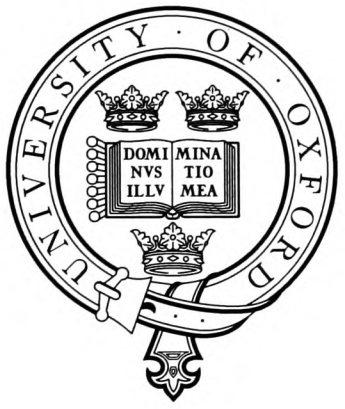
Equatorial and midlatitude circulation of Jupiter's atmosphere

Submitted for the degree of Doctor of Philosophy



David Richard Skeet
St. Edmund Hall, University of Oxford

Trinity Term, 1999



Abstract

This thesis describes the development of and results from a new numerical model of Jupiter's atmosphere. The model is an adaptation of a general circulation model (GCM) of the Earth's atmosphere. It is a high resolution, limited-area model of the upper troposphere and lower stratosphere of Jupiter, with realistic stratification and simple parametrizations of appropriate physical processes.

We have used the model to study two areas of jovian atmospheric dynamics: firstly, midlatitude vortices (exemplified by the Great Red Spot) and their interaction with baroclinic and barotropic jetstreams, and secondly, the properties of equatorial waves and their effect in driving a superrotating equatorial jet.

Before describing the main GCM, we present a time-independent model of a balanced axisymmetric vortex. The model performs a potential vorticity inversion, predicting the vertical variation of winds and temperatures required for balance, given a specified PV anomaly. We show that both shallow eddies (confined above about 20 bars) and deep ones (extending barotropically below 100 bars) could theoretically exist. The two types would be almost indistinguishable if observed only above the cloud-tops, i.e. at pressures $p \leq 300$ mb.

We then describe the primitive-equation GCM, concentrating on the modifications necessary to represent Jupiter. In order to run the GCM, basic-state temperature and zonal wind profiles must be chosen, for these are not adequately constrained by observations. Numerical solution of the vertical structure equation allows the calculation of the eigenmodes of any atmosphere, thus predicting the approximate vertical structure of Rossby waves. This simplified 1D analysis gives an insight into suitable initial profiles and the results of the full model.

The first GCM experiments examine midlatitude vortices in the absence of a background flow. Anticyclonic vortices migrate equatorward, while shrinking and intensifying by a process akin to frontogenesis. After intensifying, some eddies become unstable and disintegrate, generating daughter eddies. All eddies seem to adjust their vertical structure toward the first baroclinic normal mode. When eddies are placed in an enviroment of latitudinally alternating zonal jetstreams, meridional migration is blocked by the jets. However, most eddies intensify until they overcome the barrier presented by the jet. Mixed baroclinic/barotropic instability in the jets can generate eddies of either shallow or deep structure, echoing the results of the PV inversion model.

Moving on to equatorial dynamics, we analyse the properties of equatorial waves on Jupiter, and their interaction with the zonal mean flow. We find that mean-flow acceleration by dissipation of upward propagating Kelvin waves can (in principle) quantitatively account for the observed superrotating equatorial jet. By considering the effect of a range of possible wave sources below the clouds, we find that the troposphere filters waves primarily by their phase speed, suggesting that the speed and shape of the jet may be less sensitive to the source than to the basic atmospheric stratification. We propose a mechanism of Kelvin-wave forcing, plus weak zonally symmetric solar heating to account for the observed velocity and meridional structure of the equatorial jet. The vertical structure is less realistic, with an underlying easterly jet, but this may be because our sources are insufficiently deep in the atmosphere.

Acknowledgements

Many people have helped me with this work over the course of four years. Firstly I wish to thank my supervisor, Dr Peter Read, and my academic advisor, Dr Steve Lewis. They not only gave me sterling support during my doctorate but also arranged for funding of my postdoctoral work after the completion of (most of) this thesis. I am grateful to my examiners, Dr Patrick Irwin and Dr Michael Allison of NASA Goddard Institute for Space Studies, for giving me a fair and interesting if not always enjoyable examination.

Other direct academic help has come from Dr Warwick Norton of AOPP, in the shape of a PV inversion program, and from Dr Terry Davies and colleagues at the UK Met. Office, who supplied the Unified Model together with much advice. I would like to thank the other members of AOPP, particularly Dave Stainforth whose great knowledge and love of the UM is remarkable, and Sophia Oliver for drawing my attention to the role of the Borneo Highlands in influencing the weather. Also Adam Stephen who gave me his LaTeX thesis style file.

I should also thank all my students at St Anne's College and the inmates of St Edmund Hall for making my life interesting but tolerable, and all my Oxford friends both in Teddy Hall MCR and elsewhere, for seven fun years here. My family have been remarkably tolerant and helpful despite my best efforts. Finally, financial support was provided by the UK Particle Physics and Astronomy Research Council.

Contents

1 Introduction **1**

1.1 Overview 1

1.2 Observations of Jupiter 5

1.3 Contents of this thesis 13

2 Potential vorticity inversion models of giant eddies **15**

2.1 Mathematical formulation 18

2.1.1 Overview of the approach 18

2.1.2 The model equations 18

2.1.3 Boundary conditions 21

2.1.4 Numerical scheme for PV inversion 22

2.1.5 The model domain 24

2.1.6 Validation of the method 25

2.2 Observations of Jupiter's giant eddies 26

2.2.1 Voyager observations of the Great Red Spot 26

2.2.2 Galileo observations of the Great Red Spot 28

2.2.3 Observations of White Oval BC 29

2.2.4 Static stability profiles 30

2.3 A diagnostic model of the Great Red Spot 31

2.3.1 Potential vorticity distributions 31

2.3.2 Model results 32

2.4 A model of White Oval BC 38

2.5 Discussion of results 42

3 A simple general circulation model for Jupiter **44**

3.1 Introduction to the UK Met. Office Unified Model 44

3.2 Overview of the Unified Model numerical scheme 45

3.3 A Jovian version of the Unified Model 50

3.3.1 Physical parametrizations 58

3.3.2 Lower boundary forcing 59

3.4 Energy diagnostics 60

4 Vertical structures, normal modes and model initialization **65**

4.1 Introduction 65

4.2 Calculation of the vertical eigenfunctions 67

4.3 Vertical modes of a realistic atmosphere 74

4.4 Superposition of the normal modes 78

4.5 A method for more general boundary conditions 82

5 Midlatitude jets and vortices **90**

5.1 Introduction 90

5.2 Theories and models of jovian dynamics 94

5.3 Gaussian vortices in an atmosphere at rest 101

5.3.1 Initialization 101

5.3.2 Evolution of a shallow midlatitude anticyclone 102

5.3.3 A shallow cyclone in midlatitudes 107

5.3.4 Shallow anticyclones with arbitrary vertical structure 107

5.3.5 A deep anticyclone 112

5.3.6 Summary: effects of varying the vertical structure 112

5.3.7 Dependence of migration rate on thermal damping 114

5.3.8 Vortices of different horizontal shape 115

5.3.9 Vortex initialization by heating 117

5.3.10 Dependence on horizontal diffusion 117

5.3.11 Conclusions: vortices in an atmosphere at rest 120

5.4 Vortices in alternating jetstreams 125

5.4.1 Introduction 125

5.4.2 Jet initialization 126

5.4.3 Shallow vortex in a pair of jets: dependence on vortex strength . . . 127

5.4.4 Shallow vortex in a pair of jets: dependence on diffusion 132

5.4.5 Shallow vortex and an isolated easterly jet 135

5.4.6 Instability of the easterly jet 136

5.4.7 Energy analysis of the easterly jet instability 139

5.4.8 Conclusions: vortices embedded in jets 143

6 Equatorial waves and superrotating jets **145**

6.1 Introduction 145

6.1.1 Motivation 145

6.1.2 Overview 146

6.2 Observations of equatorial waves 148

6.3 Equatorial superrotation 150

6.3.1 Hide's theorem and equatorial superrotation 151

6.4 Theory of equatorial waves 153

6.4.1 The horizontal tidal equation 153

6.4.2 The equatorial Kelvin wave 154

6.4.3 The mixed Rossby-gravity wave 155

6.5 Forcing functions for Kelvin waves 156

6.6 Implementation and choice of wave parameters 157

6.7 Model domain, resolution and drag parameters 158

6.7.1 Stratospheric drag coefficient 158

6.7.2 Changes to resolution 161

6.7.3 Models extending higher into the stratosphere 162

6.8 Fourier analysis of waves in the model 165

6.9 Experiments with other travelling equatorial waves 168

6.9.1 A mixed Rossby-gravity wave 168

6.9.2 A superposition of Kelvin and MRG waves 171

6.9.3 Kelvin and MRG waves with different amplitude ratios 174

6.10 Wavenumber selection for travelling waves: superpositions of equal frequency 176

6.10.1 Wave spectra in a resting atmosphere 177

6.10.2 Wave-mean flow interaction 179

6.10.3 Superpositions with equal phase speed 181

6.11 Waves forced by a single stationary source 186

6.11.1 Review of some relevant terrestrial results 186

6.11.2 Stationary sources on Jupiter 187

6.11.3 A monotonic, symmetric source on the equator 187

6.11.4 A bitonic symmetric source 191

6.11.5 An antisymmetric source 192

6.12 Linear models of the Hadley circulation 197

6.12.1 An overturning circulation, conserving angular momentum 197

6.12.2 Results of the modified Held-Hou model 198

6.13 Symmetric forcing of a Hadley cell in the GCM 200

6.14 Hadley and Kelvin-wave forcing superposed 201

6.15 Conclusions 203

7 Conclusions and future work **209**

7.1 Conclusions 209

7.1.1 Limitations of our model 209

7.1.2 Midlatitude vortices and jets 211

7.1.3 Waves and the equatorial jet 214

7.1.4 Implications for Jupiter 215

7.2 Future work 215

7.2.1 Depth of model domain 216

7.2.2 Thermodynamics: moist convection and hydrogen conversion 217

7.2.3 Radiative transfer 217

7.2.4 The proposed radiation scheme 218

7.2.5 Preliminary results of the radiative-transfer model 221

A Running the Portable UM **225**

Chapter 1

Introduction

La dernière chose qu'on trouve en faisant un ouvrage, est de savoir celle qu'il faut mettre la premiere. (Pascal)

1.1 Overview

Jupiter's atmosphere exhibits a wide range of interesting meteorological phenomena, perhaps more than on any other planet. Unfortunately, quantitative information on atmospheric dynamics is much more limited than for the Earth: the data are generally restricted to remote sounding observations of cloud-top horizontal winds and the thermal structure above the clouds. One vertical profile of the deep troposphere was taken by the Galileo Probe (Atkinson et al., 1998), but there are no other direct measurements of the lower atmosphere. This situation presents a challenge to dynamicists. Numerical models can help to interpret or explain puzzling observations, whilst observations are useful in verifying models results.

The period between the Voyager and Galileo missions saw the publication of a number of review articles, summarizing knowledge of Jupiter's dynamics (e.g. Flasar, 1986; Gierasch and Conrath, 1992; Marcus, 1993) and outlining scientific objectives for Galileo (Hunten et al., 1986). In this introduction, we shall briefly draw attention to a few of the most important features of Jupiter's atmosphere.

We begin with the global-mean radiation budget. Voyager infrared measurements (Hanel et al., 1981) show that Jupiter radiates roughly twice as much heat as it absorbs from the Sun. This implies that the planet must have an internal heat source, perhaps due to gravitational collapse of the interior. The radial-mean lapse rate is calculated to be superadiabatic (Hubbard, 1968), requiring convective regions below the observed atmo-

sphere. This does not preclude the possibility of radiative envelopes (Guillot et al., 1994a), but the interior heat flux is likely to have a major effect on the observed atmospheric flow.

This issue of the lower boundary condition is a major difference between models of the Earth and of Jupiter. The latter usually attempt to simulate a thin, stably stratified "weather layer" containing the cloud decks (and hence the observed winds), supported above the much deeper, convective interior, but below the very stable stratosphere. This thesis describes models with a more realistic stratification, so we can refer precisely to the troposphere and stratosphere, rather than to an idealized weather layer. Nevertheless, the principle of modelling only part of the atmosphere is the same.

Since the vertical convective motions are weak in comparison with Jupiter's rapid rotation (its radio period is around 9 hours 55 minutes), motions are approximately two-dimensional (Marcus, 1993). Therein lies the justification for the assumption that Jupiter's weather layer may be treated as a thin envelope, like the terrestrial atmospheres but for its lower boundary condition. Most past and present Jupiter models, including ours, make this fundamental assumption. The exceptions (e.g. Zhang and Schubert, 1996) often treat Jupiter's interior using methods of stellar astrophysics, and do not model the atmosphere in detail. The coupling of these two types of model may be an interesting prospect for the future.

Seen at visible wavelengths, Jupiter is characterized by alternating bright zones and dark belts (see, for example, colour images in Rogers, 1995). These albedo variations are strongly correlated with the zonal-mean zonal wind, measured at the cloud-tops by tracking of features (Limaye, 1986). The wind pattern is one of alternating easterly and westerly zonal jets, with cloudy zones occuring in regions of anti cyclonic shear (particularly on the equatorward edge of westerly jets), and relatively clear belts in cyclonic regions. The position and magnitude of the jets remained remarkably constant over the four months between the two Voyager flybys. Peek (1958) cites historical ground-based data, showing that the jet system may be considered permanent.

Although the winds vary little, the associated cloud zones do brighten and darken intermittently (see e.g., Rogers, 1995). This suggests that the cloud thickness is not purely a function of dynamics, but depends also on the abundance of condensibles. In support of this idea, Gierasch et al. (1986) found that the cloud opacity was indeed strongly correlated with ammonia abundance. This fits with thermodynamic models (West et al., 1986) which predicted three cloud decks: water at around 4 bars, the visible ammonia cloud at around 700–300 mb, and one of ammonium hydrosulphide between the two. With regard to water, the Galileo results (see detailed discussion in section 1.2) have been rather confusing, but do not unequivocally contradict this picture.

The vertical structure of the jets above the clouds has been inferred from Voyager infrared data (Hanel et al., 1979a,b). Assuming hydrostatic and thermal-wind balance, the winds

decrease with height above the cloud-tops, generally reaching negligible speeds by about 10 mb. The same is true of the local winds in large features embedded in the jets, such as the Great Red Spot (Flasar et al., 1981) and White Ovals.

The Great Red Spot (GRS) is a giant eddy or vortex, situated at about 22°S, on the edge of the South Tropical Zone (STropZ). This region is an anticyclonic shear zone between two jets, and the Red Spot circulation is anticyclonic too. It measures roughly 12° in longitude by 6° in latitude, with the highest winds concentrated in a narrow "collar" around a quiescent interior. This interior is not without interest; it has some coherent structure, probably including cyclonic counter-rotating regions (Sada et al., 1996). The White Ovals are about half the size of the GRS, with no quiescent core.

The GRS has attracted the attention of observers for 300 years. Cassini made a continuous series of observations over the 29 years from 1665, although we cannot be sure that his spot was the same we see today. Certainly the GRS has existed for most of the 20th century, although sometimes it fades to reveal the "GRS Hollow". It is likely that these fadings occur due to (quite small) changes in the local ammonia abundance, while the dynamics remain similar. The GRS also drifts erratically in longitude, with no clear explanation.

Jupiter's equatorial region contains another interesting set of features. On the largest scale, there is a superrotating global jet centred on the equator (but with a small local minimum in zonal-mean wind there—see Fig. 6.1, or Limaye (1986)). This equatorial jet is of intrinsic interest, because it represents a maximum value of angular momentum for the whole planet, and therefore cannot arise from purely down-gradient transport of angular momentum. We observe a series of twelve plumes, spaced evenly around the planet in the dark North Equatorial Belt at 5°N, and correlated with a thermal wave (Hunt et al., 1981). Interspersed with the plumes are several '5-μm hot-spots', one of which was visited by the Galileo Probe (Young, 1998).

All these equatorial features may be interpreted (although not uniquely) in terms of the influence of equatorial waves. These are a class of inertia-gravity waves whose amplitude theoretically decreases with latitude, and which are therefore trapped close to the equator. They include Kelvin, equatorial Rossby, and mixed Rossby-gravity modes (Matsuno, 1966). Chapter 6 presents a study of the interaction of these waves with the mean flow, and the consequent structure and speed of the model equatorial jet.

In this thesis we describe two new Jupiter models and their results. One is a two-dimensional diagnostic model of a giant eddy, using the method of potential vorticity (PV) inversion. This was developed by extensive modification of an existing terrestrial model. This approach achieves very high vertical resolution and computational efficiency, at the cost of assuming axial symmetry and neglecting time dependence.

This PV inversion model suggests that balanced eddies resembling those on Jupiter could be either very deep or else rather shallow, without much effect on the observable flow above the clouds. Other surprisingly realistic features appear naturally in the model, such as weak counter-rotation just within the main collar of the Great Red Spot.

The larger part of the thesis is concerned with a much more advanced, time dependent, primitive-equation model. This is referred to as the Jupiter general circulation model (GCM), as it was developed from the terrestrial GCM in use at the UK Meteorological Office. This model runs at high resolution (0.5° horizontally and 10 levels per decade in pressure) over a limited area of the planet. Strictly speaking our version is not a true GCM, since it does not predict the global general circulation, but the label is convenient. The atmosphere is stratified with a realistic Jovian temperature profile, spanning a pressure range from 10 bars to 1 mb.

The physical parametrizations in the current version of the model are very simple, being restricted to linear drag, Newtonian cooling and dry convective adjustment. These are sufficient to allow us to simulate some of the main observed features, such as jets and eddies, and wave-mean flow interaction in the equatorial region. The next stage of work will concentrate on replacing these parametrizations with more realistic calculations.

The GCM has been used to study two separate aspects of jovian meteorology. The first is the behaviour of midlatitude vortices and their interaction with jetstreams, showing how vortices tend to migrate and intensify by frontogenesis, unless this migration is blocked by sufficiently strong shear in the zonal jets. Barotropic and baroclinic instability of the jets can generate eddies with different vertical structures, both shallow and deep. These vertical structures tend to resemble closely the gravest one or two normal modes.

The second area is that of equatorial waves, their propagation through the atmosphere and interaction with the equatorial jet. We have studied the effects of different wave-sources, such as might exist below the cloud-tops. Finally we show how the combination of a wave source with weak symmetric solar heating could explain the meridional profile of the observed equatorial jet.

The rest of this chapter forms a survey of the literature covering observations of Jupiter's atmospheric dynamics. We present a summary of spacecraft observations in historical order, showing how each mission has advanced our knowledge. Some more specialized topics, such as the Great Red Spot and White Ovals, and also equatorial waves, will be left until later, where they are reviewed alongside our model results. We shall review published theoretical and modelling work in Chapters 5 and 6, to put our own results into context.

This introduction concludes with a plan of the rest of the thesis and introduces the contents of each chapter.

Terminological note

Throughout this thesis, the words "day" and "year" should be taken as referring to 24-hour Earth days and 365-day Earth years. This is for reasons of convenience: for example the GRaDS data visualization package assumes 24-hour days. In the current model this is not a drawback, since there is no diurnal or other temporal dependence in the forcing terms. However, in future versions incorporating a diurnal cycle, it will become necessary to adjust the frequency of output. This should be straightforward.

All latitudes referred to are planetographic, and longitudes are in System III. Vertical levels are specified in terms of mean pressure, in units of bars or millibars. We tend to refer loosely to the region between the upper cloud-tops at $p \simeq 500$ mb and the tropopause at $p \simeq 120$ mb as the upper troposphere, and the region below the clouds as the lower or deep troposphere.

The terms "westerly", "eastward" and "prograde" are used interchangeably to describe motion from west to east. We normally use "westerly" when referring to winds and "eastward" when referring to wave propagation. Similarly, "easterly", "westward" and "retrograde" share the opposite meaning. Potential vorticity means the Ertel form, usually plotted on a pressure surface. We recognise that it would be more logical to specify Ertel PV on isentropic surfaces, but for comparison purposes it is more convenient to interpolate onto pressure surfaces.

1.2 Observations of Jupiter

> *"What I want is,* facts." (Dickens: *Hard Times*)

In this section we shall briefly describe the spacecraft which have visited Jupiter, and summarize the most important observational results concerning the dynamics of the neutral atmosphere. We start with ground-based and Hubble Space Telescope observations, then look at each mission in turn, to show how each has increased our knowledge. We will not go into great detail about observations of giant eddies and equatorial waves. These are covered in Chapters 2 and 6 respectively, to provide an experimental context for each modelling study.

Ground based observations

Jupiter is the brightest body in the midnight sky except for the Moon, and has been observed with the naked eye since ancient times. The first telescopic observations were made by Galileo, who discovered the four largest Jovian moons in January 1610. Other discoveries soon followed: a description of the zone-belt structure was published in 1630 by Niccolo Zucchi. The shadows of the moons were observed in 1643 by Riccioli. The rapid rotation of the planet was revealed on 9 May, 1664 when Hooke observed "a small spot in the biggest of the three obscurer belts of Jupiter". He noted its rapid motion: "...he found that, within two hours after, the said spot had moved from east to west, about half the length of the diameter of Jupiter." (Hooke, 1665).

The following year, Giovanni Cassini began a 29-year series of continuous observations of a "permanent spot". It is likely that Cassini was observing what is now known as the Great Red Spot (GRS), although there are some discrepancies between Cassini's reports and modern data. Cassini also discovered the fast equatorial jet, and co-discovered Jupiter's oblateness (a factor of 14/15 in radius) and limb darkening. He observed changes in the zones and belts, and the formation of new spots. This was the first evidence that the visible features are actually clouds, as opposed to a solid surface.

In 1831 Schwabe published a drawing of the Great Red Spot Hollow; since then, the Spot and its Hollow have been observed continuously. The development of silver-coated telescope mirrors, later in the nineteenth century, led to the discovery of the South Tropical Disturbances, and in the 1930s, the White Ovals.

Systematic photography of Jupiter began in 1904 at Flagstaff, Arizona, and has continued through most of the 20th Century. Amateur astronomers such as Bertrand Peek (director of the BAA Jupiter Section from 1934–1949) contributed much useful data (Peek, 1958; Ingersoll and Cuzzi, 1969). The internal circulation within the GRS was discovered and photographed by Reese and Smith (1968).

The subsequent introduction of charge-coupled devices, improved colour filters, and latterly the Hubble Space Telescope, have provided images suitable for comparison with spacecraft data. Examples include the HST images of the Galileo Probe entry site and of the impact of comet Shoemaker-Levy 9. Ground-based infrared observations have also been made, including investigations of stationary thermal waves (Deming et al., 1997, 1989).

The Pioneer missions

Six spacecraft have visited Jupiter so far: Pioneers 10 and 11, Voyagers 1 and 2, Ulysses, and Galileo. The Pioneers were relatively simple craft, designed primarily to investigate fields and particles in interstellar space. They were spin-stabilised and carried two instruments for observing the Jovian atmosphere: an imaging photopolarimeter and an infrared radiometer.

Pioneer 10 was launched aboard an Atlas-Centaur booster on 2 March 1972, to become the first spacecraft to negotiate the asteroid belt and reach the outer Solar System. It provided measurements of Jupiter's magnetosphere, surviving the intense flux of charged particles ($\sim 10^6$ times as intense as the Earth's van Allen belts) despite some temporary damage to its electronics. Until the final 24 hours before encounter, the quality of images was inferior to those from ground-based telescopes. However, infrared spectra and radio occultation measurements provided atmospheric temperature profiles and compositional data.

Pioneer 11 was launched on 5 April 1973. It detected oscillations in the Jovian magnetosphere, and passed closer to Jupiter (42000 km) than Pioneer 10 had done. It overflew the polar regions and then visited Saturn, discovering new moons and sounding the Saturnian atmosphere for the first time. Both Pioneers were tracked until the mid 1990s, in an attempt to detect the gravitational influence of a possible tenth planet. Reviews of the results of the Pioneer studies of Jupiter may be found in Gehrels (1976).

Voyagers 1 and 2

The Voyagers, a development of the Mariner Venus probes, were much more sophisticated spacecraft than the Pioneers. Their brief flybys of Jupiter are arguably amongst the most scientifically productive missions ever. They were inertially stabilised, which allowed more complex instrumentation to be carried. The Imaging Science Subsystem (ISS) (Smith et al., 1977) comprised a wide-angle (3.1°) and also a narrow-angle (0.424°) camera, with a selection of coloured filters to pass wavelengths from the 619 nm methane band down to 345 nm in the ultraviolet. Cloud-tracking studies of ISS images have provided data on the horizontal cloud-top winds. Limaye (1986) published estimates of the zonal-mean winds. Several studies (e.g. Mitchell et al., 1981; Dowling and Ingersoll, 1988; Sada et al., 1996) have used Voyager images to analyse the flow in the GRS. The implications of these studies will be examined in Chapter 2, Section 2.2.1.

Another useful dataset was obtained by IRIS, the Infrared Spectrometer and Radiometer Subsystem (Hanel et al., 1977). This had a field of view of 0.25° and had much higher spatial and spectral resolution than the Pioneer infrared instrument. IRIS revealed the

thermal structure of the atmosphere above the clouds (Hanel et al., 1979a,b). Through the thermal wind equation, this has provided information about the dynamics (Flasar et al., 1981; Flasar, 1986). Each Voyager also carried an ultraviolet spectrometer and photopolarimeter, and five fields and particles instruments.

Other important findings of the Voyager missions included the latitudinal variation of the infra-red emission (Pirraglia, 1984). The emitted flux seems much more uniform than the absorbed solar flux, implying either dynamical heat transport in the meridional direction, or else a meridional variation in the interior heat flux.

Radio occultation measurements (Lindal et al., 1981) provided four vertical temperature profiles. Two of these contain evidence of equatorial wave activity (see Chapter 6). The Voyager missions also discovered Jupiter's ring, three new moons, active vulcanism on Io; and they returned the first close images of Europa's surface.

Ulysses

The NASA/ESA Ulysses project was intended principally to explore the polar regions of the Sun. It was launched on 6 October 1990, and flew past Jupiter in order to slingshot out of the ecliptic plane. It was the first vehicle to do this. Ulysses is solely a "fields and particles" experiment, and so provided no data of relevance to atmospheric dynamics. However, it did make detailed measurements of the Jovian magnetosphere.

Galileo

Galileo comprised two spacecraft: the Orbiter and Probe. The mission plan was for the Probe to enter the atmosphere and measure its temperature, pressure, cloud structure, and composition. The Orbiter has spent two years in the vicinity of Jupiter, investigating the planet and its moons in much more detail than Voyager achieved. Galileo was launched, after long delays, on the Space Shuttle Atlantis on 18 October 1989. Galileo's very long trajectory meant that the journey to Jupiter lasted six years. This is because the small Inertial Upper Stage booster was only powerful enough to send Galileo to Venus. Accordingly, slingshots around Venus and then Earth were used to gain speed before sending the craft to Jupiter.

To survive a journey as close to the Sun as the orbit of Venus, the spacecraft required extra heat shielding. In the course of the modifications, it was transported across the United States by road several times, and this extra handling probably caused a loss of lubricant in three of the 18 ribs supporting its high-gain antenna. On 11 April 1994, the spacecraft was commanded to open this antenna, but three ribs stuck to the central

pole, preventing the antenna from opening properly. This meant the craft had to use its low-gain antenna only, which could only transmit from Jupiter at a nominal 10 bits per second, instead of the hoped-for 134000 bps from the high-gain antenna.

Methods of continuing the mission without the high-gain antenna were designed, including data compression, improved sensitivity of the DSN receivers, (together achieving an effective bitrate of 100 bps) and use of the onboard tape recorder. Unfortunately the tape recorder later malfunctioned, refusing to obey a command to stop rewinding. This problem was later solved, leaving the tape recorder almost fully operational.

Results from the Galileo Probe

The Galileo Probe was released from the Orbiter in July 1995, and entered the atmosphere five months later, on 7 December, at 6.53°N planetocentric latitude and 0.94°W System III longitude (Young, 1998). Nine separate experiments were conducted during the descent, complemented by Earth-based observations (Orton et al., 1996). These remote-sensing results showed that the Probe entry site was a so-called 5μm hot spot: an almost cloud-free region, characterised by enhanced thermal emission from the planet's interior, and assumed to be an area of atmospheric downwelling.

The initial scientific results from the Probe were published in *Science* **272**, pp. 837-880. A second set of papers followed in September 1998, published in *JGR Planets*. For an overview of these, and bibliography of other papers, see Young (1998). Here we shall summarize only those Probe results directly relevant to this thesis. This principally means data on winds and static stability. Other interesting results, regarding clouds, composition, lightning etc. are only mentioned in passing.

A temperature-pressure profile was obtained by the Atmospheric Structure Instrument (Seiff et al., 1996; Seiff, 1998). This was very close to a dry adiabat down to 15 bars, consistent with the lack of water suggested by the remote observations and confirmed by the neutral mass spectrometer. Reanalysis has not changed this result, although there is a tentative suggestion of increased static stability (up to 0.5 K/km) below 15 bars. This is as yet unconfirmed, because the sensor was outside its calibrated temperature range.

The nephelometer instrument (Ragent, 1998) detected what was probably the tenuous lower part of an ammonia cloud between 0.46 bar (the first valid data point) and 0.55 bar. A denser cloud was detected between 0.76 bar and 1.34 bar. This is compatible with the expected altitude for an ammonium hydrosuphide cloud (West et al., 1986). The base of the cloud was more sharply defined than the top. A third, weak structure was observed between 1.9 to 4.5 bars which may have been a water cloud. These observations are compatible with thermochemical equilibrium models (West et al., 1986). All the clouds

were less dense than predicted, but this agrees with 5 μm observations which show that the hot-spots are relatively cloud-free regions.

The helium mass fraction was measured by von Zahn and Hunten (1996) to be only slightly below the presolar value, indicating that little separation of helium from hydrogen has so far occurred, in contrast to the probable situation on Saturn.

Revised results of the Doppler Wind Experiment (Atkinson et al., 1997, 1998) showed that the wind speed increased with depth, from 90 m/s at 700 mb, to 170 m/s at 4 bars. Below this level, the winds were roughly constant down to 21 bars, implying that the zonal flow extends much deeper than the base of the water cloud. Atkinson et al. did not speculate either on the reason for the large baroclinicity at the 1–4 bar level, nor on the relative importance of internal heating. However, it seems that the downwelling at the Probe entry site extends down to at least 4 bar, which is consistent with the absence of a substantial water cloud.

Hot-spots similar to the Probe entry site cover only 1% of Jupiter's surface (but $\simeq$ 15% of the area of the NEB), so we need to ask, to what extent are these results typical of such a heterogeneous planet? Young (1998) pointed out that the radiative time constant in the deep atmosphere (of order 10 years or longer) is much greater than the lifetime of a hot spot (up to one year), and so the statically stable profile is unlikely to be restricted to the hot-spot. It may be widespread throughout the equatorial region.

The link between water depletion and a stable, radiative layer above the 400 K (or 18 bar) level was predicted by the modelling work of Guillot et al. (1994a), who found that the water abundance dominates the infrared opacity in this region. Guillot et al. concluded that the radiative "window" probably does not exist, because of the assumed presence of liquid water at this level. If the new Galileo Probe evidence for a dry atmosphere were typical of the entire equatorial region (not just a hot-spot), then Guillot's work would suggest the opposite conclusion. This seems unlikely though.

Instruments aboard the Galileo Orbiter

The atmospheric science objectives of Galileo, together with details of the instrumentation, are discussed by Hunten et al. (1986). The orbiter carries four instruments of relevance to atmospheric dynamics studies. These are: the Solid State Imaging System (SSI), the Near-Infrared Mapping Spectrometer (NIMS), the Photopolarimeter Radiometer (PPR), and Ultraviolet Spectrometer (UVS). The instruments are designed to complement each other, and together they are able to sound the atmosphere across the spectrum from the ultraviolet to the far infrared.

The SSI is a narrow-angle camera (budgetary constraints prevented the inclusion of a wide-angle imager also). Its telescope is an improved version of the design used on Mariner and Voyager, and is connected to a solid-state CCD detector, which is much more sensitive than the older cameras. Its spectral range extends to 1000 nm, in the near infrared. The SSI is intended to identify wave modes and separate these from larger scale motions, to study vertical structure of clouds, and to measure the mean vertical wind profile in certain regions.

NIMS is an infrared grating spectrometer which measures radiances in the interval 0.7μm to 5.7μm. NIMS measures temperatures in the region from the opaque clouds up to about 1 bar, maps cloud distributions and provides infrared images (for example, in the 5μm window), which complement those taken by the SSI. Work in progress in Oxford includes analysis of the scattering properties of the clouds. Composition and altitude of the clouds are retrieved from the spectra, in an attempt to fit the new data to chemical and thermodynamic models of the atmosphere.

The objectives of the PPR, which is sensitive to even longer wavelengths, are to measure the radiation budget, temperature profiles from 500 mb to 100 mb, and also cloud and haze properties. The UVS experiment, based on an Ebert-Fastie spectrometer, measures radiances at wavelengths down to 115 nm. This is used to study (among other things) stratospheric composition, and clouds and haze in the upper troposphere.

Galileo imaging results

The SSI has imaged a number of interesting features, including the Great Red Spot, the north and south equatorial regions, a hot spot, and two White Ovals. The results are presented by Vasavada et al. (1998), Banfield et al. (1998) and Simon et al. (1998). The spatial resolution is similar to the best Voyager images, but the spectral and temporal resolution is much better. We will defer detailed discussion of the Galileo observations of the GRS and White Ovals until Section 2.2.1, where we will place them in the context of the Voyager results. Here we shall briefly mention some of the other features observed.

Observations of the south equatorial jet (Vasavada et al., 1998) revealed nearly pure zonal motion, a sharp velociy peak near 6.5°S, and a maximum velocity of 150 m/s. This compares with the Voyager value of 128 m/s (Limaye, 1986). However, a reanalysis with two Voyager 1 images gave a value of 150 m/s, showing that Limaye's measurements may have underestimated the value by failing to resolve the sharp peak.

While the peak zonal wind at 6.5°S seems to have remained constant since Voyager, that on the equator appears to have slowed by more than 20 m/s (Vasavada et al., 1998). This could be due either to fluctuations in the Equatorial Zone or to Voyager and Galileo

sampling different heights.

A complicated region with a hot-spot adjacent to a cloudy cyclone was imaged around 7°N. Strong southwesterly flow was observed towards the hot spot, but little flow away from it. A possible interpretation is massive convergence and downwelling (dubbed "the mother of all downdrafts"), with vertical velocity of up to 1 m/s. Alternatively, the flow rotating in the cyclonic feature may advect clouds toward the edge of the hot-spot, where they evaporate and can no longer be tracked. This would imply some downwelling (to remove the clouds) but suggests that there is flow out of the hot spot, although it has few clouds and so cannot easily be tracked.

Galileo NIMS results

NIMS has recorded spectra from several features including the Great Red Spot, although technical difficulties with the retrieval process have meant that most of these results have not yet been published. The first four spectra to be returned (the G1 real-time spectra) were taken from a hot-spot in the NEB. These spectra were recorded across NIMS's entire wavelength range, and are the first ever measurements of near-infrared and 5 μm spectra made outside the Earth's atmosphere. Cloud height and composition data have been retrieved in Oxford (Irwin et al., 1998). Owing to the failure of the high-gain antenna, subsequent spectra were sampled over a more restricted spectral range. This caused ambiguities in the results which must be resolved by reference to the real-time spectra.

The 5 μm opacity in the real time hot-spot retrievals is best explained by a variable layer of large NH_4SH particles, underlying the main thin reflecting layer of small NH_4SH particles. This could be at a level of between 1.5 and 4 bars, so might be water instead of ammonium hydrosulphide. Future analysis of data from multiple angles should resolve the question. This cloud is consistent with one detected by the probe nephelometer as mentioned above.

NIMS found somewhat conflicting evidence for downwelling in the hotspot. On the one hand, the ammonia volume mixing ratio was anticorrelated with the 5 μm brightness, suggesting downwelling from above. On the other hand, there was a positive correlation between the relative humidity of water and the 5 μm brightness. Irwin et al. ascribed this puzzling result to the presence of a wet "collar" around the hot-spot, which was found by Roos-Serote and authors (1998), who also noted "off-centre" water distributions in two of the four spectra. The spatial distribution of water appears to be rather complex when viewed at high resolution, in contrast to the relatively simple situation suggested by Voyager IRIS results. Future work on coordinated observations by NIMS and SSI is likely to reveal more about motions within the hot spots, and indeed about other features.

1.3 Contents of this thesis

We begin with a description of the potential vorticity inversion model and its results; this is self-contained, although we frequently refer back to its predictions in later chapters. The next two chapters describe the GCM and methods for analysing vertical structures in the model. Next, we present the results in two much longer chapters, one on jets and vortices and the other on equatorial waves and jets. The final chapter has conclusions and details of the future work planned.

Chapter 2 covers the PV inversion model, its results and conclusions. At this point we review the main observations of giant eddies, to place the model in context. We compare these observations with our (idealized) simulations of the Great Red Spot and a White Oval. For each case we show both a shallow, confined vortex and a deep one extending to the base of the model at 100 bars. For completeness, we also look at a generic jovian cyclone.

Chapter 3 is a technical description of the GCM, and how it differs from the standard UKMO Unified Model on which it is based. We justify our choices of model domain and resolution, the boundary conditions and physics parametrizations. We also review a quasi-geostrophic energy budget analysis due to Lorenz, which is used in Chapter 5.

Chapter 4 goes on to describe an eigenmode analysis of the atmosphere's vertical structure. This one-dimensional model is used to identify the principal normal modes of motions in the atmosphere, and provides a useful tool for initializing the GCM fields, and for interpreting the results.

Chapter 5 is a survey of the properties of vortices and jets in the midlatitudes. We begin with the simple case of a single Gaussian vortex in an atmosphere with no background flow. The evolution of such a vortex is much more complicated than we might at first expect. It is rather difficult to tune the model parameters (such as horizontal diffusion) to achieve a vortex which drifts for a long period without either dissipating or intensifying. The process of intensification is itself complex, for there is no external energy source driving the vortex. The potential vorticity anomaly associated with the vortex tends to shrink and steepen, which accelerates the winds.

A wide range of experiments are presented, looking at the effects of different vertical structures for vortices. We find that the model tends spontaneously to produce eddies (typically when the initial one disintegrates) and that these almost always have the same vertical structure as the first baroclinic mode calculated in Chapter 4.

Having achieved some understanding of how the model handles isolated vortices in a resting atmosphere, we go on to look at the interaction of vortices with jetstreams. Initially

the jets are chosen with idealised properties (e.g. a pair of equal and opposite sinusoidal jets). The model behaviour now becomes more complicated, for the jets not only block the migration of the vortex (at least temporarily) but can also produce fresh vortices through spontaneous instabilities. Interestingly, the model is capable of generating both shallow and deep eddies (i.e. resembling the two cases studied in Chapter 2). The profiles are either the first baroclinic mode, or else similar to the basic flow, which is barotropic in the deep atmosphere but decays with height in the stratosphere). We use the energy diagnostics introduced in Chapter 3 to elucidate the processes acting to create and destroy eddies.

Chapter 6 is a study of equatorial waves and their interaction with the zonal-mean flow. The core of the work is an attempt to explain the presence and observed structure of the superrotating equatorial jet, in terms of forcing by wave sources located in the deep atmosphere.

After explaining the motivation for this work in terms of the observations, we briefly review the relevant theory of equatorial waves. We then describe the development of a version of the GCM which produces plausible simulations of the theoretical properties of monochromatic travelling waves. We show how Kelvin wave forcing can (as predicted by the theory) excite a westerly jet centred on the equator.

To discover more about any constraints there may be on the wave forcing, we look at the wave-transmission properties of the stratosphere. It turns out that waves are selected primarily by phase speed, and then by wavenumber. This implies that the observed jet may be more a function of the atmospheric filtering below it than of the properties of the wave sources.

We simulate flows generated by several different stationary, coherent wave sources, and show that a plausible source representing some form of organized convection in the deep atmosphere could, in the model, deliver zonal angular momentum to the jet seen at the cloud-top.

The last part of Chapter 6 considers the effect of a Hadley circulation on the meridional structure of the equatorial jet. We argue that axisymmetric solar heating at the clouds could force a subsidiary circulation which transports angular momentum away from the equator. Weak heating is sufficient to modify the jet structure to resemble closely the one observed on Jupiter.

Finally, Chapter 7 summarizes the conclusions from this work, and outlines future developments and experiments planned for the second stage of this project.

Chapter 2

Potential vorticity inversion models of giant eddies

The most common approach in the past, when modelling giant eddies such as the Great Red Spot (GRS), was to consider the horizontal structure in some detail, while simplifying the vertical structure. This practice is evident in all the numerical models discussed in Chapter 5.2. The model to be described in this chapter adopts the opposite approach: we simplify the horizontal structure in order to allow a more detailed description of the vertical structure. It is therefore a useful complement to other work, and points the way toward the 3-D general circulation model with which this thesis is principally concerned.

The method considers an atmospheric flow as evolving slowly through a series of balanced states, so in principle the time-dependent equations could be integrated by "inverting" the PV distribution to obtain the other fields, then advecting the PV and repeating the process. The present model only performs the PV inversion part of this procedure. It is therefore solely diagnostic, with no time dependence.

We include here only a brief discussion of PV inversion, since it is a well-known technique in the field of terrestrial meteorology. For a detailed account, see Hoskins et al. (1985). The method relies on the assumption that the atmosphere is always close to a balanced state. Taking the example of geostrophic balance, this is justified for synoptic-scale, low-frequency, extratropical flows in the Earth's atmosphere. It would seem to be appropriate also for long-lived, slowly varying features on Jupiter. If this were not true, features such as the Great Red Spot would visibly adjust toward equilibrium. The fact that they do not, but instead retain their shape over timescales much longer than the advective timescale associated with their size and the observed winds, implies that they are balanced.

The concept of balance is more precise than suggested above, where "balance" was used loosely in the sense of "equilibrium". A better expression of this rather subtle idea is

the definition of exact balance for a linear primitive-equation model. The motions may be projected onto a set of "slow" modes, e.g. synoptic scale motions, and "fast" modes, e.g. gravity waves. The motion can be said to be in exact balance if the slow basis set is orthogonal to the fast one. This means that any model state vector, consisting of a superposition of slow modes, has zero amplitude if projected onto the fast modes. Then the only fast modes present in a model result from the initialization, and cannot arise later. Geostrophic balance is the simplest example of a balance condition; of course it is not a very accurate description of the atmosphere in general. It is not known for certain whether exact balance actually exists in the general case.

This requirement for balance is one of the necessary conditions for PV inversion. Others are that the pressure on each isentropic surface is defined, the PV fields within the model domain are specified, and suitable boundary conditions are specified. The PV invertibility principle states that these conditions are necessary and sufficient for the calculation of all other dynamic and thermodynamic fields within the domain.

The crucial idealization in this approach is the assumption of axial symmetry. The model is of a cylindrical, symmetric vortex, under the 'f-plane' approximation. Under circular symmetry, the full primitive equations reduce to an exact expression for balanced motion, known as gradient-wind balance.

Two further approximations are made in our model: radial velocities are neglected, and the equations are solved for steady flow only, with no time dependence. Neither of these approximations is fundamental to the PV approach. In principle they can be relaxed, which allows the study of the divergence and of vertical motion, for which observational data exist. A jovian example is the divergence field in the Great Red Spot, measured by Sada et al. (1996). The design of the model, and the code used, is a development of work by Norton (1988) which included a diagnostic model of a circularly symmetric midlatitude terrestrial cyclone.

Norton's model is formulated using an isentropic vertical coordinate, which leads to a simple form of the model equations. This is well suited to the original problem, but poses a complication when moving to the jovian version. We require a nonuniform grid to accommodate the much larger difference in gradient of θ (i.e. static stability) between the troposphere and stratosphere. This grid-stretching has made the model more difficult to program, less efficient, and has necessitated careful monitoring of errors. Nevertheless, isentropic coordinates are a natural system for Jupiter, and have been used in a GCM (Dowling et al., 1998). The stretched grid works well in practice.

The axisymmetric PV-inversion approach has several drawbacks compared to other current methods. The lack of time dependence prevents any consideration of the processes responsible for generating or maintaining the PV anomaly structure. Important asymmetric properties, such as the lower average velocity observed in the south side of the

GRS collar compared with the north side, are beyond the scope of this model; so too are smaller features such as the small cyclones within the GRS interior, reported by Sada et al. (1996). All the relevant observations of jovian features will be covered in detail in Section 2.2.1.

However, the earlier analysis of Voyager images by Mitchell et al. (1981) makes several of the assumptions adopted here: the radial component of velocity is neglected, as is the angular dependence of the azimuthal component. Thus the horizontal structure in Mitchell et al. (1981) is suitable for comparison with the axisymmetric model. We should point out one limitation straight away. The unrealistic circular geometry imposed here means that the relation between velocity and relative vorticity is not quite the same as in the realistic, elliptical case. When initializing the model, and interpreting the results, care is taken to distinguish between the relative vorticity in the circular coordinate system, and the equivalent relative vorticity, for the same wind field in an elliptical vortex. It is the latter which should be compared to observations.

The benefits of this approach include a detailed representation of the vertical structure of the vortex, using up to 160 levels between pressure surfaces of 100 bar and 10 mb. The model is fully and realistically stratified, so the variations with height of winds and temperatures can be examined. The PV approach provides a conceptually simple and useful framework for the study of balanced flow.

There follows a description of the mathematical basis for the approach taken, and a discussion of the necessary boundary conditions and the numerical scheme used. We then describe the details of the application of the method to a model of the GRS, the methods used to validate the model and estimate errors, and a comparison between the model predictions and experimental observations. We then discuss a similar model of White Oval followed by (more speculatively) a model of a cyclone, not based on any particular Jovian feature. Cyclones are certainly observed on Jupiter, the most obvious examples being the NEB Brown Barges. Unfortunately, this approach is not really suitable for modelling a brown barge, mainly because they are far more elongated (in fact, cigar-shaped) than the giant anticyclones. The validity of the axisymmetric assumption is therefore doubtful. Another practical problem is that quantitative wind observations have not been made for the brown barges, rendering any detailed attempt to model them rather speculative. Instead, we will content ourselves with modelling a pair of hypothetical, shallow and deep cyclones, to demonstrate that the PV inversion method is not restricted to the special case of anticyclones.

2.1 Mathematical formulation

2.1.1 Overview of the approach

A review of the subject of potential vorticity (PV) inversion is given by Hoskins et al. (1985). The present work is based on a technique for isentropic inversion of Ertel PV, due to Norton (1988). The use of an isentropic coordinate system, with potential temperature as the independent vertical variable, enables inversion by the straightforward solution of an elliptic equation.

The principle of PV invertibility states that, under certain conditions, a knowledge of the PV fields within a model domain permits the calculation of all other dynamic and thermodynamic fields. The necessary conditions for isentropic PV inversion are that the PV and pressure are specified on all isentropic surfaces, and that the atmosphere satisifes a balance condition. Suitable boundary conditions are also required.

In the present work the two major constraints are that the vortex is assumed to be axisymmetric, and that time dependence is excluded, the models being a diagnostic tool only. The assumption of axial symmetry has two advantages. Firstly a simple condition for exact balance exists, namely gradient-wind balance. Secondly it allows a detailed study of the vertical structure of the flow, using only modest computer resources.

We make the further approximation of neglecting the radial component of velocity. This is not a fundamental limitation of our approach. The work described here could be extended to lift this restriction, allowing the study of divergence and of vertical motion. The radial advection of PV could be used to introduce time dependence into a future model. These assumptions are similar to those made in the earlier observational studies, e.g. Mitchell et al. (e.g. 1981). Their validity and suitability for Jupiter will be discussed in Section 2.2.1 below.

2.1.2 The model equations

The formulation, following Norton (1988), is based on the primitive equations of meteorology expressed in isentropic coordinates. For comparison with observations, we also adopt a log pressure vertical coordinate, roughly equivalent to geometric height:

$$z \equiv -H_0 \ln\left(\frac{p}{p_0}\right) \tag{2.1}$$

where H_0 is the scale height, assumed constant. In isentropic coordinates the hydrostatic equation takes the form (see, e.g., Andrews et al., 1987):

$$\frac{\partial M}{\partial \theta} = c_p \exp\left(\frac{-\kappa z}{H_0}\right) \tag{2.2}$$

where the potential temperature $\theta \equiv T(p_0/p)^\kappa$ is used as an isentropic vertical coordinate. $M \equiv c_p T + \Phi$ is the Montgomery streamfunction (Φ being the geopotential), and $\kappa \equiv R/c_p$ where R is the universal gas constant and c_p the specific heat at constant pressure. We use cylindrical coordinates (r, z) and assume axial symmetry. Differentiating (2.2) at constant r yields

$$\left(\frac{\partial^2 M}{\partial \theta^2}\right)_r = \frac{-\kappa c_p}{H_0}\left(\frac{\partial z}{\partial \theta}\right)_r \exp\left(\frac{-\kappa z}{H_0}\right) \tag{2.3}$$

It is useful to define the quantity s as

$$s \equiv \left(\frac{\partial z}{\partial \theta}\right)_r \tag{2.4}$$

Substitution of (2.4) into (2.3) gives a relation between s and M:

$$s = -\frac{H_0}{R}\exp\left(\frac{\kappa z}{H_0}\right)\left(\frac{\partial^2 M}{\partial \theta^2}\right)_r \tag{2.5}$$

where the gas constant $R \equiv c_p \kappa$. The isentropic density σ is usually defined by

$$\begin{aligned} \sigma &= -\frac{1}{g}\left(\frac{\partial p}{\partial \theta}\right)_r \\ &= \frac{p_0}{g H_0}\exp\left(\frac{-z}{H_0}\right) s \\ &= \rho s \end{aligned} \tag{2.6}$$

The quantity s is a more convenient variable than σ, and hereafter we shall loosely refer to s as the 'isentropic density'. Note that s is inversely proportional to the usual static stability parameter.

We assume a reference state that is at rest relative to a frame rotating with the planet's interior. It is convenient to separate M linearly into reference and disturbance parts: $M = M_0(\theta) + M_d(r, \theta)$. If we let $z_0(\theta)$ be the reference height of isentropic surface θ, then from (2.5), the reference isentropic density is

$$s_0 \equiv -\frac{H_0}{R}\exp\left(\frac{\kappa z_0}{H_0}\right)\left(\frac{\partial^2 M_0}{\partial \theta^2}\right)_r \tag{2.7}$$

We can then eliminate the reference Montgomery potential and write

$$\begin{aligned} s = &-\frac{H_0}{R}\exp\left(\frac{\kappa z}{H_0}\right)\left(\frac{\partial^2 M_d}{\partial \theta^2}\right)_r \\ &+ s_0(\theta)\exp\left(\frac{\kappa}{H_0}(z - z_0)\right) \end{aligned} \tag{2.8}$$

Ertel's potential vorticity (PV) in isentropic coordinates is given by

$$P = \frac{f + \zeta_\theta}{\sigma} \tag{2.9}$$

where $f = 2\Omega \sin \phi_0$ is the Coriolis parameter and ζ_θ is the relative vorticity on an isentropic surface. For the present purposes, it is more helpful to consider the related quantity Q:

$$Q \equiv \frac{f + \zeta_\theta}{s} = \rho P \tag{2.10}$$

The reason for this is that Q varies more slowly than P with height, which makes computations more straightforward. We shall make the linear separation $Q = Q_0(\theta) + Q_d(r, \theta)$, referring to Q_0 as the background or *reference* PV ($= f/s_0$), and Q_d as the PV *anomaly*. In axisymmetric cylindrical geometry, the relative vorticity is

$$\zeta = r^{-1} \frac{\partial}{\partial r}(vr) \tag{2.11}$$

where v is the azimuthal wind velocity and the radial wind is neglected. Rearranging Equation (2.10) and substituting from (2.8) gives

$$\begin{aligned} f + \zeta_\theta &= (Q_0 + Q_d)\left(-\frac{H_0}{R}\exp\left(\frac{\kappa z}{H_0}\right)\left(\frac{\partial^2 M_d}{\partial \theta^2}\right) + s_0 \exp\left(\frac{\kappa(z - z_0)}{H_0}\right)\right) \\ &= -Q\frac{H_0}{R}\exp\left(\frac{\kappa z}{H_0}\right)\left(\frac{\partial^2 M_d}{\partial \theta^2}\right) + Q_0 s_0 \exp\left(\frac{\kappa(z - z_0)}{H_0}\right) + Q_d s_0 \exp\left(\frac{\kappa(z - z_0)}{H_0}\right) \end{aligned} \tag{2.12}$$

In the undisturbed reference state, where $M_d = Q_d = \zeta_\theta = 0$, the only contribution to the vorticity is from f. Equation 2.10 therefore implies that

$$f = Q_0 s_0 \exp\left(\frac{\kappa(z - z_0)}{H_0}\right) \tag{2.13}$$

and we can use this to cancel two terms in Eq. (2.12), to leave

$$\frac{1}{r}\frac{\partial}{\partial r}(vr) + Q\frac{H_0}{R}\exp\left(\frac{\kappa z}{H_0}\right)\left(\frac{\partial^2 M_d}{\partial \theta^2}\right) = Q_d s_0 \exp\left(\frac{\kappa(z - z_0)}{H_0}\right) \tag{2.14}$$

This is the equation to be solved, constrained by a suitable balance condition. The circular symmetry of the problem allows us to use the gradient-wind balance equation, which may be written as

$$\frac{v^2}{r} + fv = \left(\frac{\partial M_d}{\partial r}\right)_\theta \tag{2.15}$$

This equation represents the balance between the Coriolis, centrifugal and pressure-gradient forces. It may be thought of as a combination of the geostrophic and cyclostrophic balance equations, and is more accurate than either alone. This is particularly useful for

our particular application, since the Rossby number in the GRS may be as large as $0{\cdot}36$ (Mitchell et al., 1981). Therefore the Red Spot is intermediate between the geostrophic and cyclostrophic cases. In three-dimenisonal geometry we would be unable to use the gradient-wind balance equation is such a convenient form, so this represents another advantage of circular symmetry.

Equations (2.14) - (2.15) are solved numerically, given suitable boundary conditions and PV distributions, to give values for M_d and v at all grid points. Knowledge of the Montgomery potential M then allows the calculation of all other thermodynamic fields within the domain.

2.1.3 Boundary conditions

Norton (1988) showed that the condition of mass continuity may be replaced by a suitable boundary condition; this simplifies the formulation of the problem. Without reproducing a fulll derivation here, we shall outline the argument. Norton considered the case in which the PV anomaly is confined to a region much smaller than the full domain ($r << R$), and the domain radius R is much larger than the largest Rossby deformation radius. The asymptotic solution was derived, taking explicit account of the mass conservation equation. Under these circumstances, the solution for r much larger than the Rossby radius has the approximate form

$$M_d = c_0 \log r + d_0 \tag{2.16}$$

and hence, by Eq. (2.15), the velocity varies inversely with r. This means that for large enough R, adopting the boundary condition $M_d(r=R) =$ constant is equivalent to inverting the PV in a much larger domain using the mass conservation condition.

With this in mind, we impose the following boundary conditions. The isentropes at the top and bottom of the model are fixed at constant pressures. Mathematically,

$$\left(\frac{\partial M_d}{\partial \theta}\right)_r = 0 \text{ at } \theta = \theta_0 \text{ and at } \theta = \theta_{max} \tag{2.17}$$

To replace the continuity equation, we also require that the flow should be quiescent at the outer boundary:

$$M_d = 0 \text{ at } r = r_{max} \tag{2.18}$$

Finally, in an axisymmetric vortex, v must vanish at the origin.

In addition to the boundary conditions, the values of the reference and anomalous PV fields on each isentrope are sufficient constraints to form a well posed problem.

2.1.4 Numerical scheme for PV inversion

The numerical scheme used for solving Equations (2.14)–(2.18) is a development of that used by Norton. The equations are discretized using a finite difference scheme. Equation (2.14) is inverted iteratively by Newton's method, until the balance condition (2.15) is satisfied to a desired precision. The use of an iterative (rather than direct) method is dictated by the non-linear term in (2.15).

The present formulation differs from that of Norton's model in two important respects. Firstly, Norton used the vertical coordinate of Hoskins and Bretherton (1972), which simplified the equations by removing the exponential factors, and allowed the linearization $s = s_0 + s_d$. Unfortunately this coordinate tends asymptotically to a constant value quite low in the stratosphere, so is not suitable for this jovian model. Instead we use the more usual log-pressure definition of z, as described in Section 2.1.

Secondly, the Norton model is based on a uniform finite-difference grid. This is inappropriate for the jovian atmosphere, whose temperature lapse rate is close to adiabatic in the deep troposphere, but extremely stable in the stratosphere. In the troposphere, the potential temperature varies very slowly with height. A "stretched" vertical grid is therefore required to obtain reasonable resolution in the troposphere without requiring an impractically large number of levels. In the horizontal direction, however, a uniform grid is used.

Finite differences on a non-uniform grid

Sundqvist and Veronis (1970) derived expressions for the first and second derivatives of a function differenced on a non-uniform grid. Consider a function $f(\theta)$, discretized on a grid where the spacing between vertical levels j-1 and j is h_j. The first spatial derivative f_j^I is given by

$$f_j^I = \frac{h_j}{(h_j + h_{j+1})h_{j+1}} f_{j+1} - \frac{h_{j+1}}{(h_j + h_{j+1})h_j} f_{j-1} - \frac{h_j^2 - h_{j+1}^2}{h_j h_{j+1}(h_j + h_{j+1})} f_j + O(h_j h_{j+1}) \quad (2.19)$$

and the second derivative by

$$f_j^{II} = \frac{2}{h_j + h_{j+1}} \left(\frac{f_{j+1} - f_j}{h_{j+1}} - \frac{f_j - f_{j-1}}{h_j} \right) + O(h_{j+1} - h_j) \quad (2.20)$$

Equation (2.19) gives f^I accurate to second order in the space increments; equation (2.20)

gives f^{II} accurate to first order in the difference between the space increments. Clearly when $h_j = h_{j+1}$ the expressions reduce to the usual forms for a uniform grid. For a stretched grid, the accuracy of (2.20) can be increased to second order by arranging the grid so that the following condition is satisfied:

$$h_{j+1} - h_j = O(h_j^2) \tag{2.21}$$

In the radial (horizontal) direction, a uniform grid is used and the usual finite difference expressions apply.

The functional form chosen for h_j as a function of j is a cubic spline, which can be adjusted to give optimal resolution. This form has several of the desirable properties listed by Kálnay de Rivas (1972). The curve is smooth, so its gradient is finite everywhere, and is zero at the base of the model. The third derivative is checked to be small, and higher derivatives are identically zero. The form is chosen also to satisfy the second-order accuracy condition, Eq. (2.21).

Discretization of the equations and boundary conditions

The coordinates are discretized by $r = (i-1)\Delta r$ and $\theta = \Sigma\Delta\theta^j$, for $i = 1 \ldots N_r$ and $j = 1 \ldots N_z$. Here $\Delta\theta^j$ is the θ-spacing between levels $j-1$ and j. Equation (2.14) is discretized as follows:

$$\frac{v^{i+\frac{1}{2}j} - v^{i-\frac{1}{2}j}}{\Delta r} + \frac{v^{i+\frac{1}{2}j} + v^{i-\frac{1}{2}j}}{2i\Delta r} + Q^{ij}\frac{H_0}{R}\exp\left(\frac{\kappa z^{ij}}{H_0}\right)M_{ij}^{II} = Q_d^{ij}s_0^j\exp\left(\frac{\kappa}{H_0}(z-z_0)\right) \tag{2.22}$$

where, from (2.20), the second derivative of M is represented by

$$M_{ij}^{II} = \frac{2}{\Delta\theta^j + \Delta\theta^{j+1}}\left(\frac{M_d^{ij+1} - M_d^{ij}}{\Delta\theta^{j+1}} - \frac{M_d^{ij} - M_d^{ij-1}}{\Delta\theta^j}\right) \tag{2.23}$$

This expression is evaluated at all grid points in r and θ. The balance condition (2.15) is discretized thus:

$$\frac{\left(v^{i+\frac{1}{2}j}\right)^2}{\left(i+\frac{1}{2}\right)\Delta r} + fv^{i+\frac{1}{2}j} - \frac{M_d^{i+1j} - M_d^{ij}}{\Delta r} = 0 \tag{2.24}$$

which is evaluated at the grid points in θ and half way between the grid points in r. The boundary condition (2.17) is implemented at the top boundary as $M^{i,nz+1} = M^{i,nz}$, and similarly at the lower boundary.

The iterative method

The method is based on the Newton-Raphson scheme used by Norton, with alterations required by the use of a log-pressure vertical coordinate. The exponential term in Equation

(2.22) requires knowledge of z^{ij} in order to calculate s^{ij}. However, z^{ij} must be calculated by summing $s^{ij\prime}$ for $j\prime < j$. This is because the equations are solved for M_d and v only. Equation (2.5) is used to calculate s^{ij} from M_d^{ij}, and we then calculate z by integrating Equation (2.4). Therefore z^{ij} must be recalculated on each iteration using the latest values of s^{ij}. The method is summarized below.

1. Initialize s from reference s_0, and set $M_d = v = 0$.
2. Calculate z^{ij} from s^{ij}
3. Use these z^{ij} to set up the matrix equation
4. Invert the matrix, solving for M_d and v
5. Calculate a better approximation to s^{ij} from Equ. (2.5)
6. Iterate steps 2 - 5 to convergence
7. Calculate z^{ij} from the final s^{ij}
8. Calculate derived fields and output the results

2.1.5 The model domain

The model domain for GRS-like vortices is a vertical cylinder of radius 4×10^4 km, extending from a pressure of 100 bar up to 10 mb. The choice of the top level is influenced by the observations, which suggest (see Section 2.2.1) that the winds above the Great Red Spot decrease with height, decaying to zero at around 10 mb. The value of 100 bar for the lower boundary is an arbitrary one, but ensures that the model domain extends well below the possible water clouds (near 5 bar) and the deepest observations made of Jupiter (22 bar). Later we present models of vortices resembling the White Ovals, in a domain of the same depth but with a radius of 2×10^4 km.

The observed peak velocity of the GRS occurs at a radius of approximately 1.2×10^4 km along the semi-major axis (Mitchell et al., 1981), so this domain allows the flow to peak at about one quarter of the maximum radius, and to vanish at $r \sim r_{max}$. This ensures the domain is large enough for the boundary condition (2.18) to be valid as a substitute for the continuity equation. The condition that the flow must vanish at large r (together with the desired axial symmetry) prevents any representation of the South Tropical Zone surrounding the Spot, or the SEB adjacent to it. Comparisons with observations are made only between the inner region of the model domain and the interior of the GRS.

The model has 38 points in the radial direction, and 160 vertical levels. The latter are specified on potential temperatures from 166 K to about 600 K (corresponding to 100 bar–10 mb). Figure 2.1 illustrates the level spacing in two forms. The left-hand panel (a)

shows the actual form of h_j in terms of model levels. The right-hand panel (b) shows the approximate geometric height spacing, on a log-pressure vertical scale, using the reference static stability profile (see Section 2.2.4). This plot shows that the geometric level spacing is rather non-uniform, but is never larger than roughly one-fifth of a scale height, and above 1 bar is much smaller than this. Our claim of good vertical resolution seems justified.

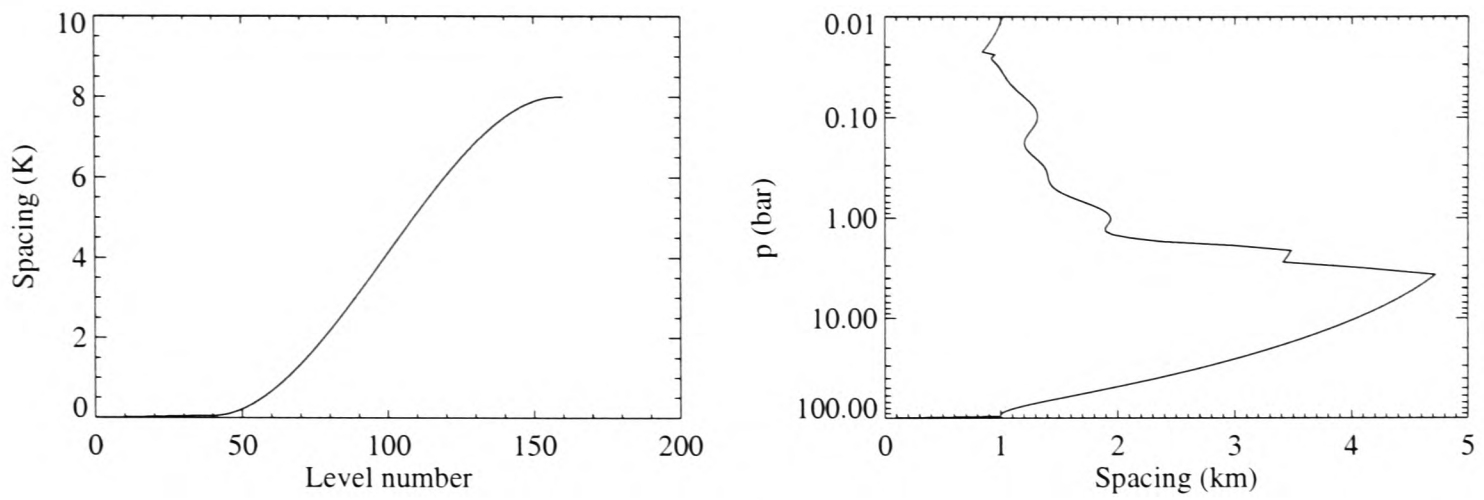


Figure 2.1: The vertical spacing of model levels, as a function of (a) level number and (b) height

2.1.6 Validation of the method

The consistency of the model solutions was checked by evaluating the predicted vertical wind shear. The thermal wind equation in isentropic coordinates is easily derived by differentiating the balance condition (2.15) with respect to θ at constant r. This gives

$$\left(\frac{2v}{r}+f\right)\left(\frac{\partial v}{\partial \theta}\right)_r = \frac{\partial}{\partial \theta}\left(\frac{\partial M_d}{\partial r}\right)_\theta = \frac{\partial}{\partial r}\left(\frac{\partial M_d}{\partial \theta}\right)_r \tag{2.25}$$

M_0 is independent of r, so can be neglected here; we can substitute from the hydrostatic equation (2.2) to obtain

$$\left(\frac{2v}{r}+f\right)\left(\frac{\partial v}{\partial \theta}\right)_r = c_p \frac{\partial}{\partial r}\left[\exp\left(\frac{-\kappa z}{H_0}\right)\right] \tag{2.26}$$

We can eliminate the exponential term by using the relation between T and θ, to arrive at the thermal wind equation:

$$\left(\frac{2v}{r}+f\right)\left(\frac{\partial v}{\partial \theta}\right)_r = \frac{c_p}{\theta}\left(\frac{\partial T}{\partial r}\right)_\theta \tag{2.27}$$

The errors in thermal wind balance in all solutions obtained are typically smaller than 1% in the troposphere, increasing to about 5% near 20 mb, and then much larger in the top 10 or so levels—above about 10 mb. It seems therefore that the results for levels below $\simeq$ 20 mb can be accepted with confidence. Above this region the truncation errors are large; but in this altitude range, the solutions simply decay to the reference state and are of little interest. In principle these errors at the top of the domain could be decreased by performing extra iterations with quadruple precision (128-bit) arithmetic.

2.2 Observations of Jupiter's giant eddies

Both Voyagers yielded high-resolution images of the giant eddies, which have been re-analysed several times. Galileo was able to improve on these with better spectral and temporal coverage, and has largely supported the Voyager results. An overview of the spacecraft design and of some observational results was given in Chapter 1. We deferred the details of observations of the giant eddies until now, when we will summarize the results relevant to our PV inversion model of the GRS and White Ovals.

2.2.1 Voyager observations of the Great Red Spot

Flow fields within the GRS have been derived from Voyager 1 narrow angle camera images (Mitchell et al., 1981; Dowling and Ingersoll, 1988). Temperatures and thermal winds above the observed clouds in the Spot have been retrieved from IRIS data (Flasar et al., 1981). More recently, Sada et al. (1996) have presented a more detailed analysis of the observations (Fig. 2.2), including relative vorticity and horizontal divergence calculated using the methods of Dowling and Ingersoll (1988). The main observations which the present model attempts to replicate are summarized below.

