epsilon}||_2$ just over twice as large as that for the ω-surface, whereas for $||\boldsymbol{\epsilon}||_1$ it is less than twice as large. The topobaric surface achieves this closeness to neutrality while still possessing a geostrophic stream function. Moreover, topobaric surfaces can be computed much faster than ω-surfaces: on a typical desktop computer, this topobaric surface is found in under 2 minutes, whereas the ω-surface requires half a day, though this can be reduced to about 30 minutes by vectorizing the ω-surface MATLAB code.

2.5 Conclusions

Given that that typical oceanic epineutral diffusivity is $\mathcal{O}(10^3\,\mathrm{m^2\,s^{-1}})$ while dianeutral diffusivity is $\mathcal{O}(10^{-5}\,\mathrm{m^2\,s^{-1}})$ in the main thermocline (MacKinnon et al., 2013), it is crucial to orient the large epineutral diffusion correctly, lest a component of it act dianeutrally and swamp the real dianeutral diffusion. This can be quantified as a fictitious dianeutral diffusivity (McDougall and Jackett, 2005b; Klocker et al., 2009). The correct orientation is given by the neutral tangent plane, which is equivalent to a surface of locally referenced potential density (McDougall, 1987a). Depth level ocean models have long had rotated diffusion tensors (Veronis, 1975; Redi, 1982), which is relatively straightforward because it is a local problem—the neutral tangent plane is well-defined everywhere, locally. However, layered models require a nearly quasi-conservative density variable for their vertical coordinate, which is a global problem and thus far harder—stitching neutral tangent planes together to form a neutral surface is commonly held to be ill-defined (McDougall and Jackett, 1988), in the sense that the depth of a neutral surface at a given latitude and longitude is not unique. I have shown that this can be true even in a zero neutral helicity ocean with islands. The islands modify the topology of the neutral surface, sometimes creating cycles in the Reeb graph of its pressure, which lead to neutral spirals.

The main topological idea presented here is that specific volume and pressure are multivalued functions of one another on a neutral surface. Of course, any function of the specific volume, such as its inverse the *in-situ* density, is a multivalued function of the pressure (or any function of it). There is an equivalent multivalued functional relation between the salinity and potential/Conservative Temperature on a neutral surface. This has been noticed before, but emphasis was placed on the Northern vs Southern hemispheres having different functional relations between specific volume and pressure (McDougall and Jackett, 2005a; de Szoeke and Springer, 2005), while in reality the ocean is replete with different branches of the multivalued function between specific volume and pressure. More importantly, these branches were argued to arise from extrema in pressure along a neutral surface (McDougall and Jackett, 2005a). This is correct in a two dimensional ocean, but not in three dimensional ocean, where different branches arise from saddle points in the pressure on a neutral surface.

These saddle points, together with the extrema, are the set of critical points of the pressure on a neutral surface. The critical points can be mapped to nodes in a graph—the Reeb graph. Two nodes in the Reeb graph have an arc between them if

there is a monotone path between them that does not pass through a contour of any other critical point. All points on contours that intersect this path are associated with this arc. These points are a region in physical space where the multivalued function is actually single-valued.

Having segmented a neutral surface into regions where specific volume is a single-valued function of pressure, these functions can be fit empirically, subject to some matching conditions determined with further assistance from the Reeb graph. A root-finding problem ensues, solving for the pressure in each water column for which the specific volume at that pressure exactly matches the single-valued function at that pressure. In this way, an iterative procedure turns an isopycnal, or any approximately neutral surface, into a new approximately neutral surface called a topobaric surface. Topobaric surfaces have an exact multivalued functional relation between specific volume and pressure, with correct domains for each branch of this multivalued function.

The name topobaric refers to the topological nature of pressure surfaces, from which the multivalued function for specific volume arises. The similarity with "orthobaric" is intentional. Orthobaric density is a pycnotropic variable—a function only of pressure and specific volume—and as such is guaranteed to possess an exact geostrophic stream function, and an Ertel potential vorticity (de Szoeke, 2000). As such, the *in-situ* density and pressure on an isosurface of orthobaric density are related by an implicit equation: one is a single-valued function of the other. Topobaric surfaces are the topologically correct extension from isosurfaces of simple orthobaric density (de Szoeke et al., 2000), in which the relation between density and pressure on the topobaric surface becomes multivalued. This allows water properties to vary geographically. Note also that orthobaric density does not explicitly take account of the neutral relation, whereas topobaric surfaces begin from it. In future work, I aim to develop a 3D density variable that builds on these topological ideas.

Topological thinking gives us a deeper understanding of why a density variable whose isosurfaces are approximately neutral should be a function of latitude and longitude. Neutral density (Jackett and McDougall, 1997) built this geographic dependence in via reference to a climatology, whereas topobaric surfaces determine the geographic dependence inherently. Density variables that are functions purely of the thermodynamic state, such as thermodynamic neutral density (Tailleux, 2016), lack this geographic dependence; this may help explain why thermodynamic neutral density performs less well in the high polar latitudes.

The exact geostrophic stream function on a neutral surface, whose existence has been known since McDougall (1989), is found in terms of the functional relation

between specific volume and pressure on a neutral surface. This geostrophic stream function is well-defined even when this is a multivalued functional relation. However, the existence of islands causes the geostrophic stream function for a particular neutral surface to become itself a multivalued function of geographic space. Of course this does not affect the flow, and particular branches of the geostrophic stream function can always be chosen so that its gradient is proportional to the geostrophic velocity.

The Lorenz convention sets the depth of a (in this case, approximately neutral) surface that has outcropped or incropped to be infinitesimally below the sea-surface or above the sea-floor, respectively (Young, 2012). In this sense, there are no holes in the surface, and thus no cycles in the Reeb graph: quite a simplification. Is this justifiable? From the topological perspective, two contours of the same level set become one contour that snakes around the hole; points on either side of the hole on the same level set are glued together in topological space. But the fundamental reason that contours of $\underset{\sim}{S}$ and $\underset{\sim}{\theta}$ are aligned is because both are materially advected by the flow, of which there is none through an island. The Lorenz convention makes sense from the perspective of one fluid column, considering also the density of air and solid Earth, but less so from the perspective of many fluid columns. For an extreme example, the Atlantic and Pacific have different water mass properties, it would be wrong to join them. The same applies to sub-marine ridges, and so on to smaller features.

I am currently updating the topobaric algorithm, particularly with regard to the the second matching condition (2.32), and the interested reader should consult a forthcoming paper.

Chapter 3

Ertel potential vorticity versus Bernoulli potential on approximately neutral surfaces in the Southern Ocean

3.1 Introduction

Vorticity-stream function correlations provide a powerful diagnostic tool for studying steady and unsteady geophysical fluid systems, including the vertical propagation of planetary waves, starting with the classic work of Charney and Drazin (1961), free-mode structures (Read et al., 2006; Marshall and So, 1990), and shear stability (Dowling, 1993; Read et al., 2006; Read et al., 2009a; Read et al., 2009b). Often, the quasi-geostrophic framework is employed, which assumes hydrostatic balance, rapid rotation and strong stratification, however, White (1990) showed Earth's atmosphere exhibits function-like correlations with some scatter between Ertel potential vorticity (PV), Q, and mass-weighted stream function, Ψ. In such a primitive-variable framework, which keeps hydrostatic balance but relaxes the assumptions of asymptotically large rotation and stratification, the Bernoulli stream function, B, is related to Q and Ψ in the steady-state limit by $Q = \mathrm{d}B/\mathrm{d}\Psi$ (Gill, 1982, pp. 231–233). Consequently, one may equally well look for correlations in scatter plots of Q versus Ψ or Q versus B. Since B is usually a sum of energy terms in hand, whereas Ψ is typically found by inverting an elliptic operator, I focus on the relationship between Q and B. The above steady-state relationship could then be inverted to obtain Ψ, if desired.

Dowling and Ingersoll (1989) determined $Q(B)$ for Jupiter's Great Red Spot and White Oval BC, and Dowling (1993) determined $Q(\Psi)$ for Jupiter's zonal winds, both in terms of primitive, shallow-water models applied to Voyager wind data. Recently,

Du et al. (2015) and Dowling et al. (2017) carried out extensive studies of $Q(B)$ on potential-temperature surfaces in the extratropical atmospheres of Earth and Mars, respectively. The goal of this chapter is to apply a similar analysis to Earth's Southern Ocean, using the Southern Ocean State Estimate (SOSE, Mazloff et al., 2010).

Homogenization of potential vorticity on a buoyancy surface has long been an attractive idea in the oceanographic (Rhines and Young, 1982) and atmospheric (Sun and Lindzen, 1994) communities. Marshall et al. (1993) first advanced this idea in the Southern Ocean. If PV is homogenized, the slope of Q vs B (each taken from a single buoyancy surface) should be essentially zero. But this slope is a dimensional quantity so this statement is meaningless. It is made meaningful by non-dimensionalising the slope, somehow. Du et al. (2015) used $(NH)^2$ as a first pass at this non-dimensionalisation, where N is the buoyancy frequency and H the pressure scale height, both taken from the definition of the U.S. Standard Atmosphere. For Mars, Dowling et al. (2017) used regional averages of the instantaneous N and H fields. For oceanographers, this non-dimensionalisation presents a considerable theoretical challenge. I will advance it by estimating the local Rossby wave phase speed by fitting the zonal velocity to the meridional PV gradient, and compare this with detailed calculations of the eigenvalue problem for the Rossby wave phase speed (which generalizes the deformation radius eigenvalue problem to include background shear). Thus, on one level this work assesses the extent of PV homogenization in the Southern Ocean.

Inhomogeneous PV on a buoyancy surface is related to instability of the flow. The Eliassen-Palm flux encapsulates both eddy momentum and eddy buoyancy fluxes into a single vector, whose divergence is the eddy potential vorticity flux[1]. Indeed, some eddy parametrisations are based on a fundamental assumption that the eddy PV flux is directed down the mean PV gradient (Green, 1970; Treguier et al., 1997). That this downgradient flux holds in an integrated sense is central to the proof of PV homogenization inside closed geostrophic contours by Rhines and Young (1982, see also the discussion in Section 1.3.3).

Dowling (2014) gives a new perspective on PV inhomogeneity and shear instability, uniting the two branches of shear stability by Arnol'd (McIntyre and Shepherd, 1987) and proposing a shock-wave type physical mechanism. Briefly, and following Dowling (2014), the first branch of Arnol'd's theorem for a quasi-geostrophic system states

[1]Specifically, this is the *poleward* eddy potential vorticity flux, and the averaging is zonal. Using temporal or ensemble averaging yields an analogous, second Eliassen-Palm flux, for the zonal direction (Young, 2012).

the flow is stable (in the sense of Liapunov, meaning perturbations may grow but not unboundedly) if

$$0 < c_1 \leq \frac{\mathrm{d}\psi}{\mathrm{d}q} \leq c_2 < \infty \tag{3.1}$$

everywhere in the domain; the second branch gives stability if

$$0 < \kappa_0^{-2} < c_3 \leq -\frac{\mathrm{d}\psi}{\mathrm{d}q} \leq c_4 < \infty \tag{3.2}$$

everywhere in the domain. Here, c_1, c_2, c_3, and c_4 are arbitrary constants, ψ is the stream function of the background state, q the potential vorticity of the background state, and κ_0^2 is the smallest eigenvalue of the eigenvalue problem relating PV perturbations to flow perturbations, i.e. $q' = -\kappa_0^2\psi'$, where primes indicate deviations from the background state. The quasi-geostrophic *intrinsic* Rossby wave phase speed for the gravest mode is then (Dowling, 2014)

$$\hat{c} = -\frac{1}{\kappa_0^2}\frac{\partial q}{\partial y}, \tag{3.3}$$

while the total phase speed is $c = U + \hat{c}$. Crucially, the intrinsic Rossby wave phase speed is proportional to the background PV gradient. With this, the Mach number[2] for Rossby waves can be defined,

$$\text{“}Ma\text{”} = -\frac{u}{\hat{c}}, \tag{3.4}$$

Noting that $\partial\psi/\partial y = -u$, then

$$\text{“}Ma\text{”} = -\kappa_0^2\frac{\mathrm{d}\psi}{\mathrm{d}q}. \tag{3.5}$$

Using (3.5) in (3.1) and (3.2) gives

$$-\infty < -c_2\kappa_0^2 \leq \text{“}Ma\text{”} \leq -c_1\kappa_0^2 < 0 \tag{3.6}$$

and

$$1 < c_3\kappa_0^2 \leq \text{“}Ma\text{”} \leq c_4\kappa_0^2 < \infty. \tag{3.7}$$

Thus, Arnol'd's first branch of shear stability is associated with negative Mach numbers, and the second branch with supersonic Mach numbers. Because sound waves are omni-directional, the concept of a negative Mach number for sound waves does not exist: the wave speed is always taken to be upstream. But Rossby waves are

[2]Successful branding by Gillette played a not-insignificant role in the decision to name this after Mach, not Froude (Dowling, 2014).

uni-directional: their intrinsic phase speed always keeps high PV to the right, as in (3.3).

It turns out that subsonic flow, $0 <$ "Ma" < 1, is potentially unstable. In this situation, Rossby waves have an intrinsic phase speed faster than the background advection, and moreover that they are "swimming" upstream. This allows the possibility for Rossby waves to "lean into the shear', to interact and amplify each other. In the supersonic case, Rossby waves are again trying to "swim" upstream, but the background flow is stronger than the intrinsic phase speed; Rossby waves are torn apart from one another, so they cannot interact long enough to grow. Likewise, when the Mach number is negative, Rossby waves are "swimming" downstream: they rip themselves apart from one another. For further discussion and schematics, see Dowling (2014). If one considers the inverse of the Mach number, the two stability branches become united within "Ma"$^{-1} < 1$, and instability is possible for "Ma"$^{-1} \geq 1$. This line of argument suggests that the physical mechanism at the onset of shear turbulence is a shock-wave, as Rossby waves drop from supersonic to subsonic (Dowling, 2014).

Seawater is a multi-component fluid with a nonlinear equation of state, which presents a new challenge not found in the previous atmospheric studies. There does not exist a density variable for the ocean that is analogous to potential temperature in a dry atmosphere. That is, it is impossible to define a density variable along whose isosurfaces fluid parcels can be exchanged without experiencing a buoyant restoring force, which is the characteristic of a so-called "neutral surface" (McDougall, 1987a). Fortunately, serviceable approximations to such a density surface do exist (particularly when restricting attention to a local region) enabling us to compare oceanic and atmospheric zonal jets directly from the point of view of non-dimensional, primitive vorticity-stream function correlations.

The chapter is laid out as follows. Section 3.2 gives further background on the relationship between potential vorticity and Bernoulli potential, its non-dimensionalisation into a Mach number, the Rossby wave phase speed, triad relations bringing the mass-weighted stream function into the mix, and orthobaric density. Section 3.3 studies the relationship between Q and B over the whole ACC, and Section 3.4 takes a more local look at the Q vs B relationship. Section 3.5 furthers the discussion and concludes.

3.2 Theory

3.2.1 Mathematical stage-setting

Suppose we have a quasi-conservative buoyancy variable b, meaning that its material rate of change, $\mathrm{Db/Dt} = \varpi$, is small and due only to similar non-conservation of salinity and potential temperature. If b is monotone with depth, it can be used as the vertical coordinate, whence the hydrostatic Boussinesq equations may be written without approximation as (Young, 2012)

$$u_{\tilde{t}} + \varpi u_{\tilde{b}} - \sigma Q v + B_{\tilde{x}} = F^x, \tag{3.8a}$$

$$v_{\tilde{t}} + \varpi v_{\tilde{b}} + \sigma Q u + B_{\tilde{y}} = F^y, \tag{3.8b}$$

$$\xi + m_{\tilde{b}} = 0, \tag{3.8c}$$

$$\sigma_{\tilde{t}} + (\sigma u)_{\tilde{x}} + (\sigma v)_{\tilde{y}} + (\sigma\varpi)_{\tilde{b}} = 0. \tag{3.8d}$$

In sequence, these are the zonal momentum equation (u the eastward velocity), the meridional momentum equation (v the northward velocity), hydrostatic balance, and the thickness equation. The velocities u, v are the same regardless of whether z or b is the vertical coordinate. The depth of the buoyancy surface, ξ, is an independent variable. The thickness is $\sigma = \xi_{\tilde{b}} = 1/b_z$, and m is the Montgomery potential. Subscripts represent partial derivatives, and the tildes on $\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t}$ remind us that these partial derivatives are taken at constant b. The potential vorticity is

$$Q = \frac{f + v_{\tilde{x}} - u_{\tilde{y}}}{\sigma}, \tag{3.9}$$

with f the Coriolis parameter. The Bernoulli potential is $B = m + K$, and $K = (u^2 + v^2)/2$ the specific kinetic energy. External forcing and viscosity are represented by (F^x, F^y). The velocity through b surfaces is equivalent to ϖ.

There are two key requirements for the surfaces upon which we'll analyse Q and B. First, the buoyancy b should be as close to conservative as possible, i.e. ϖ as small as possible. Second, the geostrophic stream function m must exist, i.e. there must be a solution m to the coupled equations $m_{\tilde{x}} = p_x/\rho_c$ and $m_{\tilde{y}} = p_y/\rho_c$.

Suppose there is a steady ($\partial_{\tilde{t}} = 0$) background flow that that solves the inviscid and unforced ($F^x = 0 = F^y$) and buoyancy conserving ($\varpi = 0$) version of (3.8). Then (3.8d) reduces to $(\sigma u)_{\tilde{x}} + (\sigma v)_{\tilde{y}} = 0$, and thus there exists a stream function Ψ such that

$$\Psi_{\tilde{x}} = \quad \sigma v \tag{3.10a}$$

$$\Psi_{\tilde{y}} = -\sigma u, \tag{3.10b}$$

and moreover, (3.8a) and (3.8b) reduce to

$$B_{\tilde{x}} = Q\Psi_{\tilde{x}} \tag{3.11a}$$

$$B_{\tilde{y}} = Q\Psi_{\tilde{y}}. \tag{3.11b}$$

This equation is of the form (1.8), thus there must be a (multivalued) functional relationship between Ψ and B. Cross-differentiating (3.11) shows the same applies between Ψ and Q. Dividing (3.11) by Q then cross-differentiating shows the same applies between Q and B.

This chapter, however, will not employ the Reeb graph to study the multivalued nature of the functional relationship between Q and B: though eddies can cause some contours of B to close, the level set of B for values in the ACC typically is dominated by one circumpolar contour (similarly for Q). In Chapter 2, the Reeb graph was used because significant differences in water properties exist between the main ocean basins, whereas this is likely less significant for the ACC.

In some cases however, such topological questions are imperative. Dowling and Ingersoll (1989) model Jupiter's Great Red Spot as a single layer of shallow water overlying meridionally varying bottom topography, which is equivalent to a two layer shallow water model in which the bottom layer motions are zonal and steady. They model the PV as a multivalued function of the Bernoulli stream function, and model the bottom topography as quartic polynomial of latitude. Then, using *Voyager* cloud-top wind data to compute the Bernoulli stream function and the vorticity, they simultaneously fit both models to the data, thereby deriving the bottom topography (equivalently, the deep layer's zonal winds). The use of a multivalued relation between PV and Bernoulli was critical, as the PV vs Bernoulli stream function relation is markedly different to the north and south of the Great Red Spot.

Of course, atmospheres and oceans are not steady and are forced, so we do not expect an exact functional relationship between the instantaneous Q and B. Nonetheless, scatter plots of Q vs B do reveal a strong correlation, suggestive of an underlying functional relationship overlain with some scatter. This functional relationship can tell us about shear stability and PV homogenization, if we also know something about the Rossby wave phase speed.

3.2.2 Rossby Wave Phase Speed

In a single layer quasi-geostrophic system with deformation radius L_d and background zonal flow u, the intrinsic Rossby wave zonal phase speed is simply

$$\hat{c} = -q_y/(k^2 + l^2 + L_d^{-2}), \tag{3.12}$$

where k and l are the zonal and meridional wavenumbers respectively, and the background meridional QGPV gradient is $q_y = \beta + uL_d^{-2}$ (Vallis, 2006). In the long wave limit, this becomes $\hat{c} = -q_y L_d^2$, of the form (3.3). This simple form makes the Mach number method particularly tractable for these systems, whence it was first applied to Jupiter (Dowling, 1995) and Saturn (Read et al., 2009b). Moreover, for these giant systems, where the Rossby deformation radius is finite and considerably smaller than the domain, the long-wave limit is essentially attainable, and in this limit the phase speed of Rossby waves equals the group speed, and thus they are non-dispersive. Physical shocks must be related to the group speed rather than the phase speed, for it is the former that propagates information. Thus our definition of the Mach number , (3.4), should likely be given using the group velocity, but in the long-wave limit, the phase speed suffices.

In the full 3D continuously stratified fluid, determination of the Rossby wave phase speed is considerably harder. As discussed in the introduction, Du et al. (2015) and Dowling et al. (2017) used the deformation radius $L_d = NH/f$, which is assigned to κ_0^{-1} in (3.3). This substitution, however, is not well-justified, because (3.3) is only valid for a single layer fluid (not to mention that the formula for the deformation radius is derived for a fluid at rest). Nonetheless, this rough scaling allowed Du et al. (2015) to clearly show that homogeneous PV does not fit the data in the troposphere or midlatitude stratosphere, though the mesosphere and the summer polar stratosphere do exhibit homogeneous PV; it allowed Dowling et al. (2017) to find PV on Mars is everywhere inhomogeneous in three seasons, but nearly homogeneous in southern summer in the southern hemisphere and the lower half of the northern hemisphere. A rough scaling such as this is less useful in locating the threshold to instability at $\text{“}Ma\text{”}^{-1} = 1$, and thus I turn to new methods.

The Rossby wave equation is a partial differential equation that is typically dealt with by separation of variables, yielding (from the vertical dimension) an ordinary differential equation that, with boundary conditions, is an eigenvalue problem for the Rossby wave phase speed and vertical modes. These are explored in Appendix 3.A, both in depth coordinates (a review) and in buoyancy coordinates (new material). From buoyancy coordinates some particularly clean results emerge. The intrinsic Rossby wave phase speed is found (as an eigenvalue problem) in (3.38) taking exactly the form of (3.3), but it becomes possible for waves to keep high PV to their left (in some layers—not in the whole fluid column). Also, with a little manipulation of this equation, the Mach number arises naturally, and poses the eigenvalue problem in a way to better capture the background state—see (3.43).

Still, the modal analysis is poorly suited for this Mach number analysis, for two reasons. First, there is ambiguity about the bottom boundary condition—essentially whether zero horizontal flow or zero normal flow should apply. The choice does affect the results significantly. So too does the formulation in buoyancy or depth coordinates. Second, which vertical mode does one select? One might select the most unstable mode, and where all modes are stable, choose the fastest mode (that keeps high PV to its right). The former has phase speed in the domain of zonal velocities in the water column, whereas the latter must have phase speed less than the minimum (i.e. most westward) velocity in the water column. The resulting phase speed field would therefore be discontinuous. Thus, a new method is needed.

Consider shifting to a slightly different reference frame, one rotating with the same axis as Earth but with an angular velocity shifted by α. That is, this observer has a period of $2\pi/(\Omega + \alpha)$. When $\alpha > 0$, this observer's rotation is faster than Earth's, and things at rest as seen from Earth (e.g. topography) appear to move westward. If the velocity measured on Earth is $\boldsymbol{v}$, this observer measures it as $\boldsymbol{v}^{(\alpha)} = \boldsymbol{v} + u^{(\alpha)}\hat{\boldsymbol{i}}$, where $u^{(\alpha)} = -\alpha r \cos\theta$, with θ the latitude and r the radial distance from Earth's centre.

In the new frame, the (Boussinesq, to use B) momentum equation is, in vector-invariant form,

$$\boldsymbol{v}^{(\alpha)}_{t^{(\alpha)}} + \boldsymbol{\omega}_a \times \boldsymbol{v}^{(\alpha)} + \nabla B^{(\alpha)} = \boldsymbol{F}. \tag{3.13}$$

Here, the Eulerian rate of change of $\boldsymbol{v}^{(\alpha)}$ in the new frame is $\boldsymbol{v}^{(\alpha)}_{t^{(\alpha)}} = \boldsymbol{v}_t + \alpha \boldsymbol{v}_\lambda$, using the fact that $u^{(\alpha)}\hat{\boldsymbol{i}}$ is independent of time t and longitude λ. The total vorticity is α-invariant, meaning it is invariant to changes in the rotation rate. Mathematically, the total vorticity is $\boldsymbol{\omega}_a = 2\boldsymbol{\Omega} + \nabla \times \boldsymbol{v} = 2(\boldsymbol{\Omega} + \boldsymbol{\alpha}) + \nabla \times \boldsymbol{v}^{(\alpha)}$, where $\boldsymbol{\Omega} = \Omega(\cos\theta\hat{\boldsymbol{j}} + \sin\theta\hat{\boldsymbol{k}})$ and similarly for $\boldsymbol{\alpha}$. To check this, confirm that $2\alpha(\cos\theta\hat{\boldsymbol{j}} + \sin\theta\hat{\boldsymbol{k}}) + \nabla \times (u^{(\alpha)}\hat{\boldsymbol{i}}) = \boldsymbol{0}$ in spherical coordinates. The Bernoulli potential in the new frame is $B^{(\alpha)}$, to be discovered. Forcing is encapsulated in $\boldsymbol{F}$, the same as in Earth's reference frame.

When deriving the primitive equations in Earth's reference frame from the Navier-Stokes equations, the centrifugal and gravitational potentials are combined into the geopotential, which motivates a new geometry, the geoid. Thereafter, this non-sphericity is usually ignored (Gates, 2004; Vallis, 2006). Having shifted from an inertial frame to Earth's frame, we now shift slightly further by α, but we want to keep the same buoyancy surfaces, so we do *not* redefine the geopotentials. Thus, the centrifugal acceleration in the new frame will have a horizontal component. In Earth's frame, the centrifugal acceleration appears in B as $-\Omega^2(r\cos\theta)^2/2$. So, in the new frame, the excess centrifugal acceleration appears in $B^{(\alpha)}$ as $[-(\Omega + \alpha)^2 + \Omega^2](r\cos\theta)^2/2$.

Similarly, the excess specific kinetic energy is $[\boldsymbol{v}^{(\alpha)} \cdot \boldsymbol{v}^{(\alpha)} - \boldsymbol{v} \cdot \boldsymbol{v}]/2$. Adding both these to B, the quadratic α terms cancel, leaving

$$B^{(\alpha)} = B - \alpha A, \tag{3.14}$$

where

$$A = \left(\Omega + \frac{u}{r \cos\theta}\right) r^2 \cos^2\theta \tag{3.15}$$

is the specific axial angular momentum. In spherical coordinates, one can directly verify that $\alpha \boldsymbol{v}_\lambda + \boldsymbol{\omega}_a \times u^{(\alpha)} \hat{\boldsymbol{\imath}} - \alpha \nabla A = \boldsymbol{0}$, thus reducing (3.13) to the momentum equation in Earth's frame. Note that contours of $B^{(\alpha)}$ are smoothly deformed from B contours when $\alpha = 0$, to latitude circles (the Ω term dominates in A) when $\alpha \to \pm\infty$.

Next, (3.13) is transformed as for the (Boussinesq) primitive equations. This includes the thin-shell approximation—thus we swap r for Earth's radius R in A—the traditional approximation, and hydrostatic balance. We define the buoyancy in the new frame as that in Earth's frame (the excess centrifugal force is ignored there), so that the buoyancy equation is identical between frames, as is the PV. The latter is because PV is the dot product of the total vorticity and the buoyancy gradient, both α-invariant, and the dot product represents a (coordinate-independent) geometric relationship between vectors. Finally, the continuity equation is also unchanged, because $u^{(\alpha)}$ is independent of longitude. All told, the new frame's equations of motion are identical to those in Earth's frame (3.8), except $\boldsymbol{v}$ is replaced by $\boldsymbol{v}^{(\alpha)}$ and B by $B^{(\alpha)}$, and coordinates are decorated by (α) supercripts.

Now, we can do as in Section 3.2.1: assuming a steady state, zero mechanical and buoyancy forcing, there arises a functional relationship between Q and $B^{(\alpha)}$. The α-variance of B gives us room to manipulate its relationship with Q. What is the meaning of this, and what is the meaning of α?

The further from stationary the system (3.8), the further Q and B become from a perfect functional relationship. Thus, choosing α to minimize the error in the functional relationship between Q and $B^{(\alpha)}$ is equivalent to selecting the reference frame in which the flow is closest to stationary, i.e. $u + u^{(\alpha)} \approx 0$. Equivalently, in Earth's frame anomalies tend to move zonally with speed $u \approx -u^{(\alpha)}$, and therefore we may interpret $-u^{(\alpha)} = \alpha R \cos\theta$ as the phase speed of long Rossby waves on the local surface (not any particular vertical mode). The intrinsic phase speed of these long Rossby waves is

$$\hat{c} = -(u + u^{(\alpha)}). \tag{3.16}$$

3.2.3 Mach number

Having studied the Rossby wave phase speed, we are now in a position to estimate the Mach number. Keeping a general form for now, suppose the intrinsic zonal Rossby wave phase speed is given by

$$\hat{c} = -\lambda \frac{\partial Q}{\partial \tilde{y}} \tag{3.17}$$

equivalent to (3.3) if $\lambda = \sigma_0 \kappa_0^{-2}$, where σ_0 is some mean layer thickness relating the primitive and QG PV gradients, $\sigma_0 Q_{\tilde{y}} \approx q_{\tilde{y}}$, where, for reminder, Q is the Ertel PV and q is the quasi-geostrophic PV. If we just knew λ, we could calculate the inverse Mach number, using the definition (3.4), as

$$\text{“}Ma\text{”}^{-1} = \frac{-\lambda Q_{\tilde{y}}}{(f+\zeta)^{-1} B_{\tilde{y}}} = -\frac{\mathrm{d}Q}{\mathrm{d}B}(f+\zeta)\lambda. \tag{3.18}$$

As I will show, the functional relationship between Q and B is close to linear over a surprisingly vast region of the Southern Ocean. Thus, we can fit a linear relation

$$Q = q_0 + q_1 B, \tag{3.19}$$

to get $\mathrm{d}Q/\mathrm{d}B = q_1$. For a well-chosen α, Q vs $B^{(\alpha)}$ is even closer to linear. Thus, we simultaneously fit a multilinear relation between Q and B and A,

$$Q = a_0 + a_1 B + a_2 A. \tag{3.20}$$

Now, taking

$$\alpha = -\frac{a_2}{a_1}, \tag{3.21}$$

transforms (3.20) to become $Q = a_0 + a_1 B^{(\alpha)}$, and hence $\mathrm{d}Q/\mathrm{d}B^{(\alpha)} = a_1$. Then, the inverse Mach number has the simple form

$$\text{“}Ma\text{”}^{-1} = -\frac{\hat{c}}{u} = \frac{u + u^{(\alpha)}}{u} = \frac{B^{(\alpha)}_{\tilde{y}}}{B_{\tilde{y}}} = \frac{\mathrm{d}B^{(\alpha)}/\mathrm{d}Q}{\mathrm{d}B/\mathrm{d}Q} = \frac{q_1}{a_1}. \tag{3.22}$$

The inverse Mach number is the ratio of the slope of the Q vs B relation in Earth's reference frame to the slope of the Q vs B relation in the frame moving with long Rossby waves.

If one desires λ, simply take the meridional derivative of Q to find $Q_{\tilde{y}} = a_1 B^{(\alpha)}_{\tilde{y}}$, recognize $B^{(\alpha)}_{\tilde{y}} = -(f+\zeta)\hat{c}$ using (3.16), then use (3.17) to obtain

$$\lambda = -[a_1(f+\zeta)]^{-1}. \tag{3.23}$$

Using this and (3.19) in (3.18) verifies (3.22).

This is susceptible to L'Hôpital's trap, where both terms in the ratio are close to zero. The chain rule overcomes this: returning q_1 and a_1 to $\mathrm{d}Q/\mathrm{d}B$ and $\mathrm{d}Q/\mathrm{d}B^{(\alpha)}$ respectively, we find

$$\text{“}Ma\text{”}^{-1} = \frac{\mathrm{d}B^{(\alpha)}}{\mathrm{d}B} = 1 - \alpha\frac{\mathrm{d}A}{\mathrm{d}B} = 1 + \frac{a_2}{a_1}A_1, \tag{3.24}$$

where we now must fit

$$A = A_0 + A_1 B \tag{3.25}$$

to obtain $\mathrm{d}A/\mathrm{d}B = A_1$. Note in this form, the fit (3.19) is unnecessary for this particular task. Admittedly (3.24) still contains a ratio in α. When working with the entire ACC (Section 3.3) both forms (3.22) and (3.24) give similar results. However, when working with local data (Section 3.4), anything can happen, and tests reveal the second form, (3.24), generally behaves better. Thus I use (3.24) henceforth.

There are three interesting limits, as follows.

1. When Q is homogeneous, $\text{“}Ma\text{”}^{-1} = 0$.
   With $Q = Q_0$ a constant, (3.20) becomes $(Q - a_0)/a_2 - (a_1/a_2)B = A$, and hence $A_1 = -a_1/a_2 = 1/\alpha$, so $\text{“}Ma\text{”}^{-1} = 0$.

2. When Q is a function only of B, $\text{“}Ma\text{”}^{-1} = 1$.
   This case gives $a_1 \neq 0$ while $a_2 = 0$, so $\alpha = 0$ and $\text{“}Ma\text{”}^{-1} = 1$.

3. When Q is a function only of A, $\text{“}Ma\text{”}^{-1} \to \pm\infty$.
   This case gives $a_1 = 0$ while $a_2 \neq 0$, so $\alpha \to \pm\infty$, and $\text{“}Ma\text{”}^{-1} \to \pm\infty$.

Thinking physically, as Q becomes uniform, the restoring force for Rossby waves goes to zero; thus their phase speed goes to zero, so the Mach number goes to infinity and the inverse Mach number to zero.

If B contours are not perfectly zonal, then a zonal Rossby wave propagates Q anomalies zonally, off the B contour, supporting the dependence of Q with latitude and diminishing its correlation with B. When $\alpha = 0$, Rossby waves are stationary, so Q should vary only with B. This predilection with zonality is because in the α-shifted reference frame, the planet spins faster or slower but its axis of rotation is unchanged.

Since A contours are basically latitude circles (the Ω term dominates), the third case harkens back to the notion of comparing the meridional PV gradient to β divided by a mean layer thickness.

These reference frame considerations have important practical applications. Arnol'd's stability criteria, (3.1) and (3.2), are coordinate-frame dependent, because the stream

function is coordinate dependent. If a reference frame is found that satisfies either stability criterion, then the flow is stable. The Mach number can be easily updated to allow this further freedom, by defining "Ma" $= -(u - \alpha' R\cos\theta)/\hat{c}$, for some other reference frame shift α' (Dowling, 2014). However, I will not pursue this further, for two reasons. First, I am not attempting to actually apply Arnol'd's stability theorems, which require the PV vs stream function relation to meet the criterion globally, but the ocean is sufficiently turbulent that there will be different patches where one or the other or neither will hold[3]. Second, Earth's reference frame is arguably more important than others: long Rossby waves, being deeply rooted disturbances, must strongly interact with the complex bathymetry of the Southern Ocean.

Read et al. (2009b) used similar reasoning to estimate Saturn's rotation period. They posited that Saturn, like Jupiter, sits on the threshold of instability, "Ma"$^{-1} = 1$ if and only if $\alpha = 0$ in Saturn's reference frame. By linearly regressing the zonal velocity to the meridional PV gradient, $u = -u^{(\alpha)} + L_d^2 q_y$, they obtained the total Rossby wave phase speed $-u^{(\alpha)} = \alpha R\cos\theta$, where here the radius R varies with latitude to account for Saturn's elliptic geometry, and L_d^2 playing the role of κ_0^{-2} in (3.3). They found that α was generally positive (eastward Rossby waves) and of a similar magnitude for all latitudes away from the equator. Redefining Saturn's angular rotation frequency by subtracting α (faster rotation period) brings Rossby waves to rest in the new reference frame. Thus, Read et al. (2009b) posited that Saturn was globally marginally stable, then adjusted Saturn's rotation rate to suit. I am accepting that Earth's rotation rate is accurately known, and allowing the flow to deviate from marginal stability.

As this method empirically fits Q to B (and A) using only data from the local approximately neutral surface, it is implicitly producing an estimate of the Rossby wave phase speed for the local surface only. If one thinks of vertical modes of Rossby waves, each mode moves, by differing amounts, Q anomalies off the B contours of this particular surface. But remember, the goal of this local empirical method was to get away from the difficulties associated with the vertical modes, and in particular, the question of which vertical mode is relevant. In fact all modes are relevant, to varying degrees, and this local empirical estimate of the Rossby wave phase speed can be

[3] One might wonder whether α can be chosen arbitrarily large, so that $|\mathrm{d}\Psi/\mathrm{d}q| > \kappa_0^{-2}$ and one of Arnol'd's stability branches is satisfied. The problem is that such a large α modifies the horizontal momentum equations such that the tendencies in the shifted reference frame are no longer small (provided the state is not perfectly zonal), thus breaking the functional relationship between the PV and the stream function.

thought of as empirically finding what weighted sum of the modal Rossby wave phase speeds is relevant to the local surface.

3.2.4 Triad relations between Q, B, and Ψ

The functional relation between Q and B, in this case (3.19), determines a triad of relations between Q, B, and Ψ. These are obtained using $Q = \mathrm{d}B/\mathrm{d}\Psi$, which follows from (3.11). Non-dimensionalising (3.19) with a characteristic value Q_0, and defining $\mu_0 = -q_1/Q_0$ and $B_0 = (q_0 - Q_0)/(\mu_0 Q_0)$, the triad relations are (Dowling et al., 2017)

$$Q/Q_0 = 1 - \mu_0(B - B_0), \tag{3.26a}$$

$$Q/Q_0 = \exp[-\mu_0 Q_0(\Psi - \Psi_0)], \tag{3.26b}$$

$$\Psi - \Psi_0 = -\frac{1}{\mu_0 Q_0}\log[1 - \mu_0(B - B_0)]. \tag{3.26c}$$

Dowling et al. (2017) chose B_0 as the area-weighted mean B, and $Q_0 = q_0 + q_1 B_0$. For the intrinsic Rossby wave phase speed, they assumed $\lambda = \sigma_0 L_d^2$ by extending the QG case, (3.3), where L_d is the first baroclinic deformation radius. This leads to "Ma" $= (\mu_0 Q_0 f_0 \sigma_0 L_d^2)^{-1}$, having replaced $(f + \zeta)$ with f_0. Finally, they non-dimensionalised μ_0 by multiplying it by $\overline{(NH)^2}$ in their notation, which is $L_d^2 f_0^2 = c_1^2$ the squared first-baroclinic gravity wave phase speed. This implicitly assumes that $Q_0 = f_0/\sigma_0$, so Q_0 is not free to be chosen, and actually Q_0 should determine B_0. I tested both methods on $\sigma_\nu = 27.1$ in the ACC (having computed L_d as in Chelton et al. (1998)). The correct choice, $Q_0 = f_0/\sigma_0$, increased "Ma"$^{-1}$ on the $\sigma_\nu = 27.1$ orthobaric layer by over 50% in SOSE's final two years, whereas the difference was slight in the first four years. If future work chooses λ as above by analogy to QG, it should use $Q_0 = f_0/\sigma_0$ (or non-dimensionalise by the less-attractive $Q_0 f_0 \sigma_0 L_d^2$). Du et al. (2015) also chose to define Q_0 based on B_0, the area-weighted B. Du et al. (2015) and Dowling et al. (2017) made rough estimates of the Mach number, not-least to distinguish homogeneous from inhomogeneous PV. I am attempting to refine these estimates with more careful analysis.

We now have all pieces in place to begin studying Q vs B on orthobaric density surfaces.

3.2.5 Orthobaric density

The nonlinear nature of the equation of state for seawater confounds many a simple analysis. Oceanographers have crafted many candidate density variables to use

analogously to potential temperature in (dry) atmospheric studies. Potential density, $\rho(S, \theta, p_{\text{ref}})$ where p_{ref} is a constant reference pressure, is a common choice but criterion (2) from Section 2.1 fails: there is no geostrophic stream function in a potential density surface (McDougall, 1989). Specific volume anomaly, $1/\rho(S, \theta, p) - 1/\rho(S_0, \theta_0, p)$ where S_0 is a reference salinity and θ_0 a reference potential temperature, possesses an exact geostrophic stream function (Montgomery, 1937), but can perform poorly at criterion (1) from Section 2.1, namely that of material conservation, in global analyses where the salinity and potential temperature depart far from their reference values. Nowadays neutral density (Jackett and McDougall, 1997) is very commonly used, but it too fails criterion (2). We employ orthobaric density (de Szoeke et al., 2000), which can be thought of as an upgraded specific volume anomaly: it maintains the existence of the exact geostrophic stream function, while improving material conservation. Our calculations of neutrality (a close cousin of material conservation; not shown) in the Southern Ocean State Estimate indicate it generally performs better than neutral density; the amount of improvement varies with depth, but can be as high as a factor of 10 (not shown). Since orthobaric density is rarely used, it is elaborated upon in the next section.

Orthobaric density, introduced by de Szoeke et al. (2000) and denoted ρ_ν, is a pycnotropic variable: it is a function only of pressure and in-situ density. As such, the existence of an exact geostrophic stream function is guaranteed (de Szoeke, 2000). In other words, the Boussinesq equations using orthobaric density as the vertical coordinate are those of Young (2012), with $b = -g\rho_c^{-1}\rho_\nu$. See also (3.8) or (4.58).

To calculate orthobaric density, one must first define a "virtual" compressibility that is a function of pressure and *in-situ* density alone, $\Gamma_0(p, \rho)$. The real compressibility[4] of seawater, $\Gamma = \partial_p \rho$, is a function of pressure, salinity, and and potential temperature (or equally well of pressure, *in-situ* density, and salinity). The goal is to fit the virtual compressibility to the real compressibility as well as possible; more will be said on this later.

The orthobaric density ρ_ν, as a function of p and ρ, is defined to be the density at a reference pressure p_{ref} when the density at p is ρ and the virtual compressibility, rather than the real compressibility, is used. That is, $\rho_\nu = R(p_{\text{ref}})$ where $R(p')$ is the solution of the ordinary differential equation

$$\frac{\partial R}{\partial p'} = \Gamma_0(p', R) \tag{3.27}$$

[4] Conventionally, $\rho^{-1}\partial_p\rho$ is called the compressibility. We choose to drop the ρ^{-1} factor for convenience.

subject to

$$R(p) = \rho. \tag{3.28}$$

This requires solving one ODE for every data point (p, ρ).

To understand orthobaric density, one can think of starting in the ocean at geographical location A at pressure p_A, where the seawater has salinity S_A, potential temperature θ_A, and hence *in-situ* density ρ_A. Now carry this seawater parcel upwards (if $p_{\text{ref}} < p_A$), while holding its salinity and potential temperature constant, to the reference pressure p_{ref}. Then $\rho_\nu(p_A, \rho_A)$ is the density at the end of this displacement, assuming the fluid parcel changed *in-situ* density according to Γ_0, not Γ.

Conceptually, one might ask whether this motion is directly vertical. It is not. As one moves up by an infinitesimal pressure change dp, the *in-situ* density changes by an amount $\Gamma(S_A, \theta_A, \bar{p})\text{dp}$, and the parcel moves some distance laterally (typically polewards) to find its neutral buoyancy at position B. Here, $\bar{p} = p_A + \text{dp}/2$. Were a parcel taken from this new location, with the ambient salinity S_B and potential temperature θ_B, and moved *down* by the same amount $-\text{dp}$, its *in-situ* density would change by an amount $\Gamma(S_B, \theta_B, \bar{p})\text{dp}$, giving an *in-situ* density different from ρ_A, and hence finding itself at a different lateral location. The use of a virtual compressibility Γ_0 corrects this problem: parcel A is moved up by dp then moves laterally until it rests neutrally at position C, where the *in-situ* density is $\rho_C = \rho_A + \Gamma_0(\bar{p}, \rho_A)\text{d}p$; this differs (only slightly, it is hoped) from ρ_B, but it does ensure that a parcel moved from C down by $-\text{dp}$ following the virtual compressibility curve will return exactly to A—or rather, to the same 1D curve described by (p_A, ρ_A).

Alternatively—and this is what I do—one can obtain the pressure of a particular orthobaric density surface ρ_ν by inverting (3.28). I calculate $R(p)$ over the full pressure range found in the ocean by solving (3.27) using a non-adaptive fourth order Runge-Kutta solver, subject to $R(0) = \rho_\nu$ (having also chosen $p_{\text{ref}} = 0$). Then for each water column i having an *in-situ* density profile $\rho_i(p)$, a root-finding method is used to find p such that $R(p) - \rho_i(p) = 0$. I use a simple bisection method. I pre-compute $R(p_i)$ on a very fine grid of evenly spaced p_i values (3000 of them) and then evaluate $R(p)$ by linear interpolation, so evaluation of $R(p)$ is cheap and bisection is efficient.

How does one construct Γ_0? I use data from the winters of 2005 and 2006, which in the SOSE data set are those producing the densest surface waters around Antarctica. I take an Eulerian time-average of p, ρ, and Γ over these winters, then exclude everything north of 55°S as well as shallow coastal regions and buoyant meltwaters around Antarctica. I then fit Γ_0 as a piecewise-linear function through this data (p, ρ, Γ) on a mesh with 200 evenly spaced data points spanning ρ, and 42 unevenly spaced data

points spanning p at the mean pressure of each model depth. This function of (p, ρ) is bilinearly interpolated to get Γ_0.

Following standard practice, I report in terms of $\sigma_\nu \equiv \rho_\nu - 1000\,\mathrm{kg\,m^{-3}}$, and frequently drop the units.

3.3 Q versus B correlations spanning the Antarctic Circumpolar Current

3.3.1 Histograms of Q vs B

We now examine the potential vorticity and Bernoulli potential in the ACC, from SOSE's 5-day archived outputs, on orthobaric density surfaces. For each 5-day archived state of SOSE's 6 year run, I compute the pressure, layer thickness, potential vorticity, and Bernoulli potential of 14 orthobaric density surfaces: $\sigma_\nu = 26.7$, 26.9, 27.0, 27.1, 27.2, 27.3, 27.4, 27.5, 27.6, 27.7, 27.75, 27.8, 27.825, and 27.85. The surfaces $\sigma_\nu = 26.7$, 26.9, and 27.85 are not circumpolar (at least in winter), and $\sigma_\nu = 27.0$ and 27.825 are just barely circumpolar, so I focus analysis on σ_ν in the range 27.1–27.8. For orientation, the depths of these surfaces are shown in Fig. 3.1 on a longitude section at 240°E from one archived time step (January 1–5, 2010); they overlay the PV shown in colour, computed on 200 orthobaric density surfaces.

