


UNIVERSITY OF READING

Department of Meteorology



Explaining Timescales Associated With Jet Stream Variability

DOMINIC WILLIAM JEAN-LUC JONES

A thesis submitted for the degree of Doctor of Philosophy

November 2021

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

-Dominic William Jean-Luc Jones

Abstract

Extreme weather events are often the result of slow-moving, large-scale wave patterns. Greater understanding of these large-scale modes of variability would allow us to anticipate how quasi-stationary modes might depend on changes to the background state in future climate. In this thesis the Empirical Normal Mode (ENM) technique, a technique for extracting dynamical modes of variability from atmospheric timeseries data, is utilised to examine the dependence of mode structure and frequency on jet latitude.

The ENM methodology is extended to include the lower troposphere spanned by isentropic surfaces that can intersect the ground. This involves careful accounting of the terms in large amplitude pseudomomentum and pseudoenergy associated with this region of the atmosphere - terms that contribute to the "boundary wave activity" in the limit of small amplitude.

In the third chapter, the implementation of the technique itself is validated by testing the characteristic 'intrinsic' phase speed of the ENMs against an empirical phase speed derived from the modes' principal component timeseries, using a set of idealised model experiments simulated using the Reading IGCM2.2. It is found that the phase speed matching conditions are met for the dominant freely propagating baroclinic modes, validating the approach to the calculation of wave activity and some approximations used in deriving relevant wave activity norms.

In the fourth chapter, a new series of idealised experiments are devised that possess a jetstream with controllable latitude such that the change in behaviour of the modes of variability with a shift in jet latitude may be examined. The initial and relaxation temperature field in thermal wind balance with a prescribed zonal wind field with jet latitudes ranging from $\sim 40^\circ$ to $\sim 65^\circ$ is found, and a sloping tropopause is added in order to maintain baroclinicity.

Subsequently, in the fifth chapter, the ENM structures of these experiments are found, and the change in the phase speed of the modes as the jet latitude changes is explored. A quasi-stationary branch of modes is identified which is associated with the most perturbation energy (for each zonal wavenumber) and therefore can propagate most strongly westwards against the background state westerly flow. As the jet is shifted polewards, the wavelength of the most energetic modes remains approximately the same, but they shift to lower zonal wavenumbers due to reduction in the latitude circle circumference.

Acknowledgements

I have been very lucky to have had the support of numerous people while undertaking this thesis, both directly and indirectly, to whom I would like to offer my heartfelt thanks.

Firstly, I would like to thank my incredible supervisory team; Prof. John Methven, whose boundless patience and continuous guidance was invaluable towards not just the completion of this thesis but towards my development as a researcher. I hope you are never again subjected to having to stay in the office until 10pm trying to hunt down bugs in FORTRAN code with a hapless 1st year PhD student. Dr. Tom Frame, whose understanding and reassurances were vital in getting me through any teething problems the project had and to whom I owe a great many hours of my life that he saved with his impressive library of Matlab code. Finally, my external supervisor Dr. Paul Berrisford at the ECMWF, whose expertise proved invaluable and whose measured but vital comments throughout the past 4 years have helped in both my understanding and the work that's been undertaken.

Extended thanks to the members of my Monitoring Committee, Prof. Maarten Ambaum, Dr. Ben Harvey, & Prof. Rowan Sutton who offered interesting discussions which lead to fruitful avenues of inquiry, and for helping to keep the project on track (even when I wasn't!). Thanks to members of the Dynamical Processes Group at Reading for showing interest in my work, and providing useful discussion and perspectives throughout.

Throughout this project I've relied on the help of Maria Broadbridge, Pawel Stasiak, & Eric Mathieu who helped solve innumerable computing issues, thank you all. On the administrative side, Wendy Neale has been a persistent pillar of support- thanks for checking up on me Wendy.

Thanks to all my amazing friends who kept my spirits up in all manner of ways; Val, Ed, Bryn, Kenny, Miguel, Anna, Polkey, & Asher whether it was skype calls, epic cycle rides, meeting for pints, brunches, walks or anything else. Special thanks to Keanu who also provided soundboarding for code and ideas. Herbert & Iestyn for your companionship over the past couple months and always. To the Cardinals; Hannah, Harry, & Miriam for letting me make wonderful music with them, & my fave desk partner Ben.

Undertaking this PhD in Reading wouldn't have been the same if it weren't for the incredible community of researchers here. Thanks to all my colleagues turned friends; "Nana" Hannah, Kaja, Sally (who provided the template this thesis was written on!), Izzy,

Alec, Jake, Mark, Meg & more. To Andrea, with whom countless tea breaks were enjoyed & who knew when to say ‘let’s a go’ and call it a day. Special thanks to shenanigans-personified Dan and Beth for their endless support and never failing to make me laugh.

My wonderful family; Mam, Friday, Sabrina, Billie, & ‘Granny’ Sheena who’ve solidly been giving me love and encouragement and have been here every step of the way even during the lowest points. I am beyond lucky to have you in my corner and appreciate everything you’ve done for me.

Thank you Saana, the ‘jetstream of my heart’ for all your love and understanding. I couldn’t have asked for anyone better to share my life with while doing this PhD.

Finally, thanks to my cat Beetlejuice, the best writing up buddy anyone could ever ask for, who was with me right till the end. I miss you cat-face, this is for you.

This was a NERC funded PhD project in the SCENARIO DTP, in partnership with the ECMWF. It’s still hard to believe that I was given this incredible opportunity, for which I will be forever grateful.

Contents

| | |
|--|------------|
| Declaration | i |
| Abstract | iii |
| Acknowledgements | iv |
| Table of contents | vi |
| List of Acronyms | ix |
| 1 Introduction | 1 |
| 1.1 Motivation | 1 |
| 1.2 Background | 2 |
| 1.2.1 Rossby Wave Propagation Mechanism | 2 |
| 1.2.2 Stationary Rossby Waves in the Atmosphere | 4 |
| 1.2.3 Baroclinic Waves and Modal Growth | 5 |
| 1.2.4 Vortex Erosion Mechanism | 7 |
| 1.2.5 Evidence for Dependence of Slow Modes on Westerly Jet Properties | 8 |
| 1.2.6 Extracting Patterns of Variability | 9 |
| 1.3 Direction of Novel Research | 10 |
| 1.4 Aims of the thesis | 12 |
| 2 Methodology: Background State, Wave Activity, and Empirical Normal Mode calculation | 14 |
| 2.1 Introduction | 14 |
| Part I: Calculation of Background State & Conserved Wave Activity Quantities | 15 |
| 2.2 Calculation of the Background State | 15 |
| 2.2.1 Mass and Circulation Integrals | 15 |
| 2.2.2 PV Equation for Inversion | 16 |
| 2.2.3 Boundary Conditions on the PV Inversion | 18 |
| 2.2.4 Numerics for PV Inversion | 20 |
| 2.3 Calculation of Wave Activity | 21 |
| 2.3.1 Defining Wave Activity | 21 |
| Part II: Novel Development of the Empirical Normal Mode Technique | 28 |
| 2.4 Calculating Empirical Normal Modes | 28 |
| 2.4.1 Fourier Filtering | 28 |
| 2.4.2 Eigenvalue Decomposition | 32 |
| 2.4.3 Calculation of \mathcal{P}_{ext} & \mathcal{H} by Projection | 34 |
| 2.4.4 Pairing ENMs | 35 |
| 2.4.5 Dynamical Phase Speed C_{int} | 36 |

| | | |
|----------|---|-----------|
| 3 | The Phase Speed Condition Test | 38 |
| 3.1 | Introduction | 38 |
| 3.2 | Experimental Design | 40 |
| 3.2.1 | Model | 40 |
| 3.2.2 | Specification | 41 |
| 3.2.3 | Relaxation | 41 |
| 3.2.4 | Drag | 43 |
| 3.2.5 | Initial Conditions | 43 |
| 3.2.6 | Experiment List | 43 |
| 3.2.7 | Experiment Overview | 44 |
| 3.3 | Evolving Background States | 47 |
| 3.4 | Empirical Normal Modes | 53 |
| 3.4.1 | Experiments without Drag | 54 |
| 3.4.2 | Experiments with Rayleigh Drag | 56 |
| 3.5 | Phase Speeds | 60 |
| 3.6 | Robustness | 66 |
| 3.6.1 | Eigenvalue Ratio | 66 |
| 3.6.2 | Sample Frequency | 67 |
| 3.6.3 | Sample Window | 68 |
| 3.7 | Conclusions | 70 |
| 4 | Relaxation Experiments With Controlled Jet Latitude Variation | 72 |
| 4.1 | Introduction | 72 |
| 4.2 | Experiment Specification | 73 |
| 4.2.1 | IGCM Parameters | 73 |
| 4.2.2 | Relaxation Field Definition | 73 |
| 4.3 | Parameter Space | 77 |
| 4.3.1 | Blending Jet Specification with Sloping Tropopause in Relaxation States | 77 |
| 4.3.2 | Relaxation Latitude | 78 |
| 4.4 | Experiment Results | 79 |
| 4.4.1 | Initial Spin Up | 79 |
| 4.5 | Background States | 82 |
| 4.5.1 | Latitude of the Jet Maximum | 83 |
| 4.5.2 | Speed | 85 |
| 4.5.3 | Modified Lagrangian Mean Background State Structure | 87 |
| 4.5.4 | Global Wave Activity | 89 |
| 4.6 | Conclusions | 90 |
| 5 | Slow Modes of Variability and Dependence on Jet Latitude | 93 |

| | | |
|----------|--|------------|
| 5.1 | Introduction | 93 |
| 5.2 | ENM Calculation | 94 |
| 5.3 | ENM Mode Distribution | 94 |
| 5.4 | Sorting Modes | 96 |
| 5.4.1 | Pair-matching | 96 |
| 5.4.2 | Structure Sorting | 98 |
| 5.4.3 | Phase Speed | 101 |
| 5.5 | Quasi-Stationary Branch | 107 |
| 5.6 | Conclusions | 112 |
| 6 | Conclusions & Future Work | 114 |
| 6.1 | Introduction | 114 |
| 6.2 | Verification of the ENM technique for Baroclinic Flows | 114 |
| 6.2.1 | Decomposition into Modes & Reduction of Dimensionality | 114 |
| 6.2.2 | The Phase Speed Condition Test | 116 |
| 6.3 | Variation of Phase Speed with Jet Latitude | 117 |
| 6.3.1 | Variation of Mode Structure with Latitude | 118 |
| 6.3.2 | Quasi-Stationary Waves | 118 |
| 6.4 | Further Research | 119 |
| 6.4.1 | Effect of Non-Linear Coupling | 119 |
| 6.4.2 | Effect of Jet Strength on the Modes | 120 |
| 6.4.3 | Seasonality Effects | 121 |
| | References | 122 |

List of Acronyms

| Acronym | Definition |
|---------|--|
| ECMWF | European Centre for Medium-Range Weather Forecasts |
| ENM | Empirical Normal Mode |
| LC | Life Cycle |
| MLM | Modified Lagrangian Mean (background state) |
| PV | Potential Vorticity |
| QSB | Quasi-Stationary Branch (of Rossby wave modes) |
| WA | Wave Activity |

Chapter 1:

Introduction

1.1 Motivation

Each year, extremes in weather cause billions of pounds of damage in Europe ([European Environment Agency, 2021](#)), the US ([National Oceanic and Atmospheric Administration, 2021](#)), Africa ([Wang et al., 2009](#)), Asia and the Pacific ([Hashim and Hashim, 2016](#)); globally, losses are estimated at \$3.47tn ([World Economic Forum, 2019](#)). Economic losses do not capture the full extent of the disruption and damage caused by extreme weather events; between 1998 and 2017 extreme weather was estimated to have caused 526,000 deaths worldwide, and primarily countries with a low economic development index bear the brunt of this cost ([World Economic Forum, 2019](#)).

With high impact weather exacting such a heavy economic and human toll, it is vital to understand the causes of extreme weather events in order to be able to predict their occurrence, such that their impacts might be mitigated. This is especially true as a result of climate change, which has already caused changes in the frequency and intensity of extreme weather events ([Arneth et al., 2019](#)).

Often, the impact of an extreme weather event is a result not of the short-term synoptic attributes such as a single rainfall event or warm day, but a long-lasting set of atmospheric conditions in a given location. These subseasonal climate variations with high persistence such as blocking events, or slow-moving waves are the result of large-scale atmospheric variability. A Rossby wave structure which maintains its position due to a low phase speed can create conditions which lead weather events to have an outsize impact. Our ability to forecast the large-scale structure of the atmosphere is then key in mitigating and preparing for high-impact weather events. For example, these large-scale stationary waves control the position of the storm tracks ([Branstator \(1995\)](#), [Chang et al.](#)

(2002)), which are linked to high-intensity weather events (Shaw et al., 2016).

In addition to extreme events, large-scale variability can have an impact on forecasting. Even now, it is possible for very large-scale forecast busts to result from the mischaracterisation of the large-scale circulation, such as blocking events (Grams et al., 2018). Low zonal wavenumber planetary structures with global spread can also be mischaracterised, such as a recent ECMWF forecast which displayed an opposite wavenumber-2 phase to the global pattern observed (Leonard, 2021) at a 3 week lead time.

1.2 Background

1.2.1 Rossby Wave Propagation Mechanism

Motions observed in the atmosphere commonly exhibit wavelike characteristics. Atmospheric waves can be broken down into three groups; acoustic waves, gravity waves and Rossby waves. Acoustic waves, or compression waves transfer energy through compressions in the atmosphere with characteristic speeds based on temperature and density. They can generally be ignored when considering motion at large scales. Gravity waves are a form of buoyancy waves, their motion controlled by the buoyant restoring force for a stratified fluid in the presence of gravity and can grow large enough to be affected by the Coriolis force (inertia-gravity waves), and can make an important contribution to large-scale dynamics, requiring consideration even if they are ultimately ignored. Rossby waves, of which planetary waves are a subset, are large-scale waves in Potential Vorticity (PV) gradients in the atmosphere. Potential vorticity is defined as;

$$q = \frac{\zeta \cdot \nabla \theta}{\rho} = \frac{\zeta_{\theta}}{r} \quad (1.1)$$

where ζ is the absolute vorticity vector, and ζ_{θ} is its vertical component in isentropic co-ordinates; θ is the potential temperature; ρ is the density of air; and r is the isentropic density; PV is a useful tracer for studying atmospheric dynamics which may be inverted to find the full flow field. Rossby waves are driven by variation in vorticity perpendicular to the flow, for example the meridional variation in the vertical component of planetary vorticity in the presence of a zonal flow. While earlier research indicated the existence of large-scale waves on a rotating sphere, it was Rossby (1939) who first described their propagation through the β -plane approximation; where latitudinal variation in the Coriolis parameter is approximated by a linear value.

Consider the barotropic vorticity equation in the beta-plane approximation;

$$\frac{D}{Dt}(\xi + f) = 0 \quad (1.2)$$

where ξ is the relative vorticity. Expanding the material derivative as;

$$\left(\frac{\partial}{\partial t} + u \frac{\partial}{\partial x} + v \frac{\partial}{\partial y} \right) \xi + \beta v = 0 \quad (1.3)$$

Where D/Dt represents the material derivative, that is the rate of change of a quantity with time as measured in an in-situ fluid parcel; f is the Coriolis parameter, the contribution to the vorticity from the Earth's rotation, and β is the approximated gradient of f in the meridional; $\partial f/\partial y$. By linearizing this equation about an appropriate background state (i.e. time-invariant, & zonally symmetric) we can examine the propagation of small-amplitude waves under this approximation. Taking u_0 as our basic flow, and u_e as the perturbation upon it (with corresponding background and perturbation PV quantities q_0 & q_e), and dropping products of perturbation quantities;

$$\left(\frac{\partial}{\partial t} + u_0 \frac{\partial}{\partial x} \right) q_e + \beta v_e = 0 \quad (1.4)$$

Consider a PV wave solution to this equation of the form;

$$q_e = q_{k,l} e^{i(kx+ly-\omega t)} \quad (1.5)$$

Where k and l here represent wavenumbers in the zonal and meridional directions respectively and ω is the frequency of the wave. As the PV is related to the flow by the streamfunction in the barotropic model by: $\nabla^2 \psi = q - f = q - \beta y$, q_e can be substituted into in the perturbation equation, and it is possible to solve for the zonal speed of the Rossby waves, giving the barotropic dispersion relation;

$$c = u_0 - \frac{\beta}{k^2 + l^2} \quad (1.6)$$

Although the scale dependence varies depending on the dynamics used, using the barotropic vorticity equation, Rossby waves derive their phase speeds from a combination of advection by the basic state (u_0), and propagation along vorticity contours westwards relative to the flow. This propagation is proportional to the square of the wavenumber of the wave; larger waves, with smaller wavenumber k & l propagate faster against the flow.

Indeed it is possible for larger waves to be stationary relative to the ground given the right background state. The specific square power of k and l are a result of the dynamics of the vorticity equation; for other dispersions, this detail might vary, but that the propagation speed of larger waves is higher is a general result.

While the idealised relations given here seem relatively straightforward nowadays, and would be insufficient to make predictions about the local behaviour of the actual atmosphere (indeed, Rossby waves can exist on boundary potential temperature gradients), they nonetheless convey the dynamics of its large-scale behaviour well, and are a good description of the “bigger picture” of the motion of planetary waves, and when first published, were the first real theoretical picture of large-scale motion on shear flows. The ideas considered here, such as advection of planetary waves by the basic state flow versus propagation ‘upstream’ against it are still key concepts when considering the speed of Rossby waves.

1.2.2 Stationary Rossby Waves in the Atmosphere

Quasi-stationary waves in the atmosphere may be characterised into distinct regimes on the basis of season, with separate regimes in winter and summer, and transitions between them during the equinoxal seasons. The northern and southern hemispheres exhibit distinct behaviour due to the vast differences in orography, as well as landmass, between them, as orography is a key driver of stationary waves ([Andrews et al., 1987a](#)).

In the northern hemisphere, the variance of planetary waves is overwhelmingly dominated by the presence of large waves forced by orography, as well as land-sea contrast at the surface. This land-sea contrast can be regarded as similar to orographic forcing in isentropic co-ordinates ([Brayshaw et al., 2008](#)). The Tibetan plateau and Rocky mountain ranges in particular are responsible for the excitation of large-scale stationary waves ([Chen \(2010\)](#), [Charney and Eliassen \(1949\)](#), [Hoskins and Karoly \(1981\)](#), & [Held et al. \(2002\)](#)). Orographic forcing of waves can be explained using simple barotropic theory; fluid parcels travelling upslope are squashed in the vertical, leading to generation of anticyclonic vorticity, and stretched over the ridge on the downslope giving rise to cyclonic vorticity. Another significant forcing for stationary wave patterns visible in the time-mean is thermally forced waves, which contribute comparably to the global circulation ([Smagorinsky \(1953\)](#), [Manabe and Terpstra \(1974\)](#)). These stationary waves result in a dominant wavenumber two pattern, with ridges above the eastern sections of the two oceans, and troughs above the eastern portions of the continental landmasses ([Lau, 1979](#)).

This wavenumber two pattern has wide-reaching effects on global weather patterns; contributing to the dryness of the continental US and Asian plateau in winter (Broccoli and Manabe, 1992); the precipitation of the East Asian monsoon region in summer (Chen, 2010); and the wintertime warmth of western Europe (Seager et al., 2002).

Due to the relative lack of orographic forcing in the southern hemisphere, the structure of waves is different to the northern hemisphere, with a bigger reliance on thermal forcing from low-latitude ocean basins, while at higher latitudes orographic forcing from the Antarctic and Andes mountain ranges still has an effect (Quintanar and Mechoso, 1995) despite their narrow band of latitudes compared to planetary scale waves.

Aside from stationary waves present in the global time-mean, more attention has been paid recently to transient slow-moving waves and the high impact weather they cause. Recently, extreme weather events such as heatwaves in North America, Western Europe, and the Caucasus; and extreme rainfall in south-east Europe and Japan; have all been linked to a dominant global Rossby-wave number seven pattern (Kornhuber et al., 2019). Due to the formation of circumpolar or nearly circumpolar stationary wave trains, conditions required for extreme weather events (such as heatwaves or floods) may co-occur at distant longitudes due to these teleconnections; this has been shown for heatwaves across Eurasia by Wu et al. (2012) & Zhou and Wu (2016).

1.2.3 Baroclinic Waves and Modal Growth

Physical systems can be characterised by ‘normal modes’; these are periodic sinusoidal motions which can be exhibited by the system, where all parts of the system move with the same frequency, or a set number of frequencies with a fixed relationship between them, and with a fixed phase relation between them. Normal modes are orthogonal, and as such changes in the amplitude of one do not affect the amplitude of the others; in fact they form a basis set from which all possible motions of the system may be derived or reconstructed. Normal modes of physical systems have been investigated since at least the 1700s, when first, Brook Taylor described the normal modes of a simple harmonic oscillator, and Daniel Bernoulli subsequently conjectured that the general solution of the equation of motion of an oscillator is a (possibly infinite) linear combination of its normal modes.

By the mid 20th century, models of the atmosphere were being developed for which modes could be calculated. The first of these models, developed independently by Eady (1949) and Charney (1947), were possible due to the simultaneous development of quasi-

geostrophic theory. In Eady’s model, a rigid upper boundary allows instability (and hence modal growth), as the upper-boundary potential temperature perturbation grows to interact with the lower boundary. A short-wave cut-off for modes exists in the Eady model because as modes increase in wavenumber and the propagation against the flow weakens, the strength of the interaction between the upper and lower level disturbance decreases, and phase-locking between the two cannot be maintained in the presence of wind-shear between the vertical levels.

Conversely, in Charney’s model, an interior PV gradient (not included in the Eady model) allows for a counter-propagating Rossby wave propagating westwards relative to the eastwards flow. The boundary thermal gradient, which is equivalent to a negative PV gradient, supports a counter-propagating wave propagating eastwards where the flow is weak. Indeed [Charney and Stern \(1962\)](#) would go on to establish that small disturbances cannot grow exponentially in a flow where the sign of the PV gradient on isentropic surfaces is everywhere-positive. An additional condition on the existence of baroclinic instability is the Fjørtoft condition ([Charney, 1973](#)), which requires a positive correlation between the zonal flow and meridional PV gradient. The first of these conditions is necessary for counter-propagating waves to amplify mutually, if their phase difference is correct, and the second is necessary for the waves to lock in phase, and therefore enable modal growth with fixed perturbation structure despite advection by the shear flow.

Eady’s approach, where there exists no meridional PV gradient in the interior, only allowed waves on the potential temperature gradients on the top and bottom boundaries, and the resultant shortwave instability cutoff were later expanded upon by [Bretherton \(1966\)](#) in a two-layer model. Bretherton’s model shows that the presence of an upper-level (for example) “differential” wave can result in the development of a phase locked “mean” wave in the lower layer, propagating in the opposite direction. The two waves exist in a normal mode configuration, and serve to re-inforce one another, resulting in mutual growth. In addition to the shortwave cutoff (and an additional, mathematically equivalent longwave instability cutoff introduced in Bretherton’s paper), instability can only occur if the PV gradients within the two layers (or on the two surfaces in the Eady model case) are of the opposite sign ([Charney and Stern, 1962](#)). While a powerful conceptual tool, Bretherton and Eady’s approaches do not account for the continuous PV gradient throughout the interior of a multi-layered atmospheric model, or indeed any kind of mean PV gradient not concentrated at the boundaries.

Due to the conceptual uses of Bretherton’s interpretation, a generalised counter-

propagating Rossby wave (CRW) model for situations where waves are not confined to the boundaries was developed by [Heifetz et al. \(2004a\)](#). Their result shows that the CRW normal mode growth mechanism applies in any situation where a PV gradient exists across a jetstream, with horizontally interacting barotropic waves resulting from a horizontal shear which they interact across, as in the much earlier model of [Rayleigh \(1880\)](#). Conversely, baroclinic waves consist of CRWs at different heights with an interaction between the lower and upper layer. Whereas in barotropic instability the reversal of the cross-stream PV gradient is on one level (or in a single-layer fluid), in baroclinic instability the “home-bases”, or opposing PV gradients are on different levels. In a mid-latitude westerly flow, this corresponds to a situation where the negative PV gradient is associated with the poleward decrease in temperature towards the pole, and the interior PV gradient is typically positive everywhere, especially at tropopause level.

For these growing normal mode configurations of CRWs, it is a requirement that the total pseudomomentum of the modes is zero, such that global conservation of pseudomomentum is not violated as the modes grow ([Methven, 2013](#)). In the CRW construction, the upper CRW has positive pseudomomentum and the lower CRW possesses negative pseudomomentum of equal magnitude such that their sum is identically zero. Therefore, the disturbances can amplify mutually while globally, pseudomomentum is conserved. As such, when extracting modes using the Empirical Normal Mode technique discussed in this thesis, only the upper CRW is picked out by design, as the ENMs are constructed such that they *maximise* pseudomomentum; specifically the pseudomomentum of the interior domain of the atmosphere. By taking the principal component associated with a mode, and projecting it onto the lower boundary to find \mathcal{P}_{ext} , the lower boundary pseudomomentum, a lower boundary wave with opposite pseudomomentum can be found. This wave on the lower boundary is phase-locked by design with the upper CRW, and as such the two make up a CRW pair behaving in a normal mode configuration.

1.2.4 Vortex Erosion Mechanism

At high latitudes, the atmosphere is dominated by high PV air, this is the stratospheric polar vortex and is separated from low-PV equatorial air by a jet along its edge which inhibits transport of air into or out of the vortex, acting as a barrier. As Rossby waves along the edge of the vortex grow and break they serve to mix the airmasses. One process through which mixing occurs is filamentation. As waves break, filaments of high PV air are stretched out into tendrils ([Polvani and Plumb, 1992](#)) through chaotic

advection (Aref, 1984) and subsequently mixed into surrounding air through turbulent and dissipative processes. This process serves to relax the vortex towards axisymmetry (Melander et al., 1987), and results in a sharper PV gradient and a stronger jet around the edge of the polar vortex.

1.2.5 Evidence for Dependence of Slow Modes on Westerly Jet Properties

Some evidence suggests that the presence, or propagation, of slow planetary waves depends on the properties of the jetstream. In the northern hemisphere, the annular mode is known as the Arctic Oscillation (AO). The variation in regimes present on the AO are closely linked to the latitude of the jetstream (Athanasiadis et al., 2010). The North Atlantic Oscillation (NAO), a localised expression of the AO (Ambaum et al., 2001), has been shown to display more persistence in the negative phase (Jia et al. (2007) Blessing et al. (2005)), which is correlated to a more equatorial jet latitude (Woollings et al., 2014).

Studies have also linked increased persistence of weather regimes to lower latitude jet positions (Barnes and Hartmann (2010) Woollings et al. (2010)), at least in the winter state. Barnes and Hartmann (2010) use stochastic stirring of the barotropic vorticity equation to produce a set of runs with jetstreams with a mean latitude of between $\sim 20^\circ$ and $\sim 50^\circ$. When the jet is close to the equator, annular mode events with both high and low phase persist for similar lengths of time, however in the mid-latitudes, events with high phase are less persistent than those with low phase. This asymmetry in persistence is attributed to spherical geometry inhibiting wavebreaking on the polar flank of the jet at higher jet latitudes. In re-analysis data, Woollings et al. (2010) use mslp and 500-hPa geopotential height and find statistical associations between negative NAO, where the transatlantic jet is in a southerly state, and persistent blocking events at higher latitudes.

These results are promising, but account only for annular modes (in the Barnes & Hartmann case) and variability in the northern hemisphere atlantic region (Woollings et al.), and other studies indicate that there may be a seasonal component to this relationship (Franzke et al., 2011).

In the southern hemisphere, the southern annular mode is more prominent due to the relative lack of orography. Model studies of the SAM have shown that lower latitude jets are associated with greater persistence (Kidston and Gerber (2010), Son and Lee (2005), Gerber and Vallis (2007), Barnes et al. (2010), & Simpson et al. (2010)).

1.2.6 Extracting Patterns of Variability

In order to study modes or patterns of variability in a physical system, it is necessary to be able to extract structures which describe that variability from the system. In classical mechanics, for a given set of dynamics used to describe a physical system it is possible to extract the normal modes directly from the equations of motion. By establishing a set of normal co-ordinates for the system which correspond to its degrees of freedom and evolve independently from one another, the motions for which these normal co-ordinates oscillate with the same frequency are the *normal modes*. The general solution of the system of equations is a linear superposition of these independent modal motions, which may be found by decoupling of the normal co-ordinates, or systematically through solution of an eigenvalue problem involving the equations of motion. In continuum mechanics and specifically geophysical fluid dynamics, characterising modal behaviour in models is more challenging, as discussed in section 1.2.3 above. Nonetheless; for a suitable description of the dynamics it is possible to extract the normal modes of oscillation *analytically* from the model used to describe the system.

In the real atmosphere, where patterns of variability must be extracted from time-series observations, statistical methods are generally used to find dominant patterns of variability. One such method is Empirical Orthogonal Function (EOF) analysis. An eigendecomposition of the covariance matrix, a matrix constructed from the inner product of the variance of a state variable with itself, is performed on a variable such as mean sea level pressure, or geopotential height. For a given state variable X , with matrix $\mathbf{X}_{s,n}$ containing the values of that state variable at each spatial point s and each time point n , the covariance matrix \mathbf{K} is;

$$\mathbf{K} = Cov(\mathbf{X}) = \frac{1}{N_t} \sum_{i=1}^{N_t} (\mathbf{X}_i - \bar{\mathbf{X}})' (\mathbf{X}_i - \bar{\mathbf{X}}), \quad (1.7)$$

of size $N_s \times N_s$ where N_s is the number of spatial points, N_t is the number of time points, and $\bar{\mathbf{X}}$ represents the time-mean of \mathbf{X} . Solving the following eigenvalue problem;

$$\mathbf{K}\mathbf{V}_x = \mathbf{V}_x\mathbf{\Lambda}, \quad (1.8)$$

yields \mathbf{V}_x , the matrix the columns of which are the eigenvectors, or spatial EOFs, and $\mathbf{\Lambda}$ is the diagonal matrix containing the eigenvalues. The eigenvectors are a set of orthogonal functions which describe the variability of state variable X . By construction,

the amount of variability described by the leading EOF (ordered by eigenvalue) is maximised, with each subsequent EOF explaining proportionally less of the spatial variability of the state variable in the input time-series.

EOF analysis is a powerful statistical tool used often in geophysical fluid dynamics to extract patterns of variability from atmospheric timeseries data (and is indeed used in the studies cited in the section above). However, EOFs do not correspond to individual dynamical modes, and do not correspond to kinematic degrees of freedom of the atmosphere (Monahan et al., 2009).

These two approaches; of finding normal modes using the dynamics of the equation set used to model the atmosphere; and of finding statistical patterns of variability in timeseries data, may be brought together by carefully choosing an inner product for the EOF calculation that is motivated theoretically by the global conservation of that product. This technique produces a set of functions, called Empirical Normal Modes, which describe the spatial variability, (as well as the amplitude timeseries, which may also be calculated for EOFs) which possess properties analogous to the normal modes of the primitive equations on a sphere, used to define the conservation quantity from which the inner product is constructed.

1.3 Direction of Novel Research

The Empirical Normal Mode technique presented here is a development on the theory presented by Brunet (1994). Brunet used data from a single level (500hPa) and a shallow water model to construct a set of ENMs. This was taken further by Zadra et al. (2002b) by considering multi-level data. Isentropic coordinates were used because the wave activity framework extends to the stratified primitive equations most readily in isentropic coordinates (Haynes (1988), Magnusdottir and Haynes (1996), Andrews et al. (1987a)). However, only the upper troposphere and lower stratosphere were considered and they did not extend their analysis down to isentropic surfaces that intersect the ground. The intersection adds considerably complication because you cannot use Fourier transforms in the zonal direction if the data does not wrap around a latitude circle. Also, there are boundary contributions to wave activity. Indeed, as discussed in Section 1.2.3, the boundary wave activity associated with disturbances to the gradient of potential temperature along the lower boundary are central to baroclinic instability, growing normal modes and more generally to transient baroclinic growth. Nonetheless, the analysis by Zadra was applied as a diagnostic study of the GEM's dynamical core Zadra et al. (2002b) and

NCEP re-analysis [Zadra et al. \(2002a\)](#).

The exact implementation used here is described in Chapter 2. However, in general the novel aspect of this implementation is the inclusion of additional parts of the wave activity quantity used for the ENM technique. Following [Methven and Berrisford \(2015\)](#), wave activity is calculated for points where isentropic surfaces used in the analysis intersect with the ground at certain latitudes, resulting in the inclusion of the ‘intersection’ and ‘exterior’ components of wave activity; the precise details are included in section 2.3 of Chapter 2. In brief, however, this has allowed us to probe through the troposphere to the lower boundary, whereas in previous implementations, upper-troposphere / lower-stratosphere variability was explored using ENMs.

In the experiments presented here, the atmosphere is relaxed to an unstable jet state with Rayleigh friction in the lower troposphere in order to obtain a set of atmospheric model timeseries data with (1) distinct dynamics (Chapter 3) and (2) a jet which can be varied in latitude (Chapters 4 & 5). A subset of the experiments presented by Zadra et al. use the Held-Suarez Newtonian relaxation field and the ENM technique is used there as an intercomparison of models with re-analysis data. In [Zadra et al. \(2002b\)](#), model experiments are done which use the ENM technique to verify the ability of the forcing regimes considered to reproduce variability and climatology seen in NCEP re-analysis.

The approach presented here offers several developments; first, the use of the Modified Lagrangian Mean as a background state differs significantly from the zonal mean approach taken by [Zadra et al. \(2002b\)](#); as shown in Chapters 3 & 4, the MLM varies more slowly, and exhibits faster zonal flow than the zonal average. As a result it is likely that the work by [Zadra et al. \(2002b\)](#) underestimates the advection term of the wave activity. The large-amplitude wave activities and their conservation properties used here are defined relative only to the MLM, not the zonal average.

Due to the aforementioned lower-boundary components of the wave activity, ENMs may possess structures that penetrate through the troposphere right to the surface, which as will be shown for the ENM structures in Chapters 3 & 5, are important for characterising baroclinic modes, and the boundary contributions to the wave activity have a significant effect on the phase speeds of the modes. The finite amplitude formulation of wave activity on isentropic surfaces intersecting the lower boundary (following [Magnusdottir and Haynes \(1996\)](#)) enables us to treat the wave activity on these surfaces in a rigorous fashion. This activity is crucial to the existence and propagation of baroclinic waves. The expression for wave activity used by Zadra is obtained by linearisation of the

wave activity about the background, zonal, mean state; here, the full non-linear wave activity is used in order to obtain values for the ‘exterior’ & ‘intersection’ domains, whereas within the interior domain, we use the same quadratic forms as Zadra for consistency with the requirements of the ENM calculation.

In addition to the Held-Suarez experiments, Chapter 4 details the development of model experiments with a latitudinally confined jetstream, controllable by a single parameter. This approach has not thusfar been undertaken in the literature and represents a novel line of inquiry. In examining the ENMs of these experiments, Chapter 5 addresses specifically the evolution of the stationary modes with jet latitude, controlled by the relaxation, to determine how they evolve.

1.4 Aims of the thesis

This thesis aims to examine how large scale patterns of variability evolve in structure and timescale as the background state of the atmosphere changes. In order to do that, the following principal hypotheses will be addressed;

H_1 The variability associated with large-amplitude disturbances on the jet can be decomposed into a small set of dominant modes with a distinct structure and intrinsic frequency which can be deduced from that structure.

H_2 This intrinsic frequency matches the frequency observed in the timeseries data of the mode

H_3 The dominant modes vary continuously with jet latitude and there is a relationship between mode frequency (for a given zonal wavenumber) and jet latitude

That atmospheric data may be decomposed into an orthogonal basis set of modes is true a priori. H_1 asserts specifically that the number of significant modes obtained is small. It is possible that a large number of modes are required to explain atmospheric variance, and that variability is widely distributed between them.

While the modes obtained may possess an intrinsic frequency, or characteristic phase speed, there are many factors which could contribute to that speed differing from the speed empirically observed in the model atmosphere. Non-linear behaviour such as wave breaking, or wave coupling, can modify the speeds such that they differ from the theoretical speed which would otherwise result from the modes’ structure. While testing H_2 , the conditions under which the speeds match are identified, and where they don’t, why that

is the case. While frequency is used to characterise the modes in past literature on ENMs (Brunet and Vautard, 1996), here the phase speed of the modes is used for convenience, as slowly moving structures with low frequencies are studied.

It is expected that should the background jet increase in speed, then longer wavelengths would be required to balance the advection (equation 1.6) to achieve slow moving modes. In the real atmosphere, it is observed in summer, when the jet is weaker, that shorter wavelengths are observed in the dominant stationary structures. Similarly, observation has shown a link between persistent structures; slow moving structures which result in persistent weather conditions, and background flow configurations where the jet is in a lower latitude state. In testing H_3 , a number of questions arise; should the modes vary continuously with latitude, do dominant waves present on more southerly jets possess lower frequency? Do a set of dominant quasi-stationary modes arise? How do the modes evolve with latitude?