322 SADA, BEEBE, AND CONRATH

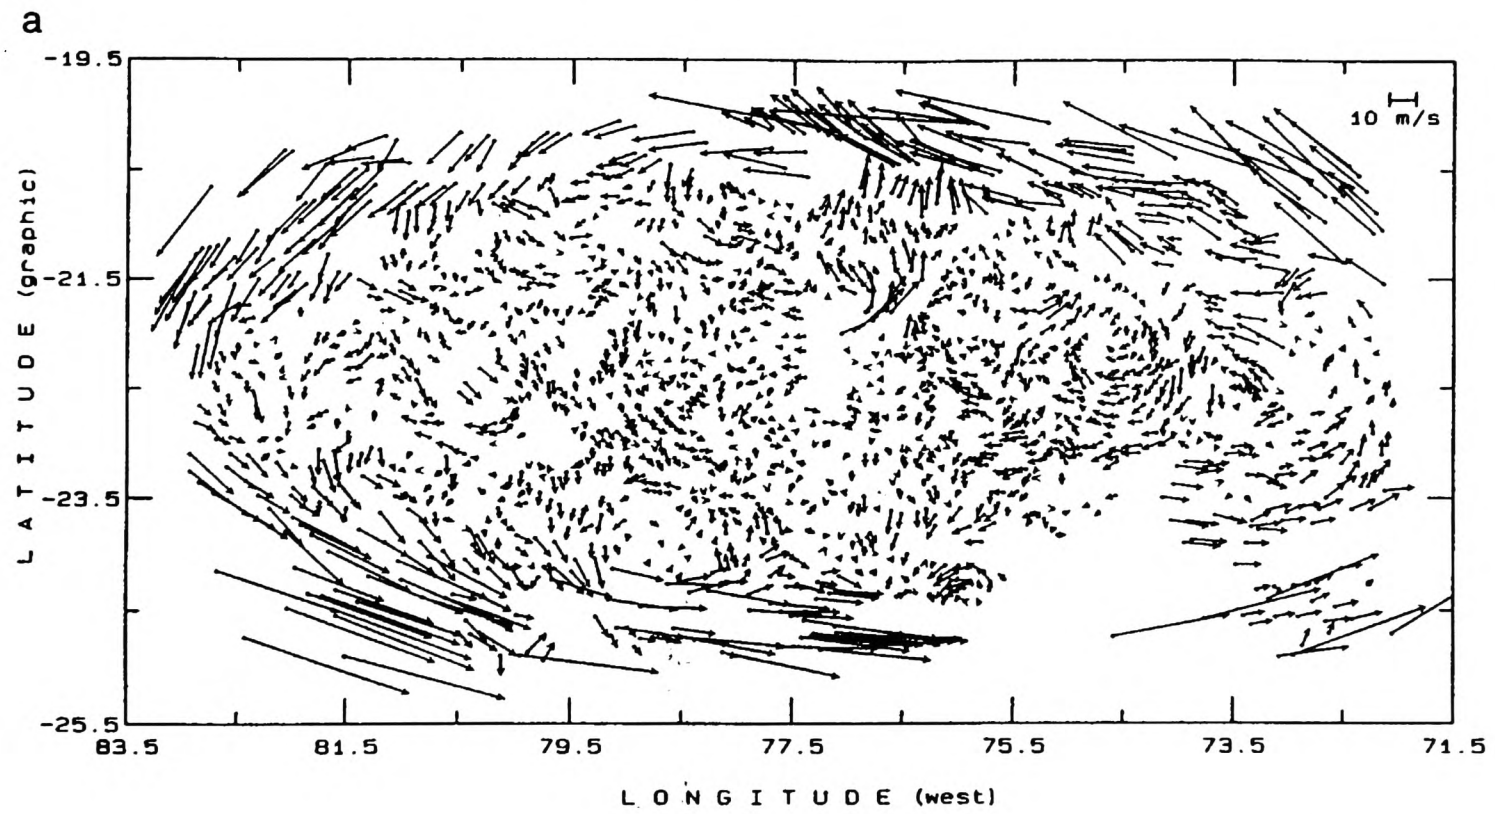
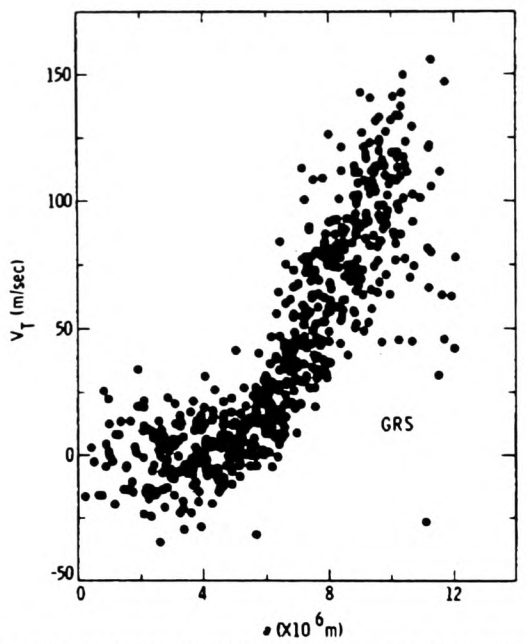


Figure 2.2: GRS wind vectors as seen by Voyager 1, from Sada et al. (1996)

Mitchell et al. (1981) neglected the angular dependence of the measured azimuthal velocities. For the purposes of comparison with this analysis, a circular model is a reasonable idealisation. However, the relation between velocity and vorticity depends upon the eccentricity; this must be taken into account when comparing model results with observa-

tions. Mitchell et al. also neglected the small radial component of velocity.



Observed tangential velocities (v_T) as a function of a within the GRS.

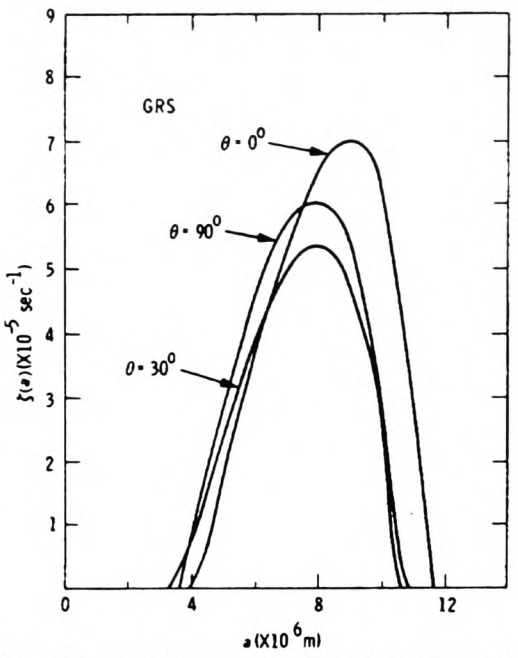


Fig. 6. Relative vorticity (ζ) as a function of a and Θ within the GRS.

Figure 2.3: Observed velocity and vorticity fields in the GRS, from Mitchell et al. (1981)

Figure 2.3 shows the one-dimensional flow fields measured by Mitchell et al.. The core of the GRS is quiescent, out to a radius of order 5000 km; the measured cloud-top velocity then increases in the peripheral jet, to $v \sim 150$ m/s at $r \sim 12000$ km. The relative vorticity peaks at about 8000 km radius (see also Fig. 2.7).

Infra-red measurements Flasar et al. (1981) suggest thermal wind shear consistent with a decay of the jet, to zero velocity at around 10 mb. We use this fact when calculating the PV structure of the vortex (Section 2.3.1), and indeed when choosing the upper boundary condition for the GCM (Chapter 4). Temperature sections reported by Flasar et al. are shown in Fig. 2.4.

The temperature fields show the following features, with which we may compare the model output. There is a cold core in the stratosphere above the Spot, up to $p \sim 10$ mb; this is surrounded by a warmer peripheral region. The IRIS east-west cross-section of temperatures over the GRS (Flasar et al., 1981, Fig. 2(a)) clearly shows that the cold core does not have a simple peak, but has an annular structure which is slightly warmer in the centre.

Sada et al. (1996) reanalysed the Voyager 2 data to give more detail about the GRS core. They discovered a large but weak cyclonic region within the slowly-moving core. This was observed by cloud tracking and corroborated by a small local maximum in the 143 mb temperatures, indicating cyclonic motion within a larger anticyclonic region. The Voyager 1 images also showed some evidence of another cyclonic feature, in the process of either formation or destruction. This was on the opposite side to the cyclone mentioned above, raising speculation on the possibility of there being several sporadic, alternating cyclones

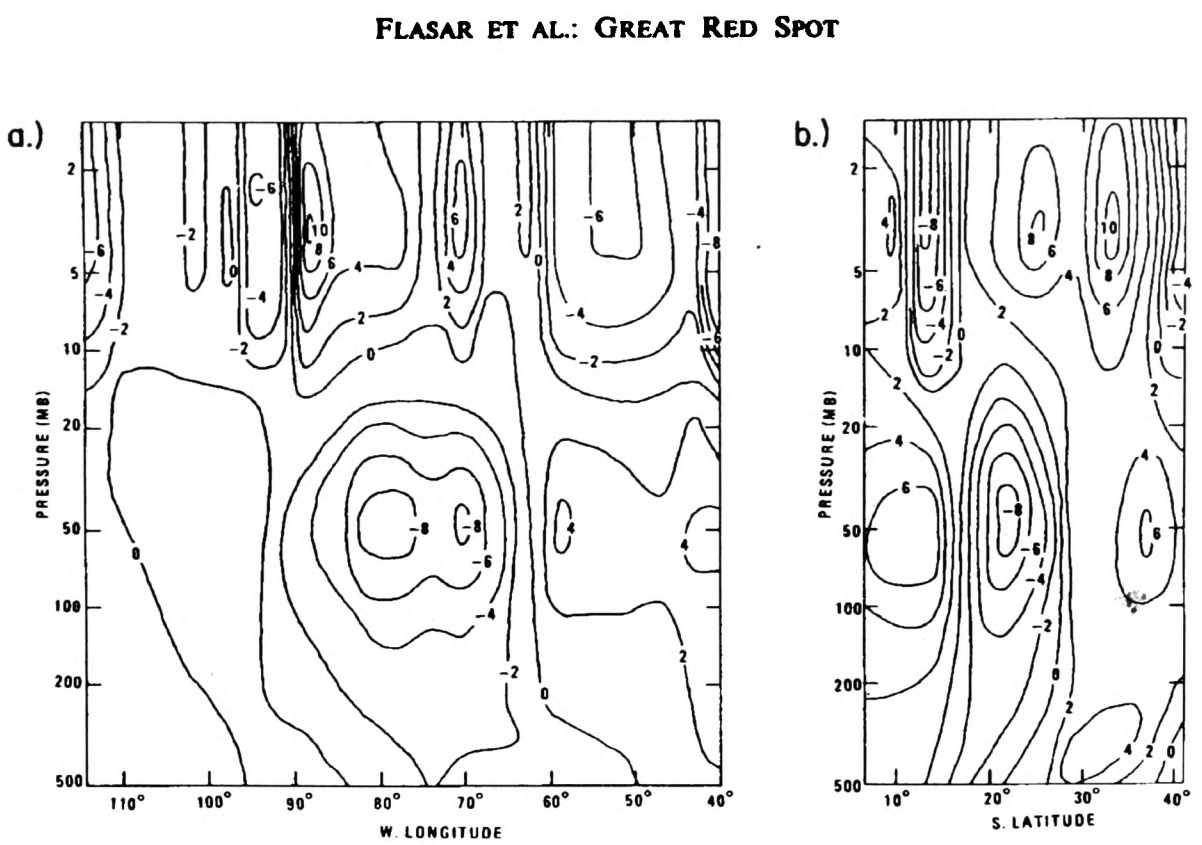


Fig. 2. (*a*) East-west vertical cross section of temperature fields over the GRS. (*b*) North-south vertical cross section of temperature fields over the GRS. The temperatures are deviations (K) from an average temperature profile.

igure 2.4: Temperature anomaly and thermal wind fields in the GRS, from Flasar et al. (1981)

rming and reforming in different areas of the GRS interior.

ada et al. also revised some of the results of Mitchell et al. (1981) with regard to velocies in the collar. The average zonal wind on the southern side was measured by both oyagers to be around 65 m/s. However, on the northern side the Voyager 1 speed was $10 \pm 23)$ m/s, compared with the Voyager 2 speed of (80 ± 25) m/s. The northern section the collar was about twice as wide as the southern section, so continuity arguments ould imply that the speed there should be half that in the south. The difficulty may be solved by noting that the depth of the flow need not be uniform. In fact, the Galileo sults (see below) conflict with these measurements.

.2.2 Galileo observations of the Great Red Spot

alileo SSI images of the GRS were taken on 26 June 1996, and the most recent published nalysis is by Vasavada et al. (1998). The zonal velocity along the central meridian was sentially unchanged since the Voyager encounters. However, Vasavada et al. suggested at the imbalance between the zonal speeds in the north and south noted by Sada et al. as due to a reliance on visual methods to determine the boundary between the GRS llar and the surrounding flow. This led to an underestimate of the true meridional width the GRS. The Galileo measurements extended further south and revealed symmetric locity peaks of nearly 150 m/s. Measurements of the ¡meridional wind across the 20.3°

latitude circle gave −94 and +98 m/s at the two edges.

The Galileo images therefore show a Red Spot which is almost symmetrical in its peak wind speeds; however, a "visual estimation of the boundary of the GRS" (Vasavada et al., 1998) showed the northern semiminor axis to be 25% longer than the southern. Of course, it is still difficult to estimate of the depth of the flow, although attempts have been made to derive the depth of the "bottom topography", for example by a vortex tube stretching analysis in a shallow-water model (Dowling and Ingersoll, 1989).

More importantly, Galileo found further evidence of large-scale cyclonic motion within the GRS. This time there was a complete circuit, counterrotating just within the collar, with a speed of 20 m/s (Vasavada et al., 1998). The velocity profiles of both automatic and manual measurements (Sada et al., 1996, Fig. 5) clearly show a reversal of the rotation. It seems that cyclonic motion within the core is a persistent feature although its extent has varied with time, much more so than the motions in the collar.

A model of the vertical structure of clouds on Jupiter generally has been retrieved by Banfield et al. (1998), using SSI data. This suggests that above the GRS, the top of the upper tropospheric haze lies at 160 mb (compared to $\sim$ 200 mb at most other locations). The visible clouds would be best fitted by optically thick but geometrically thin layers at pressures around 700 mb—probably of ammonia. Banfield et al. (1998) suggested that the clouds comprising the visible GRS may be at 700 mb, lower than previous estimates but within the errors of $\pm$200 mb.

Other interesting results were obtained, among them a deep cloud (below the 4 bar level) seen northwest of the GRS. This must be of water, unless the composition of the atmosphere is nothing like we believe. No ammonium hydrosulphide cloud was required to fit the retrieved profiles, but sensitivity tests suggested that such a cloud could not be detected if it lay beneath an ammonia cloud of uncorrelated optical thickness. However, NIMS has (unlike SSI) a complete spectral coverage of this region, and did detect a probable ammonium hydrosulphate cloud, consistent with the Probe nephelometer result (although of course the Probe entered a NEB hot-spot, not the GRS).

2.2.3 Observations of White Oval BC

The White Ovals have several characteristics in common with the Great Red Spot, and it is interesting to compare PV inversion models of each. Mitchell et al. (1981) analysed the flow within Oval BC (Fig. 2.5), as well as that in the GRS.

Oval BC lies further poleward than the GRS and has about half its horizontal scale. The most significant dynamical difference is the Oval's apparent lack of a quiescent centre.

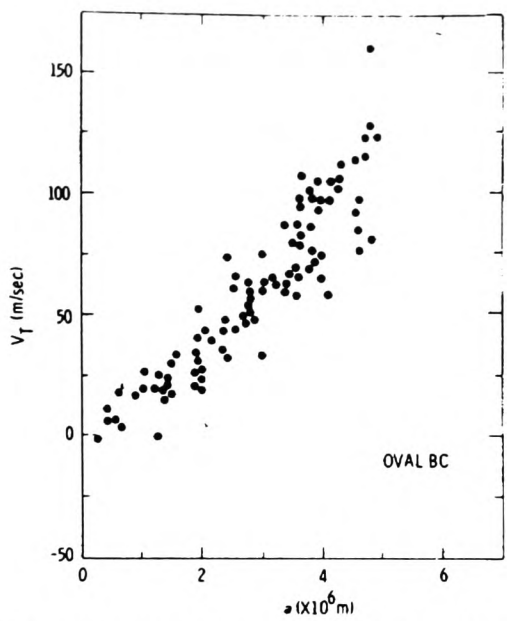


Fig. 5. Observed tangential velocities (V_T) as a function of a within Oval BC.

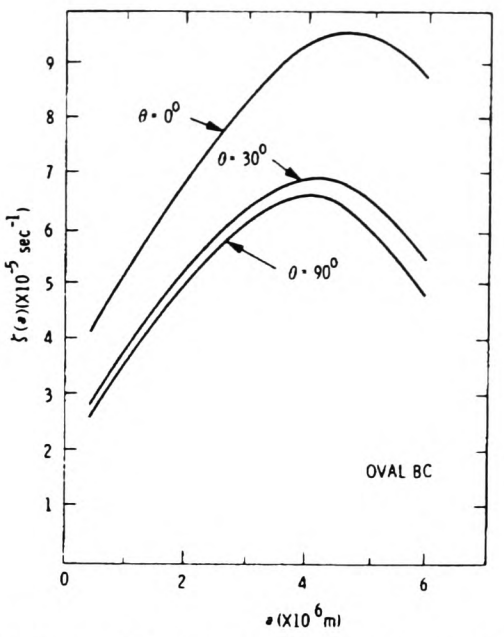


Fig. 7. Relative vorticity (ζ) as a function of a and Θ within Oval BC.

Figure 2.5: Observed velocity and vorticity fields in Oval BC, from Mitchell et al. (1981)

The azimuthal velocity increases radially from the centre in a roughly linear manner. The reason for this important difference between two otherwise similar features is unclear.

Galileo observations (Vasavada et al., 1998) have supported the Voyager analysis, with no great changes observed within the White Ovals, although Oval BC has changed shape to become slightly more circular. The interaction of the three ovals (BC and DE were recently joined by FA) with three smaller anticyclones has formed a string of cyclone-anticyclone pairs at around 31°S (Simon et al., 1998). In this chapter we are only concerned with the structure within an isolated vortex, but will return to vortex interactions in Chapter 5.

2.2.4 Static stability profiles

The profile of s (which is the inverse of the isentropic static stability) is highly idealised, but arranged so that the derived temperature profile is consistent with Voyager radio occultation measurements (Lindal et al., 1981). The reference-state profiles of $s_0(p)$ and $T(p)$ are shown in Figure 2.6.

Below $p \simeq 3$ bar, $s_0(p)$ is given a large constant value, 10^5 m/K. This gives a very weakly stratified atmosphere with a lapse rate of approximately -1.90 K/km, consistent with data taken by the Galileo probe at 5°N (Seiff et al., 1996; Seiff, 1998). At the top of the model domain, $s = 100$ m/K. Between these two extrema, the function is approximately log-linear in pressure. The effects of cloud decks on s_0 are not illustrated here; experiments have shown that the inclusion of a localised peak in static stability, representing the water clouds, has a negligible effect on the model results. The effect of the proposed stable layer around 18 bars (Guillot et al., 1994a) has not been investigated.

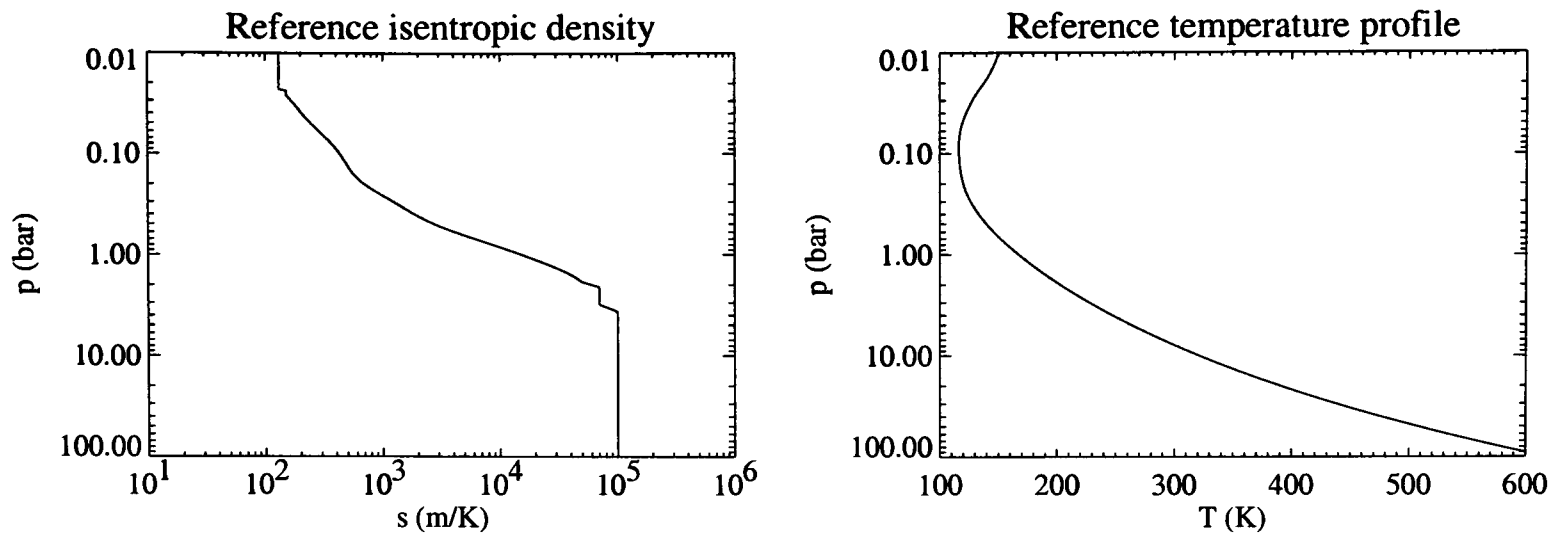


Figure 2.6: The model stratification (isentropic density) and temperature profiles

2.3 A diagnostic model of the Great Red Spot

2.3.1 Potential vorticity distributions

An approximate potential vorticity map of the region above the GRS may be obtained from the published data outlined in Section 2.2.1 above. The least-squares coefficients calculated by use of Mitchell et al. (Table 2 from 1981) are used to find the radial profile of the azimuthal wind at the cloud-tops. The expression given by Mitchell et al. for the relative vorticity at any point on the semi-major axis ($\phi = 0$) is

$$\zeta(a,0) = \frac{a}{b^2} V_T + \frac{\partial V_T}{\partial a} \tag{2.28}$$

where a and b are the lengths of the semi-major and semi-minor axes respectively, of the ellipse containing the point, and V_T is the azimuthal velocity, in our paper denoted by v. Figure 2.7 shows cloud-top velocity and vorticity profiles at a point on the semi-major axis, calculated from these quartic coefficients.

As a convenient simplification of the Voyager IRIS thermal wind measurements, the winds are assumed to decay linearly with log-pressure, reaching zero at 22 mb. We make the arbitrary assumption that the cloud-tops lie on the 173 K potential temperature surface. This corresponds to a reference-state pressure of about 400 mb, which is within the range of observational estimates of the cloud height (see Flasar et al., 1981).

Given the wind field, the relative vorticity field above the cloud tops may then be calculated using Eq. 2.28. Together with the unperturbed static stability profile we have assumed, this determines a first approximation to the PV structure above the GRS. Of course, this state has an unrealistic partitioning of the PV: the entire anomaly is manifested in the relative vorticity, and none in the static stability. However, the iterative process modifies this into a more realistic and balanced state. We shall show that the

model flows resemble the initial vorticity data at the cloud-tops.

A comparison of Equations (2.11) and (2.28) shows that if the calculated PV anomaly is inverted in an axisymmetric domain rather than an elliptical one, our geometrical approximations mean that the diagnosed vorticities are inconsistent with the observations. The winds and temperatures can usefully be compared, however.

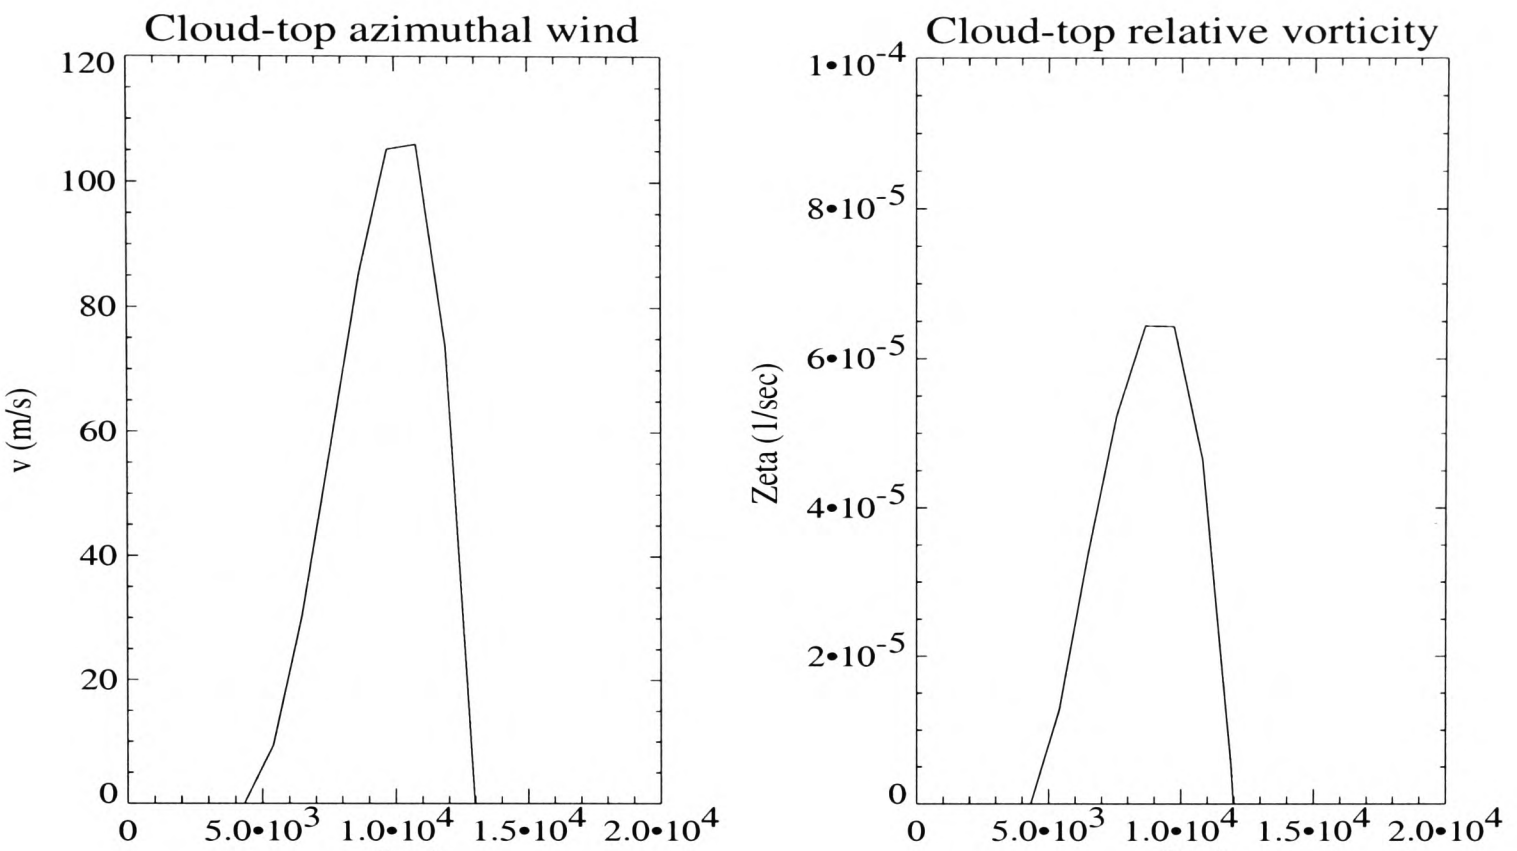


Figure 2.7: Observed cloud-top winds and relative vorticities, calculated from fitted coefficients published by Mitchell et al.

2.3.2 Model results

Little is known about the winds below the cloud tops, other than the single profile from the Galileo Probe. An aim of this work is to investigate the different possibilities. In order to specify a PV anomaly, assumptions must be made about the deep winds (below the clouds); these assumptions are then tested by the results of the PV inversion. The results of several experiments with different anomalies are given below.

Case A: inversion of an anomaly inducing a deep vortex

The first case we consider is one in which the depth of the vortex is much larger than the model domain. To obtain a suitable PV anomaly, we assume that below the clouds, the winds are constant at the cloud-top values—i.e. the flow is barotropic at depth. From this we estimate the PV distribution as outlined above. Figure 2.8 shows the resulting PV anomaly plotted on both isentropic and log-pressure coordinates.

The cross-section of the PV anomaly has a horizontally dipolar structure. The core has very small or zero anomalous PV, surrounded by a torus of anticyclonic PV surrounded by a collar of cyclonic PV. This is a consequence of the fitted wind velocities, which are small in the core, increase rapidly in the peripheral jet, and then level off towards the edge of the observed GRS. It is perhaps counter-intuitive that the anomaly does not extend into the centre of the vortex, unlike those associated with many flows in the Earth's atmosphere. It will be seen later that, for a White Oval, the anomaly does reach the centre (although with smaller values than in the periphery). The "hole" in the centre of the GRS anomaly is associated with the central quiescent region, which is not present in Oval BC.

The anomaly has only one vertical layer: the sign of the PV is constant with height. This monopolar vertical structure contrasts with the dipole anomalies we shall consider in the next section.

Figure 2.9 shows the results of a model solution obtained from this input data. The inversion predicts a cold anomaly above the Spot, which is toroidal, and weaker in the centre. This agrees qualitatively with the Voyager IRIS measurements (Flasar et al., 1981, Fig. 2a), although the IRIS temperature anomaly is rather stronger than that predicted by the model. The predicted and observed anomalies both peak at a level of 50 mb. Fig. 2 of Flasar et al. shows a double peak in the east-west cross-section of temperature, which is predicted by this model. The north-south cross-section shows only a single peak, but this is limited by IRIS resolution.

The model winds are in agreement with the measurements of Mitchell et al.. Below the cloud-tops the vortex is entirely barotropic, the winds not decaying at all, down to at least 100 bars.

It seems that, within the limitations of this model, it is possible to choose a PV anomaly which induces a vortex which is consistent with the observations above the clouds, and extends to an arbitrary depth beneath the cloud-tops. We now turn to the opposite limiting case: a shallow vortex, decaying rapidly with increasing depth below the clouds.

Case B: inversion of a vertical PV dipole, giving a shallow vortex

The results of the previous section show flows induced by a single layer PV anomaly (i.e. one whose sign does not vary with height). This anomaly was located close to the tropopause. To induce flows which are trapped in a narrower layer, we impose the PV anomaly shown in Fig. 2.11. This is a two layer, dipole distribution with anticylonic PV above and cyclonic PV below. The anticyclonic anomaly is similar to that in case A, but has a greater magnitude. This is necessary because its effect is opposed by the cyclonic lower anomaly. This lower anomaly is very much weaker than the upper part, but has a

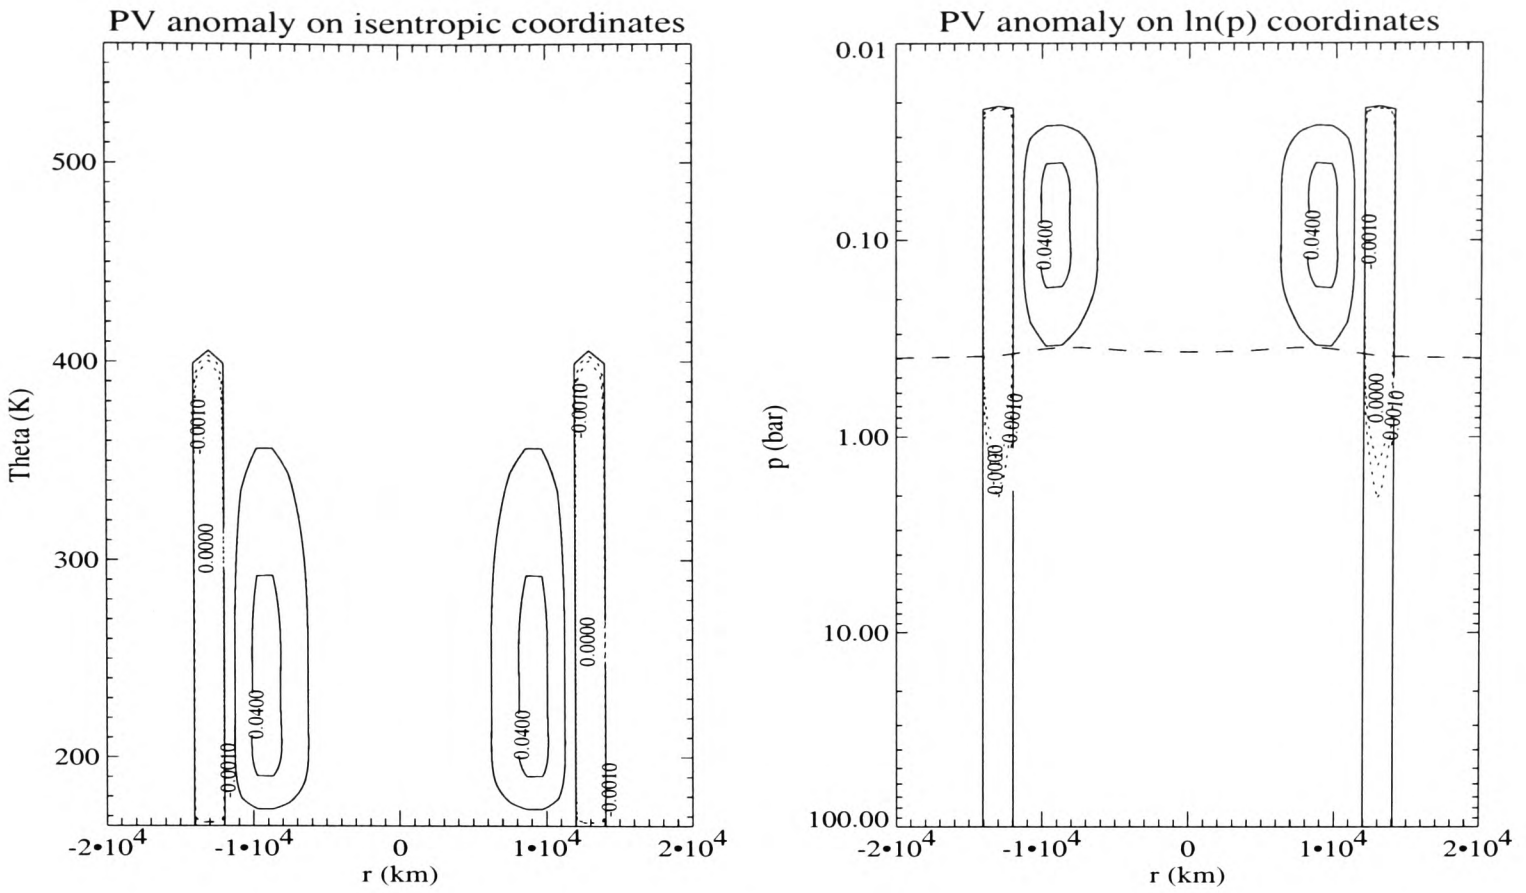


Figure 2.8: PV anomaly as a function of radius and height, for a deep GRS-like vortex. PV is shown as a function of both isentropic and log-pressure vertical coordinates. The PV units are those defined by Hoskins et al. (1985).

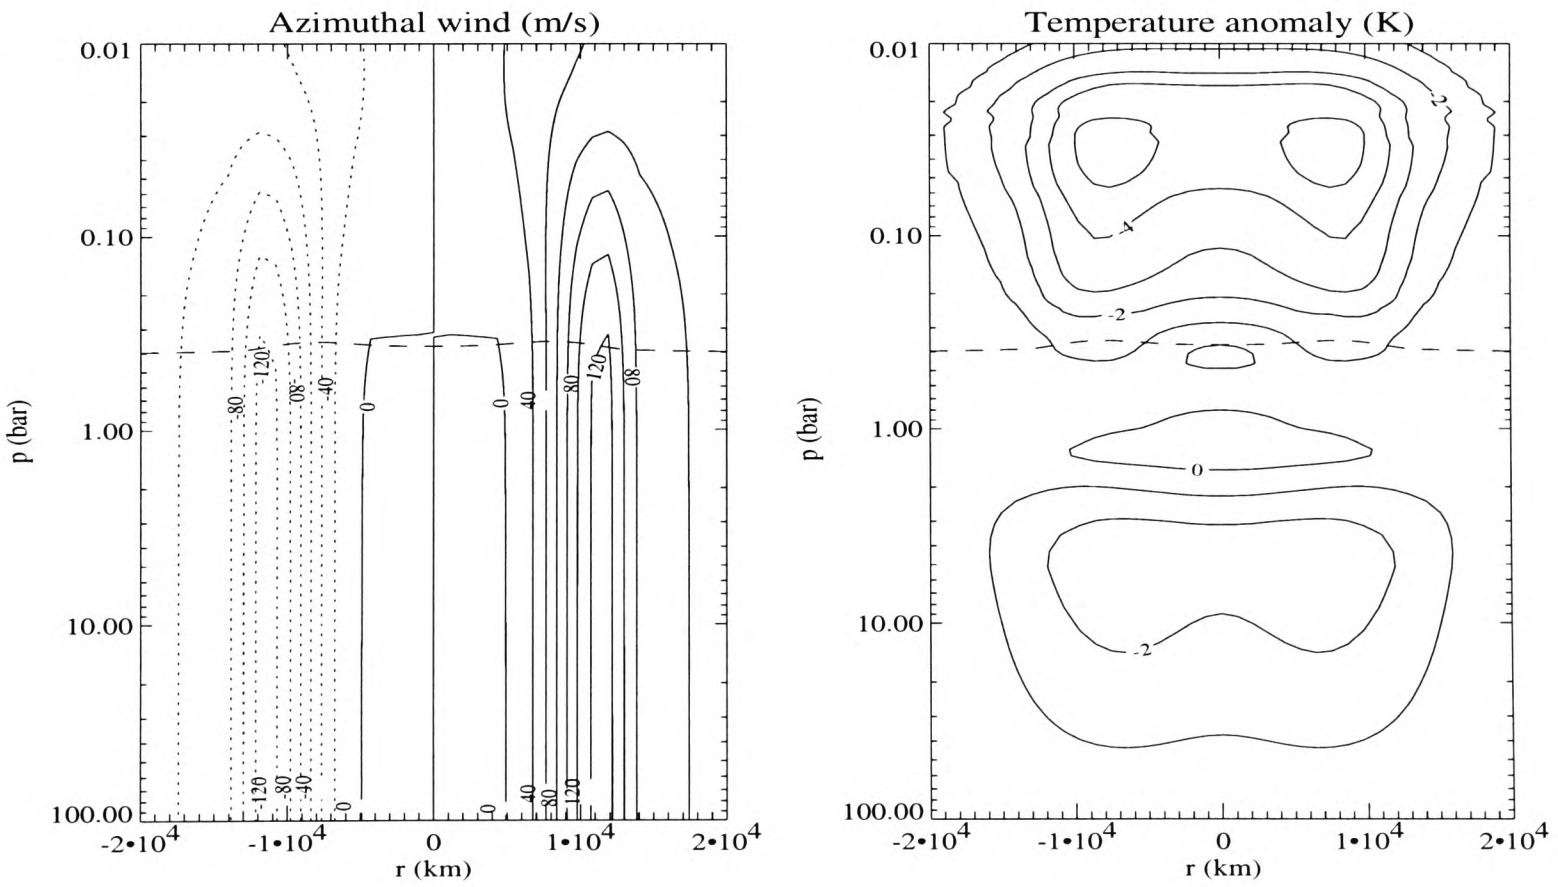


Figure 2.9: Results of a model of a deep vortex. Each panel shows a cross-section through the axisymmetric vortex. (The model domain actually extends to twice the radius shown.) The dashed line marks the approximate height of the observed cloud-tops. This is assumed to be the 178 K surface, which has a reference pressure of about 300 mb.

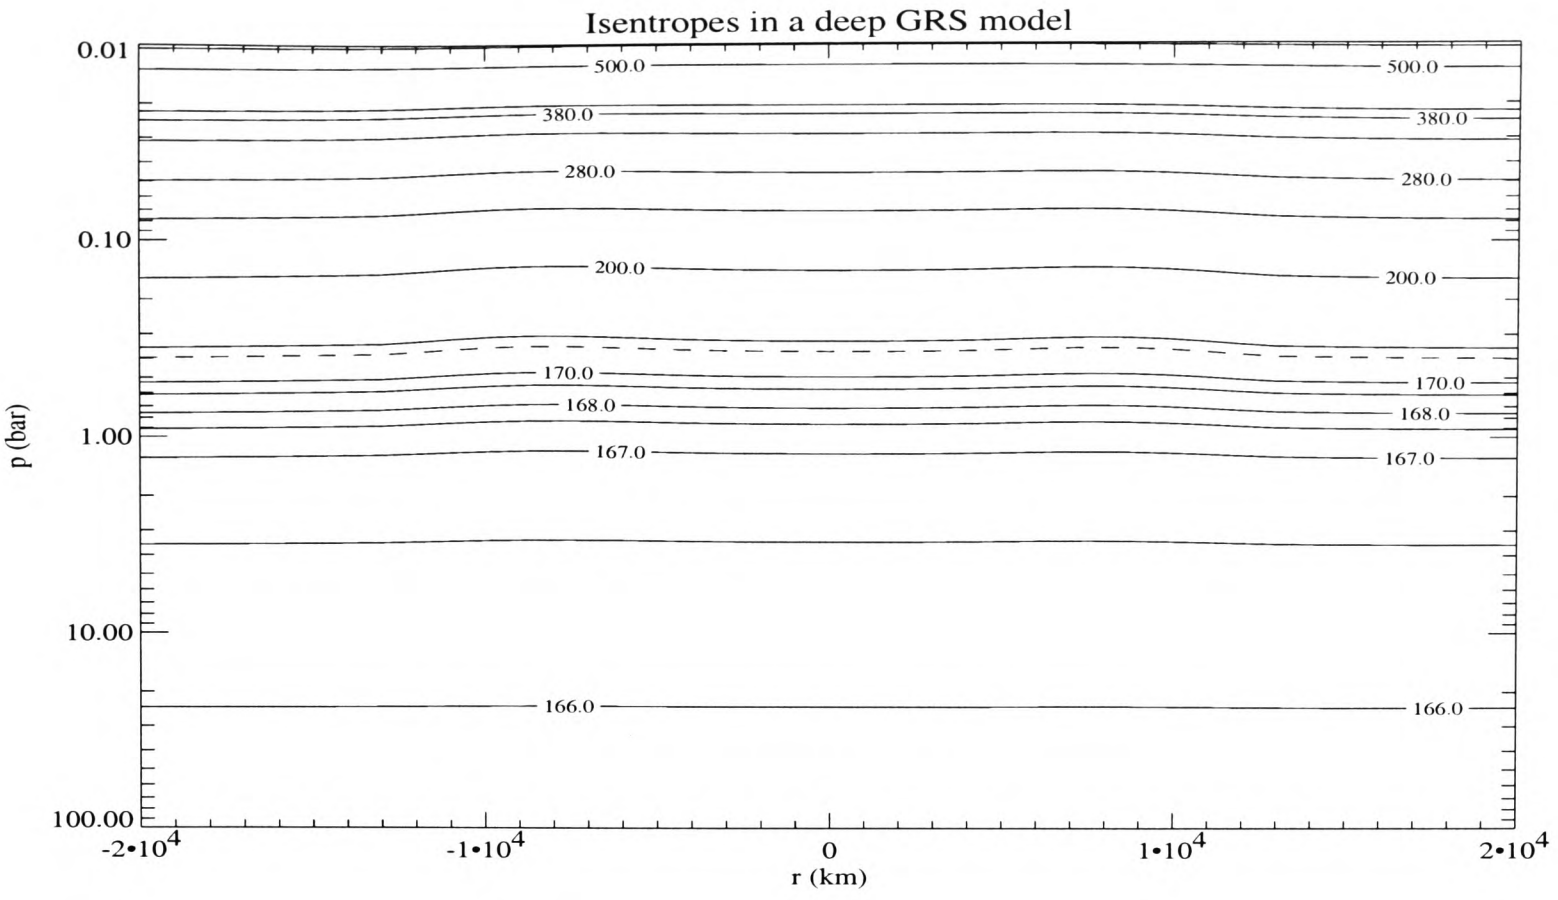


Figure 2.10: Isentropic surfaces, for the solution shown in Fig.2.12. Note the "bowing" of the isentropes. This is a much smaller effect than in Fig. 2.13 below.

significant effect on the flow, because it is placed at a level where s is much larger. The magnitudes of the two anomalies have been adjusted to give results which look realistic at the levels observed by Voyager and Galileo. The PV distribution is also smoothed in the radial direction before inversion. The results are shown in Fig. 2.12.

The cloud top winds and vorticities are consistent with Voyager measurements. The temperature anomaly field, however, shows warmer values than are observed, everywhere above the clouds. The vertical shear below the vortex maximum is so large that a small warm anomaly is predicted just above the nominal cloud-top isentropic surface. Presumably the winds decay to the zonally-symmetric flow at some depth below the observed clouds. If the flow is approximately in thermal wind balance, there must be a warm anomaly at some level below the clouds. In case B, this warmer region extends to just above the clouds, suggesting that this model vortex is shallower than the real GRS might be.

Case B2: A vertical dipole of greater magnitude

Although Case B has a GRS-like structure, is it rather weaker than the real Red Spot. The peak wind is only 100 m/s and the temperature anoamly -4 K. In case B2, the anomaly is scaled up in magnitude to produce a stronger flow. The results (Fig.2.14) are much closer to the observed flow. The winds and temperatures have realistic peak values, and the unrealistic warm region above the clouds is less evident.

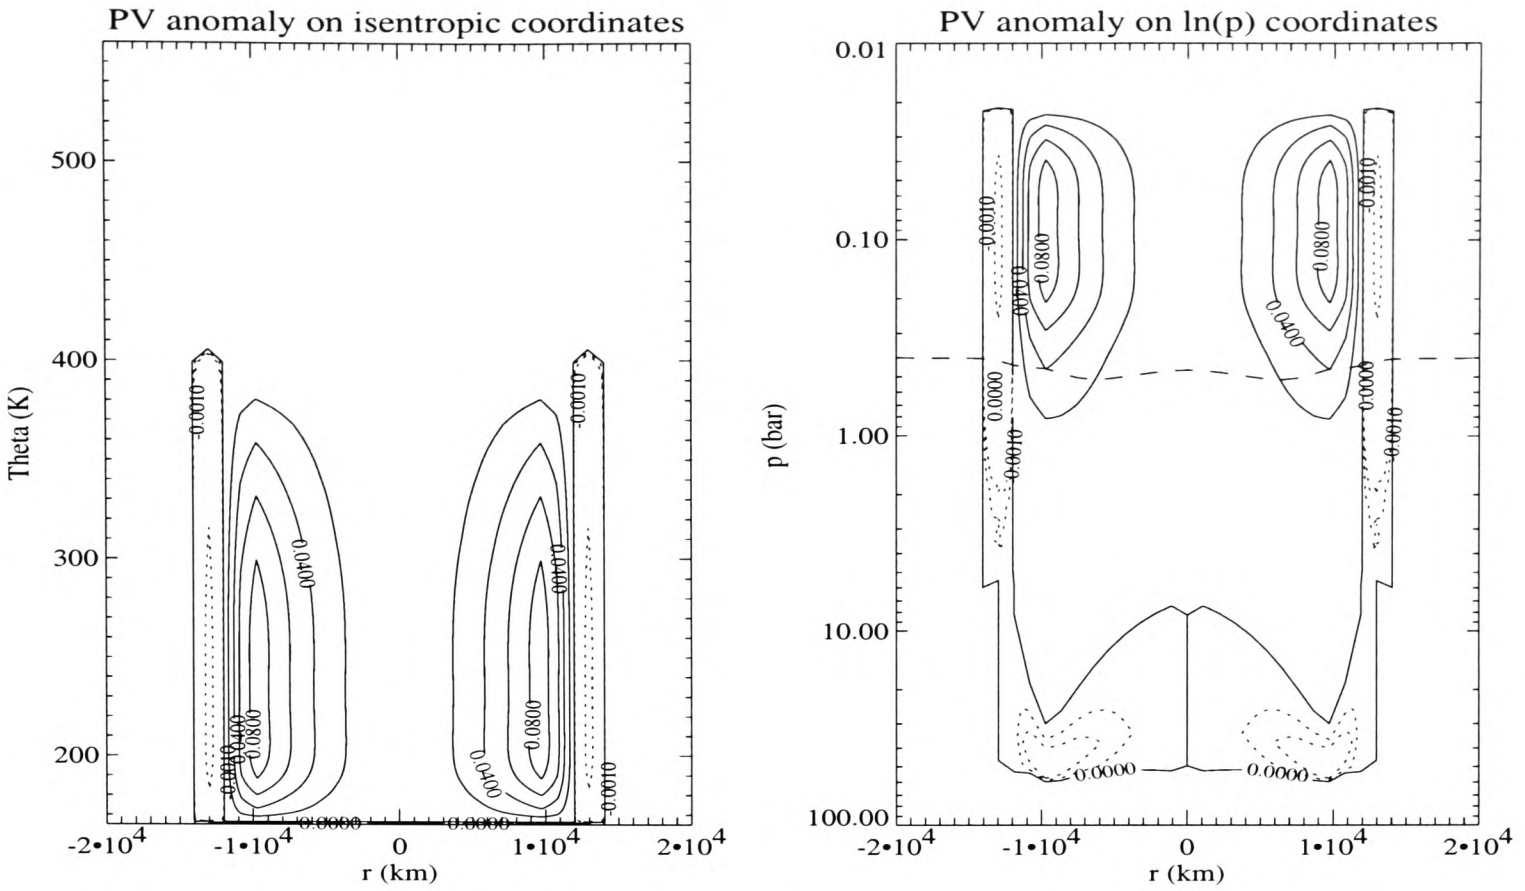


Figure 2.11: A vertical PV dipole structure. Note that the cyclonic lower anomaly is very much smaller in magnitude than the anticyclonic one above. Its effect is significant because of the much larger value of s_0 at this height. The contour interval is nonuniform: the negative contours showing the lower anomaly have values -0.0005 and -0.001 PV units.

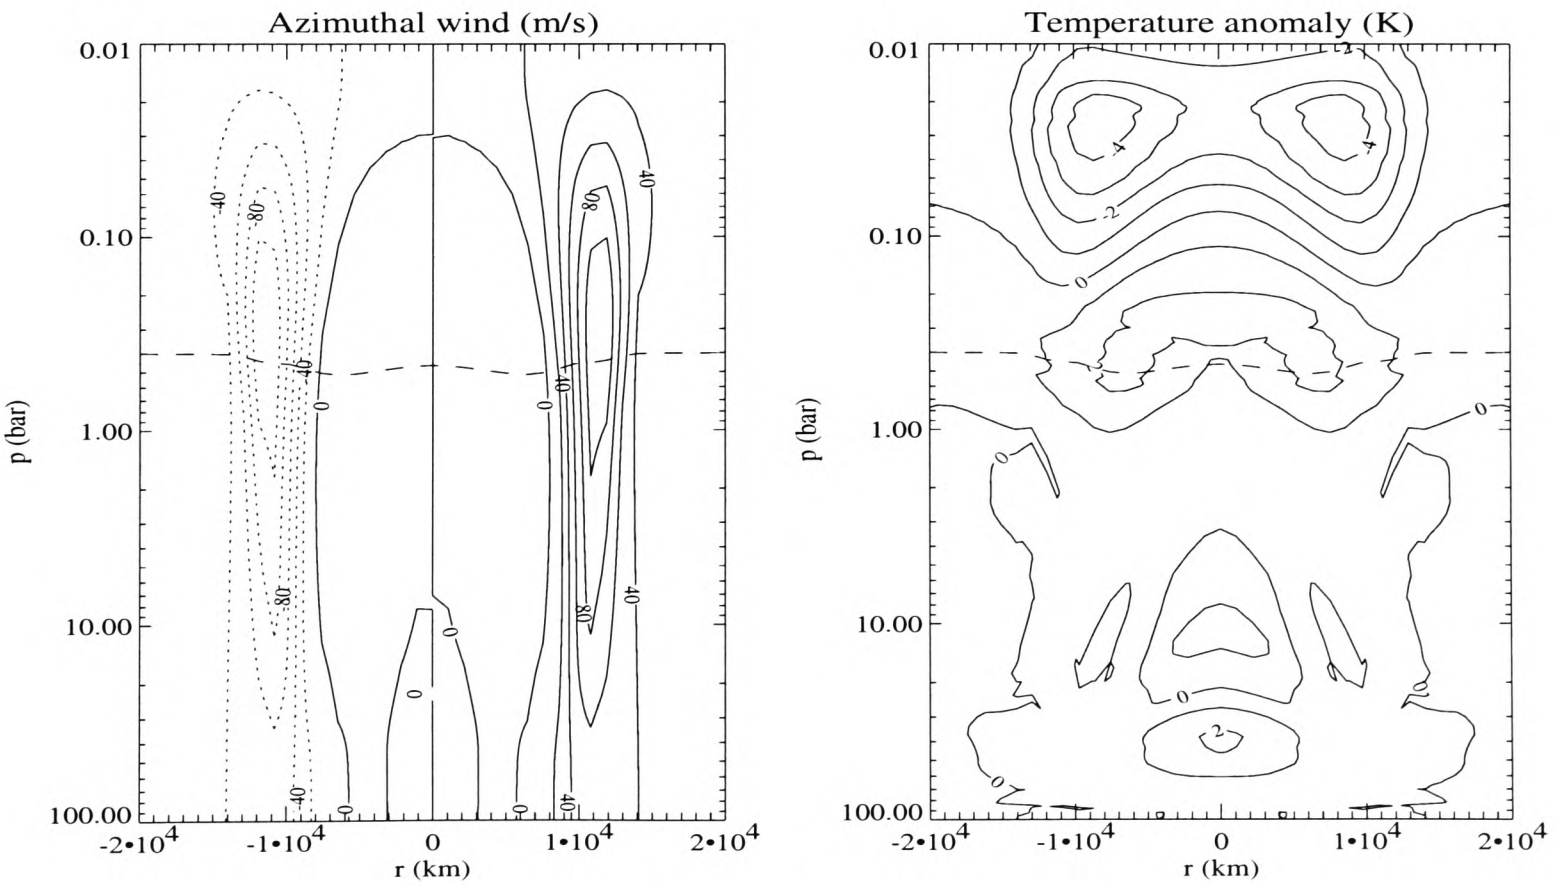


Figure 2.12: Results for Case B1: a vertical PV dipole. The warm temperature anomaly just above the clouds indicates large vertical wind shear at this height, contrary to observations. Hence the real GRS is probably deeper than this vortex, in that its circulation may not decay so rapidly with depth.

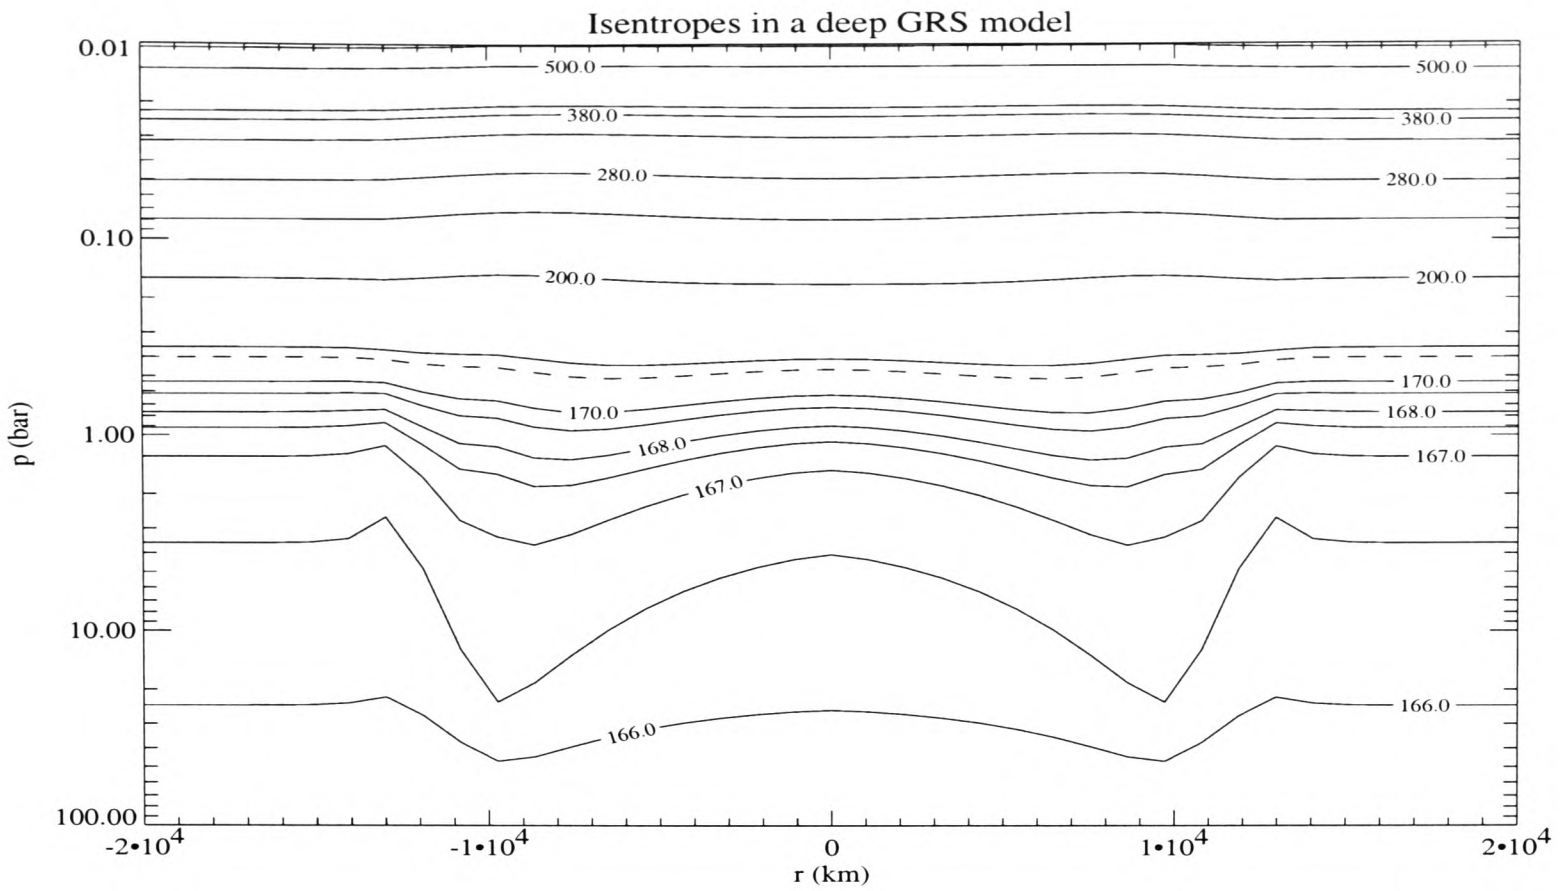


Figure 2.13: Isentropic surfaces, for the solution, B1, shown in Fig.2.12. Note the "bowing" above and (much larger) below the anomaly. This is much more pronounced than in Case A, the deep vortex.

A very interesting feature is the weak cyclonic rotation within the interior of the eddy, near the boundary with the collar. This counterrotation was not present in the wind field used to derive the PV anomaly. The inversion algorithm was presented with a strong anticyclonic PV anomaly, based on zero flow in the interior and a rapidly-moving collar. The smoothing effect of the PV inversion has smeared out the shear region, so the velocity does not merely drop to zero but becomes negative. This model predicts that a fast, narrow collar of GRS-like dimensions will be associated with cyclonic rotation inside it. The weaker vortices (Cases A and B1) did not show this because the anomaly was not large enough.

This work was originally done before the observational evidence of counterrotation within the GRS (see Section 2.2) came to light, and the cyclonic region thought to be an artifact of the model. In retrospect it seems that there is a striking agreement between model and observations.

These three model runs suggest that under the assumptions of this model, the observed cloud motions are equally consistent with both shallow and a very deep balanced flows. This demonstrates the importance of PV structures beneath the clouds, in the region about which little is known. Galileo probe Doppler wind measurements (Atkinson et al., 1996, 1998) indicate that the strong zonal winds are maintained down to a pressure level of at least 22 bar; it is interesting that a similarly deep GRS circulation can be simulated by the present model.

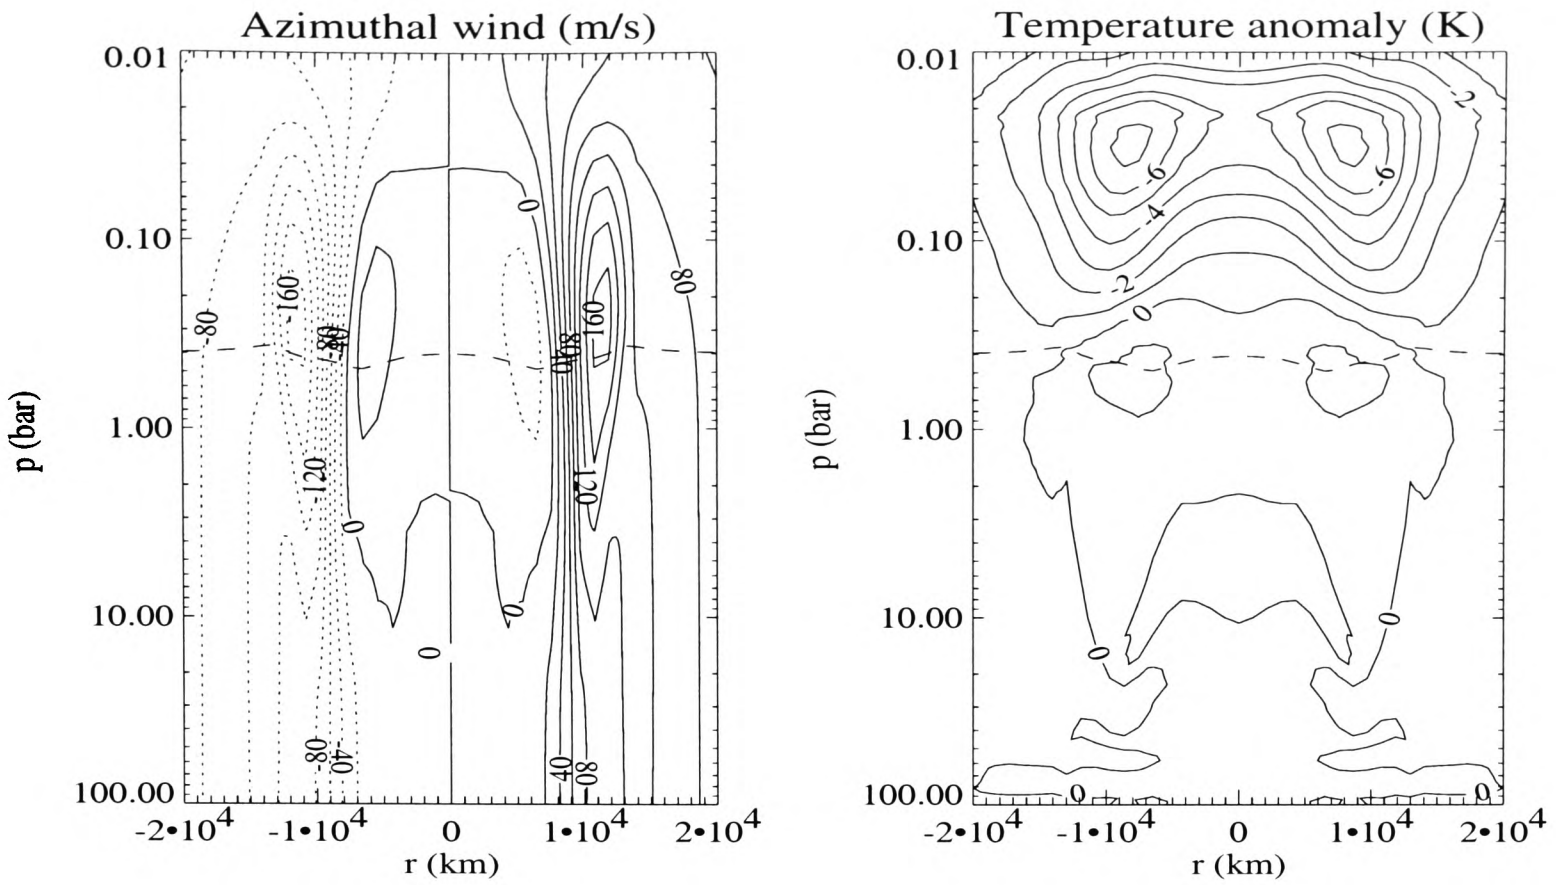


Figure 2.14: Results for a stronger vertical PV dipole. The winds now reach 160 m/s, the cold anomaly is up to 7 K, and the the warm anomaly around the cloud-tops is less pronounced. Note also the cyclonic flow in the GRS interior.

2.4 A model of White Oval BC

We shall now apply exactly the same methods to investigate the structure of White Oval BC. Two cases are presented: Case C, the limiting case of a very deep vortex where there is no vertical shear in the lower part of the model domain, and Case D, a shallow vortex. Case C has a single-layer PV anomaly, Case D a two-layer dipole.

Case C: a deep White Oval model

As in Case A, we calculate the PV anomaly by assuming the winds below the cloud-tops to be constant with depth. The PV structure obtained from this is shown, plotted in both log-pressure and isentropic coordinates, in Fig. 2.15. The anticyclonic anomaly extends into the centre of the vortex, unlike in the GRS which has a much larger "hole" in the centre. The results of inverting this anomaly are plotted in Fig.2.16, The winds are consistent with cloud-top measurements, and as in Case A, the flow is barotropic below the cloud-tops. The cold temperature anomaly above the oval is a single peak, not a torus as for the GRS. This is consistent with the fact that the winds increase linearly from the centre, with no quiescent core.