I pre-define a region of interest for the ACC, which is enclosed between two contours of the time-mean sea-surface height; the two contours are chosen to span most of Drake Passage. The left panel of Fig. 3.2 shows Q vs B in this region of interest for all orthobaric surfaces (27.1–27.8) for 22–26 March, 2006. This time is representative of the deepest winter in the record; atmospheric forcing is not expected to penetrate deeper than it does at this time. The brightness of the colour linearly maps to the number of data points within the 2D bin. All but the very deepest layers exhibit a very high (absolute) PV on the low Bernoulli side (the south side): this occurs where the layer lies just below the mixed layer, where stratification is large—recall Fig. 3.1. Aside from this, i.e. in the ocean interior, there is a remarkably linear relationship between Q and B.

There are two obvious ways to eliminate this near-mixed layer signal: one is to remove all data below a certain pressure threshold, and the other is to remove all data below a certain Bernoulli threshold. The latter is more attractive in that it does not rely on a third variable. I attempted this using the Multivariate Adaptive

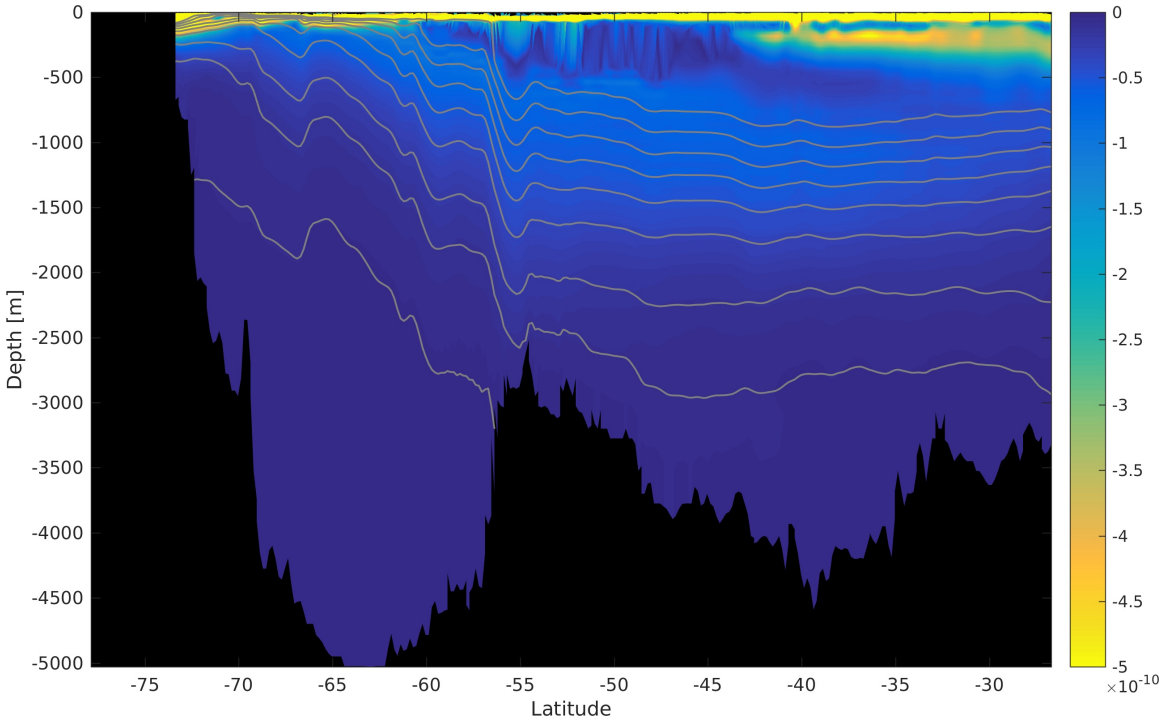


Figure 3.1: The nine σ_ν surfaces examined in this study, plotted as grey curves in the meridional plane at 240°E longitude, for January 1–5, 2010. The corresponding values, from shallowest to deepest are 27.1, 27.2, 27.3, 27.4, 27.5, 27.6, 27.7, 27.75 and 27.8. The Ertel potential vorticity, Q [$\mathrm{m^{-1}\,s^{-1}}$] for this cross-section is shown in colour. A perturbation associated with a large eddy can be seen in the vicinity of -56° latitude. The ocean floor is indicated in black. Large PV is confined to a rather small area near the sea-surface.

Regression Splines (MARS) algorithm (Friedman, 1993), as implemented in ARES-Lab (Jekabsons, 2016) to automatically minimize the error in Q as a piecewise linear function of B with one free, internal knot. This knot is the threshold for B, below which the data is ignored for further processing. This procedure seemingly produces well-chosen knots, but in the final analysis, the inverse Mach number becomes sensitive to the knot position[5]: uniformly shifting the knot to include a little more or less water noticeably changed the result.

With a pressure threshold, however, the results are much less sensitive to the particular threshold, as long as it not too shallow. Charney-type instabilities are surface-trapped and have a (self-chosen) depth scale typically shallower than 300 dbar, though can reach 500 dbar in some isolated regions (Smith, 2007). After some trial and error, I choose $p = 300$ dbar to generally ignore these shallow instabilities. The

[5] The inverse Mach is particularly sensitive to the breakpoint in the final two seasons of the SOSE data set, for $\sigma_\nu = 27.2$ and $\sigma_\nu = 27.8$; why these times and surfaces and not others is not clear.

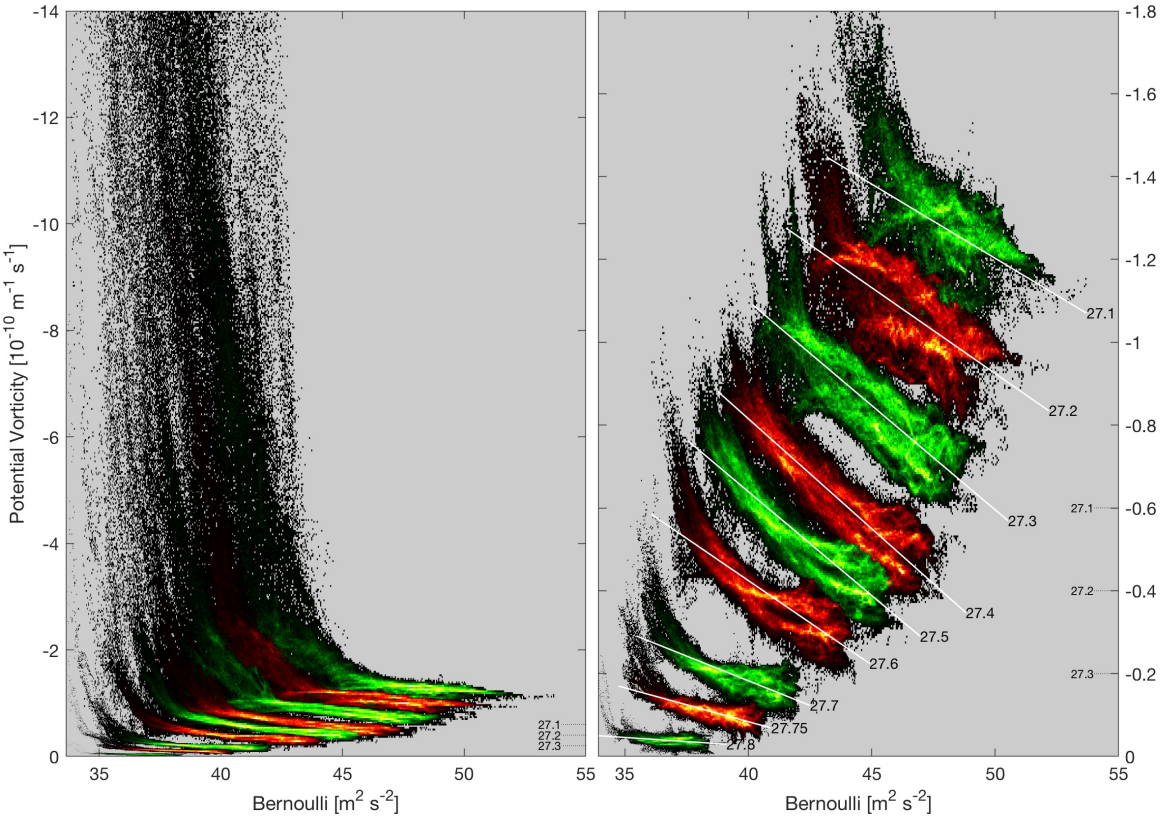


Figure 3.2: Two dimensional histograms of potential vorticity vs Bernoulli potential in the ACC on orthobaric density surfaces on 22–26 March, 2006. Nine surfaces are shown, from $\sigma_\nu = 27.1$ to 27.8. The brightness maps linearly to the number of data points per 2D bin. Surfaces $\sigma_\nu = 27.1$, 27.2, and 27.3 are progressively offset by $-1 \times 10^{-11}\,\mathrm{m^{-1}\,s^{-3}}$; the relative zero is indicated with dotted lines at right. Left panel: all data within the region of interest based on time-mean sea-surface height. Right panel: also excluding data with $p < 300\,\mathrm{dbar}$; white lines indicate the line of best fit, with σ_ν values indicated. Note the ordinate is flipped to facilitate comparison with the northern hemisphere.

right panel of Fig. 3.2 again shows Q vs B histograms, but with this cut-off. Having zoomed-in, the remarkably linear relations between Q and B may be scrutinized more closely. There are subpopulations to the Q vs B data, seen most clearly on $\sigma_\nu = 27.2$: for a given B in the ACC, there is one high $|Q|$ cluster and one low $|Q|$ cluster; they correspond to the Pacific and the Atlantic sectors respectively. This is consistent with a large drop of $|Q|$ along streamlines just downstream of Drake Passage (Gille, 1997). Chapter 4 explores this further, in the 6 year time-mean. There, the rapid drop of $|Q|$ at Drake Passage is countered by a steady gain elsewhere, whereas in this 5-day view, there is also a rather rapid gain.

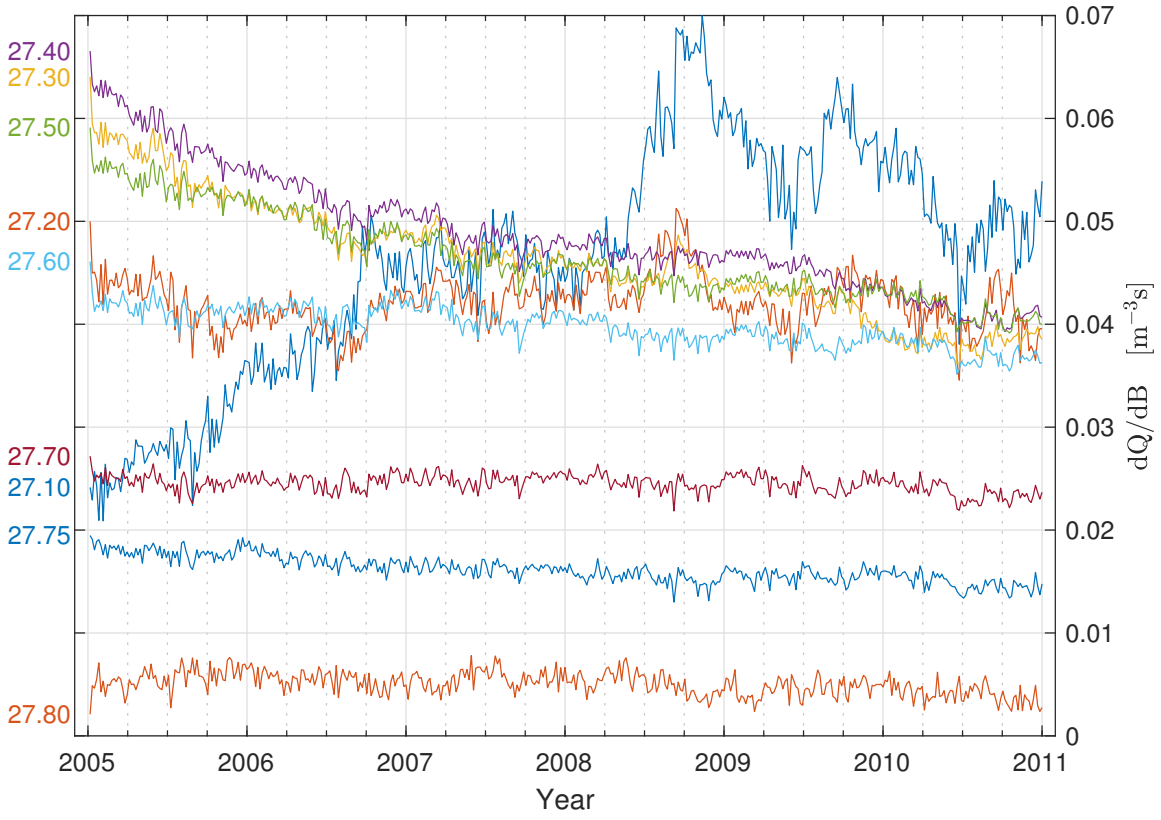


Figure 3.3: Time-series of the slope of Q vs B, i.e. q_1 in (3.19), for many orthobaric surfaces (labelled at left) over the whole ACC region of interest, excluding where the pressure is less than 300 dbar.

3.3.2 Linear regression

We now perform simple linear regressions over the whole ACC, with the region of interest and pressure threshold defined above, for each 5-day average.

Since Q and B are interdependent variables, one might argue that an errors-in-variable model should be used rather than simple linear regression with Q the dependent variable and B the independent variable. Indeed, treating B as the dependent variable, simple linear regression gives B as a linear function of Q, which is easily inverted to get Q as a function of B. However, this slope will always be greater than the slope from simple linear regression using Q as the dependent variable. Geometric mean regression simply takes the geometric mean of these two slopes, and passes the line of best fit through the mean of both variables; it is equivalent to minimizing the rectangular area created by each data point and the line of best fit (and thus is suited to variables with different units, unlike Deming regression which minimizes the orthogonal distance from the data point to the line of best fit). However, this

regression avoids horizontal (or vertical) lines of best fit, since the area formed by any data point off the line tends to infinity as the slope tends to 0 (or infinity). This is undesirable because we wish to study PV homogeneity. We shall return to this when making local fits in the next section, but for ACC-wide fits we get α by treating Q as the dependent variable.

The slope of the Q vs B relation (3.19) is $\mathrm{d}Q/\mathrm{d}B = q_1$ and is shown in Fig. 3.3. Though not used in calculating "Ma"$^{-1}$ by (3.24), it remains interesting. All layers have a positive slope, i.e. as B increases towards the equator, so too does Q, being negative. Deep layers ($\sigma_\nu \geq 27.6$) have a lesser slope, and less interannual variability, but more high-frequency variability. Middle layers ($27.3 \leq \sigma_\nu \leq 27.5$) appear somewhat locked together, with a gradual decrease in slope through the 6 year period. The slope for shallowest layer ($\sigma_\nu = 27.1$) shows marked growth, reaching a clear maximum in late winter 2008. Smaller maximums exist in late winter of other years, though much less prominently in 2007, and not at all in 2005 as the growth carries on into 2006. The second shallowest layer ($\sigma_\nu = 27.2$) shows no secular trend, perhaps counterbalancing the growth of the layer above and the decay of those below. The three shallowest layers all share the winter 2008 maximum, while only the two shallowest share the winter 2010 maximum.

Next, we determine the reference frame shift α from (3.20) and (3.21). An example is shown in Fig. 3.5 for $\sigma_\nu = 27.2$ during 22–26 March, 2006. In the new reference frame, the Q vs $B^{(\alpha)}$ relation becomes remarkably tight. At this time, $\alpha = -1.08 \times 10^{-8}\,\mathrm{s}^{-1}$ corresponding to a westward phase speed of $3.95\,\mathrm{cm\,s}^{-1}$ at 55°S. Also at this time, "Ma"$^{-1} = 2.03$. Deeper surfaces generally have "Ma"$^{-1}$ closer to 1, with α closer to 0 and the Q vs $B^{(\alpha)}$ looking more like Q vs B.

The (total) long Rossby wave phase speed at 55°S, computed as $\alpha R\cos(-55°)$, is shown in Fig. 3.4. The major late winter maximums of $\mathrm{d}Q/\mathrm{d}B$ for $\sigma_\nu = 27.1$ are evident here as fast (relative to other times) westward Rossby wave phase speeds. The long Rossby wave phase speed is westward for all surfaces, except for the deepest circumpolar layer $\sigma_\nu = 27.8$. This is in contrast to the generally accepted notion that ACC's fast eastward advection manages to Doppler-shift baroclinic Rossby waves to have an eastward phase speed (Klocker and Marshall, 2014). This discrepancy suggests the possibility that the barotropic Rossby wave, which also influences the local Q and B fields and has a fast westward phase speed even after Doppler shifting by the eastward flow, has also been incorporated, to some degree, in the local empirical Rossby wave phase speed estimate. However, there is another interpretation, as follows.

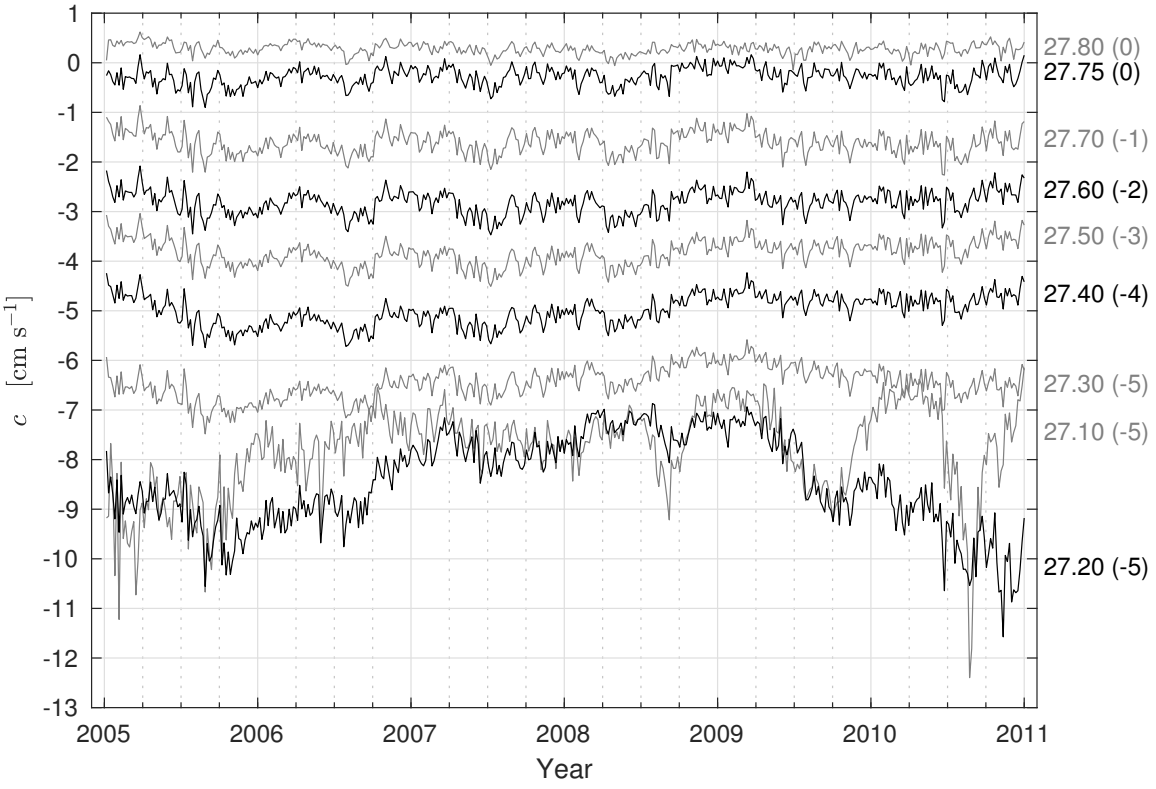


Figure 3.4: Time-series of the Rossby wave phase speed at 55°S associated with the frame shift α, for each orthobaric density surface. The orthobaric density for each surface is shown at right; next to this in parentheses is the offset value for that surface, in cm s^{-1}.

The view of shear instability as co-amplifying Rossby waves requires that the waves have the same total phase speed. Considering planetary-scale disturbances (requiring spherical geometry), the more precise condition is that the waves are both stationary in some common reference frame. Now, this reference frame must be fairly close to Earth's reference frame, otherwise topography, which moves in the common reference frame, will kill the waves quickly, before they can significantly amplify each other. This complements the shockwave explanation of why $\alpha = 0$ corresponds to "Ma"$^{-1} = 1$. The longest Rossby waves, corresponding to planetary-scale disturbances, have an intrinsically faster wave speed than short Rossby waves—see (3.12). Thus, while a planetary-scale disturbance is roughly at rest in Earth's reference frame, smaller-scale eddies get swept downstream. Jupiter provides a nice example of this, with its Great Red Spot nearly at rest and smaller eddies being advected past it. It would appear that these ACC-wide fits have somehow conglomerated the hundreds of eddies into a planetary-scale disturbance, with its own phase speed.

The (inverse) Mach number, from (3.24), is shown in Fig. 3.6. Only the deepest

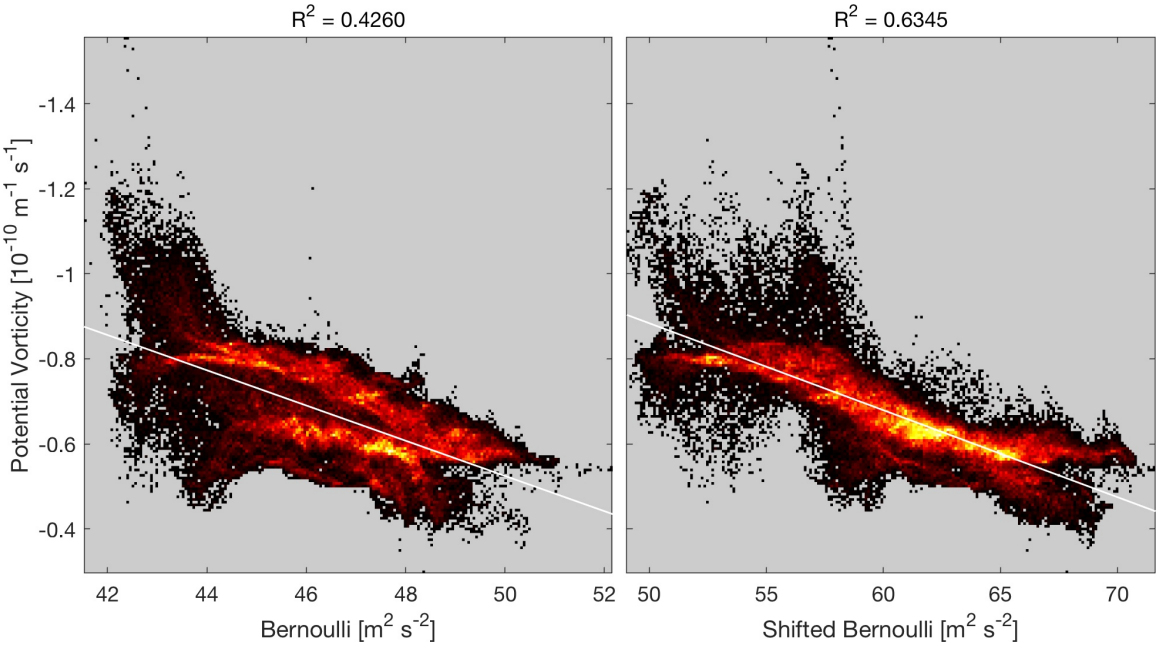


Figure 3.5: Two dimensional histograms of potential vorticity vs (α-shifted) Bernoulli potential in the ACC region of interest, excluding $p \leq 300\,\mathrm{dbar}$, on $\sigma_\nu = 27.2$, 22–26 March 2006. Left panel: as in the right panel of Fig. 3.2. Right panel: Q vs $B^{(\alpha)}$, according to (3.14) and (3.20) and (3.21). The line of best fit is shown in white, and the R^2 coefficient of determination is indicated above.

circumpolar layer, $\sigma_\nu = 27.8$, has "Ma"$^{-1} < 1$, corresponding to this being the only layer with $\alpha > 0$. However, "Ma"$^{-1}$ remains above 0.5 and hence even on $\sigma_\nu = 27.8$ the PV should not be regarded as homogeneous. The surfaces $\sigma_\nu = 27.75$ to $\sigma_\nu = 27.3$ have "Ma"$^{-1}$ between 1 and 1.5, and generally larger for shallower layers. Thus these deep layers hover at or a bit above the threshold of instability, and definitely exhibit inhomogeneous PV. The shallower layers $\sigma_\nu = 27.1$ and 27.2 have "Ma"$^{-1}$ generally between 1.5 and 2.5. They track each other quite closely during 2007–2009. At times, "Ma"$^{-1}$ is larger for $\sigma_\nu = 27.2$ than for $\sigma_\nu = 27.1$, such as in 2010 (excluding a large late-winter spike for $\sigma_\nu = 27.1$) and summer 2005/2006. The seasonal cycle is evident for the shallowest surface, $\sigma_\nu = 27.1$ with elevated winter values of "Ma"$^{-1}$, particularly for late winter 2008, 2009, and 2010 which match the peaks of $\mathrm{d}Q/\mathrm{d}B$ (Fig. 3.3). However, the strong secular increase in $\mathrm{d}Q/\mathrm{d}B$ on $\sigma_\nu = 27.1$ is not reflected in a secular increase in "Ma"$^{-1}$. Some interannual variability is evident on $\sigma_\nu = 27.2$. High-frequency variability is higher for deeper layers, as it was for $\mathrm{d}Q/\mathrm{d}B$. The high-frequency variability in 2005 on $\sigma_\nu = 27.1$ is probably due to atmospheric forcing, and the mixed layer penetrating deeper than 300 dbar.

A tight functional relationship between PV and Bernoulli potential was not *a*

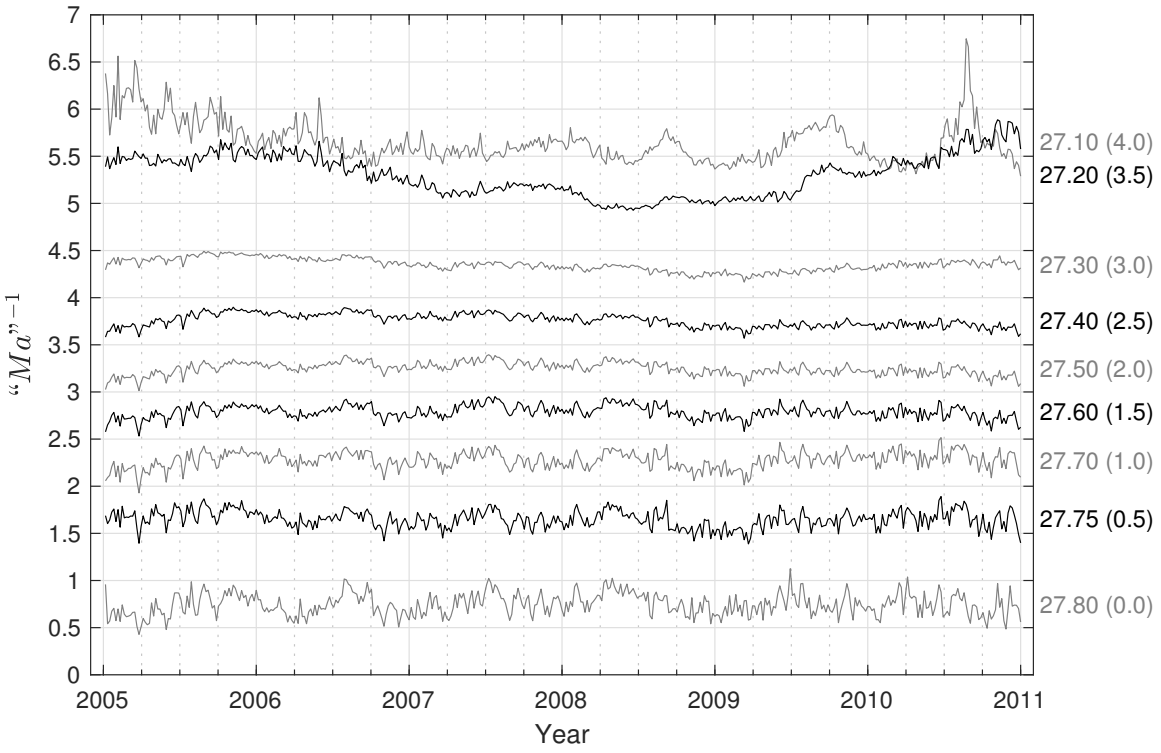


Figure 3.6: Time-series of the inverse Mach number for each orthobaric density surface, offset by 0.5 per surface. The orthobaric density for each surface is indicated; next to this in parenthesis is the offset value for that surface.

priori guaranteed: it could have been precluded by time-evolution, mechanical and buoyancy forcing of PV, or topological considerations with closed B contours supporting multiple Q values per B isovalue. Nonetheless, an impressive relation does exist over the whole ACC. This section focused on the temporal aspect of this signal, whereas the next section investigates the spatial aspect. We found the ACC has $\text{“}Ma\text{”}^{-1} > 0$ for all circumpolar orthobaric surfaces, but might that arise as a mixture of homogeneous PV in some locations (e.g. the deep northern regions) and strong PV gradients in other locations (e.g. the shallow southern regions)?

3.4 Local Q versus B correlations

3.4.1 Q the dependent variable

The methods used on the whole ACC are now applied to local regions. For each water column, a circular region 7 grid points in diameter (just over 1°, which at 60°S is 55 km zonally, a few times the deformation radius which is roughly 15 km) is considered the local region of interest; this constitutes up to 37 water columns (less

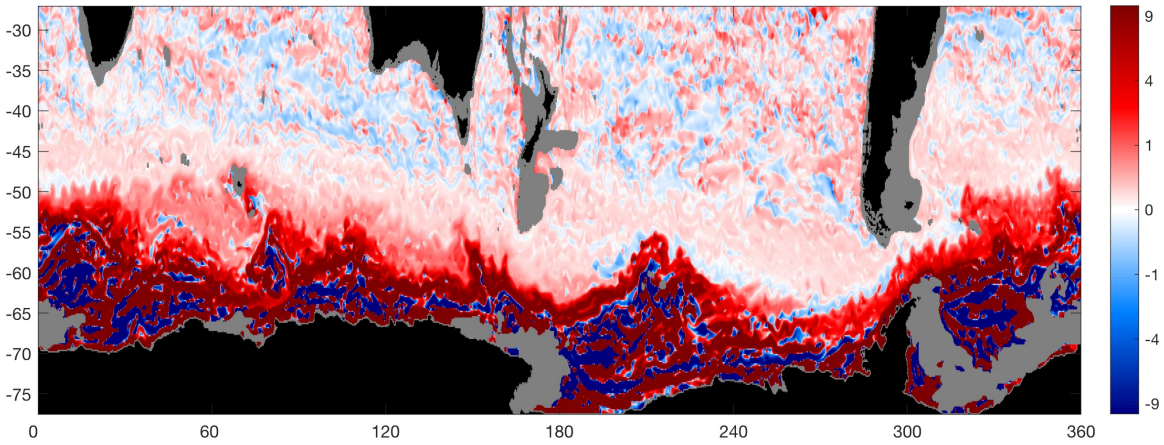


Figure 3.7: The slope, $\mathrm{d}Q/\mathrm{d}B = q_1$, from simple linear regression of Q to B using data local to each water column, on the $\sigma_\nu = 27.3$ surface on 22–26 March, 2006. Note the nonlinear colour scale.

near land).

As before, I use simple linear regression to fit Q to B, as in (3.19), on this local data. The slope, $\mathrm{d}Q/\mathrm{d}B$, is mapped in Fig. 3.7 for the surface $\sigma_\nu = 27.3$ on 22–26 March, 2006. In most of the ACC, $\mathrm{d}Q/\mathrm{d}B$ is positive (both Q and B increase northwards). At the southern flank of the ACC (where this is shallower than about 300 dbar, not shown) the slope increases drastically in magnitude. There are patches in and south of the ACC where the slope is negative. One example is in the South Pacific, where mode water formation has injected filaments of near-zero PV water, on the northern flank of which $Q_{\tilde{y}}$ and hence $\mathrm{d}Q/\mathrm{d}B$ are negative. The slope in the ocean basins north of the ACC is comparable in magnitude to that in the ACC, but is as much negative as positive, roughly. The flow being weaker here than in the ACC, B variations are smaller and hence Q variations are smaller.

Next, simple linear regression fits Q to $B^{(\alpha)}$, as in (3.20), on the local data, which yields α. As $\alpha = -a_2/a_1$ is computed as a ratio, some water columns will have unrealistically large values of α. Thus, α is converted to the Rossby wave phase speed, $c = \alpha R \cos\theta$, capped with a maximum absolute value of $20\,\mathrm{cm\,s^{-1}}$, and smoothed over about 1°; this is shown in Fig. 3.8's upper panel, while the lower panel shows the intrinsic Rossby wave phase speed by subtracting the local zonal flow. There is a band of slightly faster eastward Rossby waves on the north flank of the ACC, perhaps related to a weakening of the PV gradient. The waves' intrinsic phase speed is generally, but not entirely, westward. Theoretically, regions of intrinsic eastward phase speed are either due to a reversal of the meridional PV gradient, or a sign reversal of the coefficient multiplying the PV gradient, e.g. $\mathcal{M}/\mathcal{M}_{\tilde{b}\tilde{b}} > 0$ in (3.38).

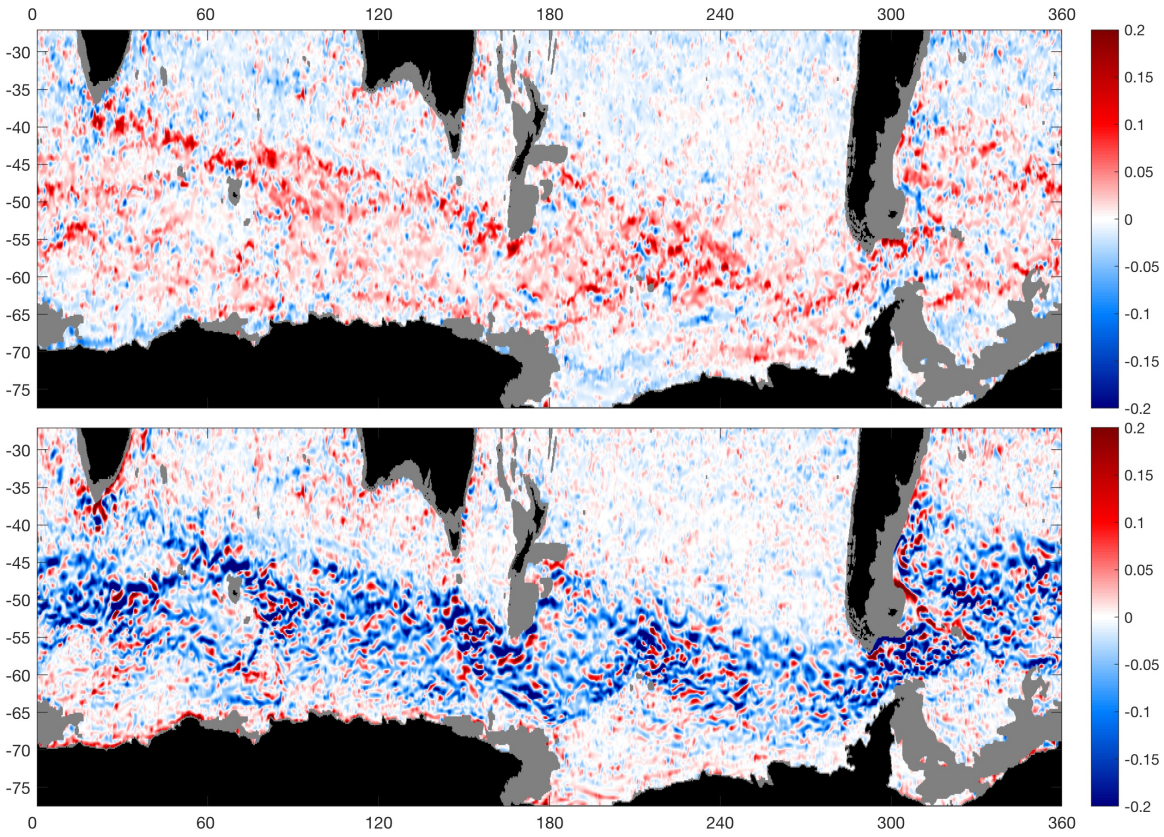


Figure 3.8: The Rossby wave zonal phase speed [m/s] on $\sigma_\nu = 27.3$, 22–26 March 2006: total (upper panel) or intrinsic (lower panel). Data local to each water column is used to fit Q to B and A as in (3.20), to obtain α.

Numerically, this could also be caused because $\hat{c}$ is calculated as $c - u$, and $c = -u^{(\alpha)}$ is calculated using a stencil of data surrounding the local grid point, whereas u is taken only as that from the local grid point.

In the local fits, the Rossby wave phase speed seems more eastward than westward, consistent with downstream advection by the strong background flow. This helps validate that this Q vs B method can reasonably estimate the Rossby wave phase speed, and suggests that the slow westward phase speed from the ACC-wide fits should not be dismissed. However, it is perhaps surprising that the Rossby wave phase speed in these local fits seems not to tend to zero near topography, as it arguably should if the reasoning of Section 3.3 is to be believed. It could tend to zero on a shorter scale than can be resolved by SOSE, though this is doubtful. It could be that this view is too myopic for the Rossby waves revealed by it to phase lock onto something, and only in a broader view can many Rossby waves phase lock together and then onto multiple bits of topography.

Finally, Fig. 3.9 shows the inverse Mach number, calculated by reverting the above

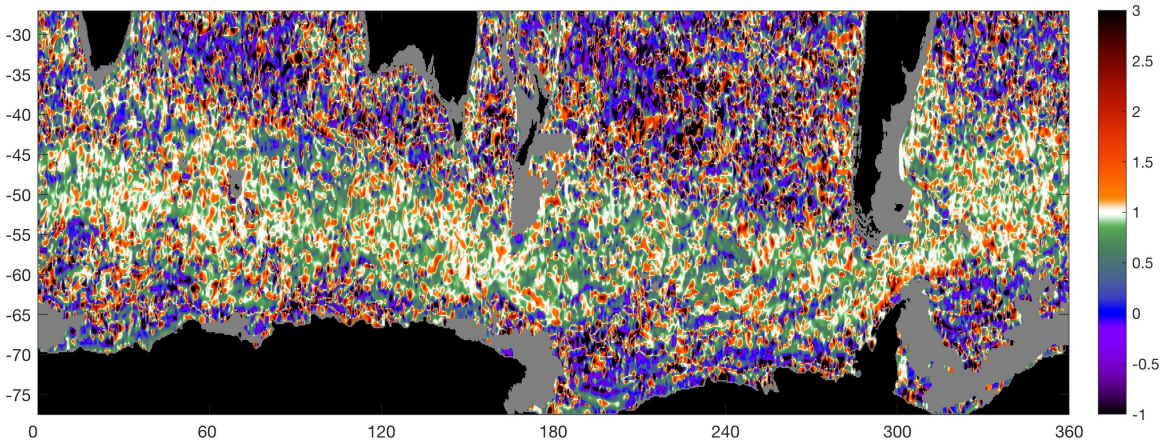


Figure 3.9: The inverse Mach number, "Ma"$^{-1}$ from (3.24), using the smoothed α field corresponding to the Rossby wave phase speed of Fig. 3.8, and simple linear regression of A to B to get A_1. Data is local to each water column, on the $\sigma_\nu = 27.3$ surface on 22–26 March, 2006.

smoothed Rossby wave phase speed to its α form, then using (3.24), with simple linear regression of local data A onto B to get A_1. The ACC lights up as a region with "Ma"$^{-1} \approx 1$, but far from uniformly-so. Indeed, the ACC is populated with hundreds of mesoscale eddy-sized regions having "Ma"$^{-1} \approx 1$, sometimes containing inside them a pocket of "Ma"$^{-1}$ clearly greater than 1. These may or may not be the mesoscale eddies themselves. Zooming in to the Cape of Good Hope, the "Ma"$^{-1}$ map does have a large patch of "Ma"$^{-1} \approx 1$ corresponding to an Agulhas eddy, but it does not stand out particularly obviously (not shown). In contrast, the Pacific, Indian, and Atlantic ocean basins are, broadly speaking, dark regions of "Ma"$^{-1} \leq 0$, though of course populated with strands of larger "Ma"$^{-1}$. The colour bar in Fig. 3.9 is circular, because "Ma"$^{-1} \to \pm\infty$ is just "Ma" (more directly related to Arnol'd's stability criteria) approaching zero, i.e. near the threshold of Arnol'd's first stability criterion.

Using (3.22) produces a map of "Ma"$^{-1}$ that is vastly noisier, due to the nature of the ratio there used. However, the basic qualitative feature of "Ma"$^{-1} \approx 1$ in the ACC and "Ma"$^{-1} < 1$ elsewhere remains (not shown).

3.4.2 B the dependent variable

In local regions of nearly zonal flow, B and A are highly correlated, whereas Q and A tend to be less correlated, because Q has much more small-scale structure . This is because Q is a differentiated variable, or physically, because of the forward enstrophy cascade; whereas B is an integrated variable, so has less small-scale structure. Thus,

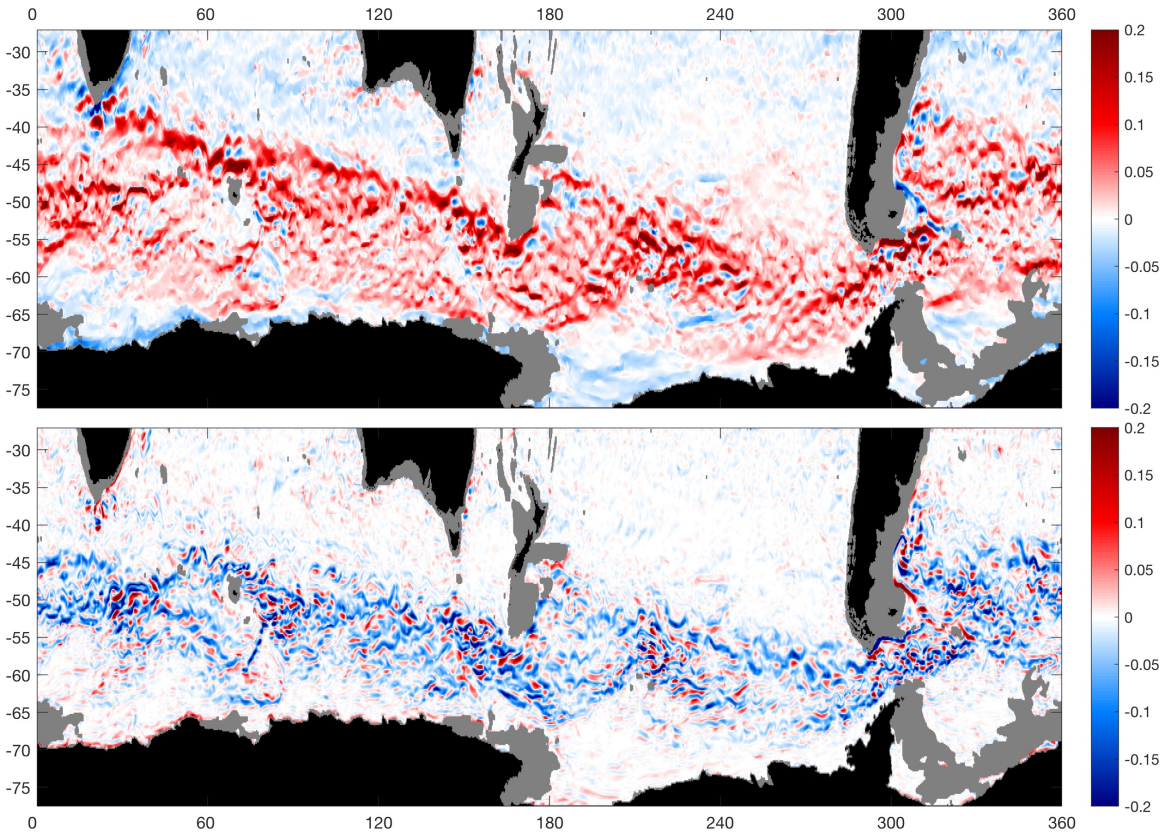


Figure 3.10: The Rossby wave zonal phase speed [m/s] on $\sigma_\nu = 27.3$, 22–26 March 2006: total (upper panel) or intrinsic (lower panel). Data local to each water column is used to fit B to Q and A as in (3.29), to obtain α.

the above method that gets α by simple linear regression of Q to B and A, (3.20), suffers from multicollinearity: one predictor variable can be well-predicted from the other(s), and the model coefficients are highly sensitive to fluctuations in the predictor variables. This is not an issue for the whole-ACC fits of Section 3.3, for which Q is almost as well correlated with A as is B.

This problem is alleviated by treating B as the dependent variable, i.e. use simple linear regression to fit

$$B = b_0 + b_1 Q + b_2 A. \tag{3.29}$$

Note that $\alpha = b_2$. Thus $\mathrm{d}B^{(\alpha)}/\mathrm{d}Q = b_1 = a_1^{-1}$, so the inverse Mach number in the form (3.22) becomes

$$\text{“}Ma\text{”}^{-1} = q_1 b_1. \tag{3.30}$$

The ratio has been expunged, and this produces smoother results.