Firstly, the implementation of the ENM technique, as well as the background material required for its use, will be discussed in Chapter 2. Part I of this chapter will deal with the currently published material concerning the MLM background state, and the conserved wave activity quantities which are used in the analysis. Part II consists of the novel development of the ENM technique. The veracity of the implementation of this technique will be discussed in Chapter 3 using idealised numerical model experiments. This is done in order to determine the constraints on the kinds of structures whose behaviour can be accurately captured by the ENM technique; a test of H_2 . This will also help to address H_1 ; it is expected that a dynamically simple model experiment should trivially give rise to a small set of modes which dominate variability.

Following this, using a set of model experiments defined in Chapter 4, the relationship between mean jet latitude and mode frequency are explored using the ENM technique in Chapter 5, in order to address H_3 , while providing more cases to answer H_1 . Finally, in Chapter 6 the conclusions from the previous chapters are summarised, and details of a proposed scheme of future work are laid out.

Chapter 2:

Methodology: Background State, Wave Activity, and Empirical Normal Mode calculation

2.1 Introduction

There are many steps involved with calculating Empirical Normal Modes from gridded timeseries data. In this chapter, Part I will first detail how the background state is determined, followed by the calculation of wave activity for large amplitude disturbances, described by the primitive equations on the sphere, using the previously derived background state, referring to the appropriate literature. Secondly, in Part II I will describe the novel development of the Empirical Normal Mode (ENM) technique. The specific calculation of the background state is key in order to ensure that the conservation properties that the ENM calculation relies on are applicable, and is therefore important background for the remainder of the thesis.

In Part I, the first two sections [2.2](#) & [2.3](#) are an essential summary of the relevant parts of [Methven and Berrisford \(2015\)](#) & [Methven \(2013\)](#) respectively, detailing firstly how the Modified Lagrangian Mean (MLM) background state used in the calculation is derived, and then the forms of the wave activity quantities used in the ENM calculation. This process of calculating the background state first requires an adiabatic rearrangement of the full flow into a zonally symmetric state such that integrals of mass and circulation are contained within a set of ‘equivalent’ latitudes (section [2.2.1](#)). Following this, these integrals are used to acquire the background state mass density and zonal flow via a PV inversion, simultaneously calculating the equivalent latitudes. Details of this PV

inversion are given in section 2.2.2. Subsequently, in section 2.3.1 wave activity quantities with conservation laws valid for large-amplitude disturbances are calculated using the background and perturbation variables provided.

In Part II section 2.4 details the Empirical Normal Mode technique; firstly by describing the Fourier filtering of perturbation data into zonal wavenumbers, and then the eigendecomposition in terms of one of the derived wave activity quantities, the pseudo-momentum, is shown in section 2.4.2, yielding the Empirical Normal Modes. This development of the ENM technique is novel and there is currently a manuscript in preparation with authors John Methven; Thomas Frame; Carlo Cafaro; Lina Boljka; and myself as contributors.

Part I: Calculation of Background State & Conserved Wave Activity Quantities

2.2 Calculation of the Background State

2.2.1 Mass and Circulation Integrals

The MLM background state is zonally symmetric and defined with respect to isentropic vertical co-ordinates and PV contours. As such, data, initially in gridded time-series form (from the IGCM, or else ECMWF re-analysis) with the terrain-following sigma vertical co-ordinate σ , is first interpolated to θ -levels. Once this vertical interpolation is complete, Ertel PV ($q = \zeta_\theta/r$) is calculated at each point; the isentropic vorticity with respect to theta is calculated through a centred finite difference formula.

The determination of isentropic density is somewhat more involved; firstly, the Montgomery potential is calculated through integration of the hydrostatic balance relation in isentropic co-ordinates 2.7. Following this, the first and second derivatives of M with respect to theta are calculated, which are used to calculate r (equation 2.8).

The MLM state is defined using volumes of integration, integrals of both mass $\mathcal{M}(Q, \Theta)$ and circulation $\mathcal{C}(Q, \Theta)$ found using the full 3D fields of the PV (q) and the isentropic density (r). The integration is performed over the area bounded by the contour $q = Q$, i.e. anywhere where the PV exceeds Q , including isolated cut-off ‘islands’ of positive PV. This is done over each isentropic surface.

$$\mathcal{M}(Q, \Theta) = \frac{1}{\Delta\theta} \int \iint_{q \geq Q} r a^2 d\lambda d\mu d\theta, \quad (2.1)$$

$$\mathcal{C}(Q, \Theta) = \frac{1}{\Delta\theta} \int \iint_{q \geq Q} \zeta a^2 d\lambda d\mu d\theta, \quad (2.2)$$

Where $\Delta\theta$ is the spacing between the chosen isentropic surfaces, a represents the radius of the earth and $\mu = \sin \phi$.

2.2.2 PV Equation for Inversion

With the constraint that the zonally symmetric background state is constrained by the mass and circulation integrals of the full flow, each contour Q of PV in the full flow will have an “equivalent” latitude in the background state (as on each isentropic surface, each Q contour has only one latitude, due to the zonal symmetry of the background and a requirement that the arrangement of PV is monotonic with respect to latitude). PV inversion can be used to obtain the zonal flow and density. The procedure described below explains how an iteration is required to find the equivalent latitudes of the PV contours in the MLM background state on every isentropic surface.

As the background state is zonally symmetric, Stokes’s theorem can be applied to the circulation \mathcal{C} in order to find an expression for the zonal flow associated with the background state, $U_0 = (u_0/a) \cos \phi$.

$$\mathcal{C}(Q, \Theta) = 2\pi a^2 [U + \Omega (1 - \mu_e^2)] \quad (2.3)$$

Where $\mu_e = \sin \phi_e$. In order to solve this equation for U , it is necessary to know the values of equivalent latitude $\mu_e(Q, \Theta)$ in the background state for each PV value. In a flow which is horizontally non-divergent, such as the Quasi-Geostrophic shallow water equations, the equivalent latitudes can be calculated directly from the area enclosed by a given PV contour, as the density is uniform in each isentropic layer. However, here, density can vary, and as such the mass enclosed by a given PV contour cannot be calculated from the area. Nonetheless, using an area integral, a sufficient ‘first guess’ at the equivalent latitudes can be made, even in the case where density is non-uniform;

$$\mu_e^1(Q, \Theta) = 1 - \frac{\mathcal{A}(Q, \Theta)}{2\pi a^2} \quad (2.4)$$

where;

$$\mathcal{A}(Q, \Theta) = \frac{1}{\Delta\theta} \iiint_{q \geq Q} a^2 \, d\lambda d\mu d\theta \quad (2.5)$$

however, to fully calculate the equivalent latitudes, a simultaneous solution of U and r is required.

In order to calculate both the background zonal wind field, and the background isentropic density, it is necessary to use the invertibility property of PV $q = \zeta/r$. By deriving appropriate expressions for both the vertical component of absolute vorticity (ζ) and isentropic density r , it is possible to describe a PV equation in terms of only the circulation \mathcal{C} and the Montgomery potential M . It is also necessary to specify a balanced flow for this equation; in this case, hydrostatic & gradient wind balance on the sphere. Solving for M allows us to subsequently calculate the background zonal flow and isentropic density. The Montgomery potential is defined as;

$$M = c_p T + gz. \quad (2.6)$$

Where c_p is the specific heat capacity of air, T is the temperature, g is the acceleration due to gravity and z is the geopotential height. Taking the derivative of M with respect to θ , hydrostatic balance can be expressed in terms of M and θ ([Andrews et al. \(1987b\)](#));

$$\frac{\partial M}{\partial \theta} = c_p \left(\frac{p}{p_{00}} \right)^\kappa \quad (2.7)$$

Taking the second derivative of M , $\partial^2 M / \partial \theta^2 = R/P_{00} (p/p_{00})^{\kappa-1} \partial p / \partial \theta$, an expression for r in terms of M can be derived;

$$r = -\frac{1}{g} \frac{\partial p}{\partial \theta} = -\frac{1}{g\Gamma} \frac{\partial^2 M}{\partial \theta^2} \quad (2.8)$$

Where the function;

$$\Gamma = \frac{R}{p_{00}} \left(\frac{1}{c_p} \frac{\partial M}{\partial \theta} \right)^{\frac{\kappa-1}{\kappa}}. \quad (2.9)$$

Once an expression for r has been derived, it may be used in combination with the definition of ζ on a sphere with angular velocity Ω with zonally symmetric flow;

$$\zeta = 2\Omega\mu - \frac{\partial U}{\partial \mu}, \quad (2.10)$$

in order to obtain a PV equation. By substituting both expressions into the definition of PV, we obtain the PV equation;

$$2\Omega\mu - \frac{\partial U}{\partial\mu} = -\frac{q}{g\Gamma} \frac{\partial^2 M}{\partial\theta^2}. \quad (2.11)$$

It is possible to replace $\partial U/\partial\mu$ in order to write equation 2.11 solely in terms of Montgomery Potential and \mathcal{C} . Firstly, an expression for $-\partial U/\partial\mu$ can be found by taking the derivative of the gradient wind balance on the sphere,

$$\left(2\Omega \sin\phi + \frac{U \sin\phi}{\cos^2\phi}\right) Ua + \frac{\cos\phi}{a} \frac{\partial M}{\partial\phi} = 0, \quad (2.12)$$

with respect to μ . Substituting this expression into the PV equation (2.11), it then becomes;

$$2\Omega\mu_e + \frac{\frac{U\Omega}{\mu_e} + \frac{U^2}{2\mu_e} \left(\frac{\sin^2\phi_e+1}{\cos^4\phi_e}\right) + \frac{1}{2a^2\mu} \frac{\partial}{\partial\phi_e} \frac{\partial M}{\partial\phi_e}}{\Omega + \frac{1}{\cos^2\phi_e}} + \frac{q}{g\Gamma} \frac{\partial^2 M}{\partial\theta^2} = 0. \quad (2.13)$$

Stokes' theorem (equation 2.3) can then be used to substitute for U in equation 2.13 in order to obtain an expression for the PV equation (2.11) which depends only on second order derivatives of M in *both* ϕ_e and θ , (aside from the nonlinear factor Gamma), and the circulation (which is known).

$$1 + \frac{1}{(2\Omega a \sin\phi_e \cos\phi)^2} \left[\frac{(1 + \sin^2\phi_e)^3}{2 \sin^2\phi_e \cos\phi_e} \frac{\partial}{\partial\phi_e} \left(\frac{\cos\phi_e}{1 + \sin^2\phi_e} \frac{\partial M}{\partial\phi_e} \right) - \frac{(1 + 3 \sin^4\phi_e)}{2 \sin^2\phi_e \cos\phi_e} \frac{\partial}{\partial\phi_e} \left(\cos\phi_e \frac{\partial M}{\partial\phi_e} \right) + \frac{q\mathcal{C}(q, \theta) \sin\phi_e}{\Gamma g \pi} \frac{\partial^2 M}{\partial\theta^2} \right] = 0. \quad (2.14)$$

With this equation, given the PV distribution (q) and Circulation functional ($\mathcal{C}(Q, \Theta)$) it is possible to solve for the Montgomery potential M , and from M , both the background zonal wind and the background isentropic density can be calculated using equations 2.11 and 2.8.

2.2.3 Boundary Conditions on the PV Inversion

In order to perform the PV inversion, it is necessary to establish the boundary conditions for the domain such that equation 2.14 can be solved.

Equator & Pole

As the method had been designed for a hemispheric domain, a lateral boundary condition is required at both the pole and the equator. As $\cos(\phi)$ reduces to 0 at the pole,

$$\cos \phi \frac{\partial M}{\partial \phi} = 0, \quad (2.15)$$

we can see from 2.12 that $U = 0$ at the pole. At the equator, $\sin(\phi) = 0$, and so, again from 2.12, we can see that that 2.15 also applies here. This is to ensure that gradient wind balance applies all the way to the equator, despite the Coriolis parameter tending to zero there. Refer to Methven and Berrisford (2015) section 2.2.1 for full details.

Top Boundary

The pressure on the isentropic surface defining the top boundary to the inversion domain is imposed from the Eulerian zonal average of the full state data. This pressure is used to specify the vertical derivative of M at the top boundary, using hydrostatic balance 2.7.

Lower Boundary

Following Bleck (1973), the lower boundary condition is;

$$M = gz + \theta \frac{\partial M}{\partial \theta}. \quad (2.16)$$

which is an identity obtained through substitution of hydrostatic balance 2.7 into the definition of M 2.6. This combination of M and its first derivative is a general form of boundary conditions for elliptic equations like 2.14; more commonly used are the Dirichlet boundary condition (where the solution at the boundary is prescribed) and Neumann boundary conditions (where the first derivative at the boundary is prescribed).

The lower boundary condition is met by fitting a quadratic curve, $M = a + b\theta + c\theta^2$, vertically through the two isentropic levels above the boundary subject to the lower boundary constraint;

$$M = gz_{\text{LB}} + \theta_{\text{LB}} \frac{\partial M}{\partial \theta} \quad (2.17)$$

Where z_{LB} and θ_{LB} are the geopotential height and potential temperature on the

lower boundary at a given position (ϕ, t) . However, as some isentropic surfaces intersect the ground, it is necessary to provide a lateral lower boundary condition in addition to the vertical lower boundary condition provided by 2.17. The lower boundary condition is used in conjunction with the fitted quadratic to extrapolate M below the surface which provides the lateral lower boundary condition.

Once M has been calculated, this quadratic curve also defines the pressure p_{LB} and density ρ_{LB} at the lower boundary of the background state. These quantities are used to calculate the geostrophic wind, defined below;

$$u_{\text{gLB}} = -\frac{1}{2\Omega a \sin \phi} \left(\frac{1}{\rho_{\text{LB}}} \frac{\partial p_{\text{LB}}}{\partial \phi} + g \frac{\partial z_{\text{LB}}}{\partial \phi} \right) \quad (2.18)$$

u_{gLB} can then be used to calculate the lower boundary background state wind, using;

$$u_g = u \left(1 + \frac{u}{2\Omega a \cos \phi} \right) \quad (2.19)$$

2.2.4 Numerics for PV Inversion

The inversion of the PV equation is made more complicated by the presence of the Γ function, which introduces weak non-linearity, as without it, equation 2.14 would be linear and elliptical. In addition, the lateral lower boundary condition defined above also introduces non-linearity. It is done via a two step process which first uses a linear inverter (routine D03EBF in the NAG library (NAG)) which makes an assumption about the form of Γ , followed by a step to recalculate Γ , as well as the lower boundary, from the results of the linear inverter, before feeding them back into it.

This is followed by an ‘outer’ iteration where differences in \mathcal{M} & \mathcal{C} between the inverted (i.e. background) state and the full flow are reduced by recalculation of the equivalent latitudes μ_e from the first-guess latitudes from equation 2.4. This is done by adjusting the equivalent latitudes at each value of $q = Q$ using estimates from both \mathcal{M} & \mathcal{C} ;

$$\delta \mu_e^{\mathcal{M}}(Q, \Theta) = -(\mathcal{M}_0 - \mathcal{M}) \frac{\partial \mu}{\partial \mathcal{M}} \mu_e^{\frac{1}{4}} \quad (2.20)$$

and

$$\delta \mu_e^{\mathcal{C}}(Q, \Theta) = -(\mathcal{C}_0 - \mathcal{C}) \frac{\partial \mu}{\partial \mathcal{C}} \mu_e^{\frac{1}{4}} \quad (2.21)$$

the contributions to the shift in μ_e are averaged from both the mass and circulation. Iterations are then performed until the difference in equivalent latitudes between iterations fall below a tolerance threshold.

This method does not guarantee the uniqueness of the solution obtained; it is possible that results obtained are local solutions, based on the position of the first-guess equivalent latitudes.

In summary, when performing the PV inversion, the following iteration procedure is followed;

1. Estimate the equivalent latitudes (equation 2.4)
2. Calculate q and then the integrals \mathcal{M} & \mathcal{C} (equations 2.1 & 2.2) on the (ϕ, θ) inverter grid
3. Calculate a first-guess of M using equation 2.14 given the circulation \mathcal{C} and the estimate of the equivalent latitudes provided in point (1). Isentropic density is then calculated using 2.8.
4. Start the non-linear iteration; calculate Γ (equation 2.9) and the lower lateral boundary condition (section 2.2.3) using the latest values of M
5. Perform linear iteration to recalculate M until the residual is sufficiently reduced that the required precision is achieved
6. Return to step 4. and iterate non-linearly, followed by the liner iteration 5., until the residual in M is sufficiently reduced that the required precision is achieved between non-linear iterations
7. Return to step 1. and perform the outer iteration by adjusting the equivalent latitudes (equations 2.20 & 2.21) until the equivalent latitude residuals is reduced to within the required precision.

2.3 Calculation of Wave Activity

2.3.1 Defining Wave Activity

A key requirement when defining a wave activity quantity is that it follows a conservation law. Such a law may be expressed mathematically by;

$$\frac{\partial A}{\partial t} + \nabla \cdot F = S \quad (2.22)$$

Where A is the given wave activity density quantity, the F term represents fluxes of that quantity, associated with propagation or advection, and S represents any sources or sinks of the wave activity. Globally, wave activity is conserved, as long as $S = 0$ (there are no sources or sinks), and there is no flux of wave activity through the lower boundary (assuming the domain to be unbounded above).

Invariants such as the global zonal angular momentum, or energy, which arise from rotational and time symmetries via Noether's theorem, usually form the basis for wave activity quantities due to their conservation properties. However, as wave activities measure the perturbation of a system relative to its background flow, these quantities taken alone are not invariant within the perturbation (or indeed the background state) as they may be exchanged between the perturbations and the background state.

Following the method presented in [McIntyre and Shepherd \(1987\)](#) it is possible to systematically construct further conservation laws through a combination of the invariants described above, with arbitrary functions of potential temperature and PV, $C(q, \theta)$, known as Casimirs. Both θ and PV are *materially* conserved in adiabatic and frictionless flows meaning that a fluid parcel does not change θ or q , though the total θ and q . This makes any function $C(q, \theta)$ a conservative quantity. Specific functions C may be chosen and combined with invariants in order to derive wave activities which are second order and globally conserved, however in order to utilise this method, the chosen background state must possess rotational and time symmetry in order for the zonal angular momentum and energy invariants to be used in these wave activity constructions.

Pseudo-Angular Momentum Conservation

Following [Haynes \(1988\)](#), the pseudo-zonal angular momentum density (henceforth abbreviated to pseudomomentum) is defined as

$$P(\lambda, \phi, \theta, t) = -r(Z + C) + r_0(Z_0 + C_0), \quad (2.23)$$

where r is the isentropic density, C is the Casimir density, an arbitrary function of θ and PV, and Z is the specific zonal angular momentum at latitude ϕ on a sphere rotating at rate Ω , defined as

$$Z = (u + a\Omega \cos \phi) \cos \phi, \quad (2.24)$$

where q_0 denotes background state PV, and C_0 specifically means $C(q_0, \theta)$. As Z and C are both conserved in the full flow and the chosen background state (in this case, the MLM), P is therefore also globally conserved, and thus C can be chosen so as to ensure that P is second order in disturbance quantities. It is then possible to use the relationship between the full flow, the background state, and the perturbations to derive an equation for P expressed in terms of perturbation quantities. Taking a Taylor expansion of the Casimir function about the background PV value q_0 and gathering all terms second order and higher into a residual term C_2 ,

$$C = C_0 + \left(\frac{\partial C}{\partial q} \right)_0 q_e + C_2(q_0, q_e, \theta), \quad (2.25)$$

where q_e refers to the perturbation PV, i.e. the difference between the full flow and the background state. After applying the identity $r q_e = \zeta_e - r_e q_0$, equation (2.23) can then be expressed in terms of both background state q_0 and perturbation q_e quantities.

$$\begin{aligned} P = & -rC_2 - r_e u_e \cos \phi - r_0 u_e \cos \phi \\ & - \left(\frac{\partial C}{\partial q} \right)_0 \zeta_e - r_e \left[Z_0 + C_0 - q_0 \left(\frac{\partial C}{\partial q} \right)_0 \right] \end{aligned} \quad (2.26)$$

ζ_e can be expanded in terms of the zonal and meridional velocity perturbations; $\zeta_e = (1/a \cos \phi) \partial v_e / \partial \lambda - (1/a \cos \phi) \partial (\mu_e \cos \phi) / \partial \phi$. Applying this and substituting it into 2.26 gives:

$$\begin{aligned} P = & -rC_2 - r_e u_e \cos \phi \\ & - \frac{1}{a \cos \phi} \frac{\partial}{\partial \lambda} \left[v_e \left(\frac{\partial C}{\partial q} \right)_0 \right] + \frac{1}{a \cos \phi} \frac{\partial}{\partial \phi} \left[u_e \left(\frac{\partial C}{\partial q} \right)_0 \cos \phi \right] \\ & - u_e \left[r_0 \cos \phi + \frac{1}{a} \frac{\partial}{\partial \phi} \left(\frac{\partial C}{\partial q} \right)_0 \right] - r_e \left[Z_0 + C_0 - q_0 \left(\frac{\partial C}{\partial q} \right)_0 \right] \end{aligned} \quad (2.27)$$

Here, the first line consists of quantities that are second order or higher; the C_2 term (i.e. the 2nd order and higher terms of the Casimir function) and the second term containing only perturbation quantities. The second line is a horizontal divergence. Hence, the terms of the final line must equal zero in order for the global integral of P to be

second order (as the divergence terms will vanish in the global integral). This places a constraint upon the Casimir function which can be suitably chosen to ensure that both terms are zero.

$$\frac{\partial}{\partial \phi} \left(\frac{\partial C}{\partial q} \right)_0 = -ar_0 \cos \phi \quad (2.28)$$

$$Z_0 + C_0 - q_0 \left(\frac{\partial C}{\partial q} \right)_0 = 0 \quad (2.29)$$

By integrating the first of these constraints with respect to ϕ , we can express it in terms of the mass integral (equation 2.1) from earlier.

$$\begin{aligned} \left(\frac{\partial C}{\partial q} \right)_0 &= -a \int_0^\phi r_0(\tilde{\phi}, \theta, t) \cos \tilde{\phi} d\tilde{\phi} \\ &= \frac{1}{2\pi a} [\mathcal{M}(Q, \theta) - \mathcal{M}_s(\theta)], \end{aligned} \quad (2.30)$$

Where \mathcal{M}_2 is the total mass of the isentropic shell in the hemisphere. This enables the calculation of the Casimir function using the previously calculated mass integral 2.1. Rearranging 2.25 for expression for C_2 ,

$$C_2(q_0, q_e, \theta) = C - C_0 - q_e \left(\frac{\partial C}{\partial q} \right)_0, \quad (2.31)$$

C_2 can be expressed in terms of the Mass and Circulation integrals by integrating over PV values, as shown in Thuburn and Lagneau (1999),

$$C_2 = \frac{1}{2\pi a} \left(q[\mathcal{M}]_{q_0}^q - [C]_{q_0}^q \right). \quad (2.32)$$

With P now expressed in a calculable form, the full global pseudomomentum (Magnusdottir and Haynes (1996)) can be derived by integration of 2.27 over different domains;

$$\begin{aligned} \mathcal{P} &= \int_{\bar{\mathcal{D}}} (-rC_2 - r_e u_e \cos \phi) a^2 \cos \phi d\lambda d\phi d\theta - \int_{\partial \bar{\mathcal{D}}} \left(\frac{\partial C}{\partial q} \right)_0 u_e \cos \phi a d\lambda d\theta \\ &+ \int_{(\mathcal{D} \cap \mathcal{D}_0) \setminus \bar{\mathcal{D}}} \left[-rC_2 - (r_0 + r_e) u_e \cos \phi - \left(\frac{\partial C}{\partial q} \right)_0 \zeta_e \right] a^2 \cos \phi d\lambda d\phi d\theta \\ &- \int_{\mathcal{D} \setminus (\mathcal{D} \cap \mathcal{D}_0)} r(Z + C) a^2 \cos \phi d\lambda d\phi d\theta + \int_{\mathcal{D}_0 \setminus (\mathcal{D} \cap \mathcal{D}_0)} r_0(Z_0 + C_0) a^2 \cos \phi d\lambda d\phi d\theta. \end{aligned} \quad (2.33)$$

Whereas before, P (in normal caps) referred to the wave activity density, the curly

\mathcal{P} notation will refer to the full global integral of wave activity.

In the “interior” domain, where isentropic surfaces of both the full and background flow do not intersect the ground (corresponding to the brown area on figure 2.1), the volume integral of the first two terms of 2.27 reduces to the first integral of 2.33. This is split into two terms; the Rossby wave, or advection term \mathcal{P}_w , which measures the wave activity from displacement of PV contours; and gravity wave term \mathcal{P}_g , which involves correlation between the isentropic density and the horizontal perturbation velocity- which does not measure Rossby wave activity, hence being termed the gravity wave term, as these waves are the dominant waves measured (though other waves such as Kelvin waves also possess this type of wave activity).

The boundary integral, (along ∂D in figure 2.1), reduces to the integral of the divergence term of 2.27, which is the second integral of 2.33, which we call the boundary term, \mathcal{P}_b .

In the “intersection” domain, where isentropic surfaces of both the full flow and background state are above the ground at certain longitudes (corresponding to the red area on figure 2.1) divergence terms are retained giving rise to the second line of 2.33. The resultant integral is referred to as the intersection pseudomomentum, \mathcal{P}_d .

In the exterior domains (corresponding to the yellow hatched area and green area of 2.1), where points are above the lower boundary only in either the background or full flow, the wave activity density must be calculated separately for the two domains, resulting in the final two integrals on the last line of 2.33, or the exterior domain pseudomomentum, \mathcal{P}_e .

These contributions to the pseudomomentum are summarized in table 2.1, along with the shorthands used when plotting timeseries of global pseudomomentum contributions.

In the small amplitude limit (Magnusdottir and Haynes (1996),Methven (2013)), the expression for the pseudomomentum density (equation 2.27) reduces to;

| Domain Name | Pseudomomentum Contribution | Plotting Shorthand |
|-----------------------|-----------------------------|--------------------|
| Interior Rossby Wave | \mathcal{P}_w | twatrop, twastrat |
| Interior Gravity Wave | \mathcal{P}_g | twag |
| Intersection | \mathcal{P}_d | twad |
| Boundary | \mathcal{P}_b | twab |
| Exterior | \mathcal{P}_e | twae |

Table 2.1: A summary of the contributions to Pseudomomentum.

$$\begin{aligned}
P = & \frac{r_o^2 \cos \phi_o q_e^2}{q_{oy}} - r_e u_e \cos \phi \\
& + \left(r_o^2 q_o \frac{y_{b_e}^2}{2} - r_o u_e y_{b_e} \right) \cos \phi_{b_o} \frac{\partial \theta_{b_o}}{\partial y} \delta(\theta - \theta_{b_o})
\end{aligned} \tag{2.34}$$

where $y = a\varphi$ and $q_{oy} = \partial q_o / \partial y$, ($q_{b_o} = Q$) is the background state PV contour everywhere coincident with the intersection of the isentropic layer with the ground, ϕ_{b_o} and θ_{b_o} are the latitude and potential temperature of the boundary in the background state, y_{b_e} is the value of y corresponding to $\phi_{b_e} = \phi_b - \phi_{b_o}$.

Pseudo-energy Conservation

Similarly to above, the pseudoenergy density, H can be defined as;

$$H(\lambda, \phi, \theta, t) = r(E + B) - r_o(E_o + B_o) \tag{2.35}$$

where E , the specific energy, is defined as;

$$E \equiv \frac{1}{2} (u^2 + v^2) + h(p, \theta) \tag{2.36}$$

and h is the specific enthalpy. In this case, the Casimir function is written B to distinguish it from the function used in (2.23). Similarly to before, it can be expanded in PV using a Taylor expansion;

$$B = B_o + \left(\frac{\partial B}{\partial q} \right)_o q_e + B_2(q_o, q_e, \theta). \tag{2.37}$$

Similarly, by expanding the enthalpy with respect to the pressure perturbation, an equation for h can be derived;

$$\begin{aligned}
h = & h_o + \left(\frac{\partial h}{\partial p} \right)_o p_e + h_2(p_o, p_e, \theta) \\
= & h_o + \left(\frac{\partial h}{\partial p} \right)_o p_e + \int_0^{p_e} (p_e - \tilde{p}) \left. \frac{\partial^2 h}{\partial \tilde{p}^2} \right|_{\theta} (p_o + \tilde{p}, \theta) d\tilde{p}.
\end{aligned} \tag{2.38}$$

Re-writing (2.36) in terms of the expression for enthalpy 2.38, and the Taylor expansion of the Casimir B (2.37) gives;

$$\begin{aligned}
H = & \frac{r}{2} (u_e^2 + v_e^2) + rB_2 + rh_2 + r(u_0u_e + v_0v_e) \\
& + r_e \left[E_0 + B_0 - q_0 \left(\frac{\partial B}{\partial q} \right)_0 \right] + \left(\frac{\partial B}{\partial q} \right)_0 \zeta_e \\
& - \left(\frac{\partial h}{\partial p} \right)_0 \frac{1}{2g} \frac{\partial p_e^2}{\partial \theta} + r_0 \left(\frac{\partial h}{\partial p} \right)_0 p_e
\end{aligned} \tag{2.39}$$

In the small amplitude limit, the expression for pseudoenergy density reduces (Methven (2013)) to to give

$$\begin{aligned}
H = & \frac{r_o}{2} (u_e^2 + v_e^2) + \frac{\kappa h_o}{gp_o\theta} \frac{p_e^2}{2} - \frac{r_o^2 u_o}{q_{o_y}} \frac{q_e^2}{2} + r_e u_o u_e \\
& + \left(-r_o^2 q_o u_o \frac{y_{b_e}^2}{2} + r_o u_o u_e y_{b_e} \right) \frac{\partial \theta_b}{\partial y} \cdot \delta(\theta - \theta_b).
\end{aligned} \tag{2.40}$$

These terms can be split up similarly to those in the pseudomomentum equation. The various terms are detailed in table 2.2. The first term in 2.40 represents the kinetic energy contribution to the pseudoenergy, \mathcal{H}_{KE} . The second corresponds to the available potential energy component from the energy, \mathcal{H}_{APE} . The third term is the interior Rossby Wave term, corresponding to the B_2 term in the full large-amplitude pseudoenergy, \mathcal{H}_w . This is similar to the \mathcal{P}_w term from the pseudomomentum, and corresponds to contributions to the wave activity from displacement of PV contours. The following term, \mathcal{H}_g is the 'gravity wave' term, measuring contributions from gravity and other waves such as Kelvin waves. The final two terms are the boundary (\mathcal{H}_b) and exterior (\mathcal{H}_e) corresponding to the same domains as in the pseudomomentum detailed above.

Generally, the gravity wave contributions are neglected due to their small magnitude.

| Domain Name | Pseudoenergy Contribution |
|-------------------------|---------------------------|
| Interior Rossby Wave | \mathcal{H}_w |
| Interior Gravity Wave | \mathcal{H}_g |
| Interior Kinetic Energy | \mathcal{H}_{KE} |
| Interior APE | \mathcal{H}_{APE} |
| Boundary | \mathcal{H}_b |
| Exterior | \mathcal{H}_e |

Table 2.2: A summary of the contributions to Pseudoenergy.

Part II: Novel Development of the Empirical Normal Mode Technique

2.4 Calculating Empirical Normal Modes

2.4.1 Fourier Filtering

In general, through the use of a basic state, a set of perturbations are able to be derived for key variables (u, v, p, r, q ; illustrated here using q) through a subtraction of the relevant basic state quantities from the full flow fields.

$$q_e(\lambda, \phi, \theta, t) = q(\lambda, \phi, \theta, t) - q_0(\phi, \theta, t) \quad (2.41)$$

However, in the following calculation, in order to ensure that the calculated ENMs are computed with respect to single zonal wavenumbers it is necessary to perform a Fourier filter in order to calculate the perturbations with respect to the background state at separate zonal wavenumbers. This Fourier transform is performed on the full grid-point data in (λ, ϕ, σ) space, as in θ space the lower boundary intersects with the ground at some latitudes. Thus the full data is transformed into phase space and partitioned according to its zonal wavenumber;

$$q(\lambda, \phi, \sigma) \xrightarrow{\text{FFT, partition}} \tilde{q}_m(\phi, \sigma) \quad (2.42)$$

For each desired zonal wavenumber then, the spectral space data is subsequently transformed back into gridpoint space retaining only the zonal average $m = 0$ and the desired wavenumber $m = M$. The resulting physical space perturbation $\hat{q}_m(\lambda, \phi, \sigma, t)$ is then interpolated onto theta levels, to provide $\hat{q}_m(\lambda, \phi, \theta, t)$. This quantity contains zonally symmetric components which must be removed to perform the ENM analysis; these components have two sources. One of which is differences between the MLM background state and the zonal average such that when the perturbations are calculated in equation (2.41), the difference between the two remains as a zonally symmetric component of q_e . In addition, the vertical interpolation of variables from σ to θ levels generates signals of wavenumbers other than the desired wavenumber $m = M$, especially in the atmospheric ‘underworld’ (Hoskins (1991)) where isentropic surfaces intersect the ground (as well as σ levels). As such, a final Fourier transform is performed in order to isolate the Fourier coefficients of the perturbation quantities in isentropic coordinates where they do not in-

intersect the lower boundary (the interior domain). The intersection and exterior domains are treated differently, as described in the following section.

As such, perturbations for each wavenumber are then defined by

$$q_{e_m}(\phi, \theta, t) = \hat{q}_m(\phi, \theta, t) - q_0 - [\hat{q}_m(\lambda, \phi, \theta, t)], \quad (2.43)$$

thus essentially removing the zonally symmetric part of the perturbation field and only retaining a single wavenumber at a time for the ENM calculation.

Defining Lower Boundary Pseudomomentum Contribution

Using equations (2.33) & 2.39 above to calculate \mathcal{P} & \mathcal{H} requires the analysis domain to be partitioned on each θ level into a number of parts.

- The interior domain, $\bar{\mathcal{D}}$, contains latitudes where both the background state and the full flow are above the lower boundary at all longitudes; the Rossby wave and gravity wave terms of 2.33 are calculated over this domain. The boundary integral is evaluated along its edge.
- The domain of intersection, $(\mathcal{D} \cap \mathcal{D}_0) \setminus \bar{\mathcal{D}}$, contains points where at certain longitudes, both the background state and full flow are above the lower boundary but excluded from the "interior" domain because the theta surface hits the lower boundary at some longitude.
- The background domain exterior to the intersection domain, $\mathcal{D}_0 \setminus (\mathcal{D} \cap \mathcal{D}_0)$, contains points which are above the lower boundary only in the background state. The exterior pseudomomentum term is calculated in this domain.
- The full flow domain exterior to the intersection domain, $\mathcal{D} \setminus (\mathcal{D} \cap \mathcal{D}_0)$ contains points above the lower boundary only in the full flow and not the background state. The exterior pseudomomentum is also calculated in this domain.

For finite amplitude disturbances, the pseudomomentum density is positive definite in both the interior and intersection domains, owing to the meridional PV gradient of the background state being positive. In the small amplitude limit, both exterior terms \mathcal{P}_e and the boundary term \mathcal{P}_b reduce to a single boundary term which is negative definite.