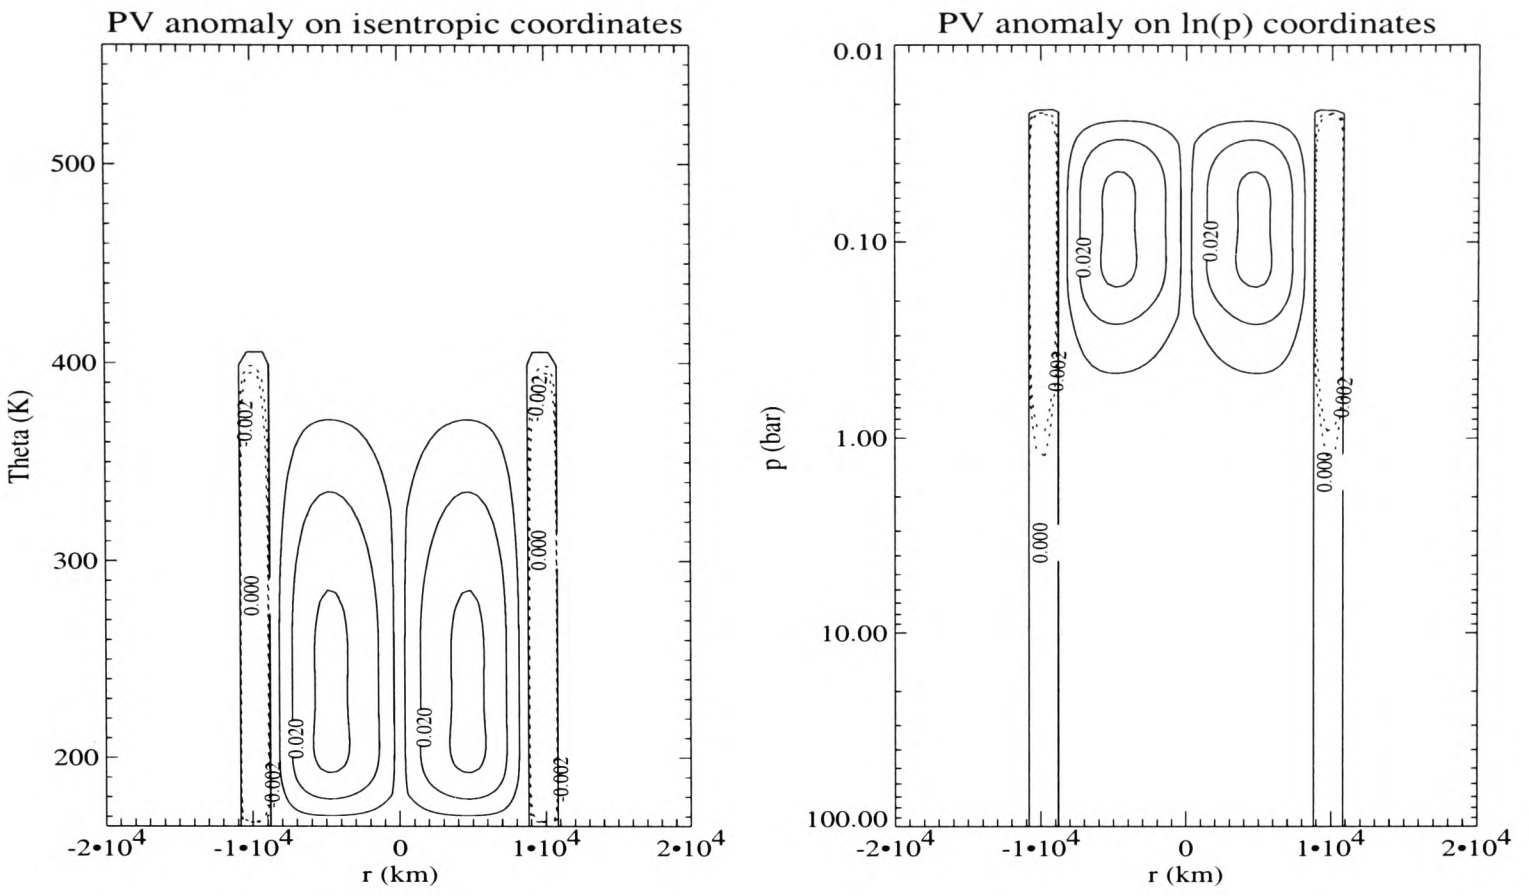


Figure 2.15: PV anomaly structure for a "deep" White Oval model

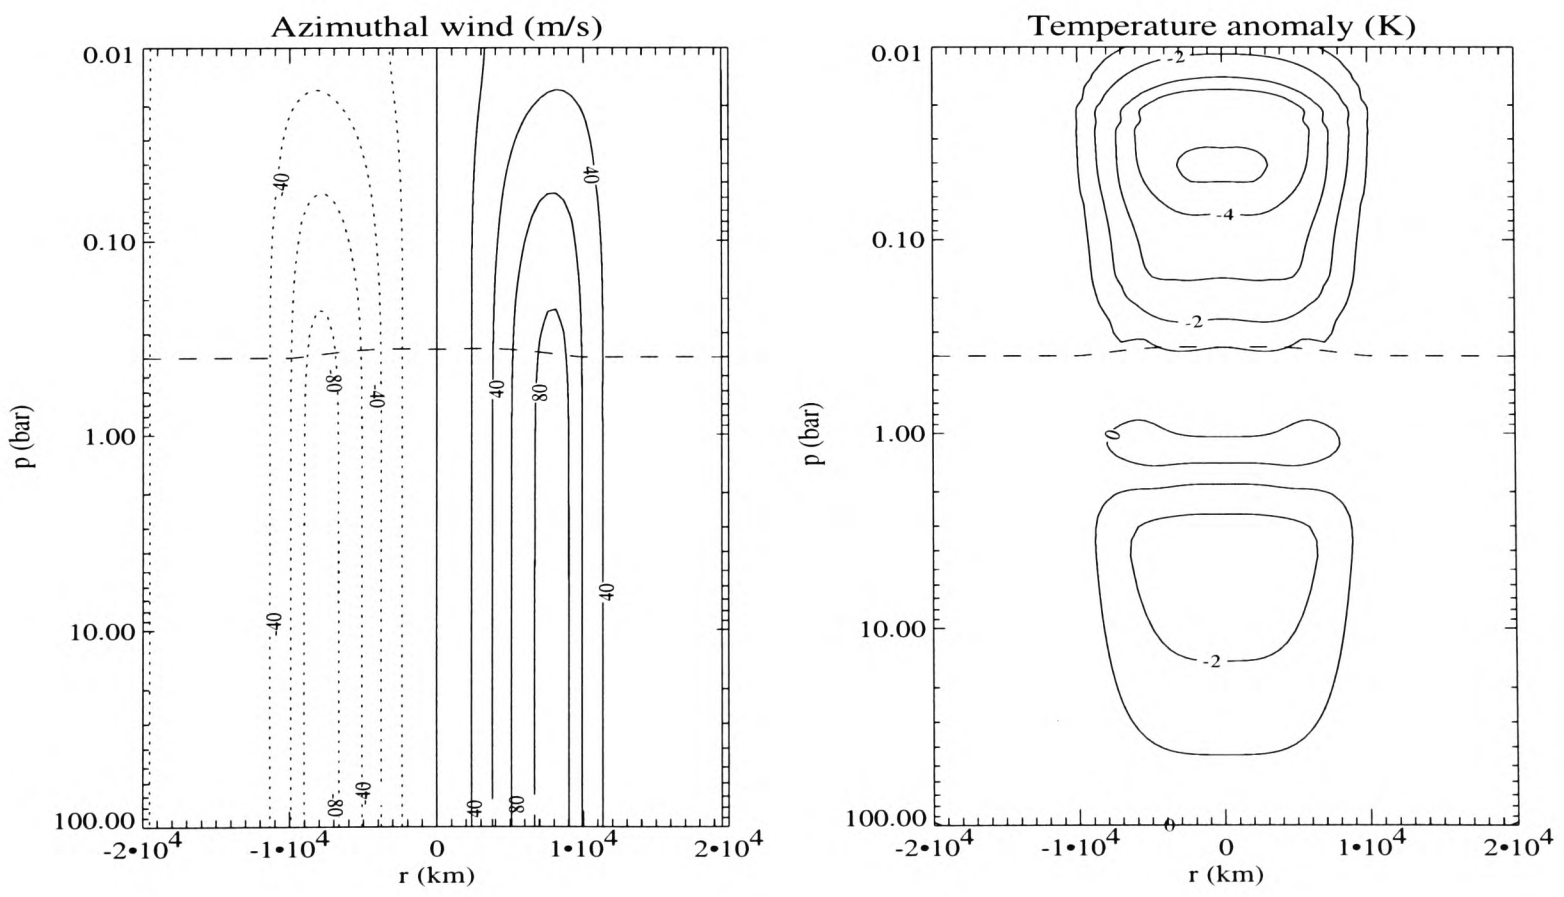


Figure 2.16: Results of a "deep" White Oval model

Case D: a shallow White Oval model

As in Case B, a cyclonic PV anomaly was placed beneath the anticyclonic one to cut off the flow below the cloud tops. The results of the inversion are shown in Fig. 2.17. Above the clouds, the winds and temperatures of Cases C and D are almost identical. In Case D the cold anomaly is slightly weaker and higher in the atmosphere, because the winds reach a maximum just above the clouds. These observable differences are very slight, and yet below the clouds, the shallow vortex decays to around 50% of its maximum value, whilst the deep vortex remains constant. We conclude that, as for the Great Red Spot, observations above the cloud-tops are insufficient to constrain the structure beneath.

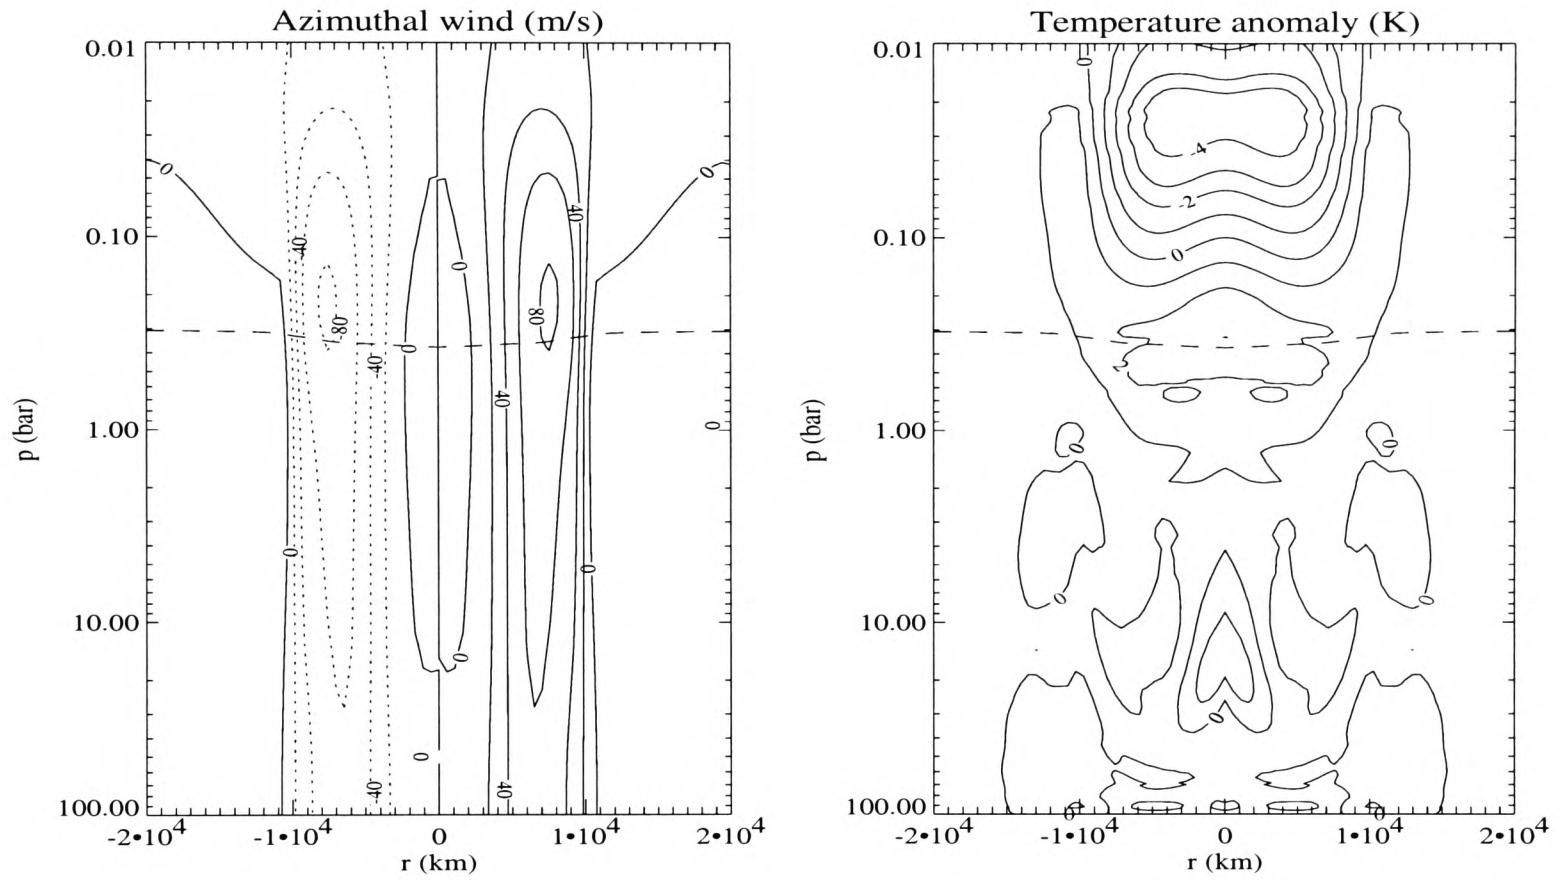


Figure 2.17: Results of a "shallow" White Oval model

Case E: a deep cyclone

The four cases presented so far have all been anticyclonic vortices. Cyclonic systems are also observed in Jupiter's atmosphere, although they are not as common as anticyclones. Examples include the "brown barges" seen in the North Equatorial Belt and in high latitudes in the northern hemisphere. These features have not been modelled in detail here, partly because data of $v(r)$ do not exist, and also because the barges are much more elongated than the giant anticyclones. However, to complete the set of results, we will present the cases of deep and shallow circular cyclones. These do not correspond in detail to any particular feature of the real jovian atmosphere, but illustrate some generic properties of cyclones on Jupiter, and demonstrate that the PV inversion method is not confined to anticyclones.

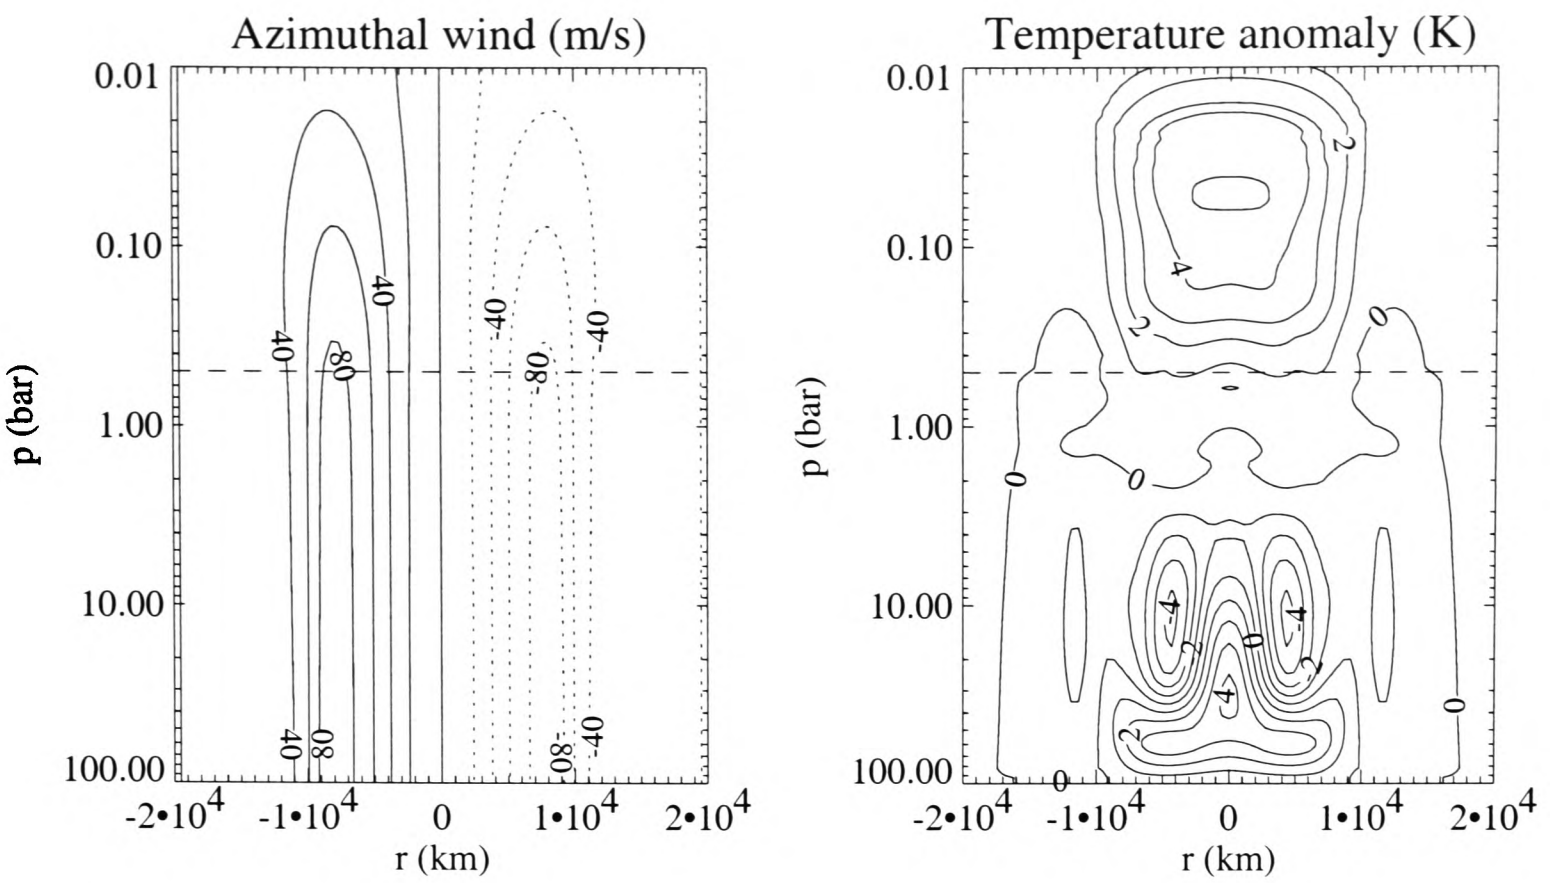


Figure 2.18: Results of a deep cyclone model

Figure 2.18 shows a cyclonic flow, similar to a White Oval but with the opposite sign of vorticity. This is obtained simply by changing the sign of the PV anomaly. This is a deep vortex, comparable with Case C. The wind is of course in the opposite sense, and the temperatures show a warm anomaly above the eddy, as expected. A cross-section of some isentropic surfaces is shown in Fig. 2.19. This shows that the isentropes bend in the opposite direction to those in the White Oval—they are depressed in the centre.

Case F: a shallow cyclone

Case E illustrated a symmetry between deep cyclones and anticyclones; they are induced by equal and opposite PV anomalies. When we come to consider the more complicated dipolar structure of a shallow eddy, the symmetry is broken by the need to satisfy the ellipticity condition $Qf > 0$.

This condition (for well-posedness of the inversion problem) is the same as that for symmetric stability. When setting up a shallow anticyclone, we can add a lower cyclonic anomaly of arbitrary magnitude without violating the ellipticity condition. However, if we wish to construct a cyclone with an underlying anticyclonic anomaly, the anomaly must satisfy $Q_d < Q_0$, which limits the strength of the vortex, if we wish to restrict the flow to a shallow layer.

Accordingly, Case F is a shallow vortex of the same horizontal size as Case E, but smaller wind speeds. Again this vortex is not considered to represent any particular jovian fea-

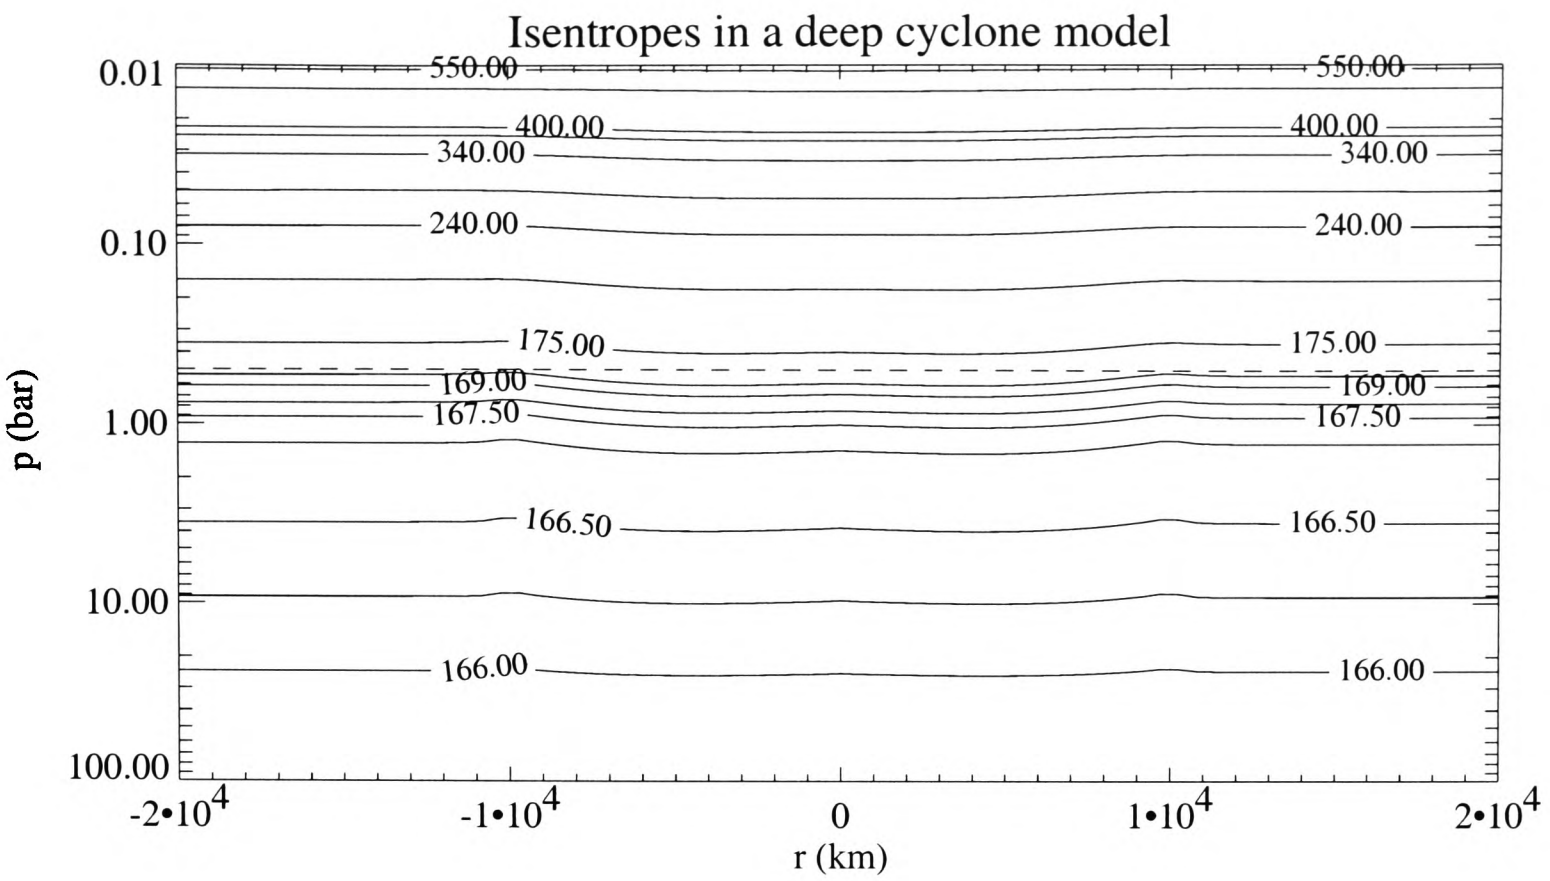


Figure 2.19: Isentropes in a deep cyclone model (Case E)

ture. It is interesting to note that on Jupiter the observed cyclones are fewer and smaller than the anticyclones (and also, most are not even approximately circular). This work does not claim to explain the imbalance, but it is suggestive that, in this model, it is more difficult to construct shallow balanced cyclones than anticyclones.

2.5 Discussion of results

The results presented in this chapter have shown that the method of IPV inversion is able to achieve quite close quantitative agreement with analyses (Mitchell et al., 1981) of the tangential flow component in the Great Red Spot and other giant eddies. Several different balanced flows, consistent with the observations, have been examined. Anomalous PV distributions must extend beneath the clouds; otherwise the winds would be weaker than observed. According to this model, even small values of PV anomaly beneath the clouds (compared to those above) have a large effect on the observed circulation. Therefore a knowledge of the thermal structure beneath the clouds is important in eliminating hypotheses.

Two types of PV anomaly have been considered, namely vertical monopoles and dipoles. The model circulations all feature an anticyclonic PV anomaly, and "bowing" of isentropic surfaces. The diagnostic model makes no direct statements about the processes generating and maintaining the Spot, but some indirect inferences are possible. A source of diabatic heating (e.g. latent-heat release from a condensible) below the clouds could adjust the isentropes in the way seen in the model: raising them above and lowering them

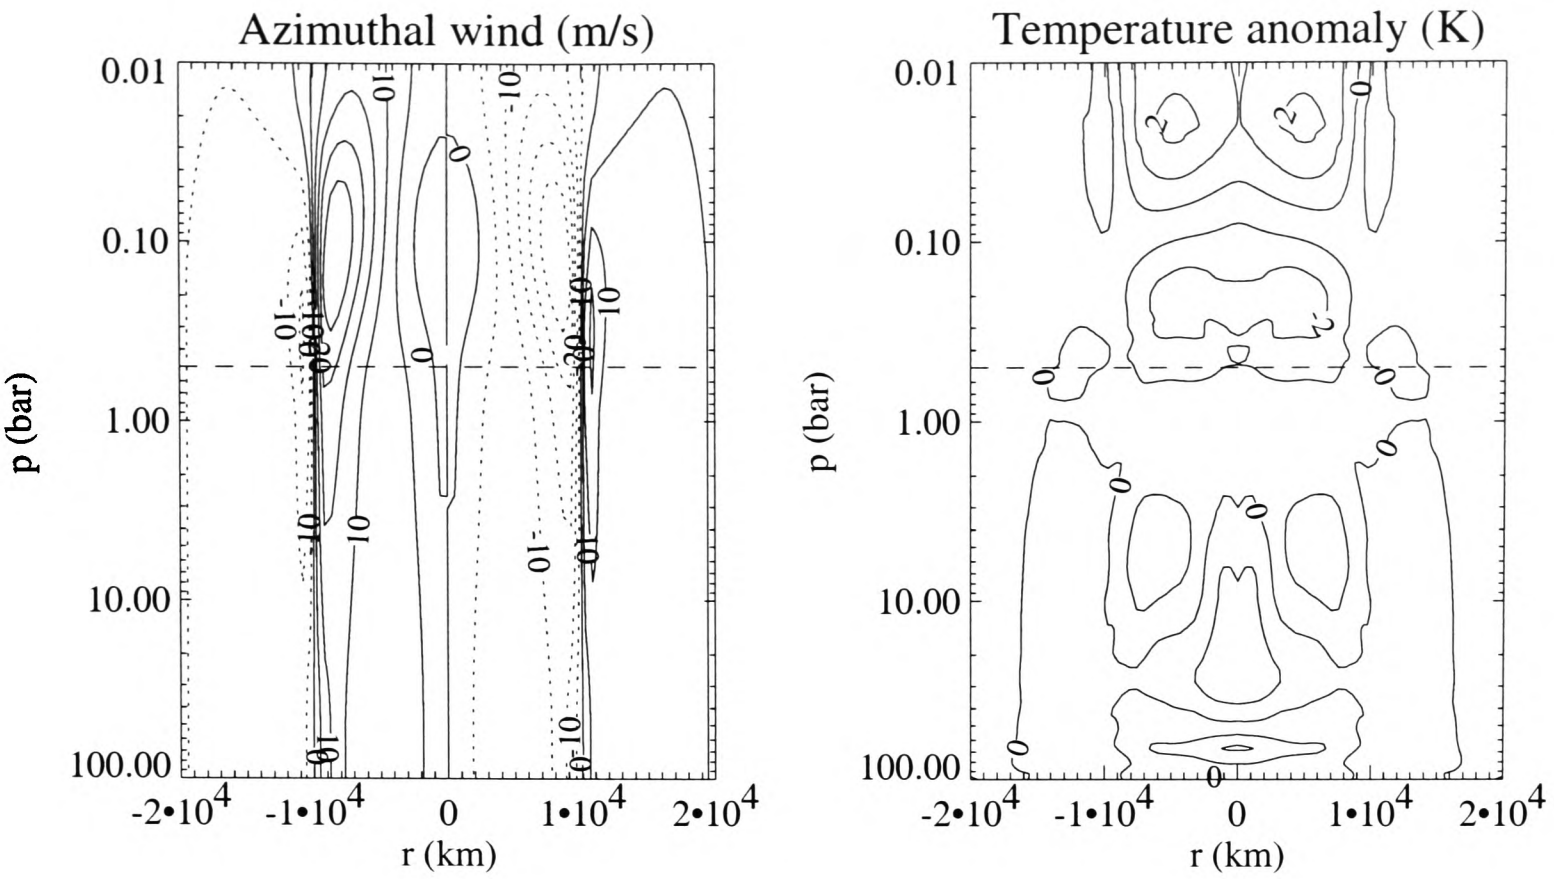


Figure 2.20: Results of a shallow cyclone model (Case F)

below. The heating would generate upwelling, with horizontal divergence above and convergence below. The Coriolis effect acting on this flow spins up an anticyclone above the heating source, with cyclonic flow beneath it. This gives rise to the vertical dipole structure of PV. Later, in Chapter 5, we shall model this process of baroclinic eddy generation, in the time-dependent 3D GCM.

This idea, loosely analogous to a terrestrial hurricane, provides a possible explanation for the presence of the PV structures (anticyclonic above and cyclonic below), which we have seen can be inverted to give realistic flows. Flasar et al. predicted a vertical scale of 100 km for the flow below the clouds. Our baroclinic model vortex is "trapped" within a height scale of similar magnitude.

On the other hand, we have seen that monopole PV anomalies (one thick layer, with the sign of PV constant with height), can also induce circulations which look similar above the clouds, although completely different beneath them. This type of anomaly could be generated by *horizontal*, essentially barotropic, as opposed to vertical mixing. Meridional transport, conserving PV, causes an air parcel to have different PV from its new surroundings. This creates an anomaly which would not change sign with height.

In Chapter 5 we shall illustrate how the Jupiter GCM can spontaneously generate both shallow and deep vortices by mixed barotropic-baroclinic instability. We shall also look at experiments with shallow vortices initialized by brief "bubble" heating events, which are localized in space and time. The GCM vortices bear a surprising resemblance to those studied in this chapter, reinforcing our confidence in the results and supporting the conclusion that Jupiter's atmosphere may feature both shallow and deep vortices.

Chapter 3

A simple general circulation model for Jupiter

All numerical models are wrong. The question is, are they uselessly wrong, or are they wrong in such a way that they can still tell us something? (from *The Economist* magazine)

3.1 Introduction to the UK Met. Office Unified Model

The starting-point for development of the present model is the UK Meteorological Office Portable Unified Model or UM. This is a large and flexible software package, used both for operational numerical weather prediction and for climate research. The portable version runs on a workstation and recently has been made available for scientific use outside the UKMO.

A general circulation model may be viewed as comprising two main scientific sections: 'dynamics' and 'physics'. The physics section includes parametrizations of processes such as thermodynamics, radiation, and unresolved dynamical processes such as gravity wave drag. The dynamics code solves the primitive equations of fluid flow for explicitly resolved motion, using forcing parameters derived from the physics. The present work to date has made use of the UM's dynamics section only; we shall not discuss the standard (terrestrial) model physics here. Section 3.2 outlines the relevant features of the model, starting with the scientific basis for the model equations. The method of solution, including the boundary conditions, is summarized. Section 3.3 then describes the changes made to the standard configuration, to produce a model suitable for Jupiter, with very simple physics.

3.2 Overview of the Unified Model numerical scheme

The UM integrates the dynamical equations derived by White and Bromley (1995). A spherical coordinate system with a hybrid sigma/pressure vertical coordinate (η) is used. The model has a finite-difference formulation with a split-explicit integration scheme. This scheme is described in detail by Cullen et al. (1993), and more briefly by Cullen and Davies (1991).

The model equations

The meteorological primitive equations of White and Bromley (1995) provide an unusually close approximation to the underlying Navier-Stokes equations in a rotating spherical domain. They include a representation of all three components of the Coriolis effect, and do not assume the atmosphere to be shallow relative to the interior. At present, the use of this system is peculiar to the UM among meteorological models, so we shall quote the equations in full and examine their novel features.

Most current models are based on the hydrostatic primitive equations (HPEs). These are obtained from the Navier-Stokes equations by a rigorous approximation and scaling procedure, designed to simplify the problem without sacrificing important conservation properties. Broadly speaking, the procedure relies on the fact that the atmosphere can be thought of as a thin spherical envelope of gas, in which the vertical terms are dominated by hydrostatic balance. The HPEs are discussed in detail by e.g. Hoskins et al. (1985), and the derivation and scaling procedures are explained in several meteorological and oceanographic textbooks, including Pedlosky (1982) and Gill (1982).

Firstly the 'shallow-atmosphere approximation' is made, replacing all radii r with the planet's mean radius a. Some small metric terms are omitted, as are terms in $2\Omega\cos\phi$ (the so-called $\cos\phi$ Coriolis terms), together with the vertical acceleration and the vertical component of friction. The main criterion for choosing terms to retain is that the resulting HPEs should implicitly conserve suitably-defined measures of energy, angular momentum and potential vorticity.

If the locally-horizontal components of the Coriolis terms are included, further changes to the equations become necessary in order to retain the consistency and conservation properties of the HPEs. In particular, a term $2\Omega w\cos\phi$ represents the height variation of planetary angular momentum; this is inconsistent with the shallow atmosphere approximation. White and Bromley rejected this approximation, replacing a with a 'pseudo-radius' r_s which varies with geopotential height. A few other, very small, terms are retained for conservation reasons. The result is a set of equations which are more accurate than the primitive equations, yet still conserve energy, angular momentum and PV.

These equations are summarized below, written in spherical geometry. In general, the Unified Model grid may be rotated with respect to the real planet, so the model poles need not be coincident with the geographical ones. This is useful when running a limited area model, because it minimizes problems due to the convergence of the meridians. However, the regions of interest here are the tropics and midlatitudes, where a rotated grid is unnecessary. Therefore we shall not consider the generalized equations for rotated grids. The full, more complicated equations are set out in the Unified Model documentation (UMDP 10).

First define the three Coriolis terms, which have simple forms in the case of a nonrotated grid (as used in our work):

$$\begin{aligned} f_1 &= 0 \\ f_2 &= 2\Omega\cos\phi \\ f_3 &= 2\Omega\sin\phi \end{aligned} \tag{3.1}$$

Also, the derivative following the motion is

$$\frac{D}{Dt} \equiv \frac{\partial}{\partial t} + \frac{u}{r_s\cos\phi}\frac{\partial}{\partial\lambda} + \frac{v}{r_s}\frac{\partial}{\partial\phi} + \dot{\eta}\frac{\partial}{\partial\eta} \tag{3.2}$$

The momentum, thermodynamic and continuity equations are then

$$\frac{Du}{Dt} + \frac{1}{r_s\cos\phi}\left[\frac{\partial\Phi}{\partial\lambda} + \frac{R}{p}\left(T_v + \mu T_s\right)\frac{\partial p}{\partial\lambda}\right] - f_3 v + f_2\tilde{w} - \frac{u}{r_s}\left(v\tan\phi - \tilde{w}\right) = \mathcal{F}_u \tag{3.3}$$

$$\frac{Dv}{Dt} + \frac{1}{r_s}\left[\frac{\partial\Phi}{\partial\phi} + \frac{R}{p}\left(T_v + \mu T_s\right)\frac{\partial p}{\partial\phi}\right] + f_3 u + \frac{u^2\tan\phi}{r_s} + \frac{v\tilde{w}}{r_s} = \mathcal{F}_v \tag{3.4}$$

$$\frac{D\theta_L}{Dt} - \frac{1}{\pi c_p T}\left[L_c q_c^L + \left(L_c + L_f\right)q_c^{(F)}\right]\frac{RT\omega}{c_p p} = \mathcal{F}_\theta \tag{3.5}$$

$$\frac{Dq}{Dt} = \mathcal{F}_q \tag{3.6}$$

$$\frac{\partial}{\partial\eta}\left(r_s^2\frac{\partial p}{\partial t}\right) + \frac{1}{\cos\phi}\left[\frac{\partial}{\partial\lambda}\left(u r_s\frac{\partial p}{\partial\eta}\right) + \frac{\partial}{\partial\phi}\left(v r_s\cos\phi\frac{\partial p}{\partial\eta}\right)\right] + \frac{\partial}{\partial\eta}\left(\dot{(\eta)} {r_s}^2\frac{\partial p}{\partial\eta}\right) = 0 \tag{3.7}$$

The quantities $\mathcal{F}$ are source terms, which include numerical diffusion. L_c and L_f are the latent heats of condensation and fusion respectively. The thermodynamic variables θ_L and q_T are given by

$$\theta_L = \theta - \frac{L_c q_c^{(L)} + (L_c + L_f)q_c^{(F)}}{c_p\Pi} \tag{3.8}$$

$$q_T = q + q_c^{(L)} + q_c^{(F)} \tag{3.9}$$

The omega equation, used to diagnose the vertical motion field ω, reads

$$\omega = \frac{-1}{r_s\cos\phi}\int_0^n\left[\frac{\partial}{\partial\lambda}\left(u r_s\frac{\partial p}{\partial\eta}\right) + \frac{\partial}{\partial\phi}\left(v r_s\cos\phi\frac{\partial p}{\partial\eta}\right)d\eta\right] + \left(\frac{u}{r_s\cos\phi}\right)\frac{\partial p}{\partial\lambda} + \left(\frac{v}{r_s}\right)\frac{\partial p}{\partial\phi} \tag{3.10}$$

and $\tilde{\omega}$, an approximation to the vertical velocity may defined as

$$\tilde{w} = -\frac{RT_s\omega}{gp} \tag{3.11}$$

The hydrostatic equation is given by

$$\frac{\partial\Phi}{\partial\eta} = -\frac{R(T_v + \mu T_s)}{p}\frac{\partial p}{\partial\eta} \tag{3.12}$$

$$= -c_p\left(\theta_v + \mu\theta_s\right)\frac{\partial\Pi}{\partial\eta} \tag{3.13}$$

where $\Pi = (p/p_*)^\kappa$ is the Exner pressure, and θ_v the virtual potential temperature. T_s is the reference-state standard temperature. There remain two other quantities: the pseudo-radius r_s is

$$r_s(p) = a + \int_p^{p_0}\frac{RT_s(p')}{gp'}dp' \tag{3.14}$$

and μ is defined as

$$\mu g = f_2 u - f_1 v + \frac{(u^2 + v^2)}{r_s} \tag{3.15}$$

The omission of the Coriolis term f_2 from the usual HPEs can lead to significant error in the solutions. To quantify this, Gill (1982) analysed a simpler system: the linear Boussinesq equations for a stratified fluid. He concluded that the two additional Coriolis terms, f_1 and f_3, may be neglected if the rotation rate Ω and buoyancy frequency N satisfy the condition $2\Omega \ll N$. White and Bromley argued that even though this condition is usually satisfied in the Earth's atmosphere and ocean, the $\cos\phi$ Coriolis terms are still important in planetary-scale flows and in diabatically driven tropical flows. The case for including them in Jovian models is even more compelling: in Jupiter's rapidly rotating, weakly stratified troposphere, $2\Omega \sim N$ and Gill's condition is seriously violated. The consequences of this are unknown, but it would be straightforward to test any hypothesis put forward. The extra Coriolis and metric terms may be removed from the model by simply clicking a button within the UM user interface. The effect could be compared against a model run which is identical in all other respects. This experiment has not been performed and we mention it merely as an interesting possibility for future work. The most difficult part of the task would be choosing a model case which would unambiguously show the importance of the horizontal Coriolis terms.

The model grid

The prognostic variables used are u, v, θ, q and Φ, being respectively the zonal and meridional wind components, potential temperature, specific humidity and geopotential. They are stored on the staggered Arakawa 'B' grid (Arakawa and Lamb, 1977). In a global model, the grid is arranged so the poles are θ points; this avoids the problems of storing wind components over the pole. The vertical coordinate η is a hybrid, being a terrain-following (σ) system at the base of the model, with a smooth transition to a pure pressure

system at the top. The variables u, v, θ, q and Φ are held on 'full levels' η_k, while $\dot{\eta}$ is held on 'half levels' $\eta_{k+1/2}$. The half-level pressures are defined in terms of the surface pressure p_* and constants A and B:

$$p_{k+1/2} = A_{k+1/2} + B_{k+1/2} p_* \tag{3.16}$$

The full-level pressures are then calculated using the pressure-weighted mean of the half-level Exner pressures $\Pi \equiv (p/p_0)^\kappa$. The lowest half-level is always a sigma level (i.e. $A_{1/2} = 0$), thus ensuring $p_{1/2} = p_*$. The number and spacing of vertical levels may be modified to suit the requirements of a particular experiment.

The split-explicit integration scheme

The concept of operator splitting (see e.g. Press et al., 1986) leads to a very efficient scheme. The operator of interest is written as a linear sum of several other (in general non-linear) operators. These may be integrated separately and the results combined after each timestep. The advantage of this method is that a different, optimal scheme may be chosen for each operator.

In the UM, the integration of each timestep is split into two parts: adjustment and advection. In the adjustment step, the pressure gradient and Coriolis terms are used in a forward-backward explicit scheme to update the pressure, temperature and wind fields. The vertical advection of potential temperature is also calculated here. In the advection step, the adjusted velocity field is used to calculate the increments due to the nonlinear advection terms. The fourth-order Heun conservative advection scheme (Mesinger, 1981) is used. A consideration of the Courant stability criterion shows that the length of the timestep is limited by the maximum group velocity of a wave solution of the equations. In the adjustment step this is the gravity wave speed, but in advection it is the highest actual wind speed. In practice the former is larger by a factor of about three, so three adjustment steps may be taken for each advection step. This allows a substantial saving in execution time. The timestep is usually chosen to give a Courant number of around 0.1.

One disadvantage of the split-explicit scheme, compared with an implicit one, is that numerical instabilities can arise. A conservative horizontal diffusion operator is applied to filter out short-wavelength noise. Diffusion may be applied to the winds, potential temperatures and humidity; and separately to the horizontal divergence. Vertical diffusion of momentum is also used. The choice of diffusion parameters requires some care, and is discussed in Section 3.3.

Boundary conditions

In the standard version of the Unified Model, the vertical (top and bottom) boundary conditions are equivalent to two rigid lids:

$$\dot{\eta} = 0 \text{ at } \eta = 0, 1 \tag{3.17}$$

As the coordinate η is purely terrain-following at the lowest level, the lower boundary condition is entirely natural for a planet with a solid surface, such as the Earth. On Jupiter, however, there is no obvious physical boundary condition: indeed, even the height of the boundary between the model atmosphere and interior is to some extent an arbitrary choice. These issues will be addressed in Section 3.3, along with boundary conditions on temperature.

The model may be run in either global mode or as a limited-area model (LAM). In the latter mode, lateral boundary conditions are required. In operational use and some scientific work, boundary tendencies (as a function of time) are input from a second model whose domain includes that of the first. A typical example is that of the British Isles operational LAM, whose boundary values are generated by a run of the global forecast model over the same time period (see the UM User Guide).

For idealised experiments or short runs, fixed boundary conditions may be used. In this case, the values of the prognostic variables at the lateral boundaries are simply not updated. They therefore retain their initial values for all time. The vertical velocity (a diagnosed quantity) is set to zero on the boundaries. These boundary conditions are unsuitable for the Jovian models described here, and Section 3.3 explains the modifications made to the scheme.

Dynamical core experiments

The Unified Model in its standard climate research configuration is a very large and complicated set of programs, with detailed physics schemes tailored to the case of the Earth. The version from which this work commenced is a much-simplified one, with all calls to the physics routines removed. They are replaced with just two simple schemes: boundary-layer drag and radiative relaxation. The former applies linear Rayleigh friction to the lowest few model levels. The latter models all radiation effects as Newtonian relaxation towards a simple zonal-mean, climatalogical temperature profile. There is no diurnal or other cycle in the radiative forcing.

Other simplifications include the removal of oceans and orography, to depict the planet as a smooth sphere with a dry atmosphere. The stratosphere is not resolved, model levels being concentrated in the troposphere. (The levels can easily be changed, as will be

described below.) There is no convective adjustment at all, no vertical diffusion or gravity wave drag, nor heat exchange with the soil. The code implementing these modifications to the UM was generously provided by Dr. T. Davies of the UKMO.

The original purpose of this so-called "dynamical core" model was to allow direct comparison between grid-point dynamics schemes (such as the UM) and competing spectral methods (Held and Suarez, 1994). The dynamical core configuration also makes a good starting point for the development of a Jupiter model. One would expect the dynamical equations to be equally valid for Jupiter as for the Earth (except for the effects of oblateness), but the physical conditions are entirely different. Sophisticated terrestrial parametrizations are therefore of little use for our purposes, without very extensive modifications.

3.3 A Jovian version of the Unified Model

Sections 3.1 and 3.2 have described the Unified Model numerical scheme and a dynamical core version of the model. The following sections will detail the original work carried out to implement a model of Jupiter's atmosphere, starting from the dynamical core terrestrial model.

We shall briefly discuss the available options for the model domain in the horizontal and vertical, and suitable gridpoint resolution. The details of (and reasons for) choices made for particular experiments will be left until Chapter 5. The rest of this chapter will discuss the implementation of boundary conditions, the simple physics so far included in the model, and the numerical diffusion scheme.

Many of the decisions made in developing the model were rather arbitrary, and in the text we will draw attention to alternatives which may be worth future investigation. In the early stages of model development, the emphasis must be on making reasonable choices which work, in the sense of generating plausible results. It is not possible to compare the relative advantages of individual features until the whole model is running. Hopefully the results presented in this thesis will provide a basis for future refinement of individual aspects of the code.

Global planetary parameters

We have seen that our model Jupiter began as a smooth sphere with a dry nitrogen-oxygen atmosphere, and terrestrial radius and rotation rate. The first task was to substitute jovian values for these global parameters. Fortunately, the UM is well designed and

these parameters are defined and referenced as symbols, not as literal numbers. Therefore it sufficed to change the declaration of each value. The table below shows the values used.

Quantity	Symbol	Value
Equatorial radius	A	71,994,888 m
Rotation rate	Ω	1.7737×10^{-4} rad/s
Gas constant	R	3745 J/mol
Specific heat cap.	C_p	13107.5 J/kg
Acceleration due to gravity	G	24.4 m/s^2

Horizontal domain and resolution

'What's the good of Mercator's North Poles and
Equators,
Tropics, Zones and Meridian lines?' So the Bellman would cry: and the crew would reply,
'They are merely conventional signs!' (Lewis Carroll: *The Hunting of the Snark*)

Jupiter's zones and belts, and their associated jetstreams and eddies, show structure on scales of up to a few degrees in latitude and longitude. For example, the latitudinal width of the easterly jet centred at 18°S (planetocentric) is around 5 degrees, while that of the westerly at 25°S is 7°. The Great Red Spot lies in the intervening shear zone; it extends roughly 6° in latitude and 15° in longitude. The White Ovals are of order half this size.

It seems clear that a gridpoint spacing of one degree or less is needed to resolve the structure of the main jets. To resolve the smaller eddies adequately, a resolution of half a degree (or better) is preferable. Current Jupiter models also need to represent the vertical structure of the atmosphere over a suitable height range (discussed below). This requires the use of twenty to forty vertical levels, depending on the vertical extent of the model. A third factor is that Jupiter's atmosphere is weakly forced and weakly dissipative by comparison with the Earth. The frictional and radiative timescales are much longer than those of the Earth's atmosphere or even the oceans. Long integrations are therefore required in order to reach a realistic equilibrium state.

These considerations restrict our scope to limited area studies. A rectangular domain can be chosen which is wide enough to include several alternating jets, and long enough so that asymmetric features such as eddies do not fill the whole domain. Running on a modern workstation, a typical model on a 60° × 21° domain with half-degree grid spacing and 30 levels, can integrate 100 model days in about 12 CPU hours. Coarser models, with a longer timestep, run correspondingly faster. Comparable recent work, such as the

model of Williams (1996), has used grids of $1° \times 1°$ or $1° \times 2°$ resolution, and 20 vertical levels. These models had domains of up to 40° latitude by 180° longitude.

Vertical domain

Jovian dynamical models differ widely in their treatment of the atmosphere's vertical structure. At one extreme are the single layer shallow-water or QG models, sometimes with bottom topography to represent the deep flow below the 'weather layer'. At the other extreme are models such as that of Williams (1996), which use several tens of levels to treat the atmosphere as a fluid of depth up to 10000 km.

The Williams (1996) model is analogous to an ocean, where the visible clouds are assumed to lie at or near the top of the model, and the atmosphere below the cloud-tops is modelled as a Boussinesq liquid with a variety of idealised stratification structures. Typically the Brunt stability parameter varies exponentially with depth, creating a shallow active layer overlying a much deeper abyssal layer. The upper layer is relatively strongly stratified, while the abyss approaches neutral static stability.

The present model is based not on a Boussinesq ocean, but on a compressible atmosphere, and the domain is chosen to reflect the region amenable to observation, i.e. the troposphere and lower stratosphere. An attempt is made to use the most realistic vertical stratification and temperature profiles, in contrast to Williams' approach of exploring a wide range of idealised structures in the deep atmosphere where there are no observations. The subject of vertical structure will be addressed in Chapter 4.

The choice of the height of the top boundary, and the boundary condition there, follow naturally from observational evidence. Imaging results (Limaye, 1986) and infrared observations (Hanel et al., 1979a) suggest that the observed cloud-top winds decay with height, towards a rest state at $p \simeq 10$–20 mb. The top half-level is therefore placed at 10 mb, and zero winds imposed on the top full-level, at around 11.3 mb. In most experiments a "sponge layer" is required at the top. This is a region in the top few levels where friction is applied to the winds, to damp spurious reflections from the top boundary.

The depth of the bottom model level is more arbitrary. The Galileo probe measured winds and temperatures down to 22 bar (Atkinson et al., 1996; Seiff et al., 1996); a level rather below the expected water clouds at $p \simeq 5$ bar (West et al., 1986). The choice of depth of the bottom boundary is a compromise between adequate coverage and computational resources. The current version of the model has a rigid lid at 10 bar, with the deepest full level at 8.3 bar, slightly below the water cloud depth. This domain, 10 bar $> p >$ 10 mb, spans roughly six scale heights. With 30 levels, uniformly spaced in $\log p$, this gives a vertical resolution of 10 levels per decade in pressure. This permits a detailed study of

the regions close to the tropopause, which lies near 100 mb.

A variety of boundary conditions have been applied at the base of the model. The bottom level is always taken to be an isentrope, so the potential temperature there is fixed. For the study of shallow, trapped vortices the lower boundary condition has little effect on the flow. Experiments on deep vortices require a barotropic boundary condition. The UM is designed with a solid surface as its lower boundary, but this may be replaced with a zero-divergence condition, the temperatures constrained to be isothermal, and the vertical wind shear to be zero. This configuration gives barotropic flow at the base of the model. The results of the Galileo Probe Doppler Wind Experiment suggested that this is likely to be a good model of the flow at the 10 bar level on Jupiter. In vortex generation experiments, mechanical forcing is applied to the wind field to simulate the effect of the deep atmosphere. The forcing coefficient is artificially large, and this is not intended to be wholly a realistic parametrization of a boundary layer. Section 3.3.2 gives details of the schemes used to force the winds.

Temperatures at the lower boundary are constrained by the assumption that the lowest level is an isentropic surface. The "surface" pressure p_* is allowed to vary, and T set appropriately so that θ is equal at all points. Alternatively, it would be possible to constrain p_* instead of u and v, which would be equivalent to imposing bottom topography. We chose to force the winds, and not the surface pressure, mainly for practical reasons. The wind and temperature fields at all levels are useful, observable quantities. In contrast, surface pressure and bottom topography on Jupiter are only analogies with no literal existence. We have been able to ignore p_* in our diagnostics, concentrating on winds instead. We note, however, that the alternative method may have desirable numerical properties and is perhaps worth investigating in the future.

The standard Unified Model is not intended to run with a bottom level deeper than about 1.1 bar, and the user interface will not allow levels to be specified deeper than this. However, altering the user interface validation code sufficed to lift this restriction. The main model code was also modified to set the new reference pressure at the start of each run. This includes continuation runs starting directly from a model dump without reconfiguration.

There seems no reason in principle why the model cannot be extended to greater depths, either by adding more levels or by making the resolution more coarse in the deep atmosphere. If, in the future, the effect of condensibles is parametrized in the model, it will be desirable to increase the depth, perhaps to 100 bar or deeper. This should allow model cloud layers to be localised well within the vertical domain. Unfortunately, in practice this may require either a coarser vertical resolution or else shorter runs, in order to keep the total program execution time reasonable.

Lateral boundary conditions

Neither of the standard UM boundary conditions is appropriate for the present work. The use of lateral boundary tendencies would require a global model to provide the boundary data. The alternative, fixed boundary values, is unrealistic and allows only stationary (not travelling) waves. A more suitable model is a channel with impermeable sidewalls to the north and south, and a periodic boundary condition at the east and west. For channels shorter than 360° in longitude, periodicity is not entirely realistic because the curvature is wrong, but it does have several advantages. It allows the zonal-mean flow to evolve freely, constrained only along the sidewalls. It allows travelling waves to propagate. Finally, observations of Jupiter show a high degree of zonal symmetry, so the results of a periodic channel model can resemble observations.

An obvious way to ensure periodicity is to copy the values of the prognostic variables, from gridpoints at each end of every row to those at the other end, after each timestep. This gives a rim region (typically of four points) at each end where the fields match those at the other end. Unfortunately, this procedure introduces discontinuities at the inner edges of the rim. These discontinuities inevitably lead to rapidly growing numerical instabilities unless very heavy diffusion is applied. If smoothing is applied to the rim regions only, true periodicity is lost. Waves reaching the ends of the channel are damped and tend to decay. A further minor disadvantage is that storage and time is wasted, because several columns of the model arrays hold identical data.

After some effort, this scheme was rejected as unsatisfactory, and a much more effective alternative was implemented. It is conceptually simple: we use the LAM fixed boundary conditions on the sidewalls, and the global model's boundary conditions on the ends. The global UM maps the sphere onto a rectangular array of points by means of a latitude-longitude grid. It ensures periodicity at the ends of the array by using a modified differencing scheme. Consider a row of points $i = 1 \ldots I$ where the longitude of point i is $\lambda = hi$. The zonal derivative of a quantity q at point i is approximated discretely by

$$\frac{\partial q}{\partial \lambda} \simeq \frac{q^{i+1} - q^{i-1}}{2h}$$

At the end of the row, the next point eastward of point I is the point $i = 1$, by periodicity. Therefore the zonal derivative at $i = I$ is given by

$$\frac{\partial q}{\partial \lambda} \simeq \frac{q^{1} - q^{I-1}}{2h}$$

A similar expression is used at $i = 1$, and advection across the periodic boundary is treated using the same principle.

In the standard Unified Model, the choice of boundary conditions is implemented by means of a Boolean flag passed to the `nupdate` source code preprocessor. Code suitable for the chosen domain (global or LAM) is inserted into the source file. When the

model runs, each operation in the adjustment and advection steps is initially applied to a whole row. In LAM mode the boundary values are then set to their original values, while in global mode the end values are recalculated as described above. The global mode also performs corrections at the Poles.

To achieve the desired changes, selected `nupdate` directives are removed, to make the model treat zonal boundaries as usual (that is, as either Poles or sidewalls) but always to treat meridional boundaries as if in the global mode. The result, run in LAM mode, is a hybrid model with LAM zonal boundary conditions but global, periodic meridional ones. This is efficient and truly periodic, causes no problems with discontinuities, and avoids wasting storage. To effect this, changes were made to twelve different subroutines, including those indirectly affected, such as interpolation routines. The modifications are all minor, due to the flexible structure of the UM.

The no-slip sidewall boundary conditions are implemented by imposing $u = v = 0$ at all points on the walls; the values are never updated from this. The continuity equation is not violated, because the vertical velocity is unconstrained. (In fact when ω is diagnosed and output, the values are set to zero on the boundaries, but ω is not a prognostic variable, so this does not affect the dynamics). The use of non-zero sidewall velocities would improve the flexibility of our model, but the implementation difficulties have not yet been solved.

A particular case is that of a free-slip sidewall placed on the equator. This would be desirable since it is allows us to model only half the equatorial jet, saving time and storage. Difficulties arise in satisfying the continuity equation within the rim region, and the flow tends to generate strings of small eddies along the sidewall. After several attempts, trying to conserve velocity or angular velocity, and experimenting with a rim region, this problem was shelved.

A simple way of avoiding the difficulty is to choose the domain so that the sidewall boundary condition $u = v = 0$ is appropriate. The sidewalls are placed on latitudes where the observed cloud-top zonal flow vanishes; the standard (fixed) boundary conditions then work well. The domain may straddle the equator without problems. The pattern of alternating jets on Jupiter means that suitable sidewall latitudes are easily found. We do implicitly assume that because the zonal wind is zero or small at the cloud-tops, it is also zero at all other levels. This seems unavoidable in a limited-area model, unless we were to employ sloping sidewalls, which seems much too elaborate.

Numerical diffusion

Diffusion is required in the model for two reasons. The first is that the explicit scheme is unstable to short wavelength disturbances. Any numerical computation is subject to

truncation error (in the use of Taylor's theorem to discretize the equations) and to round-off error (in representing variables with finite numerical precision). These errors lead to numerical noise, which in an unstable scheme will grow. This leads to problems such as grid-splitting. A typical result is that the prognostic fields show very large oscillations with a wavelength of two gridpoints. The fields have very large (and spurious) gradients, which rapidly introduce unphysical features such as negative pressure or temperature, and crash the model. Horizontal diffusion controls such instabilities.

The second reason is a physical one which justifies this apparently artificial parametrization. This is the need to represent the cascade of energy and enstrophy from large to small length scales. The direct effect of viscosity upon a synoptic-scale flow is completely negligible, and yet its net dissipative effect is not. The smallest scales, on which molecular viscosity is important, are not represented by the primitive equations. Neither are turbulent effects which draw energy from the large-scale flow and pass it to smaller scales. Numerical diffusion may be regarded as an attempt to parametrize turbulence and viscosity on the sub-grid scale. The alternative would be the impractical task of trying to represent interactions between many different scales of motion.

The horizontal diffusion scheme for the UM was described in detail by Cullen et al. (1993). Briefly, it involves the repeated application of a conservative form of the Laplacian operator ∇^2. Two applications give second-order (∇^4) diffusion and three give third-order (∇^6) diffusion. The latter is more scale selective (i.e. gives more damping on the inaccurate short scales) but is slower.

Diffusion is applied on all levels, to the variables u, v, θ and q. A different order and coefficient may be chosen for each level. The diffusion of horizontal divergence is handled separately, again with the order and coefficient specified for each level.

The Unified Model includes a scheme for vertical diffusion, but in most standard versions this is not used. However, vertical mixing is normally handled by complicated boundary layer schemes, which are not included in the Jupiter version. It seems likely that sub-grid scale vertical processes are important in Jupiter's atmosphere, given the weak stratification in the lower troposphere. This means that a suitable scheme should be included. Presently the standard UM vertical diffusion scheme is used, but with a larger coefficient than intended for terrestrial use. This can be seen qualitiatively to reduce vertical noise. The coefficient has not been tuned precisely, however. The model flows we studied are not very sensitive to the exact value. Perhaps the best way of tuning it would be to use an equatorial wave experiment (see Chapter 6) in conjunction with future observations of wave amplitudes.

Choice of diffusion coefficients

A one-dimensional analysis is used to estimate horizontal diffusion coefficients. The diffusion equation is

$$\frac{\partial u}{\partial t} = K\nabla^2 u \tag{3.18}$$

It is straightforward to discretize this equation using, for example a forward-time centred-space scheme. A von Neumann stability analysis for the case of a plane wave of wavelength $\lambda = 2\Delta x$ shows that the Courant number ξ is given by

$$\xi = 1 - \Delta t \left(\frac{4K}{(\Delta x)^2}\right)^k \tag{3.19}$$

where k is the order of diffusion. This is simply the number of times the diffusion operator of Eq. (3.18) is applied. To achieve damping on a characteristic timescale $N\Delta t$, K should be chosen such that $\xi = e^{-1/N}$. Rearrangement of Eq. 3.19 and substitution yields

$$K = \frac{(\Delta x)^2}{4}\left(\frac{1 - e^{-1/N}}{\Delta t}\right)^{1/k} \tag{3.20}$$

which allows the calculation of K for any resolution and timestep. (All dimensional quantities in the above equation are measured in S.I. units). The only free parameter is N, whose determination is a matter of trial and error. The standard Met. Office configurations use 12 hour damping ($N = 24, \Delta t = 1800$s) for climate models; 1 hour damping for the global model; and 20 minutes damping for the UK limited area model. In general, more damping is required for higher resolution models. We wish to find the lightest damping which still prevents numerical instabilities from growing. If the diffusion is too heavy, the finer structures, such as narrow jets, are smeared out and lost. For the Jupiter models, the approach has been to choose small diffusion coefficients and establish where and when numerical instabilities appear. The diffusion is then gradually increased until the instabilities are removed.

Typical values for a half-degree model are $N = 48$ with third-order diffusion; this gives $K = 2.8 \times 10^9$ with a damping time of 12 hours. This value is applied to winds and temperature diffusion on all levels except the top two or three. Here first-order diffusion with a shorter damping time is used because the fields can be noisy. $u = v = 0$ on the topmost level, and below this the first-order diffusion ensures the fields are smooth.

Model runs have tended to show an instability in the jet at 26°S, with oscillations in the vertical velocity, decaying exponentially with height. This seems to be due to the forcing scheme on the lower boundary, although the precise cause has not been identified. The problem has been cured by the use of divergence damping to smooth this field. This parametrization acts in the same way as horizontal diffusion, but is applied to the horizontal divergence field instead of to the winds. It may be thought of as horizontal diffusion of the vertical velocity. This is useful for damping model noise, because numerical

instabilities often begin with an imbalance between u and v, leading to unrealistic values of the divergence. If this can be damped before it is manifest in the winds, the instability can be stopped more easily.

A series of experiments has established the following suitable values for the divergence damping coefficients: $K = 3 \times 10^7$ on the lowest two levels, $K = 1 \times 10^7$ above. These seem to give the minimum damping sufficient to remove the numerical instability, at least in the absence of other schemes such as vertical diffusion or convection.

3.3.1 Physical parametrizations

The physical processes included in any model will depend on the scope and aims of a particular experiment. One essential feature is that the model must be stable, so disturbances (including numerical noise) cannot grow without limit. Another is that the equations must nowhere be overconstrained; this requires appropriate boundary conditions. Although exact energy conservation is difficult to satisfy, any model should represent, at least crudely, the energy sources driving the flow (solar irradiation, internal heat flux) and the sinks dissipating it (friction, including that generated by the stress from the interior flow).

In the early stages of development, the initial aim has been to create a model which, when initialised from a plausible, geostrophically balanced initial state, evolves to show Jupiter-like behaviour, with a zonal-mean flow which is not too far removed from the initial state (and hence from observations). The flow should eventually have some stable time-averaged state, to justify claims that it is energetically balanced. It is desirable to do this with the simplest possible model scheme.

This has been achieved by using a Newtonian (i.e. linear) thermal relaxation scheme, similar to that in the dynamical core models described above. After each dynamics timestep, the temperature at each point is relaxed towards the basic-state temperature. There are a range of possibilities for the form of this basic state, the simplest one being a temperature distribution varying only with height. Such an approach is reasonable, because little is known about such important variables as the latitudinal variation of the internal heat flux, or the mixing ratios of absorbing gases and aerosols. In the light of Galileo data, which may constrain some of these unknowns, we may in future be able to include a more realistic radiation scheme in the model.

To investigate the stability of the observed jets and postulated vertical structure, we use a basic-state temperature distribution which is a function of height and latitude, set to be in geostrophic balance with the basic-state winds. This is used as the initial model state, to avoid the long integration time required to reach equilibrium starting from an arbitrary

state. A method for calculating a suitable basic state will be discussed in Chapter 4.

The thermal relaxation coefficient is taken to be constant throughout the atmosphere, although this restriction may easily be lifted. Its reciprocal is the radiative time constant for the atmosphere; Conrath et al. (1990) estimated this time constant to be on the order of the orbital period, varying by a factor of two or three over the relevant height range. This is much smaller than the variation for Uranus and Neptune, as estimated by the same model, and for our purposes may reasonably be taken to be constant with height.

3.3.2 Lower boundary forcing

At the bottom boundary, we assume that the sigma surface coincides with an isentrope, and accordingly set the temperatures there so θ remains fixed. As the lowest level has to be a sigma (not pressure) level, the surface is not isothermal, but rather T varies with the surface pressure p_*.

The other physical problem to be confronted is how best to represent the interaction with the deep flow upon which the model 'weather layer' rests. The basic assumption made is that the behaviour of the thin layer does not influence the flow in the deep atmosphere. We may therefore impose the winds at the base of the model, as a parametrization of the effects of the underlying layer. A straightforward method is to take suitable zonal winds and apply these at the model base. The approach presently used is simply to sample the zonally averaged cloud-top winds (Limaye, 1986) at model gridpoints. This is used to set up a model basic state, details of which are given in Section 4.4.

The model results were found to be sensitive to the implementation of this boundary condition. For example, it was not possible merely to impose suitable winds, because this can generate numerical instabilities. A more successful scheme was one in which the model winds at each point are forced towards the observed values. At each point on the lowest one or two levels (except those at the sidewalls) an increment was added to u and v, to remove a fixed fraction of the discrepancy between the model state and the reference state. Two levels may be forced in order to guide the model towards a state of zero vertical shear at a point just above the bottom boundary. This lower boundary condition is appropriate when the basic state is barotropic at depth, which was usually true in the present model experiments.

Some results are shown in Chapter 5: the model winds and temperatures adjust from their geostrophic initial states towards a truly balanced state, which remains stable at least for several hundred days. Waves were generated, followed by giant eddies, and the flow is superficially convincing.

Unfortunately, this forcing at the lower boundary, with a relatively short time constant (large coefficient), means that any wave or eddy motions were destroyed at the lower boundary. The eddy amplitudes are constrained to vanish there. This is not necessarily realistic, for there is no reason to suppose that motion in Jupiter's deep atmosphere is entirely zonally symmetric. A different forcing scheme, which adjusts only the zonal-mean term, was used to overcome this difficulty. In this version, the zonal mean wind on the lowest level was calculated at each timestep. Then the same increment is added at all points, so that the zonal means of (u, v) were equal to the basic-state winds. Generally $\bar{v}$ is chosen to be zero. This allowed eddies to extend down undisturbed to the base of the model. Results obtained with these two forcing schemes, and combinations of them, are discussed in Chapter 5. Generally we regard the zonal-mean forcing as being more realistic, since it permits the maintenance of deep eddies which the direct scheme artificially suppresses.

3.4 Energy diagnostics

Most physical systems may be analysed either by solving the governing differential equations, or alternatively by consideration of approximately conserved quantities, such as energy or angular momentum. The latter approach can often reveal the nature of processes within the system in a relatively simple way, an obvious example being atomic spectroscopy. In the case of a numerical model of an atmosphere, energy analysis is particularly useful for distinguishing between different forms of instability and we shall describe a simple technique here. This will be used extensively in Chapter 5, for analysis of GCM results.

A quasi-geostrophic analysis of the globally integrated energy was first developed by Lorenz in the 1950s. This Lorenz energy cycle is described in many recent books, e.g. James (1994). We shall apply the method to our case of a limited-area model of Jupiter. The procedure is conceptually simple, although algebraically somewhat complicated. As the method is standard we will only summarize it here, defining the quantities necessary for later use in data analysis, and identifying the difficulties peculiar to the Jupiter model. These include the limited applicability of QG theory, and problems associated with a limited area (as opposed to global) domain.

Energy terms

We begin by identifying four energy quantities. These are the zonal-mean and eddy parts of the kinetic and available potential energy respectively. Available potential energy (APE) is defined as the gravitational PE of the atmosphere which is available for con-

version to KE, without increasing the internal energy of the atmosphere. It is thus a measure of the degree to which the isentropes slope relative to the isobars, assuming (as always in the QG theory) that static stability varies only with pressure. Following James (1994) we refer to the APE terms as AZ and AE for the zonal and eddy parts respectively, and to the KE terms as KZ and KE.

The kinetic energy per unit mass of the atmosphere, neglecting vertical motions, is just $(u^2 + v^2)/2$. The total KE is the mass-weighted integral over (λ, ϕ, p). It is convenient to express the integrals per unit surface area of the planet. We can split the motions into a zonal-mean part $\bar{u}$ (assuming $\bar{v} = 0$, and an eddy part $(u^*, v^*) \equiv (u - \bar{u}, v)$, and integrate these similarly.

The APE terms are most easily expressed in pressure coordinates by first defining an Exner function

$$h(p) = \frac{R}{p}\left(\frac{p}{p_0}\right)^{\kappa} \tag{3.21}$$

and a reference-state static stability

$$s^2 = -h\frac{\partial \theta_R}{\partial p} \tag{3.22}$$

We then have

$$\mathrm{AZ} = \int \frac{h^2}{2s^2}\theta_A^2 \tag{3.23}$$

where θ_A is the anomaly potential temperature, after subtraction of the constant reference $\theta_R(p)$. The eddy part AE is obtained by replacing θ^* with θ_A.

Conversion terms

We can learn much simply by looking at correlations between timeseries of the energy terms. However, we can make further progress by considering the conversion rates between the four partitions of energy. The analysis proceeds in principle by taking the scalar product of the QG momentum and thermodynamic equations with the wind vector. This leads to two energy equations, containing terms identifiable as rates of change of the energy terms. Other terms may be identified as conversion rates. These are labelled as follows. CZ is the conversion of zonal-mean energy, from available to kinetic (or the reverse if negative). This is a "thermally direct" process, an example of which is the Hadley circulation. CE is the conversion of eddy energy in the same sense. CA is the conversion rate of APE, from zonal-mean to eddy, while CK is the kinetic equivalent but conventionally taken to be in the oppositie sense, i.e. positive for transfers from eddy to zonal-mean KE. The cycle is summarized in Fig. 3.1.