The Rossby wave phase speed (total and intrinsic) is shown in Fig. 3.10. Unlike for Fig. 3.8, no smoothing was used or needed. It agrees, qualitatively, very well

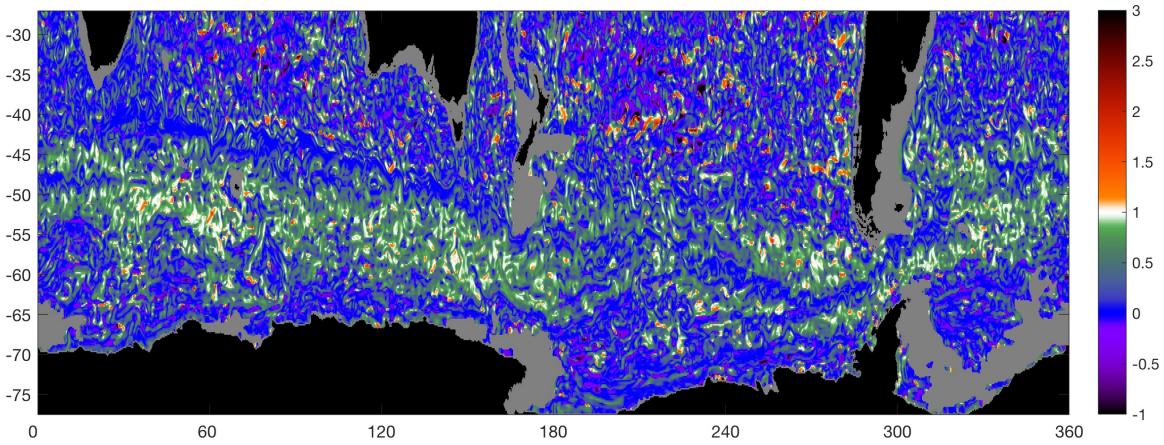


Figure 3.11: The inverse Mach number, "Ma"$^{-1}$ from (3.24), using simple linear regression of B to Q and A to get α, and of A to B to get A_1. Data is local to each water column, on the $\sigma_\nu = 27.3$ surface on 22–26 March, 2006.

with the phase speed of the most unstable mode from the eigenvalue problem, either in buoyancy or depth coordinates (not shown, see Appendix 3.A). A more detailed comparison is shown in Fig. 3.12, zooming into Drake Passage. Panels (a)–(d) use the unstable modes' phase speeds for various methods (Appendix 3.A). Panels (e) and (f) use the local methods to estimate α, treating either Q as the dependent variable as in (3.20), or B the dependent variable as in (3.29), respectively. Using B as the independent variable evidently produces a very reasonable Rossby wave phase speed.

Finally, with B the dependent variable, Fig. 3.11 shows the inverse Mach number. As expected, this field is smoother and takes less extreme values than that of Fig. 3.9. Again, the ACC lights up with many small scale features taking "Ma"$^{-1} \approx 1$, in between which are regions of $0 \leq$ "Ma"$^{-1} \leq 1$, roughly. The ocean basins to the north, more distinctly than before, have "Ma"$^{-1} < 0$ generally. A long feature of "Ma"$^{-1} \approx 0$ (homogeneous PV) appears on the northern flank of the ACC between the Cape of Good Hope and Campbell Plateau. Evidently, there exists here a reasonably strong flow but a weak PV gradient, possibly associated with strong eddy activity causing down-gradient PV fluxes. Another such feature appears in the southern flank of the ACC upstream of Drake Passage, probably related to mode water.

It must be noted that using B as a dependent variable to get b_1, but Q as the dependent variable to get q_1, produces a conservative (i.e. biased low) estimate of "Ma"$^{-1}$. This is because the slope of Q regressed onto $B^{(\alpha)}$ is always less, in magnitude, compared to regressing $B^{(\alpha)}$ onto Q then inverting. That is, $|a_1| < |1/b_1|$ is guaranteed. Thus, "Ma"$^{-1}$ estimated by (3.30) is always less than that estimated by (3.22).

3.5 Discussion and conclusion

A functional relationship between potential vorticity and Bernoulli potential is a powerful concept. It ties in the mass-weighted stream function into a triad of functional relationships, where any of Q, B, or Ψ can be determined from two others. This triad completely specifies a solution of the unforced, steady primitive equations. This is invaluable because it enables, by linearising the unsteady primitive equations about this "background state", one to study Rossby waves. Classic studies of Rossby waves pick a background state at rest or with a purely zonal flow; with the Q, B, Ψ triad, more realistic background states may be available. One could go further to study non-linear instabilities. Of course, the background state is modified by Rossby waves and non-linear eddies, so this is not a closed problem.

Dowling (2014) identified the inverse Mach number for Rossby waves with Arnol'd's shear-stability criteria. The two stability branches of Arnol'd coalesce at "Ma"$^{-1} = 0$, for which the flow is stable and the PV is homogeneous. Arnol'd's second branch of shear-stability is limited at "Ma"$^{-1} = 1$, above which instability is possible and the PV is said to be "choked" (by reference to choked flow in aeronautics, where flow undergoes a subsonic to supersonic transition and the upstream flow becomes disconnected from the downstream conditions). Dowling (2014) speculates that the threshold of shear instability is itself the shock of Rossby waves.

The Southern Ocean is commonly viewed as (a) highly unstable, and (b) having homogeneous PV on buoyancy surfaces (Marshall et al., 1993). However, one can't have it both ways. If PV is truly homogeneous, then "Ma"$^{-1} = 0$ and the flow must be stable. The near alignment of PV surfaces with buoyancy surfaces is not enough to justifiably say that the PV on buoyancy surfaces is homogeneous. By estimating "Ma"$^{-1}$ on orthobaric density surfaces in the ACC, I conclude that the PV is far from homogeneous in the ACC. In a coarse view of the whole ACC, the flow configuration is above, but not wildly beyond, the stability threshold for all but the deepest circumpolar layers, though even there PV remains about halfway between homogeneous and choked. From a local perspective, patches of subsonic flow, (beyond the stability threshold) appear sprinkled throughout the ACC, amidst a calmer background of PV part way between homogeneous and choked.

The calculation of α to best fit Q vs $B^{(\alpha)}$ is novel, but potentially confusing. Mechanical and buoyancy forcing has been neglected in this analysis. Is this, rather than time-tendency, the key to the remarkable linear relationship between Q and $B^{(\alpha)}$ in Fig. 3.5? Is it the key to the westward Rossby wave phase speed in the

ACC-wide fits of Section 3.3 but the generally eastward phase speed in the local fits of Section 3.4? The α shift aligned the two subpopulations of (Q, B). The next chapter shows that such forcing is important in the time-mean PV budget very near topography. Along B contours that stay roughly half a degree away from topography, eddy buoyancy fluxes and vertical advection piercing buoyancy surfaces are the major PV forcings. Relating eddy forcing in the time-mean state with tendency in the instantaneous state, I thus expect most B contours to exhibit a PV budget dominated by PV tendency rather than mechanical or buoyancy forcing. The subpopulations of (Q, B) found in Fig. 3.2 (right panel) would be smeared out in the time-mean, in which a rapid drop in $|Q|$ at Drake Passage is balanced by steady growth of $|Q|$ at other longitudes. This is evidently not the case in a 5-day average. Nonetheless, interesting questions remain for further research, and this empirical α method should be tested in a more controlled setting than SOSE, such as a simple channel model.

Most B contours in the ACC are circumpolar, so the full topological machinery of the Reeb graph was not used, for this initial study. Working with local data per water column in Section 3.4 also alleviated this issue. However, with $\alpha > 0$, the northward increase of B is countered by a northward decrease in $-\alpha A$, so contours of $B^{(\alpha)}$ become less circumpolar. In the extreme case that α matches the mean ACC zonal speed, hundreds of eddies that populate the ACC appear as closed $B^{(\alpha)}$ contours. Tight functional relationships between Q and $B^{(\alpha)}$ would then require the Reeb graph. Choosing α to thus complicate the topological setting may, assisted by the Reeb graph, ultimately yield far deeper insights of the statistical mechanics of turbulence than is possible in the Earth's reference frame. Indeed, from their barotropic channel model David et al. (2017) found the probability distribution function of Q becomes far cleaner when viewed in a Lagrangian frame of reference.

Appendix

3.A Rossby wave vertical modes

The Rossby wave phase speed is critical to the Mach number estimate, and we've seen that different methods yield somewhat different results for both the Rossby wave phase speed and, ultimately, the Mach number. Thus, a closer look at the Rossby wave phase speed is warranted: how do the estimates in this section (Fig. 3.8 and Fig. 3.10) compare with estimates of the Rossby wave phase speed from standard

theory? That is reviewed first, followed by a derivation in buoyancy coordinates possessing several pleasing properties.

3.A.1 Derivation in depth coordinates

Depth coordinates provide the more traditional derivation of the Rossby wave dispersion relation. In the ACC, even though $u_{\tilde{y}\tilde{y}}$ can be as large as β in the jets and so is important for small-scale features, the large-scale features of the PV gradient are mostly determined by variations of horizontal stratification, rather than relative vorticity. Thus, we consider the planetary geostrophic equations (rather than the quasi-geostrophic equations), which in depth coordinates are as follows:

$$-fv = -\phi_x, \tag{3.31a}$$

$$fu = -\phi_y, \tag{3.31b}$$

$$\phi_z = b, \tag{3.31c}$$

$$u_x + v_y + w_z = 0, \tag{3.31d}$$

$$b_t + ub_x + vb_y + wb_z = 0, \tag{3.31e}$$

where $\phi = p - g\int_z^0 \tilde{\rho}(z')\,\mathrm{d}z'$, $\tilde{\rho}(z)$ is a reference density profile, and $b = -g\rho_c^{-1}(\rho - \tilde{\rho})$ is the buoyancy. In z-coordinates, Welander (1959) expressed all the planetary geostrophic variables in terms of a single variable ω[6], which is the vertical integral of the pressure. Thus, let $\omega_z = \phi$, and observe u, v, w, and b are all functions of ω: $u = -\omega_{zy}/f$, $v = \omega_{zx}/f$, $w = \omega_x\beta/f^2$, $\phi = \omega_z$, and $b = \omega_{zz}$. Thus (3.31e) becomes

$$\omega_{zzt} - \frac{\omega_{zy}}{f}\omega_{zzx} + \frac{\omega_{zx}}{f}\omega_{zzy} + \frac{\omega_x\beta}{f^2}\omega_{zzz} = 0. \tag{3.32}$$

Letting $\omega = \Omega + \omega'$ for some background state Ω having $\Omega_x = 0$, then supposing $\omega' = \Re(\mathcal{W}(z)\exp i(kx + ly - \omega t))$ leads to

$$(c-u)\mathcal{W}_{zz} + u_z\mathcal{W}_z - \frac{\beta b_z}{f^2}\mathcal{W} = 0, \tag{3.33}$$

where u and b are for the background state. with boundary conditions $\mathcal{W} = 0$ at $z = 0$ and $z = -H$, assuming the vertical velocity vanishes at these boundaries, as for a flat-bottomed ocean with a rigid lid (Killworth et al., 1997; Killworth and Blundell, 2003). The alternative boundary condition, corresponding to zero horizontal flow (Tailleux and McWilliams, 2001; de La Lama et al., 2016), is $\mathcal{W}_z = 0$ at $z = -H$.

[6] This variable is typically denoted M, but I use ω to avoid confusion with the Montgomery potential, and because its vertical structure is the same as that of the vertical velocity, w.

3.A.2 Derivation in buoyancy coordinates

We now derive the Rossby wave vertical modal structure and phase speed in buoyancy coordinates. In the ACC, even though $u_{\tilde{y}\tilde{y}}$ can be as large as β in the jets and so is important for small-scale features, the large-scale features of the PV gradient are mostly determined by variations of horizontal stratification, rather than relative vorticity. Thus, we consider the planetary geostrophic equations (rather than the quasi-geostrophic equations), which are (c.f. (3.8)),

$$-fv + m_{\tilde{x}} = 0 \tag{3.34a}$$

$$fu + m_{\tilde{y}} = 0 \tag{3.34b}$$

$$\xi + m_{\tilde{b}} = 0 \tag{3.34c}$$

$$\sigma_{\tilde{t}} + (\sigma u)_{\tilde{x}} + (\sigma v)_{\tilde{y}} = 0 \tag{3.34d}$$

where ξ is the depth of the buoyancy surface, the thickness is $\sigma = \xi_{\tilde{b}}$, m is the Montgomery potential, and I have further assumed buoyancy is perfectly materially conserved. Happily, the Montgomery potential serves as a single state variable from which all other fields may be determined.

Putting (3.34a)—(3.34c) into (3.34d) yields (with some cancellation)

$$- m_{\tilde{b}\tilde{b}\tilde{t}} + m_{\tilde{y}} m_{\tilde{b}\tilde{b}\tilde{x}} f^{-1} - m_{\tilde{x}} (m_{\tilde{b}\tilde{b}} f^{-1})_{\tilde{y}} = 0. \tag{3.35}$$

We now linearise this about a background state M (so $m = M + m'$).

$$- m'_{\tilde{b}\tilde{b}\tilde{t}} + m'_{\tilde{y}} M_{\tilde{b}\tilde{b}\tilde{x}} f^{-1} + M_{\tilde{y}} m'_{\tilde{b}\tilde{b}\tilde{x}} f^{-1} - m'_{\tilde{x}} (M_{\tilde{b}\tilde{b}} f^{-1})_{\tilde{y}} - M_{\tilde{x}} (m'_{\tilde{b}\tilde{b}} f^{-1})_{\tilde{y}} = 0. \tag{3.36}$$

Now, let us further assume the background flow is purely zonal, so that $M_{\tilde{x}} = 0$. Searching for a horizontal wave solution with as yet unspecified vertical form, $m' = \Re\left(\mathcal{M}(b) \exp i(kx + ly - \omega t)\right)$, we get

$$\frac{\omega}{k} \mathcal{M}_{\tilde{b}\tilde{b}} = -\frac{M_{\tilde{y}}}{f} \mathcal{M}_{\tilde{b}\tilde{b}} + \left(\frac{M_{\tilde{b}\tilde{b}}}{f}\right)_{\tilde{y}} \mathcal{M}. \tag{3.37}$$

Letting $c = \omega/k$ be the eastward phase speed and $Q = -f/M_{\tilde{b}\tilde{b}}$ the background large-scale PV, this becomes

$$c - u = \frac{Q_{\tilde{y}}}{Q^2} \frac{\mathcal{M}}{\mathcal{M}_{\tilde{b}\tilde{b}}}, \tag{3.38}$$

This result (3.38), applicable to a continuously stratified 3D fluid, has the particularly clean form of (3.3). The Rossby wave phase speed is Doppler shifted by the local background flow. The intrinsic phase speed, $c-u$, is proportional to the PV gradient,

and has high PV on its right (westward for $Q_{\tilde{y}} > 0$) when $\mathcal{M}/\mathcal{M}_{\tilde{b}\tilde{b}} < 0$, which is consistent with a wave-like vertical structure. In general, solutions can have sections of the fluid column with $\mathcal{M}/\mathcal{M}_{\tilde{b}\tilde{b}} > 0$, but these regions would grow (in the b dimension) exponentially and so are quickly countermanded. Note that c is depth-independent, but all other terms in (3.38) are depth-dependent. In general, $\mathcal{M}/\mathcal{M}_{\tilde{b}\tilde{b}}$, which replaces $-\kappa_0^{-2}$ in (3.3), is depth dependent, proving that it is incorrect to use the deformation radius L_d, which is a single value per fluid column, in place of κ_0^{-1}.

At the upper boundary, where $b = b_s$ the surface buoyancy, (3.34c) implies $m_{\tilde{b}} = 0$ at $b = b_s$. Assuming the background field must also satisfy these constraints (note the background field should not be thought of as the time-mean, which does not satisfy these constraints—see the discussion of boundary conditions in Young (2012)), then $m'_{\tilde{b}} = 0$ at $b = b_s$, hence

$$\mathcal{M}_{\tilde{b}} = 0, \quad \text{at } b = b_s. \tag{3.39}$$

Making the same assumptions for the bottom boundary condition yields $m_{\tilde{b}} = H$ at $b = b_b$ the bottom buoyancy; if the background field satisfies this constraint then the perturbation, again, has $m'_{\tilde{b}} = 0$ at $b = b_b$, so

$$\mathcal{M}_{\tilde{b}} = 0, \quad \text{at } b = b_b, \tag{3.40}$$

which is equivalent to $w = 0$ at the bottom.

However, the bottom boundary condition is open to debate. An alternative is that the bottom horizontal velocity vanishes (Tailleux and McWilliams, 2001). This produces vertical modes that generally better capture (i.e. fewer modes are needed) the oceanic flow (de La Lama et al., 2016), although an additional bottom trapped mode must be added for the modes to form a complete basis. In buoyancy coordinates, this translates to the bottom boundary condition

$$\mathcal{M} = 0, \quad \text{at } b = b_b. \tag{3.41}$$

Together (3.38), (3.39), and either (3.40) or (3.41) define an eigenvalue problem for the vertical structure $\mathcal{M}(b)$ of wave solutions to the planetary geostrophic equations (3.34).

To better capture the background state, the PV gradient (which contains much high-frequency noise) can be replaced using the functional relationship between Q^{-1} and M, so (3.38) becomes

$$c - u = -Q_{\tilde{y}}^{-1} \frac{\mathcal{M}}{\mathcal{M}_{\tilde{b}\tilde{b}}} = -Auf \frac{\mathcal{M}}{\mathcal{M}_{\tilde{b}\tilde{b}}}, \tag{3.42}$$

where $A = \mathrm{d}Q^{-1}/\mathrm{d}M$ is depth-dependent. This can be rearranged as

$$\text{“}Ma\text{”}^{-1} = Af\mathcal{M}/\mathcal{M}_{\tilde{b}\tilde{b}}, \tag{3.43}$$

so the (inverse) Mach number arises naturally.

While here, let us quickly derive the classic result of Charney and Stern (1962). Multiply (3.38) by $\mathcal{M}^*\mathcal{M}_{\tilde{b}\tilde{b}}/(c-u)$, where $\mathcal{M}^*$ is the complex conjugate of $\mathcal{M}$, then integrate vertically using integration by parts to obtain

$$[\mathcal{M}_b\mathcal{M}^*]_{b_b}^{b_s} - \int_{b_b}^{b_s} |\mathcal{M}_{\tilde{b}}|^2 \mathrm{d}b = \int_{b_b}^{b_s} \frac{Q_{\tilde{y}}}{(c-u)Q^2} |\mathcal{M}|^2 \mathrm{d}b. \tag{3.44}$$

The first term vanishes (regardless of which bottom boundary condition is used), while the second term is purely real. Thus the imaginary part of the RHS must be zero, whence

$$-c_i \int_{b_b}^{b_s} \frac{Q_{\tilde{y}}}{|c-u|^2 Q^2} |\mathcal{M}|^2 \mathrm{d}b = 0. \tag{3.45}$$

Instability requires $c_i \neq 0$, which requires the integral to be zero, which is possible only if $Q_{\tilde{y}}$ changes sign somewhere in the fluid column. Whereas the quasi-geostrophic equations in depth coordinates have four necessary conditions for instability (Vallis, 2006), the planetary geostrophic equations in buoyancy coordinates have reduced this to a single necessary condition. Here, it is implicitly assumed that the buoyancy is uniform at the upper and lower layers, with massless layers as required. This is equivalent to the approach of Bretherton (1966), introducing delta-sheets of PV in the quasi-geostrophic problem.

3.A.3 Comparison of methods

Figure 3.12 compares the Rossby wave phase speed for various methods. Panels (a)–(d) show the real part of the eigenvalue with largest imaginary part—the phase speed of the most unstable mode—for the four permutations of depth vs buoyancy as vertical coordinate and zero vertical flow vs zero horizontal flow as the bottom boundary condition. The results differ considerably, both on where an unstable solution is found, and its phase speed where it exists. The buoyancy, where used as the vertical coordinate, is $-g\rho_c^{-1}\sigma_\nu$, using the carefully defined orthobaric density σ_ν, so we should not hope a different definition of buoyancy would make the results agree. Panels (e) and (f) use the methods of this chapter, using either Q or B as the dependent variable. These methods act on a single layer to estimate the reference frame shift appropriate for stationary Rossby waves, and get away from the vertical modes. Thus we should

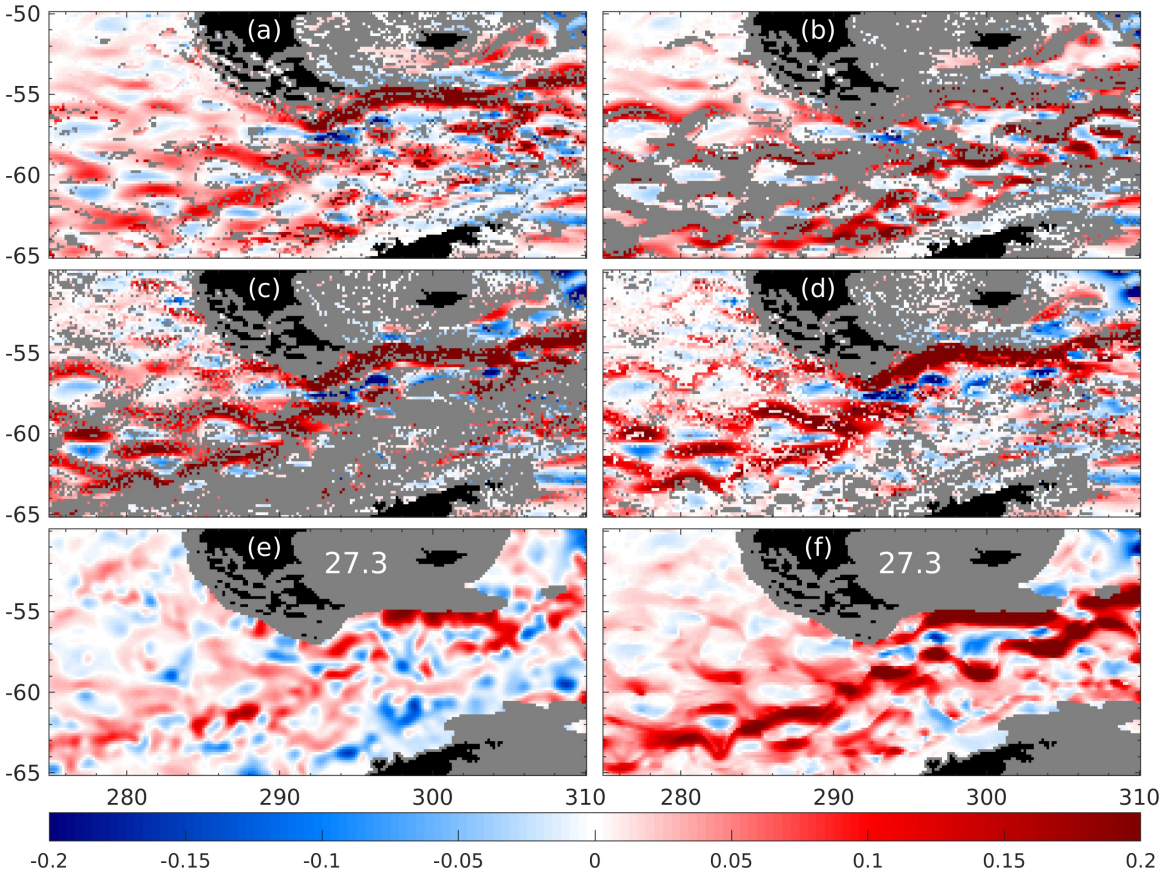


Figure 3.12: Comparison of the Rossby wave phase speed [m/s] around Drake Passage on 22–26 March 2005. The eigenvalue with the largest imaginary part (corresponding to the most unstable mode) is selected using depth (a,b) and orthobaric density (c,d) as the vertical coordinate. The bottom boundary condition is Neumann (b,c) or Dirichlet (a,d) corresponding to zero vertical flow (a,c) or zero horizontal flow (b,d). Each pixel is a water column, and grey pixels indicate all eigenvalues are real. Panels (e) and (f) use local estimates of the reference frame shift α on the orthobaric surface $\sigma_\nu = 27.3$, with Q as the dependent variable, (3.20), in (e), or B as the dependent variable, (3.29) in (f).

not expect any of panels (a)–(d) to match panels (e) or (f), because the former, as vertical modes, are influenced by the ocean state above and below the local depth with which the latter are concerned. Moreover, the most unstable vertical mode has been selected where it exists, but the new methods have no corresponding set of things to select from. Nonetheless, the Rossby wave phase speed from the new method with B as the dependent variable (panel (f)) agrees reasonably well with the vertical mode results. Using Q as the dependent variable produces less agreeable results, though still in the right ball-park. This rough agreement lends support to the methods of this chapter.

Finally, note that the vertical structure of Rossby waves should be independent of the choice of vertical coordinate. Indeed, we can relate (3.38) and (3.33) as follows. Multiply (3.33) by $f^2\sigma$ and take a vertical derivative to get

$$(c-u)f^2(\sigma\mathcal{W}_{zz})_z + (\sigma u_z)_z f^2\mathcal{W}_z - \beta\mathcal{W}_z = 0. \tag{3.46}$$

Depth and buoyancy derivatives are related by $\partial_{\tilde{b}} = \sigma\partial_z$. Supposing that $\mathcal{M} \propto \mathcal{W}_z$, this becomes

$$(c-u)f^2\sigma^{-1}\mathcal{M}_{\tilde{b}\tilde{b}} + (\sigma^{-1}u_{\tilde{b}\tilde{b}}f^2 - \beta)\mathcal{M} = 0. \tag{3.47}$$

Multiplying the above by σ/f^2 recovers (3.38) exactly, noting that

$$(Q^{-1})_{\tilde{y}} = -M_{\tilde{b}\tilde{b}\tilde{y}}/f - \beta\sigma/f^2 = u_{\tilde{b}\tilde{b}} - \beta\sigma/f^2. \tag{3.48}$$

Hence,

$$\frac{\mathcal{M}}{\mathcal{M}_{\tilde{b}\tilde{b}}} = \frac{\mathcal{W}}{\sigma(\sigma\mathcal{W}_{zz})_z}, \tag{3.49}$$

which may be useful if one solves the eigenvalue problem in depth coordinates but wishes to use the clean expression (3.38).

Chapter 4

Non-conservation of potential vorticity in the Southern Ocean

4.1 Introduction

The previous chapter explored variations of potential vorticity on isopycnals in the Southern Ocean, working on the assumption that these variations are mostly aligned with variations in the Bernoulli potential. That is, the ideal theory $Q = Q(b, B)$ holds. In this chapter I'll explore the validity of this claim in the Southern Ocean, in the time-mean. The Southern Ocean is the most interesting place to test this claim: only here are potential vorticity contours are not blocked by continents, allowing for the possibility of potential vorticity conserving circumpolar flow. However, as previously shown by Gille (1997), potential vorticity is not conserved along streamlines in the Southern Ocean, as seen both in numerical models and observations. There could be many reasons for this, including wind forcing, friction, and transient eddy-mean flow interaction, many of which might be intensified by localized interaction with topographic obstacles.

The Antarctic Circumpolar Current (ACC) deflects poleward steadily as the fluid circumnavigates Antarctica, returning equatorward abruptly after Drake Passage. This behaviour is consistent with Ekman-upwelling over the Southern Ocean stretching fluid columns, causing them to drift poleward. This is encapsulated by "Sverdrup balance" (Sverdrup, 1947; Stommel, 1957), which is an integrated and specialized form of the more general linear vorticity balance.

However, subsequent advances in understanding the residual overturning of the Southern Ocean (Danabasoglu et al., 1994; Gent et al., 1995; Marshall, 1997; Marshall and Radko, 2003) indicate that eddy-induced downwelling largely compensates the Ekman-upwelling of the Eulerian-mean Deacon cell. This eddy-induced stretching

of fluid parcels must also play a role in the linear vorticity balance: Marshall et al. (2016) give a simple illustration of this, showing that the meridional deflection of the circumpolar current is associated with both eddy- and Ekman-induced stretching.

The linear vorticity budget is closely related to the large-scale potential vorticity budget: the latter may be obtained by vertically averaging the linear vorticity budget between two nearby buoyancy surfaces[1], and then dividing by the vertical separation of these surfaces. More generally, the same procedure applied to the full vorticity budget yields the full potential vorticity budget (including relative vorticity). Given that eddies are a significant forcing in the linear vorticity budget, it therefore seems inevitable that eddies significantly force the potential vorticity, which is the main focus of this chapter.

There have been some previous studies of potential vorticity conservation in the Southern Ocean, most notably by Gille (1997). While I agree with many of the findings on a qualitative level, there are some issues with the theoretical foundation of this study, which have become clear as the theory of eddy-mean flow interaction has developed over the past two decades. Indeed, a robust theoretical foundation, namely the Thickness-Weighted Average (TWA) formalism of Young (2012), was many years away.

To get the PV equation, Gille (1997) interpolates Eulerian time-averaged isobaric model data onto isopycnal surfaces[2]. This unweighted time-average leaves very large eddy terms in her potential density equation, and these terms carry through to affect the time-mean PV equation. Her time-mean PV is defined as

$$\overline{Q} = \overline{(f+\zeta)\frac{\partial \rho_\sigma}{\partial z}} \tag{4.1}$$

where $\zeta = \partial_x v - \partial_y u$ is the vertical component of the relative vorticity and ρ_σ is the potential density.

This differs from the PV in the TWA formalism, which I now heuristically describe. A mathematical and self-contained description is given in Appendix 4.B; see Young (2012) for full details. The TWA formalism is essentially three steps: (1) the equations

[1]Here I use "buoyancy surface" to refer to an approximately neutral surface. Later I will precisely define the buoyancy field, but for the present discussion let us simply assume that a fluid parcel's buoyancy is materially conserved.

[2] Gille states that means, variances, and covariances of model fields were time-averaged online, which would be on isobaric levels. However, the PV is not (usually) a model field, nor would I expect $(\partial_x v - \partial_y u)\partial_z \rho_\sigma$ to be time-averaged and archived. Neither would I expect such statistics include the specific volume anomaly $\delta(z)$ necessary for her calculation of the Montgomery potential, since $\delta(z)$ depends on a reference salt and potential temperature that must be chosen particularly for the region of study.

are recast with buoyancy as the vertical coordinate; (2) a thickness-weighted time average is taken; and (3) the results are recast back to depth coordinates. The third step can be omitted if, as here, one's analysis is on a buoyancy surface.

Step (1): Replacing depth as the vertical coordinate with buoyancy creates a new coordinate system, which is curvilinear and non-orthogonal. It is described by coordinates $(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t})$. The coordinate transformation from Cartesian (x, y, z, t) is simply $\tilde{x} = x$, $\tilde{y} = y$, $\tilde{b} = b(x, y, z, t)$, and $\tilde{t} = t$, assuming that the buoyancy b is a monotonic function of z. For present purposes, imagine $b = -g\rho_c^{-1}\rho_\sigma$; this will be discussed carefully later. A buoyancy surface is described in two ways: as an isosurface of $b(x, y, z, t)$ in Cartesian coordinates, or as an isosurface of the "trivial" function $B(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t}) = \tilde{b}$ in curvilinear coordinates. Buoyancy surfaces are physical objects existing independently of a coordinate representation, and thus these two representations should be equivalent (Young, 2012). Note in $(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t})$ coordinates, the buoyancy $\tilde{b}$ is an independent variable while the depth ξ is a dependent variable. (Young (2012) uses ζ for depth, but I will retain ζ for the vertical component of the relative vorticity.)

Step (2): The layer thickness is defined as $\sigma = \xi_{\tilde{b}}$, and this is equivalent to $\sigma = b_z^{-1}$ in z coordinates. The thickness-weighted average of a scalar field F is defined

$$\hat{F} = \overline{F\sigma}/\overline{\sigma}, \tag{4.2}$$

where the overbar is a time-mean at constant buoyancy. The TWA potential vorticity is the TWA absolute vorticity divided by the mean layer thickness:

$$\Pi^{\#} = \frac{f + \hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}}}{\overline{\sigma}}, \tag{4.3}$$

where $\hat{u}$ is the TWA zonal velocity, $\hat{v}$ the TWA meridional velocity, and the subscripts with tildes indicate partial derivatives taken at constant buoyancy. This TWA PV is conserved following the TWA velocity, aside from buoyancy and mechanical forcing and eddying effects. This TWA PV differs from the unweighted Eulerian mean PV of (4.1) because (a) the averaging is quasi-Lagrangian, tracking the vertical heaving of the buoyancy surface (vertically Lagrangian but horizontally Eulerian), (b) the averaging is weighted by the layer thickness, and (c) the horizontal components of the relative vorticity are not neglected ($\hat{u}_z$ is embodied within $\hat{u}_{\tilde{x}}$ every bit as much as $\hat{u}_x$ is, and same for $\hat{v}$).

Since the thickness-weighted averaging eliminates eddy terms from the thickness equation and is important to obtain the residual circulation of the Southern Ocean, the TWA PV (4.3) may exhibit better material conservation than (4.1). This suggests

the question of how robust Gille (1997)'s results are. To what degree are her results of sensitive to the definition of PV? This question comes in two parts:

1. What is the structure of PV along streamlines?

2. What forces the PV to vary along streamlines?

Results from Gille (1997) for Question (1) are as follows (see Fig. 4.1, top panel, which reproduces Gille's Figure 7):

- Mean PV is not conserved along mean streamlines at 1 km depth.
- At Drake Passage there are huge fluctuations in the PV, dropping from 4.2 to 3.0 (units $1 \times 10^{-8}\,\mathrm{kg\,s^{-1}\,m^{-4}}$) on one particular streamline.
- PV is well conserved from about 0°E to 150°E.
- PV is gained from about 180°W to 60°W, between Campbell Plateau and Drake Passage. Gille attributes this to the streamline being further south in this sector and therefore beneath a negative wind stress curl (in (4.1) the stratification term is negative, so PV is positive and more negative $f + \zeta$ increases PV).
- Stratification changes and planetary vorticity changes constitute most of the changes in PV. Relative vorticity makes small contributions on scales $< 10°$ longitude, and is insignificant on larger scales.

Regarding Question (2), Gille makes the following claims about the PV budget along streamlines (see Fig. 4.1, bottom panel):

- PV variations are forced primarily by transient eddy tracer fluxes, transient eddy momentum fluxes, tracer evolution, tracer diffusion, momentum diffusion.
- Diapycnal fluxes are negligible, as are the relative vorticity tendency and ageostrophic advection.

However, I note the following:

- The solenoidal term (see (1.7) or (4.20)) is not small, suggesting potential density is a poor choice.
- The discretization error is not small, suggesting issues with the numerics.

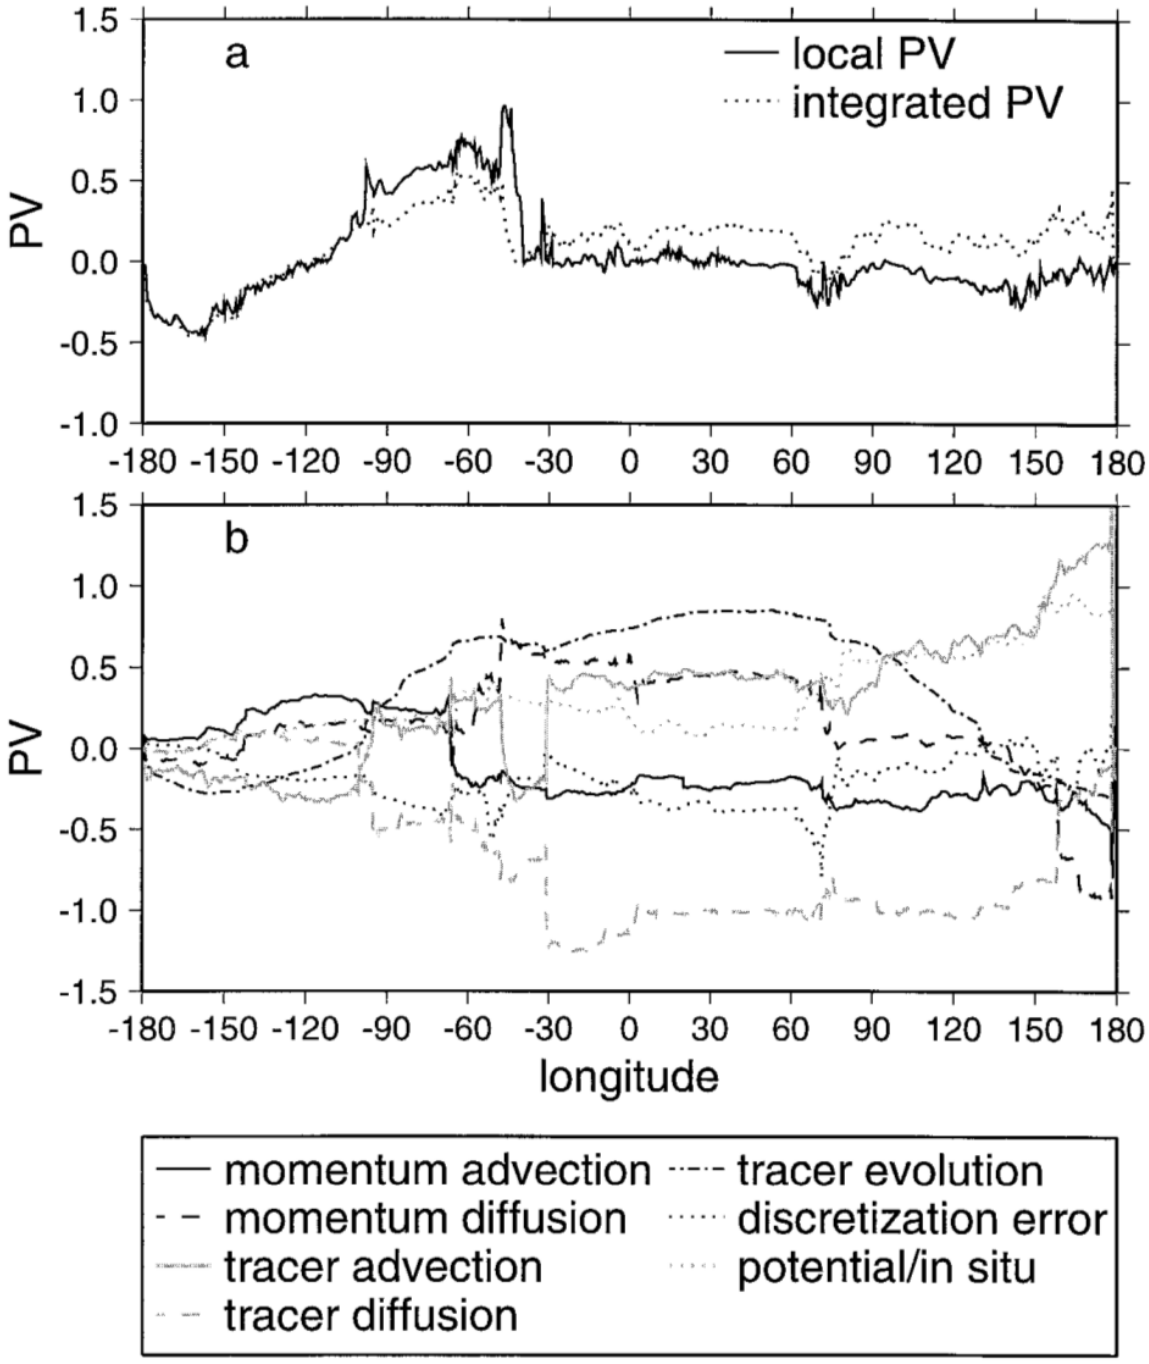


Figure 4.1: Figure 7 from Gille (1997). "Forcing terms responsible for changes in PV... The momentum advection and tracer advection terms refer to transient eddy fluxes, while the potential/in situ term results from potential and in situ density surfaces not coinciding, so that $[(\partial \bar{\boldsymbol{u}}/\partial z) \cdot \nabla]\bar{\sigma}$ is not zero. All terms are set to 0 at 180°W. Together tracer advection, tracer evolution, and momentum advection capture the leading variations in PV, while momentum and tracer diffusion play significant but smaller roles." Here σ is potential density referenced to 1000 dbar. PV is in $1 \times 10^{-8}\,\mathrm{kg\,s^{-1}m^{-4}}$. The streamline has an average depth near 1 km, and passes north of Kerguelen and deflects strongly northward downstream of Drake Passage.

- Diffusion plays a surprisingly large role at all longitudes (i.e. not just near topography), suggesting a higher resolution model may be necessary. Gille analysed the global Semtner-Chervin model with 0.4° zonal resolution, 0.4° cos(`lat`) meridional resolution, and 20 vertical levels.

I address Question (1) in Section 4.3, where the streamwise structure of PV is compared between Eulerian and TWA formalisms. Unfortunately, answering question (2) using TWA is beyond the scope of this thesis, since it ultimately would require online computation of the buoyancy coordinate momentum and thickness budgets. Using 5-day Eulerian averaged model fields, I have been unable to simultaneously close the TWA momentum *and* thickness budgets, meaning the PV budget cannot be closed. It is entirely possible to close one using offline data, but I've found it "opens" the other. The trouble stems from the thickness equation in buoyancy coordinates (before averaging), which combines the buoyancy budget and volume conservation as one. Both of these are incredibly fine balances (the sum of large terms is either tiny (buoyancy forcing) or zero (volume conservation)). Several recent studies have claimed to successfully use the TWA formalism; I believe they have fallen into this trap. I will therefore describe in Appendix 4.B the difficulties as I see them; this is a "negative result", but I hope it will be useful in furthering practical TWA analyses.

I address Question (2) in Section 4.4, using (somewhat reluctantly) Eulerian time-averages to study the PV budget. Only in this framework can I simultaneously close the momentum, buoyancy, and volume budgets, and hence the PV budget. I define PV differently to Gille (1997) however, capturing some of the thickness-weighted aspect (but not all). This is not perfect, but it is the best that can be done here, and still has plenty to teach us.

First though, I present in Section 4.2 the theoretical background for the linear vorticity and potential vorticity budgets in a pedagogical way using common parlance of Eulerian and eddy-induced velocities. The structure of PV along streamlines is explored in Section 4.3, while its forcing along streamlines is explored in Section 4.4. Conclusions are found in Section 4.5.

As before, all data in this chapter is from the Southern Ocean State Estimate (Mazloff et al., 2010).

4.2 Theoretical background

Eddies play a key role in ocean dynamics through eddy form stress, eddy buoyancy fluxes, eddy-induced velocity, and PV diffusion. But all of these are different mani-

festations of the same physical process, thus there is great scope for double-counting the effects of eddies. This confusion is evidenced by the discussion in Warren et al. (1996), and subsequent replies by Hughes (1997) and Olbers (1998). Consequently, one needs to derive the averaged equations carefully and systematically. Here I start with a pedagogical derivation of the linear vorticity budget (Section 4.2.1), in which these different eddy contributions are explicit. Standard approximations are then made to bring the linear vorticity budget into a more commonly recognized form (Section 4.2.2). I then link the linear vorticity budget to the large-scale potential vorticity budget by dividing by the mean layer thickness (Section 4.2.3).

4.2.1 Linear vorticity balance

In this pedagogical section, I will show how the linear vorticity balance for the instantaneous flow gives rise to a linear *residual* vorticity balance. By residual, I mean that it involves a velocity field that is predominantly aligned with mean buoyancy surfaces. This could be the residual velocity from the TWA, but then it would not immediately have separate mean and eddy-induced components (and indeed this is a strength of the TWA framework). Instead, and to connect with common oceanographic parlance, I use the residual velocity of Gent et al. (1995); this involves both an Eulerian time-mean and an eddy-induced linear vorticity balance, and their sum is the residual linear vorticity balance.

Let $\boldsymbol{u} = u\hat{\boldsymbol{i}} + v\hat{\boldsymbol{j}}$ be the horizontal velocity and let $\boldsymbol{v} = \boldsymbol{u} + w\hat{\boldsymbol{k}}$ be the 3D velocity, with $(\hat{\boldsymbol{i}}, \hat{\boldsymbol{j}}, \hat{\boldsymbol{k}})$ the orthonormal basis vectors pointing eastward, northward, and upward. Denote $\nabla_z \alpha \equiv \partial_x \alpha|_z \hat{\boldsymbol{i}} + \partial_y \alpha|_z \hat{\boldsymbol{j}}$ as the 2D gradient holding z constant. An overbar denotes an Eulerian time-mean (taken at a fixed point in space). The Boussinesq assumption is made from the outset for this section.

The continuity equation $\nabla \cdot \boldsymbol{v} = 0$ is multiplied by the Coriolis parameter f and re-arranged to obtain the advection of planetary vorticity:

$$\beta v = f \partial_z w + \nabla \cdot (f \boldsymbol{u}_{ag}), \tag{4.4}$$

where $\boldsymbol{u}_{ag} = \boldsymbol{u} - \boldsymbol{u}_g$ is the ageostrophic velocity, with $\boldsymbol{u}_g = (\rho_c f)^{-1} \hat{\boldsymbol{k}} \times \nabla_z p$ the geostrophic velocity having $\nabla \cdot (f \boldsymbol{u}_g) = 0$. The Boussinesq reference density is ρ_c. This, (4.4), is the linear vorticity balance without approximation.

There are no eddy terms that arise when time-averaging (4.4), so we readily have

$$\beta \overline{v} = f \partial_z \overline{w} + \nabla \cdot (f \overline{\boldsymbol{u}_{ag}}). \tag{4.5}$$

Now, suppose there is an eddy-induced velocity $\boldsymbol{v}^*$. If it is non-divergent, as it is commonly defined to be, the same treatment yields the "linear eddy vorticity balance",

$$\beta v^* = f\partial_z w^* + \nabla \cdot (f\boldsymbol{u}^*). \tag{4.6}$$

Adding (4.5) and (4.6) gives the linear residual (mean plus eddy) vorticity balance,

$$\beta(\overline{v} + v^*) = f\partial_z \left(\overline{w} + w^*\right) + \nabla \cdot \left(f\overline{\boldsymbol{u}_{ag}} + f\boldsymbol{u}^*\right). \tag{4.7}$$

This derivation has stressed that the linear vorticity balance for the residual flow, which has been postulated as the relevant balance for the Southern Ocean (Marshall et al., 2016), is in fact two separate equations rolled into one. The mean flow must satisfy its own linear vorticity balance independently of the eddying flow, and vice-versa.

However, interpreting the mean or the eddy linear vorticity balance is somewhat trying, because the mean velocity and the eddy-induced velocity both pierce through mean buoyancy surfaces. We have a much better intuition for the residual velocity, however, because it is aligned with mean buoyancy surfaces. This is particularly helpful in moving to the potential vorticity equation, which shall be obtained by manipulating (4.7). But first, the mean and eddy linear vorticity budgets shall be cast into a more usable form.