It is necessary to define points as being interior points, exterior points, or points without data for the ENM technique. In order to do this, the time-average of the back-

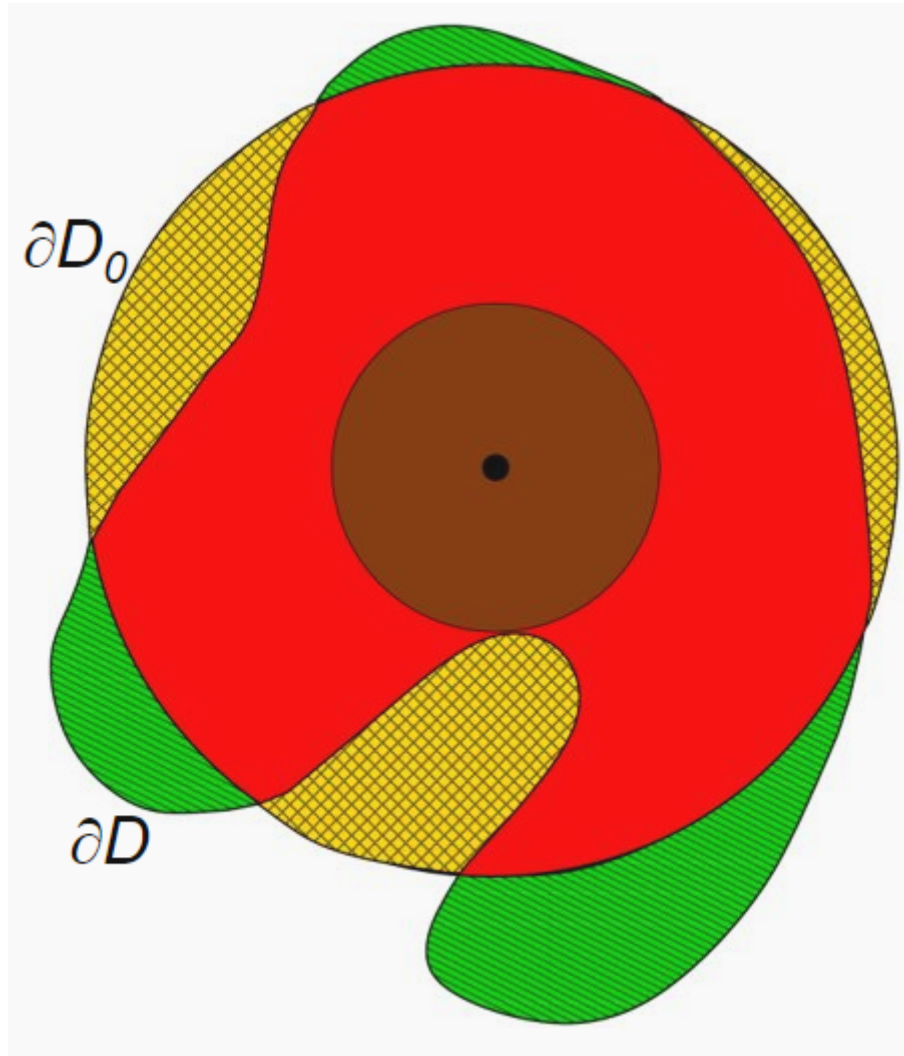


Figure 2.1: Schematic diagram of the division of the flow domain on a single isentropic surface following [Magnusdottir and Haynes \(1996\)](#). Here $\partial\mathcal{D}_0$ indicates the equivalent latitude within which the background state points are above the lower boundary at all longitudes; $\partial\mathcal{D}$ is the contour within which the full flow is above the lower boundary. Hence, the central brown area indicates the interior domain, the red indicates the intersection domain, the yellow hatching indicates points above-ground in the background flow only, and the green area outside the hatching indicates areas where the full flow only is above the lower boundary.

ground state is taken, and the ground is defined as the location of the lower boundary in this time-average state. This is because points must be assigned to either the interior *or* the exterior for the entirety of the timeseries; points cannot change designation for the eigenvector calculation used to find the ENMs. Furthermore, the conservation of pseudoenergy requires that the background state exhibits time-symmetry, i.e. that it is invariant. As the MLM is indeed slowly varying, a time-average stays near the MLM at each time point. The interior domain is defined by the brown region shown in [figure 2.1](#) for isentropic surfaces that intersect the ground. We will denote the pseudomomen-

tum density in this domain \mathcal{P}_w . If a location in the latitude, theta plane has non-zero exterior domain wave activity for greater than 20% of time points in the average, it is also assigned to the exterior mask. As these points do not change designation between boundary and interior throughout the time series, thresholding is required to determine which mask points are assigned to. In addition, all points below the lower boundary of the background state with exterior wave activity are also a part of the exterior domain.

Subsequently, at each latitude ϕ , the position of the first interior potential temperature level above the lower boundary of the time-averaged background state is $\theta_m(\phi)$. The intersection domain wave activity (term \mathcal{P}_d from equation 2.33) is integrated vertically, and assigned to the first θ -level below the ‘ground’ in the time-average background state, $\theta_{(m-1)}(\phi)$. This results in a ‘stripe’ of positive wave activity at the given θ -level, as shown in figure 2.2 (a). Similarly, the exterior pseudomomentum \mathcal{P}_e is integrated in the vertical and assigned to the θ -level two levels below the ground, $\theta_{(m-2)}(\phi)$, creating a stripe of negative wave activity below the positive stripe, shown in figure 2.2 (b).

The reason for the flattening of the intersection and exterior wave activity in this way is to ensure that the points maintain the correct sign throughout the ENM calculation (described below), and to allow the subtraction of the time-mean from the global wave activity used here.

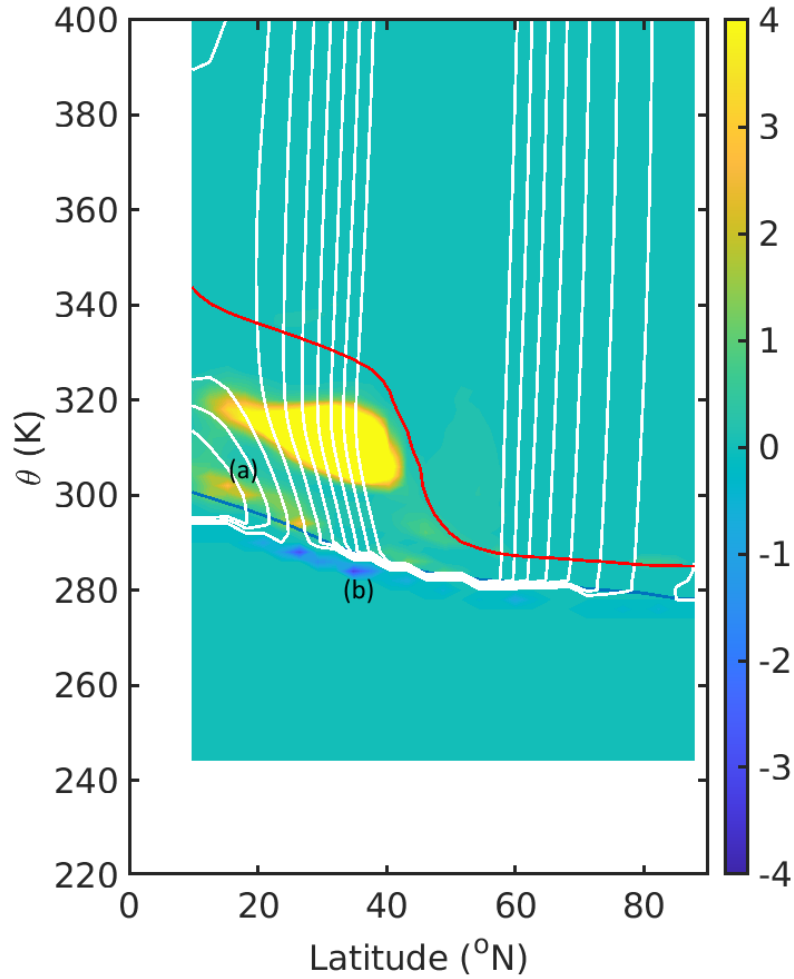


Figure 2.2: The spatial WA structure of the leading $m = 2$ ENM from the TO, as a demonstration of (a) the positive stripe of WA at $\theta_{(m-1)}$, and (b) the negative stripe at $\theta_{(m-2)}$

2.4.2 Eigenvalue Decomposition

Using the above calculation for pseudo-momentum, it is possible to construct a data matrix whose empirical-orthogonal functions (EOFs) are defined to be orthogonal with respect to pseudo-momentum. These EOFs define the Empirical Normal Modes for the data of the given wave activity and background state. To define the matrix, it is necessary to construct at each time point a wave disturbance whose inner-product is the pseudo-momentum. The amplitude coefficients for these waves come from the calculated pseudo-momentum density on a discrete (θ, μ) grid and the phase from the phase of the PV at each grid location.

To start, $P_{j,k,n}$ is defined to be the pseudo-momentum density of a grid-box of a discrete dataset defined on a (θ, μ) grid. Here j indexes θ , k indexes μ and n indexes

time. $P_{j,k,n}$ is positive on the interior and intersection points and negative on boundary points, and it is assumed that time variations in the location of the boundary have been accounted for as described in the previous section, such that the location of the different domains within the discrete grid is time invariant. The amplitude coefficient at each grid location is defined

$$\mathbf{x}_{j,k,n} = \sqrt{\Delta\mu_j \Delta\theta_k P_{j,k,n}} \quad (2.44)$$

Vectors will henceforth indexed across all of (θ, μ) space by s which runs from 1 to $N_j \times N_k$. Notably, the ENM calculation can be calculated regardless of the ordering of these vectors in space and time; as long as any transformation of spatial or temporal points is reverted after the calculation, the calculated results will be the same. These amplitudes and phases derived from the Fourier transformation are combined. Let ϵ_n be the vector whose elements are the phase of the perturbation PV at point s at each time point n for a single Fourier-filtered zonal wavenumber. We can then define a new vector \mathbf{z}_n and matrix \mathbf{Z} :

$$\mathbf{z}_n = \begin{bmatrix} \mathbf{x}_n \circ \cos(\epsilon_n) \\ \mathbf{x}_n \circ \sin(\epsilon_n) \end{bmatrix}, \quad (2.45)$$

and

$$\mathbf{Z} = \begin{bmatrix} | & | & & | \\ \mathbf{z}_1 & \mathbf{z}_2 & \dots & \mathbf{z}_{N_t} \\ | & | & & | \end{bmatrix} \quad (2.46)$$

Where \circ represents the Hadamard product. This definition of \mathbf{z}_n results in the pseudomomentum being given by $\mathcal{P} = \mathbf{z}_n^T \mathbf{z}_n$, where N_t is the number of time points. \mathbf{Z} can then be used for the computation of decomposition matrices;

$$\mathbf{P}_{time} = \frac{1}{N_t} \Re(\mathbf{Z})^T \Re(\mathbf{Z}) \text{ and } \mathbf{P}_{space} = \frac{1}{N_t} \Re(\mathbf{Z}) \Re(\mathbf{Z})^T \quad (2.47)$$

Where $\Re(\mathbf{Z})$ corresponds to the real values of \mathbf{Z} such that only the interior points are included. Either of these matrices can be eigen-decomposed to give a set of EOFs which are orthogonal with respect to $\mathbf{P}_{interior}$. We choose to calculate \mathbf{P}_{time} however a calculation of \mathbf{P}_{space} would yield the same result. The decomposition is performed by numerical decomposition software; specifically the eig function in MATLAB,

$$\mathbf{P}_{time} \mathbf{V}_t = \mathbf{V}_t \mathbf{\Lambda} \quad (2.48)$$

Where \mathbf{V}_t are the principal components, the matrix whose columns consist of the time eigenvectors of \mathbf{P}_{time} , and $\mathbf{\Lambda}$ is the matrix which contains the eigenvalues along the leading diagonal and 0 at all other points. Once the eigenvectors (and corresponding eigenvalues) have been calculated they are subsequently sorted in order of descending eigenvalue.

The temporal principal components can be used to calculate the spatial eigenvectors;

$$\mathbf{V}_x = \frac{1}{\sqrt{N_t}} \Re(\mathbf{Z}) \mathbf{V}_t \mathbf{\Lambda}^{-\frac{1}{2}} \quad (2.49)$$

These spatial eigenvectors are the structures of the ENMs, whereas the temporal eigenvectors are the principal component timeseries.

2.4.3 Calculation of \mathcal{P}_{ext} & \mathcal{H} by Projection

Once the ENMs have been calculated from the decomposition of interior (plus intersection domain) pseudomomentum, it is necessary to calculate the external contribution to pseudomomentum for each ENM through the projection of the eigenvalue matrix, \mathbf{V}_t onto $\mathbf{P}_{exterior}$;

$$\mathbf{V}_x^{exterior} = \frac{1}{\sqrt{N_t}} \Im(\mathbf{Z}) \mathbf{V}_t \mathbf{\Lambda}^{-\frac{1}{2}} \quad (2.50)$$

The total pseudomomentum density of each ENM may then be calculated through the sum of the $\mathbf{P}_{interior}$ and $\mathbf{P}_{exterior}$ ENM components.

The pseudoenergy of each ENM may be similarly calculated by projection of the variables needed to calculate each term of $\mathcal{H}(u, v, p)$ onto the principal components. Firstly, for u , we define a vector $\hat{\mathbf{u}}_n$ whose elements are given by;

$$\hat{u}_{s,n} = u_{s,n}^e \sqrt{\frac{1}{2} r_{0s,n} \Delta \mu_s \Delta \theta_s} \quad (2.51)$$

where u_n^e is the Fourier-filtered perturbation wind for a single zonal wavenumber, (obtained as for PV for equation 2.45). A vector similar to \mathbf{z}_n may be obtained for u ;

$$\tilde{\mathbf{u}}_n = \begin{bmatrix} \hat{\mathbf{u}}_n \cos(\epsilon_n^u) \\ \hat{\mathbf{u}}_n \sin(\epsilon_n^u) \end{bmatrix} \quad (2.52)$$

where ϵ_u represents the phase of the Fourier filtered component of zonal wind at μ & θ points. $\tilde{\mathbf{u}}_n$ is set to 0 at all exterior points, as the exterior does not contribute to the Kinetic Energy term of \mathcal{H} which is the term the zonal and meridional wind projections are used to calculate. A $\tilde{\mathbf{v}}_n$ term can be similarly calculated. Time averages of these terms are calculated and subtracted to yield;

$$\tilde{\mathbf{u}}'_n = \tilde{\mathbf{u}}_n - \langle \tilde{\mathbf{u}}_n \rangle \quad \& \quad \tilde{\mathbf{v}}'_n = \tilde{\mathbf{v}}_n - \langle \tilde{\mathbf{v}}_n \rangle \quad (2.53)$$

Similarly for p , using the Fourier filtered pressure perturbation \mathbf{p}_n^e , the vector $\hat{\mathbf{p}}_n$;

$$\hat{p}_{s,n} = p_{s,n}^e \sqrt{\frac{1}{2} \frac{R p_0^{(\kappa-1)}}{g p_{00} \kappa} \Delta \mu_s \Delta \theta_s} \quad (2.54)$$

Which can be used to calculate $\tilde{\mathbf{p}}_n$. $\tilde{\mathbf{u}}_n$, $\tilde{\mathbf{v}}_n$, & $\tilde{\mathbf{p}}_n$ can then each be projected along the principal components similar to equation 2.50. This projection yields the zonal wind u_i , meridional wind v_i , and pressure spatial structure p_i of each of the ENMs, with index i corresponding to the ENM mode number, which can be used to calculate the pseudo-energy density of each ENM;

$$H_{ii} = u_i^T u_i + v_i^T v_i + p_i^T p_i - \frac{u_0}{\cos \phi} P_{ii} \quad (2.55)$$

As only $\mathcal{P}_{interior}$ (the interior + intersection domain pseudo-momentum) is orthogonal as a result of the eigendecomposition, it possesses the orthogonality property; $\mathcal{P}_{interior}^{ij} = 0$ for $i \neq j$ (where j henceforth is another ENM mode index). It is expected that the full pseudomomentum and the pseudoenergy are also approximately orthogonal however this is not guaranteed, and the deviation from orthogonality of the ENMs can be measured with the quantity;

$$\frac{|\mathcal{P}_{ij}|}{\frac{1}{2} (\mathcal{P}_{ii} + \mathcal{P}_{jj})} \quad (2.56)$$

2.4.4 Pairing ENMs

In order for the ENMs to describe a propagating disturbance, it is required that they are paired into tuplets whose amplitudes vary in quadrature with one another over time.

If these pairs of ENMs are purely propagating, they possess the same spatial amplitude structure, with a phase offset by $\pi/2$. In order to find these propagating pairs, a 'matching statistic' for each possible mode pair is calculated.

The number of ENMs is equivalent to the number of time-points in the data, however most of these modes will possess very little pseudomomentum; as a result, a fractional proportion of all modes, σ_{trunc} are truncated from the spectrum before matching occurs, retaining only modes with eigenvalues λ_t which are above the threshold such that;

$$\frac{\lambda_{i_{filtered}}}{\lambda_1} > \sigma_{trunc} \quad (2.57)$$

Where λ_1 is the eigenvalue of the leading mode for that spectrum. Once truncated, the ENMs can be paired into quadrature matches, where the phases of the two constituent ENMs are $\pi/2$ out of phase. In order to find matches in quadrature, a matching statistic is defined. First, taking \mathbf{W} to be a transformation matrix designed to rotate through $\pi/2$;

$$\mathbf{W} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \quad (2.58)$$

it is possible to calculate a projection from one ENM onto a mode shifted by $\pi/2$;

$$\mathbf{M}_{\mathbf{q}} = \mathbf{V}_x^T (\mathbf{W} \mathbf{V}_x) \quad (2.59)$$

This effectively rotates each of the sin and cos components of z_n (equation 2.45) into one another, and as such the matrix $\mathbf{M}_{\mathbf{q}}$ is the matching statistic, and each element contains a measure of how well any two given modes match in quadrature. Its elements, $m_{q_{i,j}}$ are the individual match statistics of the ENMs. Similar to the amplitude truncation above, pairs below a certain matching threshold are discarded.

2.4.5 Dynamical Phase Speed C_{int}

The phase speed of a given ENM can be calculated (Brunet (1994)) from;

$$c_{int} = -\frac{\mathcal{H}}{\mathcal{P}} \quad (2.60)$$

Due to the conservation of the wave activity quantities, it is expected in general that a general atmospheric disturbance should appear steady in wave activity. It is possible

to calculate the speed of this reference frame (Held (1985)), giving equation 2.60.

Due to normalization of ENMs about interior pseudomomentum, \mathcal{P} is close to 1, however due to the contribution of $\mathcal{P}_{exterior}$, it is not exactly 1, and therefore must be included in the calculation of c_{int} ;

$$c_{int} = -\frac{(\mathcal{H}_w + \mathcal{H}_{KE} + \mathcal{H}_{APE} + \mathcal{H}_b + \mathcal{H}_e)}{\mathcal{P}} \quad (2.61)$$

With each term corresponding to a term of the the global integral of pseudoenergy for a given mode calculated by projection for each ENM. Values of c_{int} are averaged between members of a pair with equal weighting.

Chapter 3:

The Phase Speed Condition Test

3.1 Introduction

A fundamental property of dynamical normal modes is that they have an intrinsic frequency which describes how they oscillate, and a static function which describes their structure. For a wave disturbance, the time-dependence can be reflected in both the phase and the amplitude of the wave. ENMs share these modal properties even though the underpinning equations of motion, or the dynamical propagator matrix of the linearised dynamics, is not used to obtain the ENMs. As described in Chapter 2, they are obtained by an eigenvector calculation where the spatial components are the ENM structures and their time dependence is described by the principal components. Since the ENMs are eigenvectors of the pseudomomentum matrix, they are by construction orthogonal with respect to pseudomomentum. Since global pseudomomentum is a conserved property, Held (1985) showed that the analytical normal modes of the system must also be orthogonal with respect to pseudomomentum. Hence, in the case of linear dynamics, it is expected that the ENMs are a linear recombination of the analytical normal modes as they appear in the timeseries data. Furthermore, the intrinsic phase speed of each mode is given by the ratio of its pseudoenergy to pseudomomentum. Therefore, by analogy we can associate every ENM spatial structure with an intrinsic phase speed calculated from $-\mathcal{H}/\mathcal{P}$. Since each ENM also has a single zonal wavenumber, k , the phase speed yields immediately an intrinsic frequency, $\omega_{int} = kc_{int}$. This is a remarkable special property of ENMs which can be used to test the relevance of the structures, obtained simply by statistical means, as a compact description of the dynamics of flow evolution.

By finding ENMs oscillating in quadrature, and examining their principal component timeseries, it is possible also to extract, c_{emp} , an empirically measured speed of each

ENM pair, derived from tracking the phase of features across the ENM timeseries. For an atmosphere with wave-like disturbances described approximately by dynamics linearised about the background state, c_{int} and c_{emp} should match for each wave number. This is the ‘*Phase Speed Condition Test*’ and matches in phase speed indicate; (1) that the ENMs are behaving like normal modes; (2) the wave activities are approximately conserved; and (3) that large-scale behaviour is well described by linear wave dynamics, despite certain non-linear aspects to the flow.

This chapter presents a series of spectral general circulation model runs used to verify the implementation of the ENM technique for flows described by the primitive equations on the sphere. These model runs should approximately meet the requirements for the speeds to match- almost conservative flow, the absence of forcing, and a lack of nonlinear interaction between wavenumbers. As a result, performing the Phase Speed Condition Test will allow us to examine discrepancies, if there are any, between the dynamical speed and the empirical speed.

Over sufficiently long timescales, models with no seasonal cycle or other slow forcings will attain a quasi-equilibrium state, as long as other forcings, such as linear relaxation, or drag are also statistically steady. The background state of such an atmosphere could only be altered by wave-mean flow interaction and hence is expected to be approximately steady. For such states, it is expected that the ENM structures extracted from two distinct, ‘independent’ periods of the model run will exhibit the same characteristics, due to the absence of a process that might modify them. By comparing structures and speeds in this way, a test of structural consistency can be performed, called the ‘*Quasi-Equilibrium Structure test*’.

Once both the Phase Speed Condition Test, and the Quasi-Equilibrium Structure Test have been discussed, this chapter will go on to look at quantifications of the uncertainty of ENMs. Measurements of uncertainty will be discussed for both estimates of phase speed. Given the metrics of uncertainty, the effect of varying the sampling frequency on the uncertainty will be examined. Finally, a series of tests varying the analysis time-window will be presented, showing how the length of the timeseries affects the distinctness of the ENM spectra. This will serve as an underpinning for subsequent investigations with the ENM technique. It is expected that the ENMs obtained by the analysis will converge if the sampling frequency is much higher than the intrinsic frequency, and the time series length is longer than the intrinsic period of the ENMs.

3.2 Experimental Design

In order to test the phase speed condition, a set of idealized atmospheric flow data is required. There were several requirements for this data:

1. A set up where both seasonal and diurnal cycles could be removed was essential. This ensured that the model climate can enter a period of quasi-equilibrium where the background state evolves slowly in the absence of distinctive forcing, needed to perform the Quasi-Equilibrium Structure Test.
2. A dataset without orography was needed. As Earth's orography is not zonally symmetric, the presence of orography forces large-amplitude stationary waves [Dickinson \(1978\)](#) [Valdes and Hoskins \(1989\)](#). The phase speed condition test requires propagating waves without external forcing of wave activity, and as such it is convenient to remove orography for a clean experiment. This does not mean that orographic waves preclude the ENM technique being used. However, orography introduces additional complexity, such as how to define the background state's lower boundary. The choice of background state lower boundary is a decision that must be made when defining the background state (recalling that the background state is zonally symmetric). For example, the background surface geopotential can be defined as being uniform, or it could follow zonal mean orography; in either case extrapolation of data to the surface leaves a signal in θ which can modify wave-activity.
3. Waves propagating in a grid-based model are subject to numerical dispersion. Here it is preferable to utilise a spectral model so that large scale wave propagation is represented accurately. In general, spectral models better preserve low-order (and hence large-scale) structures [Simmons and Hoskins \(1975\)](#), making them favourable for our application. Ultimately, this ensures that transfer between waves can only happen as a result of non-linear interactions in the dynamics of the experiment.

3.2.1 Model

The model used was the Reading Intermediate General Circulation Model, IGCM2.2. It is a multi-level spectral transform primitive equation model on the sphere, based on the model equations from [Hoskins and Simmons \(1975\)](#), modified by [Blackburn \(1985\)](#). It includes both Newtonian cooling (as well as the potential for relaxation to variables other than temperature) and Rayleigh drag. The version used v(2.2) also possesses simple

moist parameterizations and a boundary layer scheme, neither of which were utilised in these experiments. For these experiments surface orography is zero.

3.2.2 Specification

The IGCM was specified to run at T42 resolution for the initial experiments presented in this chapter, retaining the first 42 zonal and total wavenumbers using jagged triangular truncation. The number of retained spectral coefficients at this truncation is 462; this is equivalent to about 2.79° grid-spacing at the equator on a quadratic Gaussian grid used for the spectral transforms. Following [Thorncroft et al. \(1993\)](#) there were 30 vertical levels on a stretched grid, with higher resolution in the upper troposphere, on sigma surfaces. The pressure of the uppermost sigma level is 8.4hPa, ensuring the model domain extends well into the stratosphere. The model was run with 128 semi-implicit time steps per day (roughly one time step every 11 minutes) with output fields recorded at 6 hourly intervals, and global norm data (such as global MSLP and energy) recorded at each time step. This resolution was chosen as high spatial resolution is not required to examine large-scale Rossby-wave structures, and in order to reduce computation time. High temporal resolution is required to maintain a Courant number well below 1, due to the high propagation speeds in some of the model experiments.

| Model | Horizontal Res. | Vertical Levels | Δt | Domain | τ -Relaxation | τ -Drag |
|---------|-----------------|-----------------|-------------|-------------|--------------------|--------------|
| IGCM2.2 | T42 | 30 | 1 Day / 128 | Hemispheric | 15 day | 1 day |

Table 3.1: *Summary model specification for the iGCM runs.*

A weak sixth order hyper-diffusion was applied, in order to smooth out fine structures generated by enstrophy cascades from large scale motion, with a timescale of 6 hours on the smallest retained wavenumber.

3.2.3 Relaxation

A Newtonian relaxation scheme was used to simulate radiative forcing. The relaxation timescale was 15 days. The structure of the relaxation state is an analytically prescribed, zonally symmetric function of temperature from [Held and Suarez \(1994a\)](#):

$$T_{eq} = \max\left(200K, \left[315K - (\Delta T)_y \sin^2 \phi - (\Delta \theta)_z \log\left(\frac{p}{p_0}\right) \cos^2 \phi\right] \left(\frac{p}{p_0}\right)^\kappa\right), \quad (3.1)$$

where $(\Delta T)_y = 60\text{K}$ is the temperature difference between the equator and pole; $(\Delta\theta)_z = 10\text{K}$ is the dry adiabatic lapse rate; and $p_0 = 1000\text{hPa}$ is the surface pressure. The Held-Suarez state was used (1) due to its ubiquity in the literature, and (2) due to its original purpose of comparing dynamical cores. Figure 3.1 shows the temperature field described by equation 3.1, and the zonal flow in gradient thermal wind balance with the temperature. The IGCM is used to obtain this balanced flow using the in-built balancing routine.

As the Held-Suarez state is analytically defined it is suitable for inter-comparison between models which should converge to a uniform solution. Were waves entirely conservative, the background state would be identical to the Held-Suarez state; instead the background calculated is a balance between the relaxation and nonlinear wave breaking and associated stretching, folding and dissipation of PV filaments. As studies have previously been done with this set-up, any previously run experiments, when analysed by the ENM technique, should provide the same set of modes, ensuring that the experiment is easily repeatable.




Figure 3.1: The specified temperature field used for the radiative forcing as specified in Held and Suarez (1994a) (a), with its corresponding balanced zonal wind field (b).

The relaxation could be set to relax either to this temperature field, or to *both* this temperature field and the balanced vorticity and horizontal divergence fields associated with it. This relaxation to both zonal temperature and wind fields serves to provide an additional set of experiments with distinct dynamics, where the atmosphere is forced towards a state in gradient wind balance. An exact balanced state is obtained using the IGCM discretisation by running it for one time step and utilising the balancing scheme BALANC, as described in appendix II of Simmons and Hoskins (1975) to balance from the prescribed temperature field and a uniform prescribed surface pressure field of 1000hPa

(whereas in [Simmons and Hoskins \(1975\)](#) a balance to T and p_s is achieved from a prescribed vorticity as opposed to the other way around). It should be noted that this wind field, in particular the large equatorial jet, is *not* observed within the the zonal mean once spin up is complete, as baroclinic instability establishes a mid-latitude jet.

3.2.4 Drag

Rayleigh drag, as described in [Held and Suarez \(1994a\)](#), was also implemented. The drag possesses a one day timescale and has only a vertical dependence, and its co-efficient is defined as:

$$k_v = k_f \max\left(0, \frac{\sigma - \sigma_b}{1 - \sigma_b}\right), \quad (3.2)$$

where $\sigma_b = 0.7$ is the height of the boundary layer in sigma co-ordinates, and $k_f = 1$ is the relaxation rate in days⁻¹. It is applied to all spectral co-efficients of vorticity except the first, as the first spectral co-efficient of the vorticity equals the planetary vorticity. At the surface, the drag possesses its full strength, however it decreases linearly to zero at $\sigma = 0.7$.

3.2.5 Initial Conditions

The balanced relaxation state was used as the initial condition of the simulations, with a surface pressure of 1000hPa. A random noise component was applied to all surface pressure spectral co-efficients, has a 0.1hPa amplitude and is set using a pre-determined random seed (meaning that it is the same for all runs) in order to ensure non-steady flow. See [Chapter 4](#) for a more detailed discussion of this perturbation.

3.2.6 Experiment List

Four experiments were performed using the aforementioned setup. This was in order to test the robustness of the ENM technique with a variety of model atmospheres, including extreme cases, and to ensure that results were consistent across multiple experiments. The four experiments were: ‘T-only’, with no drag implemented, and relaxing linearly to *only* the HS temperature field [3.1](#); ‘Relaxall’ with no drag implemented, and relaxing linearly to both the HS temperature field and the vorticity and divergence of its balanced wind field; ‘T-only-Drag’ with Rayleigh drag implemented and relaxing only to the HS temperature field; and finally ‘Relaxall-Drag’ with Rayleigh drag implemented

and relaxing to both the HS temperature field and its balanced vorticity and divergence fields. The experiment details are summarized in Table 3.2.

| Run Name | Relaxation | Surface Drag |
|----------------------|--------------------|--------------|
| T-only (TO) | Temperature Only | OFF |
| T-only-Drag (TORD) | Temperature Only | ON |
| Relaxall (RA) | Temperature & Wind | OFF |
| Relaxall-Drag (RARD) | Temperature & Wind | ON |

Table 3.2: *Experiment specification summary for IGCN runs.*

Relaxation to wind is usually only performed to relax models to observed or analysed states Hoskins et al. (2012) though it was seen as useful in this scenario. As the relaxation state is a solution of the primitive equations, and serves to drive the model towards gradient wind balance. This is in contrast to just the temperature field which is, on its own, not a solution of the primitive equations and would drive perturbations away from the balanced state. Temperature is often used on its own as an emulation of radiative relaxation.

Drag was introduced in order to smooth out the large pressure gradient in the T-only experiment (detailed in 3.2.7) and make the atmosphere more ‘Earthlike’. These four experiments produced distinct dynamics that allow a thorough testing of the phase speed condition in multiple contrasting scenarios.

3.2.7 Experiment Overview

All four experiments display large amplitude, long timescale oscillations, shown by the root-mean-square vorticity against time plot in figure 3.2. The overall stability of the state of the model atmosphere can be evaluated by energy norms over the entire domain. The T-only and Relaxall experiments achieve a state of equilibrium around day 200, while the T-only Drag and Relaxall Drag experiments achieve equilibrium later, between days 300 and 400. Indeed there is a slight and gradual decrease in total static energy in the experiments with drag; this is due to the imposed drag not conserving energy. The global average surface pressure (not shown) also reduces with time across all experiments, as mass continuously leaks from the experiments; for example the TO experiment loses 0.02% of its mass over the course of the 1095 day run. This could be due to the relaxation scheme not conserving mass, or mass loss due to the timestepping scheme; in either case it was considered small enough to be negligible given the length of the analysis window.

The average energy is higher after the spin up than at the model’s initial time in the

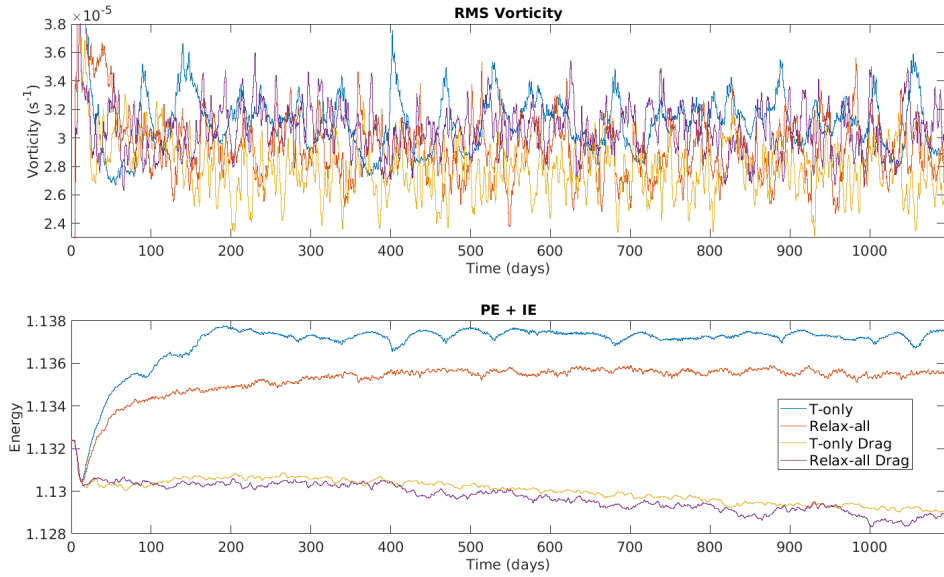


Figure 3.2: *Global integral quantities for each experiment. The top plot shows RMS vorticity in SI units, whereas the bottom gives globally integrated potential (gravitational) and internal (thermal) energy in dimensionless units as given by the IGCM. The RMS vorticity shows that each experiment undergoes periodic oscillations, the T-only experiment most clearly. After an initial transient period, the energy of the four experiments can be seen to level out, indicating they have reached a stable equilibrium.*

experiments without drag, due to relaxation adding energy to the model system. Both the energy and vorticity of the T-only experiment display quite distinct oscillations with similar periods. A simple linear dependence calculation shows that they possess a significant anti-correlation, with a simple linear regression providing a correlation coefficient of -0.58 .

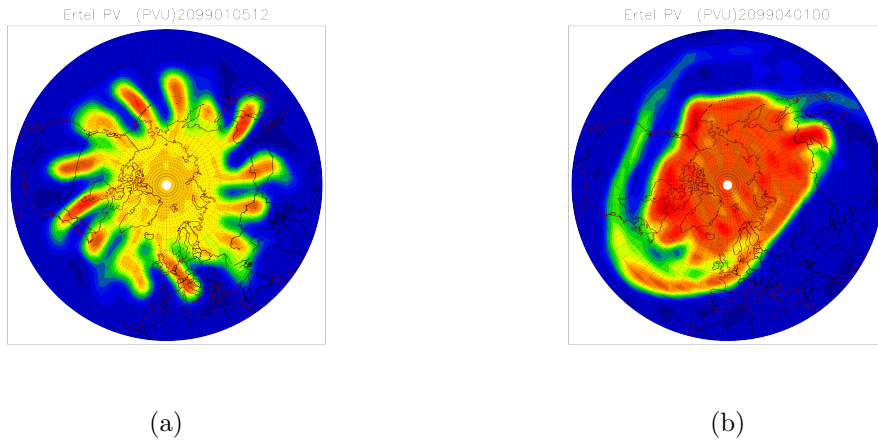


Figure 3.3: *PV on $\sigma = 0.25$ from the T-Only experiment at (a) model day 5.5 and (b) model day 90.*

Figure 3.3 (a) shows an initial baroclinic wave lifecycle with wavenumber 14 in the T-

only experiment as the initial balanced flow breaks down due to the instability introduced by the noise routine. Without this noise, the flow remains precisely zonally symmetric and steady. All experiments exhibited a similar baroclinic lifecycle, with wavenumber between 13 and 16 (inclusive). Following this lifecycle the experiments settle into a large-scale pattern of Rossby eddies, and a strong polar vortex is established. Looking at Hovmöller plots of meridional velocity mid-latitudes (figure 3.4) reveals that each experiment possesses distinct dynamics. Both the T-only and Relaxall experiments possess very distinct propagating structures. Both the T-only and Relaxall experiments show large, consistent westwards propagating modes, though these are moderated in the T-only experiment by a consistent large scale phase shift, which is absent in the Relaxall experiment. The T-only Drag experiment exhibits eastwards propagating structures with varying speeds in figure 3.4. Discontinuities in these structures propagating show their peak varying in latitude (and thus not consistently intersecting the latitude line of 60° the plot is drawn across) as they propagate longitudinally. Finally, the Relaxall Drag experiment has structures which appear almost stationary at times. Examining PV fields at 310K reveals that they correspond to positive PV structures which are deposited through filamentation in the mid-latitudes from the polar vortex where they remain, stationary in longitude, until they are subsumed by a passing wave. These ‘obvious’ differences are useful in that structures extracted by the ENM technique, and notably their phase speeds, can be compared to the structures visible in the full flow. The phase speed test is expected to work well in this situation.

The modes in the TO & RA experiments are nonetheless particularly un-Earthlike. The dominant structures are large-scale barotropic modes with very low wavenumber, reminiscent of structures found in Mars-based experiments [Michelangeli et al. \(1987\)](#). The lack of orography or any drag at the surface results in a particularly strong jet, stronger than 120ms^{-1} in some cases, which extends to the surface. The resultant pressure gradient can be seen on model day 90 (the same as in figure 3.3(b)) in figure 3.5 (a). By introducing drag, this pressure gradient is greatly reduced (Figure 3.5 (b)), and baroclinic waves dominate in the flow.

The presence of drag also serves to weaken the jet significantly. In the T-Only experiment jet peak speeds exceed 120ms^{-1} for the entire timeseries (after spin-up). This is reduced by a factor of 2 when drag is introduced to the experiments.