The definitions of the conversion terms will only be stated here, without derivation. For details see e.g. James (1994).

$$CZ = \int -h\bar{\omega}\overline{\theta_A} \tag{3.24}$$

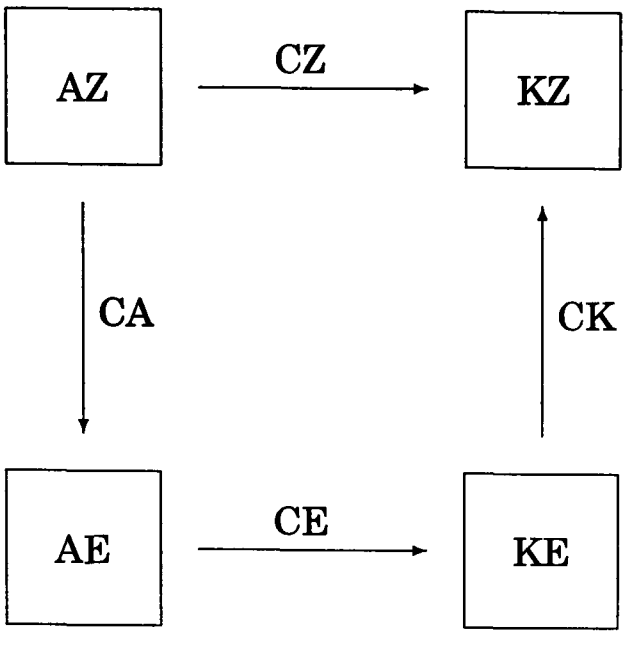


Figure 3.1: The Lorentz energy cycle, omitting the external forcing and damping terms.

$$CE = \int -h\overline{\omega^*\theta^*} \tag{3.25}$$

$$CA = \int -\frac{h^2}{s^2}\overline{\theta}_{Ay}\overline{v^*\theta^*} \tag{3.26}$$

$$CK = \int \overline{u}_y\overline{u^*v^*} \tag{3.27}$$

The importance of these conversion terms is that they are a useful indicator of the major processes at work in the atmosphere at any time. For example a large negative value of CK indicates barotropic instability, the conversion of zonal-mean KE into eddy KE. The reason the sign is negative is that the momentum flux is negatively correlated with the horizontal wind shear. Baroclinic instability is more complicated: it is a conversion from AZ (due to large horizontal temperature gradients) into KE. The conversion occurs via AE, with horizontal temperature fluxes measured by CA and then vertical fluxes measured by CE.

In fact the second part of this process is just a consequence of thermal wind balance, and model results confirm that AE is always strongly correlated with KE. Therefore we will often simply consider the total eddy energy, AE+KE. There are other drawbacks and limitations which will now be examined.

Limitations of the method

An important consideration is that we must be sure that boundary fluxes can be neglected, otherwise they would distort our results. This could be achieved by integrating

over a very large region (such as the whole planet) but this has the disadvantage that we then have no spatial resolution in the diagnostics. If several different processes are active in different regions, their contributions may cancel or lead to confusing results. In practice the latter problem is ameliorated by the fact that our model is a limited area channel. Even integrating over the entire domain, the area is so small that often only one major process is occurring at a given time (for example, a single giant eddy is breaking up).

However, modelling a channel does complicate the issue of boundary fluxes. At the eastern and western ends there is no difficulty, because the domain is periodic. Sidewalls are mathematically and physically less satisfactory than full global coverage. From a purely energetic point of view, the sidewalls do not seem to cause excessive inaccuracies. There must be zero advective flux through them, and we have some evidence to suggest that the thermal and frictional effects are not too significant. This comes from model runs where a large eddy is generated, which migrates poleward until it reaches the sidewall. On arriving at the edge of the channel, it can remain intact for a long time, advected along the edge by the flow.

This result suggests that drag and heat exchange through the walls have little effect on the energetics of the interior flow. Of course the effect on the dynamics can be very significant, for we would ideally like eddies to pass smoothly through the boundary and disappear, instead of sometimes rolling along next to the wall and sometimes disintegrating. Nevertheless we can hope that by considering the energy timeseries together with the corresponding flow patterns, we can identify the occasions where the sidewalls influence events.

This leaves the upper and lower boundaries of the model. Advective flux through these will be small (unless convection occurs at the lower boundary), but there are sometimes jets forced at the bottom, and drag at the top. Indeed we would never expect the fluxes to be zero on the real planet, for our definition of Jupiter's atmosphere, as distinct from interior, is an arbitrary division. These external fluxes (together with nonconservative terms, such as Newtonian relaxation throughout the domain) may be represented as extra quantities in the Lorenz cycle.

The physics routines are simple, and it would be straightforward to modify them, so as to record the cumulative heat exchange and kinetic energy change throughout the model run. There are two reasons why this has not been done. Firstly, it would be computationally demanding, and secondly, the forcing and damping are rather arbitrary. Neither the model nor the QG energy analysis are sufficiently sophisticated to make it worth while to calculate a detailed energy budget.

Without wishing to push the analysis too far and produce a (possibly spurious) detailed, closed budget, we can be confident that the qualitative results are reasonable, because

trends in the energy timeseries tend to be clearly defined, and it is always possible to identify large peaks or sudden changes, in terms of events observed in the model flow. Furthermore, the Unified Model has a built-in total energy diagnostic (not reliant on QG approximations), and tests have shown that this varies very little over long integrations. There seems no reason to suspect a large systematic error in the total energy of the model. If this is correct, then the QG energy quantities seem reasonable, with the sole exception of the zonal-mean available energy term.

On long timescales AZ tends to show an inexorable upward trend, which can dominate the other structure. Apart from this overall trend, AZ is correlated sensibly with the other quantities. It seems that this problem is exacerbated by the features of a giant planet atmosphere. AZ varies inversely with the square of the static stability s. In the lower part of the atmosphere, s is extremely small because the fluid is close to adiabatic. This, together with the mass-weighting, means that the integral of AZ is dominated by the potential temperature near the bottom of the model. Continuous mechanical forcing at the bottom will tend to increase the vertical windshear, and hence the horizontal temperature gradient. Although the winds are never far from equilibrium (as evinced by the kinetic energy timeseries, which have no such trend), there is a very slow increase in θ_A.

This apparently monotonic increase appears to be negligible in its impact on the model evolution. Long runs (more than 1000 days) show chaotic behaviour, with eddies repeatedly appearing and disappearing, but with no detectable trend in the primary model fields resembling that in AZ. The AZ diagnostic is extremely sensitive to the small changes in static stability, and we believe the trend to be real but to have little influence on the model behaviour. For runs with weak or no mechanical forcing, or for short integrations, the problem is not noticeable. Nevertheless, this is something to monitor in future and hopefully to resolve. It will be particularly relevant when testing more elaborate convection schemes, featuring moist convection or *ortho-para* hydrogen conversion.

Chapter 4

Vertical structures, normal modes and model initialization

4.1 Introduction

Chapter 3 has described an existing general circulation model, summarized the main observations of relevance, and discussed the implementation of a Jupiter-like version of the model. This chapter will consider the vertical structure of the basic-state model flow in more detail. We shall examine several structures consistent with observations, and show how quasi-geostrophic theory enables us to set up a dynamically consistent initial state for the model.

The available observations (see Section 1.2) are so limited as to allow a wide range of possibilities for the vertical structure. Additional criteria must therefore be applied in selecting structures. The approach taken here is to use the linear theory of small-amplitude waves, to calculate the normal modes of the vertical structure. A model may then be initialized so the fields have a known modal form, and their subsequent evolution may be analysed in terms of the normal modes. This may help to reveal the stability characteristics of different modes.

The linearized primitive equations were treated in this way by Laplace, who separated the vertical structure from the horizontal and time dependence. This leads to Laplace's tidal equations, which describe horizontal and temporal variations (Andrews et al., 1987, Chapter 4), and to a vertical structure equation. The tidal equations are identical to those governing a shallow fluid layer, whose depth and Rossby deformation radius are related to the separation constant. Each mode may therefore be viewed as behaving like a wave in a shallow-water system with a certain Rossby radius. This analogy is helpful because, broadly speaking, two isolated vortices will tend to interact, or influence each other, only

if they are separated by a distance less than the relevant Rossby radius.

More recently, Flierl (1978) applied the same technique to the linearized, flat-bottom form of the quasigeostrophic equations for a two-layer ocean. Flierl used geometric depth as a vertical coordinate, and separated the horizontal and vertical dependence of the dynamic pressure. Having calculated the vertical normal modes, Flierl went on to recast the two-layer equations in terms of a two-mode formulation. This had the same number of degrees of freedom as the two-layer models, but several additional advantages. In a layer model of a continuous system, the calibration parameters (such as layer depth and density difference) are not uniquely defined, Any particular choice of parameters is usually not optimal when more than one physical effect is being modelled—for example surface forcing combined with bottom topography. The normal-mode model, however, can be calibrated so the same parameters are optimal for several different cases. Flierl showed that a two-mode formulation can be optimal for a model with three effects: surface forcing, topography and friction. The equivalent two-layer models tended to capture only some of the physics, and Flierl concluded that at low vertical resolution, spectral methods may be superior to layer methods. This is a familiar result in the context of horizontal resolution (a tide or Rossby wave may be represented using only one harmonic, whereas one gridpoint value is of little use).

Achterberg and Ingersoll (1989) performed a similar procedure, but used a log-pressure vertical coordinate, solving for the quasigeostrophic streamfunction in a multilayer, stratified compressible atmosphere. The present model uses the same method to find the vertical modes. However, the normal-mode formulation is not used when solving the model equations, but merely to guide the choice of a basic-state profile and initial fields. It may also be used in the analysis of results, by projecting the model flow onto the normal modes.

A small disturbance in the model flow may be projected onto the calculated vertical modes, and its horizontal structure separated from this. If its vertical structure is purely that of a single mode, then we may expect its subsequent evolution not to change this, as the modes are mutually orthogonal. If the flow consists of a superposition of modes, then in general, there will be some nonlinear coupling between the modes, and the vertical structure may change over time. It is reasonable to expect disturbances to have approximately the same vertical structure as the background flow, if they are assumed to arise from an instability in that flow. It is therefore convenient to choose a background state which, while consistent with observations, has a simple composition of just one or two eigenmodes. Of course we cannot expect the primitive equation model to behave exactly as the linearized QG analysis predicts. However the nonlinearity may be insufficient to render the analysis worthless.

One reason for solving the more restrictive QG system (rather than the linearized primitive equations) is that a knowledge of the streamfunction allows us to initialize all the

model fields in a geostrophically balanced state. The streamfunction may be determined by combining the observed winds with a chosen vertical shear, for example that of the first baroclinic eigenmode. This method of initializing the model fields was used by Williams (1996). Another reason is the linearity of the PV operator in QG theory.

Section 4.2 describes the derivation of the vertical structure equation and the calculation of its eigenfunctions. Analytic solutions for the case of uniform stratification are then discussed, and compared with numerical results, as a test of the method. Then in section 4.3, we turn to the problem of finding the modes of a realistic atmosphere, and present some results. In Section 4.4 we outline a method of choosing a suitable reference state. This is done by constructing superpositions of a few modes, while satisfying rigid-lid (and sometimes zero-flow) boundary conditions. We then calculate the basic-state streamfunction and show the results. In the final section we consider an alternative technique, the so-called "shooting method", which is more flexible because it can calculate normal modes in an atmosphere with a sponge layer or other drag term. Although less elegant, this is closer to representing the GCM profile. The results are slightly different, as we would expect.

4.2 Calculation of the vertical eigenfunctions

The linearized quasi-geostrophic potential vorticity equation

The potential vorticity equation may be written using any of several different vertical coordinate systems (pressure, log pressure, geometric height, etc.) The most mathematically straightforward vertical coordinate for our present purpose would be pressure. Unfortunately there is a conflicting requirement, which is that for computational reasons it is desirable to use levels uniformly spaced in whichever vertical coordinate we choose. The reason for this is that the problem then simplifies to the inversion of a symmetric tridiagonal matrix.

Unfortunately, the second criterion rules out the use of pressure as a coordinate. The vertical domain spans up to four orders of magnitude in pressure. To achieve useful resolution in the stratosphere would require a large number of levels, if those levels were uniformly spaced in p throughout the domain. We therefore write the equations using a log-pressure coordinate z, and accept the slight increase in the number of terms. This choice has the advantage that the GCM vertical levels are also regularly spaced in z, so we can use exactly the same grid and avoid the need to interpolate our results.

In log-pressure vertical and locally Cartesian horizontal coordinates, the full quasi-geostrophic

potential vorticity equation is written (see, e.g., Andrews et al., 1987):

$$\left(\frac{\partial}{\partial t} + u\frac{\partial}{\partial x} + v\frac{\partial}{\partial y}\right)\left[f_0 + \beta y + \nabla^2\psi + \frac{1}{\rho_0}\frac{\partial}{\partial z}\left(\rho_0\varepsilon\frac{\partial\psi}{\partial z}\right)\right] = 0 \tag{4.1}$$

Here the first bracket is the substantive derivative following the geostrophic wind, and the second bracket is an expression for the potential vorticity q. The right-hand side is zero in the absence of forcing and dissipation. The notation is standard: u and v are geostrophic velocities, $f_o + \beta y = f$ is the β-plane Coriolis parameter, z a log-pressure pseudo-height, and ρ_0 is the basic-state density. The QG streamfunction ψ is given by $u = -\partial\psi/\partial y$ and $v = \partial\psi/\partial x$, and ∇^2 is the horizontal Laplacian operator. ε and ρ_0 are defined as follows:

$$\varepsilon = \frac{f_o^2}{N^2} \;; \qquad\qquad \rho_0 = \rho_s e^{-z/H} \tag{4.2}$$

where N is the Brunt-Väisäla buoyancy frequency, H the pressure scale height and ρ_s the basic-state density at the base of the model.

Following Pedlosky (1979), we may linearize the PV equation about a rest state, restricting ourselves to small-amplitude Rossby waves and neglecting products of small quantities. This procedure reduces (4.1) to

$$\frac{\partial}{\partial t}\left[\nabla^2\psi + \frac{1}{\rho_0}\frac{\partial}{\partial z}\left(\rho_0\varepsilon\frac{\partial\psi}{\partial z}\right)\right] + \beta\frac{\partial\psi}{\partial x} = 0 \tag{4.3}$$

The vertical structure equation

We can now derive the vertical structure equation, by separating the variables in Eq. (4.3). We shall write ψ as the product of a horizontal, time-dependent part, a vertical part and an exponential height factor. It turns out to be convenient to write this exponential in terms of the basic-state density $\rho_0(z)$, leading to the following separation of ψ:

$$\psi \equiv \rho_0^{-1/2}\alpha(x,y,t)F(z) \tag{4.4}$$

Substituting (4.4) into (4.3), and dividing through by ψ, we obtain a separated equation whose terms in α are independent of those in F.

$$\frac{\partial}{\partial t}\left(\frac{\nabla^2\alpha}{\alpha} + \frac{1}{\rho_0^{1/2}F}\frac{\partial}{\partial z}\left[\rho_0\varepsilon\frac{\partial}{\partial z}\left(\rho_0^{-1/2}F\right)\right]\right) = 0 \tag{4.5}$$

Equation (4.5) holds for all (x,y,z,t). Therefore, for non-trivial solutions, the terms governing the horizontal and vertical parts of the streamfunction must be separately constant. The equation for F is

$$\frac{\partial}{\partial t}\left(\frac{1}{\rho_0^{1/2}F}\frac{\partial}{\partial z}\left[\rho_0\varepsilon\frac{\partial}{\partial z}\left(\rho_0^{-1/2}F\right)\right]\right) = \text{constant} \tag{4.6}$$

The separation constant must be zero, because $F = F(z)$ only. Equation (4.6) may trivially be integrated with respect to time, and written in Sturm-Liouville form:

$$\frac{\partial}{\partial z}\left[\rho_0\varepsilon\frac{\partial}{\partial z}\left(\rho_0^{-1/2}F\right)\right] + \lambda\rho_0^{1/2}F = 0 \tag{4.7}$$

where λ is a constant to be determined. Solutions will only exist for particular eigenvalues λ; it has a further significance, in an analogy between the horizontal part of Eq. (4.5) and the single layer shallow-water equations. The analogy enables us to define a Rossby deformation radius and "equivalent depth" for each mode.

Expanding the brackets and dividing Eq. (4.7) by $\rho_0^{1/2}$ yields the vertical structure equation:

$$\varepsilon\frac{\partial^2 F}{\partial z^2} + \frac{\partial \varepsilon}{\partial z}\frac{\partial F}{\partial z} + \left(\frac{1}{2H}\frac{\partial \varepsilon}{\partial z} - \frac{\varepsilon}{4H^2}\right) F + \lambda F = 0 \tag{4.8}$$

The advantage of including the exponential factor $\rho_0^{-1/2}$ is now clear: we have removed the ρ dependence, reducing a generalized eigenvalue problem to a straightforward one. We have an expression for the normal modes of the vertical structure in terms of only the essential parameters ε and H. A comparison with the analysis of Andrews et al. (Chapter 4.2) shows that this is a slightly more complicated form of another vertical structure equation, derived from the equations for linearized disturbances to a primitive equation system.

Boundary conditions

The boundary conditions on the vertical structure equation should be the same as those used in the general circulation model, i.e. $\dot{\eta} = 0$ at the top and bottom half-levels, $z = 0$ and $z = Z$. Recall that η is the hybrid pressure-sigma coordinate used by the Unified Model. This rigid-lid boundary condition implies an infinitely large stratification and hence infinite buoyancy frequency N. By Eq. (4.2) this means that $\varepsilon = 0$ on these half-levels. Other boundary conditions are possible, for example $w = 0$, but will not be considered here. Later, in Section 4.5 we shall introduce a more complicated boundary condition, that of a sponge layer at the top of the model.

Simple analytic solutions

The vertical structure equation (4.8) may be solved analytically for simple cases. Consider an atmosphere which is uniformly stratified except for the rigid lids at $z = 0, Z$. This means that

$$\varepsilon = \begin{cases} \varepsilon_0 & : \quad 0 < z < Z \\ 0 & : \quad z = 0, Z \end{cases} \tag{4.9}$$

where ε_0 is a constant.

To find an analytic solution to (4.8), we need a boundary condition in terms of F rather than ε. By the definition (4.9), ε is discontinuous at $z = (0, Z)$, so as $z \to 0$, $\partial\varepsilon/\partial z \to \infty$.

Consider the behaviour of the vertical structure equation (4.8) near a boundary. In the limit as $z \to 0$, the terms in $\partial\varepsilon/\partial z$ dominate, and Eq. (4.8) becomes

$$\frac{\partial F}{\partial z}\frac{\partial \varepsilon}{\partial z} + \frac{F}{2H}\frac{\partial \varepsilon}{\partial z} = 0 \tag{4.10}$$

The same applies at the other boundary, where $z \to Z$. Cancelling the $\partial\varepsilon/\partial z$ (which is nonzero at the lid) gives the boundary condition on F:

$$\frac{\partial F}{\partial z} + \frac{F}{2H} = 0 \quad \text{at} \quad z = (0, Z) \tag{4.11}$$

This is of course equivalent to the condition that the vertical gradient of the streamfunction ψ be zero, which is what we require for a rigid lid. Now Eq. (4.10) is valid not only on the boundaries, but also in the interior. We can see this from Eq. (4.9), because if ϵ is constant everywhere in the interior, clearly $\partial\varepsilon/\partial z = 0$. We may therefore substitute (4.10) into (4.8) to obtain a differential equation with constant coefficients:

$$\frac{\partial^2 F}{\partial z^2} - \left(\frac{1}{4H^2} - \frac{\lambda}{\varepsilon_0}\right) F = 0 \tag{4.12}$$

The general solution of Eq. (4.12) is

$$F = A\exp(\phi z) + B\exp(-\phi z) \tag{4.13}$$

where

$$\phi^2 \equiv \frac{1}{4H^2} - \frac{\lambda}{\varepsilon} \tag{4.14}$$

Solutions with $\phi^2 > 0$ are real exponential functions; those with $\phi^2 < 0$ are sinusoidal. The simplest case is that with eigenvalue $\lambda = 0$. ϕ is real, and substitution of the solution (4.13) into the boundary conditions (4.11) shows that F is an exponential, decaying with height:

$$F = F_0 \exp(-z/2H) \tag{4.15}$$

Substitution of F into Eq. (4.4) shows that the streamfunction is

$$\psi = \rho_s \exp(z/2H)\, \alpha F_0 \exp(-z/2H) \tag{4.16}$$

which is constant with height. An inspection of Eqs. (4.7) and (4.4) confirms that for $\lambda = 0$, the vertical structure equation has a trivial solution with $\partial\psi/\partial z = 0$ everywhere; this is the so-called barotropic eigenfunction.

Baroclinic solutions with nonzero λ and imaginary ϕ have the general form

$$F = A\sin\left(\phi\frac{z}{Z}\right) + B\cos\left(\phi\frac{z}{Z}\right) \tag{4.17}$$

Substitution into the boundary conditions (4.11) gives the relations

$$B = \frac{-2\phi AH}{Z} \quad ; \qquad \phi = 0, n\pi \tag{4.18}$$

where n is an integer wavenumber. Equations (4.14) and (4.18) allow the calculation of the eigenvalue spectrum, and the full analytic solutions, which are

$$F = \sin\left(\frac{n\pi z}{Z}\right) - \frac{2n\pi H}{Z}\cos\left(\frac{n\pi z}{Z}\right) \tag{4.19}$$

Numerical solution of the vertical structure equation

For more general profiles of $\varepsilon(z)$ Eq. (4.8) can be solved numerically using a finite-difference formulation. The equation is discretized on a staggered grid to minimize roundoff error. F is held on the model full levels p^j, and ε on intermediate half-levels $p^{j+1/2}$. The levels are chosen so that the spacing $h = \Delta(\ln p)$ between adjacent full-levels is constant. There are J full-levels (typically 30) and $J + 1$ half-levels, the rigid lids being placed on the exteme half-levels. After some rearrangement, we obtain

$$F^{j-1}\frac{\varepsilon^{j-1/2}}{h^2} + F^j\left[\frac{e^{j+1/2} - e^{j-1/2}}{2hH} - \frac{\left(\varepsilon^{j+1/2} + \varepsilon^{j-1/2}\right)}{2}\left(\frac{2}{h^2} + \frac{1}{4H^2}\right) + \lambda\right] + F^{j+1}\frac{\varepsilon^{j+1/2}}{h^2} = 0 \tag{4.20}$$

This equation may be rewritten as the tridiagonal matrix equation

$$\mathbf{AF} + \lambda\mathbf{F} = \mathbf{0} \tag{4.21}$$

where λ is now a vector of eigenvalues and $\mathbf{F}$ a matrix of corresponding eigenvectors or normal modes. The general entry $\mathbf{F}^{ij}$ represents the magnitude of mode i at vertical level j. Comparing the coefficients of F^{j-1} and F^{j+1} in Eq. (4.20), we can see that the matrix $\mathbf{A}$ is symmetric. We can now also see that this would not be the case with nonuniform level spacing; hence the choice (discussed above) of $z = \ln p$ as a coordinate.

There are several efficient methods for diagonalizing a symmetric tridiagonal matrix. The one used here is the tridiagonal QL implicit (TQLI) method, implemented by Press et al. (1986). This is a robust, efficient and accurate algorithm. The only caveat is that, where matrix elements vary over many orders of magnitude, the largest elements should be in the lower right-hand corner of the matrix. This is because the matrix is reduced to bidiagonal form starting from the lower right-hand corner, and if there are very small elements in this area, large roundoff errors will result.

For the realistic Jovian N^2 profile used in this model, the elements of $\mathbf{A}$ vary by a factor of a hundred or so, with the largest values towards the centre of the matrix. The method has also been tested with much steeper vertical structures, in which the deep atmosphere is extremely close to a freely convecting, adiabatic fluid. These profiles generate matrices whose elements vary more widely, and whose smallest values appear in the lower right-hand corner.

The practical solution to this problem is to permute the matrix so that the diagonal elements are ordered by magnitude. This is done by the application of a sequence of similarity transforms, which leaves the eigenvalues unchanged. After diagonalization, the reverse sequence of transforms is applied to recover the matrix of eigenvectors. The necessary permutations are an extension of Gaussian elimination. If we interchange rows r and r', we must also interchange columns r and r'. Alternatively, if we add n times row r

to row r', we must *subtract* n times column r' from column r. Each pair of permutations is a similarity transform.

Following a sequence of interchanges, the matrix is no longer tridiagonal but is still symmetric, so may be rendered tridiagonal by applying a series of Householder reductions. Alternatively it may be diagonalized by a method such as that of Jacobi (Press et al., 1986). The extra work required to diagonalize a non-tridiagonal matrix is not a significant difficulty, because the matrix is of a reasonable size (30×30) and the process need only be done once.

Experiments have shown that, for steeper vertical profiles, the matrices are rather ill-conditioned. It is difficult to resolve the eigenspectrum cleanly. The barotropic eigenvalue should be identically zero. In practice, it is usually only slightly smaller than the other eigenvalues, which is unsatisfactory. If this proves to be a limitation in the future, a method of iterative improvement may be developed. This is not a trivial task, so has been left until it is found to be absolutely necessary. The alternative "shooting" method (see Section 4.5 below) provides an independent check.

Trial solutions: uniform stratification

The numerical method has been tested by solving the simple case where $\varepsilon =$ constant. The numerical solutions for F may be compared with the analytic solutions discussed above. The analytic forms of the first six eigenmodes are plotted in Fig. 4.1; the corresponding numerical solutions in Fig. 4.2. The numerical code outputs orthonormal eigenvectors, in the sense

$$\mathbf{F_i} \cdot \mathbf{F_j} \equiv \sum_{j=1}^{J} F^{ij} F^{i'j} = \begin{cases} 1 & : \quad i = i' \\ 0 & : \quad i \neq i' \end{cases} \tag{4.22}$$

and the analytic solutions have been normalized similarly. The numerical solutions have been sorted into order of increasing eigenvalue, the analytic solutions into order of increasing wavenumber n. The results are in very good agreement, except for the signs of the fifth and sixth modes. The reason for this is that Eq. (4.19) is antisymmetric between positive and negative wavenumbers: $F(n) = -F(-n)$. It so happens that the numerical diagonalization procedure has picked out the modes $n = -4, -5$ instead of modes $n = 4, 5$. This is not significant for our analysis of the realistic system. In any application (for example, forming a superposition to match boundary conditons) each mode is multiplied by a suitable coefficient of arbitrary sign.

Figure 4.3 shows the solutions scaled by $\rho_0^{-1/2}$; this is the complete vertical part of the streamfunction ψ. Note that the barotropic mode is constant with height, which means the winds are also constant and we are justified in identifying this as the barotropic mode.

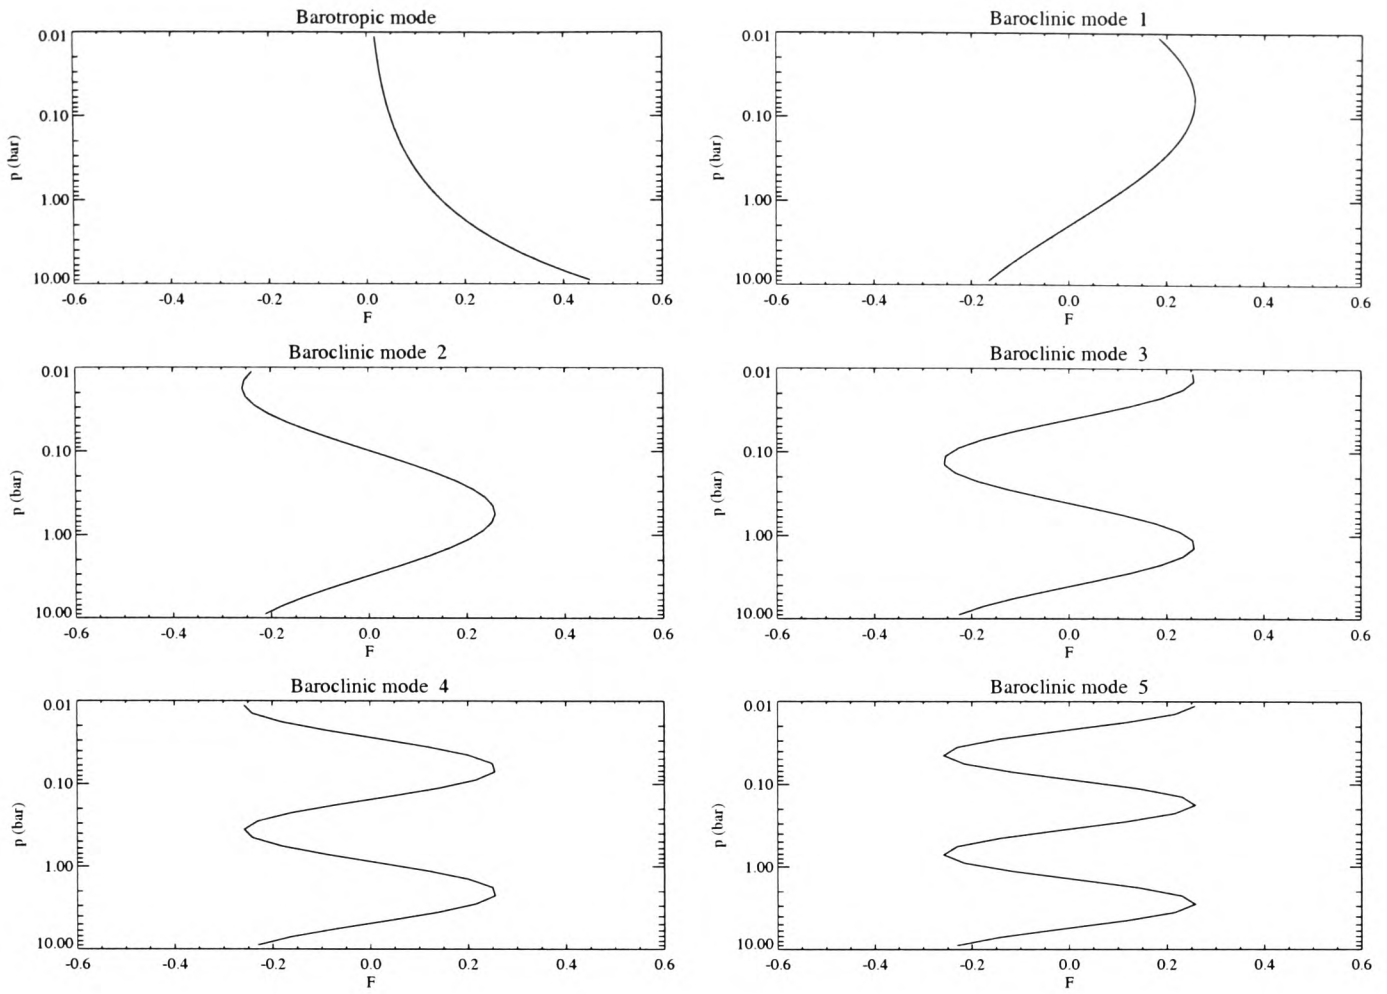


Figure 4.1: Analytic solutions of the vertical structure equation (4.8) with constant stratification

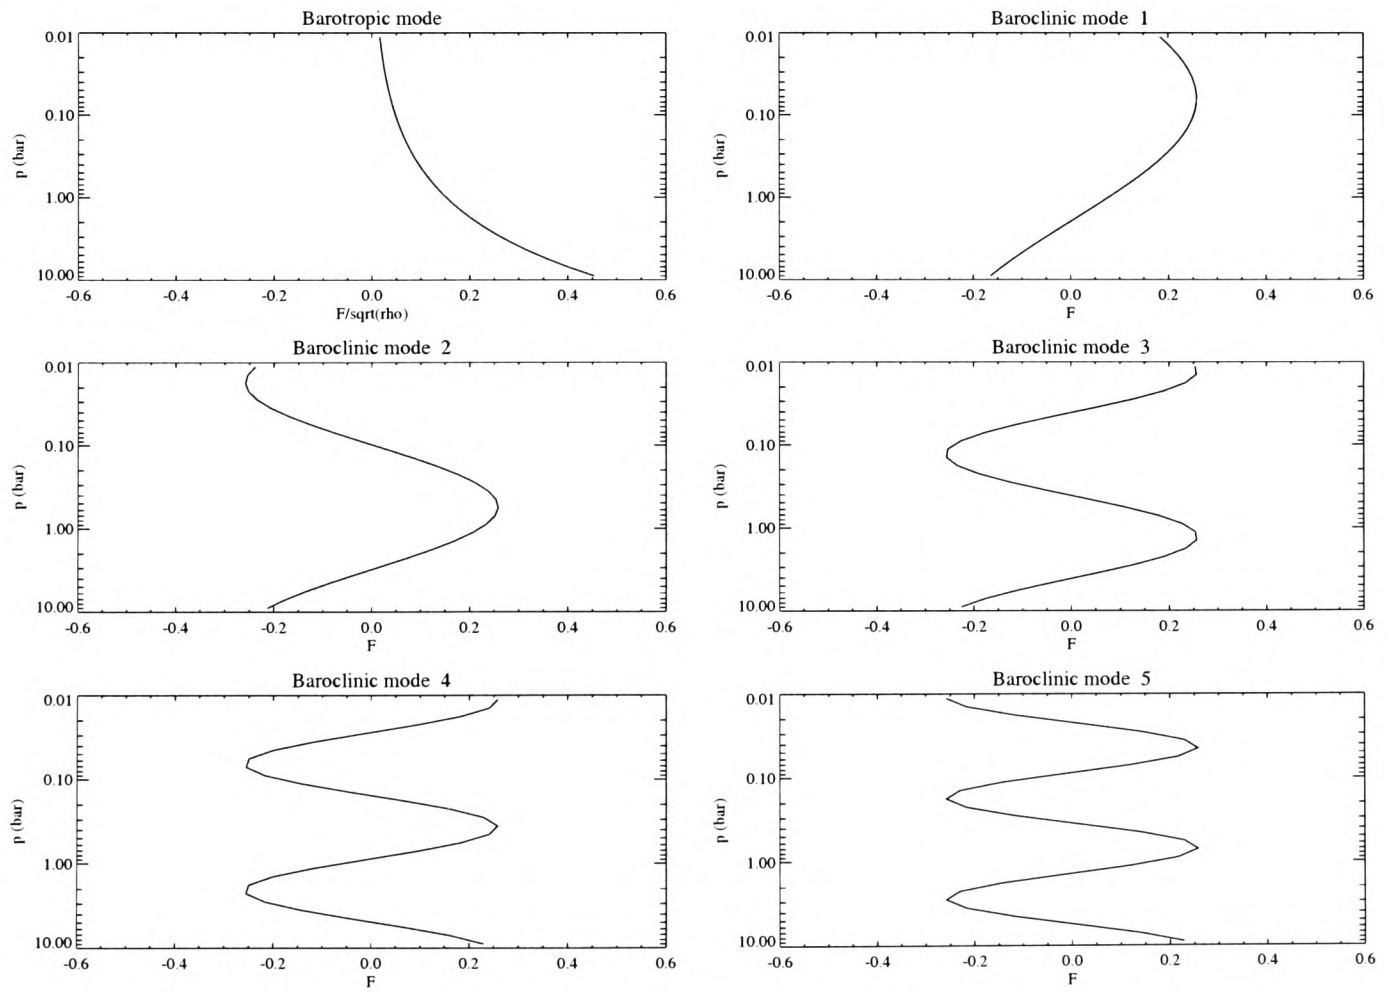


Figure 4.2: Numerical solutions of Eq. (4.8) with constant stratification

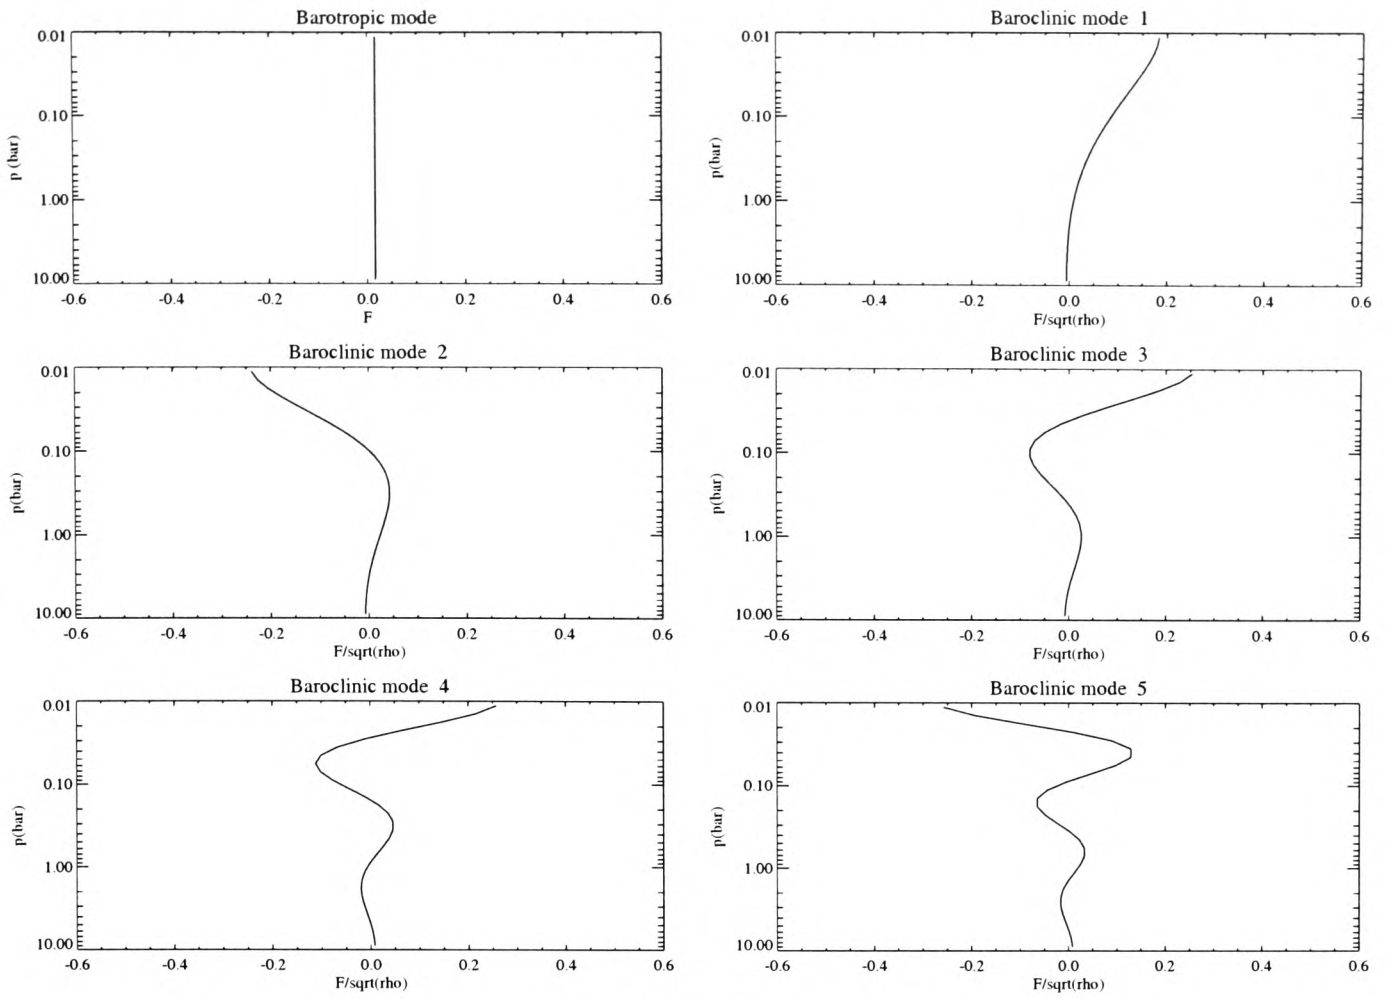


Figure 4.3: Solutions of Eq. (4.8), scaled by the density factor $\rho_0^{-1/2}$, for constant stratification

4.3 Vertical modes of a realistic atmosphere

In Section 4.2, we have developed a method for numerically calculating eigenfunctions (normal modes) of the vertical part of the linearized quasi-geostrophic PV equation. We have verified the method, by tests using the simple case of uniform stratification. Analytic solutions exist for this case, and have a few interesting properties, such as antisymmetry between wavenumbers n and $-n$.

We now turn to the problem of calculating the modes of a realistic atmosphere. Temperature and static stability profiles based on those of Achterberg and Ingersoll (1989) are used. These data were taken from Voyager radio occultation measurements (Lindal et al., 1981), with an almost-adiabatic extrapolation below the lower limit of the Voyager data. The profile was chosen to be very slightly stable. Comparison with the Galileo Probe measurements has shown that our model basic state lapse rate lies somewhere between the original Probe data (Seiff et al., 1996) and a dry adiabat. This model atmosphere is consistent with the standard view, based upon the available observations of $T(z)$ at the time this work was done. However, more recent evidence, from reanalysis of the Probe data (Young, 1998; Seiff, 1998, and other papers) and from other sources, increasingly suggests the existence of a statically stable layer below the visible clouds.

In Section 1.2 we mentioned the region of slightly positive static-stability, measured by the Galileo Probe between 6 and 12 bars. Below this was a zone of increasing stability,

from 15 to 22 bars. Although there may have been a considerable systematic error due to the very high internal temperatures within the Probe, this result is supported by other observations. For example, the Doppler Wind Experiment (Atkinson et al., 1996) reported strong vertical windshear down to about 4 bars. Theoretical dynamics implies that this would require positive static stability to support the shear, by preventing the growth of dynamical instabilities.

The Galileo Probe Mass Spectrometer (GPMS) revealed (Niemann et al., 1998) that both hydrogen sulphide and water were depleted at levels above their expected condensation levels, but became gradually more abundant with depth. This is consistent with the nephelometer's observation of the lack of thick clouds either of ammonium hydrosulphide or water. The steady increase in mixing ratio at greater depths tends to support the hypothesis of a slightly stable layer, at least within the hot-spot area. Conversely, a neutrally stable atmosphere should be well mixed by convection.

Allison (1990) proposed that the observed pattern of NEB plumes could be associated with a wavenumber-12 Rossby wave, trapped in a leaky waveguide formed by a deep stable layer. Such a stable layer could be formed by the stabilizing effect of moist convection within a thick water cloud. Ingersoll et al. (1994) analysed the gravity waves generated by the impact of comet Shoemaker-Levy 9 on Jupiter. They used the waveguide effect of a water cloud, to estimate the static stability beneath the region amenable to remote observations. The results support the existence of a water-cloud and its associated stable layer.

This range of evidence suggests that we should at least seriously consider the effects of a possible stable layer at around the 4-bar level. However, the work in this thesis predates the analysis of much of this data and is therefore restricted to the simpler case of a nearly-adiabatic deep atmosphere. We propose to investigate a stable layer, in the early stages of the planned future work. This will require the model domain to be extended, moving the lower boundary to a greater depth. For the present we will concentrate on the profile used for our experiments to date.

Achterberg and Ingersoll chose a moist adiabat for the region ($3 \leq p \leq 0.69$) bars, with a dry adiabat below. They went on to point out that Jupiter's large-scale motions may perhaps be governed by dry adiabatic processes, by analogy with the Earth's tropics (Gill, 1982). The evidence so far is inconclusive. In the present work, we have chosen a dry adiabat instead of a moist one, firstly since it is the simplest theory consistent with the facts, and secondly since the GCM has a dry convective adjustment scheme but no cloud parametrization. We therefore consistently neglect the effects of moist processes. It is hoped that the next stage of model development will be to include realistic thermodynamics in the GCM, including a stable layer generated by moist convection.

The bottom of the weather layer in Achterberg and Ingersoll's model is defined to be the

base of the water cloud, taken to be at around 4 bar. Our model extends deeper, so below the water cloud we have chosen a profile very close to a dry adiabat. The water cloud appears as a thick spike in the profile of N. Model experiments (not shown) have shown that the presence or absence of this spike has no visible effect on the normal modes, because the $\rho_0^{-1/2}$ factor in the streamfunction is very small in the deep troposphere. This is something of a deficiency in our treatment, because one would expect the cloud-top winds to depend more on the static stability just below the clouds than on that in the stratosphere. However, for consistency, we have chosen to neglect the water-cloud spike for the rest of this modelling work.

Given the observational uncertainty, there is scope for experimentation in adjusting the profile below the $\simeq$ 500 mb cloud-top level. By varying the lapse rate between a dry adiabat and a value closer to the Galileo Probe measurements (Seiff et al., 1996), it is possible to change the stability parameter σ by a factor of 10^4. This is analogous to the variation of the 'exponential structure rate' of Williams (1996). The model used here (with no changes) is comparable to the more stable of Williams' profiles, corresponding to his structure rate $N \sim 24$.

The profiles of absolute and potential temperature, buoyancy frequency, and stratification parameter ε are shown in Fig. 4.4. This model atmosphere, with a nearly-adiabatic deep atmosphere, is the one used for all of the GCM experiments described n this thesis. The alternative possibility (a deep stable layer) has been explored, to a limited extent, using both the one-dimensional normal mode model and the GCM. However, all the normal modes have very small amplitudes at depth, so are insensitive to changes to the stratification there. This is not satisfactory, but is difficult to overcome. It arises from the exponential pressure factor in the vertical structure equation, and suggests that a deeper model domain is required. While this should be achieveable in principle, implementation problems have caused us to shelve this development for the present. In its current form, our model can predict the deep winds, but cannot answer questions about the effect of a deep stable layer.

The first six eigenmodes (scaled by the density factor) of the vertical structure equation for this model atmosphere are presented in Fig. 4.5. The exponential density factor has the effect that none of the streamfunctions has large amplitude below about the 1 bar level, at least for this profile of ε. The first baroclinic mode is small in the troposphere, but begins to grow near the tropopause ($p \simeq 900\text{mb}$). It then increases monotonically with height.

The second baroclinic mode is of interest because in the stratosphere it has a dipolar form. There is a localised peak just above the tropopause, and above this a monotonic decrease with height. Modes of vertical wavenumber higher than about three oscillate rapidly with height. This suggests they are unlikely to be applicable to any large-scale motions which we can resolve in observations.

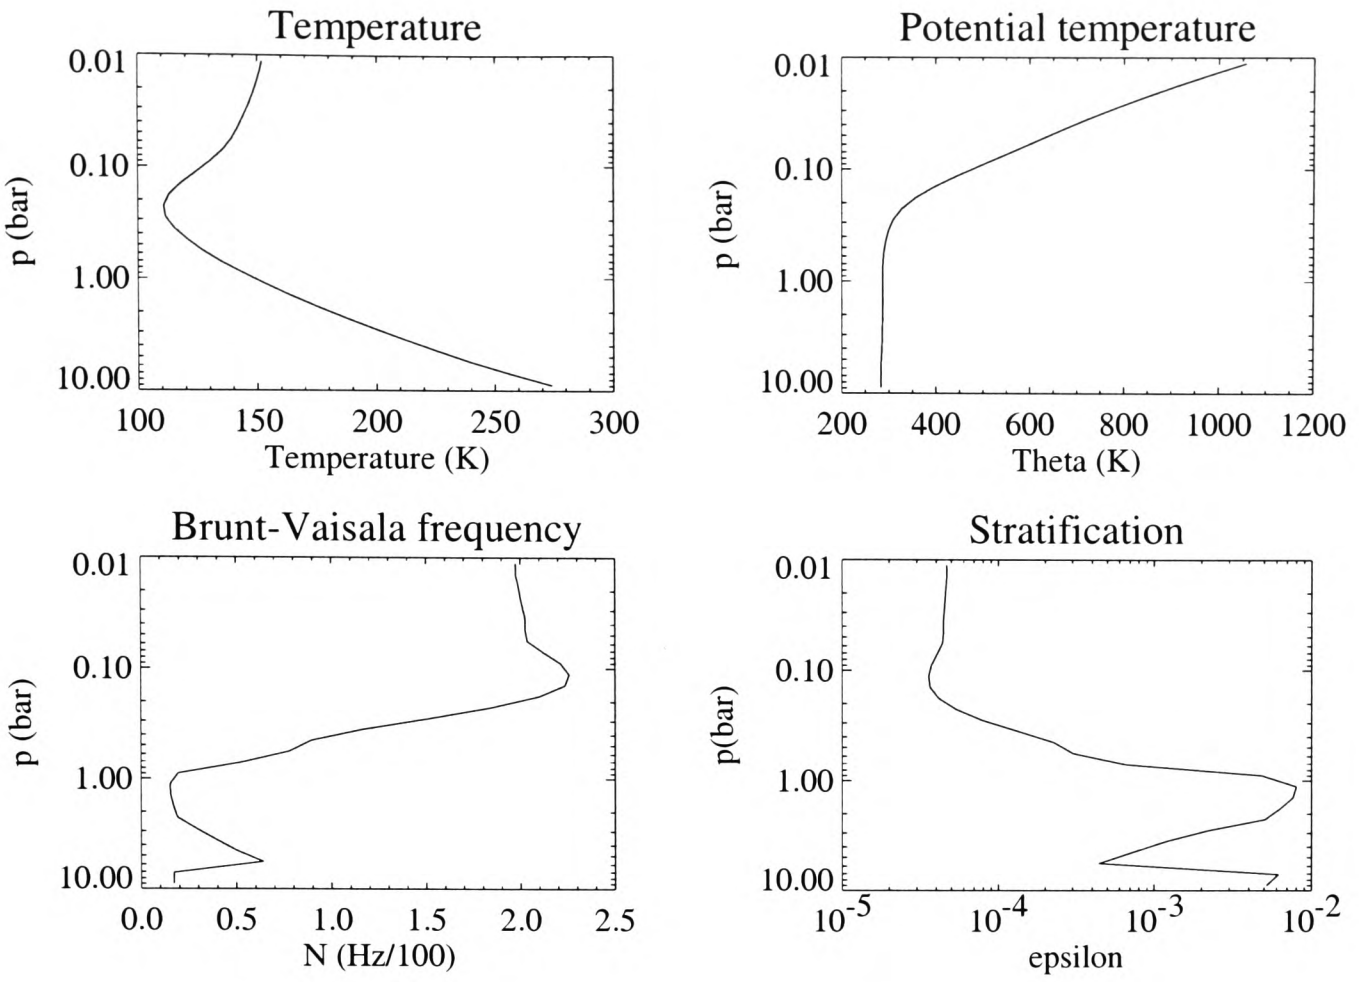


Figure 4.4: Profiles of T, θ, N and ε, for a model atmosphere based on that used by Achterberg and Ingersoll (1989)

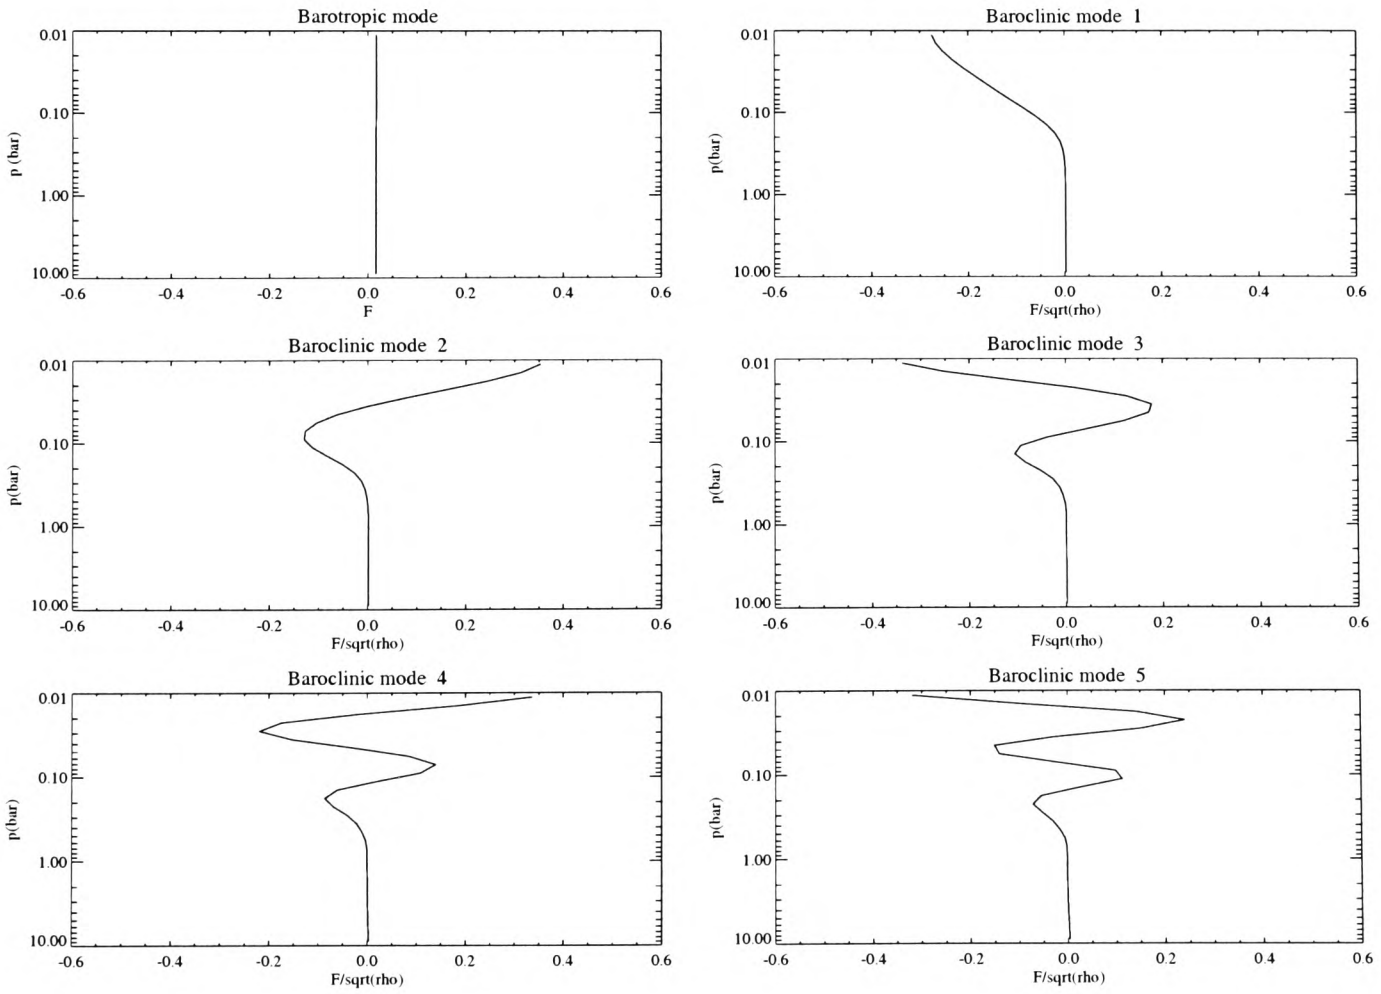


Figure 4.5: Solutions of Eq. (4.8), scaled by the density factor $\rho_0^{-1/2}$, for the realistically stratified atmosphere shown in Fig. 4.4

4.4 Superposition of the normal modes

It is possible to construct several interesting, plausible but idealised structures using simple superpositions of the first few eigenmodes of Fig. 4.5. For example, the zonal-mean basic state described in Section 3.3, which has a zero flow at the top of the model, and zero vertical shear at the base, may be obtained by superposing the barotropic and first baroclinic modes, with suitable coefficients. (This is equivalent to adding a constant to the baroclinic mode, such that the amplitude vanishes at 10 mb). The resulting streamfunction implies a vertical wind shear which is roughly uniform with height above 200 mb. This agrees well with observations, e.g. Voyager IRIS thermal winds above the Great Red Spot (Flasar et al., 1981). We thus have a plausible vertical structure for motions in Jupiter's atmosphere, which comprises only two normal modes. This result is quite robust: the shooting method (below) gives similar possibilities.

An alternative possibility is a structure function whose amplitude vanishes at both the top and bottom boundaries. This would be particularly relevant to a model of deeper extent, if one subscribes to the hypothesis that the observed motions are shallow relative to the planet's radius. In this scenario, there must be a level at which the alternating jet structure gives way to solid-body rotation or similar motions. If this is much deeper than the base of the model (as is likely if that base is placed at 10 or even 100 bars), then the two-mode basic state discussed previously may be appropriate. If not, this alternative structure may be more suitable.

Calculation of the second structure is not so trivial as the first. A general method for finding the superposition is the following. Let F_i be the mode of wavenumber i, and a_i the coefficient of this mode in the (normalized) superposition vector v. The two boundary conditions and a normalization condition form a set of three nonlinear simultaneous equations:

$$\sum_{i=1}^{I} a_i F_i(z=0) = 0 \tag{4.23}$$

$$\sum_{i=1}^{I} a_i F_i(z=Z) = 0 \tag{4.24}$$

$$\sum_{i=1}^{I} a_i^2 = 1 \tag{4.25}$$

Clearly with three equations we will need $I = 3$, i.e. three unknown coefficients a_i, to close the set and obtain a unique solution. Using the shorter notation $b_i \equiv F_i(0)$, $c_i \equiv F_i(Z)$, the solutions are

$$a_1 = \left\{ \left[\frac{b_1}{b_0} - \frac{b_2 (b_1 c_0 - b_0 c_1)}{b_0 (b_2 c_0 - b_0 c_2)} \right]^2 + 1 + \left(\frac{b_1 c_0 - b_0 c_1}{b_2 c_0 - b_0 c_2} \right)^2 \right\}^{-\frac{1}{2}} \tag{4.26}$$

$$a_2 = -a_1 \left(\frac{b_1 c_0 - b_0 c_1}{b_2 c_0 - b_0 c_2} \right) \tag{4.27}$$

which, together with Eq. (4.25) determine the coefficients a_i in terms of the boundary amplitudes b_i and c_i. Equation 4.26 has two square roots; only the negative one satisfies all of Eqns. (4.23–4.25), so this root is chosen.

The two superposition streamfunctions (i.e. $F\rho_0^{-1/2}$) are shown in Fig. 4.6. For the first solution, the coefficients are $a_0 = 0.998$; $a_1 = 0.063$. The flow is almost entirely barotropic, a consequence of the large magnitude of the first baroclinic eigenvector at the top boundary. Because it is so large, only a small coefficient is required to equal the barotropic component there.

In the second case, the coefficients are $a_0 = 0.159$; $a_1 = -0.775$; $a_2 = -0.611$. This means the streamfunction is almost entirely baroclinic. The boundary conditions may be satisfied to a good approximation using only the first two baroclinic modes. This function is localized almost entirely in the stratosphere, so is not quite appropriate to model a baroclinic vortex whose maximum amplitude occurs at or below the cloud-tops, as suggested by observations.

As the eigenmodes are mutually orthogonal, a combination of any three (or more) modes will satisfy the two boundary conditions and the normalization condition, so when looking for interesting structures it is not necessary to limit ourselves to the first three modes. However, the higher-order modes have a more oscillatory form than we would expect for large-scale, physically realistic motions, so are less suitable.

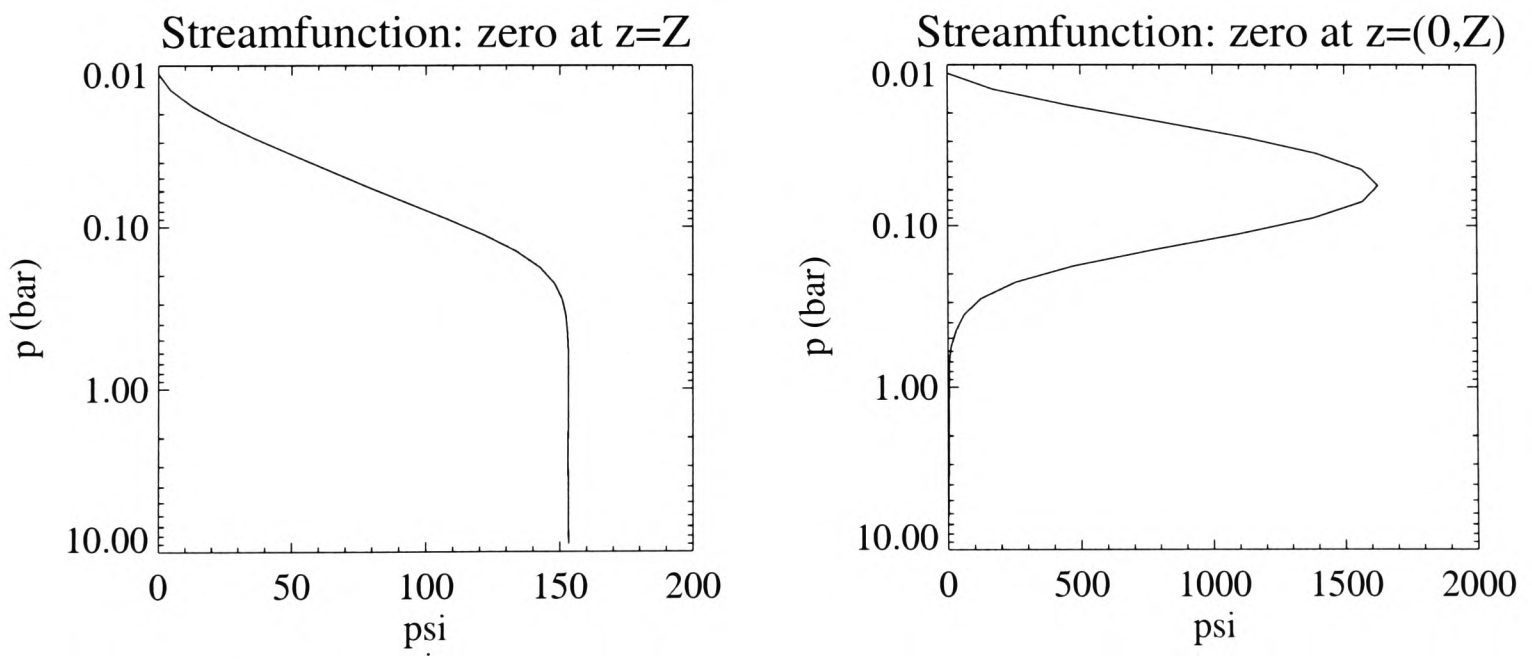


Figure 4.6: Superpositions of eigenmodes: zero at top boundary (left) and zero at both boundaries (right)

Calculation of a model basic state

Geostrophically balanced winds and temperatures may be found for a streamfunction with a given vertical structure (for example, one of those discussed above) and a horizontal structure constrained to fit the observed cloud-top winds. The balanced temperature field is determined by the vertical gradient of the streamfunction ψ, and the winds by the horizontal gradients.

We assume the basic state comprises a purely zonal flow, varying only with latitude and height. The choice of the vertical part, F, of the streamfunction ψ, has been discussed above. Horizontal winds are found by sampling the Voyager zonally averaged cloud-top winds (Limaye, 1986) at half-degree intervals. This is a crude approach, but is adequate for our purpose of obtaining an integrated quantity, the streamfunction.

A more sophisticated technique, that of smoothing by a Savitzky-Golay filter, was used by Dowling (1995) to find the potential vorticity directly. This remains an option for future work. It should also be remembered that the zonal-mean observed wind profile includes an average of waves and vortices, and is therefore not exactly the same as the hypothetical background state. The eddy terms should average to zero but we would expect them to influence the zonal-mean flow. However, for a largely zonal flow like that on Jupiter, the errors should not be too large.

The winds are integrated to find the streamfunction, using the geostrophic relation

$$u = -\frac{\partial \psi}{\partial y} \tag{4.28}$$

This integration introduces another arbitrary constant. In general this could be a function of height, but we have separated out the z-dependence and are only interested in the winds at one level. (We use the vertical structure function to determine the height-dependence.) The value of the integration constant does not affect the winds (which depend only on the gradient of ψ) but only the temperatures. We find the constant by requiring that the integral should vanish. This is reasonable (in fact, essential), as may be seen by the following argument. The temperature anomaly is proportional to the vertical gradient of the streamfunction. Therefore at every level, the meridional integral of the streamfunction must be proportional to that of the temperature anomaly. Clearly the integrated temperature anomaly must be zero (by definition of the word anomaly). Therefore the integrated streamfunction must also be zero.

To find α, it is necessary to divide by the value of $\rho_0^{-1/2}F$ at the level of the imposed winds, where F is the vertical structure function (usually a superposition of two or three eigenmodes). Then ψ is given everywhere in the domain by substitution of the numerical functions into Eq. 4.4. Hence the geostrophic winds and temperatures may be found.

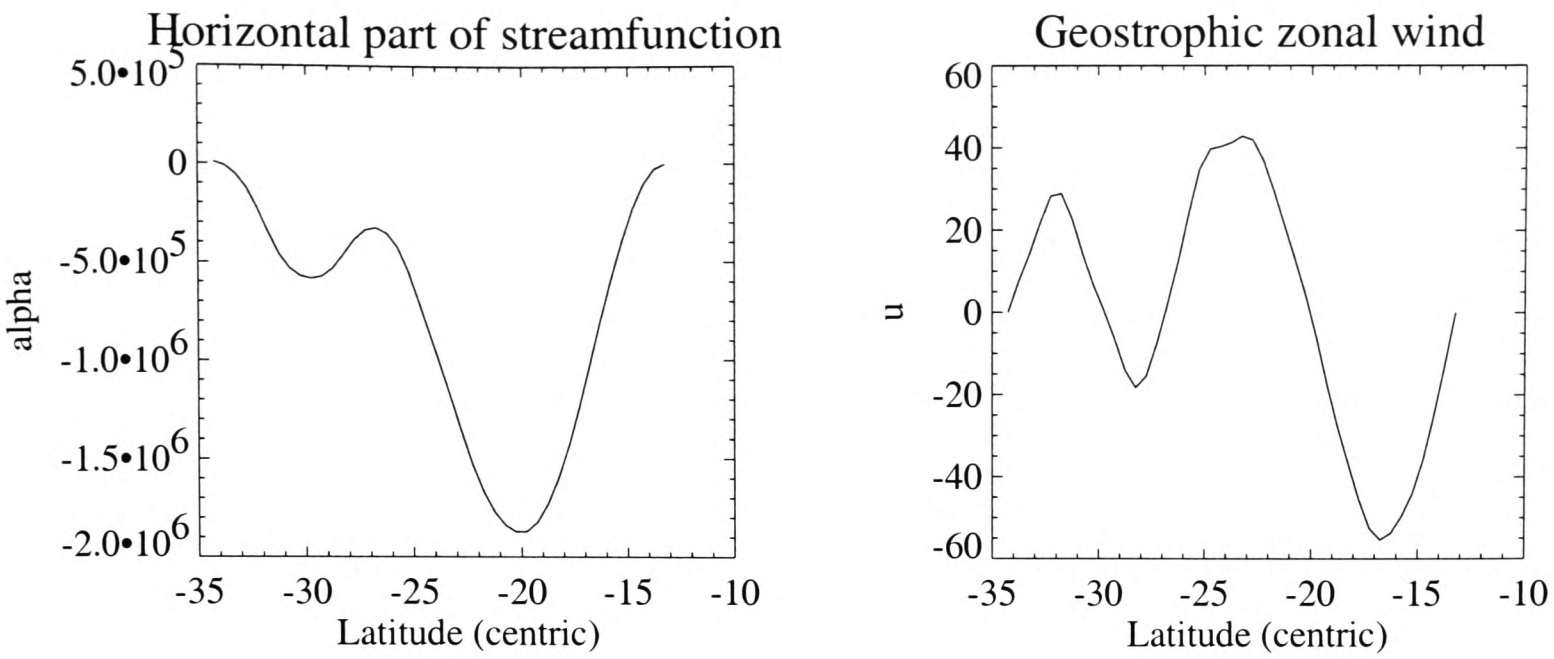


Figure 4.7: Horizontal part of the streamfunction (left) and its associated geostrophic zonal wind (right)

Figure 4.7 shows the zonal wind and streamfunction α as a function of latitude. The zonal wind is found by differentiating the streamfunction. At first sight it seems pointless to integrate a quantity and then differentiate it again. The reason is that numerical differentiation is easier and more accurate than integration, and so by deriving both the winds and the temperatures from the streamfunction, we can be sure they are consistent. The combined effect is to smooth the original wind data slightly. This wind field is taken to be representative of the flow at the cloud-tops. The winds at the base of the model, used to force it, may be found from the streamfunction at the base.

In the case shown on the left of Fig. 4.6, the flow is barotropic below the visible clouds, so the deep winds are taken to be equal to the observed values. (This is in disagreement with the Galileo Probe results—see Section 1.2). Figure 4.8 displays cross-sections of the zonal wind as a function of latitude and height, and the geostrophically balanced temperatures implied by this flow structure. It is assumed that in the basic state, $v = 0$ everywhere.

Other flows may be initialized in the same way. For example a balanced vortex may be modelled by assuming a two-dimensional horizontal Gaussian (or squared Gaussian) distribution for α, together with a suitable vertical structure function. The balanced winds and temperatures may be calculated from this and used to initialize the model.