4.2.2 Approximate linear vorticity balance

Let us write the horizontal momentum equation as

$$\rho_c \hat{\boldsymbol{k}} \times f\boldsymbol{u} = -\nabla_z p + \boldsymbol{\tau}_z \tag{4.8}$$

where $\boldsymbol{\tau}$ is the vertical stress. I have ignored the lateral viscous stresses, non-linear terms, and time evolution; if desired, these can all be carried alongside $\boldsymbol{\tau}_z$. Hence the ageostrophic velocity is simply $f\boldsymbol{u}_{ag} = -\rho_c^{-1}\hat{\boldsymbol{k}} \times \boldsymbol{\tau}_z$, and thus the approximated linear vorticity balance is

$$\beta v = f\partial_z w + \rho_c^{-1}\hat{\boldsymbol{k}} \cdot \nabla \times \boldsymbol{\tau}_z. \tag{4.9}$$

Simply add overbars to obtain the approximated mean linear vorticity balance.

While here, let us very quickly derive Sverdrup balance. Integrating (4.9) from the sea-floor at $z = -H$ to the sea-surface at $z = 0$ (a rigid lid is assumed for simplicity, but a free-surface is easily used) yields

$$\beta V = -f w_b + \rho_c^{-1}\hat{\boldsymbol{k}} \cdot \nabla \times (\boldsymbol{\tau}_s - \boldsymbol{\tau}_b) - \rho_c^{-1}(\partial_z \boldsymbol{\tau})_b \times \nabla H \tag{4.10}$$

where $V = \int_{-H}^{0} v \, \mathrm{d}z$ is the depth-integrated flow, b subscripts indicate the variable is evaluated at the sea-floor, and similarly s subscripts indicate evaluation at $z = 0$. Sverdrup balance, namely

$$\beta V = \rho_c^{-1} \hat{\boldsymbol{k}} \cdot \nabla \times \boldsymbol{\tau}_s, \tag{4.11}$$

is a special case of the barotropic vorticity equation (4.10), that neglects all but the surface wind stress.

Turning to the eddies, the Gent et al. (1995) eddy-parametrisation defines the horizontal eddy-induced velocity as

$$\boldsymbol{u}^* = -\frac{\partial}{\partial z}\left(\frac{\overline{\boldsymbol{u}' b'}}{N^2}\right), \tag{4.12}$$

where b is the buoyancy (assumed to be almost materially conserved), $N^2 = \partial_z \overline{b}$ is the mean stratification, and b' and $\boldsymbol{u}'$ are fluctuations from their means $\overline{b}$ and $\overline{\boldsymbol{u}}$. Making the standard assumptions that (a) fluctuations in the height of an isopycnal z' are related to fluctuations in local buoyancy b' by $z' = -b'/N^2$, and (b) fluctuations of the horizontal velocity $\boldsymbol{u}'$ are geostrophic, $f\boldsymbol{u}' = \rho_c^{-1} \hat{\boldsymbol{k}} \times \nabla_z p'$, then (4.12) yields

$$f\boldsymbol{u}^* = \rho_c^{-1} \partial_z \left(\hat{\boldsymbol{k}} \times \overline{z' \nabla_z p'} \right) = -\rho_c^{-1} \hat{\boldsymbol{k}} \times \boldsymbol{\tau}_z^* \tag{4.13}$$

where $\boldsymbol{\tau}^* = -\overline{z' \nabla_z p'}$ is the eddy interfacial form stress. Thus,

$$\nabla \cdot f\boldsymbol{u}^* = \rho_c^{-1} \hat{\boldsymbol{k}} \cdot \nabla \times \boldsymbol{\tau}_z^*. \tag{4.14}$$

A more careful derivation, using the TRM-II theory of McDougall and McIntosh (2001) and not making these assumptions, is given in Appendix 4.C, but the results are similar. Putting (4.14) into (4.6) yields the approximated linear eddy vorticity balance,

$$\beta v^* = f \partial_z w^* + \rho_c^{-1} \hat{\boldsymbol{k}} \cdot \nabla \times \boldsymbol{\tau}_z^*. \tag{4.15}$$

Finally, adding the approximated linear vorticity budgets, i.e. the mean of (4.9) plus (4.15), gives the approximated residual linear vorticity budget,

$$\beta(\overline{v} + v^*) = f \partial_z \left(\overline{w} + w^* \right) + \rho_c^{-1} \hat{\boldsymbol{k}} \cdot \nabla \times \left(\boldsymbol{\tau}_z^* + \overline{\boldsymbol{\tau}}_z \right). \tag{4.16}$$

4.2.3 Potential vorticity balance

The linear vorticity balance is an equation giving the material derivative of planetary vorticity f, which takes a simple form because f is dependent only on latitude. We now turn it into an equation for the material derivative of the large-scale PV, $f\overline{b}_z$,

by multiplying (4.7) by $\overline{b}_z$. This will eliminate the vortex stretching term $f\partial_z w$ that is a large forcing on the vorticity. At this point I introduce the dagger to notate the residual velocity: $\boldsymbol{v}^\dagger = \overline{\boldsymbol{v}} + \boldsymbol{v}^*$. The term $\overline{b}_z \beta v^\dagger$ is returned to its advective form $\overline{b}_z \boldsymbol{v}^\dagger \cdot \nabla f$ and manipulated through the product rule to yield

$$\boldsymbol{v}^\dagger \cdot \nabla(f\overline{b}_z) - \left(f\boldsymbol{v}^\dagger \cdot \nabla\overline{b}\right)_z + f\boldsymbol{v}^\dagger_z \cdot \nabla\overline{b} = fw^\dagger_z \overline{b}_z + \overline{b}_z \nabla \cdot (f\boldsymbol{u}^\dagger). \tag{4.17}$$

The first term on the RHS cancels part of the last term on the LHS, thus eliminating the vortex stretching term. Now, the whole point of using the residual velocity is that the mean buoyancy is nearly conserved following the residual, but not the mean nor the eddy-induced, velocity:

$$\overline{b}_t + \boldsymbol{v}^\dagger \cdot \nabla\overline{b} = \dot{b}, \tag{4.18}$$

where $\dot{b}$ is the divergence of eddy buoyancy fluxes across time-mean buoyancy surfaces, plus sources of buoyancy associated with heat or salt sources. This residual-mean buoyancy equation is used to replace the second term in (4.17). Thus,

$$\begin{aligned} \boldsymbol{v}^\dagger \cdot \nabla(f\overline{b}_z) - \left(f\dot{\overline{b}} - f\overline{b}_t\right)_z &= -f\boldsymbol{u}^\dagger_z \cdot \nabla\overline{b} + \overline{b}_z \nabla \cdot (f\boldsymbol{u}^\dagger) \\ &= \nabla\overline{b} \cdot \nabla \times (\hat{\boldsymbol{k}} \times f\boldsymbol{u}^\dagger). \end{aligned} \tag{4.19}$$

Up to here there are no approximations beyond the Boussinesq assumption. The LHS contains all the ingredients for the material derivative of the large scale PV. We now break the RHS of (4.19) into mean and eddy parts and apply approximations as in Section 4.2.2. Assuming hydrostatic balance, $\overline{p}_z = -g\overline{\rho}$, and (4.8), then

$$\nabla\overline{b} \cdot \nabla \times (\hat{\boldsymbol{k}} \times f\overline{\boldsymbol{u}}) = g\rho_c^{-1}[\overline{\rho}_x \overline{b}_y - \overline{\rho}_y \overline{b}_x] + \nabla\overline{b} \cdot \rho_c^{-1} \nabla \times \boldsymbol{\tau}_z. \tag{4.20}$$

Taking b as a linear combination of ρ and a function of z (for example, $b = -g\rho_c^{-1}[\rho - \rho(S_0, \theta_0, z)]$ for constant salinity S_0 and constant potential temperature θ_0; this is a good choice and is elaborated in Section 4.3.1.1) annihilates the term in brackets on the RHS above. The eddy part is simplified using (4.13) to give

$$\nabla \times (\hat{\boldsymbol{k}} \times f\boldsymbol{u}^*) \cdot \nabla\overline{b} = \nabla\overline{b} \cdot \rho_c^{-1} \nabla \times \boldsymbol{\tau}^*_z. \tag{4.21}$$

Putting (4.20) and (4.21) into (4.19) yields the large-scale PV equation,

$$(f\overline{b}_z)_t + \boldsymbol{v}^\dagger \cdot \nabla(f\overline{b}_z) = f\dot{b}_z + \nabla\overline{b} \cdot \rho_c^{-1} \nabla \times (\overline{\boldsymbol{\tau}}_z + \boldsymbol{\tau}^*_z). \tag{4.22}$$

On the LHS is the material derivative of the large-scale PV, and on the RHS are the PV forcings. The same structure of forcing in the Ertel PV equation (1.7) is easily

recognized: vorticity dotted with the gradient of the tracer forcing, and the curl of the momentum forcing dotted with the gradient of the tracer; we lost the solenoidal term by our assumption on the form of b.

This derivation has stressed that the large-scale PV budget is merely the linear *residual* vorticity budget, multiplied by the mean stratification $\overline{b}_z$. If we had instead multiplied the non-linear residual vorticity budget by $\overline{b}_z$, we would obtain the budget for the PV that retains relative vorticity (but still ignoring horizontal density gradients).

Another way to say this is that the PV budget is simply the vorticity budget vertically integrated between two isopycnals, in the limit of bringing the separation between those isopycnals to zero. This is made explicit in Appendix 4.C.

Let us hypothesize about the mean PV budget in the deep ACC, using (4.22). I've already assumed the relative vorticity is small compared to f (low Rossby number). I also expect $\dot{b}$ to be small compared to $\overline{b}$ advection, and eddy form stress to be significant (Munk and Palmén, 1951). By taking a long enough time-mean so the PV tendency is negligible, (4.22) retains only two terms: advection of time-mean PV and forcing by eddy form stress. The advection is the sum of the mean geostrophic advection, mean ageostrophic advection, mean diapycnal advection, and eddy-induced advection. The mean geostrophic advection doesn't actively change the mean PV along mean geostrophic streamlines. The balance among the remaining terms is *a priori* unknown, and so to a numerical model we now turn.

4.3 The structure of PV along streamlines

I now address Question (1): What is the climatological structure of potential vorticity along streamlines? Thinking of PV as a materially conserved tracer leads one to imagine PV is uniform along a given streamline. In the ACC this is particularly powerful because streamlines are circumpolar, enabling PV uniformity over a vast region.[3]

This turns out to be far from true. The mean PV, as I will show, makes large deviations on its circumpolar journey, particularly in the lee of major topographic barriers. In Section 4.4 I address why this occurs, but first let us explore the deviations, their size and their location.

[3]This chapter uses "streamlines" as shorthand for "geostrophic streamlines", or later on for "rotational streamlines", rather than curves that are tangent to the velocity field everywhere, which would spiral in or out from Antarctica.

In doing so I also assess three definitions of PV: the TWA PV, a hybrid Eulerian-mean PV to be defined, and the unweighted Eulerian-mean PV used by Gille (1997). They will be calculated on different surfaces, but they will all (nearly) coincide at (133.25°E, 54°S), south of Australia in the core of the ACC. The first surface (for the TWA PV) is chosen to have a depth of $z_0 = -1000\,\mathrm{m}$ here, on 20 September, 2008. The second surface (for the hybrid Eulerian-mean PV) is chosen to (nearly) coincide with this, and the third surface (for the unweighted Eulerian-mean PV) chosen to coincide exactly with the second surface at (x_0, y_0). These PV's and surfaces are now defined, in turn.

4.3.1 Thickness-weighted average PV

4.3.1.1 Defining Buoyancy

Before using the TWA theory, one must define the buoyancy b. Young (2012) is rather silent on this issue, stating only that "it is convenient to write the density as $\rho = \rho_0(1 - g^{-1}b)$". Of course, if ρ were the *in-situ* density and ρ_0 a scalar constant, b surfaces are identical to ρ surfaces, which due to the compressibility of seawater are nearly equal to pressure surfaces, which are nearly equal to z surfaces; this makes little progress. In fact, Young (2012) assumes a linear equation of state, so that buoyancy can simply be proportional to potential temperature. This is more complicated in the salty ocean with its non-linear equation of state.

For the TWA theory to be exact, it is crucial to be able to transform the pressure gradient force in the z coordinate momentum equations into a body force in the buoyancy-coordinate equations. That is, we need $\rho_c^{-1}p_x = m_{\tilde{x}}$ and $\rho_c^{-1}p_y = m_{\tilde{y}}$ for some scalar field m: in other words, an exact geostrophic stream function must exist on our chosen buoyancy surfaces. The competing effects of salinity and heat in the non-linear equation of state of seawater confound the definition of a single density variable superior for all purposes, so oceanographers have crafted many options. Iso-surfaces of potential density (McDougall and Klocker, 2010), iso-surfaces of neutral density (Jackett and McDougall, 1997), and omega surfaces (Klocker et al., 2009) do not posses this property, whereas iso-surfaces of specific volume anomaly (Montgomery, 1937), iso-surfaces of orthobaric density (de Szoeke et al., 2000; de Szoeke and Springer, 2009), and the topobaric surfaces of Chapter 2 do have this property. For this work I use specific volume anomaly surfaces for their ease of calculation (including the partial derivatives with respect to salinity and potential temperature);

because I only study a single circumpolar streamline at a time, these can be made quite close to neutral.

SOSE computes the *in-situ* density using $-\rho_c g z$ as the pressure rather than the actual pressure, which is appropriate for the Boussinesq equations (Young, 2010). This flips the role of pressure and depth, so the correct surface is an *in-situ* density anomaly surface, which I refer to henceforth as buoyancy, defined by[4]

$$b(x,y,z,t) \stackrel{\text{def}}{=} -\frac{g}{\rho_c}\Big[\rho\big(S(x,y,z,t),\theta(x,y,z,t),z\big) - \rho\big(S_0,\theta_0,z\big)\Big]. \tag{4.23}$$

This definition also destroys the solenoidal term in Boussinesq form of the PV equation (1.7), since b is a function of ρ and z only.

I choose a unique S_0 and θ_0 for each surface to study. To do this, I first choose a longitude x_0, latitude y_0, depth z_0, and time t_0. I select an initial guess for S_0 and θ_0 by interpolating in the water column at (x_0, y_0) to the depth z_0, then I use these to calculate the buoyancy b by (4.23). I declare the "current depth" to be $z = z_0$. I get $b_0 = b(x_0, y_0, z)$, which defines a buoyancy surface passing through (x_0, y_0, z). I compute the Montgomery potential m on this b_0 surface and select the contour of m passing through (x_0, y_0), then compute $\bar{z}$, the mean depth along this contour. If $|z_0 - \bar{z}| < 1\,\text{m}$, I declare this procedure done. Otherwise, I update the current depth z to be $z + z_0 - \bar{z}$ and return to the calculation for b_0 above. Thus I have found a b_0 such that the Montgomery streamline emanating from (x_0, y_0) on the surface b takes a mean depth of z_0. Now, I update S_0 to be the mean salinity along this streamline, and likewise for θ_0, then re-run this entire procedure twice more, as the new (S_0, θ_0) modify the surface and streamlines. I choose $t_0 =$ 18-22 September 2008 (deep winter and midway through the SOSE data set), $x_0 = 133.25°\text{E}$ (south of Australia), $y_0 = 54°\text{S}$ (in the core of the ACC), and $z_0 = -1000\,\text{m}$. (Having also studied other surfaces ranging from $z_0 = -250\,\text{m}$ to $z_0 = -2750\,\text{m}$, I will focus on the $z_0 = -1000\,\text{m}$ results.)

The resulting b_0 surfaces are typically within a factor of 2 as close to neutral as a neutral density (γ^n) surface (Jackett and McDougall, 1997), using a root-mean-square metric over the whole ACC region (not shown). Note this statement does not imply that γ^n surfaces are closer to neutral; sometimes this b_0 outperforms γ^n and sometimes not. The nearness to neutrality of the b_0 surface along the streamline for

[4] There are inevitably some overturns, where $z_1 < z_2$ but $b(z_1) > b(z_2)$, mostly in the mixed layer. A simple and effective way to deal with these is simply to sort the b field in the vertical dimension.

which it is defined will be even better; in particular, γ^n suffers because the SOSE climatology differs from the WOCE climatology upon which γ^n were defined.

For reference, the $z_0 = -1000\,\mathrm{m}$ buoyancy surface corresponds to a neutral density of about $\gamma^n = 27.825$, just above North Atlantic Deep Water (which is roughly $\gamma^n = 27.9$–28.1 in the South Atlantic). The time-mean depth of the $z_0 = -1000\,\mathrm{m}$ buoyancy surface is $-1030.6\,\mathrm{m}$ at (x_0, y_0).

4.3.1.2 Calculating PV

Having chosen the reference salinity and potential temperature (S_0, θ_0), the 3D buoyancy field b is defined by (4.23), and the iso-surface of b with value b_0 is our surface upon which the PV analysis proceeds. Recalling the TWA PV definition (4.3), we must interpolate velocities onto this isosurface, and calculate its thickness.

I use piecewise cubic Hermite interpolating polynomials (PCHIP), which preserve monotonicity of the data[5], to get the depth as a function of buoyancy for each water column $(\tilde{x}, \tilde{y})$ at each time $(\tilde{t})$ of SOSE's 5 day archived outputs.

Differentiating this interpolant gives the layer thickness $\sigma = \xi_{\tilde{b}}$, while integration gives the Montgomery potential m since hydrostatic balance is $m_{\tilde{b}} + \xi = 0$. Specifically,

$$m(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t}) = m_5 - \int_{\tilde{b}}^{b_5} \xi(\tilde{b}')\mathrm{d}\tilde{b}' \tag{4.24}$$

where $m_5 = p_5 + b_5(5\,\mathrm{m})$ is the Montgomery potential at an arbitrary depth, that happens to be the top model grid level at 5 m depth, and p_5 and b_5 are the pressure and buoyancy values there.

The zonal velocity in buoyancy coordinates $u(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t})$ is similarly interpolated using PCHIPs from the 4D zonal velocity field onto the depth of the buoyancy surface (longitudinally averaged since SOSE uses a C grid). The meridional velocity v is treated similarly. The mean thickness $\overline{\sigma}$ is not thickness-weighted, nor is the mean Montgomery potential $\overline{m}$. The TWA velocities $\hat{u}$, $\hat{v}$ are computed using (4.2), and averaging $\overline{\sigma}$ onto the U,V grids. The TWA PV is now assembled according to (4.3).

4.3.2 Hybrid Eulerian-mean PV

It is possible to define a PV out of Eulerian-mean quantities that bears some resemblance to the TWA PV. This PV is defined as

$$\check{Q} \equiv \overline{\boldsymbol{\omega}}_a \cdot \nabla \overline{b}. \tag{4.25}$$

[5] I do not use cubic spline interpolation because they often overshoot and undershoot, particularly near the pycnocline.

The check, an upside down hat, reminds us that this is *not* the TWA PV, but is meant to be a bit like it.

Consider that, almost everywhere in the ocean, almost all of the PV exists in the vertical / radial component: $Q = \boldsymbol{\omega}_a \cdot \nabla b \approx (f + \zeta) b_z$. The same is true of $\check{Q}$, so

$$\check{Q} \approx (f + \overline{\zeta}) \overline{b}_z = \overline{Q b_z^{-1}}\ \overline{b}_z. \tag{4.26}$$

Since b_z^{-1} is the layer thickness, this resembles a thickness-weighted average, (4.2), of Q. However, it differs in two ways. First, the time-average here is still performed at a constant depth, rather than following the buoyancy surface as in TWA. Second, (4.2) has $\overline{\sigma} = \overline{b_z^{-1}}$ as a denominator, which is not the same as having $\overline{b_z} = \overline{b}_z$ (the derivative and average commute) as a numerator in (4.26). In fact, $1/\overline{\sigma}$ is the harmonic mean of b_z, whereas in (4.26) we have the arithmetic mean of b_z.

These two points aside, perhaps some good properties of the TWA will be conferred upon this "hybrid" PV. However, even this is unlikely because in the low Rossby number limit where $|\zeta| << |f|$, this "hybrid" PV becomes identical to the unweighted Eulerian-mean PV (4.1). Ease of utility is an advantage of this PV definition: the Eulerian-mean quantities necessary are commonly constructed during the model run. More importantly, closing the budget for this PV simply requires closing the Eulerian-mean momentum and buoyancy budgets (Section 4.4).

I compute b from (4.23) with the same (S_0, θ_0) from Section 4.3.1.1, using S and θ for each 5-day averaged output, then take an Eulerian time-mean for $\overline{b}$. $\overline{\boldsymbol{\omega}}_a$ is constructed using the Eulerian time-mean horizontal velocities. The numerical discretization of the curl, and the dot product in (4.25), is important and described in Appendix 4.A. In particular, $\check{Q}$ (and all terms in its budget) are computed on the upper-south-west corner of the tracer cell centre. The depth of a buoyancy surface (on the horizontal tracer grid) is averaged four ways to get it on the south-west corner of the horizontal tracer grid, onto which $\check{Q}$ is interpolated from the depth of the W grid.

For b_0 corresponding to $z_0 = -1000\,\mathrm{m}$ (as in Section 4.3.1.1), the depth of the b_0 isosurface of $\overline{b}$ at (x_0, y_0) is $-1033.3\,\mathrm{m}$. This difference of $2.7\,\mathrm{m}$ from the surface for TWA is essentially the difference between McDougall and McIntosh (1996) and McDougall and McIntosh (2001): the former took an Eulerian time-mean of the density and operated on isosurfaces of that, whereas the latter defined a modified time-mean density field, the depth of whose isosurface coincides with the time-mean depth of the same isosurface for the instantaneous density field. It is a subtle difference, but important for establishing a self-consistent framework (Young, 2012).

The exact geostrophic stream function for these buoyancy surfaces is the Montgomery potential, slightly modified for the Boussinesq approximation, namely

$$m = \phi - b_0 z \tag{4.27}$$

where z is the depth of the buoyancy surface and hydrostatic balance is used to obtain the pressure anomaly on the buoyancy surface, namely

$$\phi = \rho_c^{-1} p_5 + \int_z^{z_5} b(z') \, \mathrm{d}z' \tag{4.28}$$

where $z_5 = 5\,\mathrm{m}$ is the depth at the centre of the model's top tracer cell, and p_5 is the pressure there. The time-mean Montgomery potential is obtained simply by using the time-mean fields in the above equations. This is a short-cut, because (4.28) contains a non-linearity in the integral: both z and b are time-dependent. However, I have compared the proper time-mean of m, namely that found by averaging m from 5-day mean data over the full 6 years, with the short-cut time-mean m and found the difference to be utterly negligible. Linear interpolation is used to obtain z, and the trapezoidal integration is used for the integral in (4.28).

4.3.3 Unweighted Eulerian-mean PV on potential density surfaces

The unweighted Eulerian-mean PV used by Gille (1997) has already been described: see (4.1). Here I only say that, since the ocean is generally at low Rossby number, I do not expect a great difference between that PV and the "hybrid" PV definition. Thus I compute the unweighted Eulerian-mean PV on a potential density surface as Gille (1997) did, not a buoyancy surface, to simultaneously assess whether that makes any appreciable difference. (It won't, we shall see; but if it did, I would then compute the unweighted Eulerian-mean PV on buoyancy surfaces, and bisect the problem.) I use potential density referenced, as Gille did, to 1000 dbar, and take an isovalue of 32.245834 (she used 32.3); the corresponding isosurface (of the Eulerian mean potential density) coincides in depth with the Eulerian mean buoyancy surface at the chosen longitude and latitude, (x_0, y_0).

The Zhang and Hogg (1992) modification of the Montgomery potential is used as the approximate geostrophic stream function for the potential density surface (as an exact one does not exist). This is essentially as in (4.27) but a constant reference pressure is subtracted.

4.3.4 Results: PV along streamlines

Having computed three variants of the PV on three different surfaces, with corresponding geostrophic stream functions (or approximations thereof), we can now examine the structure of PV along streamlines[6]. The PVs are linearly interpolated onto the surfaces, then bilinearly interpolated onto the streamlines.

To orient ourselves, Figure 4.2 shows the geography of many streamlines for the hybrid Eulerian PV on the $z_0 = -1000\,\mathrm{m}$ surface, as well as the relative difference of $\check{Q}$ along each streamline. Negative values (blue) mean a local reduction in $|\check{Q}|$. Already certain patterns emerge. There is a major drop of about 30% in (absolute) PV downstream of Drake Passage for more northerly streamlines, in agreement with Gille (1997). However, PV is quite well conserved through Drake Passage for southerly streamlines. There is a similar drop in $|\check{Q}|$ after Campbell Plateau for northerly streamlines. Away from these major topographic barriers, $|\check{Q}|$ tends to steadily rise.

This steady rise in $|\check{Q}|$ is seen more clearly in Figure 4.3, which also serves to compare the three definitions of PV along five streamlines crossing 133.25°E at different latitudes (top panel). The streamlines are computed from the geostrophic stream function on the surface appropriate to each PV, and thus diverge slightly from one another away from their mutual starting point along 133.25°E. Most importantly, the same qualitative features are seen for all three PVs, which is good news for the results of Gille (1997) standing the test of time, and good news for my further analysis of the PV budget using the hybrid Eulerian PV $\check{Q}$ rather than the TWA PV $\Pi^{\#}$.

This view reveals the remarkably smooth gain, for northerly streamlines, of (absolute) PV at all longitudes away from Campbell Plateau and Drake Passage. Compare this to Fig. 4.1a: Gille (1997) found essentially all the (absolute) PV gain occurs after Campbell Plateau (~ 180°E) until Drake Passage, and PV is essentially constant in the Atlantic and Indian sectors. This different result cannot be due to differences between potential density and buoyancy; I therefore attribute it to differences between the Semtner-Chervin model and SOSE. Further discussion on this is deferred until Section 4.3.

Another way to study the structure of PV along streamlines is to ask *how* it changes. Figure 4.4 shows $\check{Q}$ along four geostrophic streamlines (called A, B, C, and

[6] I obtain the streamline as follows. Fill bad data (an incropped/outcropped surface) in the stream function with $+\infty$ for the American, Australian, and African continents, with $-\infty$ for Antarctica, and $M(i_0, j_0)$ everywhere else. Contour this filled stream function on the value $M(i_0, j_0)$ and highlight the contour passing through the point (i_0, j_0); the filling process keeps the contours closed. Then contour again on the same value but on the unfilled stream function. The final streamline is all points in the last contour that existed in the highlighted contour.

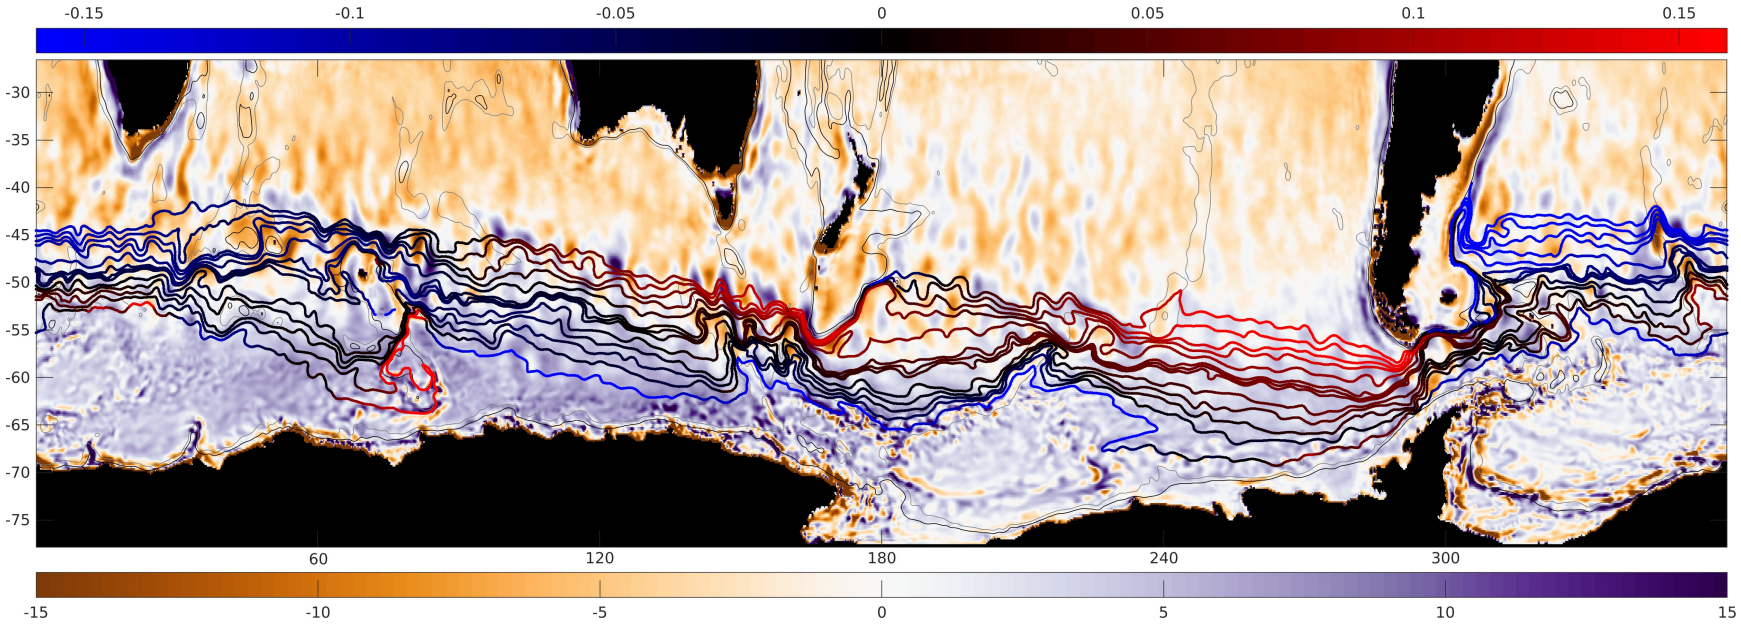


Figure 4.2: Mean vertical Ekman velocity [metres / year, orange/purple] and geostrophic streamlines on the buoyancy surface $\overline{b} = b(133.25°\text{E}, 54°\text{S}, -1000\,\text{m}, 18{-}22\,\texttt{Sept}\ 2008)$. The 13 streamlines shown here pass 133.25°E at a somewhat irregular spacing, crossing at 62, 60.5, 59, 58, 57, 56, 55, 54.5, 53.5, 52.4, 51.5, 50.7, 50°S (rounded). On streamlines, colour (blue/black/red) indicates the relative change of Eulerian mean PV from its mean on each streamline, and the colourbar is saturated. The 2000 m isobath is contoured black, and the 3000 m isobath grey.

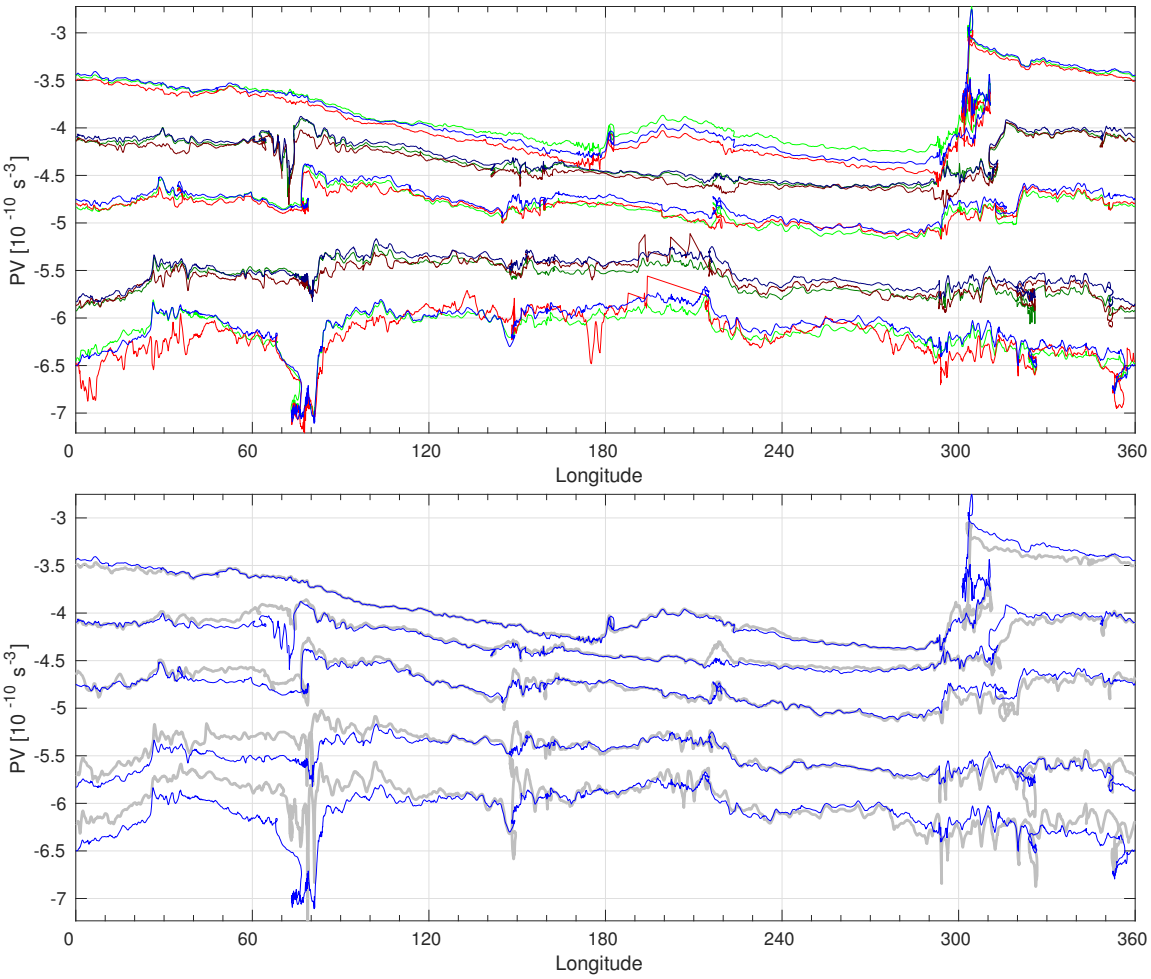


Figure 4.3: Comparison of different PV definitions and streamlines. **Top panel**: The unweighted Eulerian PV (4.1) on potential density surfaces (green), or the TWA PV, (4.3) on buoyancy surfaces (red) or the Eulerian PV (4.25) on buoyancy surfaces (blue). Five geostrophic streamlines are shown, passing 133.25°E at 60.5°S, 59°S, 57°S, 55°S, 52.4°S. The blue curves correspond to the 4th, 7th, 9th, 11th, and 12th geostrophic streamlines, from north to south, in Fig. 4.2. Alternate streamlines are darkened, for clarity. **Bottom panel**: The Eulerian PV along geostrophic streamlines (blue, as in the top panel), or along rotational streamlines (thick grey).

D) as well as its dependence upon each of three components: the planetary vorticity (Coriolis parameter), the relative vorticity, and the stratification. For southerly streamlines, particularly D, impressive compensating swings in the components of PV can be seen. In all cases, relative vorticity makes no contribution on scales larger than about 5° of longitude, but can be sizeable at these scales and smaller. Looking on these scales and larger, for the two northerly streamlines, A and B, $|\check{Q}|$ is gained due to stratification changes from Drake Passage until about Kerguelen at ~ 70°E, while

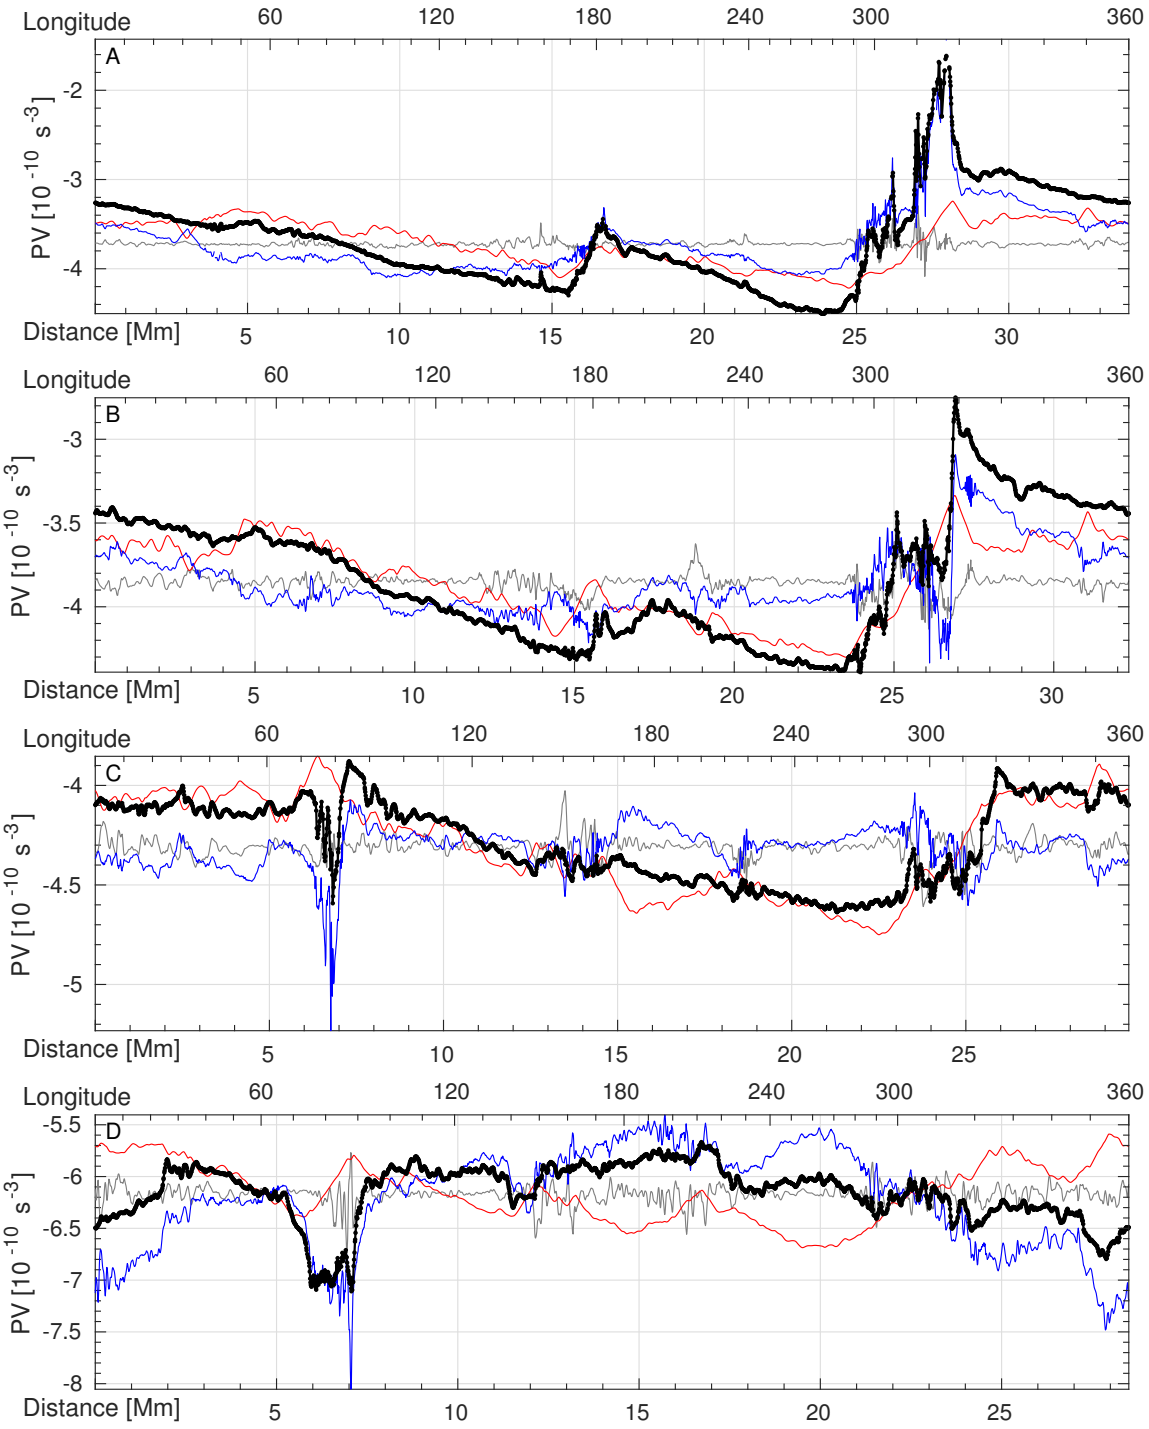


Figure 4.4: PV (thick black) along the (A) 2nd, (B) 4th, (C) 7th, or (D) 12th streamlines in Fig. 4.2 from north to south. They cross 133.25°E at (A) 50.7°S, (B) 52.4°S, (C) 55°S, or (D) 60.5°S. The PV is very well approximated by $(f + v_x - u_y)b_z$, which is decomposed as (red) only f varies, (blue) only b_z varies, or (grey) only $v_x - u_y$ varies (the other terms are set to their mean values along the streamline). The abscissa is distance along the streamline, and the first crossing of various longitudes are indicated on top.

the planetary vorticity contribution is fairly stable, as expected since these streamlines must pass north of Kerguelen and so cannot migrate much south before doing so. After Kerguelen however, $|\check{Q}|$ is gained mostly by changes in planetary vorticity as streamlines slowly migrate southwards to meet Drake Passage again, whereas the stratification component is relatively constant (barring significant changes around Campbell Plateau, which breaks the picture of slowly rising $|\check{Q}|$). Although these northerly streamlines are flirting with regions of Ekman *downwelling* in the Atlantic sector, they still manage to gain (absolute) PV by adjustments in the stratification.

Finally, we return to Fig. 4.3 (bottom panel), which shows $\check{Q}$ along geostrophic streamlines (as in the top panel) and also along "rotational streamlines", i.e. contours of constant Ψ, where Ψ is part of the Helmholtz decomposition.

$$\overline{\boldsymbol{u}} = \nabla\chi + \nabla \times \Psi, \tag{4.29}$$

where $\overline{\boldsymbol{u}}$ is the mean horizontal flow interpolated onto the buoyancy surface. The flow is close to geostrophic in most of the Southern Ocean, so PV along geostrophic streamlines is similar to that along rotational streamlines. There is a noticeable difference in Drake Passage, where rotational streamlines seem to better conserve PV, although this "improvement" seems undone along southerly streamlines which undergo rapid fluctuations along rotational streamlines.

When performing a Helmholtz decomposition, a choice must be made regarding the boundary conditions. I have chosen that both the rotational flow and divergent flow satisfy the same no-normal flow boundary condition that the full flow satisfies. However, this choice is not necessarily the best, particularly in coastal regions of strong upwelling/downwelling (Li et al., 2006). It does help rotational streamlines get through Drake Passage without intersecting topography.

The geostrophic flow does not have this property; the northernmost geostrophic streamline in Fig. 4.2 hits Tierra del Fuego, and re-emerges from South America's continental shelf around 40°S. From a topological perspective, this behaviour is interesting: if the mean PV were conserved following mean geostrophic streamlines, it would be constant on any connected component of a streamline, but could take different values between disconnected components. From this perspective, the drop in (absolute) PV for the most northerly geostrophic streamline that passes uninterrupted through Drake Passage is a consequence of different climatologies for the Pacific and Atlantic oceans (specifically, measured adjacent to the topography where this streamline almost grounds, upstream and downstream).

Henceforth I focus only on geostrophic streamlines. Having discussed *how* mean PV changes along mean streamlines, we now investigate *why* this occurs.

4.4 The forcing of PV along streamlines

This section examines the structure and forcing of the Eulerian time-mean PV in the SOSE. We shall use the full 3D Ertel PV rather than the large-scale PV of the previous pedagogical section. Thus, the derivation is similar to that in Section 1.3.1, but time-averaged. This is done in Section 4.4.1 in the context of the equations SOSE solves. Section 4.4.2 describes how to integrate the PV forcing along streamlines, then Section 4.4.3 analyses results.

4.4.1 Time-mean model equations

Our goal is the prognostic equation for the time-mean 3D Ertel PV, replete with forcing and eddying terms and the same approximations that are made in SOSE, namely the Boussinesq approximation, hydrostatic balance, and the traditional approximation. The "unified" derivation of Section 1.3.1 will help keep this large equation tidy. First we need the momentum and buoyancy equations.

4.4.1.1 Momentum equation

SOSE solves the vector-invariant form of the Boussinesq horizontal momentum equation, and the vertical momentum equation is reduced to hydrostatic balance. Together, these three equations may be written (recall $\boldsymbol{u} = u\hat{\boldsymbol{i}} + v\hat{\boldsymbol{j}}$ is the horizontal velocity)

$$\boldsymbol{u}_t + w\boldsymbol{u}_z + (f+\zeta)\hat{\boldsymbol{k}} \times \boldsymbol{u} + K_x\hat{\boldsymbol{i}} + K_y\hat{\boldsymbol{j}} + \nabla\phi = b_0\hat{\boldsymbol{k}} + \boldsymbol{F}, \tag{4.30}$$

where $\zeta = \partial_x v - \partial_y u$ is the relative vorticity's vertical component, $f = 2\Omega \sin \texttt{lat}$ is the Coriolis parameter, $K = (u^2 + v^2)/2$ is the kinetic energy per unit mass, $\phi = p\rho_0^{-1} + gz$ combines the pressure gradient and gravitational force into a single body force, $b_0 = -g\rho_0^{-1}(\rho - \rho_0)$ is the Boussinesq buoyancy (not to be confused with (4.23)), and $\boldsymbol{F}$ captures the horizontal acceleration from viscosity, drag, and wind stress. To make this amenable to the PV derivation of Section 1.3.1, first note that $(\nabla \times \boldsymbol{u}) \times \boldsymbol{v} = \zeta\hat{\boldsymbol{k}} \times \boldsymbol{u} + w\partial_z\boldsymbol{u} - \partial_z K\hat{\boldsymbol{k}}$, and of course $f\hat{\boldsymbol{k}} \times \boldsymbol{u} = f\hat{\boldsymbol{k}} \times \boldsymbol{v}$. Then in one fell swoop the time-mean of (4.30) becomes

$$\overline{\partial_t \boldsymbol{u}} + \overline{\boldsymbol{\omega}}_a \times \overline{\boldsymbol{v}} + \overline{\boldsymbol{\omega}_a' \times \boldsymbol{v}'} + \nabla\overline{K} + \nabla\overline{\phi} = \overline{b_0}\hat{\boldsymbol{k}} + \overline{\boldsymbol{F}}, \tag{4.31}$$

where the absolute vorticity is $\boldsymbol{\omega}_a = f\hat{\boldsymbol{k}} + \nabla \times \boldsymbol{u}$. This choice, rather than $2\boldsymbol{\Omega} + \nabla \times \boldsymbol{u}$, is justified in Appendix 4.A; essentially, only the former satisfies $\nabla \cdot \overline{\boldsymbol{\omega}}_a = 0$ in the thin-film approximation.