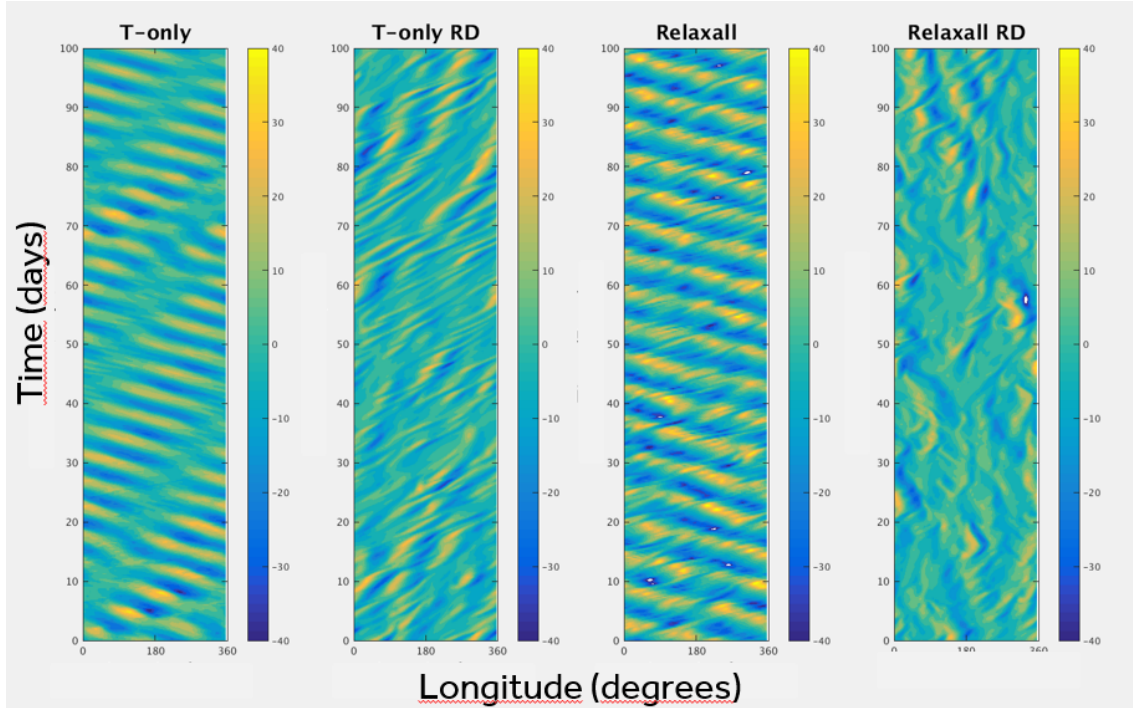


Figure 3.4: Meridional wind speed across longitudes against time at $\sigma = 0.25$, with a latitude of $60N$, for experiments *T-only*, *T-only RD*, *Relaxall*, and *Relaxall RD*, from model day 400 to 500.

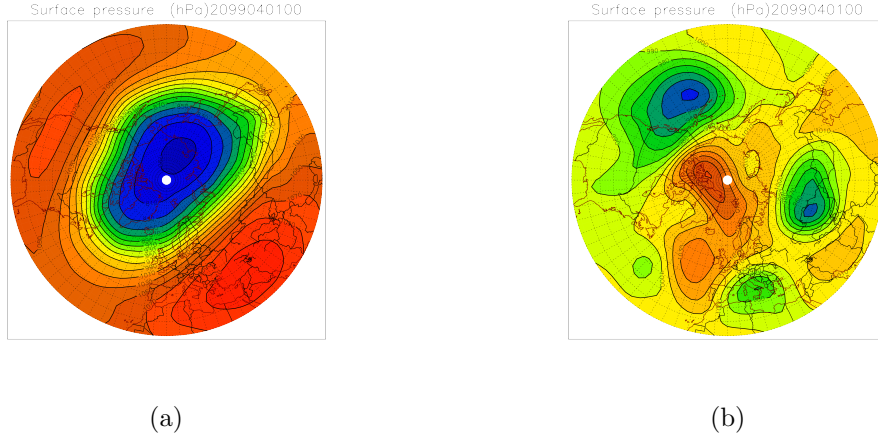


Figure 3.5: Example surface pressure field for (a) the *T-Only* experiment and (b) the *T-Only Drag* experiment. Note the different contour intervals ($20hPa$ for (a) and $10hPa$ for (b)).

3.3 Evolving Background States

As per the process detailed in chapter 2, the background state was computed for each experiment. Initially, a sampling interval of 6 hours was taken over a total sample period of 100 days, from model day 400 inclusive (timestamp 2100020500) to model day 500 noninclusive (timestamp 2100051518). This was considered to be a sufficiently long window for an initial examination of the structures seen, as it is much longer than any of the

timescales for the experiments' physical processes. Nonetheless, when results had been calculated for all four experiments, it was decided that the time-window for the experiments with Rayleigh drag applied would be increased to model day 1095 (2101123118) for a total of 695 days of analysis. This was done to ensure the longest possible timeseries for greatest accuracy when calculating c_{emp} , and in order to subdivide the analysis window for uncertainty calculations.

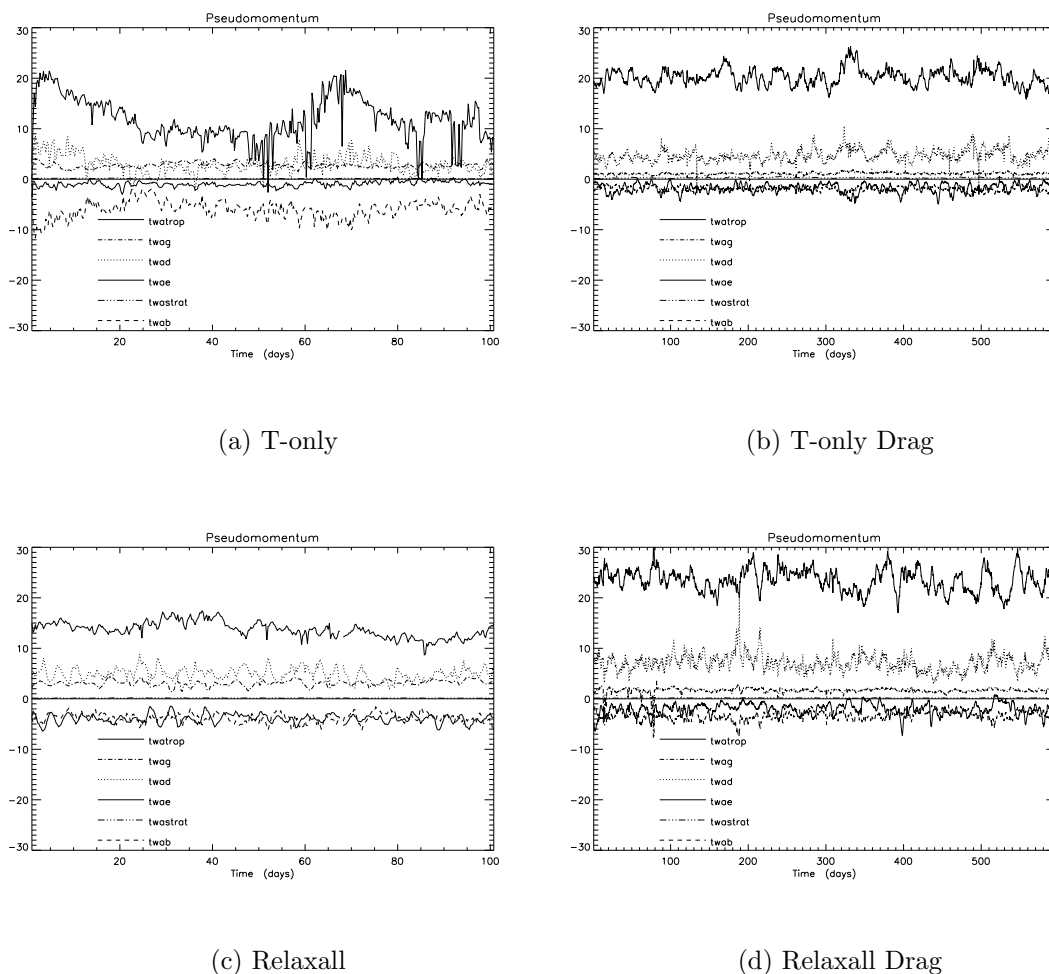


Figure 3.6: Global wave-activity for each experiment. The thick line represents the tropospheric contribution of the interior Rossby-wave WA term of \mathcal{P}_w ($twatrop$), the thin line with long dashes represents total stratospheric WA ($twastrat$) of the same; all other lines denote various contributions to the total WA. The thick dashed line represents the gravity wave contribution \mathcal{P}_g ($twag$), the thin dashed line represents the background state's intersection domain WA \mathcal{P}_d ($twad$), the thin solid line represents the WA in the exterior domain \mathcal{P}_e ($twae$), and the thin dashed line with wide spacing represents the boundary contribution from between the interior domain and intersection domain \mathcal{P}_b ($twab$). Note that the x-axis for experiments without drag applied is only 100 days, whereas it is 500 days for experiments with drag.

Examining global pseudomomentum for the four experiments (figure 3.6) highlights

some of the differences in their dynamics. The interior Rossby Wave term dominates all experiments, however the experiments with drag show a much higher relative contribution to WA from this term than those without. The positive gravity wave and intersection domain terms are approximately balanced by negative interior boundary and exterior domain terms, though the positive intersection domain term dominates somewhat in the Drag experiments.

The interior Rossby wave term of the T-only experiment shows a peak in WA at the beginning of the timeseries, with a decaying amplitude up until day 70, followed by another peak. This trend is not found in the other experiments' background states which exhibit fairly consistent Rossby-wave WA. It is possible that this is a periodic structure, however a longer analysis timeseries would be required in order to examine this fully; the wave activity for the TO and RA experiments was only calculated for the initial 100 day analysis window of model days 400-500.

Examining θ on the lower boundary for the TO experiment, it can be seen that there is a sharp θ gradient on the edge of the sub-tropics (figure 3.7) in-line with the position of the jet. This is unsurprising, as the peak of the jet (b) extends all the way to the surface, effectively creating a partition. This can be seen in the wave activity (figure 3.6(a)) as a large boundary WA, corresponding to a structure of high WA to the north of the jet (not shown).

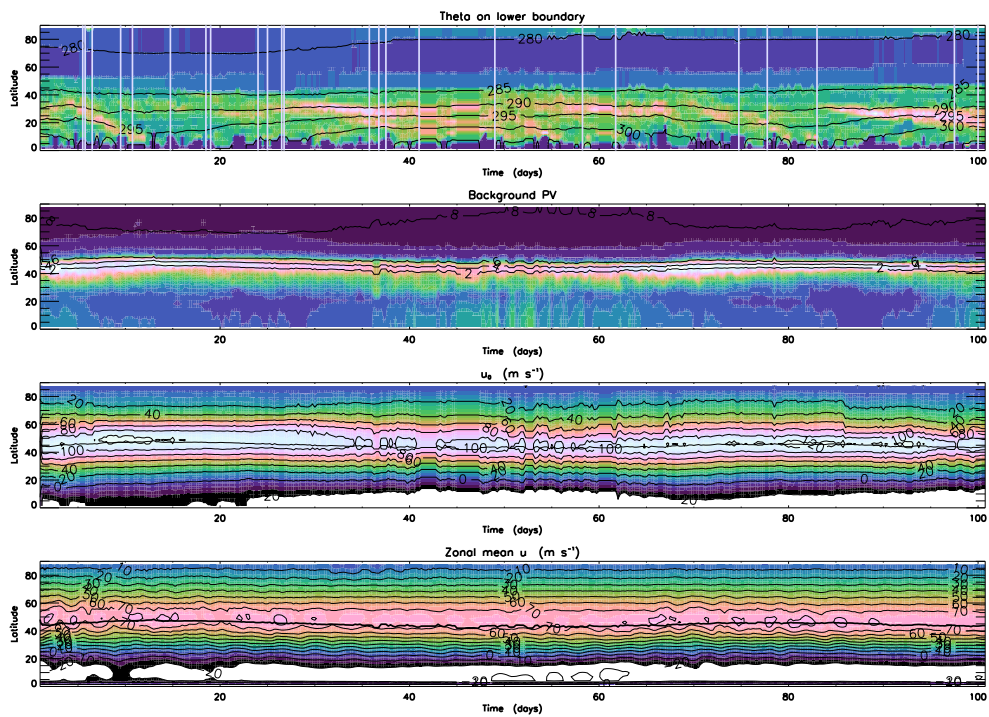


Figure 3.7: Timeseries of θ on the lower boundary, background PV in PVU, background wind, and zonal mean wind (all on the 315K θ surface) at each latitude for the TO experiment for the 100 day analysis window (model days 400 to 500). In the surface θ and background PV plots, the colour represents the meridional gradient of these quantities. Purple vertical lines indicate days where the PV inverter failed to converge and states were replaced by the previous timestep.

2100022500

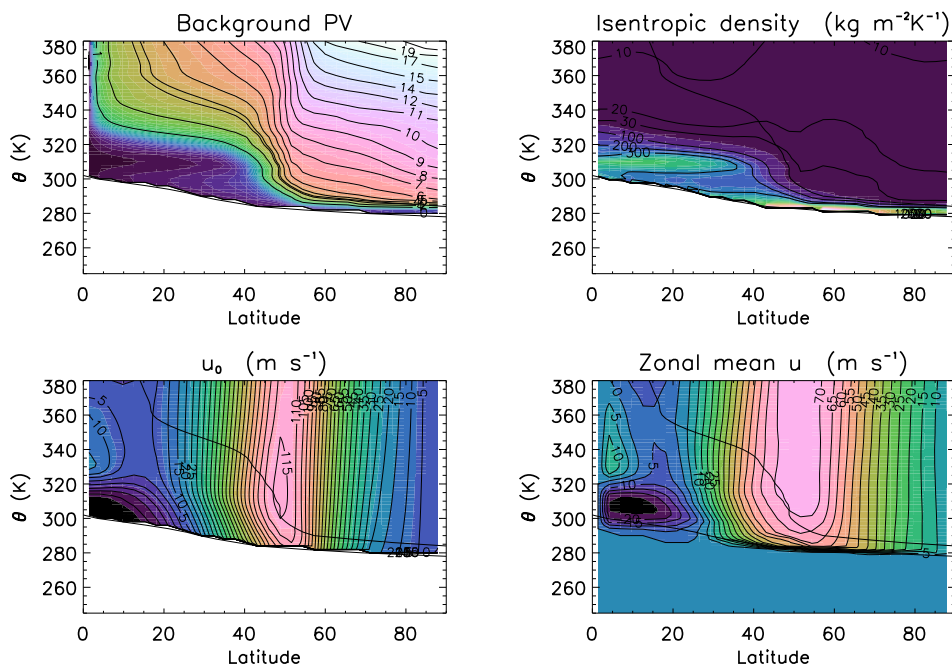


Figure 3.8: Snapshots of background PV, isentropic density \mathcal{E} and wind, and zonal mean zonal wind on midnight of day 20 of the above timewindow (model day 420).

A pattern seen in all experiments is a ‘tongue’ of surface θ gradient, beginning at the latitude of the jet and extending equatorward. The cross-sections shown in figure 3.8 show the background wind and PV for an example timestep where this pattern occurs. Of note is the presence of the jet on the poleward flank of the mid-latitude PV gradient.

The MLM state possesses much higher wind speeds than the zonal mean wind (figure 3.8) in the TO (and RA) experiment, with the jets extending to the lower boundary, due to the large barotropic vortex in these experiments. The BG state is steady in the TO and RD experiments, however even the zonal mean wind is consistent in these experiments.

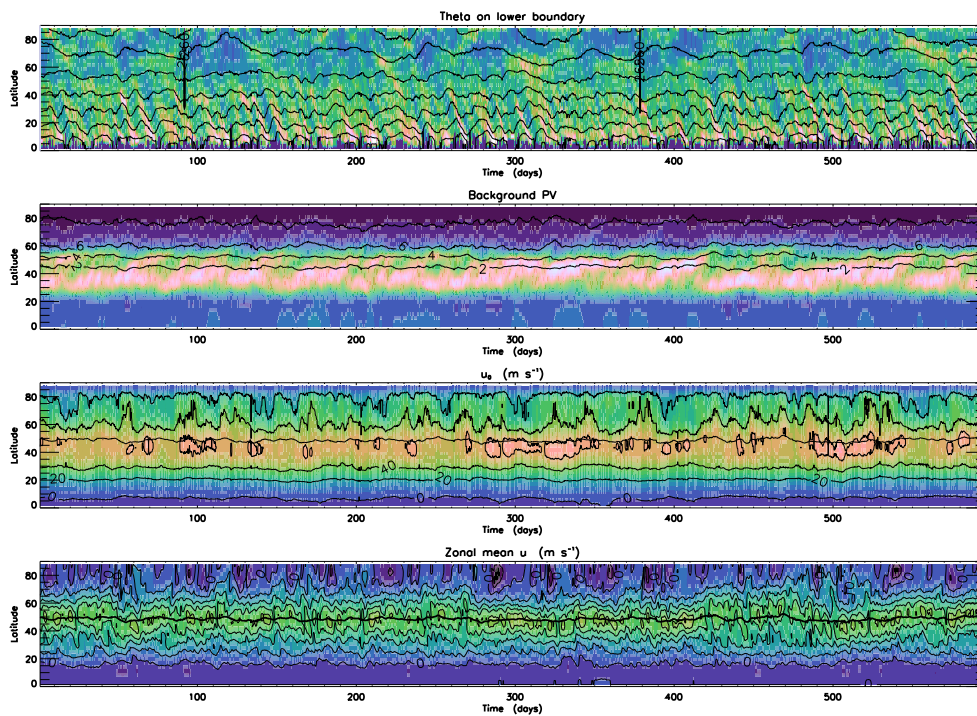


Figure 3.9: Timeseries of θ on the lower boundary, background PV in PVU, background wind, and zonal mean wind (all on the 315K θ surface) at each latitude for the TORD experiment for the 100 day analysis window (model days 400 to 500). In the surface θ and background PV plots, the colour represents the meridional gradient of these quantities. Purple vertical lines indicate days where the PV inverter failed to converge and states were replaced by the previous timestep.

2100022500

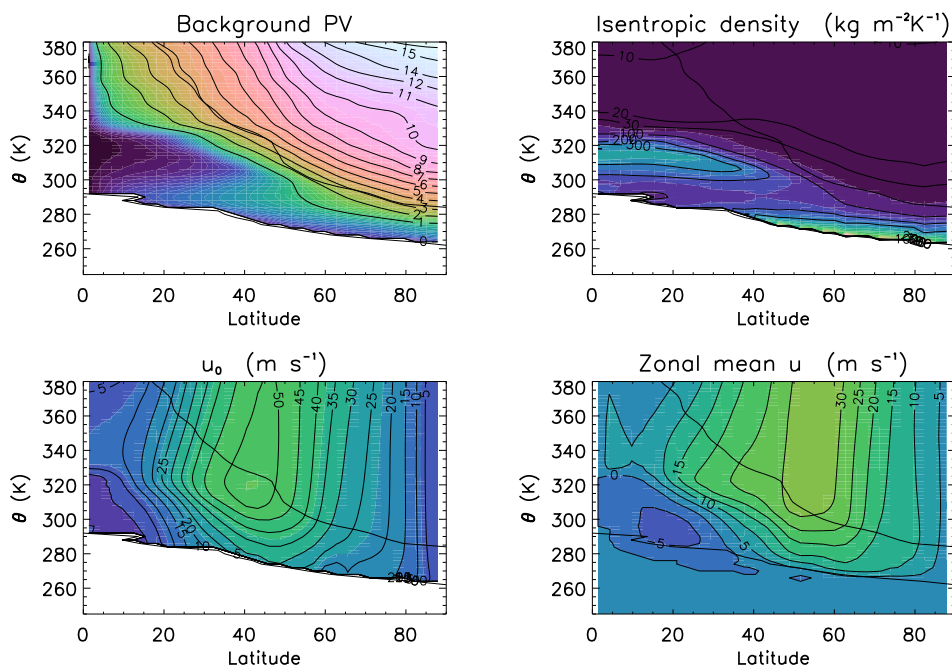


Figure 3.10: Snapshots of background PV, isentropic density \mathcal{E} wind, and zonal mean zonal wind on midnight of day 20 of the above timewindow (model day 420).

Due to the reduction in zonal wind (figure 3.10(b)) and the strength of the surface wind speed by the introduction of Rayleigh drag, the Rayleigh drag experiments’ background states correspondingly have a much weaker jet, with peak wind values closer to those found in Earth’s atmosphere. The Eulerian zonal mean jet in the TORD (and RARD) possesses a great deal more variability than in the experiments without drag, while the MLM jet changes very little with time, both in strength and variability.

The “tongues” of surface θ gradient in the drag experiments are more frequent and move faster towards the equator (though the number of inverter failure dates remains consistent). As shown in figure 3.10, the gradient of the tropopause is less in these experiments, corresponding to a weaker jet. Interestingly there is a significant variation in the latitude of the MLM jet and the Eulerian zonal mean jet, due to the presence of baroclinic wave activity.

3.4 Empirical Normal Modes

The ENM calculation was performed on the four experiments. Similar to the background state, the calculation was performed with 6 hourly resolution, over the aforemen-

tioned 100 day window. Upon completion of the initial exploratory analysis a further 595 days were calculated for the TORD and RARD experiments. ENMs were calculated for wavenumbers 1 to 12, with higher wavenumbers omitted due to low amplitude.

Two measurements were used in order to assess the relative importance of an ENM pair within a given wavenumber's spectrum; the first is prominence, how much more power is contained in the leading pair than the subsequent pair, and is defined;

$$\text{prominence} = \frac{\sum \lambda_{\text{leading pair}} - \sum \lambda_{\text{next pair}}}{\sum \lambda_{\text{leading pair}}} \quad (3.3)$$

The next is proportion, which is the amount of power in the leading pair relative to the entire ENM series for that wavenumber.

$$\text{proportion} = \frac{\sum \lambda_{\text{leading pair}}}{\sum \lambda_{\text{whole spectrum}}} \quad (3.4)$$

3.4.1 Experiments without Drag

The first two experiments shown here, T-only and Relaxall, display large, planetary-scale disturbances which do not reflect Earthlike-conditions, as discussed above. It was expected that the Empirical Normal modes would be dominated by modes with low zonal wavenumber and high amplitude, and this is largely borne out by analysis of the ENMs.

T-Only

The wave activity power spectrum, shown in figure 3.11(a) for the T-only experiment shows the dominance (i.e. relative power in pseudomomentum) of both the $m = 1$ and the $m = 2$ wave mode in pseudomomentum, reflecting the large wave structure seen in the full flow. Wave number three also has significant wave activity, but after this the power contained in structures with wavenumbers four and above diminishes.

The relative dominance of $m = 2$'s leading pair is confirmed by its high prominence and proportion. More than 80% of the total ENM spectrum for wavenumber 2 is dominated by its leading pair (figure 3.11(b)), and it is shown to be dominant even relative to the next pair in the eigenvalue series. $m = 1$ does not show quite this same dominance, however $m = 3$, despite its lower power also shows high prominence in the leading pair. As expected, higher wavenumbers show much less coherence, with the power in the spectrum spread out much more among the ENMs.

Examining the structure of the $m=2$ mode from the TO experiment (figure 3.12),

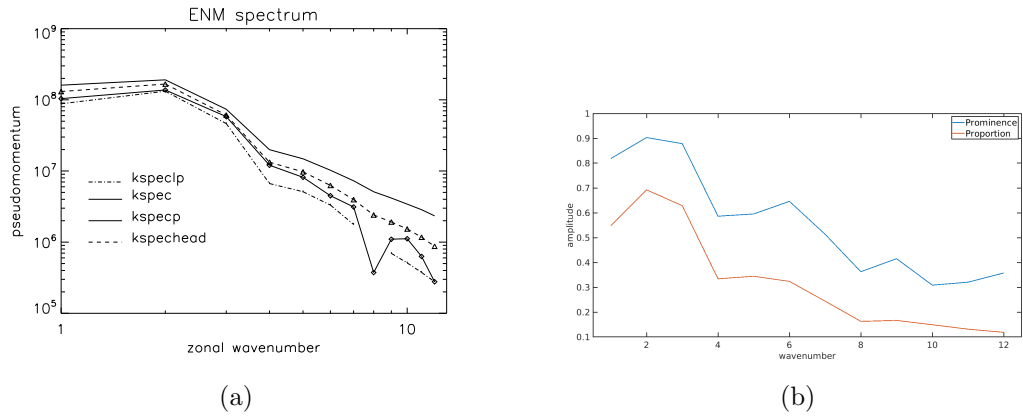


Figure 3.11: (a) The Pseudomomentum spectrum for the *T*-only experiment, and (b) Measures of prominence and proportion of spectrum for each wavenumber for the *T*-only experiment. In (a) the thin line is the total spectrum (*kspec*), the dash-dot line (*kspeclp*) contains only the leading pair of each wavenumber, and the thick line is only the paired eigenvalues (*kspecp*). The final dashed line (*kspehead*) shows the power contained in the first 10 ENMs (without pairing). By design, the total spectrum (*kspec*) is highest in pseudomomentum, as each of the other quantities is a subset of this quantity.

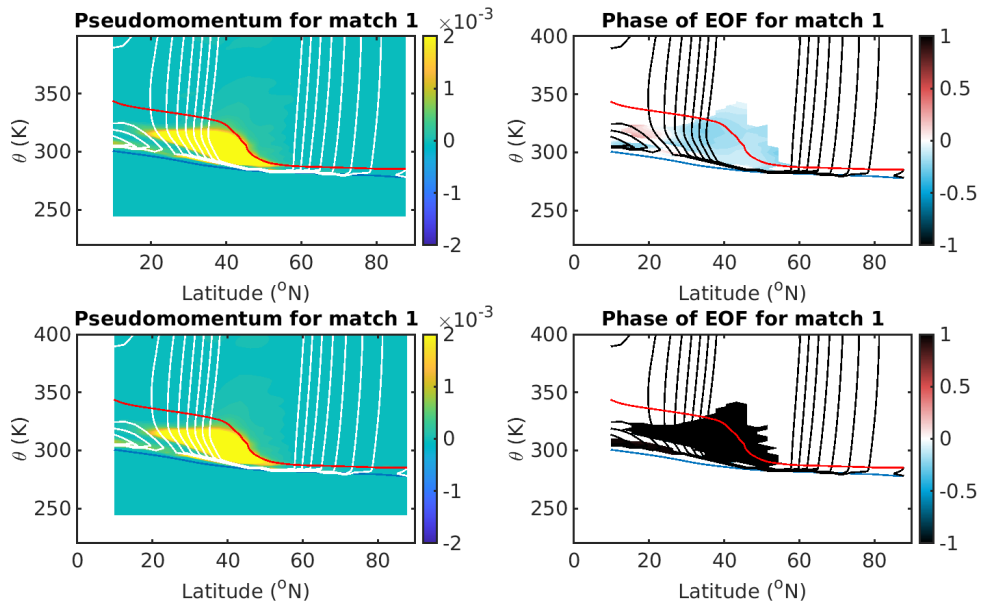


Figure 3.12: Amplitude (left hand panels) and phase (right hand panels) for the two leading modes of filtered wavenumber 2 for the *T*-only experiment. The phases are shown on the right of each subplot. The 2PVU line is shown in red and background wind is overplotted in white with a contour interval of 10m s^{-1} . The blue line shows the surface.

from the phase difference between the pair members, the leading pair are in quadrature and well-matched. The wave activity lies on the equatorward flank of the jet, and extends deeply through the troposphere from the surface to the tropopause as it increases through the jet, and tails off in the tropics. The continuity of this mode to the surface results from the lack of surface drag in the experiment; as the jet is able to extend fully downwards

to the surface, the wave activity structure is able to exist throughout the troposphere.

Relaxall

Relaxall has a highly dominant $m = 1$ mode, with the wave activity of the first wave number almost an order of magnitude larger than subsequent wavenumbers. The power (shown in figure 3.13(a)) contained in the remainder of the spectrum is lower, falling off smoothly with increased wavenumber. This step from $m = 1$ matches expectations, given the large structure seen in the Hovmöller plots.

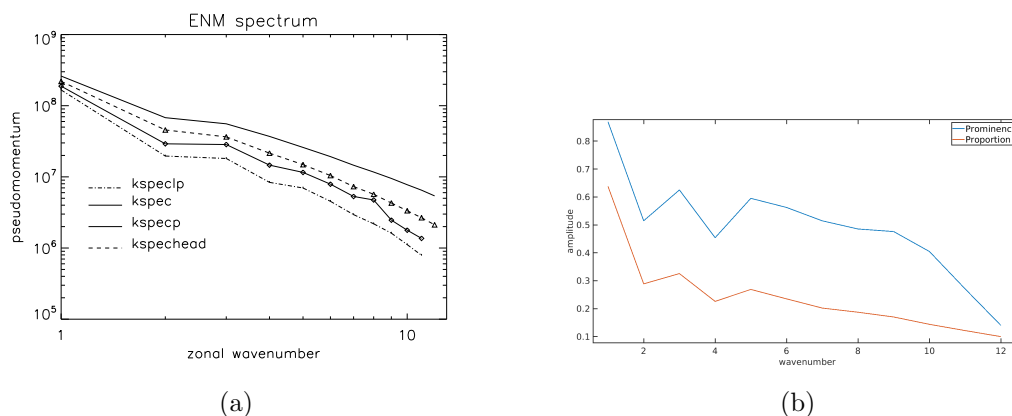


Figure 3.13: (a) The Pseudomomentum spectrum for the Relaxall experiment, and (b) Measures of prominence and proportion of spectrum for each wavenumber for the Relaxall experiment. In (a) the thin line is the total spectrum (*kspec*), the dash-dot line (*kspeclp*) contains only the leading pair of each wavenumber, and the thick line is only the paired eigenvalues (*kspecp*). The final dashed line (*kspehead*) shows the power contained in the first 10 ENMs (without pairing). By design, the total spectrum (*kspec*) is highest in pseudomomentum, as each of the other quantities is a subset of this quantity.

Examining the prominence and proportion (figure 3.13(b)), the vast majority of the power of the $m = 1$ wavenumber is contained within the first pair of the ENM spectrum, while the remainder of wavenumbers show power spread out amongst the spectrum of each wavenumber.

3.4.2 Experiments with Rayleigh Drag

The following experiments, T-only RD and Relaxall RD, were run with a Rayleigh Drag forcing added, and as a result, the large barotropic-like modes seen above were not present in these experiments. This results in conditions more similar to those seen in the literature (Zappa et al. (2010), H. Hendon and Hartmann (1985)), and a mode structure less dominant at planetary scales with more baroclinic-scale structures visible in the power spectrum.

T-only with Rayleigh Drag

There is no clear dominance of any wavenumbers in the TORD power spectrum, however for zonal wavenumbers higher than $m = 5$, the power contained in the whole ENM spectrum gradually falls (figure 3.14(a)). Of note in particular is wavenumber 3, the leading pair of which captures a much smaller amount of the total power than the rest of the large-scale spectrum due to its low proportion as shown in figure 3.14(b). While total power does fall at baroclinic scales above $m \sim 5$, there is a maximum in power within paired ENMs (kspecp of figure 3.14(a)) here, both as an absolute value (at $m = 5$) and as a fraction of total power (for most higher m) implying that at these scales, the ENM technique picks out a greater number of coherent propagating structures, due to a greater amount of power being described by paired structures as opposed to in the unsorted modes. Exceptions at $m = 8$ & $m = 11$ are due to restrictions on match criteria, as only the leading pair in these cases matches sufficiently well in phase to be considered.

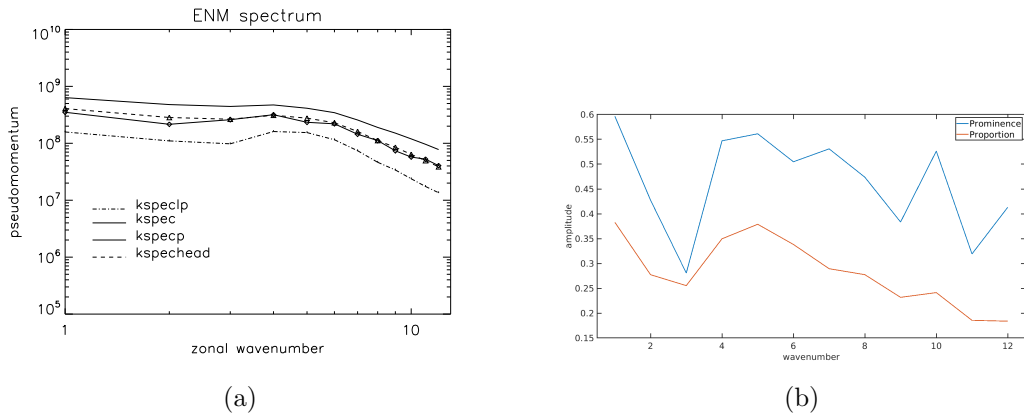


Figure 3.14: (a) The Pseudomomentum spectrum for the Tonly Rayleigh Drag experiment, and (b) Measures of prominence and proportion of spectrum for each wavenumber for the TORD experiment. In (a) the thin line is the total spectrum (kspec), the dash-dot line (kspeclp) contains only the leading pair of each wavenumber, and the thick line is only the paired eigenvalues (kspecp). The final dashed line (kspechead) shows the power contained in the first 10 ENMs (without pairing). By design, the total spectrum (kspec) is highest in pseudomomentum, as each of the other quantities is a subset of this quantity.

In addition to the maximum in power for paired ENMs at $m = 5$, the prominence (figure 3.14) for wavenumbers $m = 5$ & $m = 6$ and proportion of their respective spectra are highest for all wavenumbers, and indeed the power for both is comparable to that of waves at planetary scales. The amplitude structure of the $m = 5$ mode in the TORD experiment is shown in figure 3.15. The mode here possesses a different vertical structure to the TO experiment, with wave activity straddling the jet with a change in phase across it, indicating a varicose mode. The phases of the pair members are in quadrature with

one another, as expected, indicating a propagating pair.

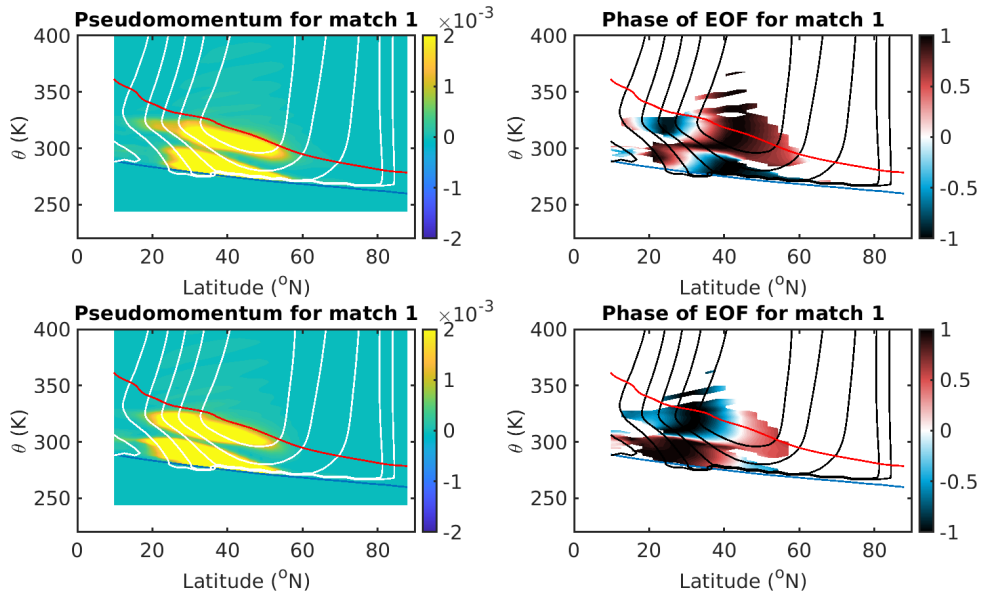


Figure 3.15: Amplitude (left hand panels) and phase (right hand panels) for the two leading modes of filtered wavenumber 2 for the T-only Rayleigh drag experiment. The phases are shown on the right of each subplot. The 2PVU line is shown in red and background wind is overplotted in white with a contour interval of $10\text{m}\cdot\text{s}^{-1}$. The blue line shows the surface.

Relaxall with Rayleigh Drag

The RARD experiment similarly possesses a flat total power spectrum across all zonal wave numbers up until wavenumber six, tailing off towards smaller scales (figure 3.16(a)). There is however a peak in the power contained within the leading pair at wavenumber 5 with the leading pair accounting for relatively more of the variance of the ENM spectrum at this scale, indicating a single coherent mode dominating the variability at this scale.

Examining the structures of the leading ENMs, the low order planetary waves $m = 1, 2, 3$ resemble varicose wave structures, with two discrete areas of wave activity straddling the background jet with opposing phases. By zonal wavenumber 4, structures (not shown) instead display an asymmetric sinusoidal mode structure, indicating a meandering jet stream.

Similar to the TORD experiment, wavenumbers $m = 5, 6$ displays the most prominence, and marks a distinct change in character from larger planetary scale waves. In both cases, it is reasonable to argue that these baroclinic waves emerge because they are the wavenumbers with the strongest growth rate for baroclinic instability (as argued in Zappa et al. (2010) & Zappa et al. (2011)).