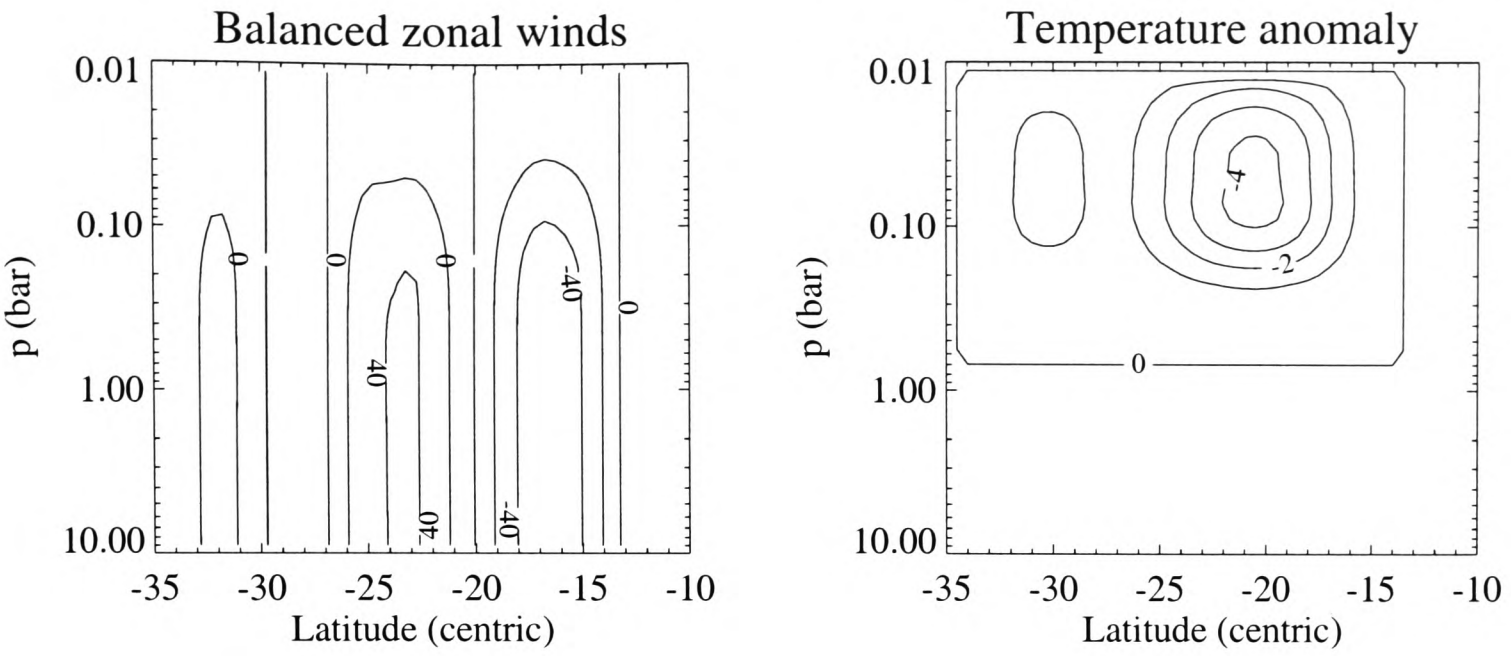


Figure 4.8: Meriodional cross-section of zonal wind (left) and geostrophic temperature anomalies (right)

4.5 A method for more general boundary conditions

Introduction

Our work so far has been restricted by the imposition of purely rigid-lid boundary conditions, and by the absence of any forcing or damping terms within the model domain. These limitations, necessary if the matrix is to be tridiagonal, have prevented us from replicating conditions in the GCM, where the top boundary condition is not always a rigid lid, but sometimes a zero-flow condition, or more often a sponge layer with heavy drag and diffusion. The equatorial-wave models presented in Chapter 6 also have weak drag throughout the stratosphere, which cannot be represented with the matrix diagonalization method.

These considerations suggest that an alternative and more flexible method is required. It is also valuable to have two different techniques for solving even the rigid-lid case, because the solutions may be checked against one another. In fact, the effect of the sponge layer on the normal modes is not significant except near the top of the model, but it was necessary to check this.

We use a standard numerical procedure—the so-called "shooting method". This is an iterative process analogous to ranging artillery, hence its name. A trial solution is integrated forward from the initial conditions (in this case, the lower boundary). When it reaches the upper boundary, the trial solution is evaluated against the required upper boundary solutions and assigned a numerical score. The initial conditions (the gun elevation) are then adjusted and an improved solution found. This process is iterated (hopefully) to convergence.

This iterative method is less efficient and elegant than the direct matrix diagonalization. Its advantages are that the boundary conditions and forcing terms may easily be represented without changing the standard method.

Implementation of the method

The code is written in C and is built around library routines by Press et al. (1986). The iterative adjustment at the heart of the algorithm is done by the Newton-Raphson method. This calls functions to interpolate the static stability ϵ and its deriviative onto the model grid, and to set up the initial conditions and evaluate the score after each iteration. The eigenvalue λ is treated as a formally dependent variable whose gradient is zero everywhere. This is necessary because the method is a special application of a general technique for solving sets of coupled differential equations.

The scoring function can easily be altered to impose different boundary conditions. The vertical structure equation is written in terms of coupled first-order differential equations, which are coded in another function. This allows us to include damping terms in a trivial way.

The grid for the shooting method is much finer than the main GCM vertical grid. More precisely, there is no fixed grid; instead the integration proceeds with an adjustable stepsize and the solution is monitored for error accumulation. The resulting profile is then interpolated onto the coarse grid.

Unlike the matrix method, which is fully automatic, the shooting method is interactive in that the user is asked to choose an initial guess for the eigenvalue. The Newton-Raphson method then finds the closest local solution. The associated eigenvector is plotted and inspected visually. It is easy to identify the mode by the number of turning points it has; each mode is identified and stored. A new guess is then made. Many initial guesses will lead to modes already found, so a certain amount of trial and error is required.

In principle there is no reason why the procedure should not be automated. The method would be to search linearly through all possible eigenvalues, at some carefully chosen interval. The number of turning points could be calculated by differentiation and modes identified accordingly. However, for the limited number of cases required, it is easier and probably less error-prone to set up each shooting attempt manually, by entering a guess and plotting the solution.

The shooting method has another idiosyncrasy: it is unable to reach the zero eigenvalue barotropic solution. Instead, when a suitably small initial value is chosen, the program latches onto a spurious growing exponential solution. This is not a practical drawback,

because the $\lambda = 0$ mode is a trivial solution, and the higher modes are found correctly. It does however illustrate the advantages of applying two different methods to the same problem.

Results: rigid lid with no damping

This is the simple case already covered with the matrix method. We shall use it as a first example, to validate the two methods against each other. Detailed results for the first baroclinic eigenmode are shown in Fig. 4.9. This shows the eigenvector F and its pressure-weighted streamfunction ψ, along with the vertical derivative of ψ. The latter vanishes at the top and bottom, as required by the rigid lid boundary condition.

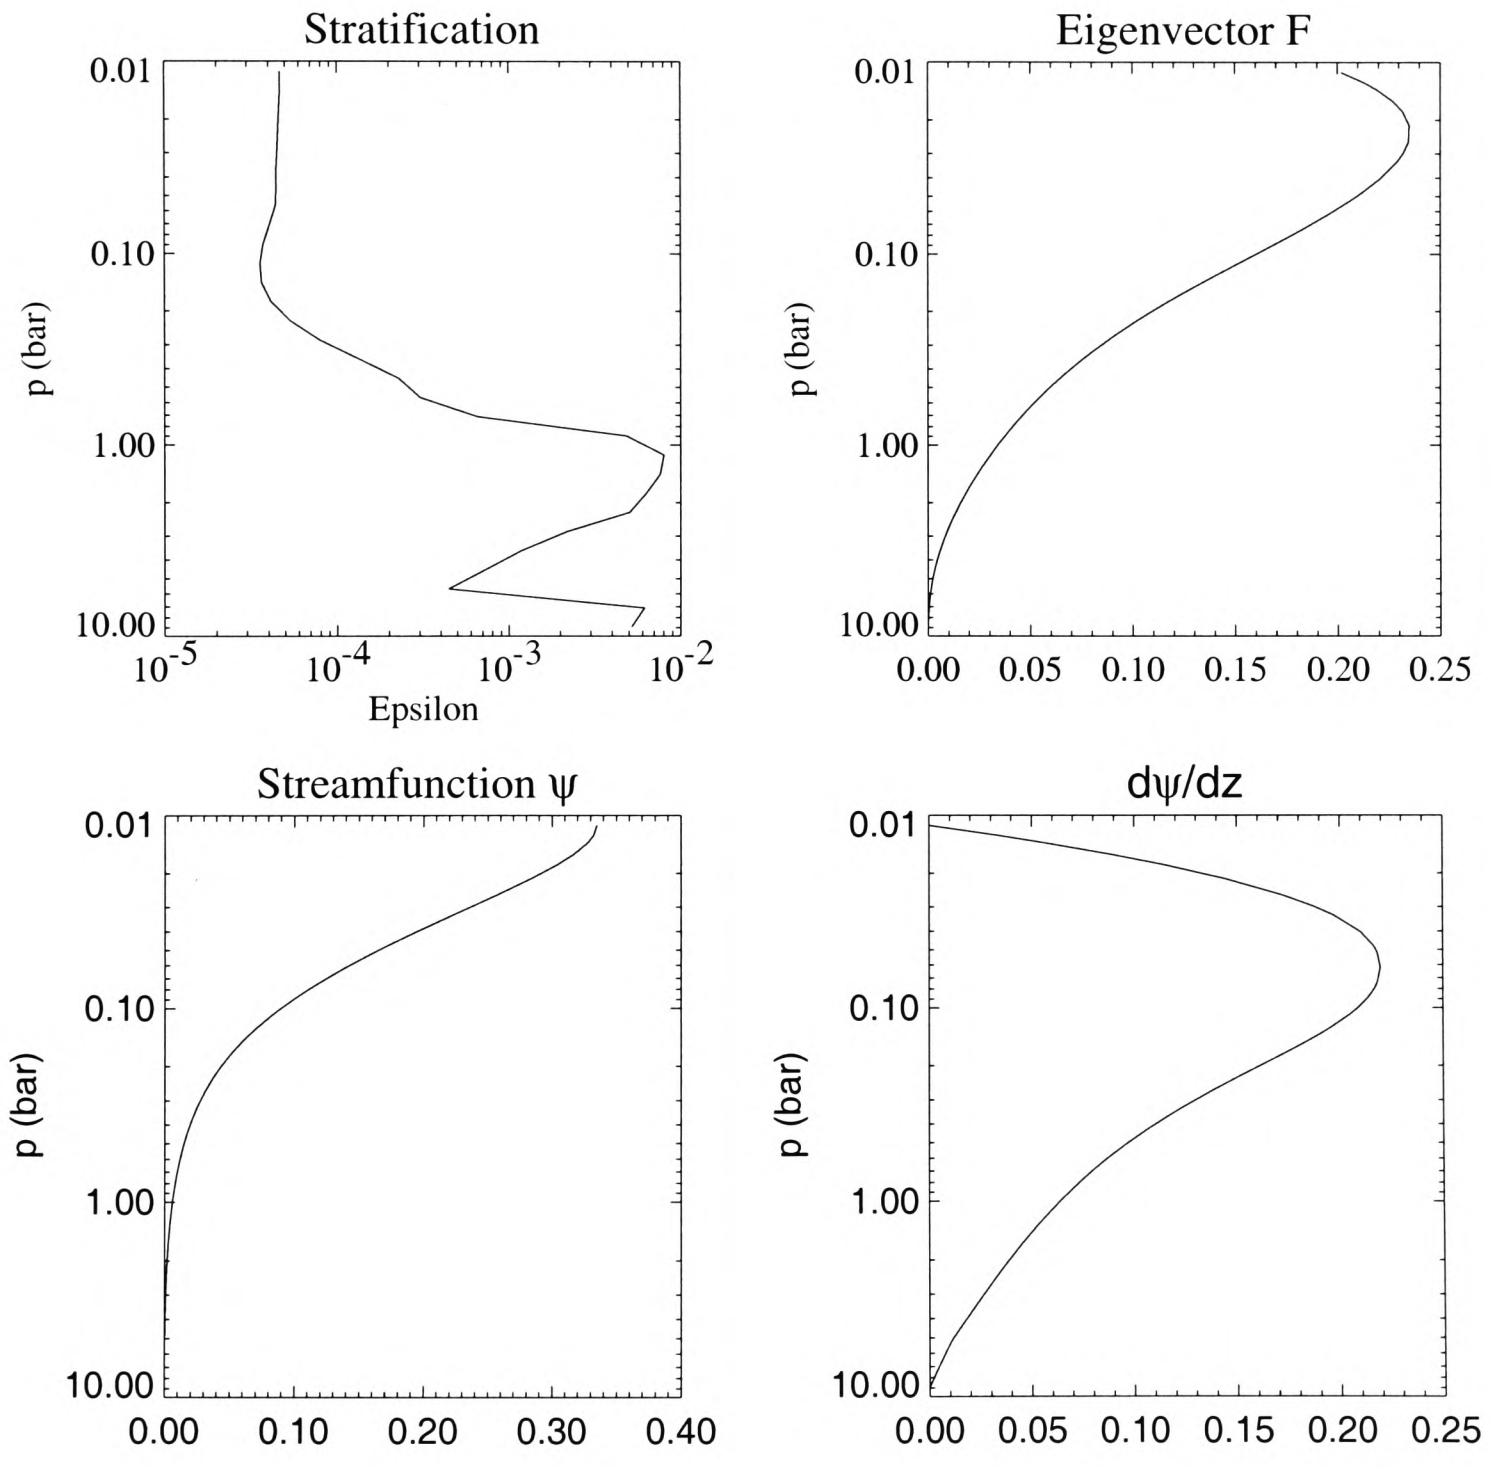


Figure 4.9: Results of the shooting method for an undamped rigid lid case, showing the stratification profile, the first baroclinic eigenvector F, the streamfunction ψ and its vertical derivative.

Figure 4.10 reveals a further advantage of the shooting method over the matrix method. The first four baroclinic modes are shown, and comparison with Fig. 4.5 demonstrates that the lower resolution matrix method fails to detect some of the modes. Mode 3 of Fig. 4.10 corresponds to mode 2 of Fig. 4.5. The first baroclinic mode of Fig. 4.5 now looks like a superposition of the first two modes. The local maximum in the true second mode is so gentle that the lower resolution technique did not detect it. In general the curvature of all the modes is less severe with the shooting method, but the gross features are clearly the same whichever technique is used. Careful examination of Fig. 4.10 shows that Mode 1 has no turning points, Mode 2 has a maximum at about 120 mb, Mode 3 has two turning points and Mode 4 has three.

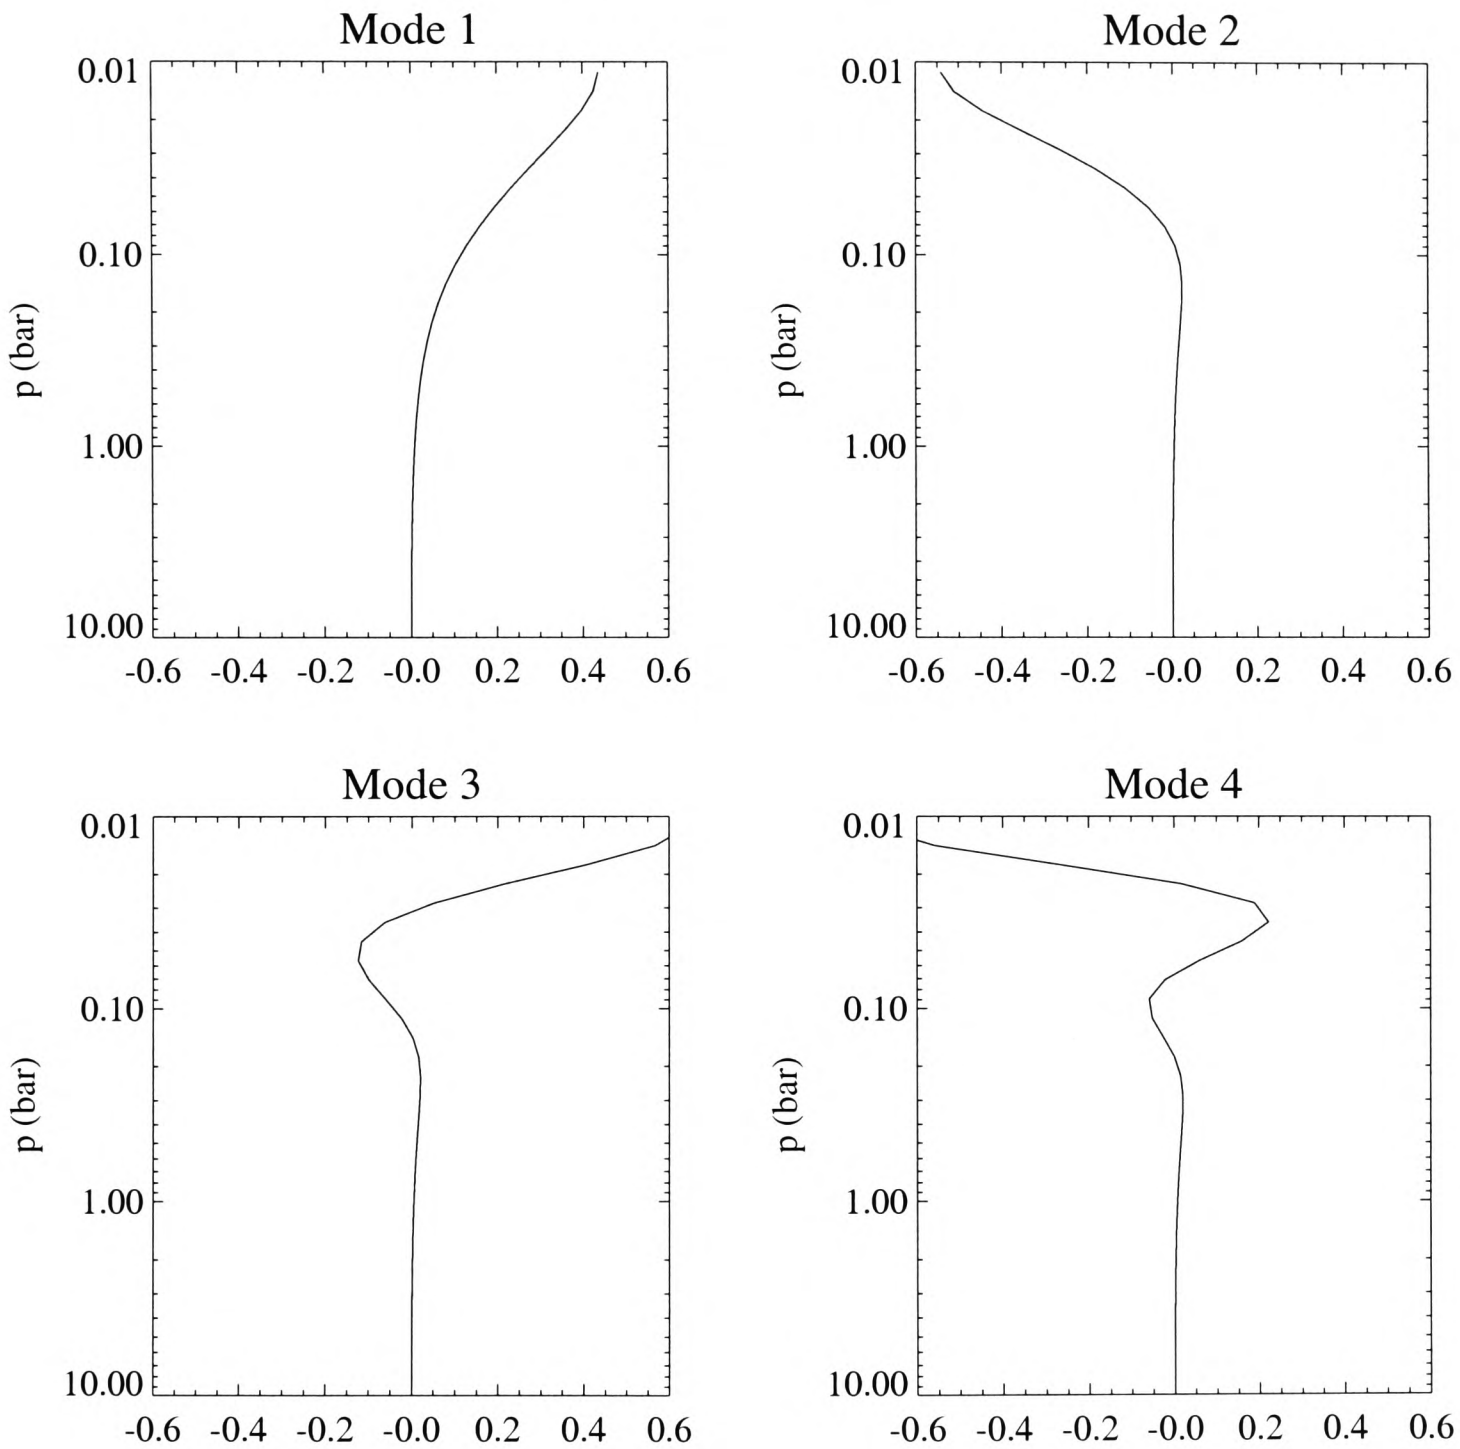


Figure 4.10: The first four baroclinic normal modes for an atmosphere between two rigid lids. Note how the shooting method has identified an extra mode missed by the matrix method (cf. Fig 4.5).

Results for zero flow at top boundary

There is clearly not a rigid lid at 10 mb in Jupiter's atmosphere. Such a boundary is unavoidable in the UM, but what about a slightly different condition, such as zero flow? Observations suggest that the cloud-top motions decay with height and the assumption of zero winds at 10 mb is perhaps no more unreasonable than a rigid lid. We will apply the shooting method to this problem and examine how the eigenvectors are modified by the new boundary condition.

Figure 4.11 illustrates the first four modes. Apart from the obvious difference that they vanish at the top, the effect of the boundary condition is to move the larger-amplitude peaks higher into the atmosphere, and to remove anything resembling the original first baroclinic mode, which reached a maximum at the top lid.

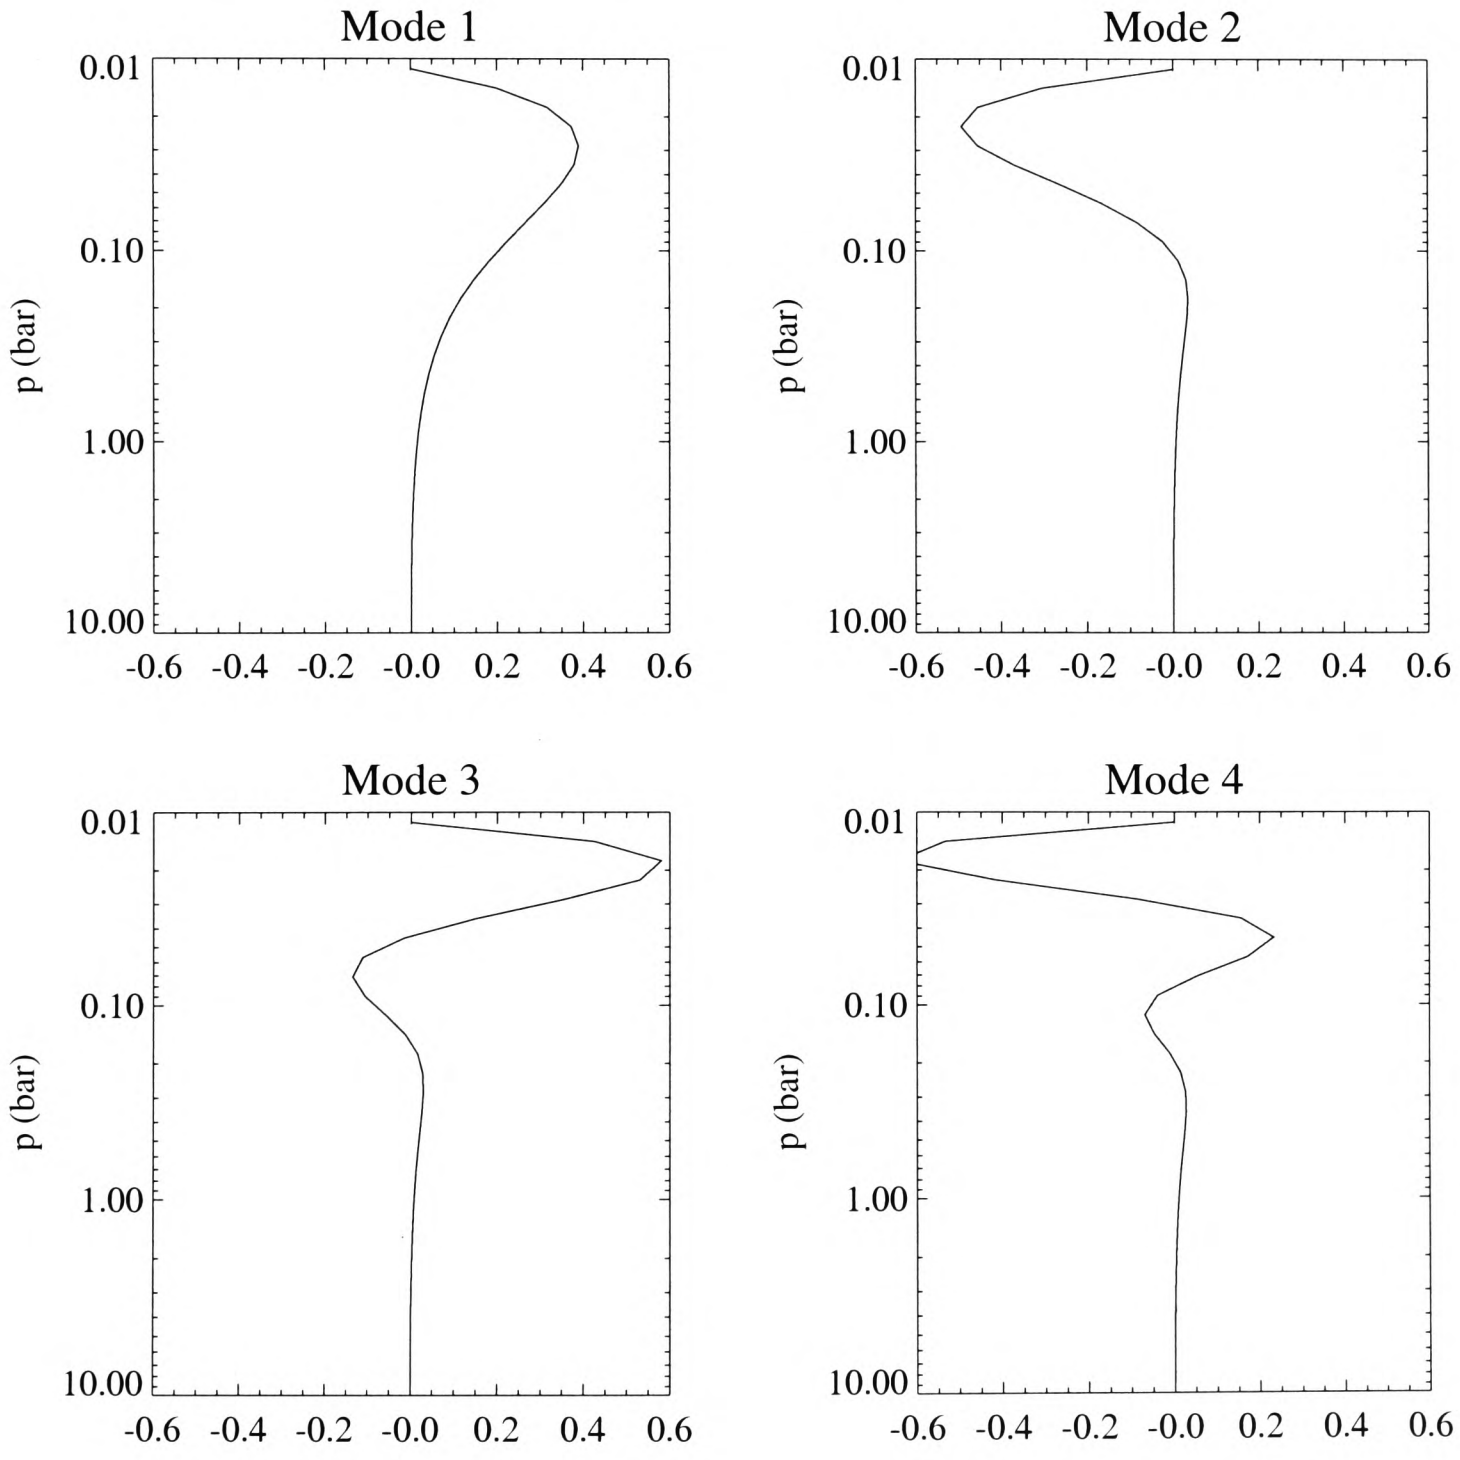


Figure 4.11: The first four baroclinic modes for the case where the flow vanishes at the top boundary.

Results with a sponge layer

To model the effects of a sponge layer, we introduce a linear drag term whose coefficient varies with height (Fig. 4.12). This profile is chosen as a linear approximation to the shape of the sponge in the GCM, which actually varies more smoothly. The other uncertainties in the model will exceed any errors in the profile. These uncertainties include the amplitude of the drag function, interpolation from the coarse GCM grid to the fine shooting grid, and the effect of horizontal diffusion, which is large on the topmost level but smaller below.

The drag is applied to the term in $\partial F/\partial z$, so the equation is analogous to that of the damped harmonic oscillator. In a time-independent model it is not possible to identify the amplitude which best represents the real sponge layer in the GCM, so we present a range of results from light to heavy drag.

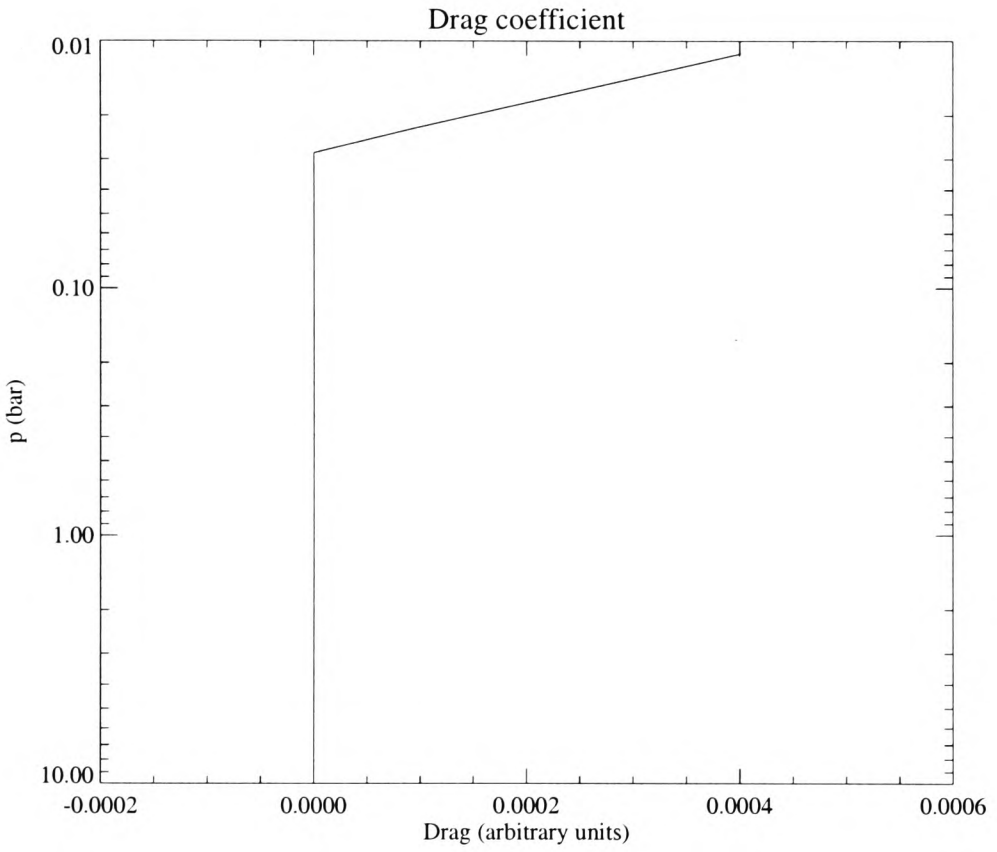


Figure 4.12: Sponge-layer drag as a function of height. The amplitude can be adjusted but the profile is always the same shape: zero below the sponge, then increasing linearly with height.

Figure 4.13 shows the effect (on the first mode only) of varying the sponge layer amplitude. The drag distribution remains the same shape in every case. We can see that the sponge only affects the top few levels, but within that region it is possible to control the amplitude precisely.

This analysis shows that at levels well below the sponge layer, the modes are unaffected, as might be expected. However, the sponge does introduce an extra turning point in the odd modes, which would otherwise have maxima at the top boundary. This will complicate later diagnosis of GCM results in terms of identifying the modes. At the top of the model

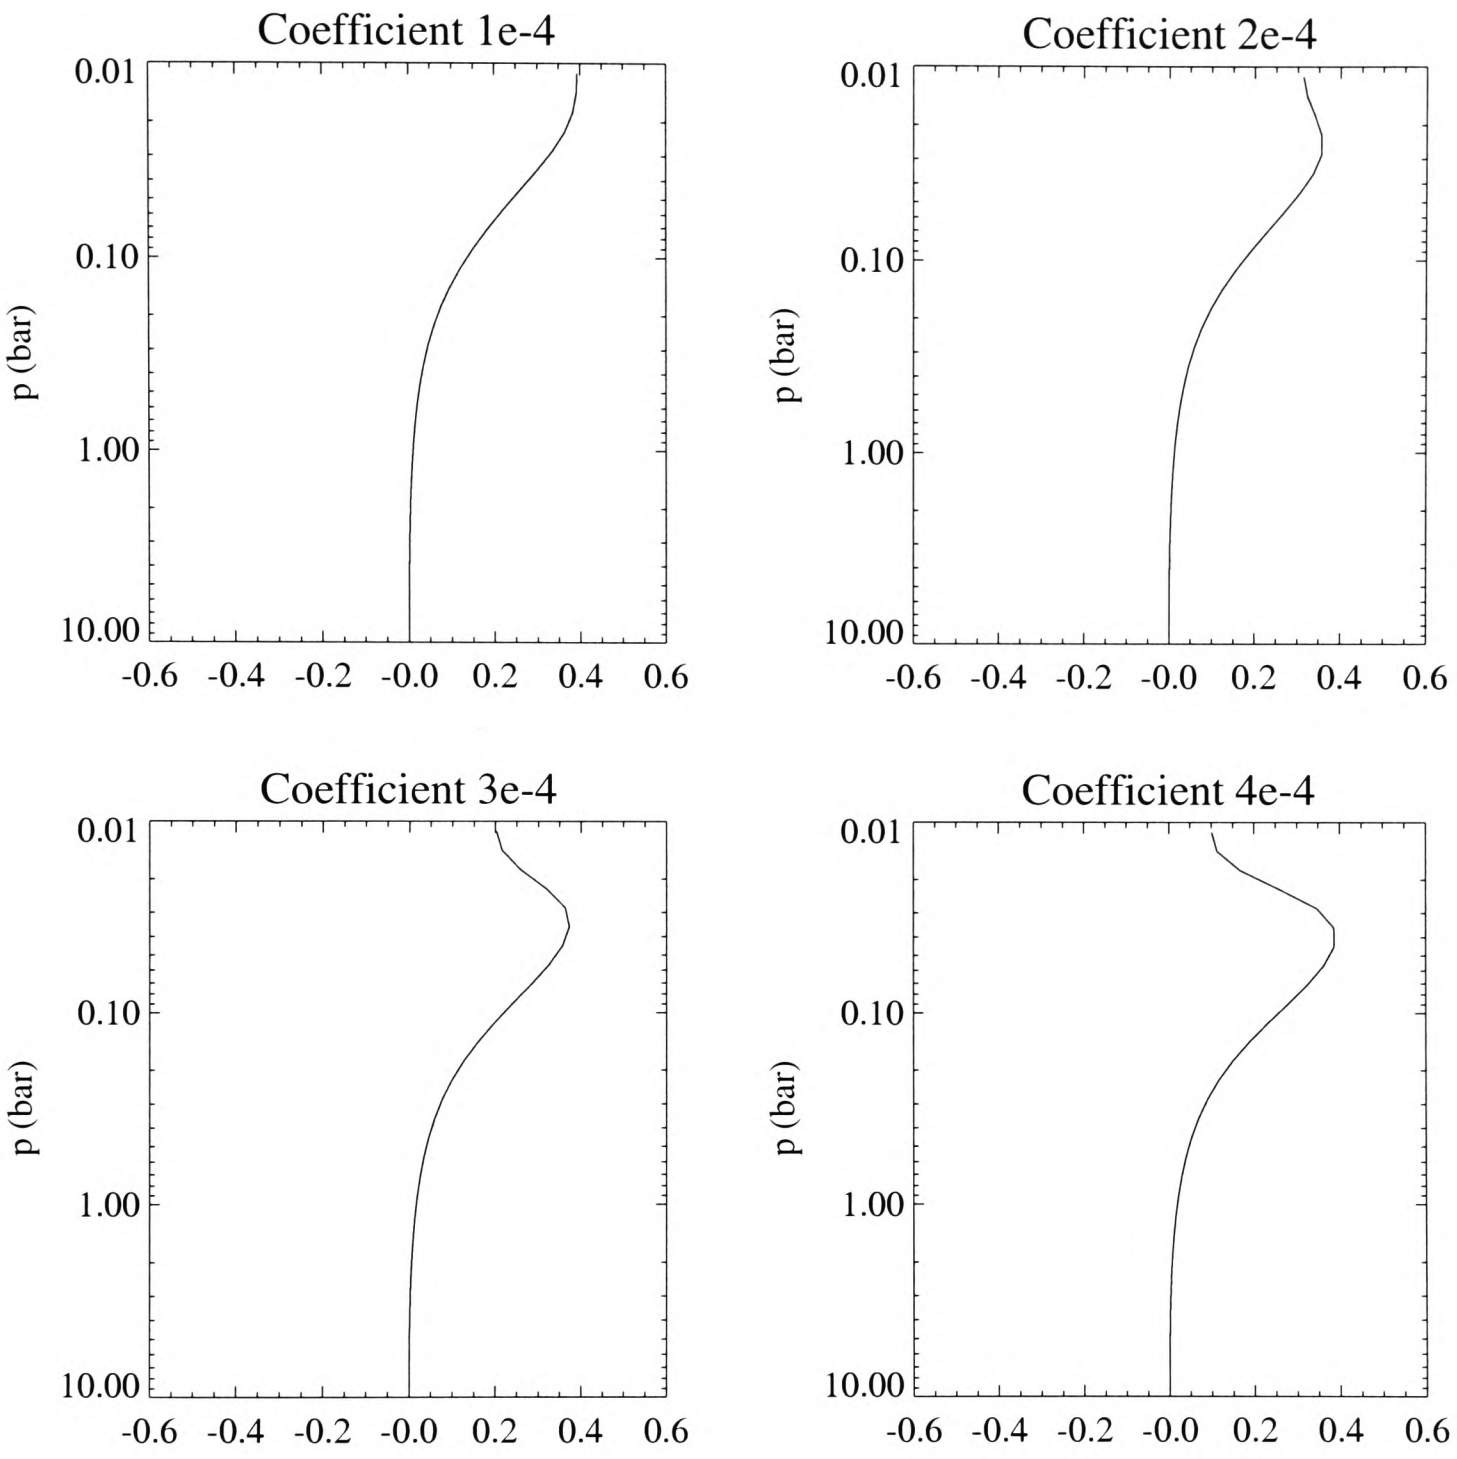


Figure 4.13: The first baroclinic mode for four different values of the sponge layer damping coefficient. The magnitude at the top lid varies approximately inversely with the coefficient.

the mode amplitude is roughly inversely proportional to the damping coefficient, and we do not know the correct value. The GCM results will show that the damping seems quite heavy, so we would expect to see modes resembling those with the larger coefficients.

Chapter 5

Midlatitude jets and vortices

Si nous ne trouverons pas des choses agréables, nous trouverons du moins des choses nouvelles. (Voltaire: *Candide*)

5.1 Introduction

The most striking visible features of Jupiter are the alternating zones and belts, associated with jetstreams, and the giant eddies embedded in these jets. This modelling study attempts to reproduce the general character of the observed behaviour. By varying parameters which are poorly constrained, such as the stratification and vertical windshear below the cloud-tops, we can try to learn more about the range of values they may take. This is done by adjusting the unknown parameters and assessing the effect on the flow above the clouds, which may be compared with observations. In particular, we look at the properties of shallow confined eddies compared to deep ones.

Many general questions spring to mind. Some of the main ones are: what is the origin of the giant eddies? Are they generated by intermittent events such as massive convection, or by instabilities in the zonal flow? What processes maintain (and eventually destroy) them? Are some structures inherently more stable than others, and how does this depend on the background flow? Is it possible to generate model vortices which are confined in latitude, like the observed ones? What constraint does this confinement condition place upon the vertical structures of the jets and eddies? Finally, if a model experiment spontaneously generates eddies, what are their preferred structures and to what extent do these resemble observations?

There is clearly scope for much more work than could be contained in one thesis, and we have not even mentioned such other basic problems as the origin of the jets—which may

well be interlinked with the properties of the vortices. We accept from the start that very little is understood about the detailed behaviour of the jovian atmosphere, although some general aspects, such as the generation of eddies by baroclinic or barotropic instability, have been simulated in several previous models (see Section 5.2). It is unrealistic to attempt to find detailed answers to all the questions listed above, for that task would be equivalent to building a completely realistic model of Jupiter.

Our aim is far more modest: it is to start with the simplest possible case, and gradually to build up a more detailed understanding of eddies in increasingly Jupiter-like environments. Although we shall touch on all of the questions mentioned above, we will focus particularly on the main one suggested by the results of Chapter 2. That is to ask which factors determine the stability and migration of vortices. In particular, are these independent of the vortex thickness, as implied by the PV inversion models? Unavoidably, we also have to address the problem of sensitivity to model parameters. This is less interesting than direct studies of physical influences, but is important in order to isolate (hopefully) real physics from mere numerical effects. A large proportion of the effort involved in the present project has gone into developing, testing and tuning the model. Future studies may be able to concentrate more on the physics.

The second factor we will consider, after the vertical structure of eddies, is their potential-vorticity environment. We do not change the stratification of the model basic state (in contrast to the work of Williams (1996, 1997) who studied a wide range of vertical structures). Instead we modify the background PV by adjusting the zonal wind. We start with a basic state at rest, and compare vortex behaviour there with that of vortices embedded in jetstreams and shear zones. Along the way we note many other interesting features in the model results. These include instability of the jets, and the structure and behaviour of the resulting spontaneous vortices.

The simplest case which can tell us anything about Jupiter's giant eddies is that of a single vortex in a resting atmosphere, with realistic vertical stratification but no background jets. Even this system is capable of surprisingly complex behaviour, and we shall begin by studying it in detail.

We focus on two aspects of the vortex behaviour: the meridional migration rate and the vortex intensification rate. All vortices migrate up or down the vorticity gradient (the direction depending on the sense of the vortex circulation—cyclonic or anticyclonic). They also tend either to intensify, and sometimes split into new eddies, or else to dissipate. The two variables (migration rate and intensification rate) are correlated, since intense vortices migrate more rapidly. In Section 5.3 we will look at what factors influence the rates of migration and intensification.

We did not have an established model of Jupiter which is known to be correct, so did not know which effects are model artefacts and which are physical. To establish this, we

first systematically vary model parameters such as the Newtonian relaxation rate, the horizontal diffusion coefficients, and the sponge-layer Rayleigh drag. All these have a greater or lesser effect upon the migration rate.

Secondly, we examine the basic question of how the evolution of shallow vortices differs from that of deep ones. This is closely related to the problem addressed in Chapter 2, but now we use a three-dimensional, time dependent GCM. In fact, we shall show that the eddies generated in the GCM look surprisingly similar to those in the PV model (and indeed to the observed flows, at least at the cloud-top level).

We also look at the relatively minor point of how best to initialize eddies in the model. We begin with what is perhaps the simplest and most obvious approach, namely to impose a geostrophically balanced vortex as an initial condition. The horizontal form of the streamfunction is a Gaussian. We choose this distribution because it gives smooth, apparently realistic wind and temperature fields, and because it can easily be scaled to adjust the magnitude or width. Other workers, such as Williams (1996), have made the same choice.

This method of imposing QG balanced fields has some disadvantages, particularly later on when we wish to place vortices in shear zones between jetstreams. Accordingly we also look at an alternative method, which is to apply a localised heating function (with correlated cooling above) over a few timesteps. This allows the winds to adjust naturally and in balance.

We shall briefly preview the main results of the work on single vortices in a resting atmosphere. All vortices tend to shrink and intensify as they migrate. The potential vorticity anomaly usually steepens, changing from a Gaussian to a structure closer to a top-hat function. The rate at which this occurs depends mostly on scale-selective model parameters, primarily horizontal diffusion or hyperdiffusion. Scale-insensitive processes (namely Newtonian cooling) have almost no effect.

Most vortices grow and migrate on rather rapid timescales, often reaching the northern sidewall in less than 40 days. This is not satisfactory and much effort was put in to tuning the model so vortices last longer. This is not easy because if diffusion is made too heavy, the vortex simply dissipates. Stable longevity, i.e. no intensification and very slow migration seems difficult to achieve. However, we will see later that adding jets to the model helps to confine the vortex and we see lifetimes of several hundred days. This is still rather short by jovian standards and is a significant drawback of our model.

If a vortex grows quickly enough it can split into two daughters, and this process can be repeated. The structure and properties of the new eddies are similar to that of the original. This is not always the case though—deep eddies generated spontaneously by jet instability can split into new shallow eddies (see below).

The effect of vertical structure is analysed in terms of the normal modes of the resting atmosphere. We find that in general, shallow flows centred high in the atmosphere (around 100 mb) are preferred. The structure is very close to the first baroclinic mode. We interpret this as meaning that the flow tends to adjust towards the gravest mode compatible with the boundary conditions. These are heavy damping in the sponge layer at the top, and a barotropy condition at the bottom. Vortices initialized with a deep structure, i.e. with zero shear at the bottom boundary, are less stable and tend to adjust towards a more baroclinic structure. Where daughter eddies are spontaneously produced, they also tend to be shallow.

This concludes our study of isolated vortices influenced only by the planetary vorticity gradient. In Section 5.3.11 we will discuss the single-vortex results in the light of the theory of frontogenesis. The work in this first half of the chapter allows us to find reasonable values of the tunable model parameters and to establish how vortices behave under the simplest (admittedly unrealistic) conditions. The most glaring disagreement between model and observations is that model vortices rapidly migrate and intensify, whilst those observed on Jupiter remain more or less stable and confined to narrow latitude bands for long periods (years, or perhaps centuries in the case of the GRS).

We tackle this deficiency by adding zonally-symmetric jets to the model. It is straightforward to show by potential vorticity arguments that, under certain circumstances, Jupiter-style jets can act as PV barriers, preventing vortex migration. Our model is now considerably more complicated, for we need to parametrize the forcing of the jets, and to consider the eddy-mean flow interactions. The fundamental assumption is that the zonal mean flow is strongly forced by the deep atmosphere, below the base of the model. This allows us to prescribe a directly-forced jet pattern as a lower boundary condition.

We assume the atmosphere to have no reciprocal effect on the planet's interior (by which we mean the deep atmosphere below the model base). This is undoubtedly an oversimplification, but without coupling the atmosphere model to one of the interior, we can go no further. We shall mention this assumption again in Chapter 6, and in Section 7.2 we look at the prospects for future development of a coupled atmosphere-interior model.

The analysis of the jet/vortex results proceeds along similar lines to those of the single vortex. We look at the effect of different diffusion rates, and at the behaviour of vortices of different vertical structure. Different jet configurations are tried, and we succeed in confining vortices for much longer periods. Instabilities are seen to generate spontaneous eddies with various vertical structures. To help identify the processes at work, we use quasi-geostrophic energy diagnostics. These are of limited quantitative accuracy, but give a useful general guide when used with care to help interpret model results.

Before presenting our new results, we need to place them in context by reviewing some of the many previous models of jovian jets and vortices. The numerical models are presented

roughly in order of increasing sophistication, and we also mention relevant laboratory experiments, which are uniquely able to study real but controllable flows.

5.2 Theories and models of jovian dynamics

Comparisons are odorous. (Shakespeare: *Much Ado About Nothing*)

A large part of the literature on Jupiter models focuses on the generation of, and interactions between, jetstreams and giant eddies. After all, the presence of eddies embedded in alternating zonal jets is perhaps the most distinctive feature of Jupiter's atmosphere. We will briefly cover the historical aspects, and then describe more recent work, categorized by model type (single layer, "$1\frac{1}{2}$" layer, primitive-equation, and so on).

The earliest theories of the Great Red Spot

Initial attempts to explain the nature of the long-lived giant eddies on Jupiter suffered from an ignorance of their environment. Although Cassini had (in 1671) calculated the orbital radius of Jupiter, and hence its mean density, observers continued to suppose Jupiter's composition to be terrestrial, with a solid surface. Green (1887) proposed that the longevity of the Great Red Spot (GRS) implied that it was "... a portion of the more solid part of the planet". The nineteenth century explanation of the coloration of the Spot was to ascribe the red colour to its very high temperature—it was believed to be literally red-hot.

The solid-body theory of the GRS was ruled out by measurements of the drift velocity of the GRS (Denning, 1899). It was superceded by "floating raft" ideas, the first of which was put forward by Hough (1905). In this hypothesis the visible surface of the planet was a liquid ocean, and the giant eddies were elliptical solid plates floating on it. This simple theory was unable to explain the observed flow patterns within and around the GRS. It was soon rejected in favour of various models explaining the Spot, not in terms of a solid topographical feature, but as an atmospheric vortex.

Early vortex models

Kritzinger (1914) first proposed that the Spot is an anticyclonic vortex. This suggestion was based on the behaviour of the adjacent jetstreams. It was not clear that the vortex would be stable, a problem addressed by Taylor-column theory due to Hide (1961) and

later Ingersoll (1969). In this model the circulation in the GRS is induced at a deeper level, by the interaction of the zonal flow with a solid, cylindrical plateau beneath the observed clouds. In a refinement of the theory, the solid plateau was replaced by a "floating raft" (but this time below the clouds), which could account for the observed drift motions. The images of the internal circulation obtained by Reese and Smith (1968), and later by the Voyager spacecraft (see e.g., Mitchell et al., 1981), provided conclusive evidence supporting the vortex theories. Subsequent work has been largely concerned with identifying mathematically idealised systems which are able to generate and sustain realistic vortices.

Recent work: free-eddy models

The wide variety of current and recent models of giant eddies can be categorizied in several different ways. One fundamental division is between the different physical approaches. Falling into one category are those models which seek solutions of the dynamical equations, generally unforced and inviscid. The other group are models which primarily consider the energetics of the flow from different angles. We shall consider a few of the free-eddy models first, and then very briefly survey the energy-source theories.

A natural subdivision of these models is into single-layer schemes, sometimes described as barotropic; and baroclinic models with multiple layers (and hence several vertical degrees of freedom). The latter allow variability with height, which is an important feature in view of the horizontal temperature contrasts observed on Jupiter, which imply vertical wind shears, at least above the cloud layer.

Single-layer models

The idea of the GRS as a long-lived free eddy, without any floating rafts or associated deep topographic structure, was first put forward by Golitsyn (1970). Compared with the Taylor-column theories, this had the advantage of making fewer assumptions about the deep atmosphere. Ingersoll (1973) showed that the barotropic vorticity equation has unforced free eddy solutions which resemble the GRS.

A single-layer, quasi-geostrophic (QG) model was discussed by Maxworthy and Redekopp (1976). This model was not latitudinally confined, but bounded only by two infinite horizontal planes. Weakly nonlinear soliton solutions (solitary Rossby waves), were identified as responsible for the eddies. Solitons exist because of a balance between dispersion (in this case due to the β-effect) and nonlinear processes (here, advection). The waves were maintained by extracting kinetic energy from the zonal shear—a form of baroclinic instability. Korteweg-de Vries solitons do not interact with each other, so the model predicted

that colliding eddies would pass through each other undisturbed. This is contrary to observations (Mac Low and Ingersoll, 1986), which indicate that colliding spots do in fact usually merge irreversibly, at least if they are of similar sizes. The two either destroy each other or else form a larger stable spot. For this reason the original soliton theory seems inadequate.

This shortcoming was later partly alleviated by an extension to the theory (Maxworthy et al., 1977). It was shown, for the simpler case of a homogeneous atmosphere, that two Rossby solitons with different phase speeds can interact. The two waves are governed by a pair of KdV equations which are coupled through the nonlinear terms; this allows the interactions. Maxworthy et al. predicted that solitons in a baroclinic atmosphere would show still more complicated behaviour.

For a detailed account of the properties of vortices in a single-layer "reduced gravity" shallow-water system, see Williams and Wilson (1988). This study, which included no diffusion beyond that required for numerical stability, illustrated a range of behaviour. The vortices were generated by horizontal shear, and eddy absorption was suggested as a mechanism for sustaining the largest vortices. A preponderance of anticyclones over cyclones was predicted, in apparent agreement with observations of Jupiter.

This result demonstrates a possible advantage of the shallow-water (SW) equations over the QG equations used in other studies: the latter are symmetric with respect to the sign of relative vorticity, and therefore favour neither cyclones nor anticyclones. The nonlinear SW equations, however, allow the dominance of anticyclones as an effect of advection by ageostrophic horizontal velocities. This does not necessarily mean that the same mechanism operates on Jupiter, though. Our GCM can generate spontaneous anticyclones but not (so far) cyclones. This may of course be a result of our jetstream configuration. Perhaps the most we can say is that the QG equations lack an essential mechanism for favouring anticyclones. It is not particularly surprising that highly simplified equations reproduce the flow less accurately.

A single-layer model is unable to take account of the effect of the deep atmosphere on the shallow layer. This deficiency is addressed by the class of so-called "one and a half layer" models. In these systems the shallow "weather" layer is assumed to be hydrostatically supported by a much deeper layer of steady zonal flow. There is only one vertical degree of freedom, associated with the weather layer. The lower layer is assumed to be much deeper, and therefore not affected by motions in the upper layer. This system is dynamically equivalent to a single shallow-water layer with spatially varying bottom topography. The bottom topography represents the effect of the deep motions; for example the assumption of solid-body rotation implies a flat bottom, so the system reduces to the single-layer model of Williams and Wilson.

The quasi-geostrophic $1\frac{1}{2}$ layer "modon" model of Ingersoll and Cuong (1981) adopted

the QG equation set used by Maxworthy and Redekopp, but with bottom topography replacing the plane lower boundary. However, the stratification was chosen to be smaller, in order to prohibit the formation of solitary Rossby waves. The motions in the lower layer were assumed to be equal to the observed (upper-layer) flow far from the GRS, and the bottom topography was prescribed accordingly.

Vortices were again observed in the experiments, and when pairs of stable vortices collided, they merged into a larger and still stable vortex. This result is the opposite to that of the soliton models, but is confirmed by observations. In this inviscid model, the larger eddies grew continually as a result of "feeding" on the smaller ones. Ingersoll and Cuong concluded that in a viscous system, eddy merging may cancel the effect of dissipation on the GRS, resulting in its present long lifespan.

Dowling and Ingersoll (1988, 1989) assumed that a $1\frac{1}{2}$ layer formulation was valid for the GRS, and derived the bottom topography from Voyager cloud-tracking data instead of prescribing it *a priori*. This was done for both a quasi-geostrophic model (Dowling and Ingersoll, 1988) and then for the more general shallow-water system (Dowling and Ingersoll, 1989). The method was to assume that potential vorticity was a conserved tracer, which is the case for invscid adiabatic flow. The effective layer depth was calculated from the measured streamfunction, specific kinetic energy and absolute vorticity. From these data they derived the shear in the deep winds, or equivalently the gradient of bottom topography. Absolute velocities depend on the (unknown) static stability, so can be found only to within a constant coefficient.

Experiments were performed on longevity and genesis of vortices, both with and without explicit dissipation. Vortices formed continuously in the shear zone when the zonal wind was forced towards the observed, unstable profile. They always merged into a single large vortex resembling the GRS. Dowling and Ingersoll suggested that this might be maintained against dissipation by absorption of smaller vortices. This behaviour appears to be similar to that seen by Williams and Wilson using flat topography.

Models with multiple layers

A two-layer quasi-geostropic system, allowing two vertical degrees of freedom, was examined by Williams (1979). This effectively assumed a solid suface below the two shallow layers; the model was forced by a prescribed differential heating rate. The zonal flow was assumed not to extend deep into the planet. The model was able to generate an anticyclonic eddy with many of the characteristics of the GRS. These included a similar scale and shape, latitudinal uniqueness and a slow easterly or westerly drift. This "Gyre" persisted for 955 days, until the end of the model integration, and was identified as the warm high-pressure region of a baroclinic wave. Limitations of the model included an inability

to predict a longitudinally unique feature like the GRS. Such gyres were absent from an earlier, barotropic model (Williams, 1978), suggesting that baroclinicity is necessary for them to exist, at least under the assumptions of these models.

A different approach to the problem of modelling baroclinic effects was taken by Achterberg and Ingersoll (1989). Their model had a QG, beta-plane formulation, but the vertical structure of the geostrophic streamfunction was expressed as a sum of normal modes. Tests of the model showed that truncation after two normal modes is a useful approximation to a continuous system. The normal mode model can be compared to a $1\frac{1}{2}$ layer model, if we assume the flow in the deep layer to be zonal and steady, and neglect all but one baroclinic mode in the weather layer. The approach has not been developed beyond this early stage, and models of vortices have not yet been published.

Latent-heat models

The radiative time scale on Jupiter is probably of order 10 years at most (Conrath et al., 1990). This implies that a free-eddy theory must include a plausible forcing mechanism if it is to explain a feature which has existed for at least a hundred years, and possibly more than three hundred. The possible maintenance of the GRS by merging with small eddies, or by shear instability, have been discussed above. Other suggestions are based on analogies with terrestrial phenomena, particularly hurricanes.

The thermal structure above the Spot implies that it has a warm core of ascending motion. This upwelling may release latent heat by condensation of water vapour, in the same way as in a hurricane. The comparison is discussed by Flasar et al. (1981). Unfortunately, the lower boundary conditions in the two cases are very different, and the rôle of moist convection on Jupiter is not known. For these reasons it is difficult to evaluate the suitability of the latent heat theory.

Another possibile energy source might be the latent heat of *ortho-para* hydrogen conversion. The effects of this are not well understood, although on Jupiter it is believed to have only a small effect. Smith and Gierasch (1995) presented a 1D thermodynamic and convective model of the major planets' atmospheres to analyse this. They concluded that hydrogen conversion is much less dominant on Jupiter than on, say, Neptune. This is because Jupiter's radiative timescale is shorter than the hydrogen conversion equilibration time, which is of order 10^9 s at 300 mb (Conrath and Gierasch, 1984). However this time may be much shorter in the deep atmosphere, leaving the question open. Another unknown is the effect of aerosols on the conversion rate. Massie and Hunten (1982) found that aerosol surfaces can catalyse the conversion. However, Conrath and Gierasch concluded that the total aerosol surface area is smaller than had been assumed, and that the effect is not actually significant.

Laboratory models

This thesis is almost exclusively concerned with numerical models; but laboratory systems must be mentioned as well, for these provide an entirely different approach to the problems. One such piece of apparatus is the rotating fluid annulus, which may be heated or cooled either at the sidewalls or internally within the fluid. This thermal annulus can simulate flow in a latitudinally bounded channel on a differentially heated planet. The annulus has the advantage (unlike numerical models) of continuous stratification. The beta-effect and topography may be introduced by changing the shape of the bottom of the tank. Alternatively it is possible to study the shallow-water equations directly, using their literal physical realisation.

Rotating annulus experiments have been applied to the study of baroclinic eddies by, e.g. Hide (1980) and Hignett et al. (1985). Read and Hide (1984) observed a solitary eddy, confined to a limited area and closely resembling the GRS. The eddies are observed to drift in a realistic way. A wide range of phenomena can be simulated naturally by these experiments, including Jupiter's cyclonic "brown barges"; this is achieved by internal cooling and sidewall heating. The rotating annulus can be seen to be a flexible experimental tool capable of agreement with varied observations.

Primitive-equation models

The present generation of Jupiter models are much closer to the general-circulation models (GCMs) used for weather forecasting on the Earth, and for climate research on both Earth and Mars. They feature many vertical levels, allowing proper resolution in the vertical, and solve the full primitive equations, without making approxmations such as quasi-geostrophy or shallow water theory.

There are two published examples of atmospheric models comparable to ours: those of Williams (1996) and Dowling et al. (1998). These models and their results will be referred to in more detail later on, when we compare specific features with those of our model. Williams' model uses a periodic channel geometry and solves the primitive equations with a Boussinesq (incompressible) equation of state and geometric-height vertical coordinate. The channel depth varies from 100 km to a very deep 10,000 km. Williams' approach is to vary the vertical structure (i.e. static stability) in order to isolate structures permitting the existence of long-lived vortices. The results suggest that thin vortices migrate more slowly than thicker ones. Migration can be slowed by an easterly jet with the same vertical structure as the vortex. Vortices can be generated by baroclinic instability of easterlies, and multiple vortices interact, the larger absorbing the smaller ones.

The EPIC model (Dowling et al., 1998) also integrates the hydrostatic primitive equa-

tions, but with a compressible atmosphere. The vertical coordinate is potential temperature — in other words it is an isentropic model, with no adiabatic vertical motion. Radiation is approximated by Newtonian cooling (as in our model), but the latent heating effect of ortho-para hydrogen conversion is included. This model is not specific to Jupiter, and we shall later refer to an application of the EPIC model to Neptune's Great Dark Spot (LeBeau and Dowling, 1998). This study is relevant to our work on the migration of isolated vortices.

Summary

This survey has attempted to describe and compare an incomplete but representative sample of the literature aimed at understanding the processes which create and maintain giant eddies. A few points emerge from this short analysis.

It is possible to simulate the observed dynamics of belts, zones and eddies with several different simple mathematical systems. The current models can simulate several aspects of the observed behaviour of Jupiter's atmosphere. This suggests that the dynamics on the largest scales are dependent only on the features common to all these equation sets. To discriminate further it is necessary to examine smaller-scale behaviour. For example, the complicated role of the filamentary regions in creating spots and interacting with them (Mac Low and Ingersoll, 1986) has yet to be examined by modelling. There is also a need for more detailed observational data (particularly on the vertical structure) to constrain models.

The simplest models are unsatisfactory for other reasons too. The largest value of the Rossby number within the GRS and White Ovals is about $0 \cdot 36$ (Mitchell et al., 1981); this means the quasi-geostrophic approximation may be questionable here. It also has notable deficiencies in the wider context, for example the inability to favour anticyclones over cyclones. Even the non-linear shallow water models are no more than an analogue (albeit a useful one) of the real atmosphere.

The most recent models improve on this situation in two ways: firstly by treating more realistic equations (the full PEs) but more importantly by achieving useful vertical resolution with suitable jovian stratification. In the future we may expect to see these models coupled to proper radiative transfer schemes (e.g. that of Conrath et al., 1990) and perhaps even to models of the deep interior, such as those of Guillot et al. (1994b).

It is notable that, while most models show reasonable jovian-style vortices, few make any attempt to generate the jets in a natural way. Williams (1975) showed how stochastic turbulence can produce many small vortices, which eventually elongate, join together and form zonal jets. It may be much more difficult to achieve this in a stratified model, and is

unlikely to produce a superrotating equatorial jet. In the present thesis, we do not try to explain the midlatitude jets, although we will show (Chapter 6) a possible mechanism for the equatorial jet.

5.3 Gaussian vortices in an atmosphere at rest

As mentioned in the Introduction, we begin by making the assumption that a useful beginning can be made by modelling vortices in the absence of zonal jets. We should stress that in reality, the generation and maintenance of jets and vortices on Jupiter are probably interlinked. Crudely speaking, it has been shown that, at least for simpler systems of equations, random turbulence can create small eddies. These coalesce into larger vortices and ultimately into zonal jets. Conversely, barotropic and baroclinic instability of the jets can generate further eddies. The presence of the jets also acts to confine eddies to particular latitudes. For early examples of these models see e.g. Williams (1975).

Nevertheless, when developing and tuning such a novel application of the Unified Model, we prefer to approach the problem in gradual stages, on the grounds that it will be easier to understand the interaction between jets and eddies if we already know how eddies behave in the absence of any zonal flow.

The potential-vorticity environment used here is a uniform gradient. It might be argued that an even more quiescent environment would be a jet of curvature such that the PV gradient (approximately $\beta - \partial^2 u/\partial y^2$) is zero. In practice this does not occur on Jupiter, since β is relatively large. Since the planetary radius is so large, any reasonable channel width leads to large y, and hence u would need to be enormous to achieve a channel with constant PV across it. It is therefore not feasible to create such a model. Later we shall use jets in a more limited but realistic way, to create PV barriers against vortex migration.

5.3.1 Initialization

The model grid for these experiments had 90 longitudinal columns and 62 latitudinal rows, with a resolution of $0.5°$. All experiments were run in a limited-area domain within the southern hemisphere. The reason for this choice is that the most interesting jovian eddies (the GRS and White Ovals) are found in the south.

The northern edge was placed at 20°S, so the domain extended to 50°S. The channel was periodic in the east-west direction with length $45°$. The vertical domain was from 10 bar to 10 mb, on 30 levels uniformly spaced in $\ln(p)$. Linear drag was applied to the winds at the top of the model (in a three-level sponge layer) and weakly on the lowest level.

The reference temperature profile was that described in Section 4.3 and was taken to be uniform in latitude and longitude. Thermal relaxation towards this profile was applied everywhere in the domain, usually with a coefficient of 2.7×10^{-9} per 15-minute timestep. This value is the reciprocal of Jupiter's orbital period, which was estimated by Conrath et al. (1990) as a suitable radiative timescale for the lower stratosphere (see Section 3.3.1).

The basic state was an atmosphere at rest, with no mechanical forcing and no horizontal temperature gradient. The initial temperatures everywhere were set equal to the reference profile. For each experiment a single balanced vortex was superimposed on this state. As described in Chapter 3, the linearised QG potential vorticity equation may be separated into horizontal and vertical structure equations. For analysis of shallow eddies, the vertical structure is chosen such that the eddy amplitude vanishes at both the top and bottom boundaries. For the structure to be simple and useful, it is helpful if it is a superposition of only a few normal modes of the vertical structure equation.

A superposition comprising the barotropic mode and first baroclinic modes gives a plausible "trapped" structure, with only one maximum.

The horizontal shape of each vortex was calculated geostrophically. For simplicity and ease of scaling the width and amplitude, the initial streamfunction was chosen to be Gaussian in longitude and latitude. The balanced winds and temperatures were then calculated by horizontal and vertical differentiation respectively. The wind and geostrophic temperature fields for a midlatitude anticyclone are plotted in Fig. 5.1.

5.3.2 Evolution of a shallow midlatitude anticyclone

Figure 5.2 shows the evolution of the anticyclone over forty days. The 56 mb temperature fields are plotted at 10-day intervals. The winds adjust in balance with the temperatures, increasing with the temperature anomaly. The vortex immediately begins to migrate westward (retrograde). This westward motion, which will be seen to occur in all our experiments, may be interpreted in terms of Rossby waves. The Rossby or planetary wave (whose restoring force is the planetary vorticity gradient) is the dominant form of wave motion on the largest scales. An important characteristic of the Rossby wave is that its phase and group velocities are always directed westward. Therefore it is plausible that the vortex is actually a Rossby wave-packet. Later, experiments in Section 5.4 will show long waves gradually "rolling up" into vortices, which supports this interpretation. The association between vortices and wave packets is sometimes referred to as "wave-vortex duality"; see for example Read et al. (1997).

The vortex shrinks and steepens, intensifying to reach a peak speed of 100 m/s by day 20. A slow equatorward migration occurs: the vortex remains centred on a latitude of 40°S

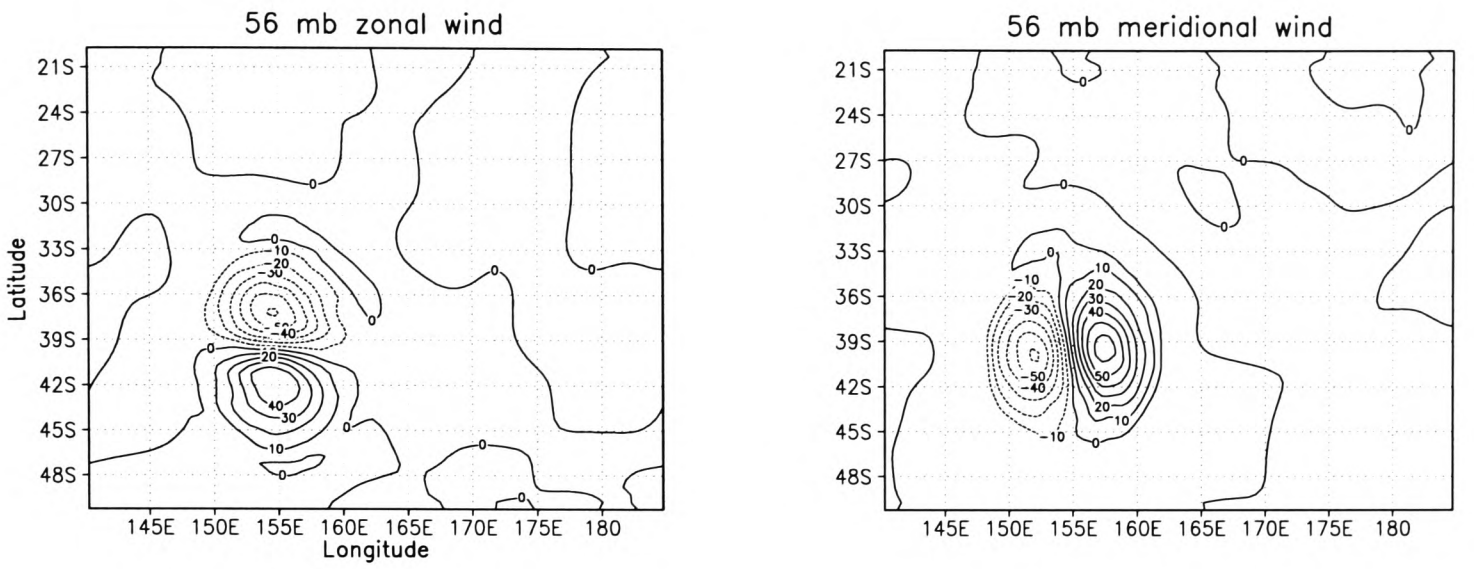


Figure 5.1: The early state (after 1 day) of a geostrophically balanced Gaussian vortex. The upper panels show the wind components at the level where they peak. The lower panels show the temperature field at this level, and a vertical cross-section through the vortex at 155°E, 42°S.

for around 20 days. By day 20 it has migrated only one degree equatorward.