4.4.1.2 Buoyancy equation

SOSE solves prognostic equations for the salinity S and potential temperature θ, rather than the buoyancy b. These are written

$$\partial_t S + \boldsymbol{v} \cdot \nabla S = \dot{S}, \tag{4.32a}$$

$$\partial_t \theta + \boldsymbol{v} \cdot \nabla \theta = \dot{\theta}. \tag{4.32b}$$

Now, the material derivative of buoyancy is

$$\frac{\mathrm{D}b}{\mathrm{D}t} = b_s \frac{\mathrm{D}S}{\mathrm{D}t} + b_\theta \frac{\mathrm{D}\theta}{\mathrm{D}t} + b_z w. \tag{4.33}$$

In the final term, $b_z = \partial_z b(S, \theta, z)$ rather than $\partial_z b(x, y, z)$. However, this term is at least two or three orders of magnitude smaller than even the forcing term ϖ, so I henceforth neglect it[7]. Then the b equation is formed from b_S(4.32a) + b_θ(4.32b),

$$(b_S S_t + b_\theta \theta_t) + (b_S \boldsymbol{v} \cdot \nabla S + b_\theta \boldsymbol{v} \cdot \nabla \theta) = \left(b_S \dot{S} + b_\theta \dot{\theta} \right). \tag{4.34}$$

The haline and thermal expansion terms b_S and b_θ are (weakly) time-dependent, so (4.34) is computed for each of SOSE's 5-day averages, and averaged over 6 years. The salt and heat tendencies and advections are computed using model output data, while the salt and heat forcings are taken as residuals. This yields the time-mean buoyancy equation, written

$$\overline{b_t} + \overline{\boldsymbol{v}} \cdot \nabla \overline{b} + \overline{\boldsymbol{v}' \cdot \nabla b'} = \overline{\varpi} \tag{4.35}$$

where $\overline{b_t} = \overline{b_S S_t + b_\theta \theta_t}$, $\overline{\boldsymbol{v}} \cdot \nabla \overline{b}$ is computed directly from $\overline{\boldsymbol{v}}$ and $\overline{b}$ using the same advection scheme used by SOSE (upwind centred in space and time, third order), $\overline{\varpi} = \overline{b_S \dot{S} + b_\theta \dot{\theta}}$, and the eddy term $\overline{\boldsymbol{v}' \cdot \nabla b'}$ taken as a residual.

[7]This considerably eases data processing if multiple buoyancy surfaces with different (S_0, θ_0) are to be analysed, because b_S and b_θ are independent of S_0 and θ_0. Only b_z depends on S_0 and θ_0.

4.4.1.3 Time-mean PV budget

In the Boussinesq approximation, $\nabla \cdot \boldsymbol{u} = 0$ so the specific volume (ν) factor in the Ertel PV (1.2) can be dropped. Following the derivation in Section 1.3.1, the prognostic equation for $\check{Q}$ is derived from

$$\nabla\overline{b} \cdot \nabla \times (4.31) \; + \; \overline{\boldsymbol{\omega}}_a \cdot \nabla(4.35) \tag{4.36}$$

This yields the following PV equation:

$$\check{Q}_t + \overline{\boldsymbol{v}} \cdot \nabla\check{Q} = -\overline{\boldsymbol{\omega}}_a \cdot \nabla(\overline{\boldsymbol{v}' \cdot \nabla b'}) - \nabla\overline{b} \cdot \nabla \times (\overline{\boldsymbol{\omega}'_a \times \boldsymbol{v}'}) + \nabla\overline{b} \cdot \nabla \times \overline{\boldsymbol{F}} + \overline{\boldsymbol{\omega}}_a \cdot \nabla\overline{\varpi}. \tag{4.37}$$

Thus the time-mean flow conserves the time-mean PV, except for eddy vorticity effects, eddy tracer effects, momentum forcing ($\overline{\boldsymbol{F}}$), and non-conservation of the tracer ($\overline{\varpi}$).

It is worth repeating that

$$\overline{\boldsymbol{v}} \cdot \nabla\check{Q} = \overline{\boldsymbol{\omega}}_a \cdot \nabla(\overline{\boldsymbol{v}} \cdot \nabla\overline{b}) + \nabla\overline{b} \cdot \nabla \times (\overline{\boldsymbol{\omega}}_a \times \overline{\boldsymbol{v}}) \tag{4.38}$$

follows from $\nabla \cdot \overline{\boldsymbol{\omega}}_a = 0$, $\nabla \times \nabla\overline{b} = 0$, $\nabla \cdot \overline{\boldsymbol{v}} = 0$, and the vector triple product $(\overline{\boldsymbol{v}} \cdot \nabla\overline{b})\overline{\boldsymbol{\omega}}_a + (\overline{\boldsymbol{\omega}}_a \times \overline{\boldsymbol{v}}) \times \nabla\overline{b} = (\overline{\boldsymbol{\omega}}_a \cdot \nabla\overline{b})\overline{\boldsymbol{v}}$. The numerical discretization that achieves the first three of these identities is described in Appendix 4.A. This discretization also successfully annihilates the solenoidal term, which should vanish because

$$\nabla b \cdot \nabla \times (b_0\hat{\boldsymbol{k}}) = \nabla b \cdot \nabla b_0 \times \hat{\boldsymbol{k}} = \hat{\boldsymbol{k}} \cdot \nabla b \times \nabla b_0 = 0, \tag{4.39}$$

with the final equality because b is a function of ρ and z alone, and b_0 is a function of ρ alone. Other discretisations, for example based on averaging to the tracer cell centre, fail to eliminate this term, and it can appear to be one of the largest terms in boundary regions where the buoyancy surface slope changes rapidly. When I test this discretization using potential density of buoyancy, the solenoidal term is non-zero but remains utterly small. Thus the large solenoidal term Gille (1997) found is likely due to numerical methods rather than inherent issues of potential density.

4.4.2 Integration of PV forcing along time-mean geostrophic streamlines

Having seen in Section 4.3 that the time-mean PV $\check{Q}$ is not conserved following the time-mean geostrophic flow, we now ask: what causes this?

To answer this, we will integrate the PV equation (4.37) along a time-mean streamline. The change in PV over a distance s along a geostrophic streamline is

$$\check{Q}(s) - \check{Q}(0) = \int_0^s \nabla\check{Q} \cdot \mathrm{d}\mathbf{s} = \int_0^s \overline{\boldsymbol{u}_g} \cdot \nabla_b \check{Q} \, \frac{\mathrm{d}s}{|\overline{\boldsymbol{u}_g}|}, \tag{4.40}$$

where $\boldsymbol{u}_g = \hat{\boldsymbol{k}} \times \nabla M/(\rho_c f)$ is the geostrophic velocity, M the Montgomery potential, and ∇_b is the horizontal gradient operator acting within a buoyancy surface.

Figure 4.5 shows $\check{Q}(s)$ (in black) along four streamlines, and its reconstruction by integrating the geostrophic advection (in faint grey), i.e. the last expression in (4.40). Away from topography the reconstruction is very accurate, capturing both the small scale wiggles of $\check{Q}(s)$ and its large-scale structure. Near topographic obstacles such as Kerguelen Plateau ($\sim$ 70°E), Campbell Plateau ($\sim$ 180°E), and Drake Passage ($\sim$ 300°E), one extreme value of $\check{Q}$ can cause the reconstruction to deviate from the original. After interpolating onto a 2D buoyancy surface, $\check{Q}$ is on the ζ grid, while $\overline{u_g}$ is on the V grid and $\overline{v_g}$ is on the U grid, which is ideal for computing $\overline{\boldsymbol{u}_g} \cdot \nabla_b \check{Q}$, using a 5 point stencil of $\check{Q}$. Say a streamline passes south and east of the central point; the value of $\check{Q}$ on the north and west points need never enter the interpolation for $\check{Q}(s)$, but it does affect $\overline{\boldsymbol{u}_g} \cdot \nabla_b \check{Q}$ at the central point, and hence the reconstruction.

The mean advection of PV is decomposed into advection in the buoyancy surface and vertical advection across the buoyancy surface, using the standard transformation $\nabla_b \Phi = \nabla_z \Phi + \Phi_z \nabla_b z$:

$$\overline{\boldsymbol{v}} \cdot \nabla\check{Q} = \overline{\boldsymbol{u}} \cdot \nabla_b \check{Q} + W \check{Q}_z, \tag{4.41}$$

where I call

$$W = \overline{w} - \overline{\boldsymbol{u}} \cdot \nabla_b z \tag{4.42}$$

the "mean vertipycnal velocity" (MVV); it is the mean vertical velocity minus that required by mean horizontal flow to stay on the mean buoyancy surface. It is a god-awful quantity having no good physical meaning (it is *not* a diapycnal velocity, nor this plus the vertical motion of the buoyancy surface), but it will be a key part of this Eulerian PV budget.

Next, advection in the buoyancy surface is further decomposed into geostrophic and ageostrophic parts:

$$\overline{\boldsymbol{u}} \cdot \nabla_b \check{Q} = \overline{\boldsymbol{u}_g} \cdot \nabla_b \check{Q} + \overline{\boldsymbol{u}_{ag}} \cdot \nabla_b \check{Q}, \tag{4.43}$$

where $\boldsymbol{u} = \boldsymbol{u}_g + \boldsymbol{u}_{ag}$ defines the ageostrophic velocity $\boldsymbol{u}_{ag}$.

Now, integrating the PV budget (4.37) along a streamline, using (4.40), (4.41), (4.43), and re-arranging gives

$$\begin{aligned}\check{Q}(s) - \check{Q}(0) = \int_0^s \Big(&- \overline{\boldsymbol{u}_{ag}} \cdot \nabla_b \check{Q} && - W\check{Q}_z \\ &- \overline{\boldsymbol{\omega}}_a \cdot \nabla \overline{b}_t && - \nabla\overline{b} \cdot \nabla \times \overline{\boldsymbol{u}}_t \\ &- \overline{\boldsymbol{\omega}}_a \cdot \nabla \overline{\boldsymbol{v}' \cdot \nabla b'} && - \nabla\overline{b} \cdot \nabla \times \overline{\boldsymbol{\omega}_a' \times \boldsymbol{v}'} \\ &+ \overline{\boldsymbol{\omega}}_a \cdot \nabla \overline{\varpi} && + \nabla\overline{b} \cdot \nabla \times \overline{\boldsymbol{F}} \Big) \frac{\mathrm{d}s}{|\overline{\boldsymbol{u}_g}|}, \end{aligned} \tag{4.44}$$

having also used $\overline{\boldsymbol{\omega}}_a = f\hat{\boldsymbol{k}} + \nabla \times \overline{\boldsymbol{u}}$ and (4.25) to break up the $\check{Q}_{\tilde{t}}$ term into the third and fourth terms inside the integral.

I could compute the ageostrophic and MVV advection integrals using (4.38), (4.43), and (4.43), yielding

$$W\check{Q}_z + \overline{\boldsymbol{u}_{ag}} \cdot \nabla_b \check{Q} = \overline{\boldsymbol{\omega}}_a \cdot \nabla(\overline{\boldsymbol{v}} \cdot \nabla\overline{b}) + \nabla\overline{b} \cdot \nabla \times (\overline{\boldsymbol{\omega}}_a \times \overline{\boldsymbol{v}}) - \overline{\boldsymbol{u}_g} \cdot \nabla_b \check{Q}. \tag{4.45}$$

This reduces (4.44) to (4.40) once the (curl of the) complete momentum equation (4.31) and the complete buoyancy equation (4.35) are recognized and cancelled. But this merely exposes the $W\check{Q}_z + \overline{\boldsymbol{u}_{ag}} \cdot \nabla_b \check{Q}$ terms to errors from the geostrophic advection, as described above. Since all the other terms have valid data everywhere that $\check{Q}$ has valid data, they do not expand the stencil as did the geostrophic advection term, and thus are more reliable. Hence, I compute the $W\check{Q}_z$ term as a residual of (4.44). I do explicitly integrate the ageostrophic advection term, even though it does expand the stencil and is not as reliable, out of interest about whether the mean advection is mostly in the buoyancy surface, or across it.

The terms from the PV equation (e.g. $\partial_t \check{Q}$) are interpolated onto the streamline exactly as was $\check{Q}$ itself. Also, $\overline{u_g}^2$ and $\overline{v_g}^2$ are interpolated onto the streamline (separately, from the V grid and U grid respectively). The latter two are used to construct the geostrophic speed, $|\overline{\boldsymbol{u}_g}|$. The distance, $\mathrm{d}s$, is computed as the great circle distance between coordinates adjacent on the streamline, and finally $G/|\overline{\boldsymbol{u}_g}|$ is integrated trapezoidally along the streamline, for each PV forcing term G in (4.44). Other methods of integration were tested, such as assuming that G and $|\overline{\boldsymbol{u}_g}|$ (rather than $G/|\overline{\boldsymbol{u}_g}|$) both vary linearly between data points, which essentially results in a trapezoidal integration of G over the time taken for the geostrophic flow to travel the distance $\mathrm{d}s$. The results are very similar.

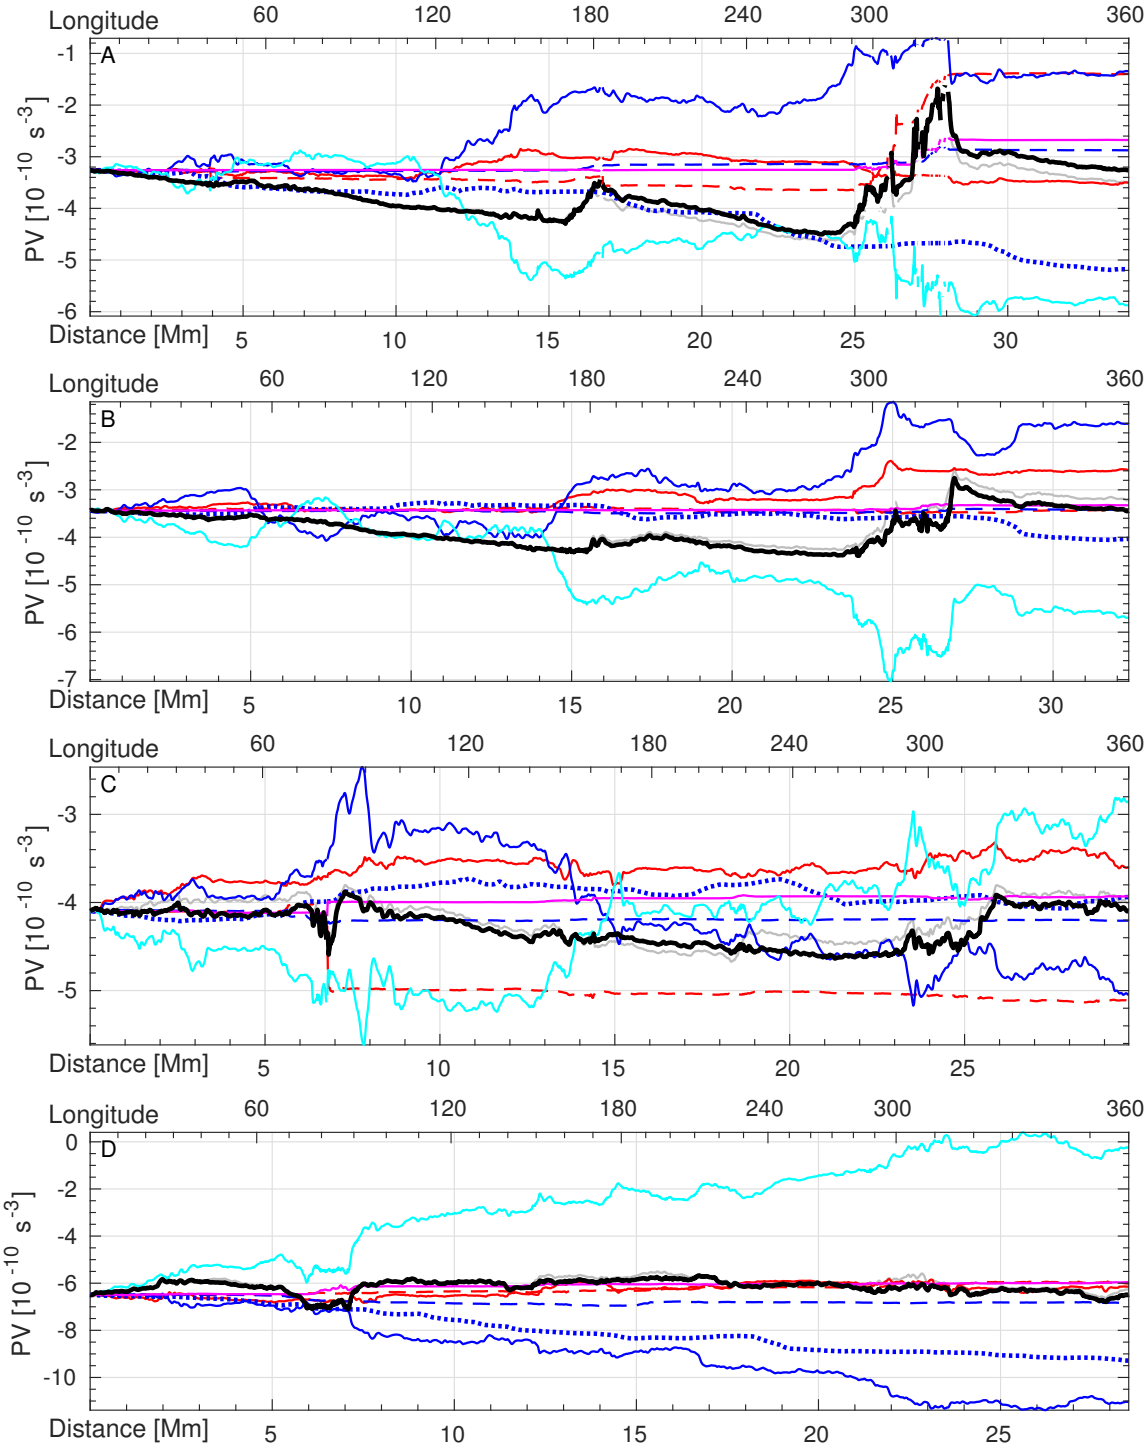


Figure 4.5: PV (solid black) and its forcing integrated along the same four streamlines as in Fig. 4.4. Streamline A is the most northerly, D the most southerly. The PV forcings are (in order of the RHS of (4.44), from top to bottom then left to right) ageostrophic advection (solid magenta), buoyancy drift (dotted blue), eddy buoyancy fluxes (solid blue), buoyancy forcing (dashed blue), MVV advection (solid cyan), momentum drift (invisible—entirely negligible and not shown), eddy momentum fluxes (solid red), and momentum forcing (dashed red). Also shown is the integrated geostrophic advection of PV (light grey, RHS of (4.40)), which should match the PV itself, apart from numerical errors.

4.4.3 Results

Figure 4.5 shows each of the integrated PV forcings along the same four streamlines from Fig. 4.4.

First, consider streamline D, the most southerly streamline (panel D). The PV (black) is fairly well conserved compared to the size of its forcings. The dominant forcings are the MVV advection (cyan) acting to decrease $|\check{Q}|$, and a combination of eddy buoyancy fluxes (solid blue) and buoyancy tendency (dotted blue) acting to increase $|\check{Q}|$. This streamline (the second from the south in Fig. 4.2, is consistently in a region of Ekman upwelling, which advects deeper, lower $|\check{Q}|$ water up onto the buoyancy surface in question, thereby decreasing $|\check{Q}|$. Helpful though this heuristic argument is, it must be remembered that the MVV advection (cyan curve) is by W, which is *not* the Ekman upwelling velocity; in addition to $\overline{w}$, W includes the upglide velocity by the mean flow staying on the mean buoyancy surface, and so W is heavily modified by the eddy field. Mathematically, $\check{Q}_z < 0$ everywhere, so $W > 0$ implies $-W\check{Q}_z > 0$, which forces $\check{Q}$ to increase ($|\check{Q}|$ to decrease) along streamlines by (4.44). This is opposite to the mechanism proposed by Gille (1997), that Ekman upwelling squishes isopycnals near the surface and increases $|\check{Q}|$.

This mean Ekman upwelling is counterbalanced in large part by an eddy-induced downwelling, bringing high $|\check{Q}|$ water from above down onto the buoyancy surface. In this Eulerian analysis, this eddy-induced downwelling is captured by the eddy buoyancy and eddy momentum fluxes.

Now consider Streamline A, the most northerly streamline that doesn't ground into South America. Here, the dominant balance between the MVV advection and the eddy buoyancy fluxes is reversed. This is consistent with being on the north end of the Residual Meridional Overturning Circulation (Marshall and Radko, 2003), where Ekman downwelling occurs (advecting high $|\check{Q}|$ water onto the surface) and is compensated by an eddy-induced upwelling (thereby decreasing $|\check{Q}|$). This signal is not quite as clean as for the southerly streamline, however, as for instance in the Indian sector between 60°E and 100°E. Northerly though this streamline is, it is not clearly in the Ekman downwelling region, but rather closer to straddling the boundary between upwelling and downwelling regions (Fig. 4.2). Once again buoyancy drift increases $|\check{Q}|$ by a surprisingly large amount. Vorticity forcing is very significant at South America; consistent with a northward flow with topography to its west (and SOSE's no-slip boundary condition), cyclonic drag increases $f + \zeta$ towards zero, thereby reducing $|\check{Q}|$. Campbell Plateau also imposes noticeable vorticity forcing of PV along

this streamline.[8] Finally, ageostrophic advection is generally insignificant but helps with the (absolute) PV drop at Drake Passage for this northerly streamline. Thus, low $|\check{Q}|$ water (such as water north of the ACC) is ageostrophically advected onto this streamline, decreasing its $|\check{Q}|$. This is consistent with the Residual Meridional Overturning Circulation flowing up and south at these mid-depths, and suggestive of a "spiralling in" of water parcels towards Antarctica as they circumnavigate it (Tamsitt et al., 2017).

Streamlines in the middle of the ACC (middle panels of Fig. 4.5) show a more mixed balance between the MVV advection and eddy buoyancy fluxes. The net signal (the value after the complete circumpolar integration) remains consistent with the notion of Ekman upwelling/downwelling in the south/north, counterbalanced by eddy fluxes. Along these streamlines, buoyancy drift is much smaller than for the northerly and southerly streamlines, and eddy momentum fluxes are more significant. Vorticity forcing is also very strong along streamline C, which is the southern-most streamline in Fig. 4.2 that passes north of Kerguelen Plateau. Here, friction accelerates the flow cyclonically, thereby decreasing PV (increasing $|\check{Q}|$). This forcing is not, though it may appear so from the zoomed-out view of Fig. 4.5, due to one malicious grid-point; the frictional forcing of PV is a robust feature along topography.

For all streamlines, diapycnal fluxes are small (dashed blue curves in Fig. 4.5, in agreement with Gille (1997).

To help understand these PV forcings, Fig. 4.6 maps the four dominant PV forcings across the Southern Ocean. Buoyancy drift is among the largest forcings in the south where the buoyancy surface is near the sea-surface; further north it would be hardly noticeable, but it is evidently still significant (Fig. 4.5). It is single-signed on much larger scales than other fields, so this small-in-magnitude term with a large "fetch" can still integrate up to something mighty.

Zooming in, Figure 4.7 shows all the PV budget terms around Drake Passage. The MVV advection and eddy buoyancy fluxes cancelling each other are still major terms, but directly adjacent to the shelf it is evident that momentum forcing and ageostrophic advection are significant, whereas away from topography and the sea-surface they are generally negligible. This also applies near other topographic obstacles, such as Campbell Plateau.

[8] As a technical note, without the discretization of Appendix 4.A, accurate analysis this close to topography seemed impossible: wild fluctuations associated with a single grid point dominated the balance.

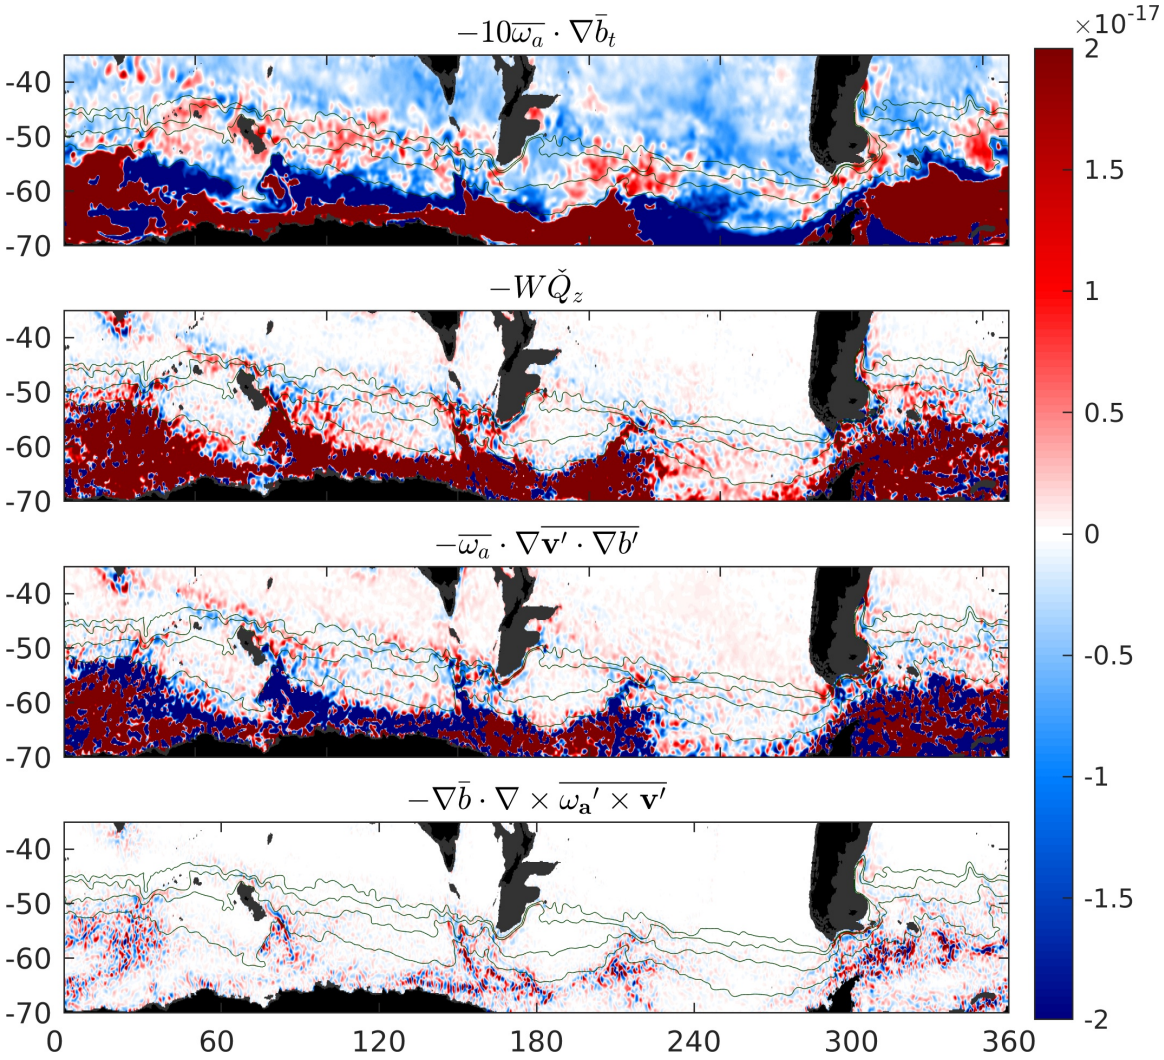


Figure 4.6: Maps of selected PV forcing terms [s^{-4}]: the buoyancy tendency, the MVV advection, the eddy buoyancy flux, and the eddy momentum flux (top to bottom). Their signs are as in (4.44); thus $\check{Q}$ increases following the geostrophic flow through red regions. Note the buoyancy tendency (top panel) has been multiplied by 10 to reveal the deep ocean. These fields are smoothed by 1/3° in both directions. The four streamlines in Fig. 4.4 and Fig. 4.5 are shown.

4.4.4 Discussion

Winds can certainly affect the deep ocean's PV structure, indirectly. Gille (1997) hypothesized that the gain of $|\check{Q}|$ is related to Ekman upwelling, which by lifting isopycnals squeezes them more closely towards the surface, thereby reducing the layer thickness and increasing $|\check{Q}|$. Of the Semtner-Chervin model, Gille says there is Ekman downwelling "to the north of about 50°S", and Ekman upwelling south of 50°S, and thus only in the Pacific sector where streamlines are further south do they

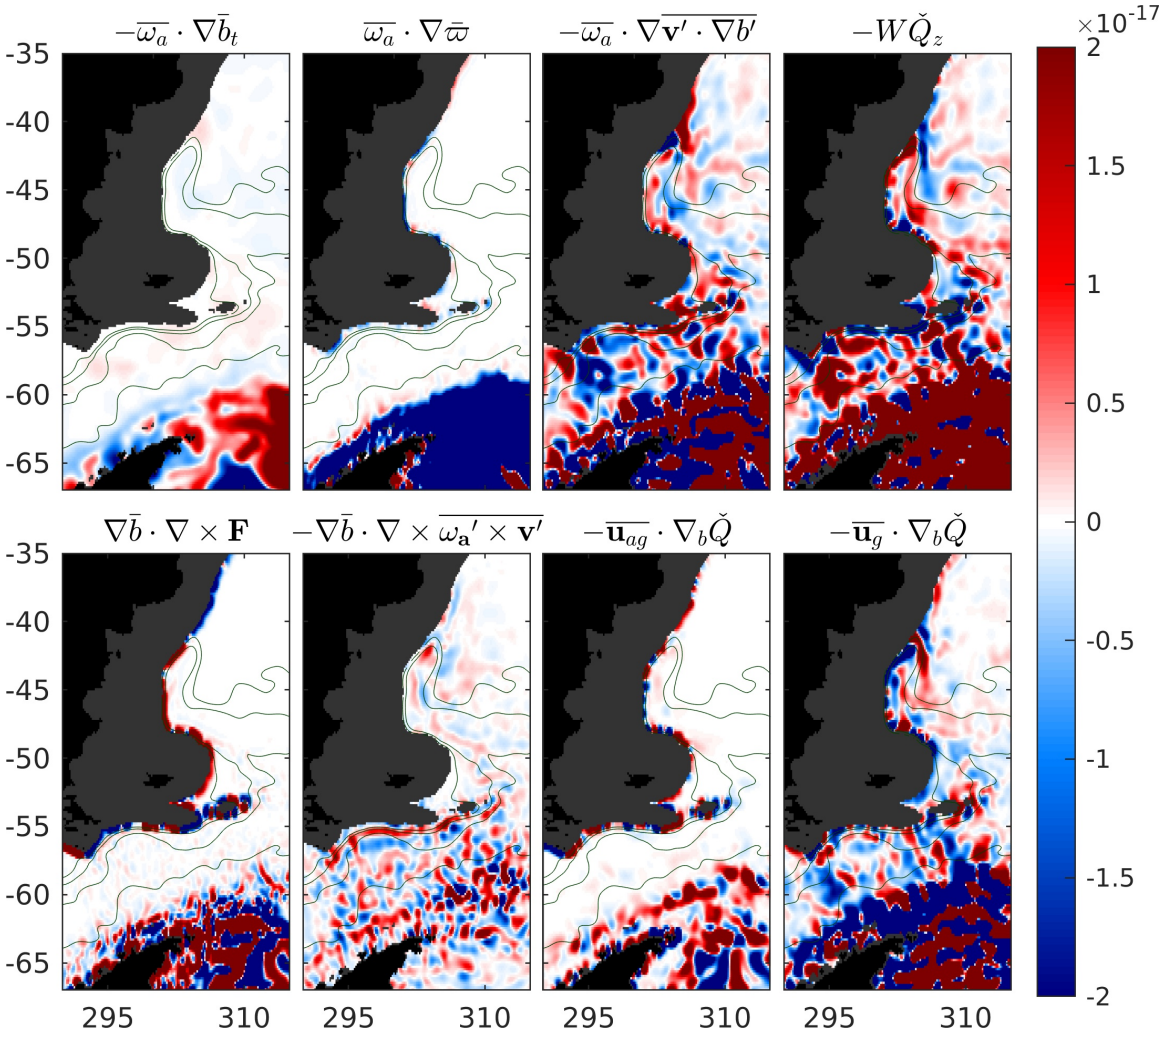


Figure 4.7: As in Fig. 4.6 but zoomed into Drake Passage and showing all terms: buoyancy tendency, buoyancy forcing, eddy buoyancy flux, MVV advection, momentum forcing, eddy momentum flux, ageostophic advection, and geostrophic advection (left to right then top to bottom).

sit beneath Ekman upwelling, increasing $|\check{Q}|$ by the squeezing mechanism. For SOSE, the time-mean Ekman upwelling velocity is shown in Fig. 4.2. The contour of zero wind stress curl, where Ekman upwelling vanishes, is clearly not at 50°S nor at other latitude: rather, it is non-zonal, taking a path similar to the ACC itself. Perhaps this is the ACC following the maximum wind stress (Allison et al., 2010). The meridional wind stress is responsible for most of the smaller scale features in the Ekman upwelling velocity (not shown). So, if winds are the main drivers of the (absolute) PV gain in the deep ocean away from topography, this tilt allows that gain to be had essentially everywhere, not just in the Pacific sector.

But do winds really influence the PV structure of the deep ocean, so strongly?

Since Earth's rotation imparts some stiffness to the fluid in the vertical direction (Taylor columns), one might argue that surface wind stress could be felt far down. This is not the right paradigm for the ACC though, which clearly exhibits strong shear (Killworth, 1992) and baroclinic compensation (Veronis, 1981). If winds did directly impact the deep PV, it would certainly fly in the face of classic oceanographic theories such as Luyten et al. (1983) that assume PV is forced in the upper ocean and conserved in the deep. Then again, this whole chapter flies in the face of such theories: PV is not conserved even in the deep.

Arguably, the winds could strongly determine the large-scale integrated PV structure, but the local PV balance will be overwhelmingly set by eddying processes. Consider the large-scale residual PV equation, (4.22). The actual vertical stress divergence is only significant in the upper few tens of metres. Beneath that, the interfacial form stress and the PV advection are the dominant terms (assuming a near steady state). The wind stress only influences the PV budget in the deep ocean by indirectly impacting these terms. It does have quite a direct effect on the mean vertical velocity, but as we have seen this is modulated in the PV equation by the mean upglide velocity. SOSE, having a higher resolution than the Semtner-Chervin model, better resolves eddies, and might exhibit less sensitivity in the deep ocean to the surface wind stress.

The spatial structure of the PV forcing by MVV advection ($-W\check{Q}_z$) is rich and interesting (Fig. 4.6). First, note that it cancels, quite strongly and locally, with the eddy buoyancy flux (middle panels of Fig. 4.6). This can also be seen in the integrated PV forcings along streamlines: the PV forcing by eddies largely cancels the MVV advection (blue and red solid curves vs. cyan curve in Fig. 4.5). This is due to the Eulerian-style averaging, which can be seen as follows. Multiplying the buoyancy equation by the layer thickness $\sigma = b_z^{-1}$ and re-arranging gives an expression for the vertical velocity; noting that $\xi_{\tilde{x}} = -\sigma b_x$ etc, in buoyancy coordinates this is

$$w = \xi_{\tilde{t}} + u\xi_{\tilde{x}} + v\xi_{\tilde{y}} + \varpi\xi_{\tilde{b}}. \tag{4.46}$$

An Eulerian average of this gives

$$-W\check{Q}_z = -\left(\overline{\xi}_{\tilde{t}} + \overline{u'\xi'_{\tilde{x}}} + \overline{v'\xi'_{\tilde{y}}} + \overline{\varpi\xi_{\tilde{b}}}\right)\check{Q}_z. \tag{4.47}$$

In the ocean interior, the eddy terms will dominate the above. Hence in this perspective eddies dominate ($-W\check{Q}_z$) as well, and it is no surprise that there is large cancellation. It is another way of viewing the cancellation between the Eulerian-mean and eddy-induced overturning circulations in the Southern Ocean.

The spatial structure of the PV forcing by MVV advection, and also by eddy buoyancy fluxes, often appears as a dipolar pattern, apparent especially from the Cape of Good Hope to the Pacific Antarctic Ridge (Fig. 4.6). Eddy buoyancy fluxes tend to increase PV on the north side of this dipole, and decrease them on the south side, and vice versa for the MVV advection. In this region, the PV actually *decreases* northward (not shown), suggesting that the MVV advection establishes this northward PV decrease, and the eddy forcing erodes it, working to homogenize PV in this region, as eddies that flux PV down-gradient do. Thus it is no surprise that this feature aligns with a region of homogeneous PV in Fig. 3.11. The structure of the $(-W\check{Q}_z)$ field is essentially all due to W (not shown); since $(-W\check{Q}_z)$ has the same sign as W, this dipole thus indicates a mean circulation pattern that is tubular from the Cape of Good Hope to the south tip of Campbell Plateau in particular, with mean upwelling piercing mean buoyancy surfaces on the south side of the tube and mean downwelling piercing mean buoyancy surfaces on the north.

This tube aligns quite well with the Dynamical Subtropical Front, which (unlike the Subtropical Front) is also a deep water mass boundary (Graham and De Boer, 2013). This tube appears similarly in other buoyancy surfaces, both shallower and deeper, with no obvious horizontal shift (not shown). It also aligns with the Ekman upwelling velocity switching sign (Fig. 4.2), and also weakens in the eastern Pacific as does the northward gradient of the Ekman upwelling velocity, but $(-W\check{Q}_z)$ is small north of this tube despite the Ekman upwelling velocity remaining large. It thus appears that $(-W\check{Q}_z)$ is related to *gradients* in the Ekman upwelling velocity. Recalling (4.42), the $\overline{w}$ part of W is closely related to the Ekman upwelling velocity, whereas the upglide part of W, $\overline{\boldsymbol{u}} \cdot \nabla_b z$, is likely related to gradients in Ekman upwelling that can maintain sloped isopycnals.

The dipole forcing of PV can be thought of in two ways: by changing the layer thickness or by changing the relative vorticity. Consider the PV forcing by MVV advection. In the first light, it is seen to thicken layers on the south of the dipole (PV increased towards zero), and thin them on the north (PV decreased), consistent with a mean overturning (Deacon) cell with Ekman downwelling on the north side, upwelling on the south. In the second light, it is seen to widen / decelerate the jet (relative vorticity gained to the south and lost to the north). The opposite is true for the eddy buoyancy fluxes: they are seen to thin layers on the south side, and thicken them on the north, opposing the Deacon cell and consistent with baroclinic instability, more intense in the upper water column than the lower, acting to flatten

isopycnals. Likewise, the eddy buoyancy fluxes are seen to sharpen / accelerate the jet.

This cancellation is built directly into the TWA perspective: there are no eddy terms in the averaged thickness equation, which derives from (4.46). Thus I expect a rather different PV balance balance to emerge in a proper TWA analysis.

I can only speculate about what such an analysis would reveal. From Fig. 4.7 and looking closely at the Drake Passage region in Fig. 4.5A, it seems likely that momentum forcing and ageostrophic advection would continue to be important in the Drake Passage region, and may even become dominant as the MVV and eddy terms are suppressed in the TWA.

4.5 Conclusions

This chapter has studied the time-mean potential vorticity along time-mean streamlines in the Antarctic Circumpolar Current, seeking to answer two questions: how does the time-mean PV change along time-mean streamlines, and why?

Precisely what is meant by "time-mean" was carefully considered, and three averaging schemes were compared: a standard Eulerian average, the thickness weighted average, and a sort of hybrid between the two that is possible specifically for PV. Time-mean PV variations were found to be qualitatively similar between the three averaging methods.

Predominantly, the time-mean PV is far from conserved along time-mean streamlines in the ACC: there is a large drop in (absolute) time-mean PV just downstream of Drake Passage for northerly streamlines, which hug the coast of South America. This (absolute) time-mean PV drop is steadily regained at all other longitudes, though other topographic barriers such as Kerguelen Plateau and Campbell Plateau also cause strong local variations in the time-mean PV along nearby time-mean streamlines. This is generally in agreement with the findings of Gille (1997), except that she found the (absolute) time-mean PV is gained only in the Pacific sector of the ACC. The gain of (absolute) PV in all sectors of the Southern Ocean is true even for northerly streamlines that might be expected to predominantly experience Ekman downwelling in the Atlantic sector. However, the non-zonal tilt of the Ekman pumping in SOSE (different from Semtner-Chervin model analysed by Gille (1997)) keeps even these northern streamlines in the Atlantic feeling a mixture of Ekman upwelling and downwelling.

Having found that time-mean PV variations were similar between the three averaging methods, the question of why the time-mean PV varies as it does was addressed by studying the time-mean PV budget using just one averaging method: the hybrid method. This was out of necessity, as the (preferable) TWA method cannot be accurately performed from the model's 5-day Eulerian averaged output. In the hybrid method, the time-mean PV budget is dominated by two terms: (1) eddy buoyancy fluxes, and (2) advection of time-mean PV by the mean vertipcynal velocity, which is the mean vertical velocity minus the upglide velocity, i.e. the vertical velocity induced by the mean horizontal velocity following mean buoyancy surfaces. These terms largely cancel each other.

To the extent that the mean vertipcynal velocity reflects the Ekman upwelling velocity, on the southern flank of the ACC this is upwards, which advects lower (absolute) PV onto the time-mean buoyancy surface, thereby decreasing the (absolute) PV there. In fact, this is opposite to the effect of the mechanism proposed by Gille (1997) for how the wind stress affects the deep PV, which is that Ekman upwelling, say, lifts up deep isopycnals and squeezes them together towards the surface, thereby decreasing the layer thickness and increasing the (absolute) PV. While it is true that Ekman upwelling positively correlates with smaller layer thickness, that does not imply causation, and one must consider the effect of the Ekman upwelling velocity on the PV structure given an already established state with these layer thickness variations.

This Ekman upwelling in the south flank of the ACC is part of the Eulerian Meridional Overturning Circulation, which is strongly cancelled by an eddy-induced Meridional Overturning Circulation. The sum of these, the Residual Meridional Overturning Circulation, is much closer to zero than either one individually. Again in the southern flank of the ACC, the eddy-induced velocity is downwards, advecting higher (absolute) PV onto the time-mean buoyancy surface, thereby increasing the (absolute) PV there. This signal is clearly seen on southerly streamlines (Fig. 4.5d). For northerly streamlines, the Eulerian and eddy-induced Meridional Overturning Circulations are reversed: Ekman downwelling tends to increase the (absolute) PV by pushing down shallower, higher (absolute) PV water onto the time-mean buoyancy surface, while the eddy-induced velocity lifts up deeper, lower (absolute) PV onto the time-mean buoyancy surface.

Having established the primary balance of the time-mean PV budget in the hybrid method, it must be said that eddy vorticity fluxes and tendency of the time-mean PV (almost entirely caused by drift in the buoyancy field, not the momentum field) both contribute significantly to the time-mean PV budget. Ageostrophic advection of the

time-mean PV and forcing (more-so of vorticity than buoyancy) can be significant near topography, but elsewhere is negligible. All these (usually second-order) forcings must be considered in the vicinity of Drake Passage.

The notable drop in (absolute) time-mean PV just downstream of Drake Passage for the second most northerly streamline studied (Fig. 4.5b) is largely caused by eddy fluxes, with both buoyancy and relative vorticity fluxes being significant. The ACC's eastward flow is sufficiently strong that eddies are advected eastwards (the flow Doppler-shifts the Rossby wave phase speed sufficiently that it becomes eastward) (Hughes, 1995; Hughes et al., 1998; Klocker and Marshall, 2014). However, this is not the case where the flow is northwards along the eastern shelf of South America (see Fig. 3.1 of Zajaczkovski (2017), or Fig. 1c of Klocker and Marshall (2014)). Here, eddies may propagate west, bringing them to impinge on topography and break, where they may strongly force the PV. This PV forcing generally causes a drop in (absolute) time-mean PV here, which enables the northward flow in the first place. Thus a self-supporting cycle is established. Closer to the shelf, for the most northerly streamline studied (Fig. 4.5a), the drop of (absolute) time-mean PV is no longer caused by eddy fluxes, but mostly by vorticity forcing, consistent with a northward flow along a no-slip boundary to its west.

Moreover, the (absolute) time-mean PV drop at Drake Passage is consistent with a topological way of thinking about time-mean PV being constant along individual streamlines, but where a streamline is disconnected (such as those slightly north of the most northerly circumpolar streamline studied here) the time-mean PV can take different values on each connected component of the streamline. This way of thinking naturally explains why the southerly streamlines through Drake Passage exhibit better time-mean PV conservation than the northerly circumpolar streamlines, which, being close to non-circumpolar streamlines that have a jump in time-mean PV between disconnected regions, must exhibit strong fluctuations of time-mean PV themselves if the time-mean PV field is smooth.

This way of thinking also helps explain why the ageostrophic forcing of time-mean PV is highest just downstream of Drake Passage: The flow must get past South America even though some geostrophic streamlines intersect topography, and thus strong ageostrophic flow is required to keep the flow from penetrating topography. (Rotational streamlines were built with a no-normal flow boundary condition, so forcing of PV by the divergent flow is much smaller than for the ageostrophic flow.)

Real fluid parcels do not follow the mean streamlines, geostrophic or otherwise. If a fluid parcel transiting Drake Passage is better at conserving its PV than fluid following

the streamlines, it will lose less (absolute) PV at Drake Passage and consequently move from its initial streamline to a more southerly streamline. This thinking is consistent with the meridional overturning circulation taking water, particularly in the Atlantic sector, across the ACC from north to south at these mid depths.

Finally, it is worth stressing the surprisingly large role of drift in SOSE from 2005—2010. It is significant in the time-mean PV budget, even in the deep ocean. It is caused almost entirely by drift of the stratification, rather than the vorticity. It is present also in the TWA analysis (not shown) and thus is not simply due to a certain depth drifting from one isopycnal to a shallower one, but is indeed a drift in the layer thickness. This drift in SOSE is also exhibited in gridded ARGO data (not shown), which is expected given SOSE incorporates ARGO data. Thus ARGO data, with a longer time horizon, could be used to help ascertain whether this stratification drift is due to climate change, or interannual variability.