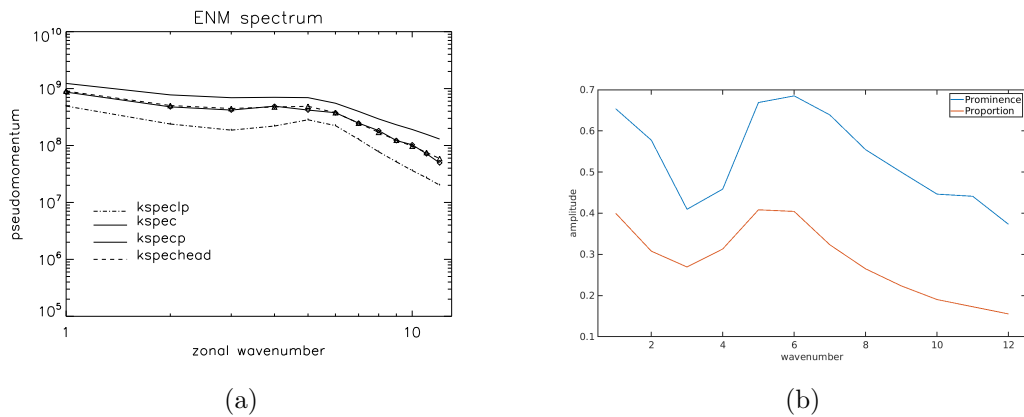


Figure 3.16: (a) The Pseudomomentum spectrum for the Relaxall Rayleigh drag experiment, and (b) Measures of prominence and proportion of spectrum for each wavenumber for the RARD experiment. In (a) the thin line is the total spectrum (kspec), the dash-dot line (kspeclp) contains only the leading pair of each wavenumber, and the thick line is only the paired eigenvalues (kspecp). The final dashed line (kspehead) shows the power contained in the first 10 ENMs (without pairing). By design, the total spectrum (kspec) is highest in pseudomomentum, as each of the other quantities is a subset of this quantity.

Variance Captured by Leading ENMs

By examining figures 3.11, 3.13, 3.14, and 3.16 (a), we can assess how well the ENM technique captures the total amount of variability in the data within the first few modes. Firstly, it’s worth noting that across all experiments, the pseudomomentum power in both the set of paired ENMs (i.e. with a matching statistic (eq. 2.59) greater than 0.7) and the set of the leading 10 ENMs in terms of eigenvalue (regardless of paired status) are similar at virtually all scales. This shows that even amongst the modes which explain the most variance, the variability is primarily dominated by paired propagating modes.

The difference then, between the pseudomomentum power across the whole spectrum and only the paired propagating modes, indicates the partition of variability between the coherent modes and the “saturated” noise spectrum. Generally, for wavenumbers below the peak in power at $m \approx 5, 6$, figures 3.11, 3.13, 3.14, and 3.16 (a) show that the coherent modes explain $\approx 70\%$ to 80% of the variability. At the smallest scales, this decreases to $\approx 50\%$, however the power generally is weaker at these scales than at the larger baroclinic and planetary scales, by an order of magnitude or more.

Examining *just* the leading pair in this fashion, we see that the leading ENM pair accounts for up to $\approx 30\%$ of this variability at its peak, generally at the baroclinic scales ($m = 4, 5, 6$). Generally we see then that the single leading modes at these scales are the most dominant; they explain the greatest amount of variability while possessing a coherent structure which describes a propagating wave.

3.5 Phase Speeds

Once the ENMs have been calculated, both the dynamical phase speed, c_{int} , and the empirical phase speed, c_{emp} can be calculated. The large-scale planetary modes found in both the TO and RA experiments resulted in a poor match for the lowest zonal wavenumbers which possessed very large westward speeds. As the calculation of c_{int} requires the addition of the various components of the pseudoenergy (Chapter 2, equation 2.61) a cancellation occurs between advection terms and propagation terms. In the case of a fast-moving jet with a large-scale mode propagating westwards against the flow, this results in large values of both the advection term H_w (due to the speed of the jet), as well as the KE and APE contributions to the wave. It is likely the case that for disturbances of exceptionally large speed, the linearized forms of \mathcal{P} and the projection onto \mathcal{H} are inaccurate; the outcome is a poor measure of c_{int} in these experiments, as can be seen in figure 3.17. Another possibility is that non-linear interaction between the waves results in higher wavenumber waves being ‘dragged around’ by the dominant $m = 1$ structure. As a result, we retain only the TORD and RARD experiments for further exploration of the phase speed.

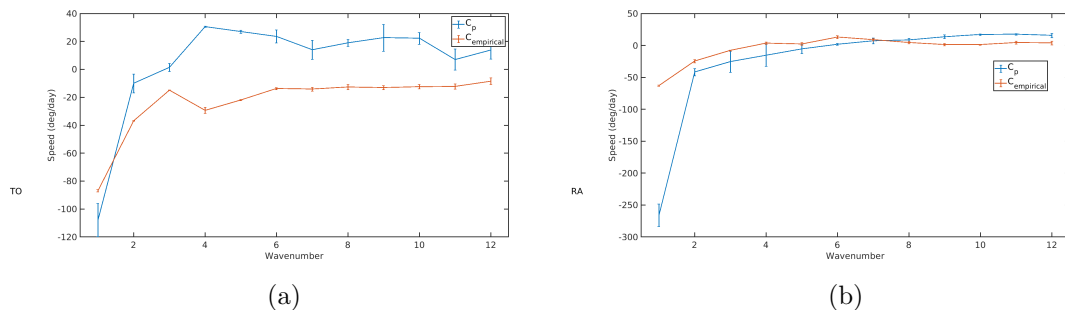



Figure 3.17: Empirical and Intrinsic phase speeds for (a) T-only, and (b) the Relaxall experiments in degrees per day, against the wave number. The error bars on c_{int} correspond to the intrinsic phase speeds of the individual members of the leading pairs, and on c_{emp} they correspond to the spread of values over the principal component timeseries.


Figure 3.18 shows the phase speeds for the leading ENMs for each wavenumber in both the TORD and RARD experiments. By sub-sampling the analysis window (see section 3.6.3), a number of samples of c_{int} may be attained in order to establish a measure of uncertainty. The estimate of c_{emp} is derived by calculating the forward difference in phase of the leading pair at each timepoint. Then, all points where the amplitude does not exceed the 40th percentile of the total amplitude over the whole timeseries are discarded. The average of the phase difference of the remaining time points is taken, and the average change in phase is used to calculate the angular velocity of the ENM in degrees/day (i.e.

c_{emp}).

The bars in figure 3.18 show the variance for both estimates of speed. The variance for c_{emp} is quite large; this represents changes in the value of c_{emp} over the timeseries as opposed to uncertainty in the measured value as calculated from subsampling. This can be seen by ‘unwrapping’ the phase of the ENM over the time-window and observing the changes in the slope over the timeseries (figure 3.19).



(a)



(b)

Figure 3.18: Empirical and Intrinsic phase speeds for (a) *T*-only RD, and (b) the Relaxall RD experiments in degrees per day, against the wave number. Box and whisker values for c_{int} are obtained from subsampling (see section 3.6.3), and on c_{emp} they correspond to the spread of values over the principal component timeseries.

These changes can come about for a number of reasons; non-linear interactions with

other modes which are not consistently present during the time-window; variation of the background flow; an inconsistency of the modes across the window; if the time-series is short compared to the timescale of variability of the mode, and so on. There are also uncertainties in the ENMs captured; such as modes picked out not being perfectly in quadrature, or failing to project well onto the "true" modes of the system being measured.

A similar picture of the phase speed match emerges from both the TORD and RARD experiments. At baroclinic scales (wavenumbers ~ 5 to 8), there appears to be a good match between the speeds, notably at wavenumbers with greater leading pair prominence and where a high proportion indicates a greater amount of power is contained within the leading pair. These are baroclinic waves, and as shown above (figures 3.14, 3.16) single ENM pairs contain a large proportion of the pseudomomentum (and as we shall see later (Chapter 5), perturbation energy). At the smallest scales (highest zonal wavenumber), c_{int} is larger than c_{emp} , the mis-match increasing with decreasing scale. At the largest scales, where prominence is low, the matching criteria is not met once more.

At the highest zonal wavenumbers, the structures picked out are small structures advected along the jet, reducing their empirically measured speed relative to their intrinsic phase speed predicted from their structure. Movies showing the evolution of the PV field with time show prominent propagating long waves which break and produce PV filaments. These smaller-scale filaments project onto all higher wavenumbers, but the structures at all wavenumbers describe the same features which are being advected by the flow. Therefore, the empirical phase speeds observed of these "modes" are similar to the phase speed of the dominant $m = 5, 6$ structures which *are* modal, as well as dominant and energetic.

Examinations of correlations between the amplitude timeseries show positive correlations between zonal wavenumber 1 and 7 & above in the RARD experiment (figure 3.20(b)). This, in combination with the similar empirical phase speeds between $m = 1$ & $m \geq 7$ (figure 3.18(b)) hints at a possible non-linear correlation between these waves. This, in combination with the cross-correlation between all waves with wavenumbers 7 and higher in both the TORD and RARD experiments (with the exception of $m = 7, 8$ in TORD) indicates that the structures are indeed filament structures which project across all high wavenumbers.

Similarly, both phase-speed matching and correlation relationships between a number of lower wavenumber ENMs indicate the possibility for Rossby wave triad interactions between the three to account for the mis-match in phase speeds between the empirical

and intrinsic measures of speed. c_{int} in this case represents the modal speed of the structure were it not for the non-linear interactions, whereas c_{emp} shows the modified speed of the waves as seen in the data.

These interactions are non-linear couplings between groups of three waves, where each combination of two out of the three waves in the triad serves to reinforce the third, such that energy is exchanged between the waves in the triad but not to other waves outside of it. These interactions are restricted to combinations of waves whose wavenumbers sum, and it is expected that two members of the triad are expected to be correlated in amplitude with time, while the other is anti-correlated, and that the modes are close in wavenumber.

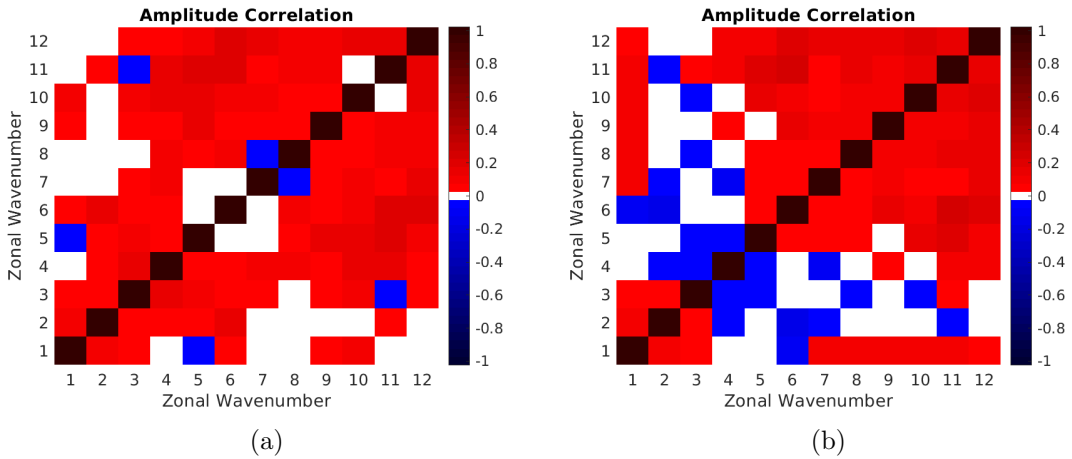
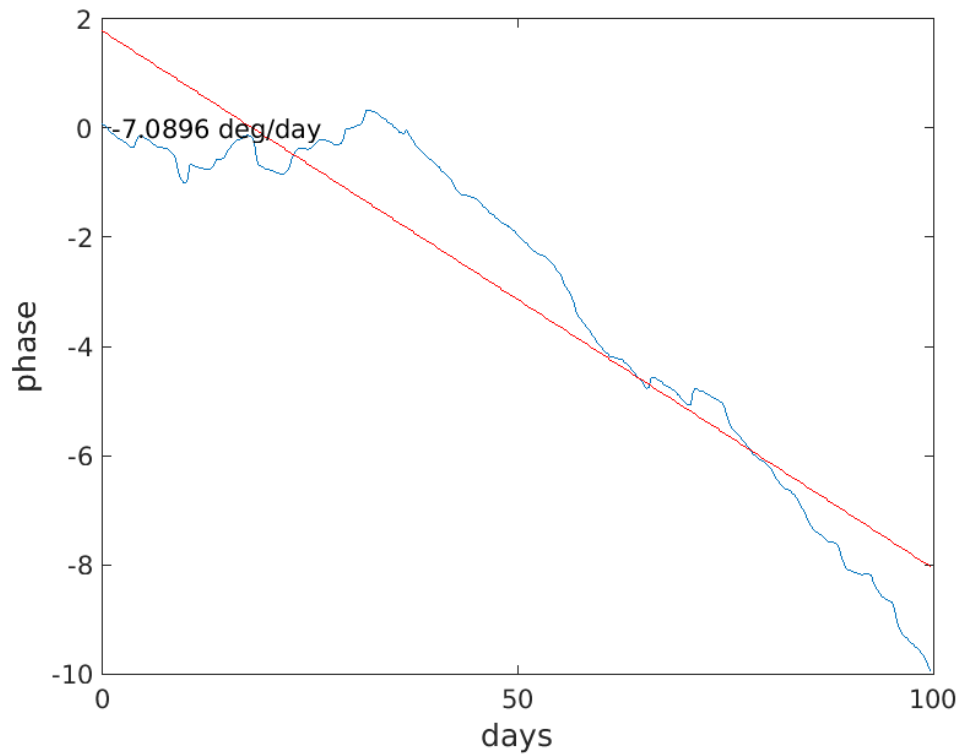


Figure 3.20: *Cross-correlations in the principal component amplitude timeseries of the leading ENM pair of each wavenumber for (a) the TORD and (b) the RARD experiment. Red denotes correlation, whereas blue indicates anticorrelation.*


Examining the cross-correlations for the TORD experiment (figure 3.20(a)), positive cross-correlations between wavenumbers $m = 1, 2, 3$ rule out this combination as a triad. Wavenumbers $m = 1, 4, 5$, display a similar empirical phase speed, and show one positive correlation (4,5), one anti-correlation (1,5), and one neutral amplitude relation (1,4) which could nonetheless imply a triad interaction between the three modes, which also share similar power (figure 3.14) Of note is the lack of correlations between the dominant baroclinic wavenumbers $m = 5, 6, 7$ implying that these waves do not interact with one another, instead only projecting onto less dominant modes.

In the RARD experiment, correlations between $m = 2, 3, 5$, as well as their close empirical phase speeds indicate the presence of a triad interaction between these three modes. The modes also possess similar power (figure 3.16), and $m = 2, 5$ possess prominent leading pairs, implying that the combination of these wavenumbers possibly project

onto $m = 3$. In either case, further investigation is required in order to test these hypothesized triad interactions.



(a)



(b)

Figure 3.19: (a) Phase of the leading ENM for $m=5$ in the RARD experiment against time. As the absolute phase of ENMs is arbitrary, discontinuities in phase can occur during a drop in amplitude where phase is ‘re-established’ in another part of the domain. The overplotted red line shows a linear fit, unsuitable as an estimate of the empirical phase speed due to discontinuities in phase over the timeseries. (b) Principal component timeseries for the mode shown in (a) with the leading pair plotted in blue and orange, and the total amplitude of the ENM in yellow. Points where amplitude is low (due to poor quadrature match of the pair members) correspond to parts of the phase timeseries with changing slope.

3.6 Robustness

In order to justify use of the ENM technique, examining performance in a multitude of situations is important, as this will determine its robustness. An advantage of the ENM technique is that it presents an opportunity to extract structures from timeseries where the sampling frequency is low, or where the sample window is short, as the dynamical phase speed c_{int} is derived from the structure of the modes and hence does not require a lengthy or regular timeseries from which to extract phase information. In order to examine how well the technique functions under these conditions, and to characterise any uncertainties in the results of the technique, experiments were performed where both the sampling frequency, and the sampling window were varied.

3.6.1 Eigenvalue Ratio

In order to describe a propagating wave, two ENMs varying in quadrature are required. The best description can be given when the members of the ENM pair possess a similar amount of power associated with them, i.e. their eigenvalues are well-matched. ENM pairs with dissimilar eigenvalues are likely to describe structures less well, as pairs propagating exactly in quadrature should possess identical eigenvalue. This is because the pseudomomentum matrix which goes into the eigendecomposition (equation 2.48, Chapter 2) depends on the background state (equation 2.34, Chapter 2) which is zonally symmetric and steady, resulting in modes propagating in quadrature to possess the same pseudomomentum structure, and hence eigenvalue. As a result, their measures of wave properties such as phase speed are likely to be poorer. The empirical phase speed c_{emp} is measured through an unwrapping of the ENMs' phase over the timeseries (figure 3.19, and a finite difference formula is applied during periods where the amplitude is above the 40th percentile. Figure 3.21 shows the number of wave cycles in the analysis timewindow, $T_{analysis}/\tau_{empirical_i}$ (which is calculated from the empirical phase speed), against the difference between the eigenvalues of the leading pair normalized by the amplitude from the experiments described in Chapters 4 & 5. It can be seen that waves with a smaller number of complete cycles over the analysis window exhibit larger possible discrepancies between the eigenvalues of the leading pair. As $T_{analysis}/\tau_{empirical_i}$ approaches 1 the extraction of the ENM propagating modes is not as clean, as indicated by the spurious difference in eigenvalues.

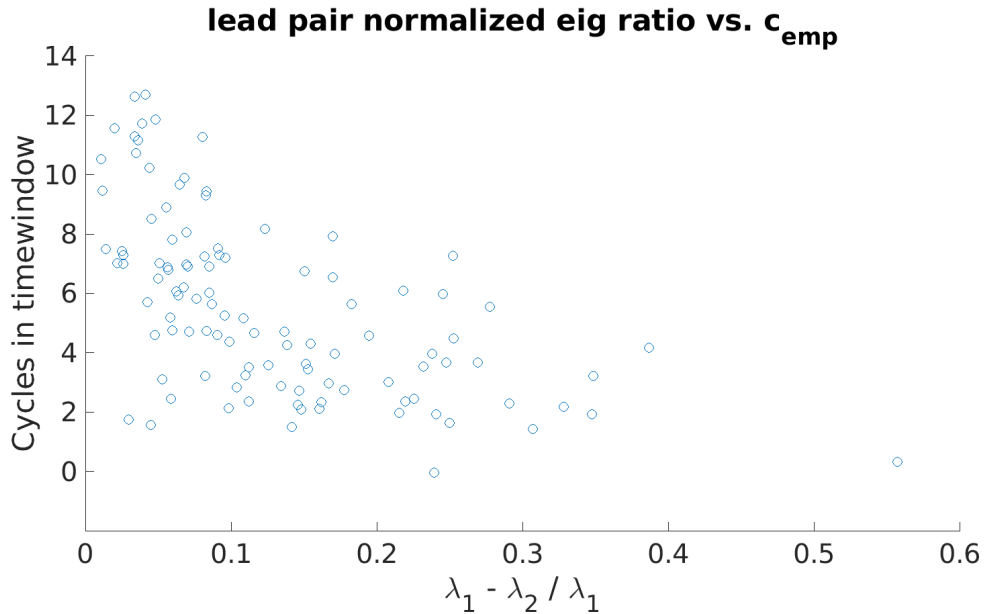


Figure 3.21: Number of cycles ($T_{analysis}/\tau_{empirical_i}$) in the 400 day time-window of the ENM analysis of the jet-latitude control experiments performed in Chapters 4 & 5 vs. the normalized eigenvalue difference of the leading pair chosen by match & structure (see chapter 5 for details). Points are taken from zonal wavenumbers $m = 1 - 12$ for each of the 9 experiments for 108 points total.

3.6.2 Sample Frequency

Sample frequencies of 6 hours (the minimum possible based on the output of the IGCM), 12 hours, 24 hours and 48 hours were used in order to examine the effect on the phase speed distributions.

Across all zonal wavenumbers the measured phase speeds do not depart significantly from the values measured for a 6 hourly sample interval. At least up to 48 hours, the ENM technique does not appear to be sensitive to the sampling interval. This is due to the calculation of c_{int} relying solely on the spatial structure of the modes and not timeseries information. Equivalent plots to figure 3.22 for c_{emp} demonstrate that it is less robust than c_{int} and requires that the timestep $\delta t \ll \tau_{empirical}$. As a result, this is not shown, as the data cannot be used.

An exception is if an ENM's period of oscillation matches exactly the sample interval. In this case, the mode itself will not contain any variability in the dataset and thus will not appear as an EOF in the eigendecomposition, and thus neither measure of phase speed may be applied.

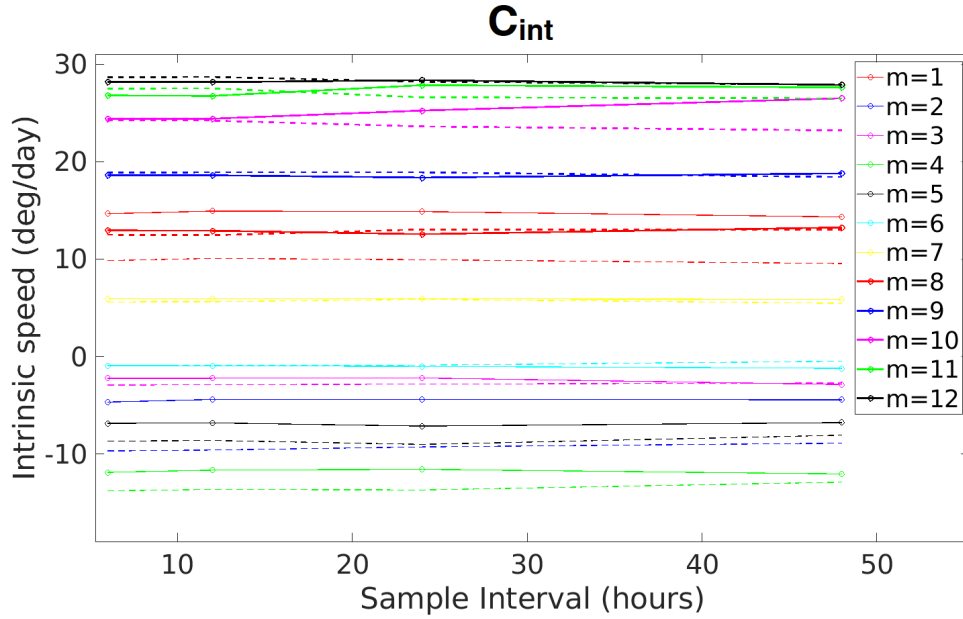


Figure 3.22: *Intrinsic phase speed c_{int} for the T -only experiment, against the sampling interval in hours. The solid line denotes the leading member of the eigenvalue series, the dashed line denoting the second member of the same pair. Each wavenumber from 1 to 12 is represented in different colours and linethicknesses as indicated in the legend.*

3.6.3 Sample Window

The study of ENMs across varying sample windows was conducted by first running the analysis across a slightly reduced 672 day window, (as this length contains a greater number of numerical factors for subdivision) before running with smaller intervals (measured as integer fractions of the total window length); one half, one quarter, one eighth, and one sixteenth. For each of the reduced intervals, analysis were performed for each sub-window and then averaged as in [Wheeler and Kiladis \(1999\)](#) for power spectra in wavenumber-frequency space. The intervals were arranged in such a way as to span the entire reduced timeseries back-to-back, with overlapping windows between each sub-window, e.g. the half-window length analysis was performed across the first half of the full window, then the second half, and then the inter-quartile half, and a spread is produced as a measure of uncertainty for the phase speed of the modes. This was done in order to ensure that the modes dominant at the point a subwindow was taken did not skew the results; by incorporating subwindows from across the entire timeseries and producing a spread, it was ensured that the exact same data was being represented in the analysis as in the full timeseries window.

Data for wavenumbers $m = 1-4$ & $m = 5-12$ were separated into figures [3.23](#) & [3.24](#) respectively for the sake of clarity, as there is a difference in dynamics (and hence phase speed) between these two sets of modes. As expected the range of values of speed are

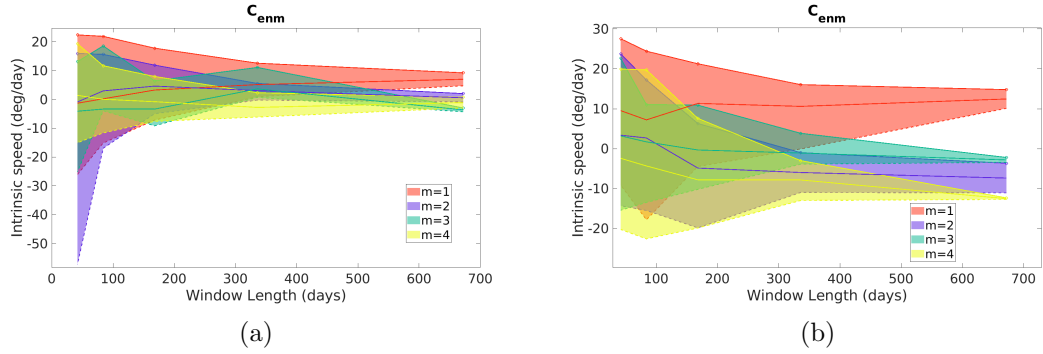


Figure 3.23: Speed (c_{int} of the leading pair for wavenumbers 1 to 4 in experiment TORD (a) & RARD (b) against the length of the subsamples window, with wavenumbers grouped into two groups for clarity. The shading indicates the range of values for speed at each subwindow.

reduced as the length of the analysis window increases, for the low wavenumbers in both experiments. Though it is important to note that the number of values used to estimate the speed also decreases as the window length increases; by the maximum window length, there is only a single sample of phase speeds (i.e. the analysis taken from 1 whole window) with two values for c_{int} (the two members of the leading pair), being used to determine the ‘range’.

For the shorter waves, shown in figure 3.24, the speeds asymptotically approach the values observed in the longest window length in some cases, and are reasonably accurate at even small window-lengths in others. It is possible that changes in speed observed as the window length changes could result from alternative structures displaying more dominance in some windows. It is not guaranteed that the same structures are extracted at the same amplitude or with the same matching strength in each sub-window, which could mean that as the timeseries progresses and different modes from within a single wavenumber’s ENM spectrum gain or lose amplitude, the ‘leading pair’ extracted is not the same mode. This could affect the average value of phase speed as the number of windows averaged over increases, though this is likely to be a small source of uncertainty. As the value of c_{int} from the full window analysis is retained, and the skew would only serve to potentially widen the uncertainty estimate of that measurement of phase speed, it is of little concern.

These conditions on the sampling interval and sample window allows us to apply a set of restrictions on the waves that we’re able to extract with the ENM technique;




Figure 3.24: Speed of the leading pair for wavenumbers and 5 to 12 in experiment TORO (a) & RARD (b) against the length of the subsamples window, with wavenumbers grouped into two groups for clarity. The shading indicates the range of values for speed at each subwindow.

$$\tau_{\text{empirical}} < \tau_{\text{window}}$$

&

$$\tau_{\text{empirical}} > \tau_{\text{interval}}$$

3.7 Conclusions

A set of experiments was designed and conducted in order to test the validity of the Empirical Normal Mode technique, applied using the MLM background state laid out in 2. The experiments possess distinctive variability in order to establish the various conditions where the ENM technique produces structures which vary in a modal way, as established using the phase speed condition test and by exploring their sensitivity to the data sampling, giving reasons for deviation from modal behaviour. Nonetheless, an important aspect of the experiment set is their simplicity such that it is possible to examine the results of the ENM analysis without contending with the influence of e.g. nonconservative diabatic effects, orography, or the diurnal cycle. In order to ensure a steady background state, seasonal variability was omitted, as time-independence is required for conservation of \mathcal{H} , crucial for the calculation of the phase speed of the extracted structures.

Of the four experiments developed, the experiments which possess surface drag possess ENMs with good phase-speed matching, with some caveats. Especially large-scale waves still match less well in phase-speed in some cases; this is likely to do with nonlinear interaction, such as Rossby wave triads though further investigation is required to confirm this. Waves at the smallest scales also propagate differently to the intrinsic phase

speed their structures would suggest; this could be due to weak nonlinear effects; another explanation is advection of increasingly small-scale waves by the jet reducing the empirical phase speed. The remaining two experiments (those without surface Rayleigh drag) display large, long-wavelength dominance and very high speeds, which are represented poorly by the intrinsic phase speeds of the technique, but well by the empirical speeds derived from tracking the phase of the extracted modes.

Subsampling this analysis allows the establishment of measures of uncertainty of the calculated intrinsic phase speed, which provide additional confidence that the phase speeds calculated are reliable. A condition on the sampling must be imposed in order to extract propagating modes with matching phase speeds; $\Delta t < \tau_{empirical} < T_{window}$. In addition, intrinsic phase speed c_{int} is easier to extract from limited data than empirical phase speed c_{emp} , where large $\tau_{empirical}/\Delta t$ and $T_{window}/\tau_{empirical}$ are required for tight estimates of phase speed.

The work in this chapter establishes that the ENM technique, implemented for the primitive equations on the sphere, is capable of extracting propagating modes and an accurate estimate of their intrinsic phase speeds. Although the phase speed matching test only works for the dominant baroclinic wave modes, there are good reasons why the longer waves ($m = 1 - 4$) and the shorter waves ($m > 8$) do not meet the phase speed matching conditions in the experiments examined - as a result of nonlinear triad interaction and Rossby wave breaking respectively. Nevertheless, we will work on the assumption that the technique extracts reliable intrinsic phase speeds associated with all modal structures and when the phase propagation diagnosed empirically from time series differs this is a result of non-modal processes (either nonlinear interaction between wavenumbers or non-conservative effects). In the following chapters, it will be used to examine the dependence of modal variability on the latitude of the jet stream. The advantages of the technique; deriving measures of speed from structures such that they can be used without spatial or temporal filtering, and on small time-scales where it might not be possible to track full oscillations of slow-moving Rossby waves are apt to potentially be applied to a myriad of studies of atmospheric behaviour.

Chapter 4:

Relaxation Experiments With Controlled Jet Latitude Variation

4.1 Introduction

In the previous chapter the use of the Empirical Normal Mode technique on a set of simple idealised experiments is demonstrated in order to conduct the ‘Phase Speed Condition Test’ and verify the technique for propagating disturbances that behave like normal modes, and also identify the circumstances where the intrinsic phase speed associated with ENM structure and its movement from its PC time series do not match. This technique will be used in Chapter 5 in order to examine changes in wave structures as the jetstream changes in latitude. In order to do this, it is necessary to develop a dataset with a range of latitudinally-confined jets, without forcing effects such as orography or seasonal variability which create large-scale non-propagating structures which are not described by the ENMs. Also desirable is a sloping tropopause in the relaxation state used in order to maintain upper-level baroclinicity in the experiment; as the end goal is to examine the relationship between wave activity and propagation as jet latitude is varied it is important to maintain Rossby wave activity on the jet.

In this chapter a set of model experiments with confined jets are developed using the Reading IGCM2.2 which can be used to conduct an analysis of the structure changes as jet latitude is varied (Chapter 5). These experiments use a series of relaxation states with varied structure in order to vary the latitude of the jet; the strength of the relaxation forcing is controlled by a relaxation timescale parameter. As the model is a nonlinear primitive equation model, it is not necessarily (or likely) the case that the prescribed relaxation state will match the equilibrium background state of the atmosphere, as wave

activity is able to drive differences between the background flow and the relaxation field used for weak forcing of the climate state. As a result, the amplitude of wave activity and the quasi-equilibrium position of the jet formed through the action of breaking Rossby waves and sharpening of the tropospheric polar vortex edge is not known a priori and must be determined through numerical experimentation.

The sensitivity of the output to model parameters is examined, and the modified Lagrangian mean background state and wave activity are calculated; both in order to perform the ENM calculation in the next chapter, and so the variability of the quasi-equilibrium state obtained between experiments can be compared.

4.2 Experiment Specification

4.2.1 IGCM Parameters

The model used in these experiments is the Reading Intermediate General Circulation Model version 2.2, as in Chapter 3. Again similar to Chapter 3 moist physics have been turned off in the model, and there is no orography in the experiments. The model is run with T85 spectral truncation and 30 vertical levels (T85, L30 resolution). A ∇^6 hyperdiffusion is used, as well as a Rayleigh drag scheme in the lower troposphere, and Newtonian temperature Relaxation, the field of which is detailed below.

Two sets of experiments were initially run; the first with a relaxation timescale of 15 days, the second with a reduced timescale of 7.5 days (and hence increased relaxation strength). This secondary set of experiments were run in order to test the sensitivity of the model response to relaxation strength to ensure that results were robust to parameters not used to control the jet latitude. Similar to Chapter 3, the Newtonian relaxation scheme (but not the particular temperature field used) matches those prescribed in Held and Suarez (1994b), with a relaxation to temperature only applied in this case. The drag scheme used in the experiments with drag in Chapter 3 is also applied in these experiments.

4.2.2 Relaxation Field Definition

We seek a zonally symmetric state with a jet which is confined both in latitude and height, with a sloping tropopause to use as a relaxation field. Constructing such a state analytically which is consistent with the zonally symmetric primitive equation gradient wind balance on the sphere required by the model is not trivial. In order to achieve this,

the balancing routine within the primitive equation model itself (Hoskins and Simmons (1975)) can be used to find a consistent balance between variables of a partially-defined analytic state which, e.g., only defines the vorticity, or else the temperature & surface pressure, not both. We specify the zonally symmetric zonal wind (and thus the associated vorticity structure from its meridional gradient). The model balancing routine is used to calculate the temperature & surface pressure fields. The state in balance with the jet in the zonal flow is defined using a flat tropopause with a jump in static stability at a given sigma level. Therefore it is blended linearly with the Held-Suarez state which has a sloping tropopause, but not a meridionally confined jet.

Initial Zonal Wind Field

In order to achieve a wind field with a jet movable by a single parameter, a method following Lai (2013) was used, which combined a wind field created by Polvani and Esler (2007), with a modification (eq. 4.2) from Lai (2013) to be able to shift the latitude. This set of equations were produced by Polvani and Esler (2007) in order to reproduce the ‘canonical’ structure of the LC1 experiment described in Thorncroft et al. (1993). In the LC1 experiment, the speed of the jet, U_0 is $45ms^{-1}$, and the latitude of the jet is 45° , corresponding to $p = 2$, reducing 4.2 to the form used in Polvani and Esler (2007).

$$u(\phi, z) = U_0 F(\phi) [(z/z_T) e^{-[(z/z_T)^2 - 1]/2}] \quad (4.1)$$

Where;

$$F(\phi) = \begin{cases} [\sin(\pi\mu^p)]^3 & \text{for } \phi \geq 0 \\ 0 & \text{for } \phi \leq 0, \end{cases} \quad (4.2)$$

including the p parameter introduced by Lai (2013). This form is chosen in order to comply with the requirements set out in Polvani and Esler (2007); namely that $F(\phi)$ & $dF(\phi)/d\mu$ both go to zero at the poles. This form of F results in a zonal wind field where the latitude of the jet maximum varies monotonically with p . The value of p corresponding to a particular latitude of the jet maximum ϕ_{relax} can be calculated with the following equation obtained from the condition $dF/d\phi = 0$ at the latitude of the maximum;

$$p = \frac{\ln(\frac{1}{2})}{\ln(\sin(\phi_{relax}))} \quad (4.3)$$

Temperature & Surface Pressure Fields

Using the balance routine in the model, the vorticity field associated with the zonal wind distribution (equation 2.10, Chapter 2) defined above can be used to numerically calculate the temperature field & surface pressure field in gradient wind balance with that vorticity field. The zonal temperature field specified in these experiments is the sum of that latitudinally dependent balance temperature and an analytically defined vertical reference temperature profile assigned as part of the balance in the primitive equation model (as the balance only provides the meridional gradient of temperature). The reference profile is defined;

$$T_{ref}(\phi) = \begin{cases} T_G + \Gamma \ln(\sigma) & \text{for } \sigma > \sigma_{trop} \\ T_G - \Delta T_{trop} & \text{for } \sigma \leq \sigma_{trop}, \end{cases} \quad (4.4)$$

where Γ is a constant tropospheric lapse rate, expressed in $\ln \sigma$ which is a pseudo-height. The temperature in the troposphere falls linearly with $\ln \sigma$ from $T_G = 285$ at the surface to $T_G - \Delta T_{trop}$ at the tropopause, $\sigma_{trop} = 0.4$ (such that $\Delta T_{trop} = 50$ is the drop in temperature over the troposphere). The flat temperature profile above the tropopause is by design in order to create a jump in static stability at a given pressure level. The result is a flat tropopause, which is unrealistic and will affect wave-breaking behaviour.

In order to ensure the relaxation state maintains baroclinicity, these fields are then combined through a linear blend with the relaxation temperature and initial surface pressure specified in Held and Suarez (1994b). This temperature field is the zonally symmetric field specified in Chapter 3 equation 3.1. The initial surface pressure specified is a uniform $P_{00} = 1000hPa$ throughout the model domain.