The precise mechanism for migration is difficult to pin down for realistic (not idealized) model flows. We can view a vortex as a potential vorticity anomaly, where the PV inside the vortex is more anticyclonic than the ambient value due solely to the planetary vorticity. The system tends to evolve so as to remove the anomaly. One way to do this would be for the vortex to expand, shrink or dissipate, but it is balanced, with the Coriolis, pressure-gradient and centrifugal forces mutually cancelling. The other way is for the anomaly to migrate towards a latitude whose ambient PV is closer to that inside the vortex. This means that anticyclones will tend to travel equatorward and cyclones poleward.

LeBeau and Dowling (1998) take this argument further (in their case, referring to Neptune's Great Dark Spot), drawing an analogy with the usual view of a terrestrial cyclone. They argue that the rotation of the eddy 'drags' ambient air from poleward to the west of the vortex, and from equatorward to the east. This sets up a dipole structure, much weaker than the main monopole vortex: a so-called beta-gyre. The dipole propels the vortex equatorward. Thus the vortex drifts by causing advection of the ambient PV gradient around itself.

Over the next 20 days the vortex steepens and migrates more quickly, reaching an internal wind speed of 300 m/s at a latitude of 38°S by day 20, and 450 m/s at 35°S at day 40. At day 35, we can see the vortex shedding a subsidiary eddy behind it. The new eddy is much smaller and weaker than the original one, and hence migrates at a lower rate. It behaves the same as the first, steepening and migrating equatorwards. When each vortex reaches the northern sidewall, its equatorward migration is halted and it slowly dissipates.

Figure 5.4 shows the later evolution of the flow, starting from day 70. These are contours of relative vorticity ζ on the 56 mb pressure surface, in arbitrary units. The "shielded" structure of the eddies is clearly shown: each has a core of anticyclonic vorticity surrounded by a ring of cyclonic vorticity of the opposite sign. Note that our model is in the southern hemisphere; were it in the north, the signs would be reversed.

By this stage the second "daughter" eddy is of similar size and structure to the first. Several smaller features have begun to appear. A third anticyclone spins up in the wake of the second, shown in closeup in Fig. 5.5. Animated sequences demonstrate that this is initially projected poleward. Its poleward motion is soon arrested, and it begins to drift equatorward like the others. All three major anticyclones have similar vertical structure, peaking at the 56 mb level. No higher-order modes are excited, and the flow retains its initial vertical structure.

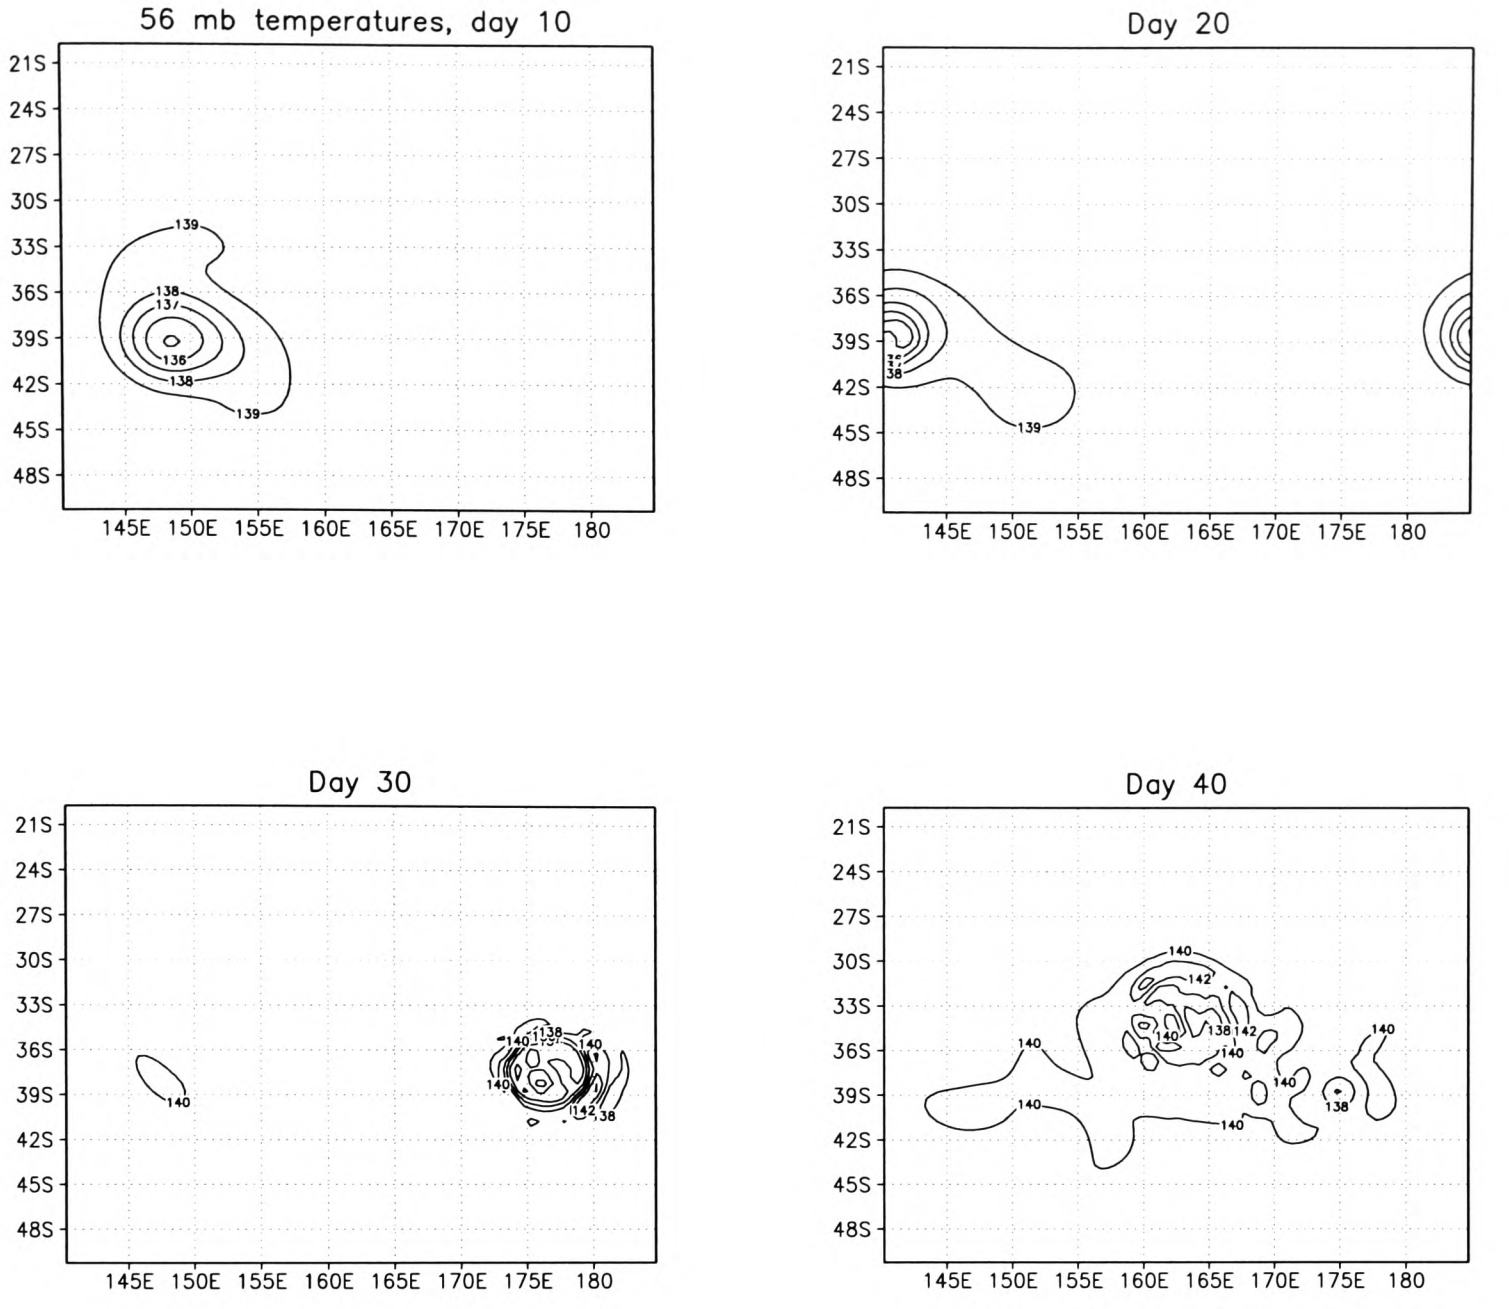


Figure 5.2: Evolution of a midlatitude anticyclone. The vortex migrates westward and equatorward, diminishing in size but growing in magnitude. After 40 days we see a smaller vortex detach from the original (see Fig. 5.3). The units are Kelvin in each case, and this is the absolute (not anomaly) temperature field.

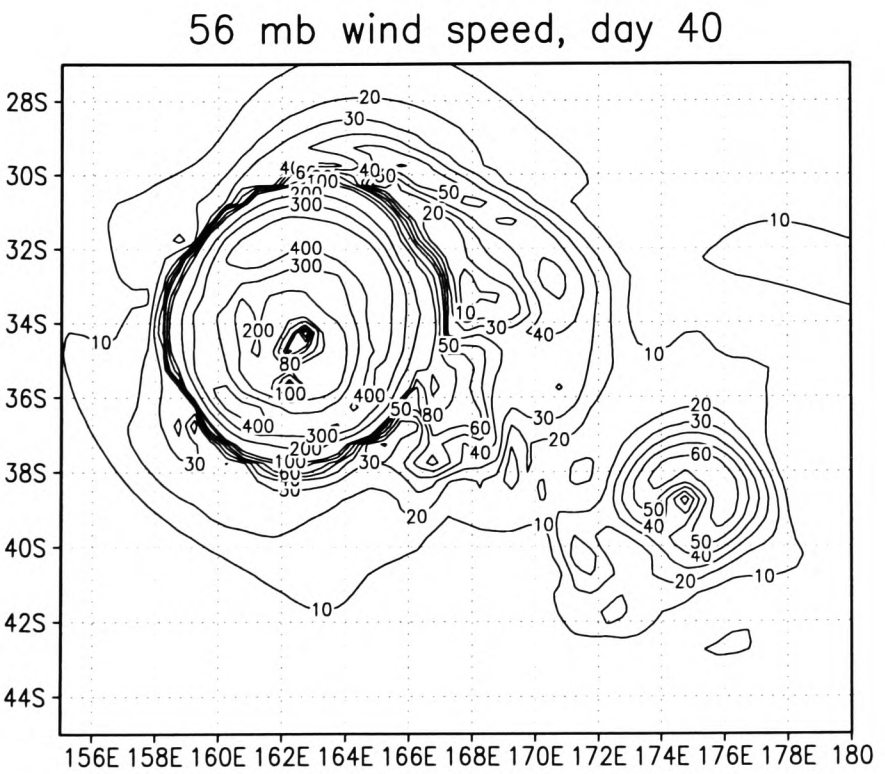


Figure 5.3: Closeup of wind speed, showing the large vortex shedding a smaller one.

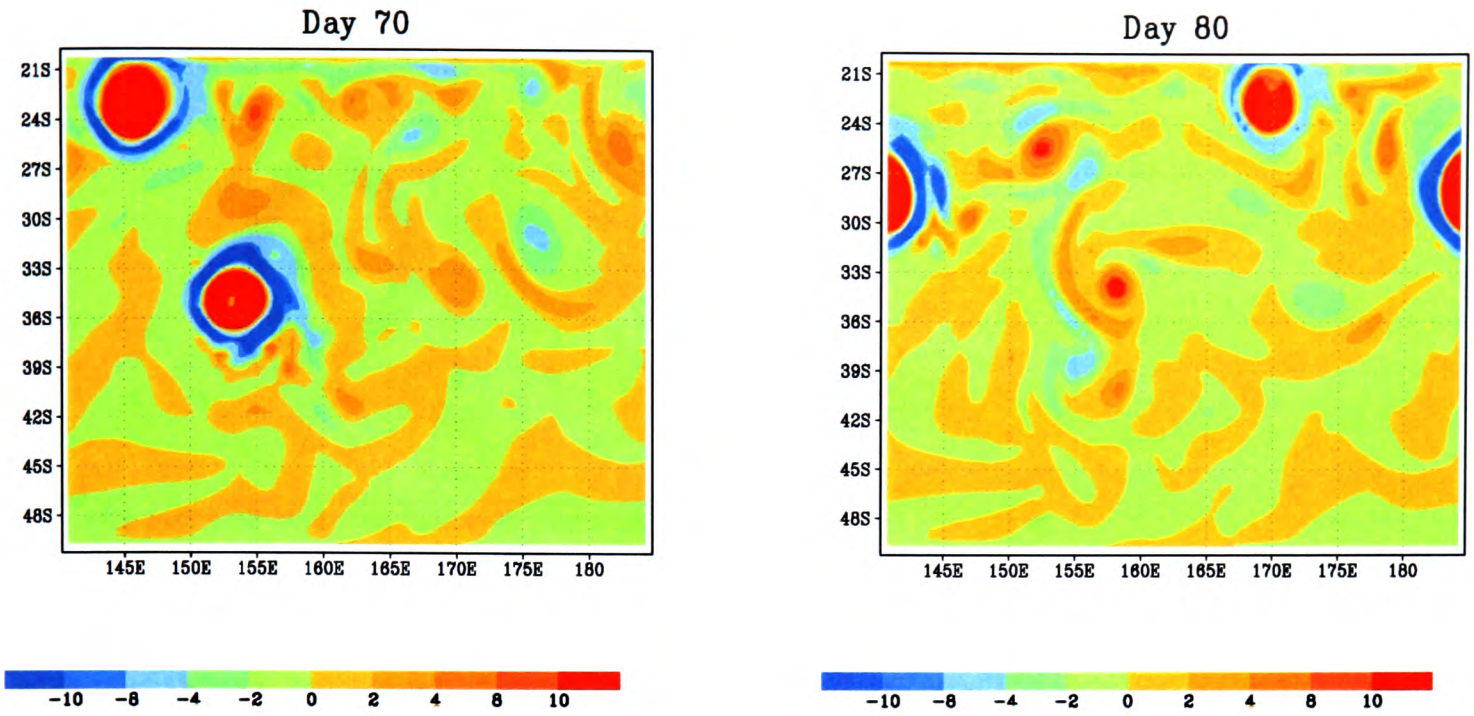


Figure 5.4: Evolution of the flow during days 70-100, shown by plots of relative vorticity ($\times 10^5$) at 56 mb, and 10-day intervals. (The contour interval is chosen to reveal the weaker features, and saturates within the strongest vortices). On day 70, only two large vortices are evident. The small features are not clearly distinguishable. By day 80, two minor vortices are visible. The one at (157°E,34°S) strengthens up to day 100, migrating equatorwards, to 34°S by day 100. By this stage the initial vortex has dissipated but the second remains.

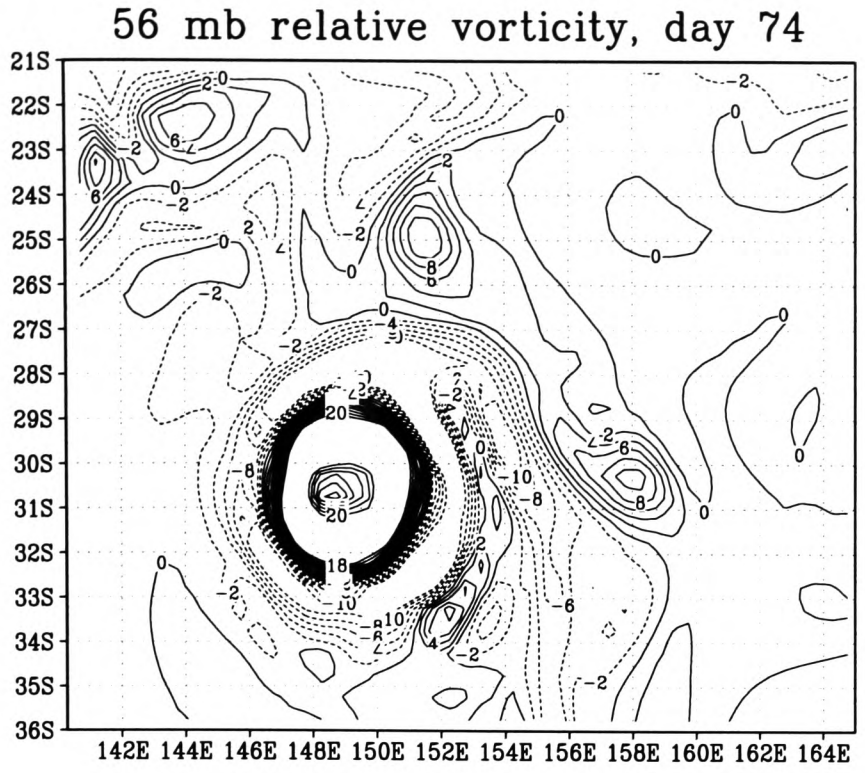


Figure 5.5: Initial appearance of a third vortex.

5.3.3 A shallow cyclone in midlatitudes

The experiment of Section 5.3.2 was repeated with a vortex of identical structure but opposite sign, i.e. cyclonic. In potential vorticity terms, the background PV gradient encourages a negative PV anomaly to drift in the opposite meridional direction to a positive (anticyclonic) one. We therefore expect the cyclone to migrate westward and poleward. The vortex was placed further north in the channel than usual, at 30°S, to allow it to migrate southwards.

The evolution of the 56 mb temperature field, over 40 days,is shown in Fig. 5.6. Its meridional migration rate, of about two degrees in 40 days, is similar to that of the anticyclone (Fig. 5.2). The cyclone changes little in size, but becomes considerably weaker. The peak windspeed is only 30 m/s (about half the original) after 40 days. After 60 days it has weakened to the point where it is only just visible. There is no evidence of the vortex shrinking and strengthening, in contrast to the anticyclonic case. It is tempting to see this as evidence of a cyclone-anticyclone asymmetry as observed on Jupiter. However, this is only one limited experiment and in an unrealistic (zero) background flow. We do not wish to draw hasty conclusions, but note that it would be interesting to repeat many of our later experiments with cyclones instead of anticyclones.

5.3.4 Shallow anticyclones with arbitrary vertical structure

In both the previous cases, the vortices were initialised with a normal-mode vertical structure. We expect that a flow comprising only two or three modes would have a relatively simple evolution, with little mixing into other modes. This was borne out by the results:

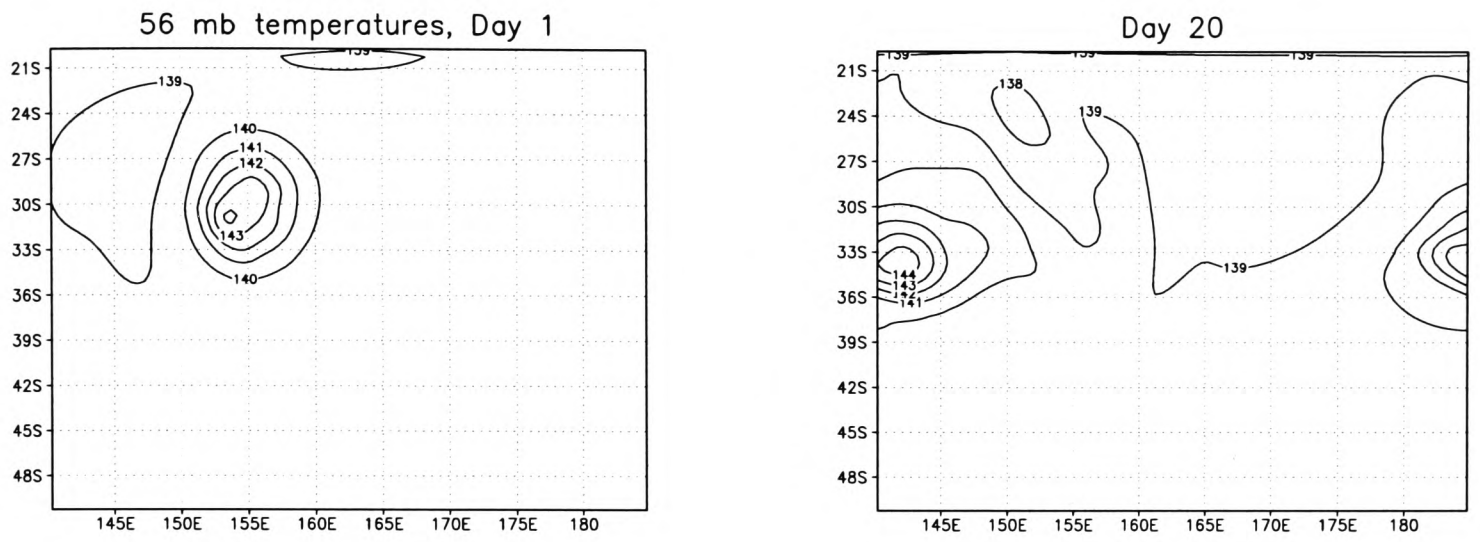


Figure 5.6: Evolution of a midlatitude cyclone. The vortex migrates westward and also poleward, but much more slowly than a similar anticyclone (Fig. 5.2)

eddies with this structure retained their vertical profiles, even when splitting into smaller daughter eddies. The results are different for eddies with an arbitrary vertical structure. The normal-mode profile in Fig. 5.1 resembles a Gaussian, but is asymmetric, with a wide "tail" into deeper pressure levels. We therefore chose Gaussian structures for our arbitrary structure functions. These may be placed at any convenient level in the atmosphere. The first case (Fig. 5.7) features a function designed to resemble closely the profile used in Section 5.3.2. It peaks at the 56 mb level but is of course symmetric in log-pressure height.

The results show that the vortex does keep its original shape, remaining close to the normal-mode profile. Its evolution is actually simpler than that in Section 5.3.2. The vortex merely migrates equatorward until it reaches the sidewall. There is no strengthening and no daughter eddies are produced.

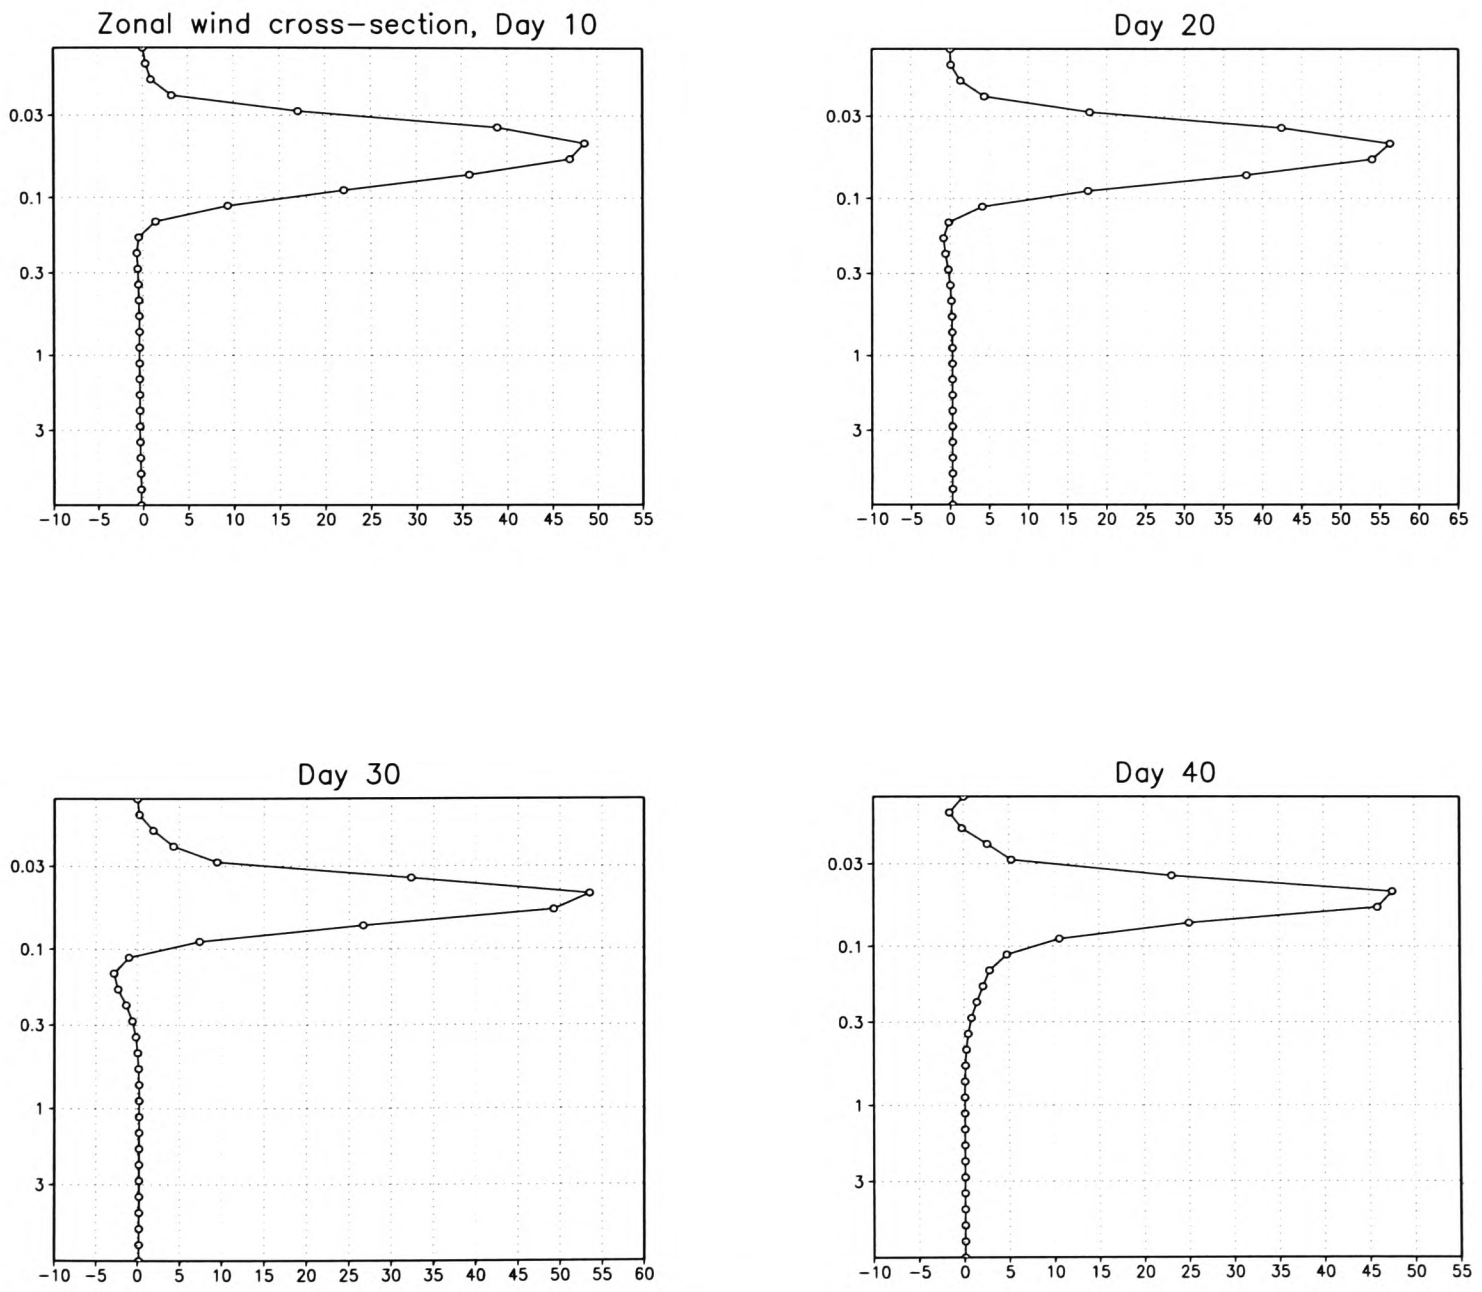


Figure 5.7: Profiles of zonal wind at the peak of a vortex (with Gaussian vertical structure), at 10-day intervals. The vertical axis is in bars and the horizontal in m/s. The shape is largely retained because this structure is very close to the normal-mode structure of Section 5.3.2.

To examine the behaviour of a vortex of more arbitrary profile, a Gaussian structure

peaking at 450 mb was chosen. The half-width was chosen so that the vortex was approximately as thick as in the previous case. The time evolution of the profile is shown in Fig. 5.8, and shows the vortex strengthening, becoming thinner and also moving higher in the atmosphere. It seems to be adjusting towards a profile closer to the original normal-mode superposition chosen for the previous experiments.

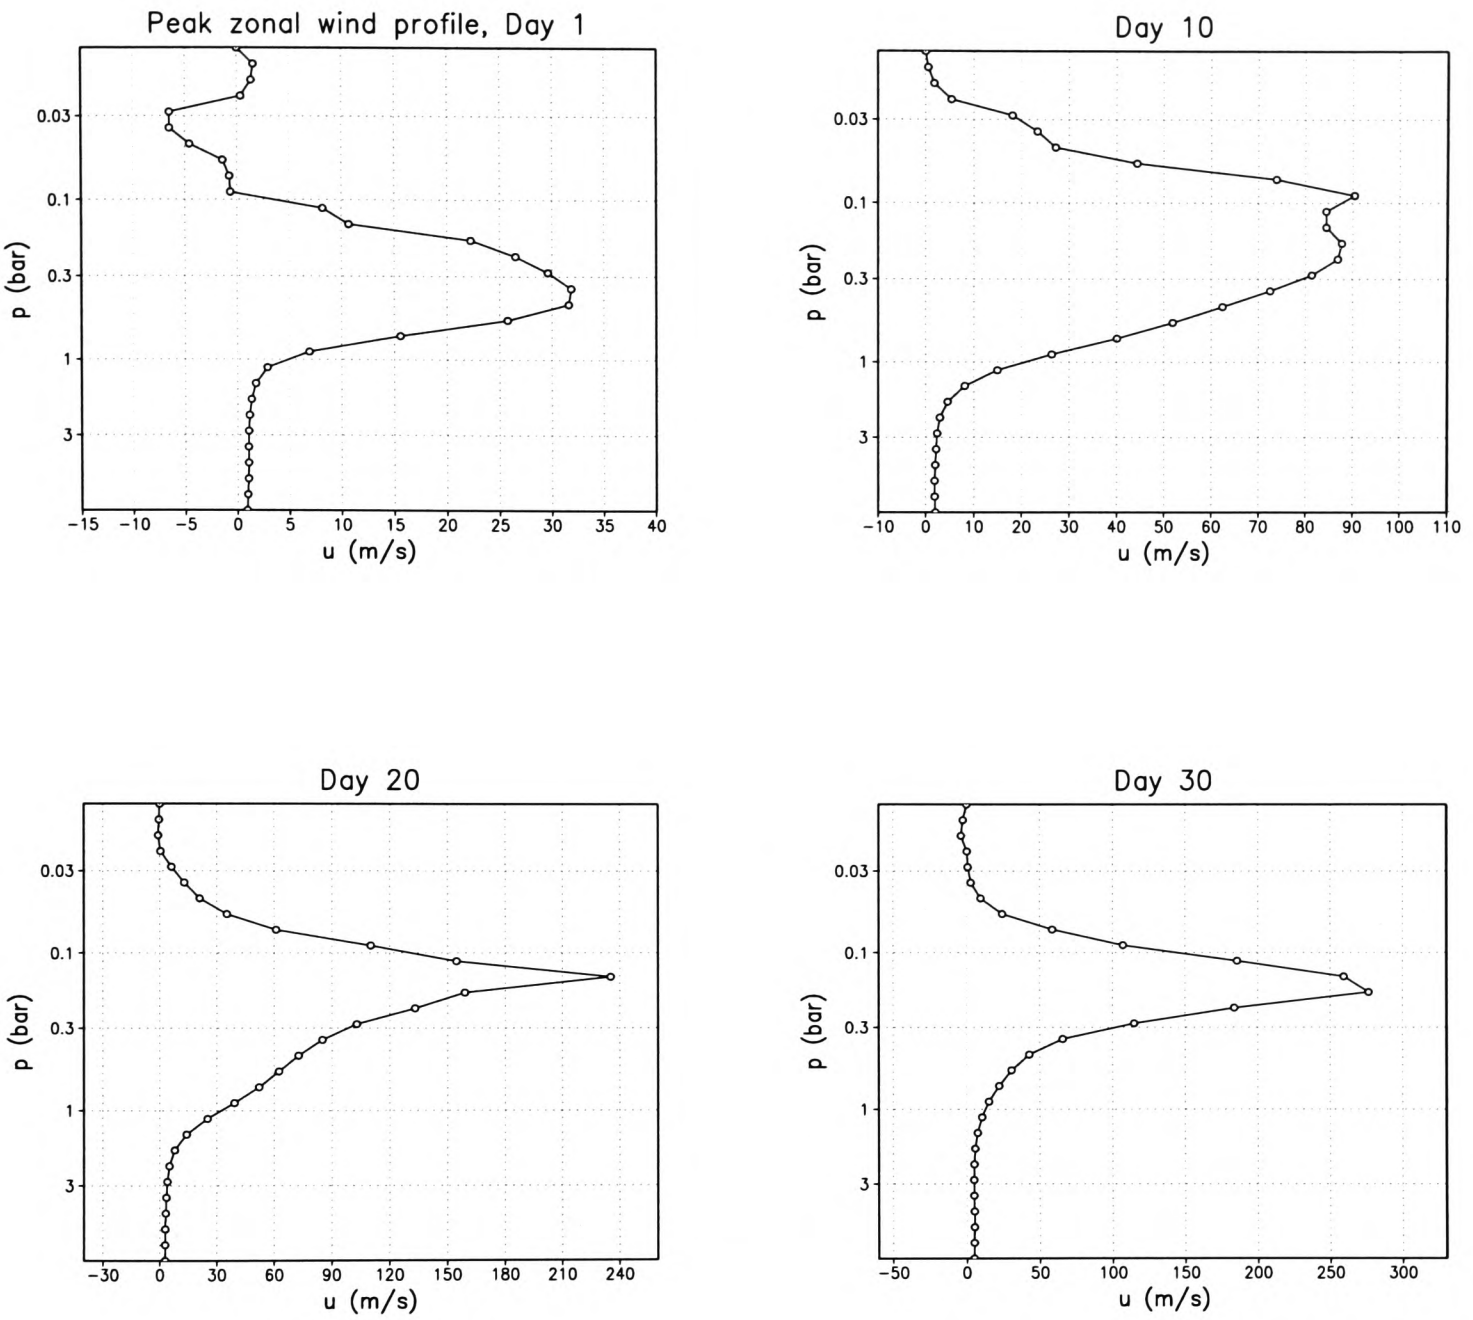


Figure 5.8: Profiles of zonal wind at the peak of a vortex initially centred on level 14 (450 mb). The vortex strengthens and becomes thinner, while moving to a higher level in the atmosphere.

A slight modification to this experiment used a Gaussian peaking at the same height but with a much narrower half-width. This initially gave noisier wind and vorticity fields but eventually showed similar behaviour: the vortex increased in speed and moved upward slightly. The vortex grew (and migrated) more slowly than in the previous case, reaching 37° afer 40 days.

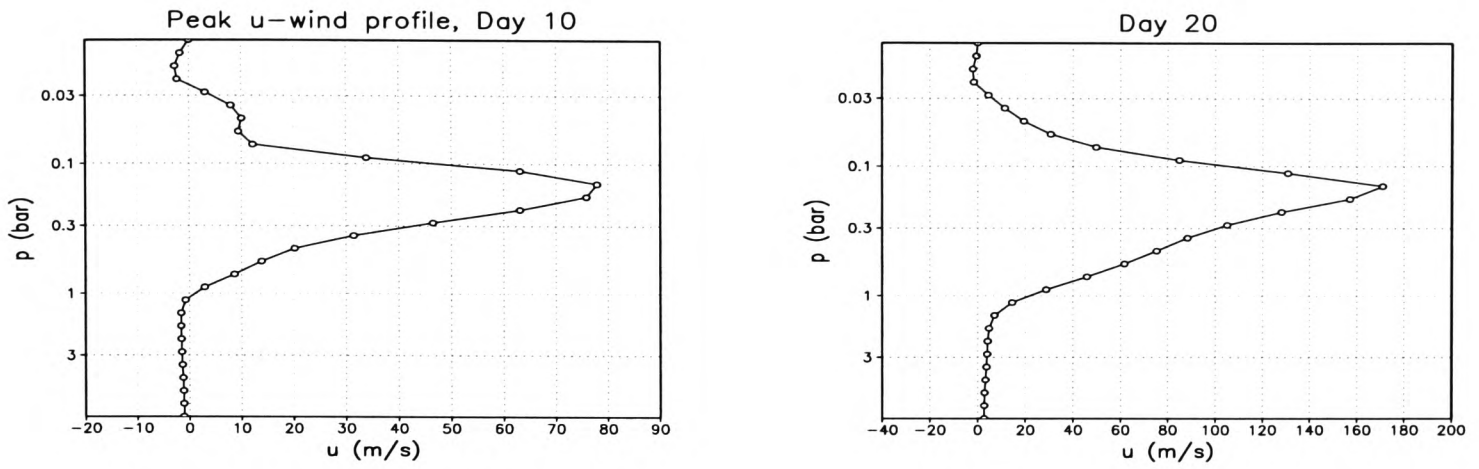


Figure 5.9: Profiles of zonal wind at the peak of a vortex initially centred on level 14 (450 mb). This structure function was much narrower than that in Fig. 5.8, confining the vortex more closely near the 450 mb level

5.3.5 A deep anticyclone

So far we have looked at the properties of shallow, trapped eddies, modelled with boundary conditions damping the flow at the top and bottom of the model. We shall now extend this to include deep eddies, and compare their behaviour with that of shallow ones. The potential vorticity inversion study of Chapter 2 suggested that deep eddies might appear indistinguishable from shallow ones, with very similar stability, if observed only above the clouds. We now test this prediction in the time-dependent, 3D GCM by following their evolution in time.

The anticyclone was initialized with an arbitrary vertical structure: constant velocity up to around 300 mb (i.e. near the cloud-tops), then a linear decrease to zero at the top boundary. This structure can actually be obtained by a superposition of the barotropic and first baroclinic modes in the undamped (no sponge layer) case. Of course in the GCM there is damping, but it appears that this structure is nevertheless close to a natural mode of the system. Later we shall show how eddies of this type are generated spontaneously by instability of a jet with the same structure. However, spontaneous eddies with this deep structure only appear when there is deep forcing.

The bottom boundary condition is barotropic: on the lowest level, the temperatures are constrained to be isothermal and the wind-shear to be zero. The winds on the lowest level are set equal to the winds on the second level; with zero horizontal temperature gradient this forces the vertical shear to be zero. Convective adjustment and vertical diffusion of winds should remove any numerical discontiuities introduced by the boundary adjustments. The horizontal shape is the same as that of the shallow vortex of Section 5.3.2

The results (Fig. 5.10) show that, although the eddy shrinks as it rapidly migrates, it does not become stronger. The 500 mb winds remain at 60 m/s, while below this level, the vortex decays to only half this speed. This suggests that a deep vortex is less stable than a shallow one. This vortex also generates a spontaneous daughter, but interestingly, its structure is very like the shallow, stratospherically confined flows seen earlier—i.e. the first damped baroclinic mode. The most likely explanation is that the vortex adjusts to the gravest (largest-scale) vertical mode available. The barotropic mode is ruled out by the sponge-layer damping, so the new vortex adopts the first baroclinic mode, as shown in Figs. 4.11 and 4.13.

5.3.6 Summary: effects of varying the vertical structure

We have seen that a cyclone whose vertical structure consists of a two-eigenmode superposition migrates polewards, without changing very much. An anticyclone of the same form migrates equatorwards, shrinking, steepening and increasing in speed. The rate of

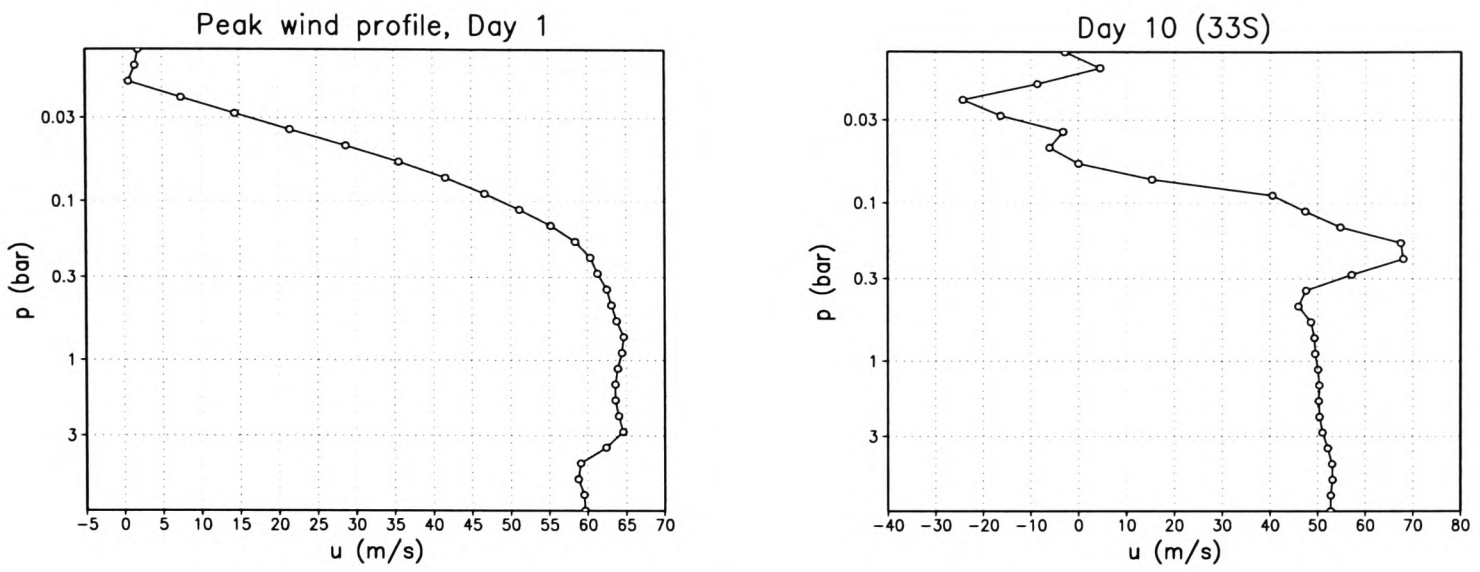


Figure 5.10: Profiles of zonal wind at the peak of a deep anticyclone. The first three plots show the development of the original vortex. It migrates equatorward more rapidly than the shallow flows. The cloud-top winds remain of similar speed, but below the clouds the flow decays, suggesting that the original structure is unstable. After 20 days the eddy has reached the northern sidewall. The final plot shows a new, spontaneously generated eddy, which is shallow and confined to the stratosphere.

migration increases as the vortex steepens. Eventually, smaller eddies are projected polewards, later following the same life cycle as the parent eddy. An anticyclone of Gaussian vertical profile, but resembling the previous one, retains its shape but does not steepen. A vortex centred deeper in the atmosphere moves upwards and adjusts towards the eigenmode form. A deep anticyclone, with barotropic flow below the clouds, migrates in a similar way to the others, changing its vertical profile slightly, and weakening. It also produces daughter eddies but these have the baroclinic structure of Fig. 5.1, not that of the parent eddy.

We conclude that there is a natural scale and position of eddy favoured by this model configuration: a small circular eddy, of radius order 3 degrees, and confined in the vertical around the 56 mb level. The wind speeds are rather high, up to 250 m/s. Any experiment (except for the case of the Gaussian centred on 56 mb) which introduces a vorticity anomaly in this region will lead to these eddies being generated. This result raises a number of questions. What causes the vortices to shrink and spin up? Why are all the spontaneous vortices thin, and confined to the middle stratosphere with a structure very like the first baroclinic mode? Do the results depend on the initial shape of the vortex? What is the sensitivity of these results to the main model parameters, such as background temperature profile and time-constants for radiative damping and diffusion? We shall address the final question first.

5.3.7 Dependence of migration rate on thermal damping

The runs discussed so far used a Newtonian cooling coefficient of 1.08×10^{-8} per hour, in line with the estimate of Conrath et al. (1990). It is possible that this figure is unsuitable for use in the present model. It assumes purely radiative thermal adjustment and takes no account of sub-grid scale mixing processes. In the Earth's oceans, the effective eddy diffusivity (due to unresolved eddies) can be much greater than the kinematic diffusivity. We expect similar considerations to apply to Jupiter's atmosphere, which in this respect resembles the Earth's oceans more closely than its atmosphere.

Some experiments were repeated with the Newtonian relaxation coefficient increased by a factor of 100. This had no significant effect on the evolution and migration of the vortex. Newtonian relaxation is not scale-selective, but operates at each point independently. This result therefore suggests that the steepening and strengthening process is scale dependent, probably linked to the inverse cascade of enstrophy observed in 2-D turbulence. (The latter is discussed for example by Pedlosky (1982)). We can test this hypothesis by examining the effect of a scale-selective model parameter, namely the horizontal diffusion. First, however, we shall digress to discuss two other factors—the shape of the vortex and its initialization method.

5.3.8 Vortices of different horizontal shape

The fact that all the vortices tend to steepen and form fronts suggests that a 'top-hat' PV profile could be a more stable initial shape. Without attempting to invert the Ertel PV in spherical geometry, we crudely approximated this shape by considering the quasi-geostrophic PV q, on an f-plane. Now in the separated horizontal equation, $q \propto \nabla^2\psi$, so a paraboloidal streamfunction will give constant PV in the region where $\psi \neq 0$. Replacing the Gaussian with a paraboloid and setting $\psi = 0$ elsewhere yields the vortex shown in Fig. 5.11.

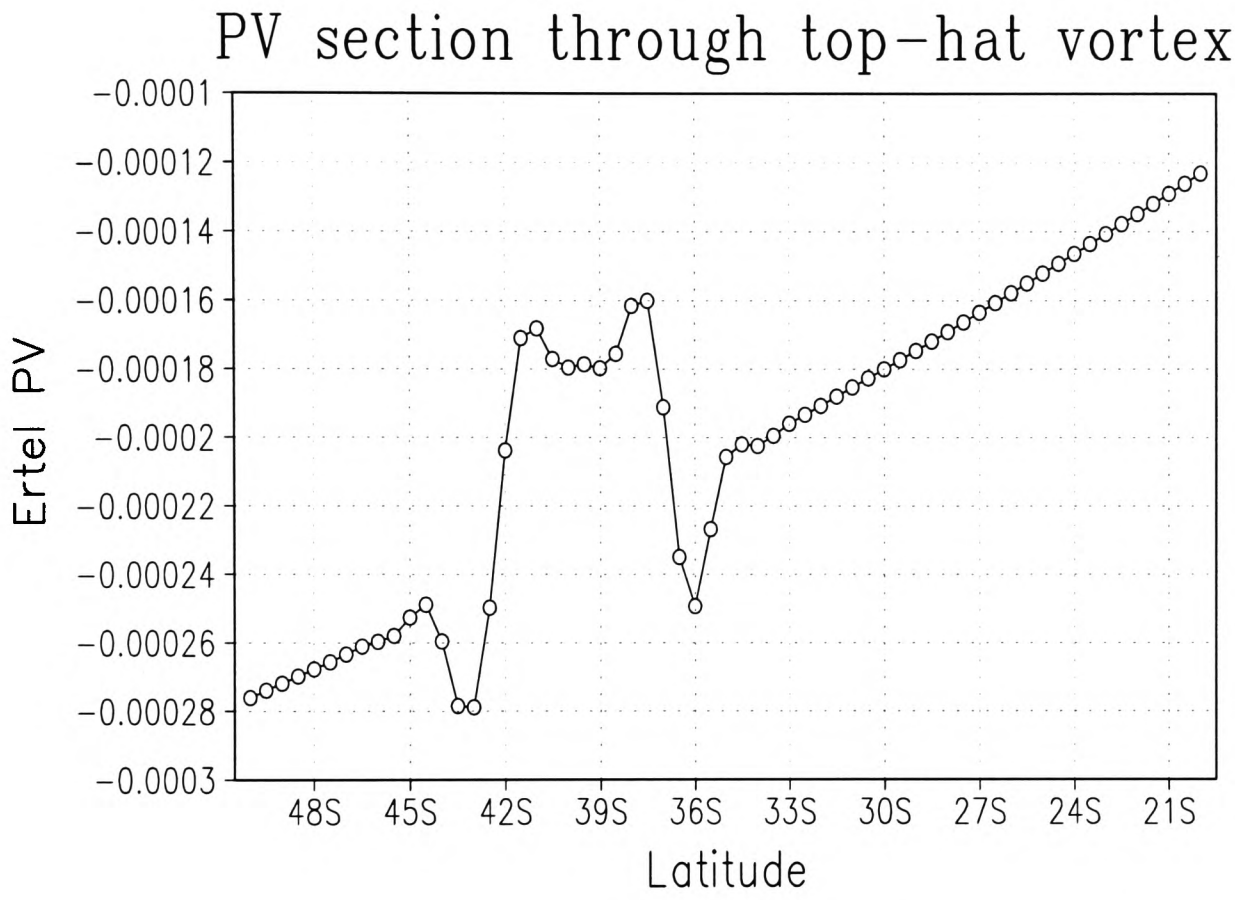


Figure 5.11: Longitudinal cross-section through the centre of a top-hat vortex, after 2 days. The field shown is PV at 56 mb, the peak level of the flow. In the absence of a background flow, the ambient PV gradient is due to planetary vorticity.

The results of this experiment (Fig. 5.12) are interesting: the top-hat shape is less stable than expected. The cyclonic collar is thrown off in the first few days, curling up into daughter cyclone lying to the southeast of the original. The main vortex then shrinks as usual, losing its initial shape. By day 40 the vorticity structure has become a sharp, roughly Gaussian peak. Equatorward migration has begun, although the rate is slower than in previous cases so we could argue that this eddy is longer lived. This problem of adjustment from top-hat to Gaussian was noted by LeBeau and Dowling (1998), in the context of a model of comparable realism but applied to Neptune's Great Dark Spot. In the concluding part of this chapter, we shall try to explain these results in terms of simple ideas of frontogenesis.

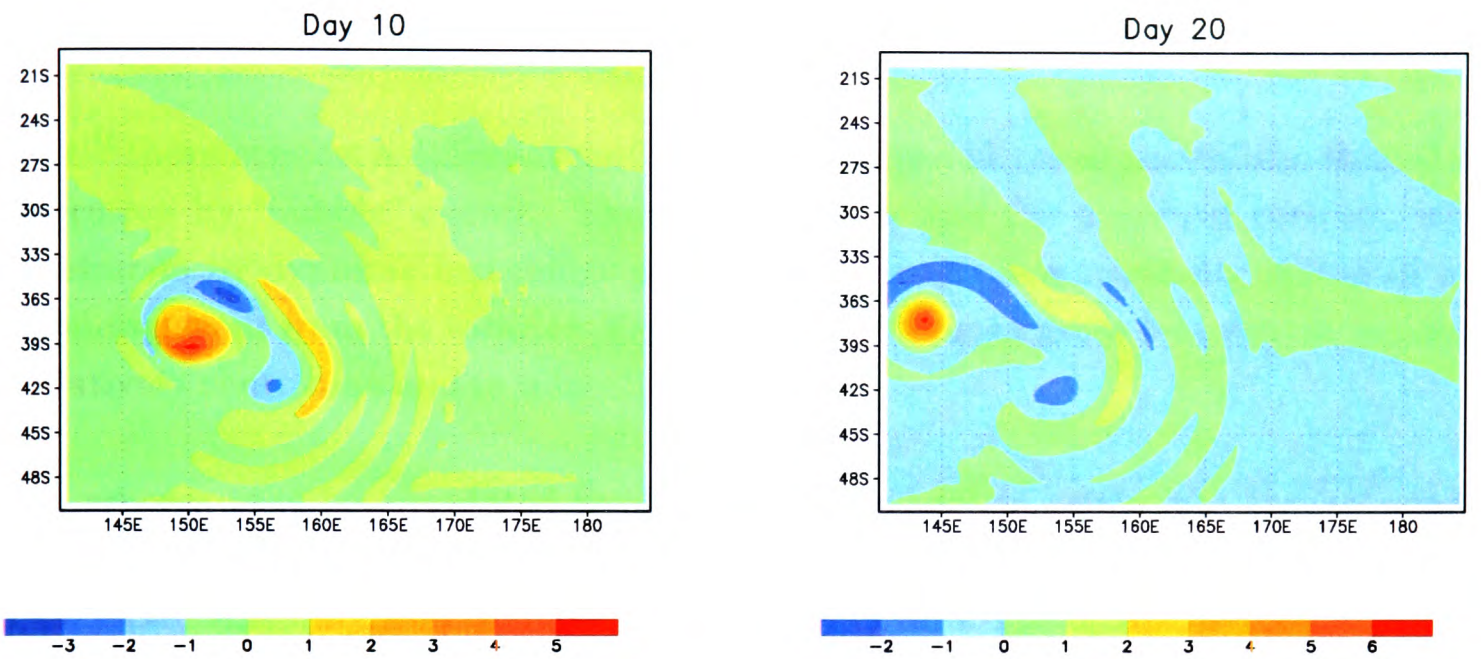


Figure 5.12: Migration of a 'top-hat' vortex over 40 days, shown by plots of relative vorticity ($\times 10^5$) at 56 mb, and 10-day intervals. The initial top-hat shape shrinks to an approximately Gaussian shape as the vortex intensifies.

5.3.9 Vortex initialization by heating

Later we will need to address the problem of initializing balanced eddies in a background flow with (as on Jupiter) both horizontal and vertical shear. The method of imposing approximate, quasi-geostrophic wind and temperature fields is less successful when the basic state is in motion. A linear superposition of the vortex on the background is not necessarily balanced under the primitive equations. Also the results are easier to interpret if the initialization method is not too sensitive to the particular basic state chosen for each experiment.

We will therefore test a different method, inspired by a physical possibility: that of vortex generation by "bubble" events. The hypothesis is that the observed vortices, whether maintained by dynamic instability or otherwise, may have originally arisen in a huge upwelling event from the interior. This would have caused great latent-heat release over a relatively short time and in a localised area.

This process may be simulated by allowing the model to settle into the axisymmetric basic state, then applying local heating for just one or two days. The heating coefficient is increased smoothly, and then decreased again to zero. The spatial pattern is chosen for convenience: a Gaussian spheroid (circular in the horizontal) of heating localized around the level of the cloud-tops (500 mb), with another spheroid of cooling just above it. This vertical dipole structure spins up an anticyclone confined to the region between the two spheroids. The relatively slow growth of the eddy (96 timesteps compared with just one) reduces wave radiation and leads to a balanced flow. When the heating stops after one day, the vortex evolves freely.

This method was first tested in a basic state at rest, for comparison with the results of the previous section. The results (Fig. 5.13) show that, although of similar size and shape to the previous examples, the vortex migrates rather more slowly, reaching the northern boundary only after about 80 days. The vertical structure (Fig. 5.14) is also preserved. This vortex seems to be more well-balanced initially, so it steepens (and migrates) more slowly. In retrospect this is reasonable, for we have constrained only the temperature field, while allowing the winds to adjust. Direct application of the thermal-wind and geostrophic approximations have not been necessary. In Section 5.4 we shall see how this "bubble" method provides a simple, effective way of initializing vortices in a background flow.

5.3.10 Dependence on horizontal diffusion

The results so far lead us to believe that the vortex behaviour is largely insensitive to non scale-selective processes, because these cannot prevent steepening of PV gradients.

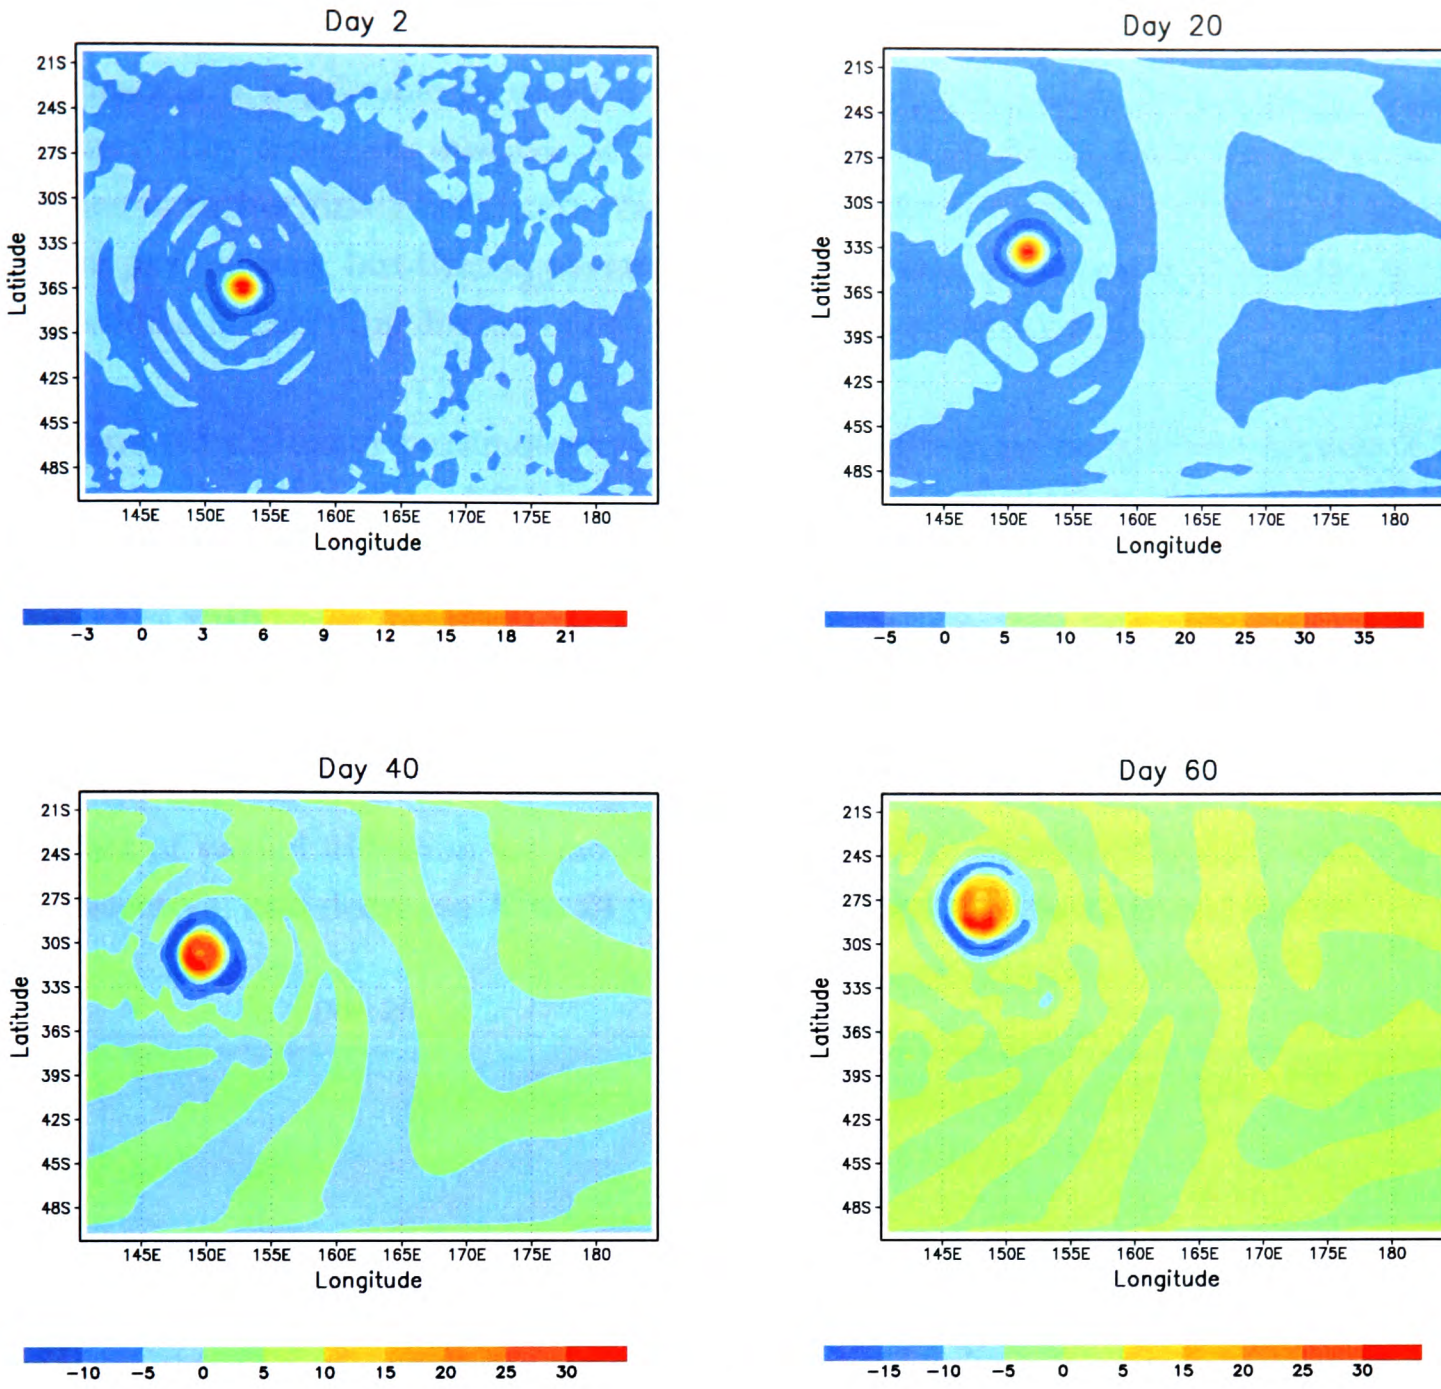


Figure 5.13: Migration of a thermally-induced vortex (in a resting background state) over 60 days, shown by plots of relative vorticity ($\times 10^5$) at 225 mb, at 20-day intervals. The vortex grows, and migrates, slightly slower than in the previous cases.

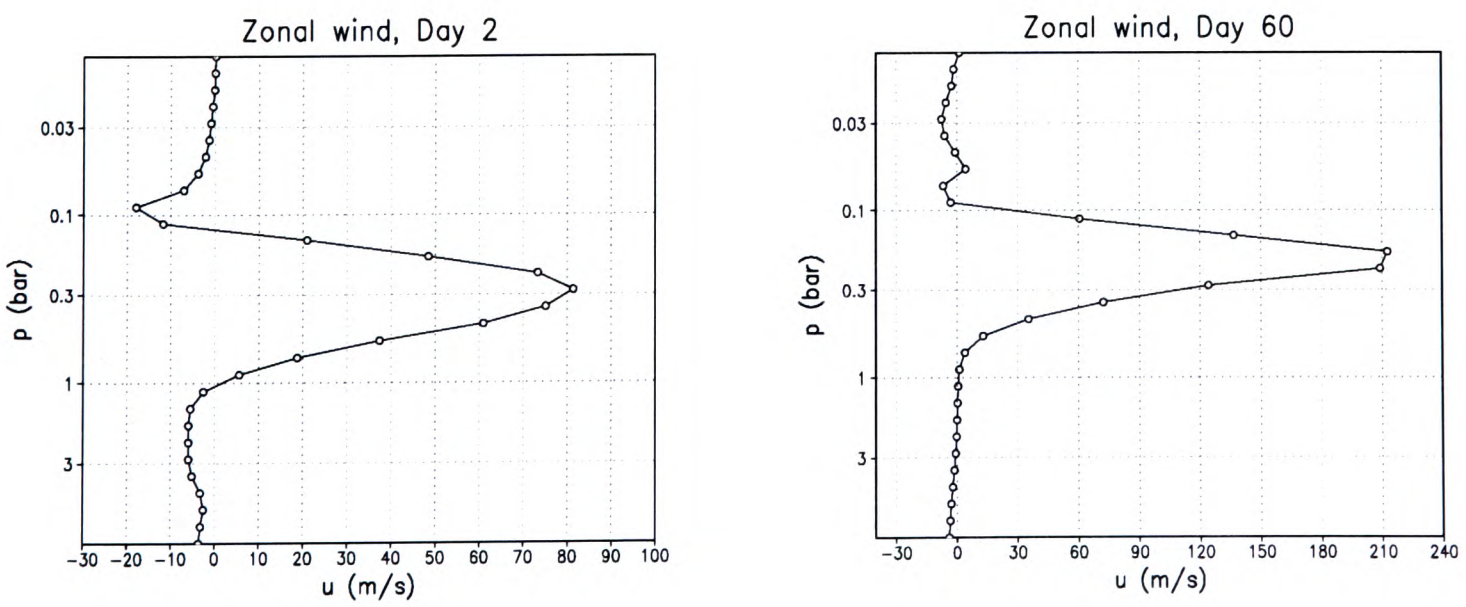


Figure 5.14: Initial and final cross-sections through the vortex, showing that the vertical structure is largely retained, but strengthened.

However, the model is sensitive to changes in horizontal diffusion which is scale selective. The results of varying the diffusion parameters will be presented below. By varying the timescale it is possible to move from a regime where eddies grow uncontrollably, to one where they dissipate quickly. It is unfortunate that the equilibrium between these extremes is rather finely balanced. We would hope for a robust result that is valid for a range of parameters, but that apparently is not the case. However, it is possible to isolate single eddies resembling Jovian ones over periods of several months.

The standard diffusion parameters used in the half-degree model (see Section 3.3) are $N = 48$, with third-order diffusion; this gives $K = 2.8 \times 10^9$ with a damping time τ of 12 hours for a two-gridpoint wavelength. Figure 5.15 shows the results of a control experiment with these parameters. The vortex is initialized by direct imposition of winds and temperatures, not by the localized-heating method. To simplify the experiment, the initial vortex is rather smaller than we have used before, so its size changes much less during the migration. Even this small vortex reaches the northern wall by day 40. To see the effect of varied diffusion we move to a fourth order "hyperdiffusion", which is more scale-selective, and decrease N to 24 ($\tau = 6$ hours) and then to 12 ($\tau = 3$ hours).

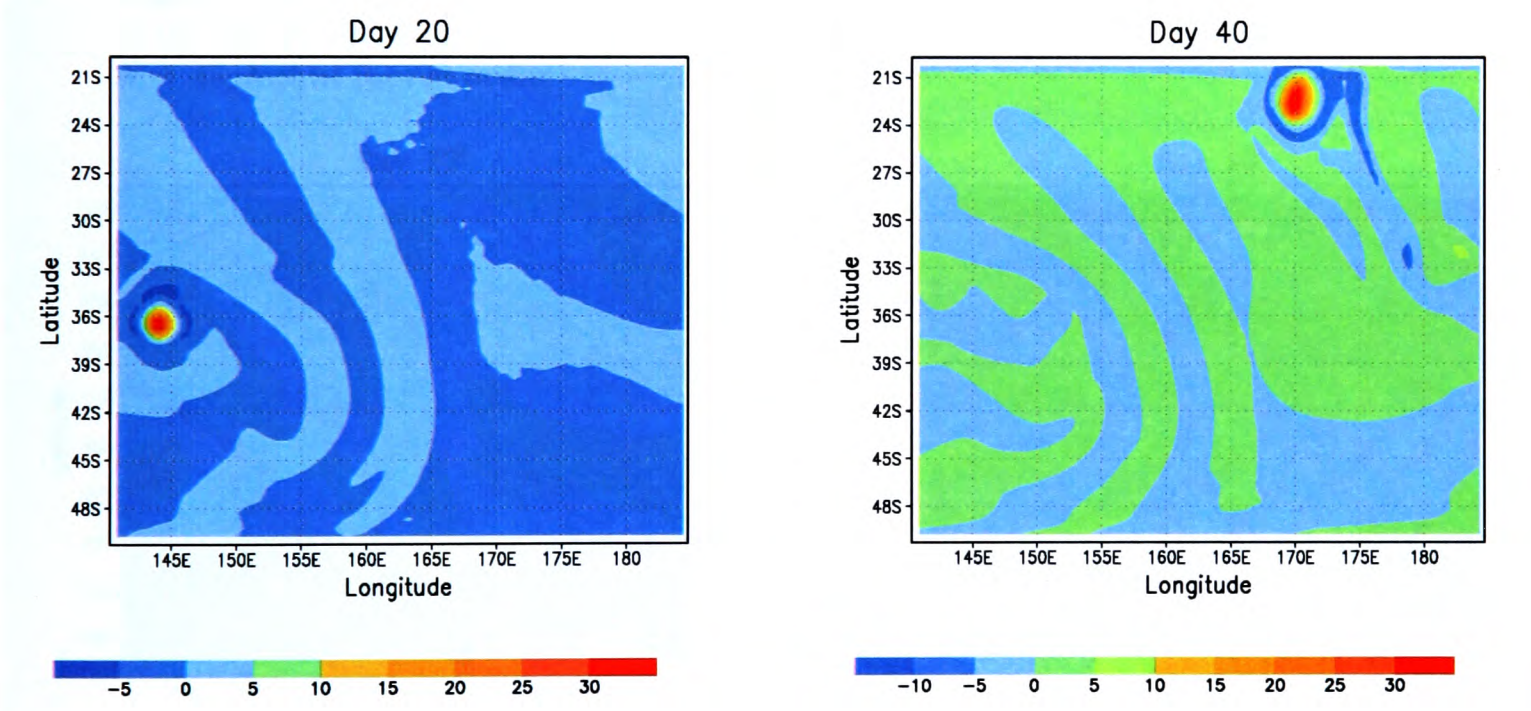


Figure 5.15: Migration of the vortex over 40 days, shown by plots of relative vorticity ($\times 10^5$) at 56 mb, after 20 and 40 days. This is a control experiment with the usual third-order diffusion but a smaller vortex.