The results of this chapter are arguably cleaner than those of Gille (1997). Some of this improvement could be due to SOSE being a much more advanced model than the Semtner-Chervin model. However, some of it must be due to the theoretical framework used here and the numerical methods. In particular, using buoyancy surfaces rather than potential density and defining the PV on the upper south-west corner of the tracer cell perfectly annihilates the solenoidal PV forcing, which was not negligible in the analysis of Gille (1997).

Still cleaner results might be obtained through the TWA method, though this requires a considerable effort to do properly, and thus I can only speculate about the PV balance it would reveal. Since the averaging is done following the buoyancy surface, the equivalent of the MVV in the TWA method is just the diapycnal velocity, which should be far smaller than the MVV. Thus the dominant cancellation found in the time-mean PV budget from the hybrid method is probably filtered out of the corresponding TWA budget. There have been several recent papers using TWA on Eulerian-averaged output from z level models. These studies either focus exclusively on the momentum equation (Aoki et al., 2016), or neglect the buoyancy forcing (Cessi and Wolfe, 2013). I argue that more care must be taken: with this data, the momentum equation is closed in a TWA only at the expense of the thickness equation, which could be fixed by massively increasing the buoyancy forcing (equivalent to the diapycnal velocity in buoyancy coordinates). Ringler et al. (2016) take the necessary step of building the TWA analysis into the model online, but more work is needed to do this with less idealized models, for example with a fully non-linear equation of state. This discussion is further developed in Section 4.B.5.

As Gille (1997) pointed out, this non-conservation of time-mean PV along time-mean streamlines bears on many theoretical models of the ACC. For instance, the equivalent barotropic structure of Killworth (1992) relies on a functional relationship between PV and the mass-weighted stream function on layers, which is broken by this non-conservation of time-mean PV. Marshall (1995b) took a further simplifying step by assuming the PV is uniform on layers. This assumption is probably why his modelled ACC does not veer northwards, as the real ACC does, along the coast of South America downstream of Drake Passage: with PV uniform, the flow can only move northwards if the layer thickness shrinks, and mass conservation puts strong limits on the extent of this. These types of models are illuminating, but of course do not capture the full dynamics of the ACC. Their goal is to capture the first-order dynamics. Given that the time-mean PV can vary by over 50% from its mean value along a time-mean streamline (not to mention variations between different stream lines, as per Chapter 3), it is not entirely clear that PV-conserving flow should be considered the first-order dynamics. The perspective of Hughes and Killworth (1995) on the ACC is helpful here: the ACC generally obeys certain first-order dynamics (such as conservation of PV) but has small regions where second-order dynamics (that break PV conservation) become of first-order importance, and thereby affect the dynamics of the entire ACC.

Appendix

4.A Where does potential vorticity live on a C grid?

Just as there are many possible ways to define the potential vorticity, given a definition there are many possible ways to numerically compute it. In this section I describe my computation of the Ertel PV in a z coordinate, C-grid model such as the MITgcm. Since PV is not a model variable, there is no pre-defined grid upon which it is defined, but I argue it lives best on the corners of the 3D tracer cell.

Figure 4.8 shows a tracer cell for some tracer T living at its centre. It is commonly known that the relative vorticity (ζ but labelled ω_k in Fig. 4.8) lives on the south-west corner of the tracer cell, at the depth of T. When the GCM uses the vector invariant form of the momentum equations, ζ is computed by the model, so this is non-negotiable (to the offline data-analyst). Its computation appeals to Stokes' theorem: the area integral of ζ equals the line integral of the velocity bounding the

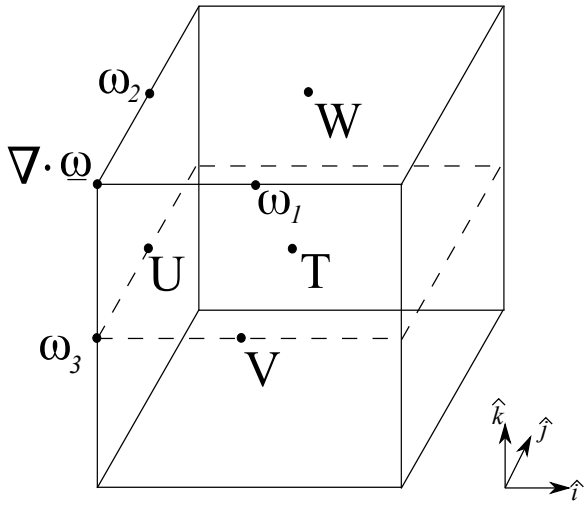


Figure 4.8: Layout of a tracer cell on the C-grid. At the centre is the tracer grid, T. To the west, the south, and above are the zonal U, meridional V, and vertical W velocities. The components of the curl $\boldsymbol{\omega} = \nabla \times \boldsymbol{u}$ are on the upper-south, upper-west, and south-west edges of the tracer cell, and the divergence of the curl, $\nabla \cdot \boldsymbol{\omega}$ is on the upper south west corner of the tracer cell. This is where I define PV. The $\hat{i}$ (zonal), $\hat{j}$ (meridional), and $\hat{k}$ (upwards) unit vectors are indicated.

area. With ζ on the south-west corner at the tracer depth, this area can be the four T grid points surrounding the ζ grid point; then the velocities are positioned so the line integral is computed naturally.

To compute the large-scale PV $(f + \zeta)b_z$, one might average ζ back to the tracer cell centre where b lives; for b_z one could average the stratification $N^2 = b_z$ (which lives on the W grid) onto the T grid, or differentiate a PCHIP to $b(z)$, or take a centred difference of b, as some examples. One could extend this approach to the Ertel PV, but it will not correctly preserve the total PV; in particular, the solenoidal term—see (4.39)—will not be annihilated.

The model grid is much better suited to calculating divergences than to inner products of vectors. This is why, in the centred second order advection scheme, $\boldsymbol{u} \cdot \nabla T$ is computed as $\nabla \cdot (T\boldsymbol{u}) - T\nabla \cdot \boldsymbol{u}$. Similarly, I compute Ertel PV as a divergence:

$$Q = \boldsymbol{\omega}_a \cdot \nabla T = \nabla \cdot (T\, \boldsymbol{\omega}_a)\,. \tag{4.48}$$

This requires that the absolute vorticity $\boldsymbol{\omega}_a$ is divergence free: $\nabla \cdot \boldsymbol{\omega}_a = 0$ identically.

Working in spherical coordinates, one finds $\nabla \cdot 2\boldsymbol{\Omega} = 0$, where $2\boldsymbol{\Omega} = 2\Omega(\cos\theta \hat{\boldsymbol{j}} + \sin\theta \hat{\boldsymbol{k}}) = f'\hat{\boldsymbol{j}} + f\hat{\boldsymbol{k}}$, where $\Omega = 2\pi/86\,164\,\mathrm{s}$ and θ is latitude. Also, $\nabla \cdot \nabla \times \boldsymbol{u} = 0$ by vector identity. So, ought the absolute vorticity $\boldsymbol{\omega}_a$ be $(2\boldsymbol{\Omega} + \nabla \times \boldsymbol{u})$?

No, because the traditional approximation (which SOSE makes) sets $f' = 0$, so $\boldsymbol{\omega}_a = f\hat{\boldsymbol{k}} + \nabla \times \boldsymbol{u}$ should be used when working with such model data. But then, working in spherical coordinates, one finds $\nabla \cdot \boldsymbol{\omega}_a = 2fR_E^{-1}$. This problem is resolved by realizing that the model does not take divergences in spherical coordinates, but rather in a distorted geometry created by the thin-film approximation (which for this reason must go hand-in-hand with the traditional approximation). The thin-film approximation approximates the actual radius r with the Earth radius R_E . Imagine a box that is $1^\circ \times 1^\circ$ (or any shape in latitude-longitude space) and bounded between two radii. The area on the outer side of this box is larger than the area on the inner side. Considering the divergence theorem to relate $\nabla \cdot \boldsymbol{\omega}_a$ to $\boldsymbol{\omega}_a$ on the boundary of this box, this is how $\nabla \cdot \boldsymbol{\omega}_a \neq 0$ in spherical coordinates, as above. Under the thin-film approximation, these areas are made equal, so $\nabla \cdot \boldsymbol{\omega}_a = 0$. Dellar (2011) discusses these issues in greater detail.

With this understanding, the numerical discretization appropriate for (4.48) is given in Listing 4.1. This calculates the curl of a "staggered" vector on the (U, V, W) grid (variables `O1`, `O2`, `O3` in the code give this curl), averages a tracer on the T grid onto the vorticity grid (`TonO1`,`TonO2`, `TonO3` in the code), multiplies these together, and finally takes the divergence (`D` in the code). The code describes the numerical discretization as well as any mathematical formula could. Though this is never computed online by the model, it is defined analogously to the divergence on the tracer cell, for terms like $\nabla \cdot (T\boldsymbol{u})$, that are computed online.

This discretization preserves the vector identity $\nabla \cdot \nabla \times \boldsymbol{a} = 0$ for any vector $\boldsymbol{a}$. This happens naturally when the divergence of the curl, and hence PV, is placed on the upper south west corner of the tracer cell. The tracer T needs to be averaged twice to get it on the grid points of $\nabla \times \boldsymbol{u}$, as opposed to the single average needed to get T onto the grid points of $\boldsymbol{u}$ for advection. The curl computed by this code also preserves $\nabla \times \nabla\phi = 0$ for any scalar field ϕ, which is vital to annihilate the pressure and kinetic energy terms (together, Bernoulli).

This method applies equally well to the instantaneous PV or to the "hybrid" time-mean PV of (4.25). Even better, the terms in the PV budget can be computed the exact same way. For example, $\nabla \cdot (\overline{T_t}\ \overline{\boldsymbol{\omega}}_a)$ gives the forcing of mean PV due to tracer tendency, and $\nabla \cdot (\overline{T}\ \nabla \times (\overline{\boldsymbol{\omega}}_a \times \overline{\boldsymbol{u}}))$ gives the forcing of mean PV due to momentum advection by the mean flow, etc. Since the terms in the tracer equation exist everywhere the tracer exists (non-land grid cells), and similarly for the velocities, the PV forcing terms are guaranteed to exist everywhere that the PV exists: masking issues are handled naturally.

```
function [D, O1, O2, O3] = gradT_dot_curluvwplusf(g, T, U, V, W, f)
% Compute gradient of tracer dotted with the curl of a velocity, as
    a divergence. Optionally include f to also add dT/dz * f.
% INPUT:
% g: a class defining the properties of the grid and also the
    functions  im1(), jm1(), and km1() to (circularly) shift data one
     grid point east, north, and down, respectively. Since the
    meridional and vertical directions are non-periodic, the second
    argument to jm1() and km1() provides the fill-value.
% T [array size [nx,ny,nz]]: tracer of size [nx,ny,nz] on the (i,j,k
    ) grid.
% U [array size [nx,ny,nz]]: eastward velocity size on the (i-1/2,j,
    k) grid.
% V [array size [nx,ny,nz]]: northward velocity size on the (i,j
    -1/2,k) grid.
% W [array size [nx,ny,nz]]: upward velocity size on the (i,j,k-1/2)
     grid.
% f [array size [1,ny,1]], optional (default 0): Coriolis parameter;
     added to the vertical component of curl(U,V,W); on the (j-1/2)
    grid.
% OUTPUT:
% D [array [nx,ny,nz]]: gradient of T dotted with (curl of (U,V,W) +
     (f k)), computed as a divergence, on the (i-1/2,j-1/2,k-1/2)
    grid.
% O1 [array [nx,ny,nz]]: 1st component of curl(U,V,W), on the (i,j
    -1/2,k-1/2) grid.
% O2 [array [nx,ny,nz]]: 2nd component of curl(U,V,W), on the (i
    -1/2,j,k-1/2) grid.
% O3 [array [nx,ny,nz]]: 3rd component of curl(U,V,W), on the (i
    -1/2,j-1/2,k) grid.

if nargin < 6 || isempty(f)
    f = 0;
end

% Multiply velocities by edge length, so they become circulations:
DRCz = permute(g.DRC(:), [3 2 1]);
uT = U .* g.DXC;
vT = V .* g.DYC;
wT = W .* DRCz;

% The three components of curl(U,V,W), on the upper south edge,
% the upper west edge, and the south west edge, respectively:
O1 = ((wT - g.jm1(wT,0)) - (g.km1(vT,0) - vT)) ./ (g.DYC * DRCz);
O2 = ((g.km1(uT,0) - uT) - (wT - g.im1(wT))) ./ (g.DXC .* DRCz);
O3 = ((vT - g.im1(vT)) - (uT - g.jm1(uT,0))) ./ g.RAZ;

% Add the Coriolis parameter:
O3 = O3 + f;

% Average T onto grid of curl(u,v,w):
good = isfinite(T);
T(~good) = 0;
```

```
good = double(good);

TonO1 = T + g.jm1(T);
TonO1 = TonO1 + g.km1(TonO1);
goodO1 = good + g.jm1(good,0);
goodO1 = goodO1 + g.km1(goodO1,0);
TonO1 = TonO1 ./ goodO1;

TonO2 = T + g.im1(T);
TonO2 = TonO2 + g.km1(TonO2);
goodO2 = good + g.im1(good);
goodO2 = goodO2 + g.km1(goodO2,0);
TonO2 = TonO2 ./ goodO2;

TonO3 = T + g.im1(T);
TonO3 = TonO3 + g.jm1(TonO3);
goodO3 = good + g.im1(good);
goodO3 = goodO3 + g.jm1(goodO3,0);
TonO3 = TonO3 ./ goodO3;

% Multiply vorticity vectors by tracer and face area:
uT = O1 .* TonO1 .* g.DYC .* DRCz;
vT = O2 .* TonO2 .* g.DXC .* DRCz;
wT = O3 .* TonO3 .* g.RAZ;

% Divergence of  T * (curl(U,V,W) + f hat{k}):
D = uT - g.im1(uT) + vT - g.jm1(vT) + g.km1(wT) - wT;
D = D ./ (g.RAZ .* DRCz);
end
```

Listing 4.1: MATLAB function to compute $\nabla T \cdot (\nabla \times \boldsymbol{u} + f\hat{\boldsymbol{k}}) = \nabla \cdot (T\ (\nabla \times \boldsymbol{u} + f\hat{\boldsymbol{k}}))$ on the upper south west corner of a tracer cell.

4.B Thickness-weighted average momentum, thickness, and PV budgets

This Appendix presents the TWA theory essential to understanding the PV equation in buoyancy coordinates—for full details see Young (2012). Then, the numerical treatment of this theory is discussed in detail.

4.B.1 Theory

A brief introduction to the TWA theory was given in Section 4.1, which is further developed here. Using buoyancy as the vertical coordinate puts us into a curvilinear, nonorthogonal coordinate system; the coordinates are labelled $(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t})$. The superscript tilde reminds us that derivatives (usually denoted by a subscript), for

example $\partial_{\tilde{x}}$, are taken at constant buoyancy and hence, from a Cartesian perspective, are taken along a sloping surface. In buoyancy coordinates, the depth of a buoyancy surface, denoted $\xi(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t})$, is an independent variable. Central to the TWA theory is the layer thickness, defined as

$$\sigma(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t}) \stackrel{\text{def}}{=} \zeta_{\tilde{b}}. \tag{4.49}$$

As the name "thickness-weighted average" suggests, this is used as a weighting when taking averages. The average of a field $\alpha(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t})$ is denoted with an overbar, $\overline{\alpha}(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t})$. This is the same symbol describing the Eulerian average at fixed z, because both are simply averages at fixed points in their respective coordinate systems; describing α with these coordinates $(\tilde{x}, \tilde{y}, \tilde{b}, \tilde{t})$ naturally leads to a "buoyancy-following average". The average satisfies three properties: $\overline{\overline{\alpha}} = \overline{\alpha}$ (idempotent), $\overline{\overline{\alpha}\overline{\beta}} = \overline{\alpha}\overline{\beta}$, and commutativity with the partial derivatives, e.g. $\overline{\partial_{\tilde{x}}\alpha} = \partial_{\tilde{x}}\overline{\alpha}$. The thickness-weighted average of a quantity α is defined as

$$\hat{\alpha} \stackrel{\text{def}}{=} \frac{\overline{\sigma\alpha}}{\overline{\sigma}}. \tag{4.50}$$

Note that the TWA does not commute with partial derivatives: $\partial_{\tilde{x}}\hat{\alpha} \neq \widehat{\partial_{\tilde{x}}\alpha}$. The notation $\hat{\alpha}_{\tilde{x}}$ will always imply taking the TWA, then the derivative, i.e. $\partial_{\tilde{x}}\hat{\alpha}$. Two types of Reynolds decompositions are used: fluctuations from the average,

$$\alpha' \stackrel{\text{def}}{=} \alpha - \overline{\alpha} \tag{4.51}$$

and fluctuations from the thickness-weighted average

$$\alpha'' \stackrel{\text{def}}{=} \alpha - \hat{\alpha}. \tag{4.52}$$

Transforming between Z-coordinates and buoyancy coordinates employs the following transformations of the partial derivatives:

$$\partial_x = \partial_{\tilde{x}} + b_x\sigma\partial_z \tag{4.53a}$$

$$\partial_y = \partial_{\tilde{y}} + b_y\sigma\partial_z \tag{4.53b}$$

$$\partial_z = b_z\partial_{\tilde{b}} \tag{4.53c}$$

$$\partial_t = \partial_{\tilde{t}} + b_t\sigma\partial_z. \tag{4.53d}$$

We have applied (4.53c) to ζ to find

$$\sigma = b_z^{-1} \tag{4.54}$$

and used this in (4.53a), (4.53b), and (4.53d).

4.B.1.1 Buoyancy-coordinate thickness equation

In z coordinates the buoyancy equation is

$$b_t + ub_x + vb_y + wb_z = \varpi, \tag{4.55}$$

where ϖ would be zero if a materially conserved density variable existed; suitable choices for b render ϖ small. The buoyancy equation, together with the continuity equation $\nabla \cdot \boldsymbol{v} = 0$ in Cartesian coordinates, are jointly transformed into an equation for the thickness. To do this, multiply (4.55) by σ (to get w alone), then take the the vertical derivative ∂_z, then multiply by $-\sigma$. The forcing term easily becomes $-\sigma(\sigma\varpi)_z = -(\sigma\varpi)_{\tilde{b}}$. The tendency term becomes $-\sigma(\sigma b_t)_z = -\sigma b_t \sigma_z + \sigma_z = \sigma_{\tilde{t}}$. The advective terms take some more manipulation but work out to be $-\sigma(\sigma u b_x)_z = (\sigma u)_{\tilde{x}} - \sigma u_x$ and similarly for v, whereas the w term is immediately $-\sigma(\sigma w b_z)_z = -\sigma w_z$. Here, it is crucial that $\sigma b_z = 1$ so w stands alone before taking the z derivative, which enables the continuity equation $u_x + v_y + w_z = 0$ to eliminate three terms. This completes the transformation to buoyancy coordinates, yielding the thickness equation

$$\sigma_{\tilde{t}} + (\sigma u)_{\tilde{x}} + (\sigma v)_{\tilde{y}} + (\sigma\varpi)_{\tilde{b}} = 0. \tag{4.56}$$

Because the buoyancy equation and continuity equation are both extremely finely balanced equations (the sum of large terms equals a small term), there are severe numerical difficulties in combining these equations while changing coordinates to produce the thickness equation. These will be discussed in Section 4.B.2, but first we finish the necessary TWA derivations.

4.B.1.2 TWA thickness equation

Noting that $\overline{\sigma}\hat{u} = \overline{\sigma u}$ and similarly for v and ϖ, averaging (4.56) immediately yields

$$\overline{\sigma}_{\tilde{t}} + (\hat{u}\overline{\sigma})_{\tilde{x}} + (\hat{v}\overline{\sigma})_{\tilde{y}} + (\hat{\varpi}\overline{\sigma})_{\tilde{b}} = 0. \tag{4.57}$$

The essential and beautiful aspect of the thickness-weighted average formalism is that there are no eddy terms in the thickness equation. All eddy terms are confined to the momentum equation, as we shall see.

4.B.1.3 Buoyancy coordinate momentum equation

Next, we require the thickness-weighted momentum equations in buoyancy coordinates. To start, the (Boussinesq, hydrostatic, traditional) z coordinate momentum

equation was given in (4.30), but is re-written here for just the zonal component:

$$u_t + wu_z - (f + v_x - u_y)v + \Phi_x + K_x = \mathcal{X} \tag{4.58}$$

where all terms are as before, except the forcing is now denoted by $\mathcal{X}$. This transforms into buoyancy coordinates as

$$u_{\tilde{t}} + \varpi u_{\tilde{b}} - v(f + v_{\tilde{x}} - u_{\tilde{y}}) + m_{\tilde{x}} + K_{\tilde{x}} = \mathcal{X}. \tag{4.59}$$

Since SOSE is a z level model we have access to the terms in (4.58). We do not directly have access to the terms in (4.59), but get them via (4.53) and (4.54) as:

$$u_{\tilde{t}} = u_t \qquad - \sigma b_t u_z, \tag{4.60a}$$

$$\varpi u_{\tilde{b}} = wu_z \qquad + \sigma b_t u_z \quad + \sigma b_x u u_z + \sigma b_y v u_z, \tag{4.60b}$$

$$-(f + v_{\tilde{x}} - u_{\tilde{y}})v = -(f + v_x - u_y)v + \sigma b_x v v_z \qquad - \sigma b_y v u_z, \tag{4.60c}$$

$$K_{\tilde{x}} = K_x \qquad - \sigma b_x v v_z - \sigma b_x u u_z \tag{4.60d}$$

$$m_{\tilde{x}} = \Phi_x, \tag{4.60e}$$

$$-\mathcal{X} = -\mathcal{X}. \tag{4.60f}$$

The first term on the RHS of each of the above equations is a term in the z coordinate zonal momentum equation. The other terms sum to zero, hence proving (4.59) is the correct transformation of (4.58). Note that stating (4.60e) (and its analogous part in the meridional momentum equation) requires the existence of m, an exact geostrophic stream function on b surfaces, as discussed in Section 4.3.1.1.

4.B.1.4 TWA momentum equation

Now, considering that the thickness-weighted average zonal velocity is $\hat{u} = \overline{\sigma u}/\overline{\sigma}$, the required operation is to multiply (4.59) by σ, average it, then divide by $\overline{\sigma}$. Prior to the final division by $\overline{\sigma}$, this yields:

$$\overline{\sigma}\hat{u}_{\tilde{t}} \qquad = \overline{\sigma u_{\tilde{t}}} \qquad - \overline{\sigma}_{\tilde{t}}\hat{u} \quad + \overline{\sigma_{\tilde{t}} u}, \tag{4.61a}$$

$$\overline{\sigma}\hat{\varpi}\hat{u}_{\tilde{b}} + \left(\overline{\sigma}\widehat{\varpi'' u''}\right)_{\tilde{b}} = \overline{\sigma\varpi u_{\tilde{b}}} \qquad - (\overline{\sigma}\hat{\varpi})_{\tilde{b}}\,\hat{u} + \overline{(\sigma\varpi)_{\tilde{b}}\, u}, \tag{4.61b}$$

$$-\overline{\sigma}(f + \hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})\hat{v} + \left(\overline{\sigma}\widehat{v'' u''}\right)_{\tilde{y}} = -\overline{\sigma(f + v_{\tilde{x}} - u_{\tilde{y}})v} - (\overline{\sigma}\hat{v})_{\tilde{y}}\hat{u} \quad + \overline{(\sigma v)_{\tilde{y}} u} \quad + C, \tag{4.61c}$$

$$\overline{\sigma}\tfrac{1}{2}\left(\hat{u}^2 + \hat{v}^2\right)_{\tilde{x}} + \left(\overline{\sigma}\widehat{u'' u''}\right)_{\tilde{x}} = \overline{\sigma K_{\tilde{x}}} \qquad - (\overline{\sigma}\hat{u})_{\tilde{x}}\hat{u} \quad + \overline{(\sigma u)_{\tilde{x}} u} \quad - C, \tag{4.61d}$$

$$\overline{\sigma}\ \overline{m}_{\tilde{x}} \qquad + \overline{\sigma' m'_{\tilde{x}}} \qquad = \overline{\sigma m_{\tilde{x}}}, \tag{4.61e}$$

$$\qquad - \overline{\sigma}\hat{\mathcal{X}} \qquad = -\overline{\sigma\mathcal{X}}. \tag{4.61f}$$

where $C = (\overline{\sigma}\hat{v})_{\tilde{x}}\,\hat{v} + \left(\overline{\sigma}\widehat{v''v''}\right)_{\tilde{x}} - \overline{(\sigma v)_{\tilde{x}} v}$. Equation (4.61) is arranged into columns. On the RHS, the first column is the time-average of σ times the buoyancy-coordinate zonal momentum equation (4.59). The second column is $\hat{u}$ times the thickness equation (4.56). The third column is the time-average of u times the thickness equation (4.56). The second through fourth columns on the RHS thus sums to zero. The LHS of (4.61) gives the buoyancy-coordinate thickness-weighted average zonal momentum equation in vector invariant form, with the second column collecting the forcing terms:

$$\hat{u}_{\tilde{t}} + \hat{\varpi}\hat{u}_{\tilde{b}} - (f + \hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})\hat{v} + \left(\overline{m} + \tfrac{1}{2}\hat{u}^2 + \tfrac{1}{2}\hat{v}^2\right)_{\tilde{x}} = \hat{\mathcal{X}} - E^u, \tag{4.62}$$

where the eddying terms are collected as

$$E^u = \overline{\sigma}^{-1}\left\{\left(\overline{\sigma}\widehat{u''u''}\right)_{\tilde{x}} + \left(\overline{\sigma}\widehat{v''u''}\right)_{\tilde{y}} + \left(\overline{\sigma}\widehat{\varpi''u''}\right)_{\tilde{b}} + \left(\tfrac{1}{2}\overline{\xi'^2}\right)_{\tilde{x}} + \left(\overline{\xi' m'_{\tilde{x}}}\right)_{\tilde{b}}\right\}. \tag{4.63}$$

The very important term $\overline{\sigma' m'_{\tilde{x}}}$ has been re-expressed in E^u as

$$\overline{\sigma' m'_{\tilde{x}}} = \left(\tfrac{1}{2}\overline{\xi'^2}\right)_{\tilde{x}} + \left(\overline{\xi' m'_{\tilde{x}}}\right)_{\tilde{b}}. \tag{4.64}$$

Here, the second term on the RHS above is the eddy interfacial form stress; it transports momentum across buoyancy surfaces, and dominates the first term on the RHS. The first three terms in E^u are Reynolds stresses. The first, second, and fourth terms in E^u transport momentum within buoyancy surfaces, whereas the third and fifth terms in E^u transport momentum across buoyancy surfaces. Young (2012) shows how E^u is actually the divergence of a vector, namely the Eliassen-Palm flux vector $\boldsymbol{E}^u$, but for our purposes, we will leave the equation in this form.

The derivation for the TWA meridional momentum equation is similar, and the final result is

$$\hat{v}_{\tilde{t}} + \hat{\varpi}\hat{v}_{\tilde{b}} + (f + \hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})\hat{u} + (\overline{m} + \tfrac{1}{2}\hat{u}^2 + \tfrac{1}{2}\hat{v}^2)_{\tilde{y}} = \hat{\mathcal{Y}} - E^v, \tag{4.65}$$

where

$$E^v = \overline{\sigma}^{-1}\left\{+\left(\overline{\sigma}\widehat{u''v''}\right)_{\tilde{x}} + \left(\overline{\sigma}\widehat{v''v''}\right)_{\tilde{y}} + \left(\overline{\sigma}\widehat{\varpi''v''}\right)_{\tilde{b}} + \left(\tfrac{1}{2}\overline{\xi'^2}\right)_{\tilde{y}} + \left(\overline{\xi' m'_{\tilde{y}}}\right)_{\tilde{b}}\right\}. \tag{4.66}$$

4.B.1.5 TWA Potential Vorticity equation

With the TWA thickness equation and the TWA momentum equations in vector invariant form (which joins the relative vorticity $\hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}}$ to the Coriolis parameter f),

the TWA potential vorticity equation is within reach. The TWA PV was defined in (4.3). Cross-differentiating (4.62) and (4.65) yields

$$(\overline{\sigma}\Pi^{\#})_{\tilde{t}} + (\hat{\varpi}\hat{v}_{\tilde{b}} + \overline{\sigma}\hat{u}\Pi^{\#})_{\tilde{x}} - (\hat{\varpi}\hat{u}_{\tilde{b}} - \overline{\sigma}\hat{v}\Pi^{\#})_{\tilde{y}} = \left[\hat{\mathcal{Y}} - E^v\right]_{\tilde{x}} - \left[\hat{\mathcal{X}} - E^u\right]_{\tilde{y}}. \quad (4.67)$$

From this we subtract $\Pi^{\#}$ times the TWA thickness equation (4.57); then dividing by $\overline{\sigma}$, yields the TWA potential vorticity equation:

$$\Pi^{\#}_{\tilde{t}} + \hat{u}\Pi^{\#}_{\tilde{x}} + \hat{v}\Pi^{\#}_{\tilde{y}} + \hat{\varpi}\Pi^{\#}_{\tilde{b}} = \frac{1}{\overline{\sigma}}\left(\left[\hat{\mathcal{Y}} - E^v - \hat{\varpi}\hat{v}_{\tilde{b}}\right]_{\tilde{x}} - \left[\hat{\mathcal{X}} - E^u - \hat{\varpi}\hat{u}_{\tilde{b}}\right]_{\tilde{y}}\right) + \frac{1}{\overline{\sigma}}\left(\overline{\sigma}\hat{\varpi}\Pi^{\#}\right)_{\tilde{b}}. \quad (4.68)$$

Since the LHS is simply the advective derivative of $\Pi^{\#}$ by the TWA velocity, we see $\Pi^{\#}$ would be conserved following the flow were it not for momentum forcing and dissipation, non-conservation of buoyancy, and eddying effects.

4.B.2 Methods

4.B.2.1 TWA numerics for Z level data

While the theoretical framework of the TWA equations is simple and attractive, applying this framework to z level model data presents its own challenges, which I now discuss.

I use PCHIPs to get z as a function of b, which then yields the depth of the surface ξ, and also its thickness $\sigma = \xi_{\tilde{b}}$ by differentiating the interpolant. (Linear interpolants have discontinuous first derivatives, so σ would exhibit unacceptable discontinuities where ξ crosses model z levels.) I also use PCHIPs to interpolate the z level horizontal velocities on the depth buoyancy surface ξ, appropriately averaged onto the U and V grids.

Linear interpolation applies a function defined solely in terms of the underlying grid to data on that grid. Hence if the data satisfies some equation (such as the buoyancy equation) then so too will its linear interpolation onto a buoyancy surface (or anywhere). This is also true of cubic spline interpolation, but not of PCHIPs. Thus, when interpolating the constituent parts of an equation (the buoyancy equation and the momentum equations), I only use linear interpolation (onto ξ previously calculated). I do not use cubic spline interpolation at all because they will often overshoot and undershoot (particularly near the pycnocline). PCHIPs preserve monotonicity of the data and are the preferred interpolation method.

4.B.2.2 Unaveraged thickness equation

Balancing the thickness equation (4.56) is the biggest challenge when using z level data to the TWA. First the buoyancy equation (4.55) must be balanced in z coordinates. Section 4.4.1.2 describes how to do this, using the salinity and potential temperature equations and thus using model diagnostic output that incorporates all the complexities of the non-linear advection scheme and time leapfrogging schemes. Call this Method A.

Method B is to calculate the buoyancy advection directly, using the 5-day archived velocity field and buoyancy field (calculated from the salinity and potential temperature). The buoyancy tendency is computed as in Method A, i.e. $b_S S_t + b_\theta \theta_t$, taking the tendency of S and θ from the 5-day model diagnostics.

In both Method A and Method B, the buoyancy forcing is taken as the residual.

Now, recall (Section 4.B.1.1) that the thickness equation is obtained by multiplying the buoyancy equation (4.55) by the thickness σ, taking a vertical derivative, and multiplying again by $(-\sigma)$. Call this manipulation Σ. So, Σ(4.55) should yield (4.56). Indeed, Σ is used to define $\sigma_{\tilde{t}} = \Sigma(b_t)$ and $(\sigma\varpi)_{\tilde{b}} = \Sigma(-\varpi)$. Moreover, $\Sigma(b_t) + \Sigma(\boldsymbol{v} \cdot \nabla b) - \Sigma(\varpi) = 0$ is guaranteed. But numerically,

$$(u\sigma)_{\tilde{x}} + (v\sigma)_{\tilde{y}} \neq \Sigma \left\{ (u\overline{b}^{x})_x + (v\overline{b}^{y})_y + (w\overline{b}^{z})_z \right\} \tag{4.69}$$

when the LHS above is computed using the PCHIP interpolation described earlier (but other methods for calculating u,v were tested). There are large errors here due to the C grid, and this is the key problem I am addressing.

Consider that SOSE, working on the C grid, actually solves[9]

$$b_t + (u\overline{b}^{x})_x + (v\overline{b}^{y})_y + (w\overline{b}^{z})_z = \varpi. \tag{4.70}$$

The coordinate-decorated overbar is not really an average like in constructing the TWA, but simply indicates averaging neighbouring data points onto a different grid: e.g. $\overline{b}^{x}$ averages b onto the u grid. Multiplying (4.70) by $\sigma = b_z^{-1}$ will not get w standing alone for the continuity equation to later annihilate.

For two variables α and β on the same grid, observe that

$$(\alpha\beta)_z = \alpha_z \overline{\beta}^{z} + \overline{\alpha}^{z} \beta_z, \tag{4.71}$$

[9]This also is a simplification. It ignores the curvilinear grid and partial-cells, but moreover it is the second-order centred advection scheme, whereas SOSE uses a third order upwind advection scheme. But the zeroth order problem is that the grid is not amenable to getting the thickness.

regardless of grid spacing, or even whether the derivatives are forward, backward, or centred differences. Thus,

$$(w\overline{b}^z)_z = w_z\overline{b}^{zz} + \overline{w}^z\overline{b}^z_z, \tag{4.72}$$

which reveals we have two problems. The first term combines with similar terms $u_x\overline{b}^{xx} + v_y\overline{b}^{yy}$ which the continuity equation cannot eliminate.

As for the second term, one could multiply the buoyancy equation (on the T grid) by $1/\overline{b}^z_z$, also on the T grid, which gets $\overline{w}^z$ alone. However, the continuity equation can only eliminate w_z, not $\overline{w}^z_z$.

That is one of *many* possible ways to calculate either of these expressions. The results do differ noticeably. After much testing (to minimize errors in (4.69)), I instead settled on the following method for Σ. Average the buoyancy equation onto the W grid and divide by b_z calculated by first difference onto the W grid. Then the vertical derivative, a first difference, brings the result back to the T grid. Finally, this is interpolated (linearly) onto the buoyancy surface, and multiplied by $\sigma = \xi_{\tilde{b}}$.

4.B.2.3 TWA momentum and thickness equations

The horizontal momentum equations in z coordinates are balanced, using SOSE's 5-day archived diagnostics. These are transformed to buoyancy coordinates using (4.60), except σb_x is computed as $-\xi_{\tilde{x}}$ and similarly $\sigma b_y = -\xi_{\tilde{y}}$.

Next, we take thickness-weighted averages. The TWA thickness equation (4.57) follows from Σ(4.55) easily. The TWA of the buoyancy coordinate momentum equation gives the first column on the RHS of (4.61). Further averages yield $\hat{u}$, $\hat{v}$, $\frac{1}{2}\overline{\xi^2}$, and (averaging (4.56) onto the U and V grids) the third column of (4.61). The vorticity and kinetic energy terms (4.61c) and (4.61d) are summed together; hence C is not computed, but more importantly, the advective part of the thickness equation, $(\sigma u)_{\tilde{x}} + (\sigma v)_{\tilde{y}}$, need not be split, thereby enabling use of the RHS of (4.69). (This staves off errors, but only for a while.)

The $\hat{u}_{\tilde{t}}$ term is computed as the residual from (4.61a), and the Reynolds terms are computed as residuals from (4.61b)–(4.61d). The interfacial form stress is obtained from (4.64), with $\overline{\sigma' m'_{\tilde{x}}}$ the residual of (4.61e).

For (4.61b) we require $\hat{u}_{\tilde{b}}$. Having computed $\hat{u}$ only as 2D fields on a particular $\tilde{b}$ surface, not as 3D fields, $\hat{u}_{\tilde{b}}$ is not readily available. One option is to expand $\hat{u}_{\tilde{b}}$, wary that the thickness-weighted average operator does not commute with the partial derivatives, to get $\hat{u}_{\tilde{b}} = \overline{\sigma}^{-1}(\overline{\sigma_{\tilde{b}}u} + \overline{\sigma u_{\tilde{b}}} - \overline{\sigma_{\tilde{b}}}\hat{u})$. This requires $\sigma_{\tilde{b}} = \xi_{\tilde{b}\tilde{b}}$; since the PCHIP interpolation is only C1, the second derivative $\xi_{\tilde{b}\tilde{b}}$ will be discontinuous. Reluctantly, I compute $\hat{u}$ and $\hat{v}$ on a second, very nearby buoyancy surface and take a first difference.

All parts of the TWA PV budget (4.68) are now assembled. For interpolation of forcings along streamlines, some fancy handling is advisable to keep vorticity things on the vorticity grid and thickness things on the tracer grid. However, we won't get that far, because the TWA PV budget will not come close to balancing.

To illustrate this, let's simply say $\Pi^\#$ is on the tracer grid, i.e. the vorticity is averaged four ways, then divided by the thickness. Now, the in-buoyancy surface advection is, in theory,

$$\hat{u}\Pi^\#_{\tilde{x}} + \hat{v}\Pi^\#_{\tilde{y}} = \frac{\big((f+\hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})\hat{u}\big)_{\tilde{x}} + \big((f+\hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})\hat{v}\big)_{\tilde{y}}}{\overline{\sigma}} - \frac{\Pi^\#}{\overline{\sigma}}\Big((\overline{\sigma}\hat{u})_{\tilde{x}} + (\overline{\sigma}\hat{v})_{\tilde{y}}\Big). \quad (4.73)$$

On the LHS, the first term is from the vorticity equation (4.67), and the second from the thickness equation (4.57). Critical to PV conservation is that the stretching terms in these equations cancel. However, in the zonal momentum equation on the C grid u gets averaged (4 point stencil) onto the V grid (denoted $\overline{u}^{xy}$) as does vorticity (2 point stencil), and similarly for the meridional equation. By continuity, $u_x + v_y = -w_z$ and this stretching term is part of a very fine balance, which expanding the stencil destroys. Taking care with the C grid, (4.73) becomes

$$\begin{aligned}\hat{u}_{rot}\Pi^\#_{\tilde{x}} + \hat{v}_{rot}\Pi^\#_{\tilde{y}} + \hat{u}_{div}\Pi^\#_{\tilde{x}} + \hat{v}_{div}\Pi^\#_{\tilde{y}} + \epsilon^{(Stretch)} + \epsilon^{(Advec)} \\ = \frac{1}{\overline{\sigma}}\overline{\Big(\overline{(f+\hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})}^{x}\overline{\hat{u}}^{xy}\Big)_{\tilde{x}} + \Big(\overline{(f+\hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})}^{y}\overline{\hat{v}}^{xy}\Big)_{\tilde{y}}}^{xy} - \frac{\Pi^\#}{\overline{\sigma}}\left((\overline{\overline{\sigma}}^{x}\hat{u})_{\tilde{x}} + (\overline{\overline{\sigma}}^{y}\hat{v})_{\tilde{y}}\right).\end{aligned} \quad (4.74)$$

The TWA velocity has been decomposed into rotational and divergent parts and these separately applied to the PV advection, and room has been left for two errors. The first is the numerical failure of the stretching cancellation, namely

$$\epsilon^{(Stretch)} = +\frac{1}{\overline{\sigma}}\overline{\overline{(f+\hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})}^{xx}\overline{\hat{u}_{\tilde{x}}}^{xy} + \overline{(f+\hat{v}_{\tilde{x}} - \hat{u}_{\tilde{y}})}^{yy}\overline{\hat{v}_{\tilde{y}}}^{xy}}^{xy} - \frac{\Pi^\#}{\overline{\sigma}}\left(\overline{\overline{\sigma}}^{xx}\hat{u}_{\tilde{x}} + \overline{\overline{\sigma}}^{yy}\hat{v}_{\tilde{y}}\right), \quad (4.75)$$

having used (4.71). The second, $\epsilon^{(Advec)}$, is taken as a residual of (4.74) and thus captures all sorts of errors, but is predominantly due to the thickness advection error (4.69).

4.B.3 Results

We now examine the terms in the TWA PV budget (4.68), with the in-buoyancy-surface advective terms broken as (4.74). Figure 4.9 shows $\Pi^\#$ and all these forcings. All terms except the rotational advection are moved to the RHS, so they force $\Pi^\#$ to increase moving along a rotational streamline. The rotational advection minus

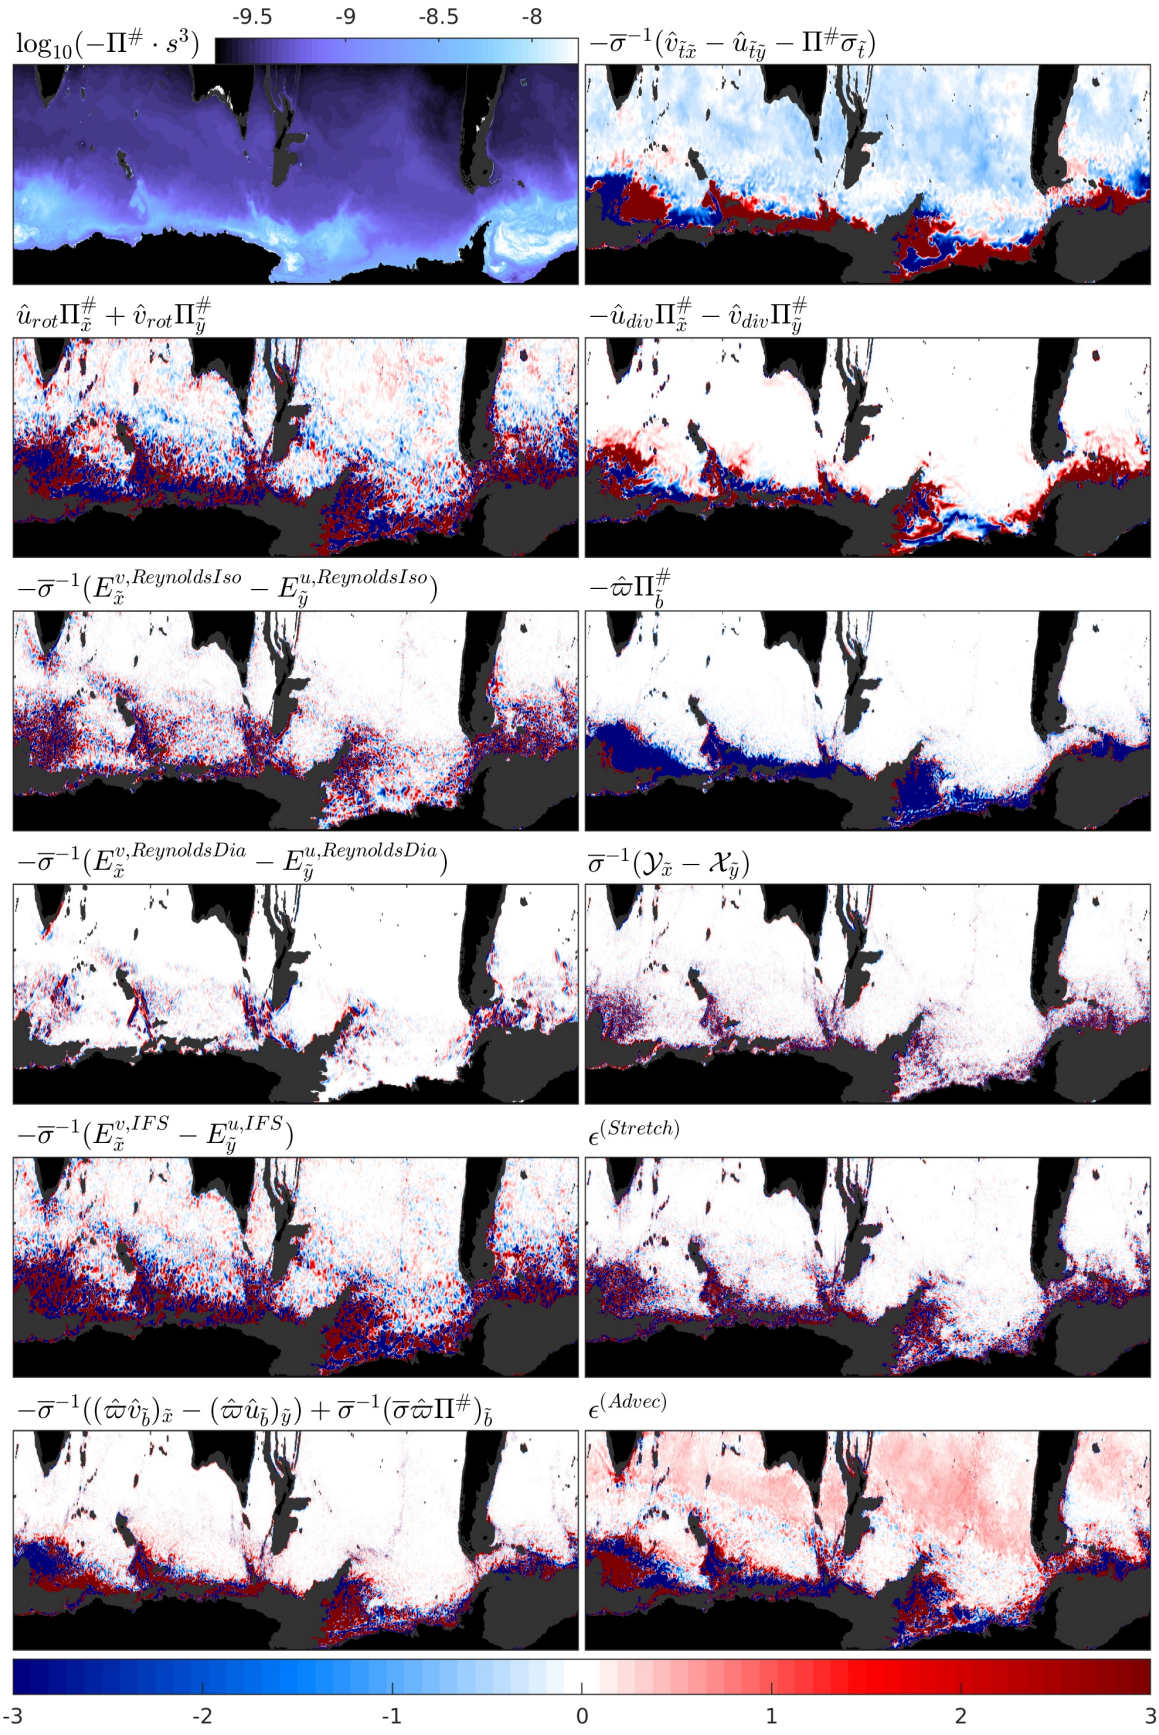


Figure 4.9: Caption on following page.