While the Held-Suarez (henceforth HS) state has been used in many experiments such as Polvani and Kushner (2002), Boer and Denis (1997), Franzke (2002), Harnik and Chang (2004), Seager et al. (2003), and Williams (2003), it was considered unsuitable for use on its own for these experiments. The zonal wind profile in gradient wind balance with the HS temperature field increases all the way to the equator, resulting in strong westerlies aloft at low latitudes (figure 3.1, Chapter 3). This does not allow for the generation of a latitudinally-confined jet. The HS state does, however, provide a more realistic sloping tropopause than the temperature field in gradient wind balance with Lai (2013)'s modification of the Polvani and Esler (2007) zonal wind field. Therefore, the two states are blended in order to utilise the properties of both the HS state and Lai (2013)'s

variable latitude reproduction of the LC1 state.

The blend results in a state with two key properties: (1) The jet is meridionally confined, and its latitude can be controlled. (2) it possesses a sloping tropopause which maintains upper level baroclinicity in the experiments. In addition, due to the resulting field being of a notably lower temperature than Earth’s atmosphere, an adjustment parameter is introduced to raise the resultant temperature, as latitude-independent adjustments do not affect the balance (Polvani and Esler (2007)). The temperature field is then calculated as;

$$T_{JC}(\phi, \sigma) = \alpha T_J(\phi, \sigma) + (1 - \alpha)T_{HS}(\phi, \sigma) + T_{adj}. \quad (4.5)$$

Where T_J is the temperature field obtained by gradient wind balance from the analytically prescribed wind field, T_{HS} is the relaxation temperature field defined in Held and Suarez (1994b), α is the parameter used to blend them, T_{JC} is the resultant temperature field used in the experiments, and T_{adj} is the temperature adjustment parameter, 5K in this case, chosen such that surface temperatures in the tropics match observations. The surface pressure from the Lai and HS states are combined in the same way, though the adjustment parameter for pressure was zero in these experiments.

Initial Conditions

The resultant surface pressure and temperature fields are used as the initial state of the experiment model runs, along with the zonal wind field obtained by balance *back* from the blended temperature & surface pressure. In addition, the temperature field is used for the relaxation forcing. The state is designed to be reproducible, as each component; the HS state, and the T_J state are specified analytically in the temperature and wind variables respectively. However, numerical solutions of balance for the primitive equations on a sphere are required to derive other prognostic fields, and to construct the initial state in temperature and surface pressure (by balancing from the prescribed wind field) and the initial vorticity field (by balancing back from the constructed temperature and pressure fields). The state obtained then is in balance up to the precision of the discrete equations represented by the IGC. Due to these states’ balance, if they are used for the initial conditions, the flow remains exactly steady, as the zonally symmetric balance state is a solution of the discrete model.

As such, a small perturbation is required in order to “kick-start” the model into

non-steady flow. A random noise component is added to the surface pressure, using the IGCM's NOISE subroutine, a common approach used for example by [Polvani and Esler \(2007\)](#) and also by [Magnusdottir and Haynes \(1996\)](#). The noise is applied to all surface pressure spectral co-efficients, has a 0.1hPa amplitude and is set using a pre-determined seed (meaning that it is the same for all runs). This is the same as was used for Chapter [3](#).

Test of the Balancing Routine

It is expected that the action of the forward and inverse balance routines obtains the result that you start with, but due to approximations in the numerical balance it does not exactly. This is because only the lower spectral coefficients are used in balancing in order to control noise. In order to conduct a brief test of the routine, we first specify the zonal flow using equation [4.1](#) and then balance to temperature and pressure with the IGCM, and we are then able to balance *back* from temperature and pressure (without performing a blend) to zonal flow. When doing this, the resultant flow is not exactly equal to the initial zonal flow field, and the jet is somewhat weaker. Regardless, the quasi-equilibrium jet state of the nonlinear experiments is quite different to the initial and relaxation state, so these small differences in the initial state will not have much influence, and the overall structure of the state is preserved. Similarly, it would be possible to have started with the HS temperature field, balance to vorticity, and balance back and consider the difference between the analytic temperature field and the field resulting from balance, however here we chose to instead examine the wind field to ensure the most important requirement of the field; a latitudinally confined jet; is preserved.

4.3 Parameter Space

4.3.1 Blending Jet Specification with Sloping Tropopause in Relaxation States

In examining the parameter space, the balance between these two key properties; latitudinal confinement of the jet, and a change in static stability across the jet, was desired for all latitudes of the jet maximum in the relaxation states. As such, a range of α were examined from 0.5 to 1, shown in figure [4.1](#). A value of $\alpha = 0.7$ was determined to capture the required variation in static stability, as well as preserving the shape and latitude of the prescribed jet. This was done both in order to pursue the experimental aim

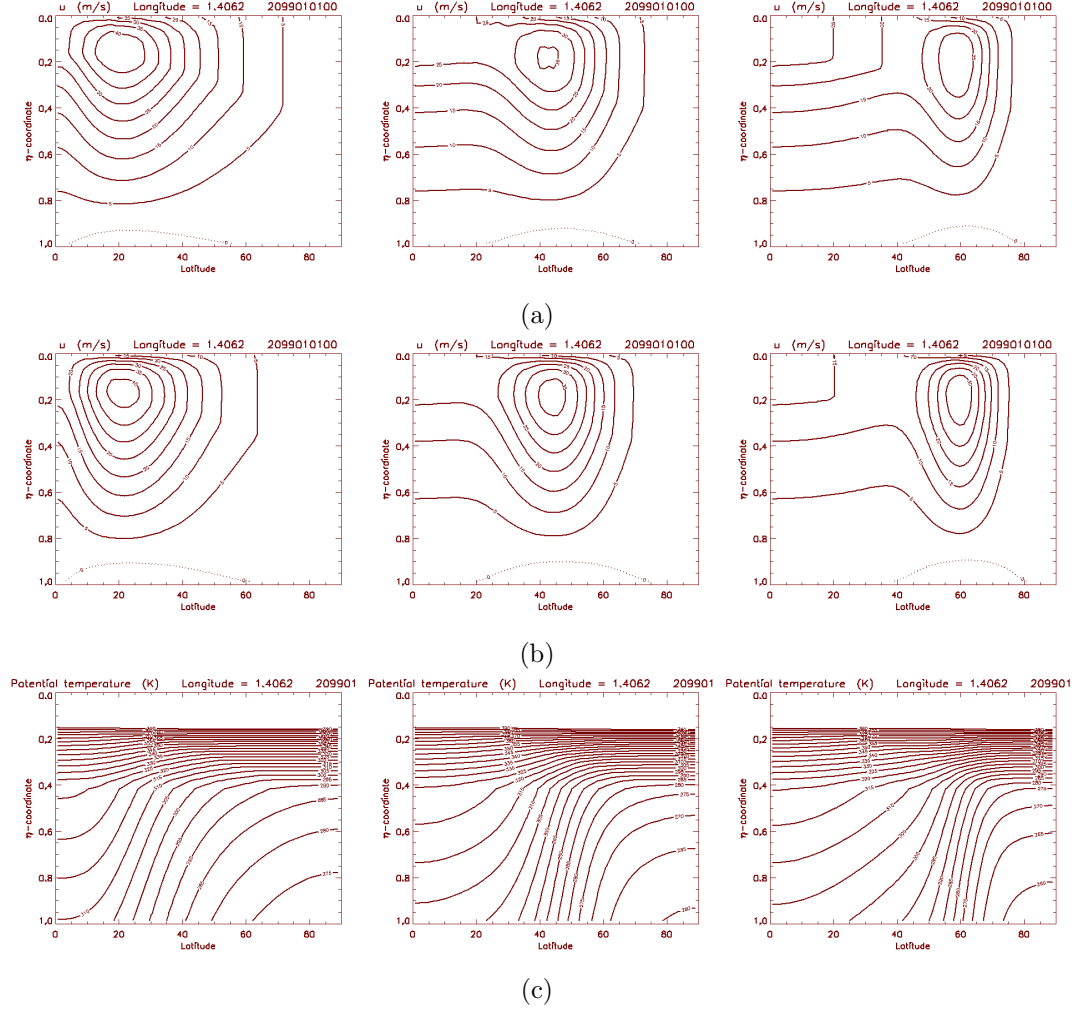


Figure 4.1: Zonal wind field for a subset of experimental setups; (a) $\alpha = 0.5$, (b) $\alpha = 0.7$, and p (calculated using 4.3) corresponding to relaxation jet maxima of 22° , 45° , and 60° from left to right. (c) θ field for $\alpha = 0.7$, in balance with the zonal flows in (b).

of a latitudinally confined jet, as well as to attempt to minimize the tropical westerlies aloft in the experiments.

4.3.2 Relaxation Latitude

The specification of the relaxation jetstream defined in 4.1 and 4.2 requires ϕ_{relax} to exceed 21° in order to comply with the conditions on $F(\phi)$ outlined above, providing a lower limit to the latitude of the relaxation state matching that of a subtropical jet. $p = 2/3$ is the minimum for desired limiting behaviour of the first and second derivatives of F approaching the equator (corresponding to $\phi_{relax} = 20.7^\circ$), however, due to discretisation of the field at the model resolution, a stronger constraint of $p \geq 1/\sqrt{2}$ was used. At high latitude, 60° was used as an upper limit to relaxation jet latitude, to match the upper limit on mean observed polar jet latitudes. Between these values, a spacing of 5°

was used, (with the exception of the lowest latitude values of 22° and 25°). Equation 4.3 is used to calculate the corresponding values of p which are then specified in the model, providing nine separate experiments with distinct relaxation latitudes in a range of values between the subtropical and polar regions.

The experiments are designated as $p22$ through $p60$; indicating the value of ϕ_{relax} set by varying the p parameter in the analytic specification of the wind field. Note that the value of p itself is not equal to the value in the experiment identifier, but is calculated from ϕ_{relax} using equation 4.3.

4.4 Experiment Results

4.4.1 Initial Spin Up

All experiments were run for a period of 500 days, with each experiment requiring a spin-up period of between 20-100 days to reach a ‘quasi-equilibrium’ flow, characterised by an almost steady modified Lagrangian mean background flow and also a steady magnitude of nonlinear wave activity. All experiments exhibited strong anti-cyclonic wave breaking, leading to filamentation, which served to sharpen the tropopause boundary, and shift it polewards relative to the jet maximum in the relaxation state. This is consistent with the LC1 experiment of [Thorncroft et al. \(1993\)](#) which describes ‘troughs’ filamenting equatorward anticyclonically off the jet itself, and which was the paradigm which [Polvani and Esler \(2007\)](#) were attempting to emulate with their simulations.

The initial wave breaking period where this poleward migration occurs is followed by steady flow at a consistent equilibrium latitude in the MLM state (and Eulerian zonal mean state).

Some experiments were run with an additional surface pressure perturbation of a large magnitude, giving rise to a corresponding barotropic zonal flow in balance with the surface pressure distribution, adding to the baroclinic flow specified by the blending procedure described above. A number of supplementary experiments were run with perturbations varying sinusoidally in the meridional direction possessing amplitudes of 10hPa, 20hPa, 40hPa, & 100hPa added to the initial conditions. This is motivated directly by the LC2 experiment from [Thorncroft et al. \(1993\)](#), which only differs from the LC1 (which the [Polvani and Esler \(2007\)](#) is a reproduction of) in the meridional surface pressure variation in the initial state, but the growing mode breaks cyclonically rather than anticyclonically. The barotropic component of the flow changes the initial state

substantially, adding a strong cyclonic meridional wind shear across the mid-latitudes. However, this component is independent of the temperature field and therefore cannot be maintained by relaxation forcing of temperature. Although the initial Rossby wave breaking behaviour is changed, the background state does not evolve to a different equilibrium structure after spin-up, indicating that the relaxation temperature field and associated baroclinic zonal flow component is the primary influence on the quasi-equilibrium jet structure obtained, together with the Rossby wave activity on that jet.

Examining zonally averaged timeseries of the upper level wind, shown in figure 4.2, it is clear that experiments with a more equatorward initial and relaxation state exhibit a longer transition period to quasi-equilibrium flow, and undergo further poleward migration than the states with a more poleward relaxation state. These states with more equatorward jets show less variability, and possess a more latitudinally confined jet with longer-scale variability. This is due to the quasi-equilibrium state of the higher latitude jets being closer to that of the relaxation state, and as such the relaxation forcing is weaker.

This increased spin-up time can be seen in the global integral of energy, $\langle c_p T \rangle$; while all experiments settle on a similar value of total internal + potential energy (i.e. neglecting kinetic energy), figure 4.3 shows the time taken for the energy to flatten out into an equilibrium is much longer for the lower latitude $\phi_r = 22^\circ$ experiment than for $\phi_r = 55^\circ$, which after the initial dip (representing the initial baroclinic lifecycle) is almost

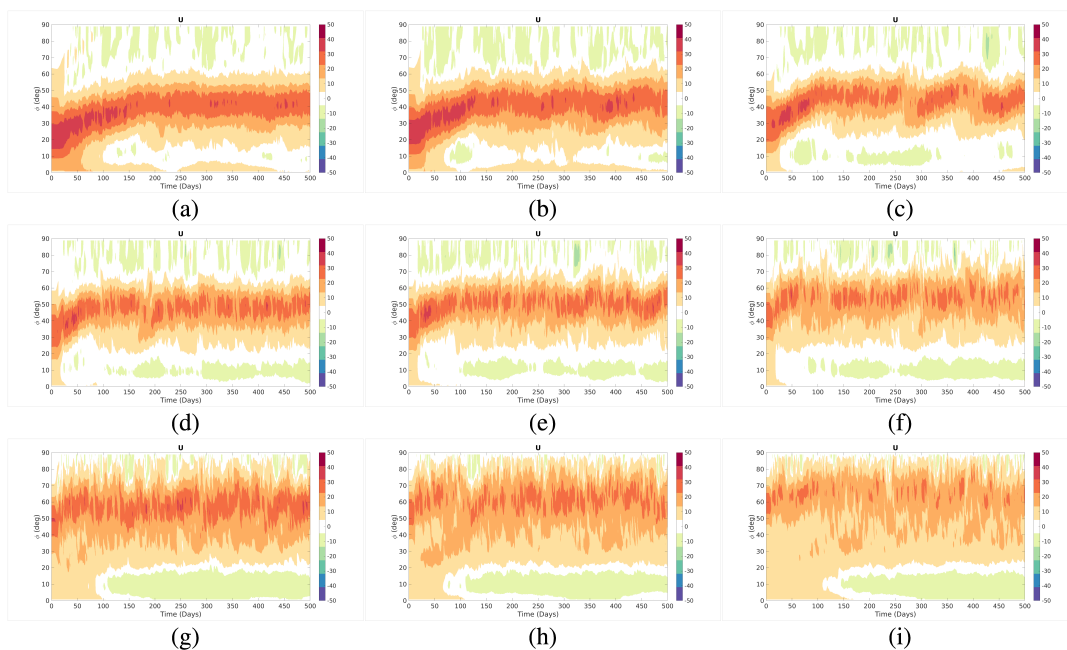


Figure 4.2: Eulerian zonal average zonal wind on $\sigma \approx 0.25$ against time for relaxation latitudes 22° (a) through 60° (i).

immediately level and consistent over the rest of the experiment's run.

Note that as shown in figure 4.4, in each case following the initial spin-up, when the model has achieved quasi-equilibrium, the vertical gradient of the 2PVU contour of the Eulerian zonal average increases sharply on the equatorward flank of the jet, maintained by the blending of the relaxation states and the sharpening of the tropopause through the vortex erosion mechanism.

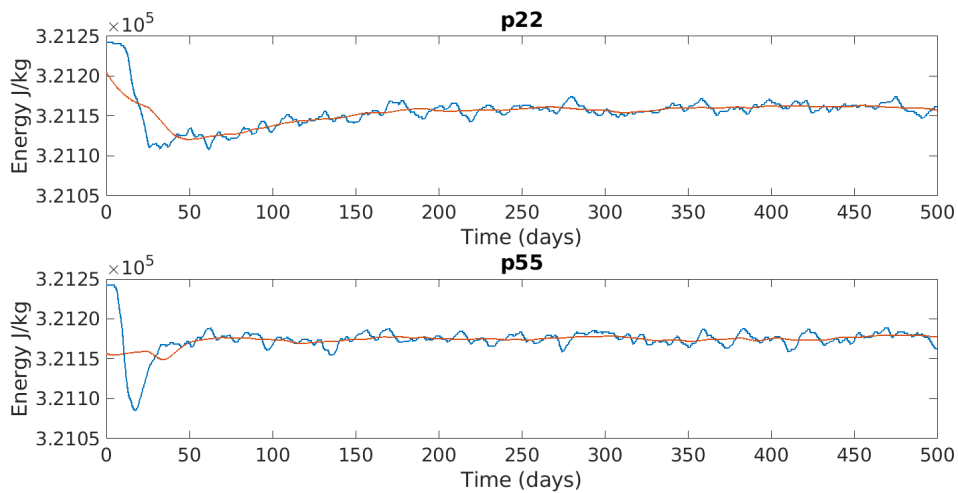


Figure 4.3: Globally integrated energy norms for experiments *p22* and *p55* (blue) with 50-day running means overplotted (orange).

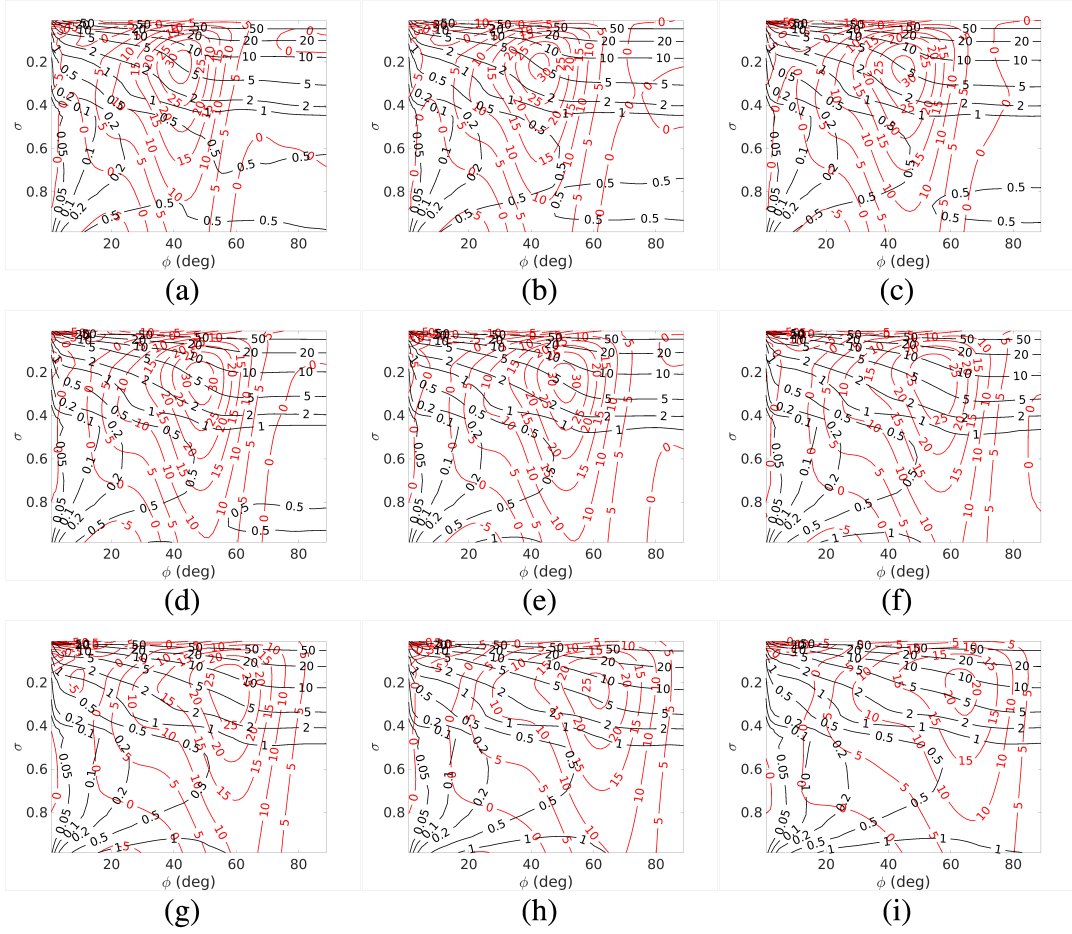


Figure 4.4: Time mean Eulerian Zonal average wind (red) and PV (black) averaged over model days 200-500 for relaxation latitudes 22°(a) through 60°(i).

4.5 Background States

The Modified Lagrangian Mean background state for each of the nine experiments was calculated over the 500-day window of each experiment, at each 6-hour interval, using 79 θ levels spaced 2K apart from 244K to 400K. The latitude grid used by the calculation is a gaussian quadrature grid which matches that used in the IGCM. This small $\Delta\theta$ is required (at least in the troposphere) in order to resolve the lower boundary, as calculation of the background state is easier when the sloping lower boundary in isentropic coordinates is well resolved.

There were a number of dates when the PV inverter failed to converge; these dates were filtered out by replacement, substituting the missing date with the previous timestep. The failure to converge is likely related to the change in position of the lower boundary as part of the outer iteration and the update of the lower boundary condition in the nested nonlinear iteration described in Chapter 2. Failures can occur if strong meridional gradients in θ on the lower boundary approach too near the equator where the elliptic

PV inversion problem is very stiff (due to the constraints on gradients in Montgomery potential approaching the equator). As before, in addition to replacement of failed dates by previous timesteps, methods were employed to filter out spurious solutions from the background state calculation using the PV inverter. The first is a filter removing any timesteps outlying in terms of global average surface pressure change over time. This filter was performed by inspection, though a first-guess algorithm is used to narrow the range down. This calculates $|dp/dt|$, the magnitude of the change in global average surface pressure in the background state at each point, and re-orders the points in terms of descending $|dp/dt|$ on a plot. As the points to be filtered are spuriously large, by normalizing the axis, and calculating the point closest to $(0,0)$, the place where the discrete step between spurious values and normal solutions may be found and adjusted by eye in order to filter states out as necessary.

The second filter removes any timesteps with integral isentropic density greater than $1600kgm^{-2}K^{-1}$. Both filters acted as markers for dates which had failed to converge to a reasonable state. Due to the higher (T85) resolution of these experiments, a smaller proportion ($\sim 10\%$) were filtered relative to the previous experiments.

4.5.1 Latitude of the Jet Maximum

The modified Lagrangian mean of the experiments shows similar quasi-equilibrium jet latitudes to the Eulerian zonal mean, though the MLM exhibits slower time variation. By calculating the mean latitude of the maximum of the jet after spin-up in both the zonal average and the MLM state (as well as the zonal average of the experiments with half the relaxation timescale) it is possible to examine the relationship between the latitude of the jet seen in the experiment and the latitude of the jet in the relaxation field used in that experiment (figure 4.5).

In each case, there is an approximately linear relationship between the latitude of the jetstream and the latitude of the relaxation state of the experiment. Despite the increased strength of the relaxation in the reduced relaxation timescale experiments, this linear relationship is maintained, with a slightly steeper gradient, as opposed to the jets being more closely confined to the relaxation state.

It is possible that the faster relaxation timescale's similar overall dynamics to the standard set of experiments is an indirect effect where the background state changes so that it is a little more unstable (closer to the relaxation state) and then eddy growth rate




Figure 4.5: Mean latitude of the jet maximum (*i.e.* maximum in zonal wind) for the MLM state (blue), Eulerian zonal average (orange), and the Eulerian zonal average of the reduced relaxation timescale experiments (yellow), versus the latitude of the jet in the relaxation state.

(associated with instability) increases to match the enhanced damping rate. As a result, a more dynamic equilibrium is attained with a balance between higher rates of wave activity generation and dissipation. This is the eddy saturation hypothesis and has been studied previously in the context of the ocean. Studies (Böning *et al.* (2008)) show that the Antarctic Circumpolar Current (ACC), which is driven by wind-stress also exhibits this eddy saturation behaviour. Increasing wind-stress in models of the ACC does not alter the mean flow, and instead increases eddy energy, as the energy of the mean flow is ‘saturated’ and released by eddies (Marshall *et al.* (2017)).

Examining the timeseries of global total internal + potential energy, and enstrophy, for the p40 relaxation jet (figure 4.6) it is clear that both are consistently higher in the shorter relaxation timescale experiment, indicating a larger amount of energy contained within the experiment eddies. This finding is consistent across latitudes. This evidence indicates that the eddy saturation hypothesis merits further investigation.




Figure 4.6: Time-series of both the global integral total internal + potential energy, as well as the root mean square vorticity (enstrophy), for the $\phi_{relax} = 40$ experiment with relaxation timescales $\tau_{relax} = 7.5days$ (blue, with 50 day running mean overplotted in orange), & $\tau_{relax} = 15days$ (yellow, with 50 day running mean overplotted in purple)

4.5.2 Speed

The speed of the jetstream is expected to be higher in the background state than in the Eulerian zonal flow in the presence of disturbances; this is true in the case of each experiment presented here. See [Brunet and Methven \(2019\)](#) for an argument of the inequality between u_0 & $[u]$. In addition, it is anticipated that the Eulerian zonal average will fluctuate with the amplitude of the waves in the timeseries. In each case, following the overturning of the initial wave, the jetspeed quickly attains its quasi-equilibrium value; the jets do not increase in speed during the northward migration described above, which is the case in each experiment.

For some other measures of wave activity, such as [Nakamura and Zhu \(2010\)](#)'s A^* wave activity, or previously introduced measures of wave activity, such as those derived from a Generalized Lagrangian Mean, non-acceleration theorems have been shown ([Solomon and Nakamura \(2012\)](#)) to indicate a relationship between the reference state, the wave activity, and the strength of the full flow;

$$[u] + A \approx u_0 \quad (4.6)$$

However, this relationship is a result of the specific choice of wave activity quantity. In the case of [Nakamura and Zhu \(2010\)](#), the relationship holds for a single-layer fluid studied in that paper, however, this is difficult to generalize to the full 3D case. Examination of the results from the experiments, such as those shown in figure 4.7, provides a slightly

different relation;

$$u_0 - [u] \geq A \geq 0, \quad (4.7)$$

such that the Wave Activity sets a lower bound on the difference between the background and zonal average flow.




Figure 4.7: A comparison of the maximum zonal angular velocity from the Modified Lagrangian Mean (blue) state and the Eulerian zonal average (red), their difference (purple), and the global pseudomomentum (yellow), all measured in degrees per day, for the p45 experiment.

4.5.3 Modified Lagrangian Mean Background State Structure

The timeseries of the equilibrium background flow of all experiments is relatively steady. The zonal flow of the background is steadier and more consistent than in the zonal average, as seen in figure 4.7. Examining the gradient of theta on the lower boundary (figure 4.8), there are several ‘tongues’ of sharp theta gradient which migrate equatorward at similar rates in each case; when these structures reach the equator they are associated with timesteps where the PV inverter used to calculate the background state fails to converge.

The latitude of the tropopause at 2PVU remains steady throughout the timeseries, however the meridional gradient of PV, shown at the 315K isentropic level in figure 4.8, shows clear episodes of poleward migration and then days when the gradient is large at two latitudes - indicative of ‘split jet’ states.

This split in PV gradient can be explained more clearly by examining a cross-section of background PV at an appropriate point in time (figure 4.9). There are two distinct

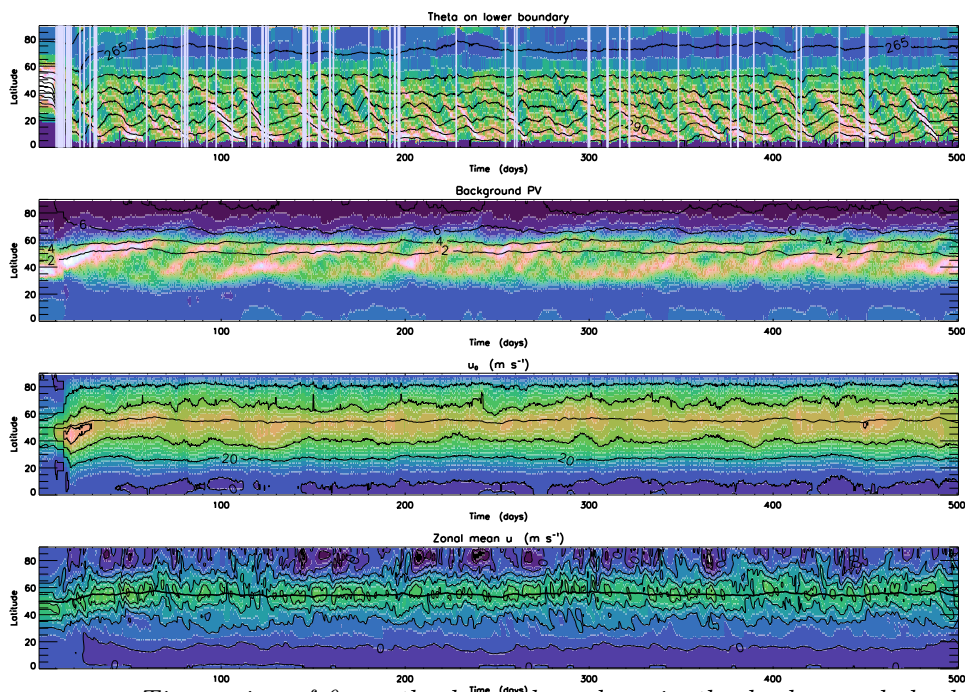


Figure 4.8: Timeseries of θ on the lower boundary in the background, background PV, background wind, & Eulerian zonal mean wind at all latitudes on the 315K θ surface (with the exception of the first plot), with meridional gradients shown in colour, for the p45 experiment. 315K is chosen due to its intersection with the 2PVU tropopause near the latitude of the jet maximum. Vertical lines on the lower boundary θ plot indicate days where the calculation of the background state did not converge and as a result data was replaced by the previous timestep.

slopes in PV, at 60° and 35° , which match both to an upwards slope of the tropopause on the equatorward flank of the main jet as well as further south of the jet seen both here and in figure 4.4 in the time-mean-zonal-mean wind. We expect latitudes with strong PV gradients to be associated with stronger zonal flow through the properties of PV inversion (e.g., the PV staircases described by Dritschel and McIntyre (2008)). The strong meridional PV gradient at the tropopause and above at around 60°N is associated with a maximum in the MLM zonal wind at the same latitude. The second region of stronger PV gradient on the edge of the subtropics (at about 35°) is associated with a plateau in the MLM zonal flow strength (although not quite enough to produce a second maximum). Similar strong winds can be seen in the zonal mean, zonal wind at this time. However, the magnitude of the zonal mean wind is much smaller than the MLM zonal flow because the wave activity has large amplitude at this time (and this results in a much smaller zonal mean for reasons given above - 4.5.2). The source of this upper-level jet is possibly the high tropospheric winds in the relaxation state due to the blend with Held-Suarez, or else a distinct sub-tropical jet forced polewards by the relaxation.

Comparing the background wind with the Eulerian zonal average by taking figures 4.7, 4.8, and 4.9 into account, the background wind is more coherent; despite the strong

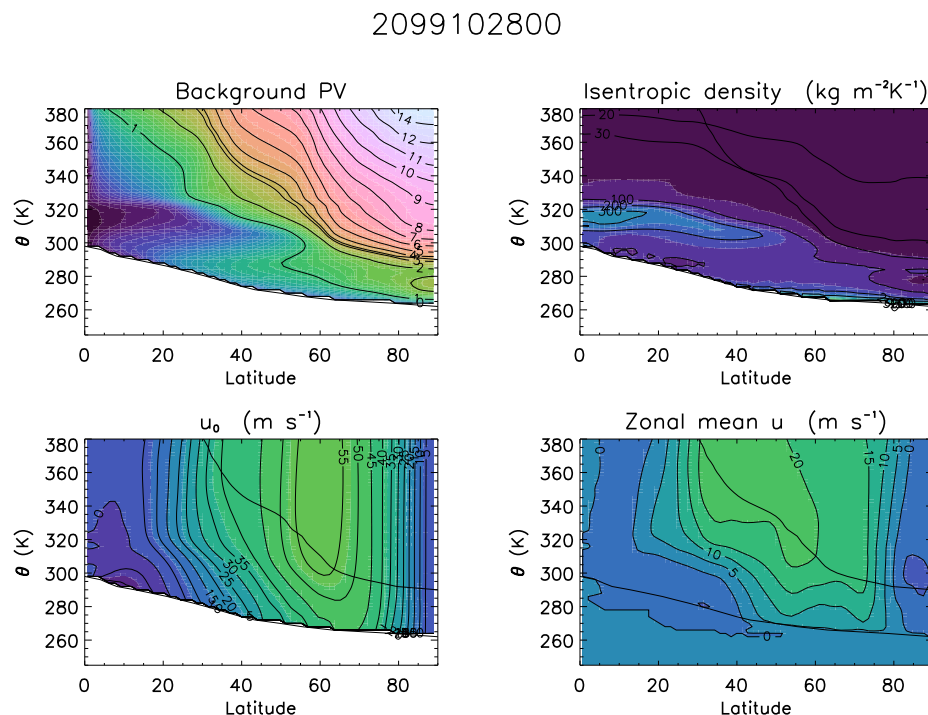


Figure 4.9: Cross-sections of background PV, background isentropic density, background zonal flow, and Eulerian zonal mean from from day 300 of the p45 experiment, chosen as a ‘typical’ timestep. The background state was not averaged over time for this plot due to its slow variation throughout the timeseries.

similarities in latitude when taking the maximum, it is clear that the MLM is in fact a more robust measure of jet latitude due to its simpler structure. Of note is that these experiments do not have a seasonal cycle or the meridional gradient in solar forcing that results in a strong Hadley circulation. Therefore they do not have a distinct subtropical jet that is prominent in global reanalysis data (Zolotov et al. (2018)). Therefore, a single jet is dominant most of the time in all the experiments (although the zone with highest meridional PV gradient can split on occasion associated with Rossby wave breaking and wave-mean flow interaction).

4.5.4 Global Wave Activity

In each experiment, total wave activity increases, as expected, during the initial overturning wave resulting from the baroclinic instability of the initial state used, before settling at an equilibrium after ~ 50 days, as seen in the timeseries shown in figure 4.10. The modal growth stage of the lifecycle can be seen by the increase of intersection domain wave activity (i.e. low-level WA, \mathcal{P}_d & \mathcal{P}_b) which peaks around day 20. At this point, the low-level baroclinic waves become nonlinearly saturated and thus the wave-activity at these levels begins to decay. The higher level wave activity continues to grow for several more days before peaking. This behaviour is seen in both the p45 experiment (figure 4.10) and the p50 experiment. Across all experiments, wave activity is settled down to an equilibrium value before the MLM jet completes its poleward migration to equilibrium, as these migratory jets nonetheless support waves. In experiments with a relaxation latitude lower than 35° , there is no large peak in wave activity. Boundary and exterior terms \mathcal{P}_e & \mathcal{P}_b terms are negative due to the poleward decrease in potential temperature at the lower boundary of the background state, and are an essential part of baroclinic growth of the disturbances that continue to grow and decay on the jet. The initial overturning increases global WA until it reaches the equilibrium maintained for the remainder of the timeseries.

The tropospheric interior domain pseudomomentum is by far the largest contribution to the total wave activity, followed by \mathcal{P}_d , the intersection domain wave activity. Both gravity wave and stratospheric terms are very small, with little variation or contribution throughout the entire timeseries, as expected. Both exterior and boundary pseudomomentum terms contribute a small but not insignificant amount of negative wave activity.

Further to this, there is strong anticorrelation between \mathcal{P}_d and \mathcal{P}_b and also $\mathcal{P}_b + \mathcal{P}_e \approx -\mathcal{P}_d$. This is a signature of baroclinic waves and fluctuations in amplitude associated

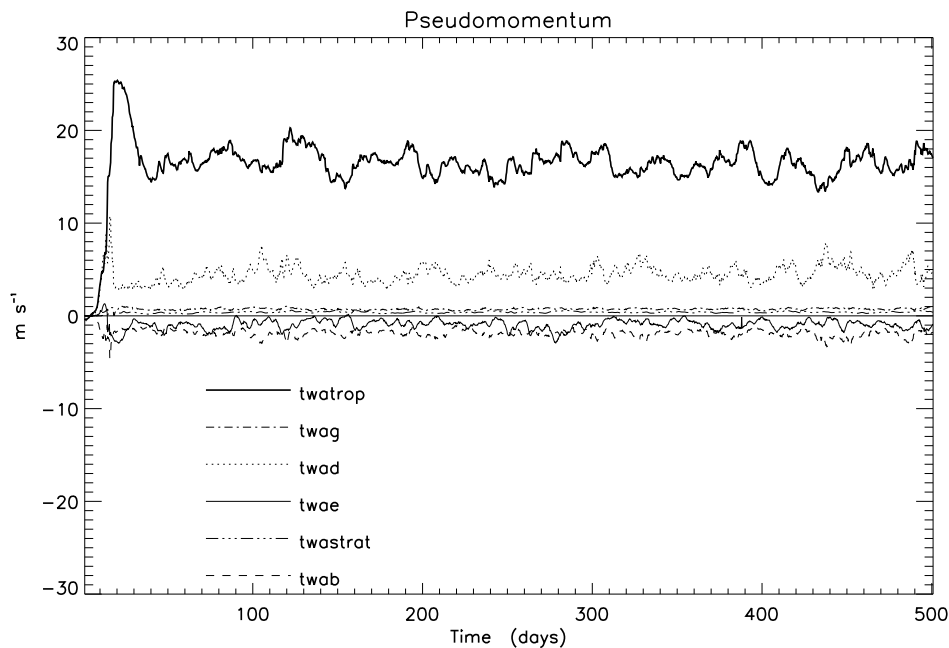


Figure 4.10: Timeseries of globally integrated wave activity (Pseudomomentum) for the *p45* experiment. The thick line represents the tropospheric contribution of the interior Rossby-wave WA term \mathcal{P} (*twatrop*), the triple-dashed line represents total stratospheric WA (*twastrat*); all other lines denote various contributions to the total WA. The dot dash line represents the gravity wave contribution \mathcal{P}_g (*twag*), the dotted line represents the background state's intersection domain WA \mathcal{P}_d (*twad*), the thin solid line represents the WA in the exterior domain \mathcal{P}_e (*twae*), and the dashed line with wide spacing represents the boundary contribution from the interior domain \mathcal{P}_b (*twab*).

with behaviour similar to growing normal modes and their nonlinear life cycles. In a growing normal mode, it is the case that global $\mathcal{P} = 0$ (section 1.2.3, Chapter 1) and therefore the negative boundary contributions are equal and opposite to the positive interior pseudomomentum. In the experiments shown here, the interior p-mom is much larger and this has also been noted with global analysis data (e.g., Methven (2013)). This indicates that there is significant wave activity in addition to unstable baroclinic waves.