The results for $N = 48$ and $N = 12$ are shown in Figs. 5.16 and 5.17 respectively. Moving to fourth-order has clearly been beneficial; the migration is much slower although N is unchanged. The effect of changing the damping time to $\tau = 3$ hours, on migration and growth is not overwhelming so we continue to reduce the damping time.

At this point a further decrease of N leads the model to crash, so the next run is conducted with the same damping value $N = 12$ but the timestep halved to $\Delta t = 450$ s. This gives a damping time of $\tau = 1.5$ hours, and the results (Fig. 5.18) are significantly different. The migration rate hardly changes but the growth rate is much smaller. The vorticity peak

only intensifies by about 30%, instead of more than 400% as before. Another run with $\tau = 1$ hour has the same result, so it appears that further reductions in damping time are not effective. In later runs we stick to $\tau = 1.5$ hours, to keep diffusion as light as possible.

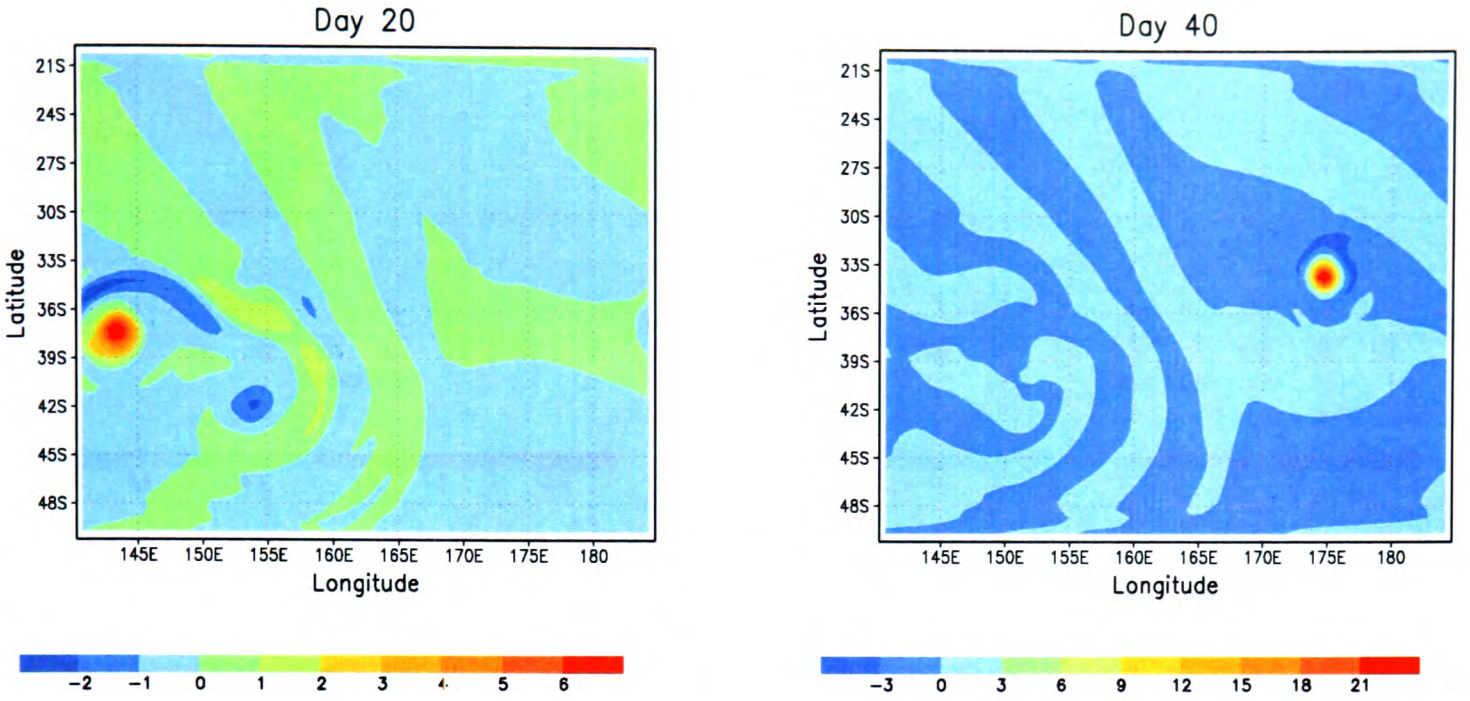


Figure 5.16: Migration of the vortex over 40 days, shown by plots of relative vorticity ($\times 10^5$) at 56 mb, after 20 and 40 days. The diffusion parameter was set at the usual $N = 48$, but at fourth order.

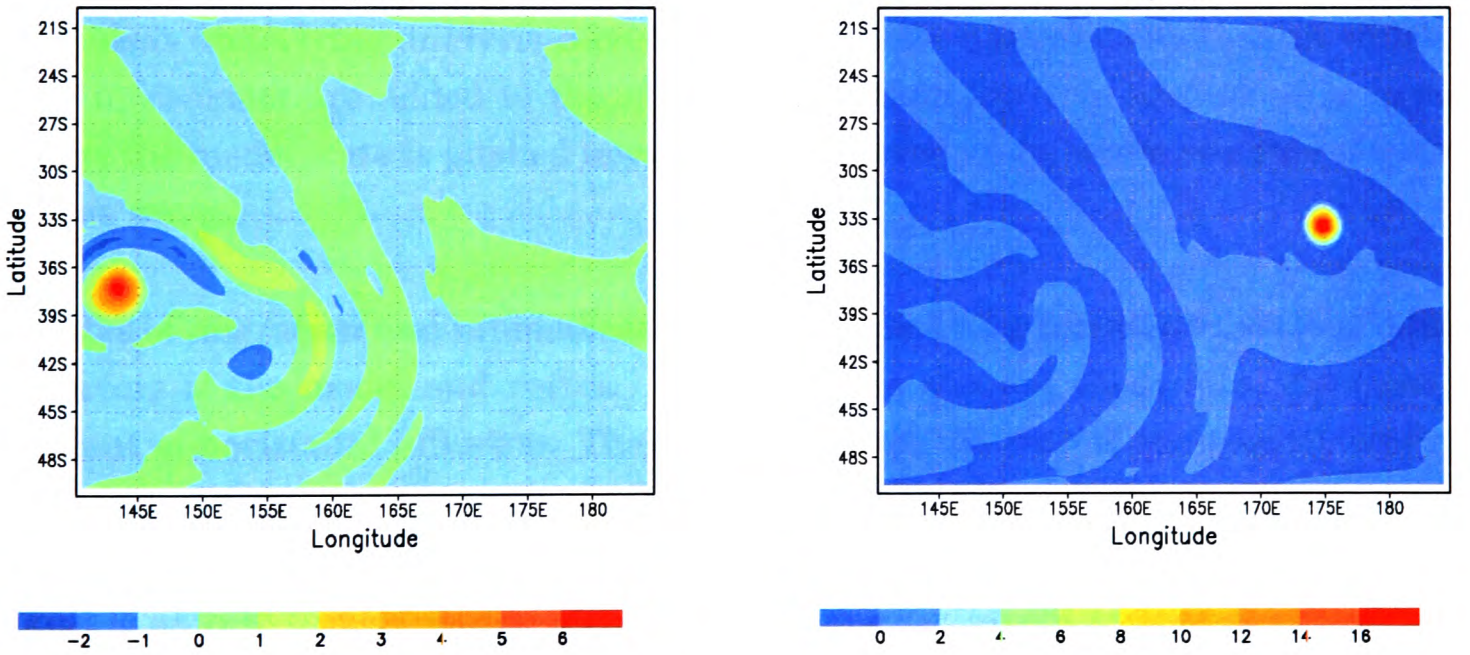


Figure 5.17: Migration of the vortex over 40 days, shown by plots of relative vorticity ($\times 10^5$) at 56 mb, after 20 and 40 days. The diffusion parameter was decreased to $N = 12$, at fourth order.

5.3.11 Conclusions: vortices in an atmosphere at rest

We have conducted model experiments on isolated vortices in a resting atmosphere. Two alternative methods of vortex initialization were used. One was pre-calculation of the geostrophic winds, followed by their instantaneous imposition. The other used distributed

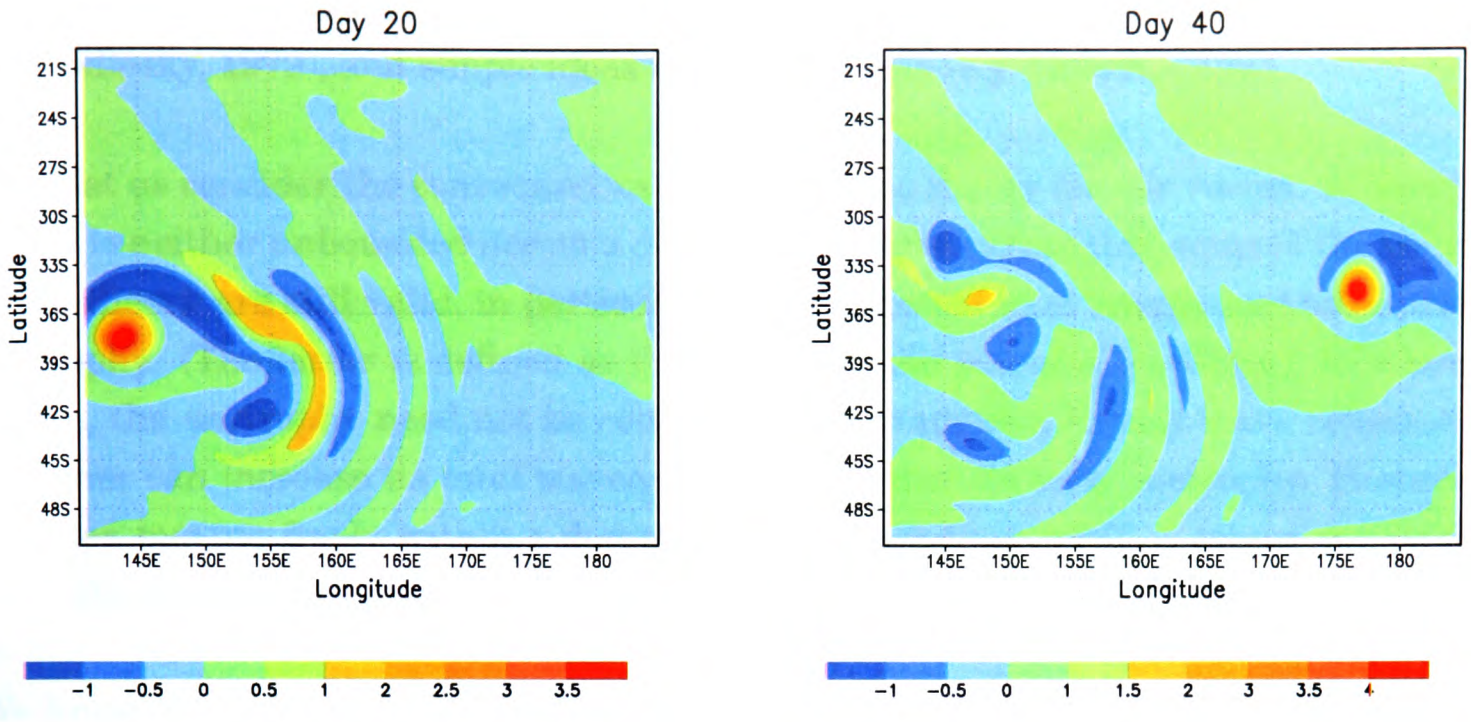


Figure 5.18: As Figs. 5.16 and 5.17 but with $N = 12$ and the timestep halved. The migration rate has slowed slightly, but more importantly the vortex has not intensified significantly.

heating below, and cooling above, the vortex. The heating rate was increased smoothly, then decreased again over the period of a few hours. The former method produces vortices with a more easily controlled shape, but the latter ensures a more accurate balance, because only temperature is forced, while the winds adjust freely. This method is based on the ideas underlying analysis-correction schemes for data assimilation, in which necessary increments are added to the model fields gradually over several time steps. In this way the model state is pushed towards the observed state, while avoiding shocks and spurious waves.

The rates of migration and intensification were studied as a function of several different parameters of the model and vortex. The most important of these were the order and coefficient of horizontal diffusion. These factors affect the rate of steepening, whilst the Newtonian cooling rate does not. Diffusion is a scale-selective process, unlike thermal relaxation. The modelled intensification is actually a process of steepening, which we will suggest is akin to frontogenesis. This means the horizontal shear scale is changing, so it is logical that scale-selective damping will be effective in controlling it. In this section we discuss the theoretical ideas which can help us qualitatively interpret the results.

We have seen that all vortices initialized in a resting atmosphere migrate meridionally. The migration direction depends on the sign of the vortex PV anomaly, and the planetary potential vorticity gradient. Each vortex attempts to minimize its PV anomaly by migrating towards a latitude where the ambient PV is closer to that in the vortex. This means equatorward travel for anticyclones, and poleward for cyclones.

Almost all vortices in our model shrink, steepen and intensify in the horizontal. Meanwhile their vertical structure adjusts towards the first baroclinic normal mode. We can go

some way towards understanding this by use of the theory of an unbounded β-plane (see e.g. Pedlosky, 1979), and simple ideas of frontogenesis (e.g. Holton, 1992).

First let us consider the consequences of QG β-plane theory for our vortex. Although the model is neither unbounded nor on a β-plane, we might expect that some of the principles of the theory are still valid, in particular conservation of total energy and total potential enstrophy. (The latter is defined as the square of the potential vorticity.) In a bounded region, the enstrophy need not be conserved: for example a Rossby wave reflected from a barrier can increase its total wavenumber and hence its total enstrophy. However, for a vortex moving freely within a domain much larger than itself, it seems reasonable to treat the domain as unbounded.

We know that some turbulent dissipation, however weak, must be acting in the atmosphere. It is essential in order to damp the large-scale motions and achieve a steady average state. Turbulence can be observed directly in a variety of features, from aircraft contrails to the wakes of the Jovian eddies. In our model, turbulence is represented by high-order hyperdiffusion. This smooths out features on the smallest resolvable scales, transferring energy to the sub-grid scale. A theorem due to Fjortoft shows that if energy is transferred to smaller scales, the conservation of enstrophy requires that a greater amount of energy be transferred in the opposite sense, to larger scales. For a derivation and discussion of this result, see Pedlosky (1982).

Under quasi-geostrophy, the upshot is that if energy is present on the large scales, it will tend to remain there. Therefore synoptic features are able to persist despite the energy cascade. However, this argument does not explain the steepening of the model vortex. We tentatively ascribe this to ageostrophic processes such as frontogenesis. The evidence for breakdown of geostrophy in the model is reasonably convincing. We mentioned in Section 5.4.3 that the Rossby number in the developed model vortices is usually at least 0.3. This is despite the fact that each vortex is initialized in geostrophic balance, and indicates that the flow does have an ageostophic component.

Frontogenesis is complicated to analyse, but up to a point the kinematics can be considered qualitatively. In a developing baroclinic wave, the geostrophic flow can affect the temperature gradient by means of two mechanisms: horizontal shear deformation and stretching deformation. These are quantified in QG theory by means of the Q-vector (Holton, 1992), which we will use to discuss qualitatively the effect on a vortex.

Consider a parcel of air moving into a shear zone. The parcel will tend to rotate as a result of shear vorticity, and also to deform as it is stretched (sheared out) along the shear direction, and shrunk along the perpendicular direction. If, in a developing eddy, the isotherms are not parallel to the isobars, horizontal shear deformation tends to rotate isotherms and concentrate them along the line of greatest shear. Similarly, stretching deformation will concentrate isotherms along the axis of deformation. These effects combine to intensify

the vortex by narrowing the shear zone.

The forcing term in the omega equation is sometimes known as the frontogenesis function. Hoskins and Pedder (1980) showed that a useful form for the frontogenesis function is the advective derivative of squared temperature gradient. This is related to the Q-vector by

$$\left(\frac{\partial}{\partial t} + \mathbf{v}_g \cdot \nabla\right) \mid \nabla_h T|^2 = 2\mathbf{Q} \cdot \nabla_h T \tag{5.1}$$

where $\mathbf{v}_g$ is the quasi-geostrophic wind vector, and ∇_h indicates the horizontal Laplacian. Also, the divergence of Q is related to the vertical velocity. This is of less relevance here, since the UM can diagnose ω with useful accuracy. But Eq. 5.1 implies that frontogenesis may be diagnosed by plotting Q-vectors superimposed on isotherms. Where the vectors cross isotherms we expect frontogenetic forcing.

We have done exactly this in Fig. 5.19, which shows a close-up view of a developing vortex. Although many of the Q-vectors are small or parallel with the isotherms, many do cross isotherms, particularly on the southern and eastern sides. This lends weight to our claim that the steepening process is a form of frontogenesis.

The steepening effect is also illustrated by the PV cross-section, which evolves from a smooth Gaussian to a sharp top-hat shape. The narrowing of the shear zone tends to increase the winds, because the isobars become closer. To test this interpretation, the obvious experiment was to initialize the model with a non-Gaussian 'top-hat' vortex. This was expected to be more stable against frontal steepening because it was already steep. What happens in practice is that the vortex splits almost immediately, throwing off the cyclonic collar and assuming a more Gaussian shape. It then evolves like the other Gaussian ones.

The result of this experiment is the opposite of the prediction. We speculate that the reason for this lies in the initialization method. It is probable that the initial state is not sufficiently balanced, which is why the vortex disintegrates so quickly. The surviving core is a Gaussian like the others.

In the vertical, in most cases the flow tends to adjust toward the first baroclinic mode. Vortices produced by spontaneous shedding from parents always appear to adopt this mode structure, even if the parent vortex had the deep-barotropic form.

On Jupiter the observed eddies do not migrate rapidly, as in these models; they remain at the same latitude for decades or centuries. PV considerations suggest at least one reason for this, which is that the alternating jetstream pattern modifies the ambient PV gradient and introduces a local PV maximum which prevents migration. The eddy is localized at a latitude where $\partial q/\partial y = 0$. In the second part of this chapter, we shall look at vortices embedded in a simple zonal flow and compare the results with those of the

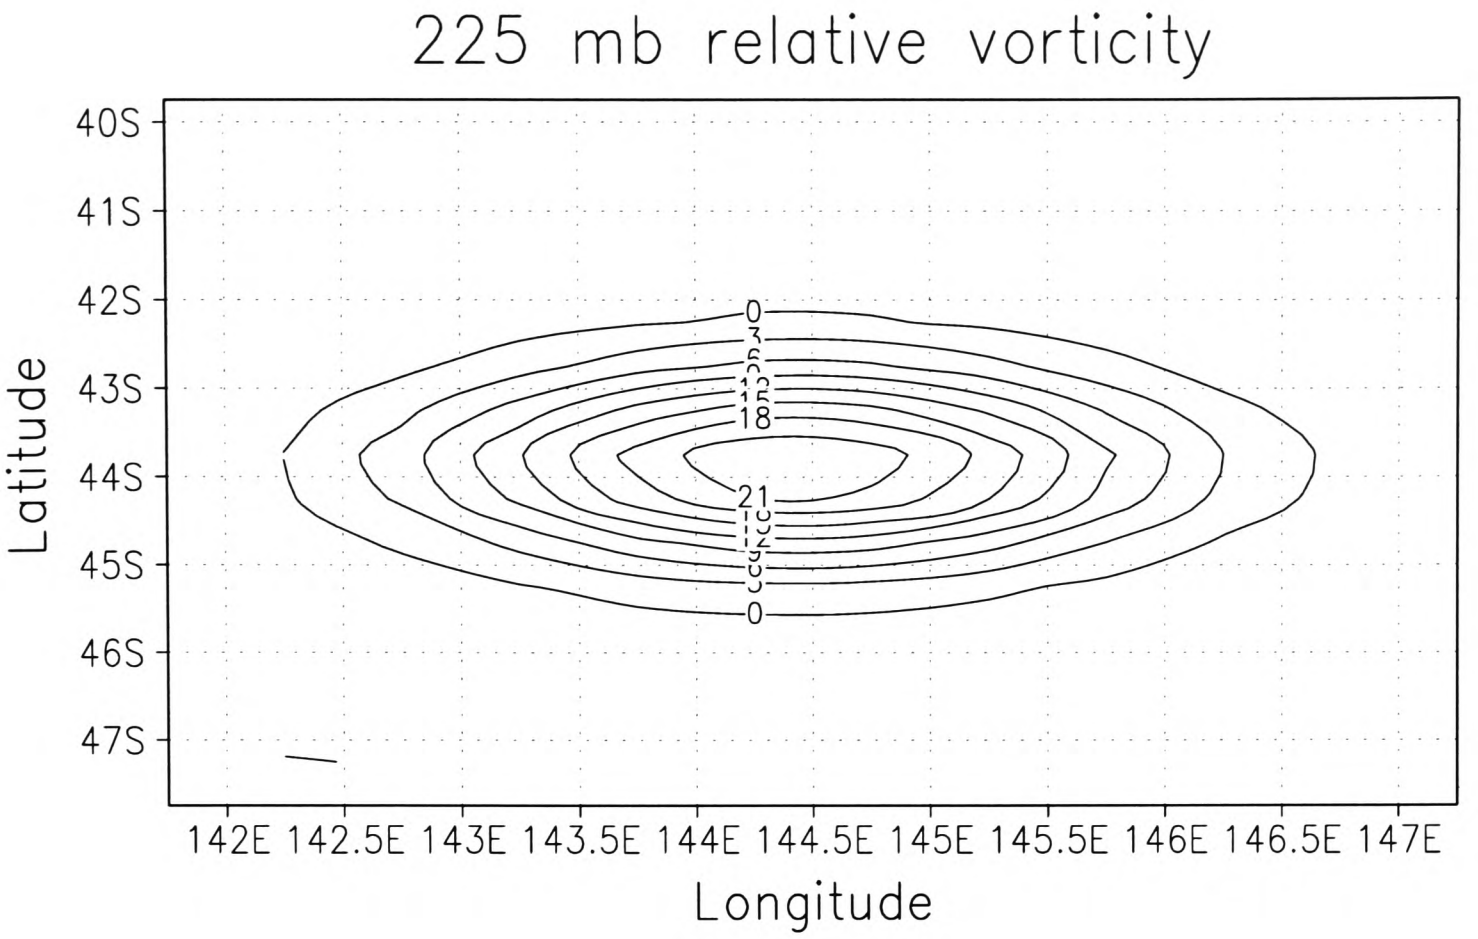


Figure 5.19: Relative vorticity (above), isotherms and Q-vectors (below) in an intensifying vortex. Note how the Q-vectors cross the isotherms in several places in the vortex collar.

resting atmosphere.

5.4 Vortices in alternating jetstreams

5.4.1 Introduction

The observed zonal-mean wind pattern on Jupiter consists of many narrow jets, aligned mostly zonally but alternating in direction (Limaye, 1986). From the point of view of the present study, the significance of the jets and their associated shear zones is the environment they create for the eddies. As we mentioned before, we are here ignoring the question of the origin and maintenance of the jets themselves.

Perhaps the simplest conceptual view is to consider the jets as changing the meridional gradient of background potential vorticity. In quasi-geostrophic theory, with purely zonal barotropic jets, this gradient is given by $\bar{q}_y = \beta + u_{yy}$. (The subscripts indicate meridional differentiation). In other words, the planetary vorticity gradient is modified by the curvature of the mean flow. This will have a corresponding effect on eddy migration, although the picture is complicated by eddy-mean flow interactions such as baroclinic instability.

In the present investigation, we have run models with a realistic jovian pattern of horizontal winds, but this pattern has been found to have certain disadvantages. The main one is that we have been unable to implement a satisfactory boundary condition at the equator. We attempted to impose the observed equatorial wind speed, so the model depicts a channel with one sidewall at rest and the other in motion. Unfortunately, an unrealistic boundary layer always developed, with the flow rolling up into a string of small vortices. Various modifications failed to resolve this problem. We were forced to fall back on the simpler boundary condition of stationary walls, requiring the flow to vanish at the boundaries.

This restriction is not a problem in principle, since Jupiter's pattern of alternating jets has many latitudes where $\bar{u} = 0$. In practice, the equatorial jet is so wide that it is too expensive to model its entire width at the resolution required for vortex experiments. (In Chapter 6 we do model this jet, but at half the horizontal resolution).

Therefore we cannot choose a sidewall latitude further north than 15°S if we are to use wholly realistic jovian winds. This turns out to be a more severe restriction than expected, because eddies generated at the GRS latitude (around 22°) can easily reach the sidewall. It is therefore difficult to judge whether they are confined realistically by the jets, or unrealistically by the wall. One possibility which was tried was to include part of the equatorial jet, but allow it to decay to zero at the equator. This produces interesting eddies

but is an unfortunate compromise, being neither particularly realistic nor convenient.

Another factor militating against the use of realistic jets is that, south of the STB, the jets are rather narrow and sharp. They are hard to simulate in a rather strongly diffusive model of half-degree gridlength. In view of all these considerations, it was decided to perform most of the experiments with simplified, wider jets whose strength, width, curvature etc. can be designed for convenience. We hope that the essential processes occurring on Jupiter do not depend sensitively on the precise positions and strengths of the jets, and that this is a sensible starting point. Indeed the approach of idealized jets is not a new one; it has been used extensively, for example by Williams (1996) in a primitive equation model, and by Maxworthy and Redekopp (1976), Ingersoll and Cuong (1981), Williams and Wilson (1988), and Dowling and Ingersoll (1989).

5.4.2 Jet initialization

In this section we begin by examining a few simple, idealized flows which resemble Jupiter (although with wider jets) in which the horizontal PV structure is chosen with the aim of confining the eddies. The vertical structure is more a matter of technical convenience and is based on the assumption that the jets extend to levels deeper than the bottom of the model.

We initialize a zonally symmetric flow in geostrophic balance, then run the model for a period (typically 80 days) allowing the flow to adjust to non-geostrophic equilibrium. The model is forced mechanically on the lowest level, using the zonal-mean forcing scheme outlined in Section 3.3.2. This generates a stable flow which is close to the reference state calculated using QG theory. Because we force only the zonal mean winds, the eddy part is unconstrained, and deep vortices (barotropic in the lower part of the model) can be supported. Brief localized heating is then applied to initialize an eddy within the flow, and the run continues with the same zonal forcing scheme as initially. The dissipation in the model is the usual sponge layer at the top, with Newtonian cooling throughout the domain.

Two sinusoidal jets

An example of the adjustment stage is illustrated in Fig. 5.20. The initial flow is a pair of sinusoidal jets of amplitude 100 m/s. The two jets fill the model domain, vanishing at the sidewalls. The vertical structure is the same as that of the deep vortices seen earlier, so the flow is barotropic below the clouds, then decreases with height to zero at the upper boundary. The final flow is shown in the two upper panels: note that the easterly jet has a sinusoidal shape but the westerly has been slightly eroded. The lower two panels show

timeseries illustrating the adjustment towards this profile. There are large oscillations over the first few days (evidence that geostrophic balance is not quite valid) which then die away as more accurate balance is reached.

Figure 5.21 demonstrates how jets can modify the meridional structure of the potential vorticity. Without them the PV has a uniform gradient (see e.g. Fig. 5.11) causing eddy migration. Here the PV is almost uniform between 30° and 36°S. We shall discuss the evolution of vortices of varied magnitudes in Section 5.4.3 below. We shall see that the PV peak is rather too flat to confine large vortices, although it does confine weaker ones. We would therefore like to find an alternative configuration with a sharper PV maximum, which we will expect to act as a barrier against vortex migration. Looking again at Fig. 5.21, we see that the PV barrier is positioned on the poleward flank of the easterly jet, the westerly having little effect. This suggests we should consider a flow dominated by one strong easterly jet.

One easterly jet

The reference jet pattern now has all positive (westerly) winds set to zero and the easterly increased. As before, the flow equilibrates before day 80. This equilibrium state is determined by the model and so is partly beyond our control. It includes a weak westerly, because the easterly is strong (140 m/s) and the atmosphere cannot support a shear zone as narrow as required. There is an overshoot, into weak westerly winds, but this helps to extend the anticyclonic shear zone. The PV features a large peak at 34°S. This peak has a similar value to that in Fig. 5.21 but is much sharper, and has a deeper trough to the north. There is a wide band between 26°S and 34°S where the PV gradient is mostly negative, in which we expect an anticyclone to drift polewards.

5.4.3 Shallow vortex in a pair of jets: dependence on vortex strength

In this section we use fourth-order hyperdiffusion with a fixed damping time of 3 hours ($N = 12$), and vary the strength of the vortex. Then in Section 5.4.4 we shall compare these results with those obtained using a shorter timestep and damping time of 90 minutes. It is necessary to compare models with a range of parameters, since these are chosen for numerical reasons with little physical justification for particular choices. As always, we need to try to separate physical effects from numerical ones.

We begin with the model state shown in Fig. 5.20, i.e. a pair of equal and opposite jets. At day 80, the vortex heating is switched on for one day, increasing from midnight to noon, then decreasing back to zero by midnight. This generates an anticyclone which quickly adjusts into an elliptical shape, reminiscent of the Great Red Spot. Here is an

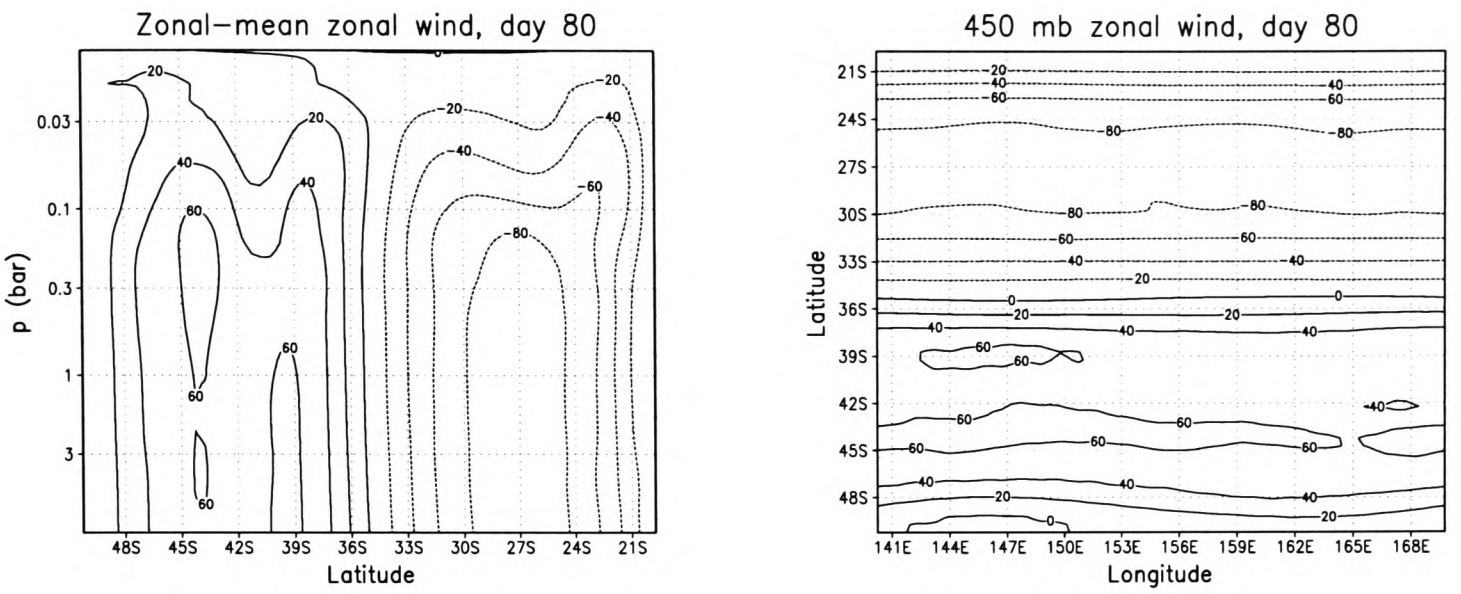


Figure 5.20: Results of an experiment to generate a stable baroclinic shear flow with alternating jets. The first panel shows the final zonal-mean zonal wind cross-section. The second shows the cloud-top winds at the same time. The lower pair show the time evolution of the winds, at the base of the model and at the cloud-tops. The initially large oscillations die away to leave a stable flow similar to the reference state.

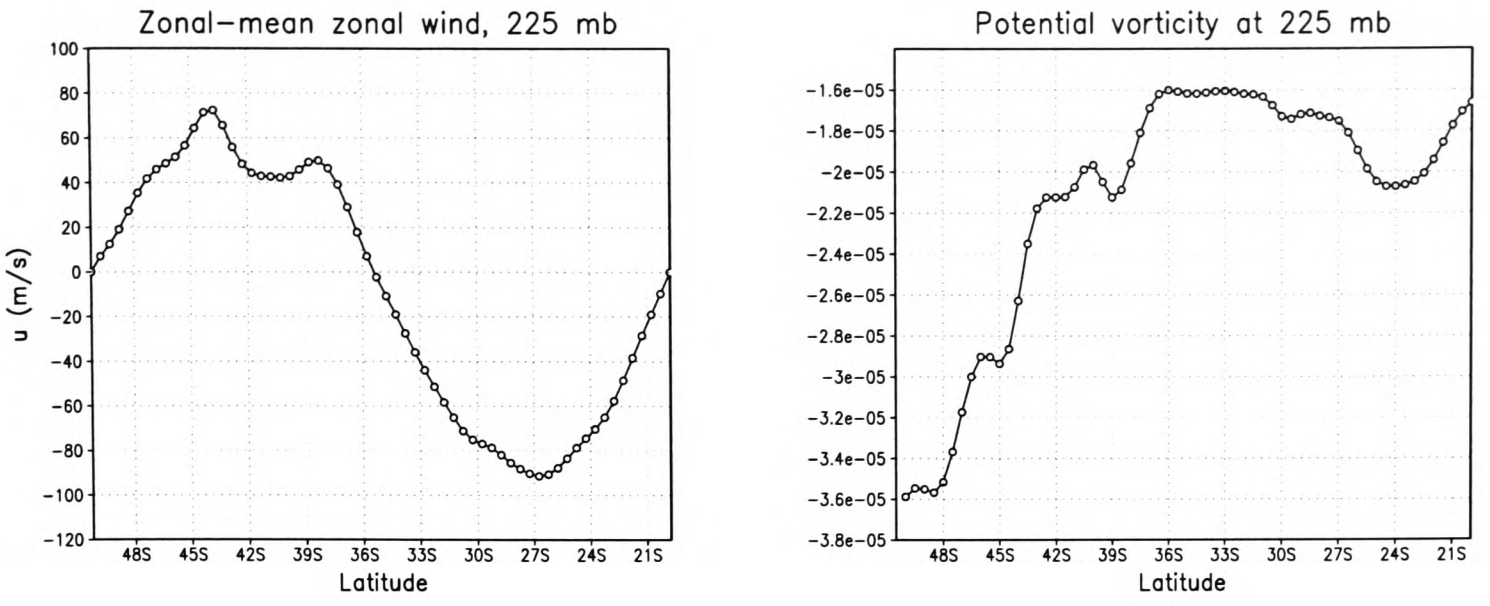


Figure 5.21: Zonal-mean wind and potential vorticity at 225 mb, after 80 days' run. The jets have altered the PV profile from its rest state (a constant gradient) and introduced several local maxima. The 225 mb level is the peak of the jet, and in later runs, vortices will be centred at this level.

obvious but important difference from the first part of the chapter. Vortices in a resting atmosphere assume a circular shape, whereas the elliptical one characteristic of observed jovian features is due to shear in the background flow. This has implications for the relevance of our PV inversion work in Chapter 2. The assumption of axisymmetry was quite well satsified by the GCM vortices until now. Henceforth that assumption is not strictly valid, but our vortices still have much in common with those in the earlier work.

The subsequent evolution of the vortex depends on the global model parameters and on the strength of the initial vortex. In this section we vary only the vortex strength, characterised by the heating coefficient ΔT_0. The shape, position, and time-dependence of the heating function are fixed so a larger ΔT_0 produces a stronger vortex. A model run with $\Delta T_0 = 0.4$ is found to generate a meridional wind of 25 m/s, once the vortex has spun up (at day 4). By day 20 this decreases to 15 m/s and the vortex is no larger than the general noise and small spontaneous eddies in the westerly jet.

With $\Delta T_0 = 0.6$, the eddy is the same shape but the v-wind is 50 m/s at day 4. This is sufficient for the vortex to survive and actually grow. At day 20 the v-wind has increased slightly, to 60 m/s. The vortex migrates slowly equatorward and continues to grow. Figure 5.23 shows a sequence of horizontal PV maps at 20-day intervals. There are several features of interest. Firstly, the background PV gradient is clearly small compared to gradients within the vortex. Secondly, the vortex is elliptical not circular, probably because of the background shear. Comparison of the top two panels (Days 10 and 30) shows that during this period, the eddy has not migrated at all but has grown *in situ*. It later becomes strong enough to overcome the PV barrier, migrating to 31°S. By this stage the winds are unrealistically high ($\sim$ 400 m/s) and the vortex very large. The vertical structure (not shown) also changes in the usual way seen in previous runs: the vortex rises and approaches the first baroclinic form. Either the physical mechanism responsible for

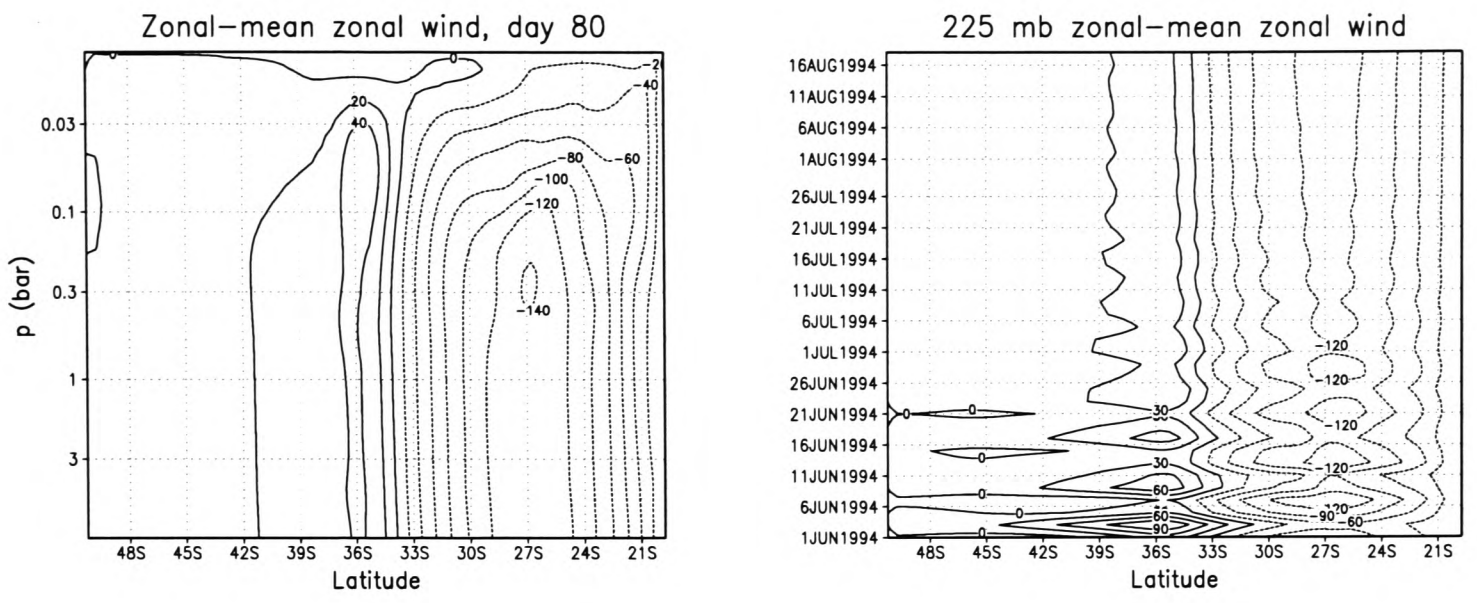


Figure 5.22: An experiment to generate a single easterly jet as a background state. Shown are a cross-section of the final zonal wind; the 225 mb wind as a function of time; and the final 225 mb wind and PV profiles. The 225 mb level is chosen because the jet peaks at that level, and it clearly shows the local maximum in PV.

limiting growth on Jupiter has not been included in the model, or the background flow is unrealistic in some crucial but unknown respect.

These model results certainly suggest that barrier penetration would not occur if the vortex intensification could be curbed, and a vortex of stable size and magnitude might be confined indefinitely. In the next few sections we will look at ways in which we can make some progress toward this goal, by means of modifying the general model parameters.

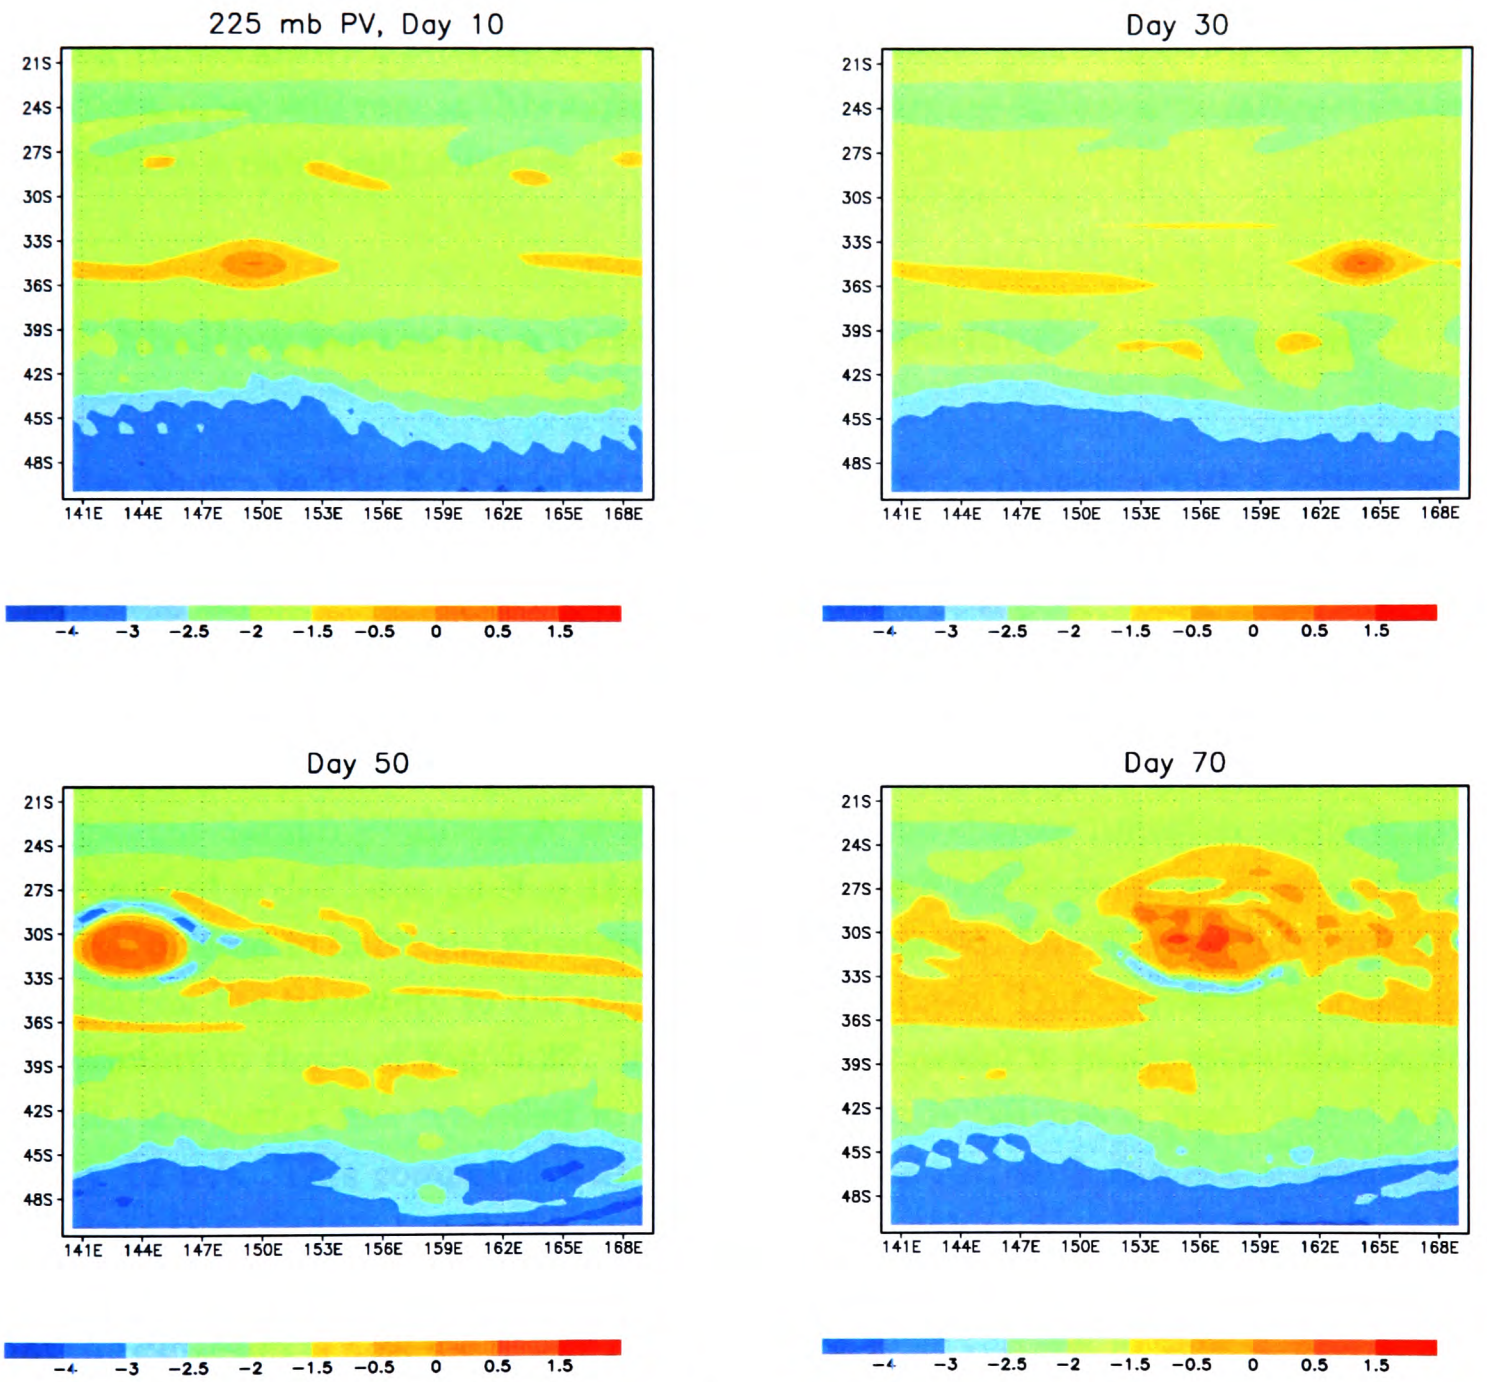


Figure 5.23: Potential vorticity ($\times 10^5$) at 225 mb, at 20-day intervals, for a vortex in a twin-jet flow with $\Delta T_0 = 0.6$. For 30 days the vortex grows without migrating, then suddenly begins to migrate and grows dramatically.

In a third experiment, $\Delta T_0 = 0.8$. The resulting eddy has a meridional wind of 100 m/s at day 4. Migration begins almost immediately, the eddy growing and strengthening rapidly. The life cycle is similar to the previous case but more rapid.

From these three experiments we may conclude the following: (a) varying ΔT_0 is a simple and effective way to regulate the peak wind velocity within a vortex while keeping its size and shape constant. This is a useful measure of the strength of the vortex. (b) Long-lived vortices exist only in a narrow window around $\Delta T_0 = 0.6$. Weaker ones are sheared

apart, while stronger ones are not confined by the PV barrier. (c) in the intermediate case, the vortex in its early state *is* effectively confined, and grows relatively slowly. Raising the barrier seems likely to expand the regime in which vortices are confined. Later we will look at similar experiments in the same basic flow as shown in Fig. 5.22—a strong easterly jet with a much weaker westerly to the south. This flow will be seen to provide a better barrier against vortex migration.

However, the experience of Section 5.3.10 showed that hyperdiffusion can be effective in limiting vortex growth, at least in a resting atmosphere. Before moving on to a different jet structure, we will repeat this experiment with varying diffusion parameters to extend the result to a more realistic case.

5.4.4 Shallow vortex in a pair of jets: dependence on diffusion

The flow shown in Fig. 5.23 was obtained using a 900 s timestep with a damping value $N = 12$ (at fourth order). Three more runs were done with a 450 s timestep, and a range of damping values $N = 18, 24, 30, 46$. The associated damping times, given by $\tau = N\delta t$, are 2.25, 3.0, 3.75, and 5.75 hours respectively.

The first experiment attempts to change the timestep while leaving other factors unchanged; the damping value is $N = 24$ which, with the shorter timestep, ought to give the same amount of diffusion as $N = 12$ originally, since $\tau = 3$ hours in each case. The model code was altered to halve the Newtonian drag coefficients. The thermal damping depends explicitly on the timestep, so did not need to be changed. This should have given results very similar to those of Fig. 5.23. In fact the new model is much more dissipative—at Day 80, the vortex has travelled to the same latitude but has a peak meridional wind of only 12 m/s. This compares with 250 m/s in the run with the 900 s timestep! The difference is worrying, but we know that vortex growth is very sensitive to the diffusion parameter. The new result is certainly closer to observed Jovian behaviour. It means that the diffusion is now too heavy for stable long-lived vortices to persist, and so another experiment was run, using a larger value, $N = 30$.

This second run (Fig. 5.24) was integrated over 250 days because the vortex was expected to be longer-lived. At Day 80 the v-wind reached 60 m/s, i.e. roughly the same as near the beginning of the run and rather greater than in the $N = 24$ experiment. Figure 5.24 shows a sequence of 225 mb potential vorticity. The behaviour of the vortex is quite similar to earlier cases, but the timescale is much longer. After 200 days the vortex is still at 30°S, the latitude of maximum ambient PV. It has grown but not strengthened noticeably (note that the contours are the same as in the top panels). However, after this its growth is rapid and it reaches the equatorial sidewall by Day 240. This experiment is very promising in that we see a vortex unchanged over a period of several (Earth) months. It remains

to be seen, though, whether we could identify a growth-limiting mechanism more physically justifiable than heavy diffusion. We do not take the procedure any further, since it is computationally expensive to perform long runs with this shorter timestep.

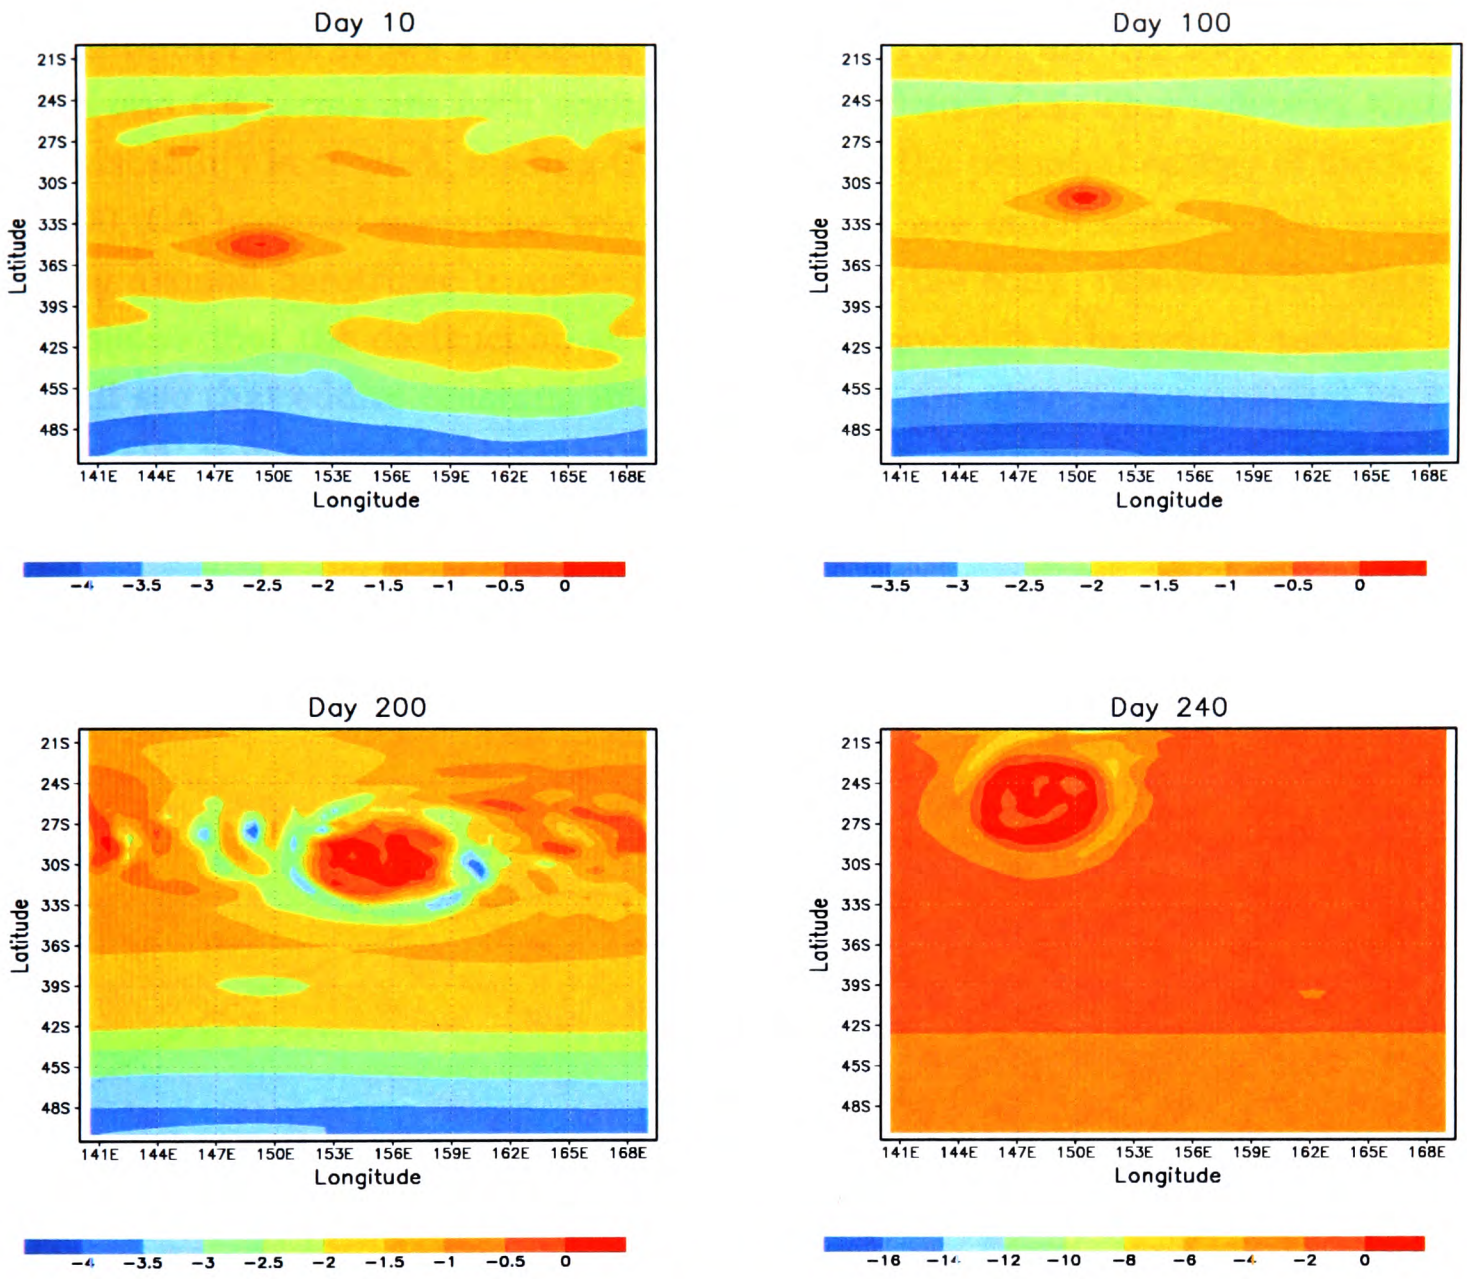


Figure 5.24: Potential vorticity ($\times 10^5$) at 225 mb, for a vortex in a twin-jet flow with $N = 30$. The vortex slowly migrates to the PV barrier at 30°S, where it subsequently grows larger and finally breaks through the barrier.

To identify the processes governing the growth and decay of the vortex, we use the quasi-geostrophic energy diagnostics defined in Section 3.4. There are several caveats: one is that the small vortices are not truly quasi-geostrophic. The Rossby number, $\mathrm{Ro} = U/(f_0 L)$, is around 0.1 for the initial phase of the spontaneous vortices, but when the winds reach several hundred metres per second, Ro rises to ~ 0.3. Secondly, the available potential energy depends on the squared reciprocal of the static stability. In the troposphere the static stability is often very small, so convective motions in the deep atmosphere dominate the integrated potential energy. Nevertheless the gross features of the energy timeseries indicate what is happening.

Figure 5.25 shows the kinetic energy terms on the left, and three of the conversion rates on the right. There is a slow growth in the zonal-mean kinetic energy KZ, because the

flow is being forced from below. Around Day 200, the KZ stops growing while the eddy energy KE grows more rapidly. This is the point at which the vortex strengthens rapidly. Finally, both energies peak and KZ drops away as the vortex reaches the sidewall.

The right-hand plot shows a possible interpretation. From around Day 160 to Day 200, the CA and CE terms are both positive and larger than CK. This indicates that baroclinic instability is at work, feeding the vortex from the potential energy of the jet. After Day 200, CA becomes enormous while CE and CK are much smaller. CK is negative, showing a small barotropic transfer from the jet to the eddy. However, the large value of CA shows that the destruction of the vortex is probably a baroclinic process. Later, we shall see that eddies crashing into the sidewall are always accompanied by a large spike in CA. It is a useful pointer to model processes: large peaks tend to mean that the sidewalls are affecting the results so we need to proceed with caution.

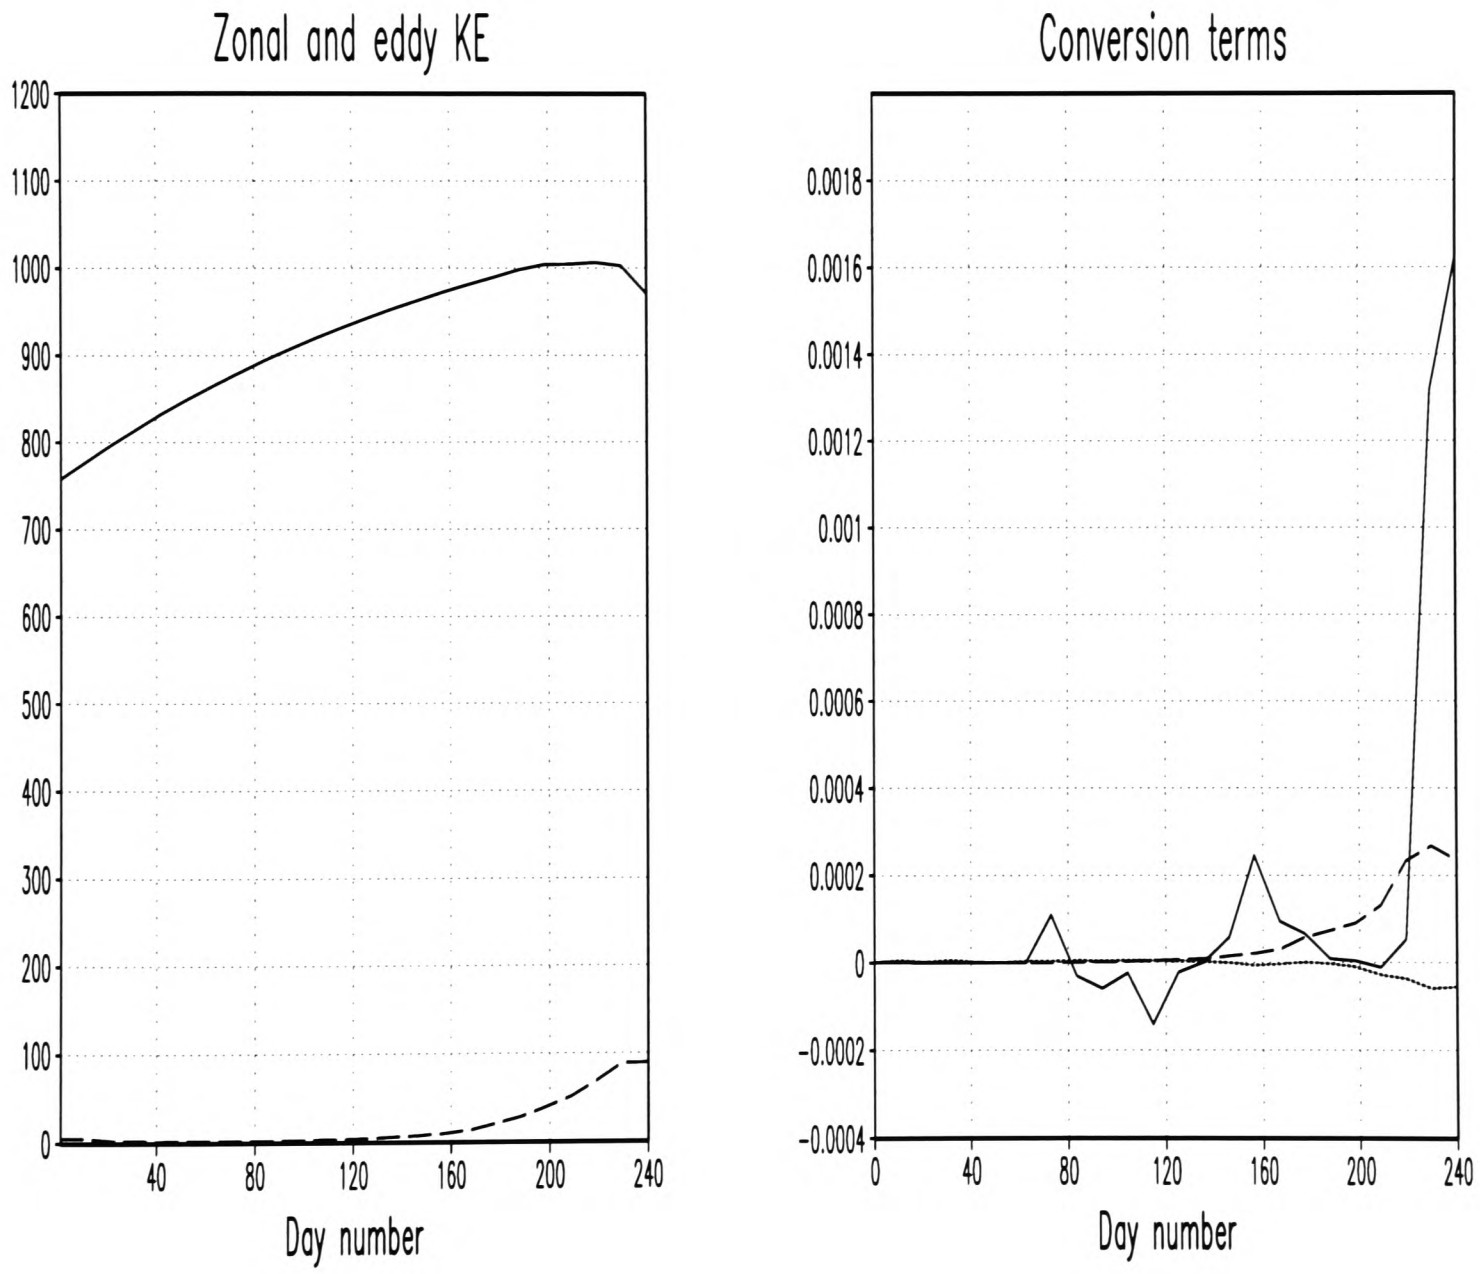


Figure 5.25: Energy diagnostic timeseries for the flow shown in Fig. 5.24. Left: KZ (upper, solid line) and KE (lower, dashed line). Right: CA (solid line), CE (dashed line), and CK (dotted line).

5.4.5 Shallow vortex and an isolated easterly jet

We now consider a vortex facing the steeper PV barrier presented by a strong easterly jetstream. We set $\Delta T_0 = 0.5$, a little larger than in Section 5.4.3 because the background winds are stronger. The timestep was 450 s, and horizontal diffusion time of $\tau = 6$ hours—rather weaker diffusion than we found necessary before. Fig 5.26 shows a relative vorticity sequence over the first 120 days. The vortex does not oscillate in latitude, but remains on the flank of the jet. It intensifies steadily, and by day 120 it has become strong enough to migrate past the barrier. Experiments with a shorter diffusion damping time led to the vortex dissipating, rather than (as we would hope) remaining stable. Once again it seems difficult to find an equilibrium state in which the vortex intensity remains constant.