Figure 4.9: Thickness-weighted average potential vorticity (upper left) and its budget (all else, $10^{-18}\,\mathrm{s}^{-4}$), on the buoyancy surface $z_0 = -1000\,\mathrm{m}$. Top to bottom then left to right, the panels show the following: the PV, rotational advection of PV, "isopycnal" Reynolds stresses, "diapycnal" Reynolds stresses, interfacial form stress, vertical shear, PV tendency, divergent advection of PV, "diapycnal" advection of PV, momentum forcing, stretching error, and advection error. See text for full definitions. PV budget terms are smoothed with a Gaussian filter of 5 by 5 grid points.

all other terms plotted is not identically zero: a small residual remains from the Montgomery potential, $\overline{\Phi_{x\tilde{y}}} - \overline{\Phi_{y\tilde{x}}} \neq 0$, recalling from (4.60e); by our choice of buoyancy (4.23), however, this is error is very small, of order $0.01 \times 10^{-18}\,\mathrm{s}^{-4}$ away from the sea-surface, which is much less than 1% of the dominant terms in the PV equation (Fig. 4.9). The eddy terms are split as in-buoyancy-surface Reynolds stresses ($E^{u,ReynoldsIso}$ the first, second, and fourth terms in (4.63)), across-buoyancy-surface Reynolds stresses ($E^{u,ReynoldsDia}$ the third term in (4.63)), and interfacial form stress ($E^{u,IFS}$ the last term in (4.63)).

Though the TWA follows the same buoyancy surface from start to finish, the thickness of this layer can still drift throughout the model run, and evidently this remains significant. Advection by the divergent (think ageostrophic) flow remains small in ocean interior away from boundaries. Reynolds stresses within the buoyancy surface are surprisingly significant, but the interfacial form stress remains dominant. Unfortunately, the errors in PV advection are among the largest terms in the PV budget. The stretching error is significant (it is an unbearable mess with no smoothing) but the thickness error is larger.

Figure 4.10 shows a subset of the panels in Fig. 4.9 but for a deeper buoyancy surface, having $z_0 = -2000\,\mathrm{m}$. The advection error shows an incredible bipolar pattern throughout the north side of the ACC, reminiscent of the MVV advection (Fig. 4.6) but cleaner and spanning almost all longitudes. Unlike in the Eulerian diagnostics, this bipolar pattern appears only in the deep ocean below $-1000\,\mathrm{m}$.

Up to here, buoyancy advection used Method A, using SOSE's archived salt and temperature fluxes. Now consider Method B, calculating $\boldsymbol{v} \cdot \nabla b$ directly using the 5-day archived velocity and buoyancy. With this, the advective error seen in Fig. 4.11 (right panels) has been considerably reduced, suggesting the numerics I have used for the coordinate transformation are not too awful. Some structure remains in the dipole region. But this is not a panacea: this breaks the delicate balance in the buoyancy equation (4.55); the buoyancy forcing (as a residual) is increased by over

an order of magnitude. As such, the PV errors have been shifted to the shear terms (Fig. 4.11 left panels).

This implies that a proper TWA analysis cannot be done with SOSE's Eulerian 5-day averaged archive data. Higher frequency output is needed, and ideally a TWA analysis would be done online. The $z_0 = -1000\,\mathrm{m}$ buoyancy surface, for example, can change by up to 200 m from one 5-day window to the next, in SOSE's ACC.

Even still, some errors would remain, due to the transformation into buoyancy coordinates. These errors ($\epsilon^{(Advec)}$ in Fig. 4.11) are smaller than the errors caused by 5-day averaging *and* this transformation ($\epsilon^{(Advec)}$ in Fig. 4.9 and Fig. 4.10), but are still considerable (much larger than the PV forcing by diapycnal terms, for example).

4.B.4 Helmholtz for the thickness equation

One creative approach might be to use a Helmholtz decomposition to fudge the velocity field so as to satisfy the thickness equation, i.e. to make (4.69) an equality. This can be done with some Helmholtz trickery, first by decomposing

$$\sigma \boldsymbol{u} = \nabla\chi + \nabla \times \psi \tag{4.76}$$

where $\boldsymbol{u}$ is the horizontal velocity interpolated (with PCHIPs, say) onto the buoyancy surface. If (4.69) were an equality, then the divergence of (4.76) gives a Poisson equation,

$$\Sigma(\boldsymbol{v} \cdot \nabla b) = \nabla^2\chi. \tag{4.77}$$

Solve this Poisson equation for a new χ', using Dirichlet boundary conditions[10] specified by the original χ wherever the buoyancy surface has outcropped or incropped. Then define a new velocity field,

$$(\sigma u, \sigma v) = \nabla\chi' + \nabla \times \psi. \tag{4.78}$$

This differs slightly from the original velocity field because $\chi \neq \chi'$, but most of both velocity fields is carried by the rotational term $\nabla\times\psi$. Fortunately, this doesn't change the rotational streamlines at all, but again we won't get that far.

[10] A Neumann boundary condition, $\nabla\chi \cdot \hat{\boldsymbol{n}} = 0$ where $\hat{\boldsymbol{n}}$ is the unit vector normal to the domain, specifies zero flow across the domain boundary. This is commonly used when in z coordinates because the domain boundary is rock, and it is helpful to assume that the no-normal flow boundary condition applies to both the divergent and the rotational parts of the flow. In b coordinates however, the domain boundary is more nebulous: when a layer outcrops, it still exists but has zero thickness. If the layer has $h_{i,j} > 0$ but has outcropped at $(i-1, j)$, then $h_{i-1,j} = 0$; but $u_{i,j}(h_{i,j}+h_{i-1,j})/2 \neq 0$, so there is non-zero normal flow to the outcropping or incropping contour. Furthermore, our attempts to use Neumann boundary conditions commonly created a divergent flow $\nabla\chi'$ that was almost as large as the rotational flow $\nabla \times \psi$.

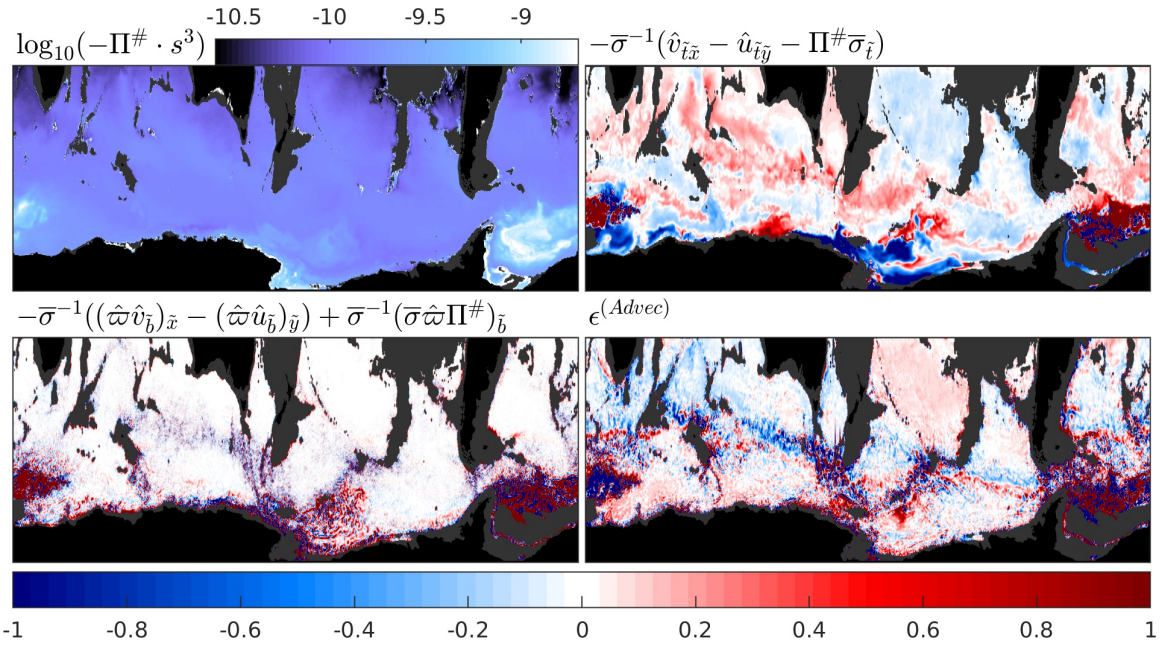


Figure 4.10: As in Fig. 4.9 but for $z_0 = -2000\,\mathrm{m}$ and only a selection of panels.

This trickery has changed u and v and hence $\hat{u}$ and $\hat{v}$. It is a tiny relative change, but when multiplied by the large term f on the LHS of (4.61c), it becomes very significant. The RHS has $-\overline{\sigma f v}$, which cannot similarly adapt without breaking the momentum budget, the first column of (4.61). So, rather than fixing the thickness advection error, this attempt has just sent it to the term in the TWA horizontal momentum equation that is computed as a residual: the Reynolds term, $(\overline{\sigma}\widehat{v''u''})_{\tilde{y}}$. This is entirely unphysical, so this attempt must be abandoned.

4.B.5 Discussion

Recent studies have used the TWA equations (Cessi and Wolfe, 2013; Aoki et al., 2016), but they do not simultaneously balance the momentum and the thickness equations. As such, one cannot properly analyse the potential vorticity. I have attempted such an analysis but found it impossible to simultaneously balance the momentum and thickness equations in buoyancy coordinates, using 5-day Eulerian averaged z level model outputs.

Despite testing many interpolation methods to get the horizontal velocity, layer depth, and layer thickness, I find that computing the buoyancy forcing term in the thickness equation as a residual typically magnifies it by at least an order of magnitude compared to its direct computation using the true buoyancy forcing of the model (computed from the model's potential temperature and salinity forcings).

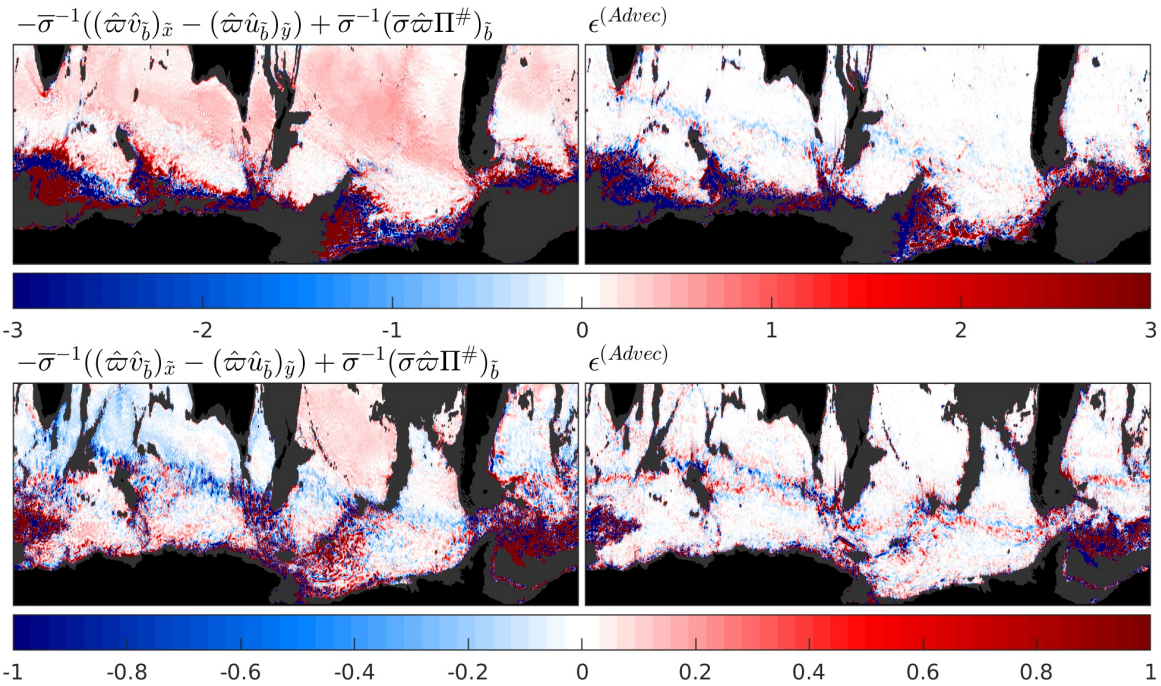


Figure 4.11: Just the shear forcing and advective error in the TWA PV budget, but here the z coordinate buoyancy advection ($\boldsymbol{v} \cdot \nabla b$) was computed using the offline velocity and buoyancy fields, on buoyancy surfaces $z_0 = -1000\,\mathrm{m}$ (top row, for comparison with Fig. 4.9), or $z_0 = -2000\,\mathrm{m}$ (bottom row, for comparison with Fig. 4.10).

Both Cessi and Wolfe (2013) and Aoki et al. (2016) construct the 3D TWA velocity, $\boldsymbol{u}^{\#}$, to be divergence-free: the horizontal TWA velocity is computed, its divergence calculated, and the vertical velocity obtained by vertical integration. This does not, however, translate into buoyancy coordinates to give a diapycnal velocity ϖ that is as small as the model's buoyancy forcing. Cessi and Wolfe (2013) recognize a "striking cancellation" between the mean and eddy-induced diapycnal velocity in their model eastern boundary current. While this may be true, correct online computation of the TWA variables would likely have revealed an even more strikingly small diapycnal velocity. They develop a TWA PV budget but only by ignoring the diapycnal velocity (among other terms), which could be far larger in their budgets than they anticipate, and thus not necessarily negligible. Increasing the diapycnal diffusion by a factor of 10-100 in an ocean model produces a qualitatively different circulation (Bryan, 1987; Vallis, 2000; Park and Bryan, 2000; Saenko and Merryfield, 2005; Stanley and Saenko, 2014). Therefore, TWA manipulations of z level model data which do not accurately balance the thickness equation should be regarded with suspicion.

Ringler et al. (2016) take a different approach for their idealized model, setting explicit buoyancy forcing to zero in the ocean interior. Thus, they can diagnose the

diabatic velocity as due to surface forcing only, impacting a buoyancy layer only when it outcrops. As they note, this misses numerical errors due to advection in the interior (implicit mixing). Nonetheless, their effort to compute the TWA online must be followed further if a TWA analysis is to be trusted.

The TRM-II theory of McDougall and McIntosh (2001) may be worth investigating further, as it is built to work with Eulerian averages, yet it captures much of the spirit of the buoyancy-following averages. I have not pursued this here primarily because TRM-II requires the variance of a quasi-conservative density variable, which should be computed online. This is far easier to do than to implement the full TWA analysis online, but nonetheless is not available in the standard SOSE outputs (perhaps not least because it raises the perennial question of which quasi-conservative density variable to choose). Computing such a variance offline, for example from SOSE's 5-day Eulerian averages, may lead us into a similar quagmire as for the TWA computed offline. Moreover, care must be taken when discretising the TRM-II theory for the C grid: the quasi-Stokes stream function is defined in terms of quantities on the tracer grid, yet is used to evaluate horizontal velocities. Some of the TRM-II theory is discussed in the following Appendix.

Finally, a potential issue with TWA is that the PV becomes asymptotically large when a layer is outcropped or incropped for all but a single time step, since the layer thickness is taken as zero for other time steps, following the boundary conditions suggested by Young (2012).

4.C Alternative derivation of the eddy linear vorticity balance

4.C.1 Style 1: Integration between isopycnals

In this section we show how the instantaneous linear vorticity balance, (4.4), may be vertically integrated between two surfaces, time-averaged, and then vertically differentiated to produce a residual linear vorticity balance. For notational convenience, let $\boldsymbol{u} = u\hat{\boldsymbol{i}} + v\hat{\boldsymbol{j}}$ be the horizontal velocity, and $\boldsymbol{v} = \boldsymbol{u} + w\hat{\boldsymbol{k}}$ be the 3D velocity.

The surfaces we integrate between come from the Temporal-Residual-Mean (TRM) framework (McDougall and McIntosh, 2001). I will not fully review the TRM theory but give enough that this derivation is self-contained. Defining the surfaces requires considerable care, and is as follows. For every point in the ocean (x, y, z), find $\widetilde{\gamma}(x, y, z)$ such that the depth of the isopycnal $\widetilde{\gamma}(x, y, z)$ in water column (x, y)

has a time-mean depth of z. In other words, $\xi = \xi(x, y, z, t)$ is the instantaneous depth of the $\widetilde{\gamma}(x, y, z)$ isopycnal in water column (x, y), and the perturbation depth, $z'(x, y, z, t) = \xi(x, y, z, t) - z$, is such that $\overline{z'(x, y, z, t)} = 0$ for all (x, y, z). With a chosen z, $\xi(x, y, z, t)$ indeed defines a surface, but it is not the depth of a particular isopycnal, and nor is it a fixed depth; rather it is a sort of hybrid between these two modes of thinking.

To help familiarize the reader with the notation, a tilde over a variable indicates that variable is time-averaged following the undulating surface, whereas the overbar indicates an Eulerian time-mean taken at fixed depth. The curvy nature of the tilde vs the straight overbar might serve as a reminder to this difference. Happily, when the instantaneous neutral density field $\gamma(x, y, z, t)$ is time-averaged following the surface, we indeed get the value $\widetilde{\gamma}(x, y, z)$. Observe that $\widetilde{\boldsymbol{\tau}}(z_1) = \overline{\boldsymbol{\tau}_1}$ because the RHS has its z dependence specified to be ξ_1 before the (Eulerian) time-average is taken.

Let us choose two fixed depths z_1 and z_2, then find associated time-dependent surfaces $\xi(x, y, z_1, t)$ and $\xi(x, y, z_2, t)$. As shorthand, write $\xi_1 \equiv \xi(x, y, z_1, t)$, $p_1 \equiv p(x, y, \xi_1, t)$, $\boldsymbol{\tau}_1 \equiv \boldsymbol{\tau}(x, y, \xi_1, t)$, and similarly for subscripts "2". Capital letters denote vertically integrated variables, e.g. $\boldsymbol{U} = \int_{\xi_1}^{\xi_2} \boldsymbol{u}\, \mathrm{d}z$, and $P = \int_{\xi_1}^{\xi_2} p\, \mathrm{d}z$. For brevity, the x, y, and t arguments are sometimes dropped, but generally the z argument is maintained (except where subscript "1" or "2" handles that).

The first step in this derivation is to integrate (4.8) from ξ_1 to ξ_2, obtaining[11]

$$\rho_c f \hat{\boldsymbol{k}} \times \boldsymbol{U} = -(\nabla P - p_2 \nabla \xi_2 + p_1 \nabla \xi_1) + (\boldsymbol{\tau}_2 - \boldsymbol{\tau}_1) + \boldsymbol{A}, \tag{4.79}$$

where $\boldsymbol{a}$ (which defines $\boldsymbol{A}$) has been included on the RHS of (4.8) to capture all terms that were neglected there. Second, cross-differentiate (4.79) to obtain the vertically integrated linear vorticity equation:

$$\rho_c \beta V = -\rho_c f \nabla \cdot \boldsymbol{U} + \nabla \times (p_2 \nabla \xi_2 - p_1 \nabla \xi_1) + \nabla \times (\boldsymbol{\tau}_2 - \boldsymbol{\tau}_1) + \nabla \times \boldsymbol{A}. \tag{4.80}$$

Third, time-average (4.80), which yields

$$\rho_c \beta \overline{V} = -\rho_c f \nabla \cdot \overline{\boldsymbol{U}} + \nabla \times (\boldsymbol{\tau}^+(z_2) - \boldsymbol{\tau}^+(z_1)) + \nabla \times (\widetilde{\boldsymbol{\tau}}(z_2) - \widetilde{\boldsymbol{\tau}}(z_1)) + \nabla \times \overline{\boldsymbol{A}}. \tag{4.81}$$

Here, the horizontal eddy form stress is $\boldsymbol{\tau}^+(z_1) = \overline{p_1' \nabla z_1'}$ where $p_1' = p_1 - \overline{p_1}$ and $z_1' = \xi(z_1) - z_1$. (Note, p_1' is not the instantaneous p at (x, y, z_1, t) minus the Eulerian time-mean of p at (x, y, z_1, t).)

[11] While variables such as ξ_1 are functions of time t, the gradient operator ∇ only acts on spatial dimensions.

The total northward transport between our surfaces is

$$\overline{V} = \overline{\int_{z_1}^{z_2} v \, \mathrm{d}z} = \int_{z_1}^{z_2} \overline{v} \, \mathrm{d}z + \Psi^y(z_2) - \Psi^y(z_1), \tag{4.82}$$

where

$$\boldsymbol{\Psi}(x, y, z) = \overline{\int_{z}^{\xi(x,y,z,t)} \boldsymbol{u} \, \mathrm{d}z} \tag{4.83}$$

is the quasi-Stokes stream function (McDougall and McIntosh, 2001), capturing the net eddying horizontal transport in the water column (x, y) between the mean-depth z of an isopycnal—namely $\widetilde{\gamma}(x, y, z, t)$ as previously discussed—and its instantaneous depth ξ.

The first term on the RHS of (4.81) is the convergence of flow into the water column between our surfaces. Using the continuity equation $\nabla \cdot \boldsymbol{u} + \partial_z w = 0$, this is

$$-\nabla \cdot \boldsymbol{U} = -\boldsymbol{u}_2 \cdot \nabla \xi_2 + \boldsymbol{u}_1 \cdot \nabla \xi_1 - \int_{\xi_1}^{\xi_2} \nabla \cdot \boldsymbol{u} \, \mathrm{d}z = F_2 - F_1, \tag{4.84}$$

where, $F_1 = w_1 - \boldsymbol{u}_1 \cdot \nabla \xi_1$ and $F_2 = w_2 - \boldsymbol{u}_2 \cdot \nabla \xi_2$. If the boundary were fixed in time, F_1 is the flow into the volume across the bottom boundary, and F_2 is the flow out of the volume across the top boundary. It is apparent that $-\nabla \cdot \boldsymbol{U}$ represents the stretching of vortex tubes. In the TRM theory, the eddy-induced vertical velocity w^+ is defined to be the horizontal gradient (at constant z) of $-\boldsymbol{\Psi}$; that is, $w^+ \equiv -\nabla_z \cdot \boldsymbol{\Psi}$. Observe that

$$\nabla_z \cdot \boldsymbol{\Psi} = \overline{\boldsymbol{u}|_\xi \cdot \nabla \xi} + \overline{\int_z^\xi \nabla_z \cdot \boldsymbol{u} \, \mathrm{d}z} = \overline{\boldsymbol{u}|_\xi \cdot \nabla \xi} - \overline{w(\xi)} + \overline{w(z)}. \tag{4.85}$$

The last term on the RHS above is the time-mean vertical velocity at depth z, i.e. the Eulerian-mean vertical velocity $\overline{w}$; the second last term is the time-mean vertical velocity following the undulating isopycnal, $\widetilde{w}$. Thus, we see the sum of the Eulerian-mean and eddy-induced vertical velocities gives the flow through the ξ surface:

$$\overline{w} + w^+ = \widetilde{w} - \overline{\boldsymbol{u}|_\xi \cdot \nabla \xi}. \tag{4.86}$$

The RHS, evaluated at z_1, is the time mean of F_1, and likewise for z_2. Thus, we find

$$-\nabla \cdot \boldsymbol{U} = \left(\overline{w}(z_2) + w^+(z_2)\right) - \left(\overline{w}(z_1) + w^+(z_1)\right). \tag{4.87}$$

Collecting (4.82), (4.83), and (4.87) into (4.81) yields

$$\begin{aligned} \rho_c \beta \left(\int_{z_1}^{z_2} \overline{v} \, \mathrm{d}z + \Psi^y(z_2) - \Psi^y(z_1) \right) =& \rho_c f \left(\overline{w}(z_2) - \overline{w}(z_1) + w^+(z_2) - w^+(z_1) \right) \\ &+ \nabla \times (\boldsymbol{\tau}^+(z_2) - \boldsymbol{\tau}^+(z_1)) \\ &+ \nabla \times (\widetilde{\boldsymbol{\tau}}(z_2) - \widetilde{\boldsymbol{\tau}}(z_1)) \\ &+ \nabla \times \overline{\boldsymbol{A}}. \end{aligned} \tag{4.88}$$

Finally, we divide (4.88) by $(z_2 - z_1)$ and take the limit as $z_2 - z_1 \to 0$, yielding a "residual linear vorticity balance":

$$\beta\left(\overline{v} + v^+\right) = f\frac{\partial}{\partial z}\left(\overline{w} + w^+\right) + \frac{1}{\rho_c}\nabla\times\left(\widetilde{\boldsymbol{\tau}}_z + \boldsymbol{\tau}_z^+\right) + \frac{1}{\rho_c}\frac{\partial}{\partial z}\nabla\times\ \overline{\boldsymbol{A}}. \tag{4.89}$$

This reprises (4.16), but with a different formulation of the eddy-induced velocity, and the forcing terms.

4.C.2 Style 2: Fast, separable derivation

The following is gives a faster route to the residual vorticity balance using the TRM-II theory. Having defined the quasi-Stokes stream function (4.83), the eddy-induced horizontal velocity is $\boldsymbol{u}^+ = \partial_z\boldsymbol{\Psi}$, and the eddy-induced vertical velocity is $w^+ = -\nabla_z\cdot\boldsymbol{\Psi}$. Hence the 3D eddy-induced velocity $\boldsymbol{v}^+ = \boldsymbol{u}^+ + w^+\hat{\boldsymbol{k}}$ is divergence-free, and thus

$$\beta v^+ = f\partial_z w^+ + \nabla_z\cdot(f\boldsymbol{u}^+), \tag{4.90}$$

as in Section 4.2.1. The momentum equation (4.8) is used to deal with $\nabla_z\cdot(f\boldsymbol{u}^+)$ as follows:

$$\begin{aligned}
\nabla_z\cdot(f\boldsymbol{u}^+) &= \frac{\partial}{\partial z}\nabla_z\cdot\overline{\int_z^\xi f\boldsymbol{u}\,\mathrm{d}z} \\
&= \frac{1}{\rho_c}\frac{\partial}{\partial z}\nabla_z\times\left(-\overline{\int_z^\xi\nabla_z p\,\mathrm{d}z} + \overline{\int_z^\xi\boldsymbol{\tau}_z\,\mathrm{d}z} + \overline{\int_z^\xi\boldsymbol{a}\,\mathrm{d}z}\right) \\
&= \frac{1}{\rho_c}\frac{\partial}{\partial z}\nabla\times\left(\left[-\nabla_z\overline{\int_z^\xi p\,\mathrm{d}z'} + \overline{p(\xi)\nabla_z\xi} - \overline{p(z)\nabla_z z}\right] + \left[\overline{\boldsymbol{\tau}(\xi)} - \overline{\boldsymbol{\tau}(z)}\right] + \mathcal{A}\right) \\
&= \frac{1}{\rho_c}\nabla\times\left(\boldsymbol{\tau}_z^+ + \widetilde{\boldsymbol{\tau}}_z - \overline{\boldsymbol{\tau}}_z + \boldsymbol{a}^+\right).
\end{aligned} \tag{4.91}$$

Again, the overtilde notation is used to represent the time-mean following the undulating surface; for example, $\widetilde{p}(x,y,z,t) = \overline{p(x,y,\xi(x,y,z,t),t)}$, or in short-hand, $\widetilde{p} = \overline{p(\xi)}$. The eddy form stress is $\boldsymbol{\tau}^+ = \overline{p(\xi)\nabla_z\xi}$, which, if desired, may be written as $\boldsymbol{\tau}^+ = \overline{p'\nabla_H z'}$ where$p' = \underline{p(\xi) - \widetilde{p}}$ and $z' = \xi - z$. The extraneous terms in $\boldsymbol{a}$ are handled by defining $\mathcal{A} = \int_z^\xi\boldsymbol{a}\,\mathrm{d}z$, analogous to the quasi-Stokes stream function (4.83), and defining $\boldsymbol{a}^+ = \partial_z\mathcal{A}$ analogous to $\boldsymbol{u}^+ = \partial_z\boldsymbol{\Psi}$. It is possible to decompose this term but, considering $\boldsymbol{a}$ of secondary importance, that is not pursued that here.

Combining (4.90) and (4.91) gives the linear vorticity relation for the eddy-induced velocities, that is, a "linear eddy vorticity balance":

$$\beta v^+ = f\partial_z w^+ + \rho_c^{-1}\nabla\times\left(\boldsymbol{\tau}_z^+ + \widetilde{\boldsymbol{\tau}}_z - \overline{\boldsymbol{\tau}}_z + \boldsymbol{a}^+\right). \tag{4.92}$$

Adding (4.5) and (4.92) gives the a linear residual (mean plus eddy) vorticity balance:

$$\beta(\overline{v} + v^+) = f\partial_z(\overline{w} + w^+) + \rho_c^{-1}\nabla \times \left(\widetilde{\boldsymbol{\tau}}_z + \boldsymbol{\tau}_z^+ + \overline{\boldsymbol{a}} + \boldsymbol{a}^+\right) \tag{4.93}$$

This is identical to (4.89) but with the extra terms in $\boldsymbol{a}$ presented slightly differently.

Again, this derivation has stressed that the linear vorticity balance for the residual flow is in fact two separate equations rolled into one. The only connection between the two is the cancellation of the Eulerian time-averaged vertical stress, $\overline{\boldsymbol{\tau}}$.

Chapter 5

Conclusion

In this thesis, I have explored three common assumptions in oceanography. These assumptions are the following: (1) the flow is nearly confined to neutral surfaces; (2) potential vorticity is uniform on neutral surfaces; and (3) potential vorticity is conserved following the flow—a less restrictive case of assumption (2). I set out to investigate the nature of neutral surfaces, and, focusing on the Southern Ocean, explored the degree of potential vorticity homogenization on approximately neutral surfaces as well as its variation along streamlines on such surfaces. A topological perspective proved extremely useful for the study of neutral surfaces, and the same ideas can apply to the studies of potential vorticity, though this has not been pursued to its full extent.

5.1 Key results

Chapter 2 explored the topological nature of neutral surfaces. The defining equation of neutral surfaces dictates that gradients of salinity and potential temperature on a neutral surface are aligned, or equivalently that gradients of specific volume and pressure on a neutral surface are aligned. This implies there is a functional relation between specific volume and pressure, because their contours are aligned. However, as there can be multiple disjoint contours for a given value of the pressure, say, the specific volume is actually a multivalued function of the pressure. Thus, the topological analysis of neutral surfaces focuses on the connected components of level sets of pressure on a neutral surface. The different branches of this multivalued function are encoded in the Reeb graph, which is formally defined as a topological quotient space but in practice is used in a graph theoretical context. The arcs of the Reeb graph correspond to different regions in physical space upon which the specific volume is a single-valued function of pressure, on the neutral surface. The nodes of the Reeb

graph correspond to critical points—maxima, minima, and saddles—of the pressure on the neutral surface. An arc exists between two nodes when there is a continuous path through physical space connecting the two nodes' corresponding critical points, that is monotonic in pressure and that does not intersect any pressure contour emanating from any other critical point. Thus, the Reeb graph encodes all the necessary topological information about how specific volume is a multivalued function of pressure on a neutral surface: namely, the domains of the different branches (where it is single-valued), and which branches meet and where, nesting into a global structure.

This topological understanding of neutral surfaces is the foundation of the geographic dependence of density variables whose isosurfaces are approximately neutral. Knowing the regions where specific volume is a single-valued function of pressure, this function can be fit empirically, subject to matching conditions where different branches of the multivalued function meet. A new surface, called a topobaric surface, can then be determined wherein these functions are satisfied exactly. Topobaric surfaces can be rapidly computed and are approximately neutral to high degree. Topobaric surfaces also possess an exact geostrophic stream function, namely the exact geostrophic stream function for neutral surfaces, which McDougall (1989) proved must exist but a closed form expression for it had not been known, until now.

The defining equation of neutral surfaces is of a general form that appears elsewhere in geophysical fluid dynamics. For example, for an inviscid and unforced (mechanically and thermodynamically) fluid in steady state, this general relation also applies between the potential vorticity and Bernoulli potential. Their gradients are aligned, and thus there is a multivalued functional relation between the two. Chapter 3 explored this functional relationship in the Southern Ocean, though treating it as a single-valued function for a first study. In particular, the slope of potential vorticity as a function of Bernoulli potential was non-dimensionalised to yield a Mach number for Rossby waves. This Mach number gives insight about potential vorticity homogenization and shear instability. The statistics of potential vorticity on an approximately neutral surface cannot alone characterize homogenization, as they are dimensional. In contrast, potential vorticity homogenization is characterized by the Mach number tending to infinity, which occurs when potential vorticity gradients, being the restoring force for Rossby waves and thus proportional to their phase speed, tend to zero. Moreover, shock waves are generally associated when a Mach number crosses unity, and Dowling (2014) speculates that shocks of Rossby waves are the onset of shear instability, though this remains to be proven. Nonetheless, the limits of Arnol'd's two branches of shear stability correspond to the Mach number crossing

zero and unity, further stimulating interest in this type of analysis. However, the phase speed of Rossby waves is considerably harder to calculate in a three dimensional fluid than a quasi-geostrophic system. The traditional method involves solving an eigenvalue problem in each fluid column, but difficulties arise about which eigenvalue to select out of many. I developed a new method to estimate the Rossby wave phase speed, by empirically determining the reference frame in which the functional relation between potential vorticity and Bernoulli potential is tightest. With this, I analysed the Mach number for Rossby waves in Antarctic Circumpolar Current using the Southern Ocean State Estimate (Mazloff et al., 2010) on orthobaric density surfaces (de Szoeke et al., 2000). Studying the ACC as a whole, I find the potential vorticity is far from homogeneous in the ACC, and generally hovers just above the threshold to instability at an inverse Mach number of unity. A local perspective supports this claim, with the inverse Mach number generally larger than 0.5 in the ACC and speckled with pockets where it surpasses unity, the latter perhaps characterizing instability at work, though this has not been proven. Outside the ACC, the Mach number analysis suggests potential vorticity is much closer to uniform and the flow much further from the instability threshold.

Finally, Chapter 4 explored the degree to which this functional relationship between potential vorticity and Bernoulli potential on an approximately neutral surface breaks down. In other words, it explored how uniform potential vorticity is on streamlines on an approximately neutral surface, focusing on an approximately neutral surface around 1000 m depth in the core of the ACC. This chapter focused on the time-mean, with attention to how such time-means are constructed. The thickness-weighted average (Young, 2012) is an ideal theoretical framework for such an analysis, but ultimately the potential vorticity budget is not robustly computable from 5-day Eulerian mean model outputs. Nonetheless, the potential vorticity itself is computable in the thickness-weighted average framework, and in the deep ACC the structure of time-mean potential vorticity appears very similar between the thickness-weighted average and an Eulerian-style average. Along time-mean streamlines in the ACC, the time-mean potential vorticity is far from uniform. In particular, a large drop in (absolute) time-mean potential vorticity occurs downstream of Drake Passage, in agreement with (Gille, 1997), but only for northerly streamlines. Southerly streamlines of the ACC exhibit a much more uniform potential vorticity in the vicinity of Drake Passage. By analysing the potential vorticity budget in this Eulerian-style time-mean, it is found that eddy buoyancy fluxes and mean vertical advection through time-mean approximately neutral surfaces are the dominant forcings of the time-mean

potential vorticity, and that these cancel each other to first order, consistent with the Deacon cell and eddy-induced overturning cell largely cancelling each other. Near the South American shelf in particular, friction causes a significant drop in (absolute) potential vorticity, consistent with a no-slip boundary condition, while ageostrophic advection also contributes though to a lesser degree than friction, consistent with an intake of northern water. There is a similar drop in (absolute) potential vorticity at Campbell Plateau. Outside of these regions, (absolute) potential vorticity is steadily gained at all longitudes, in contrast to the results of Gille (1997) who found the gain occurs predominantly in the Pacific sector. This gain is remarkably linear with distance, despite the forcing terms exhibiting far more wild behaviour. Time-evolution of the stratification during SOSE's 6-year run (2005–2010) plays a surprisingly large role in the potential vorticity budget, though it cannot be discerned from SOSE alone whether this is related to climate change or interannual variability.

5.2 Future work

The topological analysis of neutral surfaces opens up similar possibilities in many other areas. Two obvious avenues of further work along these lines have been discussed. First, topobaric surfaces might be extendible to isosurfaces of a three dimensional density variable. However, the increase in dimensionality requires topological calculations to be performed on a multivariate, rather than scalar-valued, function. Multivariate computational topology is as a field is just beginning. The joint contour nets (Carr and Duke, 2014) that generalize the contour tree to multivariate functions may be useful here. They must be further generalized to allow for islands (though no longer outcroppings or incroppings) in the ocean, as going from the Contour Tree to the Reeb graph.

A second obvious avenue of further work is to extend Chapter 3, applying the Reeb graph to potential vorticity and Bernoulli potential in the ACC. This was not done for a first pass, in part because the level sets of either variable in the ACC are dominated by one circumpolar contour. However, in the reference-frame of Rossby waves (having an intrinsic westward phase speed relative to the background flow) this breaks down, and the Reeb graph becomes necessary.

Along these lines, the statistics of potential vorticity in the presence of geostrophic turbulence can reveal information about complex phenomena such as mixing barriers (David et al., 2017). These statistics are bound to become cleaner in topologically

distinct regions. Perhaps even the statistics, or other measures, of the topology itself will yield new insights.

A fourth avenue involves studying the geostrophic flow at the sea-floor, starting with the theory of Marshall (1995a), which relates the bottom geostrophic flow to gradients of a conservative density on the sea-floor, and a quantity that is related to a level of no motion. In fact, what arises is a multi-valued functional relationship between the Montgomery potential at the sea-floor and the buoyancy (or any conservative density variable) at the sea-floor. Approximating the bottom buoyancy as the bathymetry is equivalent to requiring the bottom vertical velocity be zero. The bathymetry being constant, this will yield a relatively simpler problem for Reeb graphs in oceanography, as no iterative procedure will be required. Such a theory of the bottom velocity will be especially useful as Deep ARGO measurements of the bottom (or near-bottom) density become prevalent.

Finally, there is a growing interest in the oceanographic community to perform thickness-weighted average (Young, 2012) analyses, but doing so from Eulerian averaged model data is dubious. Implementing accurate diagnostics of thickness-weighted averages into an ocean model (online) would likely support many future oceanographic studies.

5.3 Concluding remarks

Functional relations between variables arise throughout (geophysical) fluid dynamics, and often provide great insight. For example, the Prandtl-Batchelor theorem (Batchelor, 1956) is based on a functional relation between the vorticity and the stream function; it states that two-dimensional, steady, inviscid, and non-divergent flow has uniform vorticity inside regions of closed streamlines. Rhines and Young (1982) built upon this in their theory of potential vorticity homogenization inside closed geostrophic streamlines, where instead it is potential vorticity and the geostrophic stream function that are functionally related. A long line of shear stability theories was encapsulated by Arnol'd (1957), again based on a functional relation between potential vorticity and geostrophic stream function. McIntyre and Shepherd (1987) note a distinction arises when the functional relation between potential vorticity and stream function becomes multivalued compared to single valued: in the former case, Lagrangian information is required to express their finite-amplitude conservation theorem for the wave activity, whereas in the latter case, Eulerian quantities suffice. Dowling and Ingersoll (1989) derive the topography of Jupiter's convectively

adjusted deep interior around the Great Red Spot by fitting a multivalued function between potential vorticity and Bernoulli potential. Killworth (1992) argued that the Antarctic Circumpolar Current should posses an equivalent barotropic mode—that is the flow direction is independent of depth, though its magnitude may change—by using functional relationships between the potential vorticity, pressure, and stream function (see also Hughes and Killworth, 1995). Salmon (1994) generalized two-layer equations for the ocean circulation by using a functional relationship between temperature and potential vorticity. Similarly, Marshall (1995a) and Marshall (1995b) uses potential vorticity as a simple function of buoyancy alone to create an insightful quasi-analytic model of the ACC. This list could go on, and is not meant to be exhaustive.

A functional relationship between variables is a powerful idea. Our understanding of (geophysical) fluid dynamics is considerably aided by these relationships. When these relationships arise due to an alignment of the variables' gradients, topological analysis is required to correctly use these relationships to their fullest potential.

Bibliography

Agarwal, P. K., H. Edelsbrunner, J. Harer, and Y. Wang (2004). "Extreme Elevation on a 2-Manifold". ACM Press, p. 357. DOI: 10.1145/997817.997871 (cit. on pp. 14, 18, 49).

Allison, L. C., H. L. Johnson, D. P. Marshall, and D. R. Munday (2010). "Where Do Winds Drive the Antarctic Circumpolar Current?" *Geophysical Research Letters* 37.12. L12605, n/a–n/a. DOI: 10.1029/2010GL043355 (cit. on p. 131).

Andrews, D. G. (1983). "A Conservation Law for Small-Amplitude Quasi-Geostrophic Disturbances on a Zonally Asymmetric Basic Flow". *Journal of the Atmospheric Sciences* 40.1, pp. 85–90. DOI: 10.1175/1520-0469(1983)040<0085:ACLFSA>2.0.CO;2 (cit. on p. 9).

Aoki, K., A. Kubokawa, R. Furue, and H. Sasaki (2016). "Influence of Eddy Momentum Fluxes on the Mean Flow of the Kuroshio Extension in a 1/10° Ocean General Circulation Model". *Journal of Physical Oceanography* 46.9, pp. 2769–2784. DOI: 10.1175/JPO-D-16-0021.1 (cit. on pp. 137, 154, 155).

Arnold, V. I. (2006). "From Hilbert's Superposition Problem to Dynamical Systems". *Mathematical Events of the Twentieth Century*. Springer, pp. 19–47 (cit. on p. 14).

Arnol'd, V. I. (1957). "On the representability of a function of two variables in the form $\chi\phi(x) + \psi(y)$". *Uspekhi Matematicheskikh Nauk* 12.2, pp. 119–121 (cit. on pp. 14, 165).

— (1966). "An a Priori Estimate in the Theory of Hydrodynamic Stability". *Izvestiya Vysshikh Uchebnykh Zavedenii. Matematika* 5, pp. 3–5 (cit. on p. 9).

Azaneu, M., R. Kerr, and M. M. Mata (2014). "Assessment of the Representation of Antarctic Bottom Water Properties in the ECCO2 Reanalysis". *Ocean Science* 10.6, pp. 923–946. DOI: 10.5194/os-10-923-2014 (cit. on p. 19).

Batchelor, G. K. (1956). "On Steady Laminar Flow with Closed Streamlines at Large Reynolds Number". *Journal of Fluid Mechanics* 1.02, p. 177. DOI: 10.1017/S0022112056000123 (cit. on p. 165).

Beal, L. M., W. P. M. De Ruijter, A. Biastoch, R. Zahn, M. Cronin, J. Hermes, J. Lutjeharms, G. Quartly, T. Tozuka, S. Baker-Yeboah, T. Bornman, P. Cipollini, H. Dijkstra, I. Hall, W. Park, F. Peeters, P. Penven, H. Ridderinkhof, and J. Zinke (2011). "On the Role of the Agulhas System in Ocean Circulation and Climate". *Nature* 472.7344, pp. 429–436. DOI: 10.1038/nature09983 (cit. on p. 3).

Bethel, E. W., O. Rübel, K. Wu, G. H. Weber, V. Pascucci, H. Childs, A. Mascarenhas, J. Meredith, S. Ahern, and others (2009). "Modern Scientific Visualization

Is More than Just Pretty Pictures". *Numerical Modeling of Space Plasma Flows: ASTRONUM-2008*. Vol. 406, p. 301 (cit. on p. 15).

Biasotti, S., D. Giorgi, M. Spagnuolo, and B. Falcidieno (2008). "Reeb Graphs for Shape Analysis and Applications". *Theoretical Computer Science* 392.1-3, pp. 5–22. DOI: 10.1016/j.tcs.2007.10.018 (cit. on p. 15).

Bremer, P.-T., G. Weber, J. Tierny, V. Pascucci, M. Day, and J. Bell (2011). "Interactive Exploration and Analysis of Large-Scale Simulations Using Topology-Based Data Segmentation". *IEEE Transactions on Visualization and Computer Graphics* 17.9, pp. 1307–1324. DOI: 10.1109/TVCG.2010.253 (cit. on p. 15).

Bretherton, F. P. (1966). "Critical Layer Instability in Baroclinic Flows". *Quarterly Journal of the Royal Meteorological Society* 92.393, pp. 325–334. DOI: 10.1002/qj.49709239302 (cit. on p. 98).

Bryan, F. (1987). "Parameter Sensitivity of Primitive Equation Ocean General Circulation Models". *Journal of Physical Oceanography* 17.7, pp. 970–985. DOI: 10.1175/1520-0485(1987)017<0970:PSOPEO>2.0.CO;2 (cit. on p. 155).

Bryan, K. and M. D. Cox (1972). "The Circulation of the World Ocean: A Numerical Study. Part I, A Homogeneous Model". *J. Phys. Oceanogr.* 2.4, pp. 319–335. DOI: 10.1175/1520-0485(1972)002<0319:TCOTWO>2.0.CO;2 (cit. on p. 3).

Carr, H., T. Moller, and J. Snoeyink (2006). "Artifacts Caused by Simplicial Subdivision". *IEEE Transactions on Visualization and Computer Graphics* 12.2, pp. 231–242. DOI: 10.1109/TVCG.2006.22 (cit. on p. 45).

Carr, H. and D. Duke (2014). "Joint Contour Nets". *IEEE Transactions on Visualization and Computer Graphics* 20.8, pp. 1100–1113. DOI: 10.1109/TVCG.2013.269 (cit. on p. 164).

Carr, H., J. Snoeyink, and U. Axen (2003). "Computing Contour Trees in All Dimensions". *Computational Geometry* 24.2, pp. 75–94. DOI: 10.1016/S0925-7721(02)00093-7 (cit. on pp. 18, 44–46).

Carr, H., J. Snoeyink, and M. van de Panne (2010). "Flexible Isosurfaces: Simplifying and Displaying Scalar Topology Using the Contour Tree". *Computational Geometry* 43.1, pp. 42–58. DOI: 10.1016/j.comgeo.2006.05.009 (cit. on pp. 15, 16, 44, 47–49).