4.6 Conclusions

The aim of this chapter was to design a set of model experiments with a small number of control parameters, describing latitudinally-confined jetstreams. These experiments are presented above, showing that the method used met the criteria set out at the start. A small number of control parameters were used; a parameter to control the latitude of the jetstream in the experiment *p* (corresponding to the relaxation latitude of the jet as

shown in equation 4.3), a parameter to control the strength of the relaxation τ_{relax} , a parameter controlling the magnitude of the zonal wind speed U_0 , the blending parameter α controlling the ratio of the mix between the balanced temperature of the prescribed jet and the HS state; the parameters involved in the control of the HS state, and the change in static stability across the tropopause as determined in the jet state prescription.

Both the variability of the background state within each experiment, looking at the spin-up time and differences between the Eulerian zonal mean and the MLM, as well as the intervariability between different experimental runs have been explored. We see that increasing the latitude of the jet maximum of the wind-field, corresponding to the balanced relaxation temperature field, increases the latitude of the quasi-equilibrium state of the jet after spin-up. These two quantities (ϕ_{relax} & ϕ_{eqm}) display a linear relationship with a constant offset in both the Eulerian ZA and the MLM.

The variability of the jets in latitude increases with the latitude, with jets further south displaying more tightly confined meridional structure, whereas jets further poleward exhibit larger variability, likely as relaxation forcing is weaker when the equilibrium jet is closer to the relaxation state, such as in the higher latitude experiments. A series of experiments was run where initial background state was made different from the relaxation state by the additional of a barotropic component to the flow associated with a meridional pressure gradient. The lack of dependence of the equilibrium behaviour of the experiments on the initial perturbations was shown. The equilibrium behaviour was also shown to be insensitive to the timescale of the applied relaxation. When the relaxation rate is increased (i.e. a shorter relaxation timescale), the quasi-equilibrium jet latitude remains similar (though a little closer to ϕ_{relax} at lower latitudes), but the eddy amplitude, as indicated by global enstrophy (figure 4.6), and the total internal + potential energy is also slightly higher.

The wave activity in the experiments is large amplitude; the pseudomomentum is as much as half of the zonal average zonal flow (figure 4.7). The circulation around the tropospheric vortex is strong, as indicated by the maximum in the MLM zonal flow, but as the disturbances are large amplitude, the maximum zonal mean, zonal flow is only just over half of the MLM flow. This is consistent with the magnitude of the global pseudomomentum. In addition, a partition of terms contributing to pseudomomentum reveals that there is sustained baroclinic wave activity (associated with negative pseudomomentum on isentropic surfaces that intersect the lower boundary) and strong correlation in the amplitude of the interior and boundary activity, which is behaviour associated with

baroclinic wave growth (and decay). However, there is large interior pseudomomentum that is associated with Rossby waves that are not described by growing baroclinic wave structures.

In the next chapter, the Empirical Normal Modes of these experiments will be calculated in order to examine the effect of jet latitude on the large-scale modes of the model runs.

Chapter 5:

Slow Modes of Variability and Dependence on Jet Latitude

5.1 Introduction

In global atmospheric flows with large amplitude wave activity, the background flow and the eddies propagating relative to that background are strongly coupled and their partition depends on the theoretical and diagnostic framework used. However, [Barnes et al. \(2010\)](#) (see Chapter 1 for a more detailed discussion) have shown from the statistics of reanalysis and simulation data that the frequency of dominant modes of variability may depend on the structure of the background state jet, especially its latitude. In this chapter, we use the ENM technique (Chapter 2) to find the dominant modes of variability in the global IGCM experiments of Chapter 4 where the latitude of the quasi-equilibrium jet can be controlled primarily through the variation of a single parameter. The ENM technique not only extracts the structures of the dominant modes, for each zonal wavenumber, but also predicts their intrinsic phase speeds (and therefore frequencies) from their spatial structures using the ratio of pseudoenergy to pseudomomentum (see Chapter 3 for tests of this technique using global simulations). This is used to test the following hypotheses:

1. The variability associated with large-amplitude disturbances on the jet can be decomposed into a small set of dominant modes with distinct structures and frequencies.
2. The dominant modes vary continuously with jet latitude and there is a relationship between dominant mode frequency (for a given zonal wavenumber) and jet latitude

In this chapter, the above hypothesis are addressed. The ENM calculation is per-

formed, and the resultant modes are paired and sorted. A small number of modes dominate the variability. The relationship of these dominant modes to the latitude of the jetstream is explored, and a set of stationary dominant modes are identified. In particular, the focus on slow-moving structures, and understanding the factors that contribute to their dynamics, is motivated by a desire to better understand the causes of extreme weather events which, as described in Chapter 1, can often be attributed to persistent, slow-moving waves.

5.2 ENM Calculation

The empirical normal modes for the jet-control experiments were calculated over a window of 400 days, starting at model day 100 to allow for spin-up as discussed in chapter 4. The analysis used 6-hourly data, and a time-averaged background state spanning the window is calculated using the modified Lagrangian mean for each instant. A set of ENMs is calculated for each zonal wavenumber 1 to 12; smaller scales were not chosen in order to focus on the largest-scale structures. Once all modes have been calculated, all modes with pseudomomentum less than 1% of the leading mode of each spectrum are neglected in order to limit the number of eigenvectors returned for use in analysis.

5.3 ENM Mode Distribution

Upon calculating the ENMs, it is possible to see that the greatest amount of variability is contained in only the first few modes by examining the mode distribution ordered by amplitude for a single zonal wavenumber, such as the $m = 5$ distribution for the p30 experiment shown in figure 5.1. This distribution is distinct from the ENM power spectra shown in figures 3.11 through 3.16- here we examine the variation in eigenvalues, phase speed, and pseudoenergy components for a single set of ENMs at one scale (i.e. one zonal wavenumber) as opposed to comparing the dominant modes across scales. The eigenvalues descend in paired steps (as expected), up until mode 7 and beyond, indicating the presence of three propagating pairs of ENMs.

Total pseudoenergy over pseudomomentum represents the intrinsic phase speed of the modes, c_{int} . In all spectra, the interior Rossby wave component of the pseudoenergy \mathcal{H}_w provides the largest westerly component of the mode pseudoenergy and is balanced by the energy components, \mathcal{H}_{KE} and \mathcal{H}_{APE} , propagating the structures against advection by the flow. The exterior component of pseudoenergy provides a small but not negligible




Figure 5.1: *ENM Mode distribution for zonal wavenumber $m = 5$ in the p30 experiment. The blue line corresponds to the phase speed of the mode, $-\mathcal{H}/\mathcal{P}$. The following lines correspond to contributions to the pseudoenergy: \mathcal{H}_{KE} (orange), \mathcal{H}_{APE} (with the interior contribution in yellow, and the APE on the top boundary in purple), \mathcal{H}_w (green), \mathcal{H}_e (cyan). The eigenvalues (red line) have been normalized and scaled to the “Doppler term”, \mathcal{H}_w . Finally, the total pseudomomentum \mathcal{P} is shown in blue, and is ≈ 1 for all modes in the spectrum.*

easterly contribution to intrinsic phase speed.

5.4 Sorting Modes

Using the ENM technique allows for a reduction in dimension of the input timeseries dataset. In order to do this however, a choice needs to be made about which modes are retained (and more importantly, analysed). Often when undertaking principal component analysis, only the leading few modes are retained. However, it is important to ensure that we retain all modes which might possess some physical significance. The section below details how the ENMs are selected and sorted for analysis.

5.4.1 Pair-matching

As the ENM technique is an eigenvalue decomposition of the pseudomomentum covariance matrix, this decomposition aims to orient the data to maximise variance in pseudomomentum in the leading modes. As a result, the output of the technique is ordered in terms of eigenvalue (i.e. the square root of that mode's pseudomomentum), ordered from large to small. While this amplitude-sorting is relevant in that modes which explain more variance are considered first, it is essentially arbitrary in that the modes are calculated from the dataset as opposed to dynamical normal modes, which are calculated from the eigenvector problem applied to the linear propagator matrix of the dynamical model.

However, in this analysis we are attempting to reconstruct (and hence simplify) information about the dynamics of the system from which the data comes. In order to do this, we attempt to match the ENMs' output into pairs whose amplitudes vary in quadrature, which are able to describe a propagating disturbance. If the variability is dominated by these propagating modes, then it is expected that the pairs extracted will possess the same pseudomomentum amplitude structure, but with phases in quadrature (like a sin/cosine pair, or complex number representation of a propagating wave). If the modes found are purely propagating, then the two ENMs must have equal eigenvalues, such that they possess an identical structure in amplitude with only phase differing between the two, offset by $\pi/2$. As such, it is expected that the eigenvalue spectrum (organized in terms of amplitude) for a given zonal wave number is paired into steps until the part of the spectrum constructed of "noise" associated with non-modal structures; this can be seen, for example, in the ENM spectrum for $m = 5$ in the p30 experiment (5.1).

A match statistic, M_q (equation 2.59 in Chapter 2) measures the quality of the quadrature match of any two ENMs on the spectrum; an example set of matching statistics is shown in figure 5.2 for zonal wavenumber $m = 5$ in the p30 experiment, showing that

modes $i = (1, 2)$, $i = (3, 4)$, & $i = (5, 6)$ (the bright yellow squares) display a strong quadrature match, corresponding to the steps seen in figure 5.1.




Figure 5.2: The matching statistics m_{qi} for modes $i = 1 - 23$ from the $m = 5$ ENM spectrum of the p30 experiment. Both x and y axis represent the same spectrum, with each square representing the match between the modes $i = (x, y)$, resulting in the plot being mirrored along the $y = x$ line.

Generally, these pairs are also the ‘leading modes’ (i.e. the modes with the highest amplitude in pseudomomentum) of their respective ENM spectra, however this is not always the case. Figure 5.3 shows that the vast majority of modes selected by optimising for matching-statistic (described in Chapter 2) are indeed the leading two modes in terms of amplitude from their respective ENM mode spectra. However, for larger, planetary scale waves (i.e. those shown in figure 5.3 with zonal wavenumber less than 4), there are a small number of cases where the ENMs with the strongest quadrature match are not the ENMs with the highest amplitude. These cases are also generally less-well matched implying that at this scale, the ENM technique is less adept at finding propagating Rossby wave structures.




Figure 5.3: Matching statistic m_q (coloured shading) for the best-matched ENM pair plotted against zonal wavenumber and relaxation latitude ϕ_{relax} . Each value is labelled by order in terms of amplitude.

5.4.2 Structure Sorting

When we examine these pair-matched modes, it is possible to examine their pseudo-momentum structures in latitude and potential temperature, in order to compare modes with similar dynamics across experiments. The structures of high-variance pairs in the ENM mode spectra are strictly orthogonal to one another *within* a given experiment and zonal wavenumber, (with intra-pair variability dominated by the variation in phase, i.e. the quadrature separation in phase between pair members), and so one way for different ENM pairs to have an orthogonal structure is to have different wave-like variation in the meridional and/or vertical directions. For the dynamics of a fluid on a plane, this would entail different meridional wavenumbers. On a sphere it can be achieved through projection on different spherical harmonics with latitudinal variation given by associated Legendre functions. Given the large number of degrees of freedom for the ENM structures

(at T85 L30 truncation), we characterise the meridional structure of the leading pairs by a meridional node number determined by visual inspection (the count of pseudomomentum maxima in the meridional direction in the troposphere). For example, for the ENMs in figure 5.4 the meridional node number of 1, 2, 3 and 4 in panels (a)-(d).

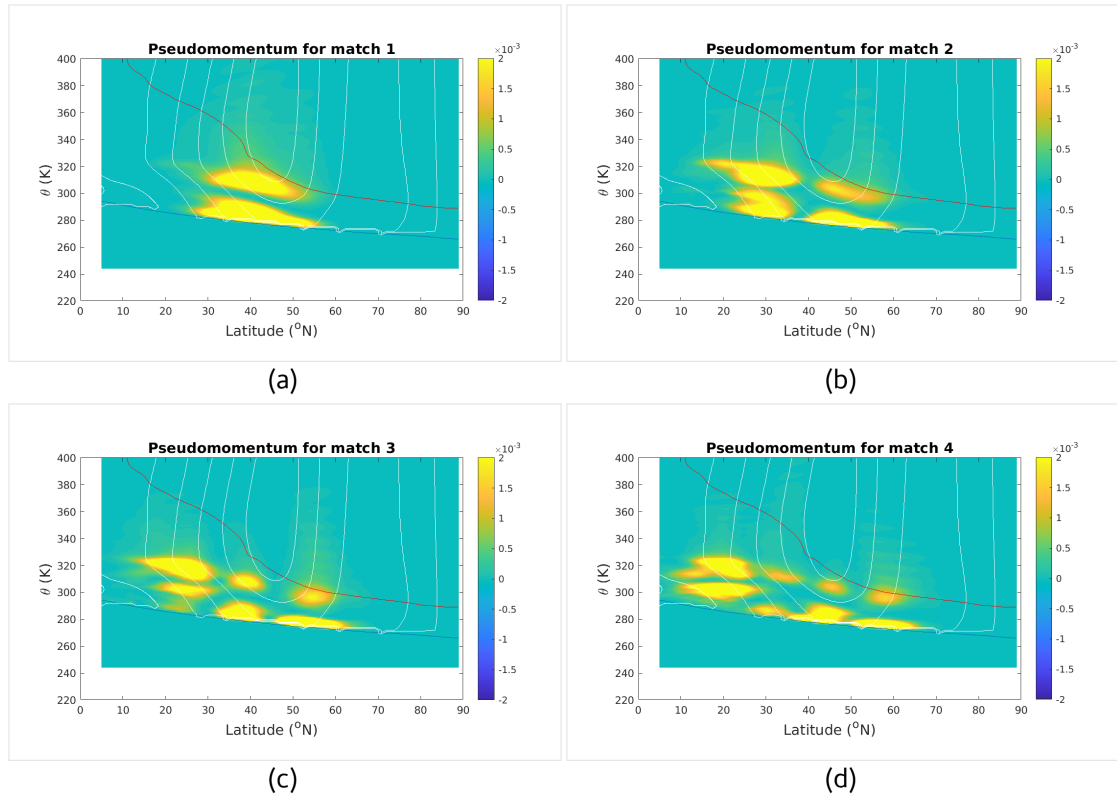


Figure 5.4: The pseudomomentum amplitude structure (in ms^{-1}) for the first member of each of the four first leading pairs of ENMs in the $\phi_{relax} = 30^\circ$ experiment for zonal wavenumber $m = 6$. As ENM pair number increases ((a) \rightarrow (d)), so does meridional node number. The background state 2PVU line is overplotted in red, and the background state wind field is overplotted in white (intervals of $10ms^{-1}$).

Figure 5.5 shows the meridional node number for the first four propagating ENM pairs (ordered by the quality of the match between them). By construction, only modes above a certain match-threshold (0.7) are retained, however a single match with a lower matching-statistic are also retained if there is no suitable match in that ENM mode spectrum. In all cases, the leading matched pair for wavenumbers four and higher possesses a simple $l = 1$ node structure. Larger, planetary-scale waves generally have more complicated structures often taking the shape of a varicose mode (Leib and Goldstein (1989)) which straddles the lower flank of the jet. Some matches possess non-integer node numbers; nodes were determined by visual inspection, and as non-integer node numbers were chosen to represent a number of situations. Firstly, where the structure was ambiguous and not all parts of the structure were separated into a clearly discrete number of nodes that could

be distinguished by eye. In these cases, one or two distinct nodes were present alongside the ambiguous structures and as such non-integers were used to best approximate the node number. Additionally, for some pairs, the node number was clear and discrete but different between the members of the ENM pair, and the average of their node numbers was used.

Of note is the vertical structure of the modes presented. All modes possess structure with vertical node number two, and as such vertical structure was not considered for sorting the modes. These nodes correspond to the CRW structure presented by [Heifetz et al. \(2004a\)](#). Nonetheless, the maxima near the lower boundary and in the upper troposphere are suggestive of counter-propagating Rossby wave structures coupled in a baroclinic mode structure as discussed in Chapter 1 section 1.2.3. In each case (both here and pseudomomentum amplitude structures shown subsequently in this chapter such as in figure 5.12) the upper node is only near the tropopause at the jet maximum and polewards of this. On the equatorward flank of the jet the upper mode appears to be a similar distance above the lower boundary, but below the tropopause (as expected for Charney modes for example, [Heifetz et al. \(2004b\)](#)).

By selecting the pair with the lowest available node structure at each point in zonal wavenumber/jet latitude space, we construct a set of ENMs which;

1. Consist of propagating pairs
2. Possess similar spatial structures across experiments & zonal wavenumbers
3. Have the least complex structures

Examining figure 5.5, it can be seen that relatively few changes are present in this set relative to the set of modes with the highest matches shown in the top left panel of figure 5.5. No modes for zonal wavenumbers 4 and above are filtered. In fact, only Wavenumber 3 of the p25 experiment and wavenumber 2 of the p50 experiment are replaced, both by the second match.




Figure 5.5: Pseudocolor plot showing the number of meridional nodes in the pseudomomentum structure of leading quadrature ENM pairs 1 (top-left), 2 (top-right), 3 (bottom-left), and 4 (bottom-right), in each case over each experiment (x -axis) and over each zonal wavenumber (y -axis). Points where there is no match above the minimum matching statistic threshold are greyed out.

5.4.3 Phase Speed

With modes above the amplitude & matching statistic cut-off now arranged in terms of their meridional structure and quadrature matching statistic, we can examine the intrinsic phase speeds associated with the structures of each ENM, calculated as;

$$c_{int} = -\frac{\mathcal{H}}{\mathcal{P}}. \quad (5.1)$$

Examining c_{int} across wavenumber and experiment (figure 5.6) there are several features of note. Firstly, it is possible to identify two bands where phase speed is approximately zero; the first of these ranges from zonal wavenumber $m = 5$ in the more equatorward relaxation scenarios, $m = 4$ in the mid-latitudes, descending to $m = 3$ at the

most poleward latitudes. This ‘branch’ of similar structures I will call the quasi-stationary branch (QSB). An initial cursory glance at figure 5.6 would imply that the lowest speeds in p55 & p60 are at lower wavenumbers; however, in order to examine a consistent set of structures across experiments, $m = 2$ in experiment p55 & p60 were avoided due to their higher meridional node structure. This choice is validated by further investigation shown in the next section (5.5) which shows that classifying the QSB as $m = 3$ in these makes the most sense when considering components of pseudoenergy. The second set of stationary dominant modes appears for $m=2$ at mid-latitude ϕ_{relax} , and notably has a distinct structure from the quasi stationary branch.




Figure 5.6: *Intrinsic phase speed across zonal wavenumber and experiments, here with symbols overlaid indicating the number of peaks in meridional structure for each data point; black circles, blue triangles and cyan crosses respectively indicate meridional node numbers 1, 2 and 3. The quasi-stationary branch is outlined in red.*

Beyond the QSB, structures at shorter scales (higher wavenumbers) possess higher c_{int} ; however, at these scales this measure of speed is overestimated relative to c_{emp} as shown in Chapter 3 due to structures with a given intrinsic phase speed being advected by the jet (see Chapter 3). Comparing 5.6 with a similar plot instead showing the empirically

derived measure of phase speed (figure 5.7(a)), we see that the overall distribution of speeds is similar, and the difference (figure 5.7(b)) is mostly in its magnitude, as well as a ‘lobe’ of high speeds in c_{int} seen at $\phi_{relax} \sim 30$ to 45 for the very highest wavenumbers.




Figure 5.7: (a) Empirical phase speed c_{emp} across zonal wavenumber and experiments, with the quasi-stationary branch outlined in red. The pairs shown here match those in figure 5.6. (b) the intrinsic phase speed c_{int} minus c_{emp} for the same set of modes. Symbols in each box refer to the meridional node number as in Fig. 5.6.

Recalling equation 5.1, the intrinsic phase speed c_{int} is made up of contributions from the pseudomomentum and pseudoenergy. Due to the orthogonality imposed by the ENM technique, the interior pseudomomentum $\mathcal{P}_{interior} = 1$, and the other contributions to \mathcal{P} are small, making \mathcal{P} near to 1, as seen in figure 5.1. As such, insight into the intrinsic phase speed can be gained by partitioning the \mathcal{H} terms (which are all individually divided by \mathcal{P} to obtain c_{int}). It is possible to divide the contributions to \mathcal{H} into those which advect modes along the jet, and energy terms which contribute to propagation against it;

$$\mathcal{H} = \underbrace{\mathcal{H}_w + \mathcal{H}_e}_{\text{Advection Terms}} + \overbrace{\mathcal{H}_{KE} + \mathcal{H}_{APE} + \mathcal{H}_b}^{\text{Propagation Terms}} \quad (5.2)$$

Negative terms in \mathcal{H} provide positive contributions to the phase speed, hence ‘‘advection’’ terms, whereas positive terms \mathcal{H} instead serve to reduce the phase speed, hence ‘‘propagation’’ terms which serve to move waves against the flow. This partition mirrors the dispersion relation, equation 1.6 of Chapter 1, with wave speed dictated both by advection by the basic state, and propagation against the flow. Figures 5.10 and 5.11 show, respectively, the advection and propagation pseudoenergy values of the selected modes at each zonal wavenumber, across all the experiments. In figure 5.11 it can be seen that the

mode with the greatest propagation in each experiment is that identified with the QSB. Near-zero phase speed is achieved in the QSB because these modes have sufficient energy to maintain position against advection by the jet (which can be seen for the QSB in figure 5.10).

Looking at modes not on the QSB, the propagation energy is generally strongest for the mid-latitude jets and weaker for jets established further south and north. Zonal wavenumber 2 stands out as a notable feature, as between $\phi_{relax} = 22^\circ$ and $\phi_{relax} = 45^\circ$, it shows both a higher energy than adjacent wavenumbers, and a more complex meridional structure.

The advection speed peaks for the mid-latitude jets at a given zonal wavenumber. As shown in figure 5.8 for $m = 10$, this is due to the relative positions of the wave activity and the core of the jet. The largest contribution to the advection, H_w/\mathcal{P} , the Rossby wind-weighted term represents the background zonal flow weighted by the wave activity. As such, the magnitude of this term is controlled by the cross-section of the background zonal wind, and the location where the wave activity amplitude is greatest.

While the jet's latitude increases across experiments, the location where the wave activity is strong moves relatively little, and as a result the cross section of the wave activity and the jet changes. In the experiments with the lowest ϕ_{relax} , the core of the jet is initially above the strongest wave activity (5.8 (a)), until ϕ_{relax} is in mid-latitudes (b), where the core of the background state jet has descended, but its poleward position relative to the wave activity means that it only partially intersects the jet. When ϕ_{relax} is in high latitudes, the jet is far enough poleward that it does not intersect the wave activity significantly at all, despite its continued descent. Thus, the ϕ_{relax} at mid-latitudes represents the greatest advection of the waves by the jet due to this peak in intersection.

Additionally, it can be seen in figure 5.10 that from experiment $\phi_{relax} = 30$ upwards, the advection speed gets stronger as wavenumber reduces from 9 down to the QSB across




Figure 5.8: Wave activity structures for $m = 10$ in experiments (a) $p22$, (b) $p40$, and (c) $p60$, with time-mean background zonal flow overplotted in white, and the 2PVU line in red, and the lower boundary in blue.

all experiments. A similar argument can be made to that of the variation across latitude. As wavenumber increases, as shown for the $\phi_{relax} = 40^\circ$ experiment in figure 5.9, the position of the strongest wave activity decreases in latitude. This does not result in a strong variation in advection speed in the experiments with lowest phi_{relax} as the core of the jet is above the wave activity, and thus the intersection of WA and the jetstream does not change. However, in the experiments with $\phi_{relax} = 30^\circ$ and upwards, the core of the jet has descended sufficiently that the equatorward change in position of the wave activity at higher wavenumber does serve to decrease the strength of the advection.

The propagation terms (figure 5.11) show most clearly the dynamical origin of the QSB; as the branch is a local maxima in wave energy, which acts to counteract advection by the jet, thus keeping the waves stationary. In each experiment, the ENM energy and therefore propagation rate increases as wavenumber decreases from 12 to the QSB where it attains a maximum. For each zonal wavenumber there is also a weak maximum in energy as jet latitude is varied, being slightly weaker for the lowest and highest jet latitudes. There is less variation at smaller scales than in the advection terms, however propagation is noticeably weaker at the highest latitudes. This effect is insufficient to counteract the weakening propagation at high latitude and maintain wave speed, however.




Figure 5.9: Wave activity structures for the $p40$ across wavenumbers (a) 5, (b) 7, and (c) 9, with time-mean background zonal flow overplotted in white, the $2PVU$ line in red, and the lower boundary in blue.




Figure 5.10: Advection ($\mathcal{H}_w/\mathcal{P}$) and exterior $\mathcal{H}_e/\mathcal{P}$ term contributions to phase speed (c_{int}) across zonal wavenumber and experiments. Symbols in each box refer to the meridional node number as in Fig. 5.6.




Figure 5.11: Propagation term contributions (i.e. $\mathcal{H}_{KE}/\mathcal{P}$, $\mathcal{H}_{APE}/\mathcal{P}$ & $\mathcal{H}_b/\mathcal{P}$) to phase speed (c_{int}) across zonal wavenumber and experiments. Symbols in each box refer to the meridional node number as in Fig.5.6.

5.5 Quasi-Stationary Branch

The structure of the pseudomomentum of the ENMs in the QSB are shown in figure 5.12 for each experiment. While the relative position of the WA and the jet changes across latitudes the individual structures remain similar. Each possesses a bimodal vertical distribution, with wave activity along the lower boundary, as well as near the tropopause near the latitude the jet core; this lobe of WA extends equatorwards on isentropic surfaces in the upper troposphere. All modes in the QSB possess a single dominant node in the meridional direction close to the jet core, although experiments with $\phi_{relax} = 50^\circ$ and above also possess a second much weaker node of wave activity in the subtropics, far from the jet.

These structures possess similar shapes, and can be identified as a similar mode changing continuously as the latitude of the background state jet is varied. While the

5.5. QUASI-STATIONARY BRANCH

relation between the wave activity and the jet does change slightly, as across the lower jet latitude experiments the core of the jet descends to a lower isentropic level, and in higher jet latitude experiments, the wave activity is on the equatorward flank of the jet. Nonetheless at the highest latitudes, (figure 5.12 (g) \rightarrow (i)) an additional ‘lobe’ of wave activity does appear at the surface in the tropics; in addition the tail found in the structure aloft becomes more pronounced. These changes in structure accompany changes in the dynamics of these modes.




Figure 5.12: *The pseudomomentum structures for the leading mode of each member of the QSB, from p22 (a) through to p60 (i). The average background zonal wind of each experiment is overplotted in white, the tropopause (2PVU line) in red & the surface in blue.*

The phases of these ENM structures, shown in figure 5.13 also display continuity across latitudes, however this is somewhat more difficult to see. For each mode in the QSB, phase is coherent across the main node of wave activity, which is not the case for other ENMs with more complex meridional structure. Phase also changes markedly between the lower and upper maxima in wave activity, indicating the phase difference in

5.5. QUASI-STATIONARY BRANCH

PV between the upper and lower counter-propagating Rossby wave. For the high ϕ_{relax} p50, p55, and p60 states, it is clear that the weak subtropical lobe of wave activity has a phase difference from the main node of wave activity on the jet.

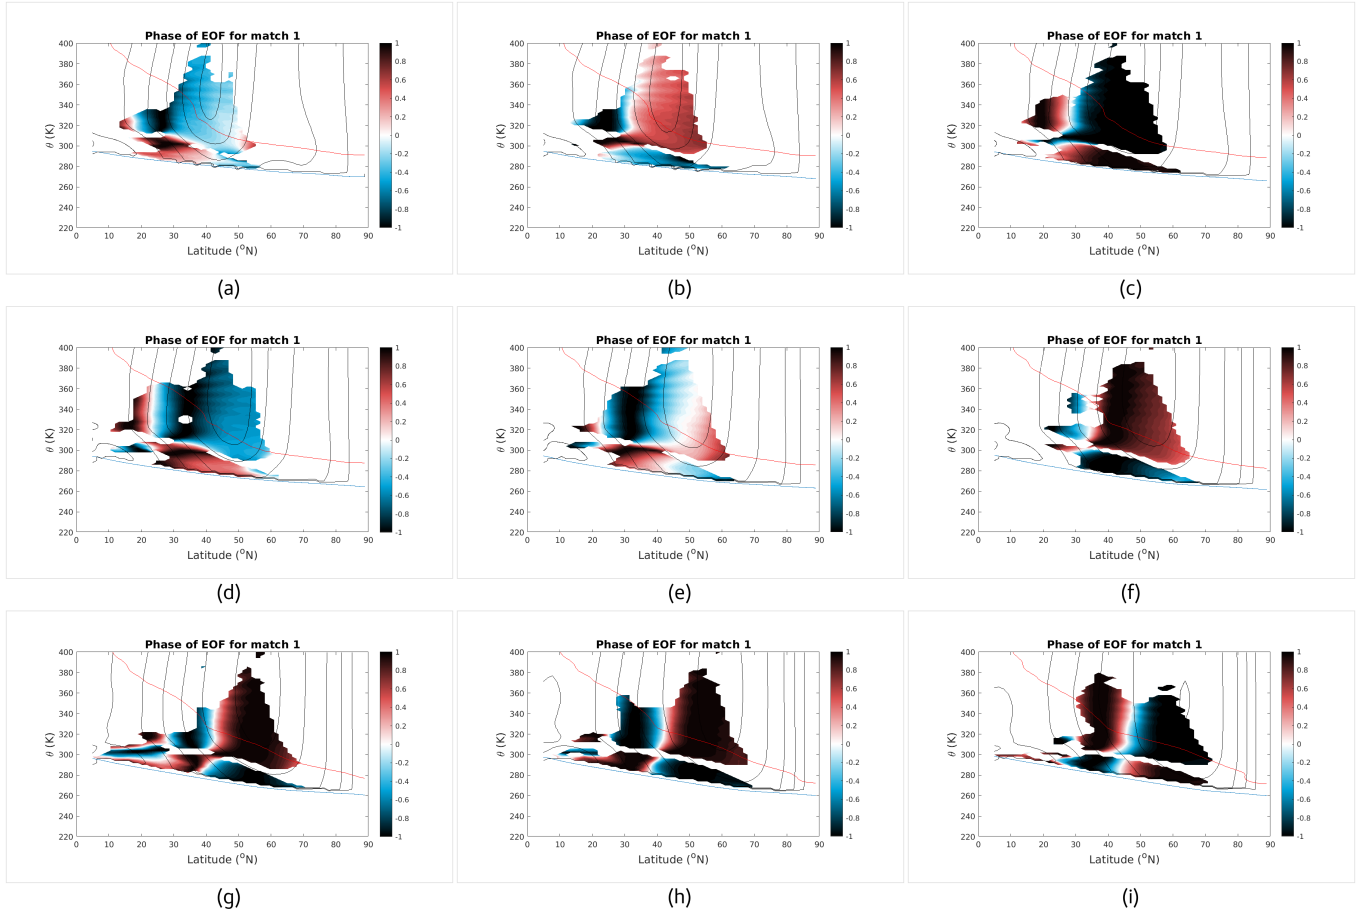


Figure 5.13: *The cosine of the phase for the leading mode of each member of the QSB, from p22 (a) through to p60 (i). The average background zonal wind of each experiment is overplot in black, the tropopause (2PVU line) in red & the surface in blue. The values of the phase vary greatly between experiments as they are essentially arbitrary; only the phase differences within a single structure are important. Note that as $\cos \epsilon$ is used, as opposed to the phase itself, that phases can appear to ‘wrap around’.*

While the zonal wavenumber of the experiments on the QSB decreases as the jet shifts polewards, by taking the latitude of the jet maximum in the background flow, it is possible to use the zonal wavenumber as well as the circumference of that latitude circle to calculate and compare the physical wavelengths of modes across experiments, shown for the QSB in figure 5.14. The physical wavelength stays within the relatively narrow range of 5500-7300km for all experiments, with variations imposed by the integer values of zonal wavenumber.




Figure 5.14: *Physical wavelength in metres of waves at each zonal wavenumber for each experiment, calculated using the background state latitude of the jet maximum. The first three points have a zonal wavenumber $m = 5$, the second three $m = 4$, and the final three $m = 3$.*

Figure 5.15 breaks down the phase speed c_{int} of the QSB into its contributions from advection and propagation, and further breaks down propagation contributions into \mathcal{H}_{KE} vs \mathcal{H}_{APE} . It can be seen that at the equilibrium jet latitude of 55° , there is a change in behaviour of both the advection and propagation contributions which level out as opposed to continuing to grow with latitude. In the case of the propagation contribution, this is due to a reduction of the KE term, which begins to reduce at $\sim 55^\circ$. APE on the other hand continues to grow with latitude.

At this point, the advection term also levels out, whereas it had been following the zonal angular velocity of the jet as expected at lower jetstream latitudes (the zonal angular velocity continues to grow despite reduced jet speed due to further poleward shifting).




Figure 5.15: Contributions to c_{int} (blue) in degrees latitude/day versus latitude of the background state jet maximum in each experiment. The orange line shows the maximum in zonal angular velocity of the jetstream. The yellow line shows the advection contributions to the mode speed (the same as Fig. 5.10), the purple line shows the energy contributions (the same as Fig. 5.11), with the dashed green and cyan lines representing the \mathcal{H}_{KE} & \mathcal{H}_{APE} components of the energy, respectively.

The point at which this departure of mode advection speed from the Zonal Angular Velocity occurs, as highlighted in figure 5.16, matches that at which the structures in the QSB depart from those at low latitude; implying the presence of a critical point around $\sim 55^\circ$ where the behaviour of the modes changes subtly. This could be due to the equatorward position of the WA at these latitudes relative to the jet, which causes a reduction in the advection term. This reduction is balanced by the reduction in the KE term, ensuring the modes' low phase speed despite the shift in the underlying dynamics.




Figure 5.16: Ratio between the Zonal Angular Velocity and the advection contribution to C_{int} .

5.6 Conclusions

The ENM technique successfully pulled out a set of modes which explain the variance of the data in the experiments described in Chapter 4. I was able to use the modes in order to examine how the large-scale structure changes as the structure of the background flow changes, with the jet moving polewards. Several observations have been made.

Firstly, and most importantly, is that there exists a set of slow modes which vary continuously with jetstream latitude. Due to the absence of features fixed relative to the surface in the experiments (i.e. orography), all structures must propagate, however a set of quasi-stationary disturbances are seen with near zero phase speed. This quasi-stationary branch (QSB) of slow modes corresponds to the dominant propagating modes, selected by the ENM technique, with the highest energy in each experiment.

Each ENM quadrature pair in the QSB has a single dominant maximum in wave activity in the meridional direction, spanning or flanking the jet maximum in latitude. In the vertical they possess two separate maxima near the lower boundary and near the tropopause, related to the counter-propagating Rossby waves of baroclinic instability theory, discussed in more depth in Chapter 1.

The ENM technique enables a partition of intrinsic phase speed of each mode into contributions from advection by the background zonal flow and propagation relative to it. The stronger energy of the QSB modes enables them to propagate westwards relative to

the eastward background zonal flow, sufficiently fast to achieve their characteristic zero phase speed.

As jet latitude varies between experiments, the advection rate of the QSB modes is approximately a constant fraction of the maximum zonal angular velocity of the jet, until the equilibrium jet latitude reaches $\sim 55^\circ$, when the behaviour undergoes a change.