Cessi, P. and C. L. Wolfe (2013). "Adiabatic Eastern Boundary Currents". *Journal of Physical Oceanography* 43.6, pp. 1127–1149. DOI: 10.1175/JPO-D-12-0211.1 (cit. on pp. 137, 154, 155).

Charney, J. G. (1947). "The Dynamics of Long Waves in a Baroclinic Westerly Current". *Journal of Meteorology* 4.5, pp. 136–162. DOI: 10.1175/1520-0469(1947)004<0136:TDOLWI>2.0.CO;2 (cit. on p. 9).

Charney, J. G. and P. G. Drazin (1961). "Propagation of Planetary-Scale Disturbances from the Lower into the Upper Atmosphere". *Journal of Geophysical Research* 66.1, pp. 83–109. DOI: 10.1029/JZ066i001p00083 (cit. on p. 65).

Charney, J. G. and M. E. Stern (1962). "On the Stability of Internal Baroclinic Jets in a Rotating Atmosphere". *Journal of the Atmospheric Sciences* 19.2, pp. 159–172. DOI: 10.1175/1520-0469(1962)019<0159:OTSOIB>2.0.CO;2 (cit. on p. 98).

Chelton, D. B., R. A. deSzoeke, M. G. Schlax, K. El Naggar, and N. Siwertz (1998). "Geographical Variability of the First Baroclinic Rossby Radius of Deforma-

tion". *Journal of Physical Oceanography* 28.3, pp. 433–460. DOI: 10.1175/1520-0485(1998)028<0433:GVOTFB>2.0.CO;2 (cit. on p. 77).

Cole-McLaughlin, K., H. Edelsbrunner, J. Harer, V. Natarajan, and V. Pascucci (2003). "Loops in Reeb Graphs of 2-Manifolds". ACM Press, p. 344. DOI: 10.1145/777792.777844 (cit. on p. 18).

Cunningham, S. A., S. G. Alderson, B. A. King, and M. A. Brandon (2003). "Transport and Variability of the Antarctic Circumpolar Current in Drake Passage". *Journal of Geophysical Research: Oceans* 108.C5. DOI: 10.1029/2001JC001147 (cit. on p. 3).

Cunningham, S. A. (2000). "Circulation and Volume Flux of the North Atlantic Using Synoptic Hydrographic Data in a Bernoulli Inverse". *Journal of marine research* 58.1, pp. 1–35. DOI: 10.1357/002224000321511188 (cit. on p. 31).

Danabasoglu, G., J. C. McWilliams, and P. R. Gent (1994). "The Role of Mesoscale Tracer Transports in the Global Ocean Circulation". *Science* 264.5162, pp. 1123–1126. DOI: 10.1126/science.264.5162.1123 (cit. on p. 101).

David, T. W., D. P. Marshall, and L. Zanna (2017). "The Statistical Nature of Turbulent Barotropic Ocean Jets". *Ocean Modelling* 113 (Supplement C), pp. 34–49. DOI: 10.1016/j.ocemod.2017.03.008 (cit. on pp. 94, 164).

De La Lama, M. S., J. H. LaCasce, and H. K. Fuhr (2016). "The Vertical Structure of Ocean Eddies". *Dynamics and Statistics of the Climate System*, dzw001. DOI: 10.1093/climsys/dzw001 (cit. on pp. 95, 97).

De Szoeke, R. A. (2000). "Equations of Motion Using Thermodynamic Coordinates". *Journal of Physical Oceanography* 30.11, pp. 2814–2829. DOI: 10.1175/1520-0485(2001)031<2814:>2.0.CO;2 (cit. on pp. 27, 63, 78).

De Szoeke, R. A. and S. R. Springer (2005). "The All-Atlantic Temperature-Salinity-Pressure Relation and Patched Potential Density". *Journal of Marine Research* 63.1, pp. 59–93. DOI: 10.1357/0022240053693752 (cit. on pp. 41, 56, 62).

— (2009). "The Materiality and Neutrality of Neutral Density and Orthobaric Density". *Journal of Physical Oceanography* 39.8, pp. 1779–1799. DOI: 10.1175/2009JPO4042.1 (cit. on pp. 26, 28, 41, 112).

De Szoeke, R. A., S. R. Springer, and D. M. Oxilia (2000). "Orthobaric Density: A Thermodynamic Variable for Ocean Circulation Studies". *Journal of physical oceanography* 30.11, pp. 2830–2852 (cit. on pp. 27, 33, 63, 78, 112, 163).

De Boor, C. (1978). *A Practical Guide to Splines*. Vol. 27. Springer-Verlag New York (cit. on p. 51).

Dellar, P. J. (2011). "Variations on a Beta-Plane: Derivation of Non-Traditional Beta-Plane Equations from Hamilton's Principle on a Sphere". *Journal of Fluid Mechanics* 674, pp. 174–195. DOI: 10.1017/S0022112010006464 (cit. on p. 140).

Donohue, K. A., K. L. Tracey, D. R. Watts, M. P. Chidichimo, and T. K. Chereskin (2016). "Mean Antarctic Circumpolar Current Transport Measured in Drake Passage". *Geophysical Research Letters* 43.22, pp. 11,760–11,767. DOI: 10.1002/2016GL070319 (cit. on p. 3).

Doraiswamy, H. and V. Natarajan (2013). "Computing Reeb Graphs as a Union of Contour Trees". *IEEE Transactions on Visualization and Computer Graphics* 19.2, pp. 249–262. DOI: 10.1109/TVCG.2012.115 (cit. on pp. 18, 44, 46).

Dowling, T. E., M. E. Bradley, J. Du, S. R. Lewis, and P. L. Read (2017). "Ertel Potential Vorticity versus Bernoulli Streamfunction on Mars: Q versus B on Mars". *Quarterly Journal of the Royal Meteorological Society* 143.702, pp. 37–52. DOI: 10.1002/qj.2916 (cit. on pp. 66, 71, 77).

Dowling, T. E. (1993). "A Relationship between Potential Vorticity and Zonal Wind on Jupiter". *Journal of the Atmospheric Sciences* 50.1, pp. 14–22. DOI: 10.1175/1520-0469(1993)050<0014:ARBPVA>2.0.CO;2 (cit. on p. 65).

Dowling, T. E. (1995). "Estimate of Jupiter's Deep Zonal-Wind Profile from Shoemaker-Levy 9 Data and Arnol'd's Second Stability Criterion". *Icarus* 117.2, pp. 439–442. DOI: 10.1006/icar.1995.1169 (cit. on p. 71).

— (2014). "Saturn's Longitude: Rise of the Second Branch of Shear-Stability Theory and Fall of the First". *International Journal of Modern Physics D* 23.04, p. 1430006. DOI: 10.1142/S0218271814300067 (cit. on pp. 12, 66–68, 76, 93, 162).

Dowling, T. E. and A. P. Ingersoll (1989). "Jupiter's Great Red Spot as a Shallow Water System". *Journal of the Atmospheric Sciences* 46.21, pp. 3256–3278. DOI: 10.1175/1520-0469(1989)046<3256:JGRSAA>2.0.CO;2 (cit. on pp. 65, 70, 165).

Dritschel, D. G. and M. E. McIntyre (2008). "Multiple Jets as PV Staircases: The Phillips Effect and the Resilience of Eddy-Transport Barriers". *Journal of the Atmospheric Sciences* 65.3, pp. 855–874. DOI: 10.1175/2007JAS2227.1 (cit. on p. 9).

Du, J., T. E. Dowling, and M. E. Bradley (2015). "Ertel Potential Vorticity versus Bernoulli Streamfunction in Earth's Extratropical Atmosphere". *Journal of Advances in Modeling Earth Systems* 7.2, pp. 437–458. DOI: 10.1002/2014MS000420 (cit. on pp. 66, 71, 77).

Eady, E. T. (1949). "Long Waves and Cyclone Waves". *Tellus* 1.3, pp. 33–52. DOI: 10.1111/j.2153-3490.1949.tb01265.x (cit. on pp. 8, 11).

Edelsbrunner, H., D. Letscher, and A. Zomorodian (2000). "Topological Persistence and Simplification". IEEE Comput. Soc, pp. 454–463. DOI: 10.1109/SFCS.2000.892133 (cit. on p. 48).

Eliassen, A. and E. Palm (1961). "On the Transfer of Energy in Stationary Mountain Waves". *Geofysiske Publikasjoner* 22, pp. 1–23 (cit. on p. 9).

Ertel, H. (1942). "Ein Neuer Hydrodynamischer Wirbelsatz". *Meteorol. Z.* 59, pp. 277–281 (cit. on p. 9).

Farneti, R. and T. L. Delworth (2010). "The Role of Mesoscale Eddies in the Remote Oceanic Response to Altered Southern Hemisphere Winds". *Journal of Physical Oceanography* 40.10, pp. 2348–2354. DOI: 10.1175/2010JPO4480.1 (cit. on p. 2).

Ferrari, R., M. F. Jansen, J. F. Adkins, A. Burke, A. L. Stewart, and A. F. Thompson (2014). "Antarctic Sea Ice Control on Ocean Circulation in Present and Glacial Climates". *Proceedings of the National Academy of Sciences* 111.24, pp. 8753–8758. DOI: 10.1073/pnas.1323922111 (cit. on p. 2).

Forget, G. (2010). "Mapping Ocean Observations in a Dynamical Framework: A 2004–06 Ocean Atlas". *Journal of Physical Oceanography* 40.6, pp. 1201–1221. DOI: 10.1175/2009JPO4043.1 (cit. on p. 36).

Friedman, J. H. (1993). "Fast MARS". *Lab.for Computational Statistics, Stanford Univ., Tech.Report* LCS110, pp. 1–18 (cit. on p. 81).

Fyfe, J. C. and O. A. Saenko (2006). "Simulated Changes in the Extratropical Southern Hemisphere Winds and Currents". *Geophysical Research Letters* 33.6. DOI: 10.1029/2005GL025332 (cit. on p. 2).

Gates, W. L. (2004). "Derivation of the Equations of Atmospheric Motion in Oblate Spheroidal Coordinates". *Journal of the Atmospheric Sciences* 61.20, pp. 2478–2487. DOI: 10.1175/1520-0469(2004)061<2478:DOTEOA>2.0.CO;2 (cit. on p. 72).

Gent, P. R., J. Willebrand, T. J. McDougall, and J. C. McWilliams (1995). "Parameterizing Eddy-Induced Tracer Transports in Ocean Circulation Models". *Journal of Physical Oceanography* 25.4, pp. 463–474. DOI: 10.1175/1520-0485(1995)025<0463:PEITTI>2.0.CO;2 (cit. on pp. 101, 107, 109).

Gill, A. E. (1982). *Atmosphere–Ocean Dynamics*. Academic Press (cit. on p. 65).

Gille, S. T. (1997). "Why Potential Vorticity Is Not Conserved along Mean Streamlines in a Numerical Southern Ocean". *Journal of physical oceanography* 27.7, pp. 1286–1299. DOI: 10.1175/1520-0485(1997)027<1286:WPVINC>2.0.CO;2 (cit. on pp. 27, 82, 101, 102, 104–106, 112, 116, 117, 124, 128–130, 134, 135, 137, 138, 163, 164).

Gorban, A. N. (2013). "Thermodynamic Tree: The Space of Admissible Paths". *SIAM Journal on Applied Dynamical Systems* 12.1, pp. 246–278. DOI: 10.1137/120866919 (cit. on pp. 14, 15).

Graham, R. M. and A. M. De Boer (2013). "The Dynamical Subtropical Front: The Dynamical Subtropical Front". *Journal of Geophysical Research: Oceans* 118.10, pp. 5676–5685. DOI: 10.1002/jgrc.20408 (cit. on p. 133).

Green, J. S. A. (1970). "Transfer Properties of the Large-Scale Eddies and the General Circulation of the Atmosphere". *Quarterly Journal of the Royal Meteorological Society* 96.408, pp. 157–185. DOI: 10.1002/qj.49709640802 (cit. on p. 66).

Hallberg, R. and A. Gnanadesikan (2006). "The Role of Eddies in Determining the Structure and Response of the Wind-Driven Southern Hemisphere Overturning: Results from the Modeling Eddies in the Southern Ocean (MESO) Project". *Journal of Physical Oceanography* 36.12, pp. 2232–2252. DOI: 10.1175/JPO2980.1 (cit. on p. 2).

Hartmann, D., A. K. Tank, M. Rusticucci, L. Alexander, S. Brönnimann, Y. Charabi, F. Dentener, E. Dlugokencky, D. Easterling, A. Kaplan, B. Soden, P. Thorne, M. Wild, and P. Zhai (2013). "Observations: Atmosphere and Surface". *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Ed. by T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, A. N. J. Boschung, Y. Xia, V. Bex, and P. Midgley. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA. Chap. 2 (cit. on p. 2).

Heine, C., D. Schneider, H. Carr, and G. Scheuermann (2011). "Drawing Contour Trees in the Plane". *IEEE Transactions on Visualization and Computer Graphics* 17.11, pp. 1599–1611. DOI: 10.1109/TVCG.2010.270 (cit. on p. 54).

Hidaka, K. and M. Tsuchiya (1953). "On the Antarctic Circumpolar Current". *J. Mar. Res* 12.2, p. 214 (cit. on p. 3).

Hoskins, B. J., M. McIntyre, and A. W. Robertson (1985). "On the Use and Significance of Isentropic Potential Vorticity Maps". *Quarterly Journal of the Royal Meteorological Society* 111.470, pp. 877–946. DOI: 10.1002/qj.49711147002 (cit. on p. 9).

Hughes, C. W. (1995). "Rossby Waves in the Southern Ocean: A Comparison of TOPEX/POSEIDON Altimetry with Model Predictions". *Journal of Geophysical Research: Oceans* 100.C8, pp. 15933–15950. DOI: 10.1029/95JC01380 (cit. on p. 136).

— (1997). "Comments on "On the Obscurantist Physics of 'Form Drag' in Theorizing about the Circumpolar Current"". *Journal of Physical Oceanography* 27.1, pp. 209–210. DOI: 10.1175/1520-0485(1997)027<0209:COOTOP>2.0.CO;2 (cit. on p. 107).

Hughes, C. W., M. S. Jones, and S. Carnochan (1998). "Use of Transient Features to Identify Eastward Currents in the Southern Ocean". *Journal of Geophysical Research: Oceans* 103.C2, pp. 2929–2943. DOI: 10.1029/97JC02442 (cit. on p. 136).

Hughes, C. W. and P. D. Killworth (1995). "Effects of Bottom Topography in the Large-Scale Circulation of the Southern Ocean". *Journal of Physical Oceanography* 25.11, pp. 2485–2497. DOI: 10.1175/1520-0485(1995)025<2485:EOBTIT>2.0.CO;2 (cit. on pp. 138, 166).

Jackett, D. R. and T. J. McDougall (1997). "A Neutral Density Variable for the World's Oceans". *Journal of Physical Oceanography* 27.2, pp. 237–263. DOI: 10.1175/1520-0485(1997)027<0237:ANDVFT>2.0.CO;2 (cit. on pp. 8, 27, 61, 63, 78, 112, 113).

Jekabsons, G (2016). "ARESLab: Adaptive Regression Splines Toolbox for Matlab/Octave" (cit. on p. 81).

Johnson, G. C. and H. L. Bryden (1989). "On the Size of the Antarctic Circumpolar Current". *Deep Sea Research Part A. Oceanographic Research Papers* 36.1, pp. 39–53. DOI: 10.1016/0198-0149(89)90017-4 (cit. on p. 5).

Kanongchaiyos, P. and Y. Shinagawa (2000). "Articulated Reeb Graphs for Interactive Skeleton Animation". World Scientific, pp. 451–467. DOI: 10.1142/9789812791993_0029 (cit. on p. 15).

Keffer, T. (1985). "The Ventilation of the World's Oceans: Maps of the Potential Vorticity Field". *Journal of Physical Oceanography* 15.5, pp. 509–523. DOI: 10.1175/1520-0485(1985)015<0509:TVOTWO>2.0.CO;2 (cit. on p. 12).

Killworth, P. D. (1992). "An Equivalent-Barotropic Mode In the Fine Resolution Antarctic Model". *Journal of Physical Oceanography* 22.11, pp. 1379–1387. DOI: 10.1175/1520-0485(1992)022<1379:AEBMIT>2.0.CO;2 (cit. on pp. 132, 138, 166).

Killworth, P. D. and J. R. Blundell (2003). "Long Extratropical Planetary Wave Propagation in the Presence of Slowly Varying Mean Flow and Bottom Topography. Part I: The Local Problem". *Journal of Physical Oceanography* 33.4, pp. 784–801. DOI: 10.1175/1520-0485(2003)33<784:LEPWPI>2.0.CO;2 (cit. on p. 95).

Killworth, P. D., D. B. Chelton, and R. A. de Szoeke (1997). "The Speed of Observed and Theoretical Long Extratropical Planetary Waves". *Journal of Physical Oceanography* 27.9, pp. 1946–1966. DOI: 10.1175/1520-0485(1997)027<1946:TSOOAT>2.0.CO;2 (cit. on p. 95).

Klocker, A., T. J. McDougall, and D. R. Jackett (2009). "A New Method for Forming Approximately Neutral Surfaces". *Ocean Science* 5.2, pp. 155–172. DOI: 10.5194/os-5-155-2009 (cit. on pp. 27, 28, 39, 61, 62, 112).

Klocker, A. and D. P. Marshall (2014). "Advection of Baroclinic Eddies by Depth Mean Flow". *Geophysical Research Letters* 41.10, pp. 3517–3521. DOI: 10.1002/2014GL060001 (cit. on pp. 84, 136).

Laney, D., P.-t. Bremer, A. Mascarenhas, P. Miller, and V. Pascucci (2006). "Understanding the Structure of the Turbulent Mixing Layer in Hydrodynamic Instabilities". *IEEE Transactions on Visualization and Computer Graphics* 12.5, pp. 1053–1060. DOI: 10.1109/TVCG.2006.186 (cit. on p. 15).

Le Quere, C., C. Rodenbeck, E. T. Buitenhuis, T. J. Conway, R. Langenfelds, A. Gomez, C. Labuschagne, M. Ramonet, T. Nakazawa, N. Metzl, N. Gillett, and M. Heimann (2007). "Saturation of the Southern Ocean CO2 Sink Due to Recent Climate Change". *Science* 316.5832, pp. 1735–1738. DOI: 10.1126/science.1136188 (cit. on p. 3).

Li, Z., Y. Chao, and J. C. McWilliams (2006). "Computation of the Streamfunction and Velocity Potential for Limited and Irregular Domains". *Monthly Weather Review* 134.11, pp. 3384–3394. DOI: 10.1175/MWR3249.1 (cit. on p. 121).

Lovenduski, N. S., N. Gruber, and S. C. Doney (2008). "Toward a Mechanistic Understanding of the Decadal Trends in the Southern Ocean Carbon Sink". *Global Biogeochemical Cycles* 22.3, n/a–n/a. DOI: 10.1029/2007GB003139 (cit. on p. 3).

Luyten, J. R., J. Pedlosky, and H. Stommel (1983). "The Ventilated Thermocline". *Journal of Physical Oceanography* 13.2, pp. 292–309. DOI: 10.1175/1520-0485(1983)013<0292:TVT>2.0.CO;2 (cit. on pp. 9, 12, 27, 132).

Lynn, R. J. and J. L. Reid (1968). "Characteristics and Circulation of Deep and Abyssal Waters". *Deep Sea Research and Oceanographic Abstracts* 15.5, pp. 577–598. DOI: 10.1016/0011-7471(68)90064-8 (cit. on p. 27).

MacKinnon, J., L. St Laurent, and A. C. Naveira Garabato (2013). "Diapycnal Mixing Processes in the Ocean Interior". *International Geophysics*. Vol. 103. Elsevier, pp. 159–183. DOI: 10.1016/B978-0-12-391851-2.00007-6 (cit. on p. 62).

Marinov, I., A. Gnanadesikan, J. R. Toggweiler, and J. L. Sarmiento (2006). "The Southern Ocean Biogeochemical Divide". *Nature* 441.7096, pp. 964–967. DOI: 10.1038/nature04883 (cit. on p. 3).

Marshall, D. (1995a). "Influence of Topography on the Large-Scale Ocean Circulation". *Journal of Physical Oceanography* 25.7, pp. 1622–1635. DOI: 10.1175/1520-0485(1995)025<1622:IOTOTL>2.0.CO;2 (cit. on pp. 165, 166).

— (1995b). "Topographic Steering of the Antarctic Circumpolar Current". *Journal of Physical Oceanography* 25.7, pp. 1636–1650. DOI: 10.1175/1520-0485(1995)025<1636:TSOTAC>2.0.CO;2 (cit. on pp. 6, 12, 13, 138, 166).

— (1997). "Subduction of Water Masses in an Eddying Ocean". *Journal of Marine Research* 55.2, pp. 201–222. DOI: 10.1357/0022240973224373 (cit. on p. 101).

Marshall, D. P. (2016). "A Theoretical Model of Long Rossby Waves in the Southern Ocean and Their Interaction with Bottom Topography". *Fluids* 1.2, p. 17. DOI: 10.3390/fluids1020017 (cit. on p. 13).

Marshall, D. P. and J. C. Stephens (2001). "On the Insensitivity of the Wind-Driven Circulation to Bottom Topography". *Journal of marine research* 59.1, pp. 1–27. DOI: 10.1357/002224001321237344 (cit. on p. 13).

Marshall, D. P., D. R. Munday, L. C. Allison, R. J. Hay, and H. L. Johnson (2016). "Gill's Model of the Antarctic Circumpolar Current, Revisited: The Role of Latitudinal Variations in Wind Stress". *Ocean Modelling* 97, pp. 37–51. DOI: 10.1016/j.ocemod.2015.11.010 (cit. on pp. 102, 108).

Marshall, D. P., M. H. P. Ambaum, J. R. Maddison, D. R. Munday, and L. Novak (2017). "Eddy Saturation and Frictional Control of the Antarctic Circumpolar Current: Eddy Saturation of the ACC". *Geophysical Research Letters* 44.1, pp. 286–292. DOI: 10.1002/2016GL071702 (cit. on p. 2).

Marshall, J. and T. Radko (2003). "Residual-Mean Solutions for the Antarctic Circumpolar Current and Its Associated Overturning Circulation". *Journal of Physical Oceanography* 33.11, pp. 2341–2354. DOI: 10.1175/1520-0485(2003)033<2341:RSFTAC>2.0.CO;2 (cit. on pp. 27, 101, 128).

Marshall, J. and D. W. K. So (1990). "Thermal Equilibration of Planetary Waves". *Journal of the Atmospheric Sciences* 47.8, pp. 963–978. DOI: 10.1175/1520-0469(1990)047<0963:TEOPW>2.0.CO;2 (cit. on p. 65).

Marshall, J., D. Olbers, H. Ross, and D. Wolf-Gladrow (1993). "Potential Vorticity Constraints on the Dynamics and Hydrography of the Southern Ocean". *Journal of Physical Oceanography*. DOI: 10.1175/1520-0485(1993)023<0465:PVCOTD>2.0.CO;2 (cit. on pp. 5, 6, 12, 66, 93).

Mazloff, M. R., P. Heimbach, and C. Wunsch (2010). "An Eddy-Permitting Southern Ocean State Estimate". *Journal of Physical Oceanography* 40.5, pp. 880–899. DOI: 10.1175/2009JPO4236.1 (cit. on pp. 18, 19, 66, 106, 163).

McDougall, T. J. (1987a). "Neutral Surfaces". *Journal of Physical Oceanography*. DOI: 10.1175/1520-0485(1987)017<1950:NS>2.0.CO;2 (cit. on pp. 6, 7, 26, 30, 62, 68).

— (1987b). "Thermobaricity, Cabbeling, and Water-Mass Conversion". *Journal of Geophysical Research* 92.C5, p. 5448. DOI: 10.1029/JC092iC05p05448 (cit. on p. 29).

— (1989). "Streamfunctions for the Lateral Velocity Vector in a Compressible Ocean". *Journal of Marine Research* 47.2, pp. 267–284. DOI: 10.1357/002224089785076271 (cit. on pp. 5, 8, 27, 31, 40, 63, 78, 162).

McDougall, T. J. and D. R. Jackett (1988). "On the Helical Nature of Neutral Trajectories in the Ocean". *Progress in Oceanography* 20.3, pp. 153–183. DOI: 10.1016/0079-6611(88)90001-8 (cit. on pp. 6–8, 26, 29, 32, 62).

McDougall, T. J. and D. R. Jackett (2005a). "An Assessment of Orthobaric Density in the Global Ocean". *Journal of Physical Oceanography* 35.11, pp. 2054–2075. DOI: 10.1175/JPO2796.1 (cit. on pp. 41, 56, 62).

— (2005b). "The Material Derivative of Neutral Density". *Journal of Marine Research* 63.1, pp. 159–185. DOI: 10.1357/0022240053693734 (cit. on p. 62).

— (2007). "The Thinness of the Ocean in $S - \Theta - p$ Space and the Implications for Mean Diapycnal Advection". *Journal of Physical Oceanography* 37.6, pp. 1714–1732. DOI: 10.1175/JPO3114.1 (cit. on p. 29).

McDougall, T. J. and A. Klocker (2010). "An Approximate Geostrophic Streamfunction for Use in Density Surfaces". *Ocean Modelling* 32.3-4, pp. 105–117. DOI: 10.1016/j.ocemod.2009.10.006 (cit. on pp. 28, 31, 34, 112).

McDougall, T. J. and P. C. McIntosh (1996). "The Temporal-Residual-Mean Velocity. Part I: Derivation and the Scalar Conservation Equations". *Journal of Physical Oceanography* 26.12, pp. 2653–2665. DOI: 10.1175/1520-0485(1996)026<2653:TTRMVP>2.0.CO;2 (cit. on p. 115).

— (2001). "The Temporal-Residual-Mean Velocity. Part II: Isopycnal Interpretation and the Tracer and Momentum Equations". *Journal of Physical Oceanography* 31.5, pp. 1222–1246. DOI: 10.1175/1520-0485(2001)031<1222:TTRMVP>2.0.CO;2 (cit. on pp. 109, 115, 156, 158).

McDougall, T. J., S. Groeskamp, and S. M. Griffies (2014). "On Geometrical Aspects of Interior Ocean Mixing". *Journal of Physical Oceanography* 44.8, pp. 2164–2175. DOI: 10.1175/JPO-D-13-0270.1 (cit. on pp. 8, 30).

McIntyre, M. E. and T. G. Shepherd (1987). "An Exact Local Conservation Theorem for Finite-Amplitude Disturbances to Non-Parallel Shear Flows, with Remarks on Hamiltonian Structure and on Arnol'd's Stability Theorems". *Journal of Fluid Mechanics* 181, pp. 527–565. DOI: 10.1017/S0022112087002209 (cit. on pp. 66, 165).

Meijers, A. J. S., E. Shuckburgh, N. Bruneau, J.-B. Sallee, T. J. Bracegirdle, and Z. Wang (2012). "Representation of the Antarctic Circumpolar Current in the CMIP5 Climate Models and Future Changes under Warming Scenarios". *Journal of Geophysical Research: Oceans* 117.C12, n/a–n/a. DOI: 10.1029/2012JC008412 (cit. on p. 3).

Menemenlis, D., C. Hill, A. Adcrocft, J.-M. Campin, B. Cheng, B. Ciotti, I. Fukumori, P. Heimbach, C. Henze, A. Köhl, T. Lee, D. Stammer, J. Taft, and J. Zhang (2005). "NASA Supercomputer Improves Prospects for Ocean Climate Research". *Eos, Transactions American Geophysical Union* 86.9, p. 89. DOI: 10.1029/2005EO090002 (cit. on pp. 18, 19, 43).

Meredith, M. P., P. L. Woodworth, T. K. Chereskin, D. P. Marshall, L. C. Allison, G. R. Bigg, K. Donohue, K. J. Heywood, C. W. Hughes, A. Hibbert, A. M. Hogg, H. L. Johnson, L. Jullion, B. A. King, H. Leach, Y.-D. Lenn, M. A. Morales Maqueda, D. R. Munday, A. C. Naveira Garabato, C. Provost, J.-B. Sallée, and J. Sprintall (2011). "Sustained Monitoring of the Southern Ocean at Drake Passage: Past Achievements and Future Priorities". *Reviews of Geophysics* 49.4. DOI: 10.1029/2010RG000348 (cit. on p. 1).

Montgomery, R. B. (1938). *Circulation in Upper Layers of Southern North Atlantic Deduced with Use of Isentropic Analysis*. Cambridge, MA: Massachusetts Institute of Technology and Woods Hole Oceanographic Institution. DOI: 10.1575/1912/1093 (cit. on p. 27).

Montgomery, R. (1937). "A Suggested Method for Representing Gradient Flow in Isentropic Surfaces". *Bull. Amer. Meteor. Soc* 18, pp. 210–212 (cit. on pp. 27, 31, 78, 112).

Munday, D. R., H. L. Johnson, and D. P. Marshall (2013). "Eddy Saturation of Equilibrated Circumpolar Currents". *Journal of Physical Oceanography* 43.3, pp. 507–532. DOI: 10.1175/JPO-D-12-095.1 (cit. on p. 2).

Munk, W. H. and E. Palmén (1951). "Note on the Dynamics of the Antarctic Circumpolar Current". *Tellus* 3.1, pp. 53–55. DOI: 10.1111/j.2153-3490.1951.tb00776.x (cit. on pp. 5, 111).

Müller, P. (1995). "Ertel's Potential Vorticity Theorem in Physical Oceanography". *Reviews of Geophysics* 33.1, p. 67. DOI: 10.1029/94RG03215 (cit. on p. 27).

Nikurashin, M. and G. Vallis (2011). "A Theory of Deep Stratification and Overturning Circulation in the Ocean". *Journal of Physical Oceanography* 41.3, pp. 485–502. DOI: 10.1175/2010JPO4529.1 (cit. on p. 3).

Nycander, J., M. Hieronymus, and F. Roquet (2015). "The Nonlinear Equation of State of Sea Water and the Global Water Mass Distribution". *Geophysical Research Letters* 42.18, 2015GL065525. DOI: 10.1002/2015GL065525 (cit. on p. 1).

Olbers, D. (1998). "Comments on "On the Obscurantist Physics of 'Form Drag' in Theorizing about the Circumpolar Current*"". *Journal of Physical Oceanography* 28.8, pp. 1647–1654. DOI: 10.1175/1520-0485(1998)028<1647:COOTOP>2.0.CO;2 (cit. on p. 107).

Olbers, D. and M. Visbeck (2005). "A Model of the Zonally Averaged Stratification and Overturning in the Southern Ocean". *Journal of Physical Oceanography* 35.7, pp. 1190–1205. DOI: 10.1175/JPO2750.1 (cit. on p. 4).

Olbers, D., D. Borowski, C. Völker, and J.-O. Wölff (2004). "The Dynamical Balance, Transport and Circulation of the Antarctic Circumpolar Current". *Antarctic Science* 16.4, pp. 439–470. DOI: 10.1017/S0954102004002251 (cit. on p. 1).

O'Dwyer, J. and R. G. Williams (1997). "The Climatological Distribution of Potential Vorticity over the Abyssal Ocean". *Journal of Physical Oceanography* 27.11, pp. 2488–2506. DOI: 10.1175/1520-0485(1997)027<2488:TCDOPV>2.0.CO;2 (cit. on p. 12).

Park, Y.-G. and K. Bryan (2000). "Comparison of Thermally Driven Circulations from a Depth-Coordinate Model and an Isopycnal-Layer Model. Part I: Scaling-Law Sensitivity to Vertical Diffusivity". *Journal of Physical Oceanography* 30.3, pp. 590–605. DOI: 10.1175/1520-0485(2000)030<0590:COTDCF>2.0.CO;2 (cit. on p. 155).

Pascucci, V., G. Scorzelli, P.-T. Bremer, and A. Mascarenhas (2007). "Robust On-Line Computation of Reeb Graphs: Simplicity and Speed". ACM Press, p. 58. DOI: 10.1145/1275808.1276449 (cit. on pp. 47, 49).

Prim, R. C. (1957). "Shortest Connection Networks And Some Generalizations". *Bell System Technical Journal* 36.6, pp. 1389–1401. DOI: 10.1002/j.1538-7305.1957.tb01515.x (cit. on p. 51).

Read, P. L., B. J. Conrath, L. N. Fletcher, P. J. Gierasch, A. A. Simon-Miller, and L. C. Zuchowski (2009a). "Mapping Potential Vorticity Dynamics on Saturn: Zonal

Mean Circulation from Cassini and Voyager Data". *Planetary and Space Science* 57.14, pp. 1682–1698. DOI: 10.1016/j.pss.2009.03.004 (cit. on p. 65).

Read, P. L., T. E. Dowling, and G. Schubert (2009b). "Saturn's Rotation Period from Its Atmospheric Planetary-Wave Configuration". *Nature* 460.7255, pp. 608–610. DOI: 10.1038/nature08194 (cit. on pp. 65, 71, 76).

Read, P. L., P. J. Gierasch, B. J. Conrath, A. Simon-Miller, T. Fouchet, and Y. H. Yamazaki (2006). "Mapping Potential-Vorticity Dynamics on Jupiter. I: Zonal-Mean Circulation from Cassini and Voyager 1 Data". *Quarterly Journal of the Royal Meteorological Society* 132.618, pp. 1577–1603. DOI: 10.1256/qj.05.34 (cit. on p. 65).

Redi, M. H. (1982). "Oceanic Isopycnal Mixing by Coordinate Rotation". *Journal of Physical Oceanography* 12.10, pp. 1154–1158. DOI: 10.1175/1520-0485(1982)012<1154:OIMBCR>2.0.CO;2 (cit. on p. 62).

Reeb, G. (1946). "Sur Les Points Singuliers d'une Forme de Pfaff Completement Intégrable Ou d'une Fonction Numérique". *CR Acad. Sci. Paris* 222.847-849, p. 2 (cit. on pp. 14, 16).

Reid, J. L. and R. J. Lynn (1971). "On the Influence of the Norwegian-Greenland and Weddell Seas upon the Bottom Waters of the Indian and Pacific Oceans". *Deep Sea Research and Oceanographic Abstracts* 18.11, pp. 1063–1088. DOI: 10.1016/0011-7471(71)90094-5 (cit. on pp. 8, 27).

Rhines, P. B. and W. R. Young (1982). "Homogenization of Potential Vorticity in Planetary Gyres". *Journal of Fluid Mechanics* 122, pp. 347–367. DOI: 10.1017/S0022112082002250 (cit. on pp. 6, 9, 11, 12, 27, 66, 165).

Ringler, T., J. A. Saenz, P. J. Wolfram, and L. Van Roekel (2016). "A Thickness-Weighted Average Perspective of Force Balance in an Idealized Circumpolar Current". *Journal of Physical Oceanography* 47.2, pp. 285–302. DOI: 10.1175/JPO-D-16-0096.1 (cit. on pp. 137, 155).

Saenko, O. A. and W. J. Merryfield (2005). "On the Effect of Topographically Enhanced Mixing on the Global Ocean Circulation". *Journal of Physical Oceanography* 35.5, pp. 826–834. DOI: 10.1175/JPO2722.1 (cit. on p. 155).

Salmon, R. (1994). "Generalized Two-Layer Models of Ocean Circulation". *Journal of Marine Research* 52, pp. 865–908. DOI: 10.1357/0022240943076939 (cit. on p. 166).

Schubert, W., E. Ruprecht, R. Hertenstein, R. N. Ferreira, R. Taft, C. Rozoff, P. Ciesielski, and H.-C. Kuo (2004). "English Translations of Twenty-One of Ertel's Papers on Geophysical Fluid Dynamics". *Meteorologische Zeitschrift* 13.6, pp. 527–576. DOI: 10.1127/0941-2948/2004/0013-0527 (cit. on p. 9).

Shinagawa, Y., T. Kunii, and Y. Kergosien (1991). "Surface Coding Based on Morse Theory". *IEEE Computer Graphics and Applications* 11.5, pp. 66–78. DOI: 10.1109/38.90568 (cit. on p. 15).

Shuckburgh, E., H. Jones, J. Marshall, and C. Hill (2009). "Robustness of an Effective Diffusivity Diagnostic in Oceanic Flows". *Journal of Physical Oceanography* 39.9, pp. 1993–2009. DOI: 10.1175/2009JPO4122.1 (cit. on p. 27).

Smith, K. S. (2007). "The Geography of Linear Baroclinic Instability in Earth's Oceans". *Journal of Marine Research* 65.5, pp. 655–683. DOI: 10.1357/002224007783649484 (cit. on p. 81).

Spence, P., O. A. Saenko, M. Eby, and A. J. Weaver (2009). "The Southern Ocean Overturning: Parameterized versus Permitted Eddies". *Journal of Physical Oceanography* 39.7, pp. 1634–1651. DOI: 10.1175/2009JPO4120.1 (cit. on p. 2).

Stanley, G. J. and O. A. Saenko (2014). "Bottom-Enhanced Diapycnal Mixing Driven by Mesoscale Eddies: Sensitivity to Wind Energy Supply". *Journal of Physical Oceanography* 44.1, pp. 68–85. DOI: 10.1175/JPO-D-13-0116.1 (cit. on p. 155).

Starr, V. P. (1945). "A Quasi-Lagrangian System of Hydrodynamical Equations". *Journal of Meteorology* 2.4, pp. 227–237. DOI: 10.1175/1520-0469(1945)002<0227:AQLSOH>2.0.CO;2 (cit. on p. 30).

Stommel, H. (1948). "The Westward Intensification of Wind-Driven Ocean Currents". *Eos, Transactions American Geophysical Union* 29.2, pp. 202–206. DOI: 10.1029/TR029i002p00202 (cit. on p. 27).

— (1957). "A Survey of Ocean Current Theory". *Deep Sea Research (1953)* 4, pp. 149–184. DOI: 10.1016/0146-6313(56)90048-X (cit. on p. 101).

Stommel, H. and A. Arons (1959). "On the Abyssal Circulation of the World Ocean — II. An Idealized Model of the Circulation Pattern and Amplitude in Oceanic Basins". *Deep Sea Research (1953)* 6, pp. 217–233. DOI: 10.1016/0146-6313(59)90075-9 (cit. on p. 27).

Straub, D. N. (1993). "On the Transport and Angular Momentum Balance of Channel Models of the Antarctic Circumpolar Current". *Journal of Physical Oceanography* 23.4, pp. 776–782. DOI: 10.1175/1520-0485(1993)023<0776:OTTAAM>2.0.CO;2 (cit. on p. 2).

Sun, D.-Z. and R. S. Lindzen (1994). "A PV View of the Zonal Mean Distribution of Temperature and Wind in the Extratropical Troposphere". *Journal of the Atmospheric Sciences* 51.5, pp. 757–772. DOI: 10.1175/1520-0469(1994)051<0757:APVOTZ>2.0.CO;2 (cit. on p. 66).

Sverdrup, H. U. (1947). "Wind-Driven Currents in a Baroclinic Ocean; with Application to the Equatorial Currents of the Eastern Pacific". *Proceedings of the National Academy of Sciences* 33.11, pp. 318–326. DOI: 10.1073/pnas.33.11.318 (cit. on p. 101).

Sverdrup, H. U., M. W. Johnson, R. H. Fleming, and others (1942). *The Oceans: Their Physics, Chemistry, and General Biology*. Vol. 7. Prentice-Hall New York (cit. on p. 1).

Tailleux, R. (2016). "Neutrality Versus Materiality: A Thermodynamic Theory of Neutral Surfaces". *Fluids* 1.4, p. 32. DOI: 10.3390/fluids1040032 (cit. on pp. 27, 63).

Tailleux, R. and J. C. McWilliams (2001). "The Effect of Bottom Pressure Decoupling on the Speed of Extratropical, Baroclinic Rossby Waves". *Journal of Physical Oceanography* 31.6, pp. 1461–1476. DOI: 10.1175/1520-0485(2001)031<1461:TEOBPD>2.0.CO;2 (cit. on pp. 95, 97).

Takahashi, S., T. Ikeda, Y. Shinagawa, T. L. Kunii, and M. Ueda (1995). "Algorithms for Extracting Correct Critical Points and Constructing Topological Graphs from Discrete Geographical Elevation Data". *Computer Graphics Forum* 14.3, pp. 181–192. DOI: 10.1111/j.1467-8659.1995.cgf143_0181.x (cit. on p. 15).

Talley, L. D. (1988). "Potential Vorticity Distribution in the North Pacific". *Journal of Physical Oceanography* 18.1, pp. 89–106. DOI: 10.1175/1520-0485(1988)018<0089:PVDITN>2.0.CO;2 (cit. on p. 12).

— (2013). "Closure of the Global Overturning Circulation Through the Indian, Pacific, and Southern Oceans: Schematics and Transports". *Oceanography* 26.1, pp. 80–97. JSTOR: 24862019 (cit. on p. 2).

Tamsitt, V., H. F. Drake, A. K. Morrison, L. D. Talley, C. O. Dufour, A. R. Gray, S. M. Griffies, M. R. Mazloff, J. L. Sarmiento, J. Wang, and W. Weijer (2017). "Spiraling Pathways of Global Deep Waters to the Surface of the Southern Ocean". *Nature Communications* 8.1. DOI: 10.1038/s41467-017-00197-0 (cit. on p. 129).

Tarjan, R. (1972). "Depth-First Search and Linear Graph Algorithms". *SIAM Journal on Computing* 1.2, pp. 146–160. DOI: 10.1137/0201010 (cit. on p. 45).

Toggweiler, J. and B. Samuels (1995). "Effect of Drake Passage on the Global Thermohaline Circulation". *Deep Sea Research Part I: Oceanographic Research Papers* 42.4, pp. 477–500. DOI: 10.1016/0967-0637(95)00012-U (cit. on p. 2).

Treguier, A.-M., I. M. Held, and V. D. Larichev (1997). "Parameterization of Quasigeostrophic Eddies in Primitive Equation Ocean Models". *Journal of Physical Oceanography* 27.4, pp. 567–580. DOI: 10.1175/1520-0485(1997)027<0567:POQEIP>2.0.CO;2 (cit. on p. 66).

Vallis, G. K. (2006). *Atmospheric and Oceanic Fluid Dynamics*. Cambridge, U.K.: Cambridge University Press (cit. on pp. 6, 11, 71, 72, 98).

Vallis, G. K. (2000). "Large-Scale Circulation and Production of Stratification: Effects of Wind, Geometry, and Diffusion". *Journal of Physical Oceanography* 30.5, pp. 933–954. DOI: 10.1175/1520-0485(2000)030<0933:LSCAPO>2.0.CO;2 (cit. on p. 155).

Veronis, G (1975). "The Role of Models in Tracer Studies". *Numerical models of ocean circulation*, pp. 133–146 (cit. on p. 62).

Veronis, G. (1981). "Dynamics of Ocean Circulation". *Evolution of Physical Oceanography*. Ed. by B. A. Warren and C. Wunsch. The MIT Press, pp. 140–183 (cit. on pp. 5, 132).

Viebahn, J. and C. Eden (2010). "Towards the Impact of Eddies on the Response of the Southern Ocean to Climate Change". *Ocean Modelling* 34.3-4, pp. 150–165. DOI: 10.1016/j.ocemod.2010.05.005 (cit. on p. 2).

Warren, B. A., J. H. LaCasce, and P. E. Robbins (1996). "On the Obscurantist Physics of "Form Drag" in Theorizing about the Circumpolar Current". *Journal of Physical Oceanography* 26.10, pp. 2297–2301. DOI: 10.1175/1520-0485(1996)026<2297:OTOPOD>2.0.CO;2 (cit. on p. 107).

Weber, G. H., S. E. Dillard, H. Carr, V. Pascucci, and B. Hamann (2007). "Topology-Controlled Volume Rendering". *IEEE Transactions on Visualization and Computer Graphics* 13.2, pp. 330–341. DOI: 10.1109/TVCG.2007.47 (cit. on p. 15).

Welander, P. (1971). "The Thermocline Problem". *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. DOI: 10.1098/rsta.1971.0081 (cit. on p. 27).

Welander, P. (1959). "An Advective Model of the Ocean Thermocline". *Tellus*. DOI: 10.1111/j.2153-3490.1959.tb00036.x (cit. on p. 95).

Werghi, N., Yijun Xiao, and J. Siebert (2006). "A Functional-Based Segmentation of Human Body Scans in Arbitrary Postures". *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)* 36.1, pp. 153–165. DOI: 10.1109/TSMCB.2005.854503 (cit. on p. 15).

White, A. A. (1990). "Steady States in a Turbulent Atmosphere". *Meteorol.Mag.* 119.1410, pp. 1–9 (cit. on p. 65).

Whitworth, T. and R. G. Peterson (1985). "Volume Transport of the Antarctic Circumpolar Current from Bottom Pressure Measurements". *Journal of Physical Oceanography* 15.6, pp. 810–816. DOI: 10.1175/1520-0485(1985)015<0810:VTOTAC>2.0.CO;2 (cit. on p. 3).

Wüst, G (1935). "The Stratosphere of the Atlantic Ocean". *Scientific Results of the German Atlantic Expedition of the Research Vessel "Meteor" 1925–27* 6 (cit. on pp. 8, 27).

Xue, H.-j. and Y.-l. Gu (2009). "Application of Topological Analysis in Ocean Feature Extraction". *Computer Engineering* 3, p. 093 (cit. on p. 15).

Young, W. R. (2010). "Dynamic Enthalpy, Conservative Temperature, and the Seawater Boussinesq Approximation". *Journal of Physical Oceanography* 40.2, pp. 394–400. DOI: 10.1175/2009JPO4294.1 (cit. on pp. 42, 113).

— (2012). "An Exact Thickness-Weighted Average Formulation of the Boussinesq Equations". *Journal of Physical Oceanography* 42.5, pp. 692–707. DOI: 10.1175/JPO-D-11-0102.1 (cit. on pp. 64, 66, 69, 78, 97, 102, 103, 112, 115, 142, 146, 156, 163, 165).

Zajaczkovski, U. (2017). "A Study of the Southern Ocean: Mean State, Eddy Genesis & Demise, and Energy Pathways". 125 pp. (cit. on p. 136).

Zhang, H.-M. and N. G. Hogg (1992). "Circulation and Water Mass Balance in the Brazil Basin". *Journal of marine research* 50.3, pp. 385–420. DOI: 10.1357/002224092784797629 (cit. on pp. 27, 31, 116).