The rate of counter propagation depends on scale (both zonal and meridional). On the QSB all the ENMs have one meridional node (of similar scale). So variation in propagation is expected to vary with zonal wavelength. Interestingly, although zonal wavenumber on the QSB decreases from 5 to 3 as jet latitude is increased, the physical zonal wavelength does not vary greatly. This indicates that the most energetic modes on the jet in this spherical domain have similar zonal (and meridional) length scales.

The increase in energy needed to counteract the increase in jet angular velocity and achieve zero phase speed, as jet latitude increases across the experiments, stems mainly from an increase in the APE of the ENMs, rather than KE. Therefore it cannot be explained by recourse for the familiar dispersion relation for Rossby waves in single layer dynamics, such as the barotropic vorticity equation (the dispersion relation of which is equation 1.6 of Chapter 1) or QGSW. Indeed, the small variation of KE for the ENMs on the QSB is consistent with the fact that the zonal wavelength does not vary much and therefore neither do the wind speed anomalies attributable to the inversion of PV anomalies in a Rossby wave (the inversion operator being scale-dependent).

The overall finding is that quasi-stationary modes depend on jet latitude in a systematic way. In the experiments studied, these waves are the most energetic propagating sinuous modes and they have a preferred zonal and meridional scale. They are baroclinic waves with distinct phase-locked upper and lower counter-propagating Rossby wave structures. Although the zonal advection contribution to propagation increases with jet zonal angular velocity, the energy of the modes also increases enabling zero net phase speed. Given the near-constancy of zonal wavelength, it can be predicted that the zonal wavenumber of the quasi-stationary modes must be lower on higher latitude jets and that the discrete values of wavenumber have an important influence on the energy of the dominant modes as jet latitude varies resulting in discontinuous steps in behaviour of the dominant waves.

Chapter 6:

Conclusions & Future Work

6.1 Introduction

Several hypotheses were posed in the introduction to the thesis;

H_1 The variability associated with large-amplitude disturbances on the jetstream can be decomposed into a small set of dominant modes, each with a distinct structure and an intrinsic frequency which can be deduced from the structure.

H_2 The intrinsic frequency matches the frequency observed in the timeseries data of the modes that are freely propagating in the absence of strong forcing by non-conservative processes or nonlinear interactions between modes

H_3 The dominant modes vary continuously with jet latitude and there is a relationship between mode frequency (for a given zonal wavenumber) and jet latitude

This chapter outlines how the results in this thesis address and answer these hypotheses; where new, unresolved questions have arisen, and in which direction further research into this area may be taken.

6.2 Verification of the ENM technique for Baroclinic Flows

6.2.1 Decomposition into Modes & Reduction of Dimensionality

In Chapters 3 and 5, the ENM technique was used to extract modes of variability from idealised IGCM simulations. The construction of the ENM technique, in contrast

to other principal component analysis approaches, allows for the classification of wave-like structures with an intrinsic frequency and phase speed dictated by the structure of the ENMs themselves due to their construction using conserved wave activity quantities. Equation 2.60 in Chapter 2 gives the speed of these disturbances for the ENMs. The extraction of these modes, and the speeds shown in the aforementioned chapters serves as confirmation of the latter half of hypothesis 1.

For each of the experiments analysed with the ENM technique, it is evident from inspection that the modes reflect the large scale structure of the model atmosphere time-series. This is most obvious in the simplified experiments performed in Chapter 3 without surface drag, where the atmospheric motion consists of large amplitude, large wavelength planetary waves, with smaller structures riding the edge of the planetary waves. In these cases, each structure's zonal wavenumber and empirical phase speed matches the observed rotation rate of the single dominant mode as seen in Hovmöller plots, and which are observed in animations of upper level PV surfaces. For the model atmospheres where the presence of drag results in a more complex structure, the Hovmöller diagrams of meridional wind 3.4 demonstrate propagating structures which on deeper examination, can be seen reflected in the dominant ENMs for these experiments.

In Chapter 3, two measures of relative dominance for leading ENMs in a given spectrum were introduced; the ENM prominence, describing the jump in power between the leading mode pair and subsequent modes; and the ENM proportion, a fractional measure of the proportion of the total pseudomomentum explained by the leading pair. As the ENMs are orthogonal with respect to pseudomomentum, the total pseudomomentum is a sum of the pseudomomentum possessed by each mode and as such the power in the full field can be decomposed rigorously in this manner. In each case, the dominant modes of the zonal wavenumbers with the most power as shown by figures 3.11 (a) to 3.16 (a) display high proportion as shown in subplot (b) of the same figure. In each experiment, local maxima in power correspond to the leading ENM pair for a given zonal wavenumber possessing a proportion of at least 0.4. As such, this means that for wavenumbers dominant in a given atmospheric timeseries, nearly all of the total variance of the timeseries is explained by just the leading ENMs. A large amount of the remainder of the variability is described by the first few modes in each spectrum. Looking back to figures 3.11 through 3.16 (a) once more, it can be seen that the leading 10 ENMs (kspehead, the dashed line) match the total power in the spectrum very closely.

All together, for these dominant structures, this evidence points to confirmation of

H_1 ; that the variability may be described by a relatively small set of modes with distinct structures. The number of co-ordinates required to describe the variability of the wave activity across the timeseries is reduced from 2000 (i.e. the number of time points in the atmospheric timeseries used in the analysis) to a few dozen at most, with a significant proportion of that variability explained by the single leading pair alone.

6.2.2 The Phase Speed Condition Test

An advantage of the ENM technique is the intrinsic phase speed which is related to the structure of the modes. However, validation of the ability of the ENMs to represent dynamical atmospheric flows is required in order to utilise these phase speeds to study the behaviour of large scale waves. The Phase Speed Condition allows for an objective measure of this ability; the empirical phase speed, as measured from the principal component timeseries, must match the intrinsic phase speed for freely propagating modes. This condition is expected to hold when the pseudoenergy conservation requirement is met, i.e. the background state is slowly evolving. In addition, it is not expected to hold when the flow is subject to strong non-conservative forcing, or when there is a significant amount of non-linear wave-wave interactions.

Across all experiments in Chapters 3 & 5, a common pattern emerges; dominant baroclinic scale waves at zonal wavenumbers $\sim 5, 6, 7, 8$ meet the phase speed condition, while larger, planetary scale waves, and smaller structures, do not always match in phase speed because the dynamics do not meet all the conditions stipulated above. While the most dominant ENMs (which explain the largest amount of variability in the atmospheric timeseries) are consistent with phase speed matching, there is a working theory, of why other structures' phase speeds do not match.

Firstly, as shown in figure 3.18, the smallest scale structures from wavenumbers ~ 8 and above show larger eastward intrinsic phase speeds compared to c_{emp} . Nonetheless, these wavenumbers possess similar c_{emp} to each other and to the largest scale planetary waves. From this, it can be hypothesized that these structures picked out by the ENM technique are subject to highly non-local (in wavenumber) nonlinear effects. Large scale Rossby wave breaking leads to filamentation, where a single dominant planetary scale wave projects onto the small scale structures. The dominant planetary scale wave undergoes Rossby wave breaking where high PV filaments are wrapped around the cat's eye (McIntyre and Palmer, 1985) structures associated with the wave. This behaviour projects onto wavenumbers higher than that of the dominant wave and it is also expected

that these structures would move in step with the primary wave. This matching of empirical speeds between high wavenumber structures and the $m = 1$ planetary wave can also be seen in the Chapter 5 experiments, in figure 5.7.

It is possible that lower wavenumbers, planetary waves with zonal wavenumber ≤ 5 (depending on the experiment) display a different type of nonlinearity, which is instead local in wavenumber. We hypothesize that these waves are undergoing Rossby wave triad interactions. While their c_{int} display different phase speeds, the c_{emp} across wavenumbers are similar. Triad behaviour is identified in Chapter 3 for the TORD and RARD experiments. It is observed that the modes of the proposed triads (1) match the selection rules in zonal wavenumber $k_1 + k_2 = k_3$ (on the basis the structures share the same meridional node structure of a single node); (2) have similar power in pseudomomentum; (3) have similar c_{emp} , even though the c_{int} do not match; (4) consist of two wavenumbers with prominent propagating modes, while the third of the triad does not, and; (5) the wavenumbers with prominent modes are anticorrelated. All of these observations are consistent with triad behaviour.

While the phase speed condition test is not met at all scales, the phase speed matching at wavenumbers with the most prominent propagating pairs, isolating baroclinic waves, serves as confirmation of H_2 for these structures, while there are good possible explanations for exceptions.

6.3 Variation of Phase Speed with Jet Latitude

In order to study the effect of a change in the structure of the background zonal flow, specifically the latitude of the jetstream, on the ENMs, in Chapter 4 a set of idealised model experiments were developed. These experiments, like the ones which came before, lack orography, moist processes and seasonality. However, through the construction of the relaxation state used in the model, the model timeseries possess an equilibrium jetstream whose latitude may be controlled by a simple parameter. As shown in figure 4.5, a nearly linear response is attained in the equilibrium jet latitude from changes to the latitude of the jet in the relaxation state.

Using these experiments, the change in modal response as equilibrium jet latitude is varied was examined using the ENMs for each experiment.

6.3.1 Variation of Mode Structure with Latitude

As the latitude of the jetstream increases, the relative position of the structure of the dominant ENMs and the jet changes. As shown in figure 5.8 (Chapter 5), as the jet latitude increases, a structure with the same zonal wavenumber, and identical meridional and vertical node numbers in its pseudomomentum amplitude structure can be observed. This structure changes the position of its wave activity from underneath the jet core, to intersecting with it, to a position on the equatorward flank.

These structural changes correspond to changes in speed of the modes, however these changes are scale-dependent, as can be seen in figure 5.6. At the smallest scales, the structures all possess an eastward intrinsic phase speed, which is smallest when the relaxation jet is at both tropical and polar latitudes as the propagation against the jet is weak across all latitudes, whereas when the relaxation state jet is in the mid-latitudes the wave activity's intersection with the jet increases the advection term, resulting in fast-moving structures when ϕ_{relax} is at mid-latitudes.

At the dominant scales for baroclinic growth ($m = 5 - 8$), large energy modes dominate, resulting in low phase speeds, due to their large perturbation energy and therefore ability to counter-propagate, which get faster the further poleward the jet shifts. All the dominant modes across the experiments for $m \geq 4$ have a simple meridional structure with a single maximum in pseudomomentum which varies in its location relative to the jet core.

6.3.2 Quasi-Stationary Waves

A 'branch' of quasi-stationary waves, or a Quasi Stationary Branch (QSB), exists across all experiments. These waves possess a similar physical wavelength, and are the modes with the greatest propagation rate relative to the jet, indicating that they have the greatest perturbation energy for each zonal wavenumber. They are baroclinic mode structures with large pseudomomentum amplitude at the tropopause and in the lower troposphere near the lower boundary, the necessary ingredients for strong baroclinic growth through the interaction of counter-propagating Rossby waves. It is hypothesised that these dominant modes are the baroclinic wave structures with the strongest baroclinic interaction strength (and strongest growth rate for normal modes). As in the Eady and Charney models the fastest growing modes have a characteristic wavenumber determined by the competition between the requirements for baroclinic interaction between the upper and lower CRWs (strongest at larger scales) versus the requirement to phase-lock in the

shear flow (counter-propagation becomes too strong for the longest waves). This even results in the neutral longwave point in the Charney model ((Heifetz et al., 2004b)). It is interesting that their wavelength does not vary much between experiments which is likely a result of the scale effect of PV inversion, relating PV waves to the velocity they induce. However, the zonal wavenumber of the sphere is discretised. So as the jet shift polewards and so does the wave activity, the length of the latitude circle $a \cos(\phi)$ reduces. The zonal wavenumber corresponding to the same wavelength must therefore decrease and it can only do so in discrete jumps ($m = 6, 5, 4$). Figure 5.14 shows that as the wavelength begins to decrease below $\sim 5.5 \times 10^3 km$, the wavelength increases sharply (corresponding to a jump in zonal wavenumber), suggesting a minimum wavelength required to maintain quasi-stationary behaviour. This shows that a shift in the background jet latitude could lead to a change in the wavenumber of stationary structures in the real atmosphere.

6.4 Further Research

There are a number of directions in which further research into this topic could proceed.

6.4.1 Effect of Non-Linear Coupling

Firstly, a number of questions have arisen while answering the hypothesis put forward at the beginning of this thesis. The first of these, is the impact of non-linear wave-wave coupling on the ENMs. As discussed in section 6.2.2, some ENMs at the largest and smallest scales do not meet the ENM phase speed condition. It is hypothesized that this is due to the effects of nonlinear interactions, both between scales, and in the form of triad interaction.

Resonant Rossby wave triads are a form of non-linear interaction between three waves where each combination of two waves interacts to produce the third. This way, energy is exchanged by the three wave types within the triad, but the waves do not interact with other waves outside of the triad. Wave triads are subject to strict selection criteria; the sum of the wavenumbers in a given spatial direction must sum to zero. In practice this allows for triad interactions between, e.g., wavenumbers 2, 3 and 5, but not between 1, 3 and 5. In addition, selection rules must apply to the frequencies, and therefore phase speeds of the waves.

A study of Rossby wave triads within the identified ENMs would require examination

of possible combinations of triads, based upon the selection rules. By examining the sum of the amplitudes of a given combination, it should be possible to see whether the total pseudoenergy and also pseudomomentum in a given triad remains constant as would be expected, due to the triads being unable to exchange energy with other waves outside the coupling. With an exhaustive examination of the triads it would be possible to see exactly how nonlinear interactions are modifying the phase speeds of the waves at the planetary scale, and examine the precise case-by-case mechanism for the phase speed matching failing at certain small wavenumbers.

For the largest wavenumber waves, triad interaction cannot explain the disparity between empirical and intrinsic phase speeds. Large scale coupling, where a single dominant wavenumber projects onto smaller scales as a result of wave breaking is considered to be the cause of this interaction. Further studies of interaction between the smallest and largest scales is required to ascertain the mechanics of this interaction. By examining the phase relationships between the lowest and highest wavenumbers, it is hypothesized that further insight might be gained.

6.4.2 Effect of Jet Strength on the Modes

The aim of chapters 4 & 5 was to establish a set of experiments where different aspects of the background state could be varied through changing a simple variable to define each of them, and examine how the modes evolve as the background flow is changed. Specifically, the experiments in Chapter 4 vary the latitude of the jetstream while keeping other parameters constant, however there are other axes within the phase space of the experimental setup along which to vary relaxation parameters in order to achieve changes in the background state. For example Chapter 4 briefly discusses a set of experiments whereby the relaxation strength was altered in order to examine the effect on equilibrium latitude.

The behaviour in the Jet Control experiments is at odds with the usual seasonal behaviour in quasi-stationary waves ([Kornhuber et al. \(2017\)](#), [Coumou et al. \(2015\)](#)) where $m=3-4$ is dominant in winter and $m=6-7$ is dominant in summer.

A variation in the speed of the jetstream is a likely candidate which would account for the observed behaviour in the atmosphere, which would provide a more complete picture of how the modes change with changes in the background flow. In the real atmosphere, as the jet moves into the weaker summer state, stationary waves of higher wavenumber appear; as we can see from the barotropic dispersion relation (equation 1.6, Chapter 1),

weaker zonal mean wind means that higher wavenumbers are able to maintain stationary phase. By focusing on changing only the strength of the jet, the effect of jet strength on modes could be isolated.

The most energetic baroclinic waves have the most negative phase speeds in the jet latitude experiments and there are no modes propagating strongly westwards. This is hypothesised to be because in the absence of orography or other forcing, all the amplitude in the waves is driven by baroclinic instability. The semi-circle theorems (e.g., (Pedlosky, 1963)) indicate that unstable modes in shear instability cannot have phase speeds outside the range of speeds present in the background state. Therefore longer waves with westward speeds are not excited in these experiments. It is in contrast to the atmosphere where westward propagating long waves are frequently observed.

6.4.3 Seasonality Effects

In the physical atmosphere, the jet latitude is not constant. The background flow varies with the seasonal cycle as a result of changes in solar insolation and wave activity, and with it the latitude of the background jet exhibits seasonal changes. It is possible that relatively abrupt transitions could occur from wavenumber 6 to 4 as the jet shifts polewards, mirroring the discrete shifts in wavenumber which can be seen in the QSB (Fig. 5.6 & similar, Chapter 5).

A series of experiments that specifically examine the transitional behaviour of the modes, using a seasonally-varying relaxation state that would oscillate the model atmosphere from conditions mimicking winter, to summer, and back, could examine both the transition of the waves in the timeseries as the season changes, as well as whether the same structures re-emerge following the return to the initial relaxation scheme.

References

- Ambaum, M. H., B. J. Hoskins, and D. B. Stephenson, 2001: Arctic oscillation or north atlantic oscillation? *Journal of Climate*, **14** (16), 3495–3507.
- Andrews, D. G., J. R. Holton, and C. B. Leovy, 1987a: *Middle atmosphere dynamics*. 40, Academic press.
- Andrews, D. G., J. R. Holton, and C. B. Leovy, 1987b: *Middle atmosphere dynamics*. 40, Academic press.
- Aref, H., 1984: Stirring by chaotic advection. *Journal of fluid mechanics*, **143**, 1–21.
- Arneth, A., and Coauthors, 2019: Summary for policymakers.
- Athanasiadis, P. J., J. M. Wallace, and J. J. Wettstein, 2010: Patterns of wintertime jet stream variability and their relation to the storm tracks. *Journal of the Atmospheric Sciences*, **67** (5), 1361–1381.
- Barnes, E. A., and D. L. Hartmann, 2010: Dynamical feedbacks and the persistence of the nao. *Journal of Atmospheric Sciences*, **67** (3), 851–865.
- Barnes, E. A., D. L. Hartmann, D. M. Frierson, and J. Kidston, 2010: Effect of latitude on the persistence of eddy-driven jets. *Geophysical research letters*, **37** (11).
- Blackburn, M., 1985: *Program description for the multi-level global spectral model*. Atmospheric Modelling Group, Department of Meteorology, University of Reading.
- Bleck, R., 1973: Numerical forecasting experiments based on the conservation of potential vorticity on isentropic surfaces. *Journal of Applied Meteorology and Climatology*, **12** (5), 737–752.
- Blessing, S., K. Fraedrich, M. Junge, T. Kunz, and F. Lunkeit, 2005: Daily north-atlantic oscillation (nao) index: Statistics and its stratospheric polar vortex dependence. *Meteorologische Zeitschrift*, **14** (6), 763–770.
- Boer, G., and B. Denis, 1997: Numerical convergence of the dynamics of a gcm. *Climate dynamics*, **13** (5), 359–374.

- Böning, C. W., A. Dispert, M. Visbeck, S. Rintoul, and F. U. Schwarzkopf, 2008: The response of the antarctic circumpolar current to recent climate change. *Nature Geoscience*, **1** (12), 864–869.
- Branstator, G., 1995: Organization of storm track anomalies by recurring low-frequency circulation anomalies. *Journal of the atmospheric sciences*, **52** (2), 207–226.
- Brayshaw, D. J., B. Hoskins, and M. Blackburn, 2008: The storm-track response to idealized sst perturbations in an aquaplanet gcm. *Journal of the Atmospheric Sciences*, **65** (9), 2842–2860.
- Bretherton, F., 1966: Baroclinic instability and the short wavelength cut-off in terms of potential vorticity. *Quarterly Journal of the Royal Meteorological Society*, **92** (393), 335–345.
- Broccoli, A. J., and S. Manabe, 1992: The effects of orography on midlatitude northern hemisphere dry climates. *Journal of Climate*, **5** (11), 1181–1201.
- Brunet, G., 1994: Empirical normal-mode analysis of atmospheric data. *Journal of the atmospheric sciences*, **51** (7), 932–952.
- Brunet, G., and J. Methven, 2019: Identifying wave processes associated with predictability across time scales: An empirical normal mode approach. *Sub-Seasonal to Seasonal Prediction*, Elsevier, 65–90.
- Brunet, G., and R. Vautard, 1996: Empirical normal modes versus empirical orthogonal functions for statistical prediction. *Journal of the Atmospheric Sciences*, **53** (23), 3468–3489, doi:10.1175/1520-0469(1996)053<3468:ENMVEO>2.0.CO;2, URL [https://doi.org/10.1175/1520-0469\(1996\)053<3468:ENMVEO>2.0.CO;2](https://doi.org/10.1175/1520-0469(1996)053<3468:ENMVEO>2.0.CO;2), [https://doi.org/10.1175/1520-0469\(1996\)053<3468:ENMVEO>2.0.CO;2](https://doi.org/10.1175/1520-0469(1996)053<3468:ENMVEO>2.0.CO;2).
- Chang, E. K., S. Lee, and K. L. Swanson, 2002: Storm track dynamics. *Journal of climate*, **15** (16), 2163–2183.
- Charney, J. G., 1947: The dynamics of long waves in a baroclinic westerly current. *Journal of Atmospheric Sciences*, **4** (5), 136–162.
- Charney, J. G., 1973: Planetary fluid dynamics. *Dynamic meteorology*, Springer, 97–351.
- Charney, J. G., and A. Eliassen, 1949: A numerical method for predicting the perturbations of the middle latitude westerlies. *Tellus*, **1** (2), 38–54.

- Charney, J. G., and M. E. Stern, 1962: On the stability of internal baroclinic jets in a rotating atmosphere. *Journal of Atmospheric Sciences*, **19** (2), 159–172.
- Chen, T.-C., 2010: Characteristics of summer stationary waves in the northern hemisphere. *Journal of climate*, **23** (17), 4489–4507.
- Coumou, D., J. Lehmann, and J. Beckmann, 2015: The weakening summer circulation in the northern hemisphere mid-latitudes. *Science*, **348** (6232), 324–327.
- Dickinson, R. E., 1978: Rossby waves–long-period oscillations of oceans and atmospheres. *Annual Review of Fluid Mechanics*, **10** (1), 159–195, doi:10.1146/annurev.fl.10.010178.001111, URL <https://doi.org/10.1146/annurev.fl.10.010178.001111>, <https://doi.org/10.1146/annurev.fl.10.010178.001111>.
- Dritschel, D., and M. McIntyre, 2008: Multiple jets as pv staircases: The phillips effect and the resilience of eddy-transport barriers. *Journal of the Atmospheric Sciences*, **65** (3), 855–874.
- Eady, E. T., 1949: Long waves and cyclone waves. *Tellus*, **1** (3), 33–52.
- European Environment Agency, 2021: Indicator assessment: Economic losses from climate-related extremes in europe. URL <https://www.eea.europa.eu/data-and-maps/indicators/direct-losses-from-weather-disasters-4/assessment>.
- Franzke, C., 2002: Dynamics of low-frequency variability: Barotropic mode. *Journal of the atmospheric sciences*, **59** (20), 2897–2909.
- Franzke, C., T. Woollings, and O. Martius, 2011: Persistent circulation regimes and preferred regime transitions in the north atlantic. *Journal of the atmospheric sciences*, **68** (12), 2809–2825.
- Gerber, E. P., and G. K. Vallis, 2007: Eddy–zonal flow interactions and the persistence of the zonal index. *Journal of the Atmospheric Sciences*, **64** (9), 3296–3311.
- Grams, C. M., L. Magnusson, and E. Madonna, 2018: An atmospheric dynamics perspective on the amplification and propagation of forecast error in numerical weather prediction models: A case study. *Quarterly Journal of the Royal Meteorological Society*, **144** (717), 2577–2591.
- H. Hendon, H., and D. Hartmann, 1985: Variability in a nonlinear model of the atmosphere with zonally symmetric forcing. **42**, 2783–2797.

- Harnik, N., and E. K. Chang, 2004: The effects of variations in jet width on the growth of baroclinic waves: Implications for midwinter pacific storm track variability. *Journal of the atmospheric sciences*, **61** (1), 23–40.
- Hashim, J. H., and Z. Hashim, 2016: Climate change, extreme weather events, and human health implications in the asia pacific region. *Asia Pacific Journal of Public Health*, **28** (2_suppl), 8S–14S.
- Haynes, P. H., 1988: Forced, dissipative generalizations of finite-amplitude wave-activity conservation relations for zonal and nonzonal basic flows. *Journal of Atmospheric Sciences*, **45** (16), 2352–2362.
- Heifetz, E., C. Bishop, B. Hoskins, and J. Methven, 2004a: The counter-propagating rossby-wave perspective on baroclinic instability. i: Mathematical basis. *Quarterly Journal of the Royal Meteorological Society*, **130** (596), 211–231.
- Heifetz, E., J. Methven, B. Hoskins, and C. Bishop, 2004b: The counter-propagating rossby-wave perspective on baroclinic instability. ii: Application to the charney model. *Quarterly Journal of the Royal Meteorological Society*, **130** (596), 233–258.
- Held, I. M., 1985: Pseudomomentum and the orthogonality of modes in shear flows. *Journal of Atmospheric Sciences*, **42** (21), 2280–2288.
- Held, I. M., and M. J. Suarez, 1994a: A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. *Bulletin of the american Meteorological society*, **75** (10), 1825–1830.
- Held, I. M., and M. J. Suarez, 1994b: A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. *Bulletin of the American Meteorological society*, **75** (10), 1825–1830.
- Held, I. M., M. Ting, and H. Wang, 2002: Northern winter stationary waves: Theory and modeling. *Journal of climate*, **15** (16), 2125–2144.
- Hoskins, B., R. Fonseca, M. Blackburn, and T. Jung, 2012: Relaxing the tropics to an ‘observed’ state: Analysis using a simple baroclinic model. *Quarterly Journal of the Royal Meteorological Society*, **138** (667), 1618–1626.
- Hoskins, B., and A. Simmons, 1975: A multi-layer spectral model and the semi-implicit method. *Quarterly Journal of the Royal Meteorological Society*, **101** (429), 637–655.

- Hoskins, B. J., 1991: Towards a p - θ view of the general circulation. *Tellus A: Dynamic Meteorology and Oceanography*, **43** (4), 27–36.
- Hoskins, B. J., and D. J. Karoly, 1981: The steady linear response of a spherical atmosphere to thermal and orographic forcing. *Journal of Atmospheric Sciences*, **38** (6), 1179–1196.
- Jia, X. J., J. Derome, and H. Lin, 2007: Comparison of the life cycles of the nao using different definitions. *Journal of climate*, **20** (24), 5992–6011.
- Kidston, J., and E. Gerber, 2010: Intermodel variability of the poleward shift of the austral jet stream in the cmip3 integrations linked to biases in 20th century climatology. *Geophysical Research Letters*, **37** (9).
- Kornhuber, K., S. Osprey, D. Coumou, S. Petri, V. Petoukhov, S. Rahmstorf, and L. Gray, 2019: Extreme weather events in early summer 2018 connected by a recurrent hemispheric wave-7 pattern. *Environmental Research Letters*, **14** (5), 054002.
- Kornhuber, K., V. Petoukhov, S. Petri, S. Rahmstorf, and D. Coumou, 2017: Evidence for wave resonance as a key mechanism for generating high-amplitude quasi-stationary waves in boreal summer. *Climate Dynamics*, **49** (5), 1961–1979.
- Lai, T. K., 2013: Emergence of tropical cyclones in baroclinic waves. M.S. thesis, Meteorology, The University of Reading.
- Lau, N.-C., 1979: The observed structure of tropospheric stationary waves and the local balances of vorticity and heat. *Journal of Atmospheric Sciences*, **36** (6), 996–1016.
- Leib, S., and M. Goldstein, 1989: Nonlinear interaction between the sinuous and varicose instability modes in a plane wake. *Physics of Fluids A: Fluid Dynamics*, **1** (3), 513–521.
- Leonard, D., 2021: Sep 21. Tech. rep.
- Magnusdottir, G., and P. H. Haynes, 1996: Wave activity diagnostics applied to baroclinic wave life cycles. *Journal of the atmospheric sciences*, **53** (16), 2317–2353.
- Manabe, S., and T. B. Terpstra, 1974: The effects of mountains on the general circulation of the atmosphere as identified by numerical experiments. *Journal of Atmospheric Sciences*, **31** (1), 3–42.
- Marshall, D. P., M. H. Ambaum, J. R. Maddison, D. R. Munday, and L. Novak, 2017: Eddy saturation and frictional control of the antarctic circumpolar current. *Geophysical research letters*, **44** (1), 286–292.

- McIntyre, M., and T. Palmer, 1985: A note on the general concept of wave breaking for rossby and gravity waves. *pure and applied geophysics*, **123** (6), 964–975.
- 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**, 527–565.
- Melander, M., J. McWilliams, and N. Zabusky, 1987: Axisymmetrization and vorticity-gradient intensification of an isolated two-dimensional vortex through filamentation. *Journal of Fluid Mechanics*, **178**, 137–159.
- Methven, J., 2013: Wave activity for large-amplitude disturbances described by the primitive equations on the sphere. *Journal of the Atmospheric Sciences*, **70** (6), 1616–1630, URL <http://centaur.reading.ac.uk/32701/>.
- Methven, J., and P. Berrisford, 2015: The slowly evolving background state of the atmosphere. *Quarterly Journal of the Royal Meteorological Society*, **141** (691), 2237–2258, doi:10.1002/qj.2518, URL <http://dx.doi.org/10.1002/qj.2518>.
- Michelangeli, D. V., R. W. Zurek, and L. S. Elson, 1987: Barotropic instability of midlatitude zonal jets on mars, earth and venus. *Journal of the atmospheric sciences*, **44** (15), 2031–2041.
- Monahan, A. H., J. C. Fyfe, M. H. Ambaum, D. B. Stephenson, and G. R. North, 2009: Empirical orthogonal functions: The medium is the message. *Journal of Climate*, **22** (24), 6501–6514.
- (NAG), T. N. A. G., ?????: The nag library. Oxford, United Kingdom, URL www.nag.com.
- Nakamura, N., and D. Zhu, 2010: Finite-amplitude wave activity and diffusive flux of potential vorticity in eddy–mean flow interaction. *Journal of the Atmospheric Sciences*, **67** (9), 2701–2716.
- National Oceanic and Atmospheric Administration, 2021: Billion-dollar weather and climate disasters: Overview. URL <https://www.ncdc.noaa.gov/billions/>.
- Pedlosky, J., 1963: Baroclinic instability in two layer systems. *Tellus*, **15** (1), 20–25.
- Polvani, L. M., and J. Esler, 2007: Transport and mixing of chemical air masses in idealized baroclinic life cycles. *Journal of Geophysical Research: Atmospheres*, **112** (D23).

- Polvani, L. M., and P. J. Kushner, 2002: Tropospheric response to stratospheric perturbations in a relatively simple general circulation model. *Geophysical Research Letters*, **29** (7), 18–1–18–4, doi:10.1029/2001GL014284, URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2001GL014284>, <https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2001GL014284>.
- Polvani, L. M., and R. A. Plumb, 1992: Rossby wave breaking, microbreaking, filamentation, and secondary vortex formation: The dynamics of a perturbed vortex. *Journal of Atmospheric Sciences*, **49** (6), 462–476.
- Quintanar, A. I., and C. R. Mechoso, 1995: Quasi-stationary waves in the southern hemisphere. part i: Observational data. *Journal of climate*, **8** (11), 2659–2672.
- Rayleigh, L., 1880: On the stability, or instability, of certain fluid motions. *Proc. London Math. Soc.*, **9**, 57–70.
- Rossby, C.-G., 1939: Relation between variations in the intensity of the zonal circulation of the atmosphere and the displacements of the semi-permanent centers of action. *J. mar. Res.*, **2**, 38–55.
- Seager, R., D. S. Battisti, J. Yin, N. Gordon, N. Naik, A. C. Clement, and M. A. Cane, 2002: Is the gulf stream responsible for europe’s mild winters? *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, **128** (586), 2563–2586.
- Seager, R., R. Murtugudde, N. Naik, A. Clement, N. Gordon, and J. Miller, 2003: Air–sea interaction and the seasonal cycle of the subtropical anticyclones. *Journal of climate*, **16** (12), 1948–1966.
- Shaw, T., and Coauthors, 2016: Storm track processes and the opposing influences of climate change. *Nature Geoscience*, **9** (9), 656–664.
- Simmons, A. J., and B. J. Hoskins, 1975: A comparison of spectral and finite-difference simulations of a growing baroclinic wave. *Quarterly Journal of the Royal Meteorological Society*, **101** (429), 551–565, doi:10.1002/qj.49710142912, URL <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/qj.49710142912>, <https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/qj.49710142912>.
- Simpson, I. R., M. Blackburn, J. D. Haigh, and S. N. Sparrow, 2010: The impact of the state of the troposphere on the response to stratospheric heating in a simplified gcm. *Journal of climate*, **23** (23), 6166–6185.

- Smagorinsky, J., 1953: The dynamical influence of large-scale heat sources and sinks on the quasi-stationary mean motions of the atmosphere. *Quarterly Journal of the Royal Meteorological Society*, **79** (341), 342–366.
- Solomon, A., and N. Nakamura, 2012: An exact lagrangian-mean wave activity for finite-amplitude disturbances to barotropic flow on a sphere. *Journal of Fluid Mechanics*, **693**, 69–92, doi:10.1017/jfm.2011.460.
- Son, S.-W., and S. Lee, 2005: The response of westerly jets to thermal driving in a primitive equation model. *Journal of the atmospheric sciences*, **62** (10), 3741–3757.
- Thorncroft, C., B. Hoskins, and M. McIntyre, 1993: Two paradigms of baroclinic-wave life-cycle behaviour. *Quarterly Journal of the Royal Meteorological Society*, **119** (509), 17–55.
- Thuburn, J., and V. Lagneau, 1999: Eulerian mean, contour integral, and finite-amplitude wave activity diagnostics applied to a single-layer model of the winter stratosphere. *Journal of the atmospheric sciences*, **56** (5), 689–710.
- Valdes, P. J., and B. J. Hoskins, 1989: Linear stationary wave simulations of the time-mean climatological flow. *Journal of the Atmospheric Sciences*, **46** (16), 2509–2527, doi:10.1175/1520-0469(1989)046<2509:LSWSOT>2.0.CO;2, URL [https://doi.org/10.1175/1520-0469\(1989\)046<2509:LSWSOT>2.0.CO;2](https://doi.org/10.1175/1520-0469(1989)046<2509:LSWSOT>2.0.CO;2), [https://doi.org/10.1175/1520-0469\(1989\)046<2509:LSWSOT>2.0.CO;2](https://doi.org/10.1175/1520-0469(1989)046<2509:LSWSOT>2.0.CO;2).
- Wang, L., S. Kanji, and S. Bandyopadhyay, 2009: The health impact of extreme weather events in sub-saharan africa. *World Bank Policy Research Working Paper*, (4979).
- Wheeler, M., and G. N. Kiladis, 1999: Convectively coupled equatorial waves: Analysis of clouds and temperature in the wavenumber–frequency domain. *Journal of the Atmospheric Sciences*, **56** (3), 374–399, doi:10.1175/1520-0469(1999)056<0374:CCEWAO>2.0.CO;2.
- Williams, G., 2003: Barotropic instability and equatorial superrotation. *Journal of the atmospheric sciences*, **60** (17), 2136–2152.
- Woollings, T., C. Czuchnicki, and C. Franzke, 2014: Twentieth century north atlantic jet variability. *Quarterly Journal of the Royal Meteorological Society*, **140** (680), 783–791.

- Woollings, T., A. Hannachi, B. Hoskins, and A. Turner, 2010: A regime view of the north atlantic oscillation and its response to anthropogenic forcing. *Journal of Climate*, **23** (6), 1291–1307.
- World Economic Forum, 2019: This chart shows where extreme weather is causing the most fatalities. URL <https://www.weforum.org/agenda/2019/12/extreme-weather-environment-climate-change>.
- Wu, Z., Z. Jiang, J. Li, S. Zhong, and L. Wang, 2012: Possible association of the western tibetan plateau snow cover with the decadal to interdecadal variations of northern china heatwave frequency. *Climate Dynamics*, **39** (9-10), 2393–2402.
- Zadra, A., G. Brunet, and J. Derome, 2002a: An empirical normal mode diagnostic algorithm applied to ncep reanalyses. *Journal of the atmospheric sciences*, **59** (19), 2811–2829.
- Zadra, A., G. Brunet, J. Derome, and B. Dugas, 2002b: Empirical normal mode diagnostic study of the gem model’s dynamical core. *Journal of the atmospheric sciences*, **59** (16), 2498–2510.
- Zappa, G., V. Lucarini, and A. Navarra, 2010: Interactions between almost unstable baroclinic waves and the tropical convection in aquaplanet simulations. *EGU General Assembly Conference Abstracts*, Vol. 12, 11842.
- Zappa, G., V. Lucarini, and A. Navarra, 2011: Baroclinic stationary waves in aquaplanet models. *Journal of the atmospheric sciences*, **68** (5), 1023–1040.
- Zhou, Y., and Z. Wu, 2016: Possible impacts of mega-el niño/southern oscillation and atlantic multidecadal oscillation on eurasian heatwave frequency variability. *Quarterly Journal of the Royal Meteorological Society*, **142** (697), 1647–1661.
- Zolotov, S. Y., I. Ippolitov, and S. Loginov, 2018: Characteristics of the subtropical jet stream over the north atlantic from reanalysis data. *IOP Conference Series: Earth and Environmental Science*, IOP Publishing, Vol. 211, 012005.