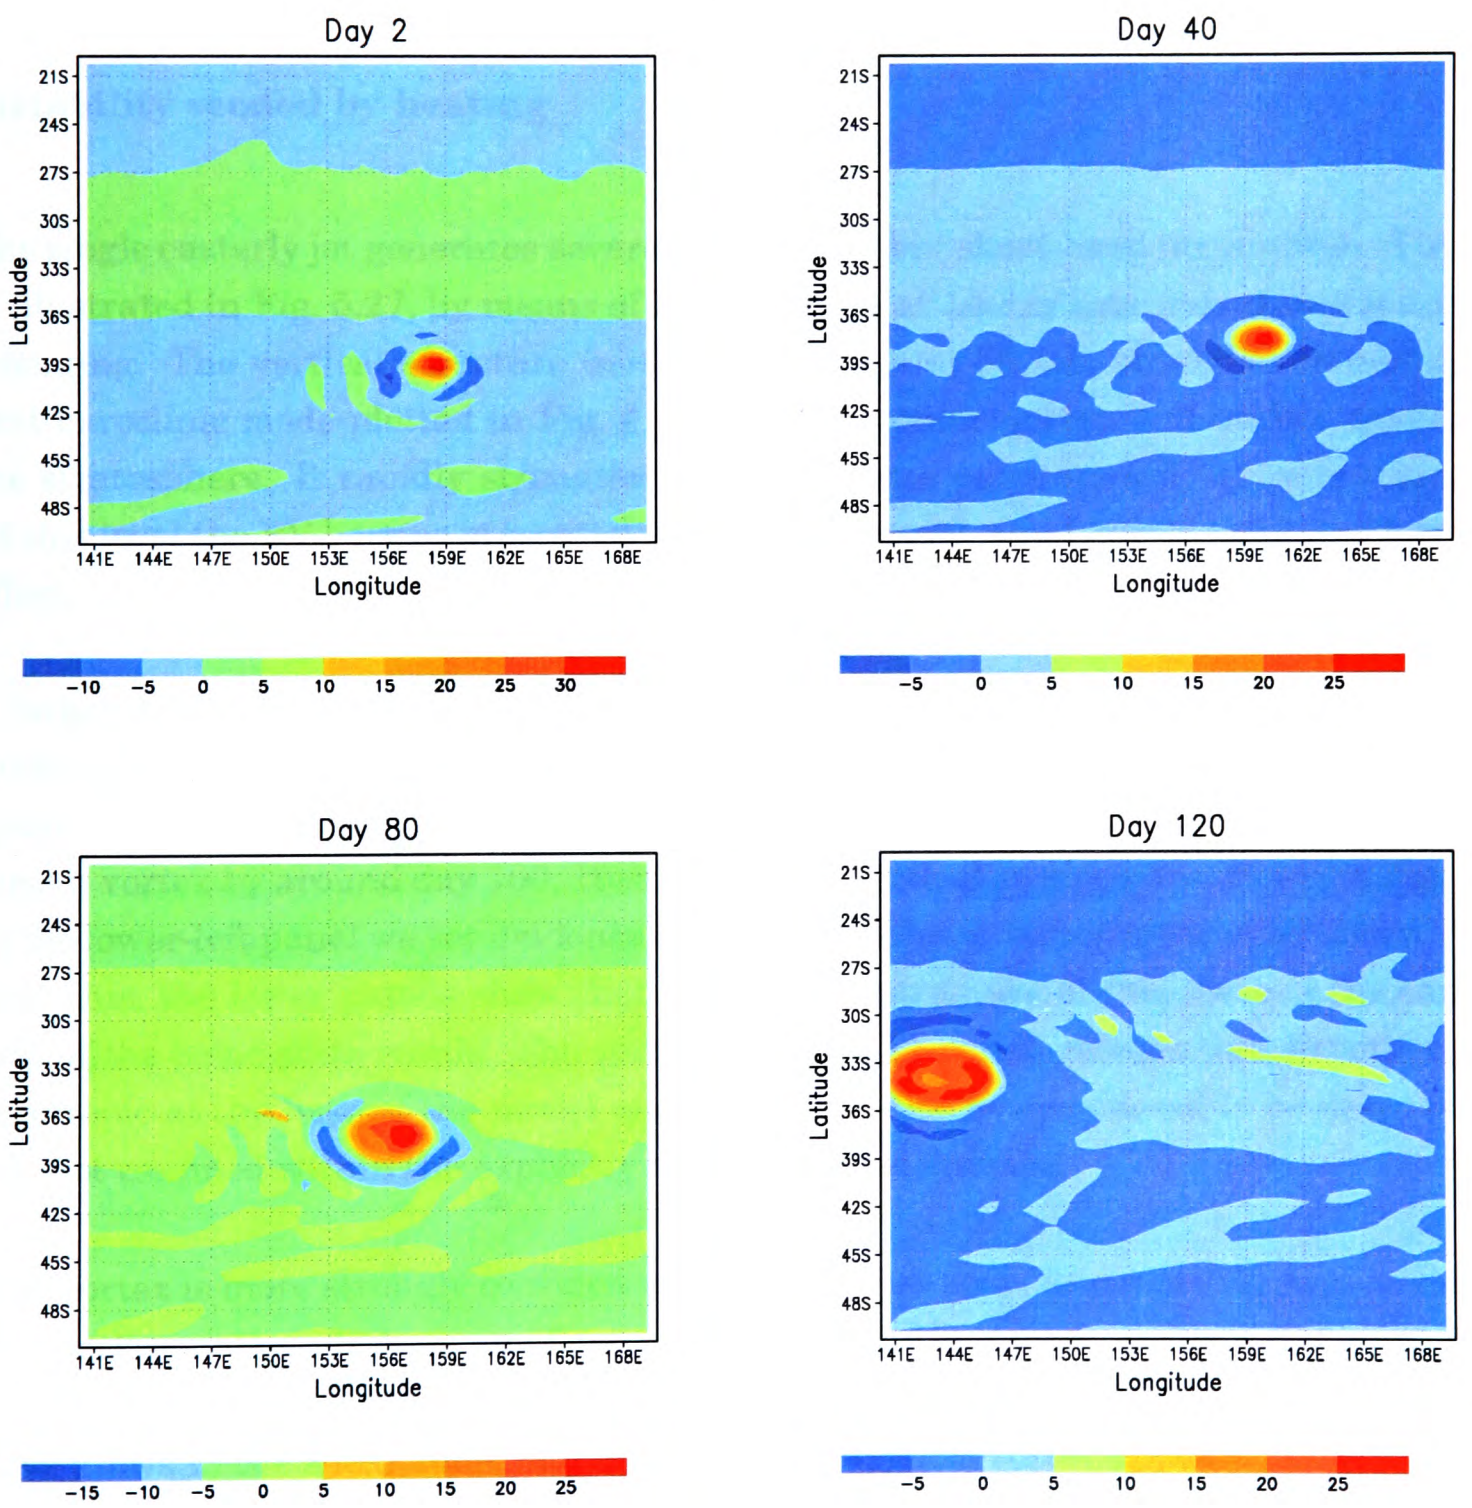


Figure 5.26: Relative vorticity $\times 10^5$ at 225 mb

5.4.6 Instability of the easterly jet

Although Fig. 5.22 shows the jet adjusting to an apparently stable state, it is in fact baroclinically unstable. This is easily demonstrated by seeding an instability with an arbitrary non-symmetric forcing, or even by allowing numerical noise to grow. A small "bubble", like those used to initialize the eddies but much weaker, was found to be suitable for forcing. It is physically plausible and, unlike alternatives such as wave forcing, does not preselect any particular wavenumber or vertical structure. We compare these results with those of a null experiment without heating. The model behaviour is similar in each case. The form of forcing turns out to be unimportant; these are two examples of model evolution, differing in detail but with the same general features.

Instability seeded by heating

The single easterly jet generates several interesting but short-lived instabilities. The first is illustrated in Fig. 5.27, by means of four PV plots at 10-day intervals showing a vortex emerging. The vertical structure (not shown) peaks at 56 mb, strongly resembling the first baroclinic mode plotted in Fig. 4.13. This baroclinic vortex is therefore confined to the stratosphere. It rapidly strengthens and migrates equatorward before breaking up. At this level the PV barrier presented by the background zonal flow is weak and has little effect.

A larger scale feature becomes noticeable by day 80. This is a zonal wave, centred on latitude 34°S, precisely in the middle of the shear zone and the point where the barotropic shear $\partial u/\partial\phi$ is largest. This wave, initially wavenumber 1, shortens and "rolls up" to become a vortex by around day 100. Horizontal and vertical sections are shown in Fig. 5.28. In the lower-left panel we see evidence of the stratospheric vortex mentioned above. Apart from this, the lower panels show that the vertical structure of this wave is the same as that of the basic-state winds. The vortex (lower-right panel) retains this structure. It is barotropic at the base of the model and must therefore be considered to be deep, in that it is not confined within the explicitly resolved model domain.

This vortex is more strongly confined within the shear zone, as one would expect (the tropospheric zonal winds form an effective PV barrier). However it continues to strengthen, and breaks through around day 115. Further disturbances are seen: another barotropic wave grows in the shear zone, and another stratospheric vortex appears.

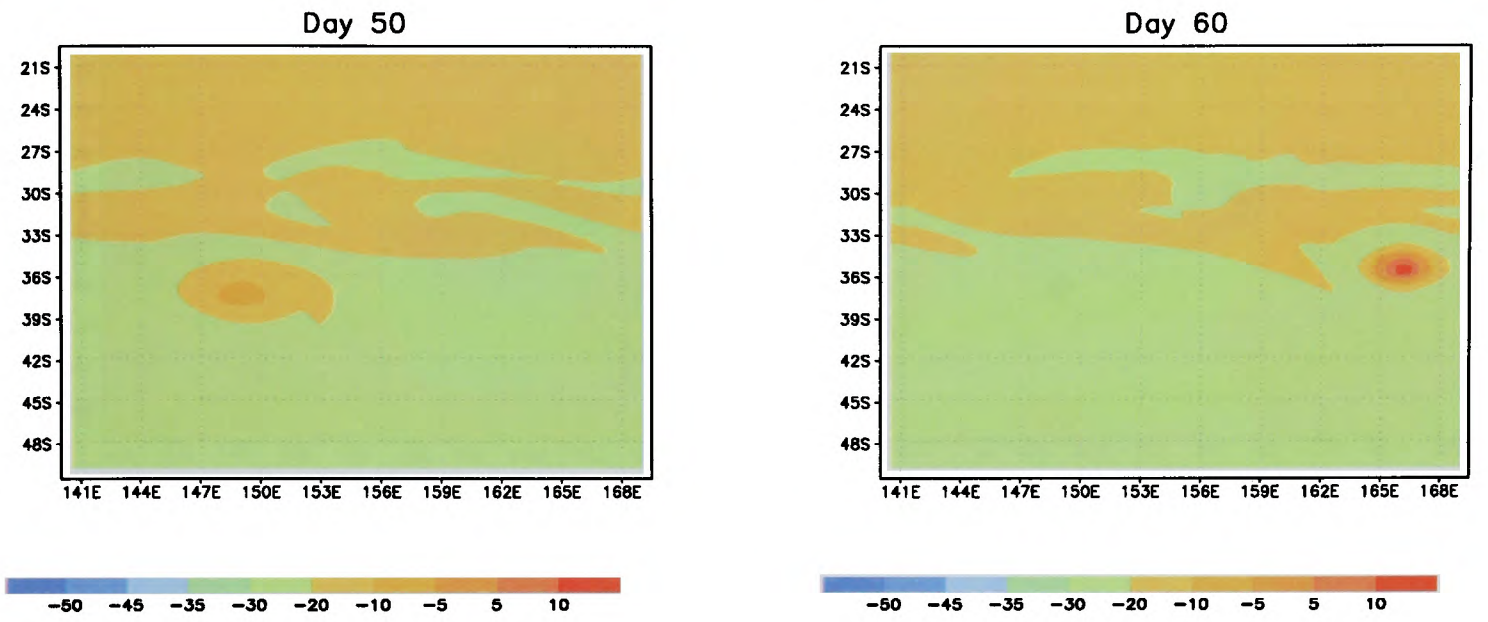


Figure 5.27: Potential vorticity ($\times 10^5$) at 56 mb, at 10-day intervals (dates relative to the heating applied on Day 0), showing the spontaneous growth of a stratospheric vortex.

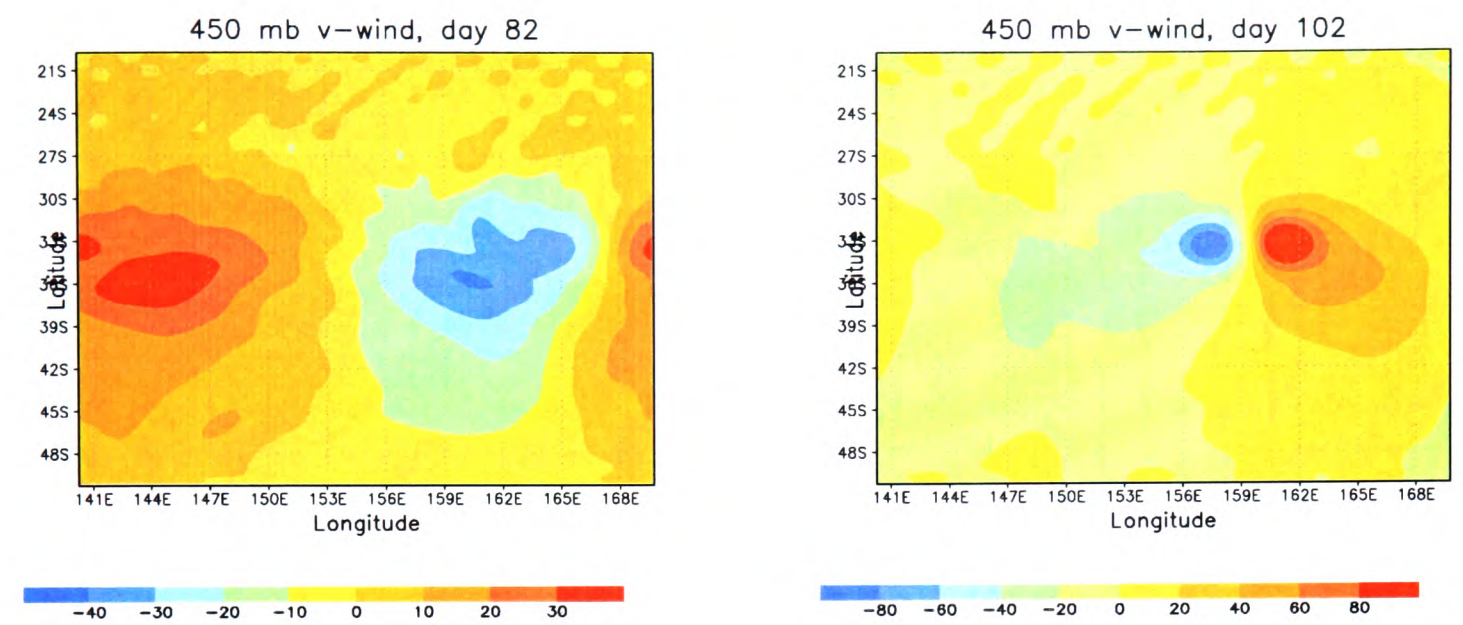


Figure 5.28: Development of an instability wave which rolls up into a vortex. The top panels show the 450 mb meridional wind. On the left a wavenumber 1 wave is obvious. On the right, 20 days later, it has shortened into a compact vortex. The lower panels show vertical sections through the flow at the same times. The vortex selects the same vertical structure as the background flow, i.e. the first baroclinic mode.

Spontaneous instability

The null experiment—identical except without a forced disturbance—shows similar behaviour. Figure 5.29 shows the development of a wave (this time of wavenumber 2). The vertical structure is largely barotropic (damped by the sponge layer at the top) but is more complicated than in the previous example.

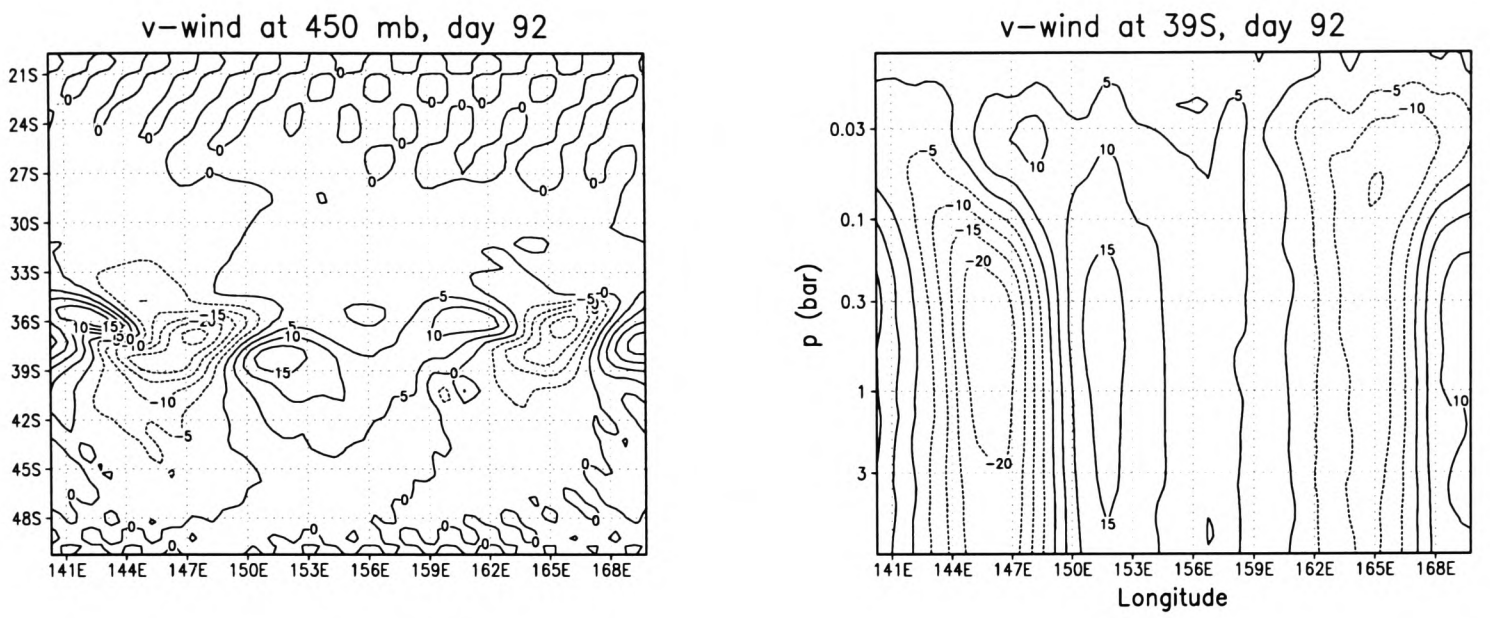


Figure 5.29: Meridional wind against latitude and height, showing a baroclinic wavenumber 2 developing in the shear zone.

This wave later develops into vortices as in the previous experiment. However, as it has wavenumber 2, it breaks into two vortices instead of one (Fig. 5.30). These are initially at similar latitudes, but one is slightly stronger. They remain at around 34°S while increasing in strength. By Day 116, one vortex has become strong enough to migrate over the barrier. The other is weaker, but later follows the first. The vertical structure of these eddies is like that in the first experiment: mostly the first baroclinic mode, with a peak near 450 mb, decaying rapidly with height. There is weak shear in the troposphere, the winds decaying by about 30% by the 10 bar level.

5.4.7 Energy analysis of the easterly jet instability

Figure 5.31 shows the zonal-mean energy terms for the period whilst the jet is spinning up. The eddy terms are negligibly small and are not plotted. Both curves have an upward trend underlying a decaying oscillation. This agrees well with the timeseries of the wind fields, in Fig. 5.22.

Figure 5.32 shows the next 160 days of the same experiment, after the spin-up phase. Kinetic energy is shown on the left and available potential energy on the right. For this stage of the experiment we also plot the eddy terms, which become large. At around day 85 (i.e. five days after the disturbance heating), the jet rapidly collapses, losing more than

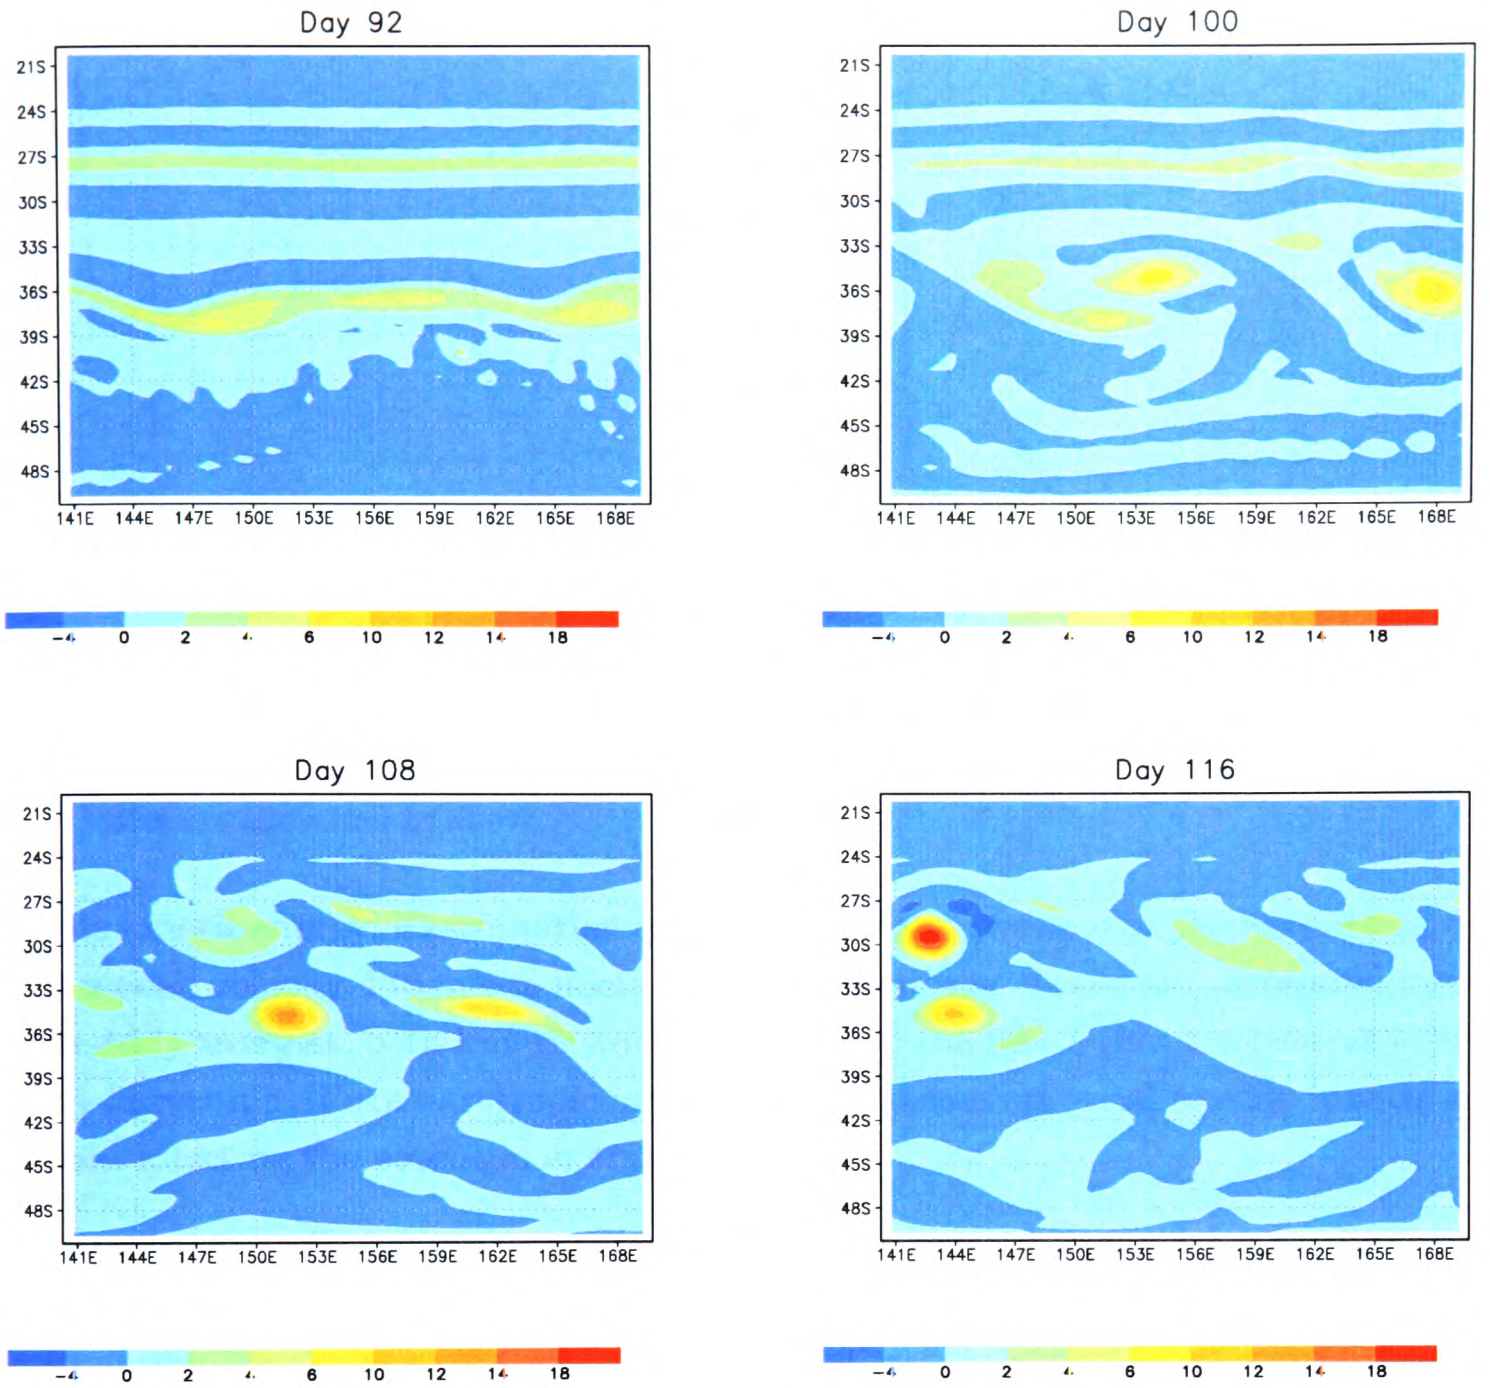


Figure 5.30: Relative vorticity ($\times 10^5$) at 450 mb, at 4-day intervals, showing the development and migration of a pair of eddies. (Relative, not potential, vorticity is plotted to show the motion of the eddies more clearly).

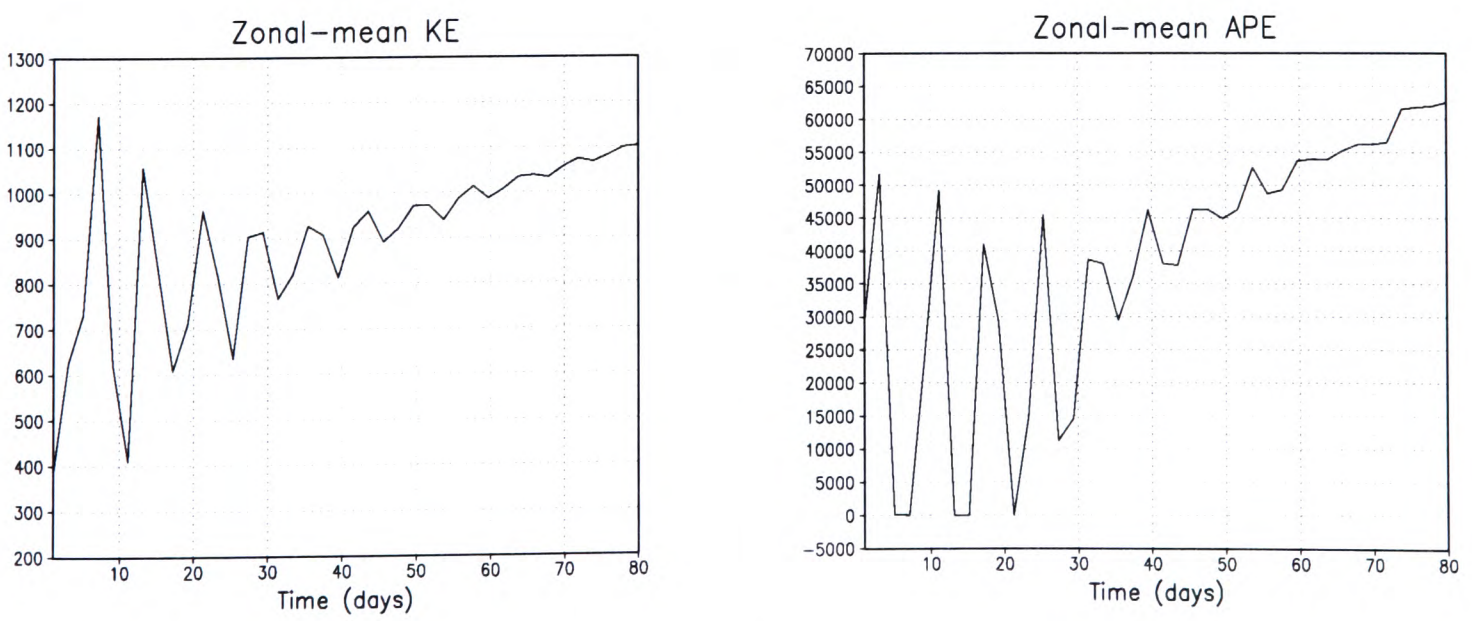


Figure 5.31: Timeseries of zonal-mean kinetic and available potential energy for the 80-day period of initializion of the easterly jet. The units are arbitrary, representing joules per grid square.

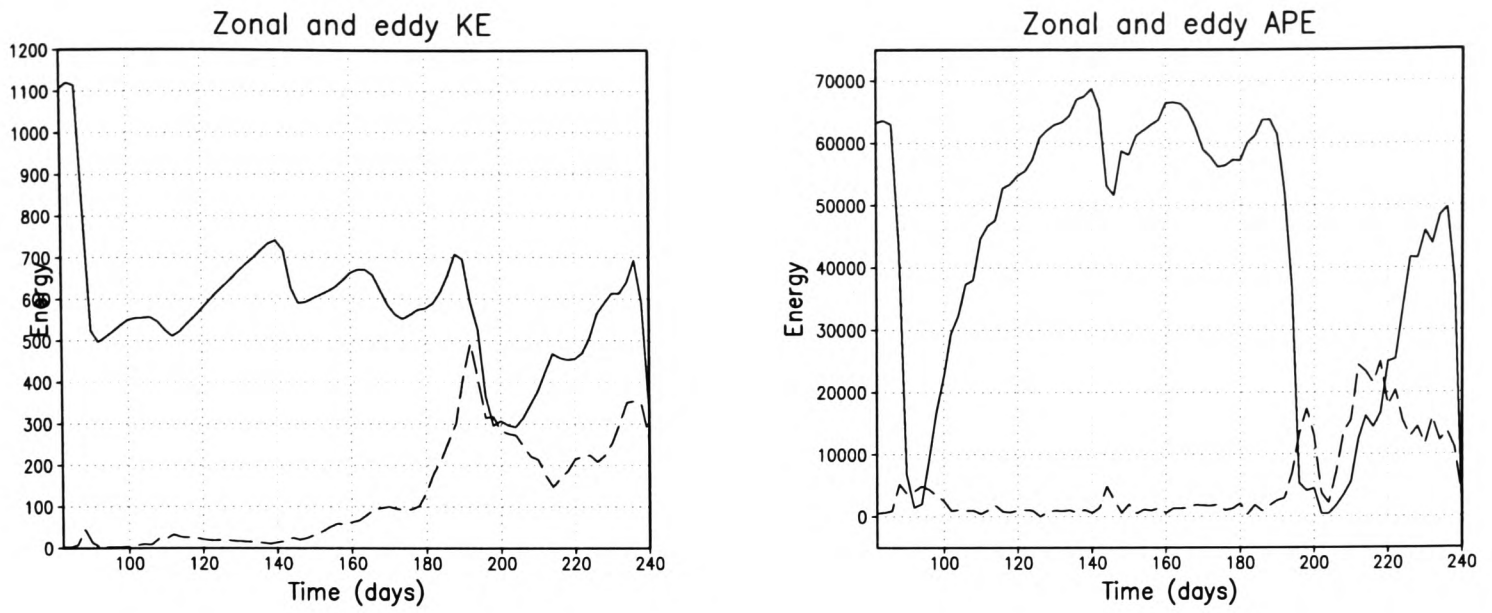


Figure 5.32: Timeseries of zonal-mean energies (solid lines) and eddy energies (dashed lines) for 160 days of the seeded-instability experiment. The time axis begins at day 80, immediately after the spin-up stage shown in Fig. 5.31. The energy units are the same as above.

half its kinetic energy. Examination of the model winds shows that the easterly jet decays from a peak velocity 140 m/s to about 90 m/s at this time. There are correlated increases in the eddy energies, both kinetic and potential. This suggests that the process is at least partly baroclinic. (Pure barotropic instability would convert zonal to eddy kinetic energy without affecting the potential energy.)

This collapse of the jet is not due to the introduction of the small heating bubble—it occurs at almost the same time in the null experiment. It happens because the bottom-boundary forcing is continually feeding kinetic energy into the jet, and this cannot continue beyond a certain point, at which the jet becomes unstable. The energy increases again after the drop, the APE recovering to its original maximum level but the KE failing to return to its maximum.

The most striking process is that taking place between days 180 and 200. During this period the eddy kinetic energy grows rapidly. Then, on day 190, it peaks and begins to drop just as rapidly. The eddy potential energy peaks a few days later. As the eddy KE reaches its maximum, the zonal-mean KE and AE both drop sharply. The explanation for all this is straightforward. A solitary deep eddy was formed at day 100, as shown in Fig. 5.28. This slowly strengthens and then begins to migrate (see Section 5.4.6 above). The rapid growth of eddy energy after day 180 represents the final stages of its explosive growth. At day 194 it reaches the northern sidewall, and this precipitates its destruction. The zonal-mean terms drop sharply too, because the vortex is so strong that it dominates the circulation. A similar event happens at the end of the run, as another vortex reaches the sidewall.

Figure 5.33 shows the corresponding energy terms for the experiment with no vortex forcing. There is a clear sawtooth pattern in the zonal-mean kinetic energy. The forcing

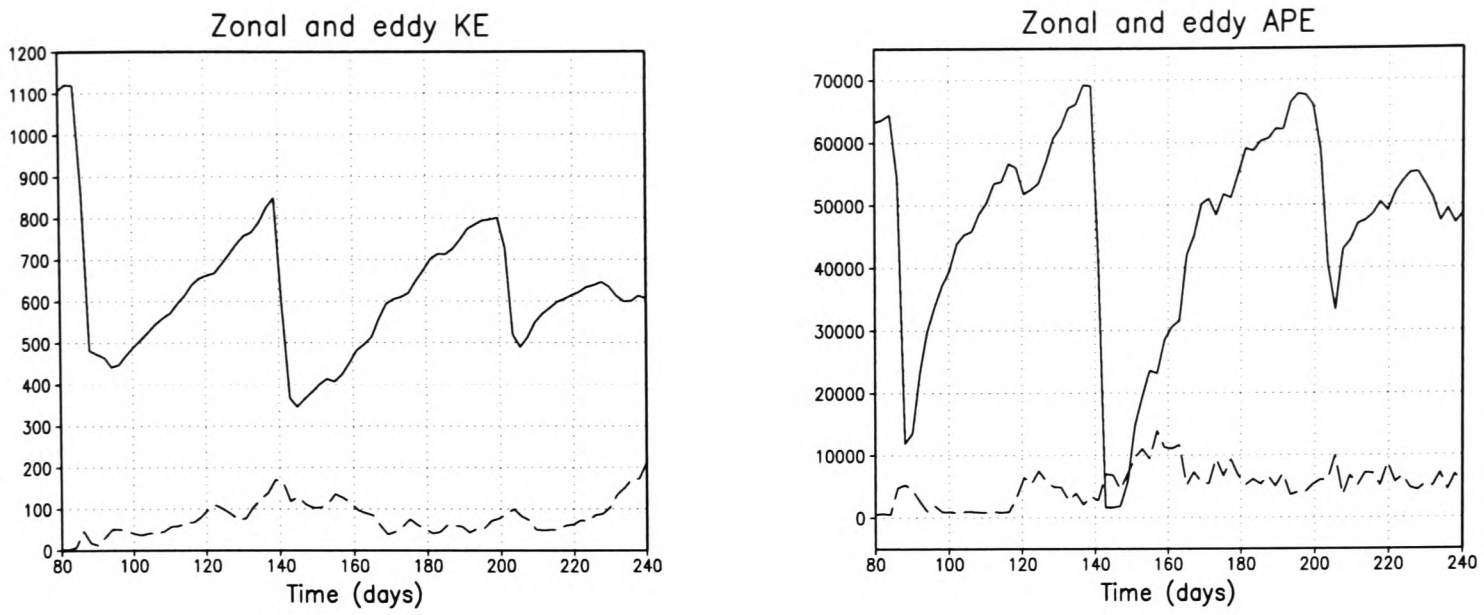


Figure 5.33: Timeseries of zonal-mean energies (solid lines) and eddy energies (dashed lines) for 160 days of the spontaneous instability experiment. The energy units are the same as above.

builds up the jet until it becomes unstable, releasing eddies. The eddies migrate to the northern sidewall and dissipate, sharply reducing the energy, and the cycle is repeated. The mechanism causing the periodic crashes in energy is revealed by explicitly calculating the energy conversion terms. Three of them have no clear pattern, but the baroclinic term (CE) has large spikes correlating with the sharp drops in kinetic energy.

The correlation between CE and KZ (plotted on the same axis) is shown in Fig. 5.34. The CE term quantifies the conversion from eddy potential to eddy kinetic energy and is an indicator of baroclinic instability. Barotropic instability would be indicated by large CK, denoting conversion from zonal kinetic to eddy kinetic energy. The timeseries of CK (not shown) does not show any obvious structure, or correlation with the other quantities. We interpret this to mean that barotropic instability is not a significant process in this case.

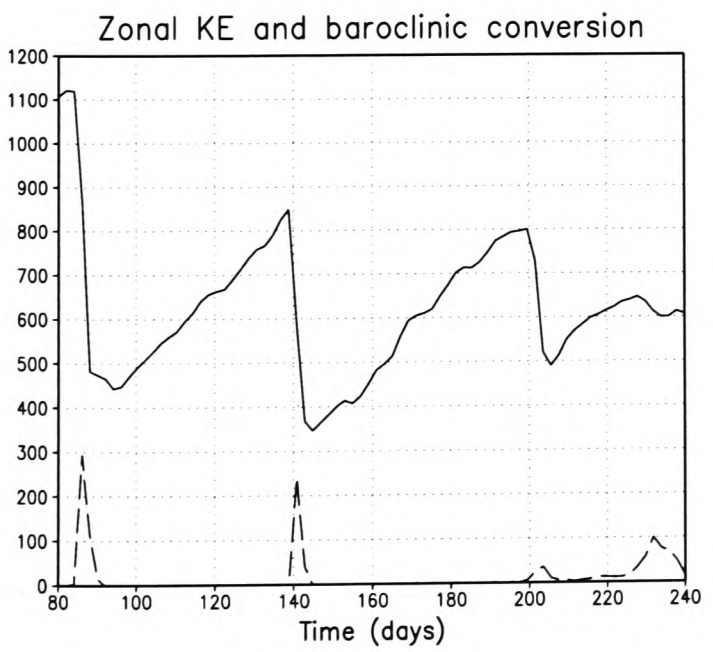


Figure 5.34: Timeseries of zonal-mean KE (solid line) and baroclinic conversion term CE ($\times 20000$ for clarity). The correlation is clear.

5.4.8 Conclusions: vortices embedded in jets

We began by generating a pair of sinusoidal jets of opposite sign, with an anticyclonic shear zone between them. The forcing method was that of direct mechanical forcing at the lower boundary, combined with thermal relaxation toward a geostrophically-balanced temperature distribution. The former had a typical coefficient of 0.001 per timestep, while the latter used a more realistic (i.e. much smaller) thermal damping coefficient. The model adjusted over 80 days or so, leaving the easterly jet almost unchanged but weakening the westerly slightly. The meridional PV gradient was steeper than the planetary value in the southern half of the channel, but flat in the centre and slightly negative toward the northern wall.

Vortices were placed in this background flow, initialized by means of spatially-localised heating over a period of 1 or 2 days. They tended quickly to adjust to an elliptical shape, due to the shear between the jets. Initial migration was blocked by the easterly jet, or in terms of PV, by the lack of a meridional PV gradient. Each vortex then steepened and the winds became stronger. As this happened, the vortex became able to overcome the PV barrier, migrating through the jet to reach the northern boundary. This experiment had some features in common with observations: elliptical shape of vortices, and qualitative evidence that jets do slow down vortex migration. However, the PV barrier was not large enough to confine the vortex for any length of time, so quantitative agreement was lacking.

Tuning the model parameters proved quite successful at controlling some aspects of the behaviour of vortices. With a shortened timestep (450 s) and large diffusive damping ($N = 30$ at fourth order), the steepening process was slowed to the extent that the vortex took more than 200 days to overcome the barrier. It maintained its appearance almost unchanged for more than 100 days, although the winds gradually increased. Next, the vortex grew larger around Day 200. By Day 240 it had penetrated the barrier: see Fig. 5.24. It did not seem possible to find a value of N which allowed a truly stable vortex; values greater than about 30 tended to cause eddies to dissipate.

Two hundred days (or about 7 months) is not a long period for a jovian vortex to persist. This is an undoubted failing in our model, despite much effort on attempts to improve it. This result is the best that could be done with the model in its present state. It is possible that more advanced physical parametrizations could be effective, although difficult to see how. It seems that the UM is prone to this phenomenon of vortex sharpening and intensification. Other models, such as that of Williams (1996), do not seem to suffer from this to the same extent. One possible reason is that our model is of higher resolution. Our choice of resolution (0.5°) does not seem excessive, indeed barely adequate given that very fine detail is observed in many places on Jupiter, for example in the filaments near the GRS. In respect of ability to resolve eddy motions, we may regard Jupiter as being

closer to the Earth's oceans than to its atmosphere. Like ours, ocean GCMs tend to make use of hyperdiffusion as a means of parametrizing sub-gridscale eddies.

Williams does note that vortex migration depends strongly on the vertical stratification in the model, and presents a number of cases with quite widely varying static stabilities. In our work we have, for better or worse, chosen a static-stability profile which is as realistic as possible, and not changed it. However, experiments with different stratification, and probably a deeper domain, will be a priority in future development (see Section 7.2). The choice of static stability profile is probably important for the results of Chapter 6 too.

A second background state was then tried. This consisted of a single strong easterly, and in the initial calculation no westerly at all. When this pattern was spun up in the GCM, a weak westerly did appear because of the "overshoot" of the shear zone from the easterly. The slightly diffusive model was unable to support the large shear present in the proposed reference state. Even the adjusted jet was rather unstable: after 90 days it suddenly weakened, immediately beginning to recover. It was also found to be capable of spontaneous eddy generation.

The PV profile of this second flow presented a much sharper barrier to the vortices, with a pronounced peak at 34°S. Even this change was not successful in trapping the eddy. Instead the latter always grew strongly, penetrating the barrier after only 100 days. This happened despite changes to the diffusion coefficient. We think this is because the baroclinically unstable easterly jet continues to feed energy into the vortex and therefore accelerates growth. Energy diagnostics tend to confirm this view. We saw in Fig. 5.25 that baroclinic energy conversion occurs as the vortex intensifies.

The horizontal structure of our vortices resembles (at least for several months) that of the elliptical giant eddies such as the White Ovals and Great Red Spot. The migration behaviour in the early stages is rather unrealistic, since the eddies tend to drift equatorwards, then "bounce" off the barrier, briefly return poleward and then resume the cycle. On the real planet the eddies remain at more or less the same latitude. Perhaps this suggests that our zonal PV structures are too steep, and yet they are inadequate to block the migration of intense developed vortices.

Perhaps the most interesting spontaneous phenomenon seen in the model is the development of vortices of each vertical structure: some confined to the first baroclinic mode, but others extending down to the base of the model. Thus in returning to our original question—shallow eddies or deep—the GCM results echo those of the PV inversion model, in that both seem to coexist and either may be generated by a deep, unstable jet.

Chapter 6

Equatorial waves and superrotating jets

What the waves were always saying. (Dickens: *Dombey and Son*, title of ch. 16)

6.1 Introduction

6.1.1 Motivation

In this chapter we move away from midlatitude dynamics, to look at the equatorial region. This poses different but equally interesting problems. The equatorial atmosphere has features not present elsewhere, including a class of waves whose amplitude decreases exponentially with latitude on either side of the equator. These equatorial waves have a major effect on the flow near the equator, both in the Earth's stratosphere and probably also on Jupiter.

Jupiter's equatorial region has received much attention in recent years, largely due to the Galileo Probe mission, which made the first direct measurements of the atmosphere, in a so-called North Equatorial Belt hot-spot (see e.g. Seiff et al., 1996). The Probe experiments were complemented by many other observations, e.g. from the Galileo Orbiter, both in the infrared (Roos-Serote and authors, 1998) and visible (Vasavada et al., 1998), and from ground-based instruments (Orton et al., 1996).

It is clear that there is great scope for interpreting the observations with a sufficiently realistic model. Modelling studies have not yet reached the stage where it is possible directly to compare model and observations in detail. However, the present model rep-

resents progress towards that aim. In order to develop a future model capable of quantitative comparison with observations, it will be necessary to understand the most basic properties of the flow rather better than we do at present. For this part of the thesis we have picked one such fundamental question: why does Jupiter have a westerly (as opposed to easterly) equatorial jet?

This question has the scientific merit of being easy to understand but difficult to answer. We shall argue (Section 6.3.1) that angular momentum considerations lead us to expect the jet to be easterly, and that therefore the observed westerly does need an explanation. This thesis is by no means the first attempt to address this question. Work on the problem dates back at least to the paper byMaxworthy (1975). We aim to build on earlier theoretical work, by making use of the more accurate dynamics and realistic vertical structure in our GCM. We approach the main question gradually, via model experiments on equatorial waves which serve the dual purposes of building our quantitative understanding of the concepts and of testing the model. This testing and development is necessary because in the chapter we use a rather different model configuration.

6.1.2 Overview

The chapter is rather long, with several different threads. All are related to the problem of the origin and shape of the superrotating equatorial westerly jet, and its interaction with equatorial waves. We present an overview here, and then throughout we will try to emphasize how each section relates to the others.

To summarize briefly: we explore the hypothesis (Maxworthy, 1975) that Jupiter's westerly equatorial jet is forced by waves propagating upward from the deep atmosphere. Dissipation deposits eddy momentum into the mean flow and thus accelerates the jet toward the zonal phase speed of the wave. The sources of the waves remain mysterious. Since visible cloud features move with the jet, the sources must be below and therefore not directly observable. One plausible candidate is organized moist convection in the troposphere, analogous to the Earth's Intertropical Convergence Zone, or ITCZ (see Section 6.11.1).

We will split up our main question and initially ask, "Can we build a model in which a plausible source generates a realistic equatorial jet?" The answer is yes, although this depends upon the strictness of our definition of "plausible". To strengthen this conclusion we then investigate the robustness of the result, trying to find out how tight are the constraints on the wave-source, such as its frequency and length scale. This is done by modelling the filtering effects of the atmosphere using a range of waves. In many experiments we force an infinite travelling wave near the bottom of the model, while in others we place a stationary oscillating source there. We have not yet explored the

third possibility, which is a travelling source, advected by the local zonal flow in the deep atmosphere.

We show that several types of source can theoretically generate a superrotating jet, with a maximum of zonal wind at the equator. We shall show that the details of the source seem to be less important than the filtering properties of the atmosphere. If a broadband source is present, the atmosphere seems to select and pass only those waves with speeds within quite a narrow band. If our model is correct, any deep source would be adequate, so long as the required phase speed (around 100 m/s) was present within its spectrum. The sources would not need to be pure or time-invariant, because other harmonic components would not be absorbed within the upper troposphere to force the jet.

The jet structure we obtain, with a single central maximum, is not quite realistic. In order to improve it, we consider the effect of solar heating in setting up a Hadley-style, axisymmetric angular momentum-conserving circulation. This secondary circulation modifies the main jet generated by wave forcing, producing a double-peaked structure of zonal wind as a function of latitude, similar to that observed. The model can generate realistic-looking flows, although this requires tuning of such loose parameters as the variation of solar heating rate with latitude. Without proper radiative-transfer calculations in the code, we can claim only that our models are plausible, but not necessarily correct, and that a reasonable heating distribution can lead to a realistic flow.

We begin in Section 6.2 by summarizing some of the observations of Jupiter's equatorial region. Then we discuss how the conservation of angular momentum places restrictions on equatorial superrotation in axisymmetric flow. This leads to the idea of wave forcing, introducing our survey of wave properties in the Jupiter GCM. Section 6.4 summarizes the theory we need to set up the model and interpret the results.

We then discuss (Sections 6.5–6.7.3) the development of the equatorial version of the model. Early versions (based on the midlatitude model) did not work at all well, so we describe the process by which the model was refined until it gave sensible results. This gives an explanation for all the choices of various parameters such as the domain, resolution, forcing form, drag parametrization, etc. We demonstrate, using straightforward Fourier analysis, that various pure travelling waveforms behave as expected (Section 6.9). The wave-mean flow interaction is noted: in accordance with theory, equatorial waves in a dispersive medium can inject momentum into the zonal-mean flow and spin up a jet.

In the next part (6.10) we explore in more detail the idea that Jupiter's equatorial jet is forced by waves from below. We investigate the filtering properties of the atmosphere by presenting the model with superpositions of waves. We find that the phase speed is the most important factor in determining which waves propagate into the stratosphere. Of many waves with equal phase speed, those of smallest wavenumber are preferred. This suggests that a source (such as organised convection) could generate a wide spectrum of

modes, but the jet and waves we observe at and above the cloud-tops depend more on tropospheric filtering than on the precise details of the source.

Section 6.11 looks at the waves and jets generated by various stationary sources: monotonic, bitonic, symmetric and antisymmetric. This is set in the context of previous work on the equatorial Pacific on Earth, in particular on Kelvin waves generated by the ITCZ (Holton, 1972). The spectra of different sources are examined, and it is found that antisymmetric sources tend to generate easterly jets and so can probably be ruled out for Jupiter. Symmetric sources near the equator can however generate suitable Kelvin waves.

In the final section, we introduce a modified Held-Hou analytic model of the Hadley cell, and show how its predictions suggest the addition of zonally symmetric heating to the GCM. This small additional forcing modifies the zonal-wind profile by means of a weak, angular-momentum conserving secondary circulation. We then tune the forcing parameters, until the model flow closely resembles the observed cloud-top zonal winds. Having achieved a variety of equatorial jets with plausible forcing mechanisms, we proceed to examine both the robustness of the result, and the properties of other waves and eddies in such an environment.

6.2 Observations of equatorial waves

There is a range of evidence for the existence of equatorial waves on Jupiter, from visible, infrared and radio-occultation observations. A striking feature of the equatorial region is the series of approximately 12 plumes, centered near 5°N and spaced nearly evenly around the planet. Visible images and IRIS measurements were taken by both Voyagers (Hunt et al., 1981), showing that the convective centre of each plume moved with the main westerly jet, whilst the surrounding material drifted westward at about 20 m/s relative to the centre. The IRIS measurements showed large-amplitude longitudinal thermal waves both in the deep troposphere and near the tropopause. There was generally no correlation between the position of the plume heads and the temperature peaks, but this could be ascribed to the low horizontal resolution of IRIS. Where the thermal structure of an individual plume could be examined, it was found to be colder above the centre of an active plume, consistent with upwelling motion.

Hunt et al. suggested that the plumes and associated thermal waves are generated by wave-CISK (Lindzen, 1974), by analogy with the Earth's tropics. The acronymn CISK refers to "convective instability of the second kind". The mechanism begins with instability triggered by low-level convergence in a conditionally unstable region. This causes air parcels to rise, and when they reach their level of instability they become convective, drawing up more air from below and establishing a cycle of feedback. Latent heat re-

lease from condensation warms the top of the plume, while evaporation cools the base and stablizes the temperature profile. Lunine and Hunten (1987) suggested a combination of narrow, rapidly-ascending moist plumes, surrounded by much larger dry regions of slow descent. This may be able to explain the apparent depletion of water in the equatorial region mentioned in Section 1.2. However, more detailed calculations by Del Genio and McGrattan (1990) later showed the mechanism to be quantitatively inadequate to produce sufficient drying.

A more familiar example of CISK is in some models of terrestrial hurricanes, in which the initial low-level convergence is caused by friction with the ground or ocean surface. In wave-CISK, on the other hand, the compression is achieved by advection due to the waves. This is perhaps responsible for establishing the Intertropical Convergence Zone in the Earth's tropics. We shall further explore the analogy between the ITCZ and Jupiter's equatorial region in Section 6.11. However, one must be wary of expecting too close a correspondence between moist convection on the Earth and on Jupiter. For example the molecular weight buoyancy contrast operates in the opposite sense (Del Genio and McGrattan, 1990). Terrestrial moist air is lighter than dry air, reinforcing the density contrast due to temperature. On Jupiter, moist air is denser than dry, which opposes the thermal effect on buoyancy.

Allison (1990) analysed the Voyager plumes with a linear model and concluded that they may be caused by an equatorial Rossby wave of wavenumber 11–13. At the time of the Voyagers, there was also a single plume in the southern equatorial region, which appeared to interact strongly with the flow at the adjacent GRS latitude. It may have been a wavenumber-1 Rossby wave. In terms of vertical structure, Allison was able to match the observations to a Rossby wave of equivalent depth $h \simeq 2$ km, trapped in a (slightly leaky) stable waveguide formed by a water cloud below about 6 bars.

The "active" plumes as seen by Voyager were always next to dark areas, known to infrared observers as the 5-μm "hot spots" since they are bright in this atmospheric window. These hot-spots have recieved much recent attention, since the Galileo probe entered one of them. Their longitudinal locations appear to be unpredictable, but probabilistic analysis (Ortiz et al., 1998) shows a correlation between regions with a high probability of containing a hot spot. These regions drift together in System III and could be explained by a Rossby wave of wavenumber 8 or 9. It therefore seems likely that two of the most conspicuous features of the equatorial region—plumes and hotspots—are linked to wave phenomena.

As direct evidence for equatorial waves, we have the Voyager radio occultation profiles (Lindal et al., 1981). The temperature-pressure profiles showed an oscillation with two components. The larger had an amplitude of approximately 10 K and vertical wavelength of $3H_0$ (H_0 being the pressure scale height). The smaller had amplitude 2 K and wavelength $0.5H_0$. A corresponding midlatitude profile showed no such waves, consistent with

the idea that these were equatorially-trapped modes. As Allison (1990) pointed out, wave motion is by far the most plausible explanation for these results.

Finally, the very existence of a superrotating equatorial jet may be good evidence for equatorial wave activity. Much of this chapter is devoted to exploring the hypothesis that the westerly equatorial jet arises primarily through momentum deposition by Kelvin waves. We shall try to identify the frequency and wavenumber of the waves which might be responsible for driving the jet, and speculate about what kind of wave sources are able to produce these waves. First, however, we outline the theoretical (and modelling) arguments in favour of wave forcing of the equatorial jet.

6.3 Equatorial superrotation

The central question of equatorial superrotation is, what causes the westerly winds seen at the equators of Jupiter and Saturn? A profile of zonal-mean zonal wind on Jupiter, as a function of latitude, is shown in Fig. 6.1. The answer to the question is less straightforward, for simple considerations suggest that easterly winds might be expected. We discuss this in physical terms and then present a more mathematical argument in the next section.

An atmosphere locally at rest (i.e. in solid body rotation with respect to an inertial frame) has a maximum of specific absolute angular momentum at the equator. Westerly jets at other latitudes may be excited by poleward motion of air parcels (as in the Hadley circulation, to be discussed later). Equatorward motion of parcels with zero zonal velocity would generate an equatorial easterly. Axisymmetric westerlies, on the other hand, have greater specific angular momentum than that of the planet (or atmosphere) beneath. What process is driving angular momentum up the gradient?

In the initial phase of the present work, this question was sidestepped. In the spirit of the rest of this thesis, the presence of the equatorial jet was assumed without explanation. The intention was to use model experiments to compare the properties of waves and eddies with those in midlatitudes, and to study wave-propagation as a function of the ambient flow. For this task, an *ad hoc* forcing method was to be used to generate the equatorial jet, just as with the midlatitude jets of Chapter 5.

This agnostic approach ran into major difficulties almost at once. It is not possible, even in the model, to drive a westerly equatorial jet with direct axisymmetric forcing, except perhaps by means of a 'body force' acting throughout the troposphere. We examine the theoretical reasons for this next.

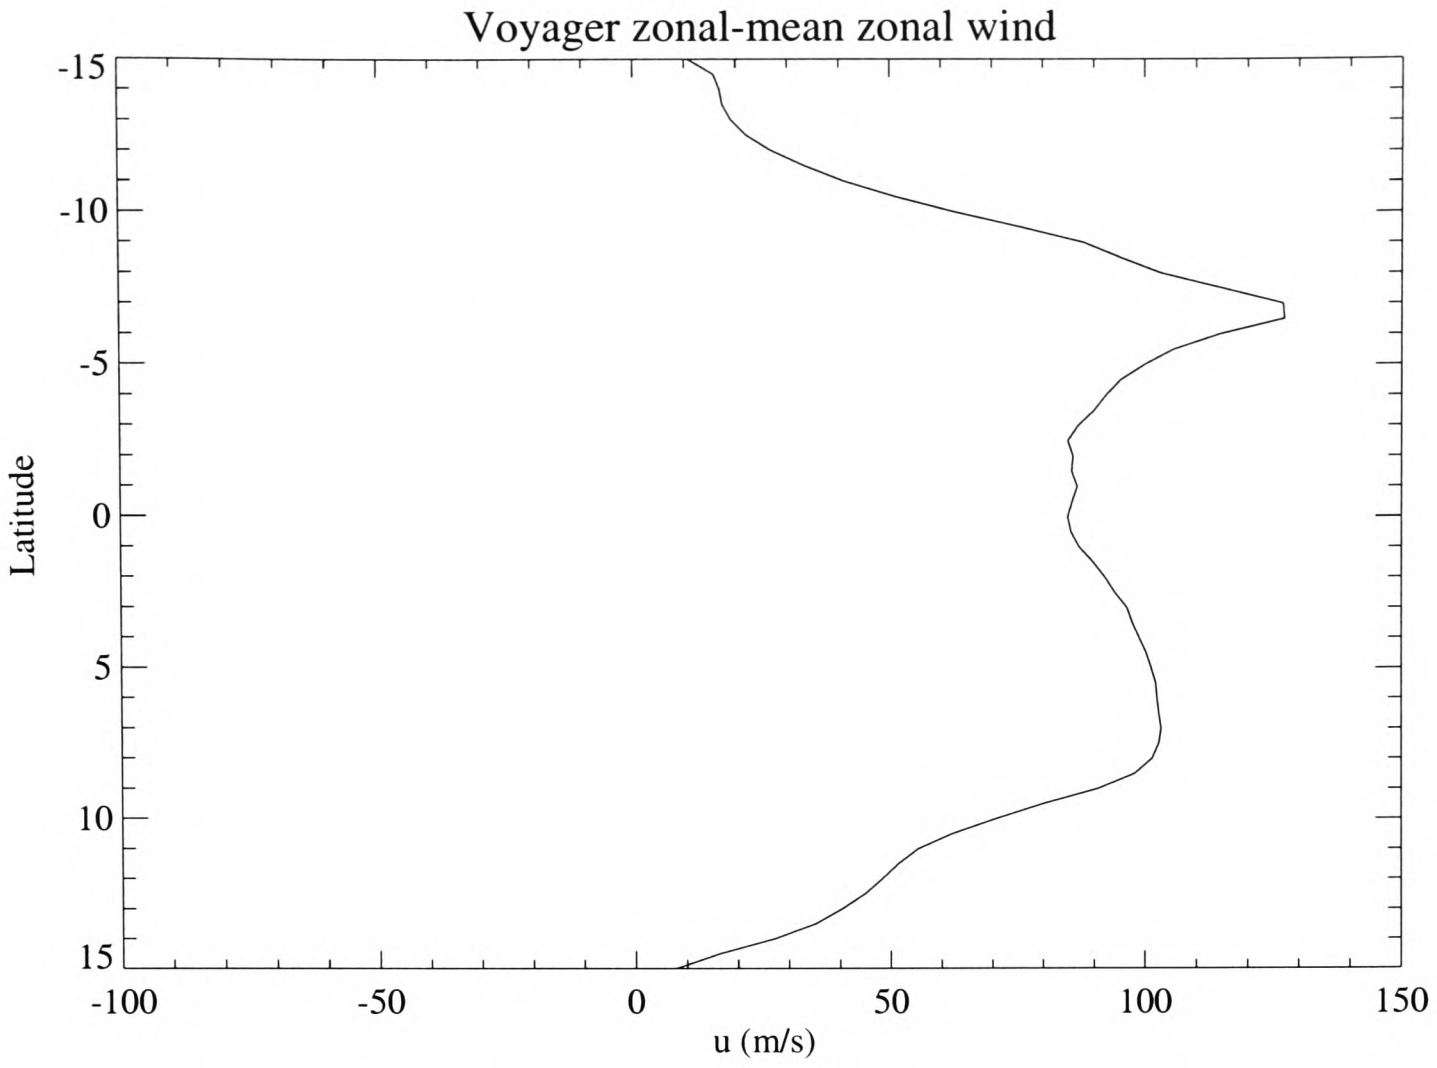


Figure 6.1: Voyager zonal-mean zonal wind, against planetographic latitude. Data from Limaye(1986).

6.3.1 Hide's theorem and equatorial superrotation

Hide's thoerem states that a prograde equatorial jet cannot be supported by purely axisymmetric processes, without particular initial conditions. The following argument was proposed by Hide (1969) and reviewed by Held and Hou (1980), and in a more general way by Read (1986). We give an outline of the analysis as it applies to Jupiter. Let the axial component of the specific angular momentum be M, where

$$M = \Omega R^2 \cos^2 \phi + uR \cos \phi \tag{6.1}$$

and $R(p)$ is the planetary radius. Note that we do not use the shallow-atmosphere approximation.

In the absence of viscosity and non-axisymmetric pressure gradients, M is conserved following the motion. Under these circumstances angular momentum cannot be created, only redistributed. Therefore both M for each fluid element, and its global integral, must be fixed at their initial values. Unless we choose an initial condition with a surfeit of angular momentum, the axial component can nowhere exceed the maximum value for solid-body rotation. This maximum value occurs high in the atmostphere, over the equator. Therefore symmetric processes cannot excite a prograde jet high in the equatorial atmosphere. We include the qualification "high in the atmosphere" since *downward* diffusion could theoretically lead to superrotation at lower levels. The decrease of density

with height means that massive downward diffusion would be required to increase the specific angular momentum significantly. This possibility seems unlikely on Jupiter.

Real atmospheres (and even our model) do not quite conform to these restrictions, primarily due to viscosity, non-axisymmetric eddies, and (in the case of the terrestrial planets) boundary-layer friction. Held and Hou (1980) argued that the diffusive angular momentum flux **F** would act down the gradient of M, and showed that a steady prograde jet (i.e. a local maximum in M) cannot exist under axisymmetric conditions.

Read (1986) pointed out the flaw in this argument, namely that in the Navier-Stokes equation, the flux **F** diffuses not down the gradient of angular momentum, but down that of angular velocity. In a deep spherical shell, such as the atmosphere of Saturn (Stevenson, 1982), we have

$$\mathbf{F} = -\mu R^2(\cos^2\phi)\nabla\gamma \tag{6.2}$$

where γ is the relative angular velocity,

$$\gamma = u/(R\cos\phi) \tag{6.3}$$

The projection P of **F** down the meridional m-gradient is

$$\begin{aligned} P &= -\mathbf{F}\cdot\nabla m \\ &= \mu|\nabla m|^2 - (2m/R)(m_R - m_\phi(\tan\phi)/R) \end{aligned} \tag{6.4}$$

Up-gradient diffusion can occur if $P < 0$, i.e. if $m_R > 0$. In a deep atmosphere, this depends on the component of **F** normal to the rotation which is locally vertical at the equator. Read also showed that in a shallow atmosphere, like that of the Earth, it is the locally horizontal component which is important. Jupiter's atmospheric aspect ratio is intermediate between those of the Earth and Saturn (Stevenson, 1982), so both components may be important.

The quasi-axisymmetric model of Mayr and Harris (1983) used very strong horizontal diffusion (as a parametrization of eddy motion) to generate prograde equatorial jets. These had a double-peaked structure resembling the observed zonal wind (Fig. 6.1), with a local minimum of $\bar{\gamma}$ at the equator. In Section 6.13 we propose the alternative mechanism of solar heating. In our model, this drives a thermally direct meridional circulation which redistributes angular momentum toward the tropics, creating the double peak.

Returning to our model for the gross structure of the westerly jet, we require a mechanism capable of driving fluxes of M up the vertical gradient. Such a model was proposed by Maxworthy (1975), who suggested that tropospheric sources generate Kelvin waves (e.g.

Holton, 1972) which propagate upward (against the M-gradient) and dissipate, depositing momentum in the stratosphere to form a steady westerly jet.

The jet's maximum speed is tuned to the phase speed of the wave. When these are similar, the vertical wavelength becomes short and the group velocity small, the dissipation has longer to act and therefore the interaction is greatest. The mechanism resembles that responsible for the Earth's quasi-biennial stratospheric oscillation (Plumb and Bell, 1982). Indeed there is evidence for a similar, quasi-quadrennial oscillation in Jupiter's stratosphere (Leovy et al., 1991; Friedson, 1998; Li and Read, 1999). An oscillation requires at least two wave modes in opposite senses; a steady jet may be achieved with just one.

We shall extend Maxworthy's work by modelling this phenomenon using the GCM. We test the hypothesis in three dimensions by investigating whether wave forcing (modified by Hadley-cell solar heating) can reproduce the observed winds in the model. Before describing this, we first revisit the theory of equatorial waves and justify our choice of parameters and experiments.

6.4 Theory of equatorial waves

The theory of Laplace's tidal equations was mentioned in Section 4.1, when we analysed the vertical structure equation in detail. We shall now outline the treatment of the horizontal part of the tidal equations and show how it leads to solutions for waves of interest, notably the equatorially-trapped Kelvin and Rossby-gravity modes. For a detailed treatment see Andrews et al. (1987). The original work was largely by Matsuno (1966). It should again be stressed that this analysis assumes a basic-state at rest ($\bar{u} = 0$). Our models are initialized in this state, but the theory is not strictly valid once a zonal-mean flow has begun. It can be extended by regarding the flow as acting to Doppler-shift the waves, but even this requires that the shear not be too great.

6.4.1 The horizontal tidal equation

Given the horizontal tidal equations, we look for sinusoidal wave solutions of the form

$$\Phi = \hat{\Phi}(\phi) \exp i(s\lambda - 2\Omega\sigma t) \tag{6.5}$$

where Φ is the perturbation geopotential, s is the zonal wavenumber, and $(2\sigma)^{-1}$ the period in days.

The meridional part of the solution $\hat{\Phi}$ satisfies the eigenvalue equation

$$\mathcal{L}\hat{\Phi} + \gamma\hat{\Phi} = 0 \tag{6.6}$$

Lamb's parameter $\gamma = 4\Omega^2 a^2/gh$ and the operator $\mathcal{L}$ is defined as

$$\mathcal{L} \equiv \frac{d}{d\mu}\left[\frac{(1-\mu^2)}{(\sigma^2-\mu^2)}\frac{d}{d\mu}\right] - \frac{1}{\sigma^2-\mu^2}\left[\frac{-s(\sigma^2+\mu^2)}{\sigma(\sigma^2-\mu^2)} + \frac{s^2}{1-\mu^2}\right] \tag{6.7}$$

where $\mu \equiv \sin\phi$.

Equation (6.6) has eigenfunctions $\Theta_n^{\sigma,s}$ with eigenvalues $\gamma_n^{\sigma,s}$. These are bounded at the poles, as required. They are known as Hough functions and have been tabulated by Longuet-Higgins (1968). Different limiting values of Lamb's parameter give different classes of waves. For example, as $\gamma^{-1/2} \to \infty$ we have gravity waves. As $\gamma^{-1/2} \to 0$ we see eastward-propagating Kelvin waves and westward-propagating Rossby waves and mixed Rossby-gravity waves. We are most interested in the latter regime, for as $\gamma^{-1/2} \to 0$ all modes are confined near the equator.

Andrews et al. (1987) analysed equatorial waves by considering the linearized primitive equations on an equatorial beta-plane. Assuming constant buoyancy frequency N, zero background wind and no dissipation or forcing, they sought solutions of the form

$$(u, v, w, \Phi) = e^{z/2H}\mathrm{Re}\left\{\left[\hat{u}(y), \hat{v}(y), \hat{w}(y), \hat{\Phi(y)}\right]\exp i(kx + mz - \omega t)\right\} \tag{6.8}$$

Substitution into the linearized primitive equations gave

$$-i\omega\hat{u} - \beta y\hat{u} + ik\hat{\Phi} = 0 \tag{6.9}$$

$$-i\omega\hat{v} + \beta y\hat{u} + d/dy(\hat{\Phi}) = 0 \tag{6.10}$$

$$ik\hat{u} + d/dy(\hat{v}) - i\omega m^2 N^{-2}\hat{\Phi} = 0 \tag{6.11}$$

The set of equations (6.8–6.11) may be solved to obtain the modes we require.

6.4.2 The equatorial Kelvin wave

Two modes particularly interest us: the Kelvin and mixed Rossby-gravity modes. These are designated $n = -1$ and $n = 0$ respectively in the notation of Matsuno (1966). One reason for concentrating on these two is that they tend to have vertical wavelengths comparable to a scale height. Modes of higher n have much shorter vertical wavelengths so are harder to resolve, both in observations and models. Another reason is that they have opposite phase velocities. In future models of stratospheric oscillations we will need a minimum of one eastward and one westward wave.

We obtain the solutions by noting that there is zero cross-equatorial flow in a pure Kelvin wave, i.e. $\hat{v} = 0$. We may substitude this into Equations (6.8–6.11) to yield

$$\omega\hat{u} = k\hat{\Phi};\quad \beta y\hat{u} = -d/dy(\hat{\Phi});\quad k\hat{u} = \omega m^2 N^{-2}\hat{\Phi} \tag{6.12}$$

Elimination of $\hat{\Phi}$ gives the dispersion relation $\omega = \pm Nk/m$ and the vertical group velocity $c_g = \mp Nk/m^2$. For upward-propagating energy, we choose the positive root for c_g, and hence $\omega = -Nk/m$.

We may find the meridional form if we eliminate u and solve for $\hat{\Phi}$: we have a first-order ODE whose solution is the Gaussian

$$\hat{\Phi}(y) = \hat{\Phi}_0 \exp(-\beta k y^2/2\omega) \tag{6.13}$$

where $\hat{\Phi}_0$ is a constant amplitude.

Only solutions decaying with y are physically acceptable for equatorially confined waves. This imposes the constraint that $\omega/k > 0$, i.e. all Kelvin waves propagate eastward. In a background flow, the frequency is Doppler shifted ($\omega' \equiv \omega - k\bar{u}$) such that the phase speed is always eastward with respect to the flow. If the basic zonal wind $\bar{u}$ exceeds the absolute phase speed $c \equiv \omega/k$ the Kelvin wave cannot propagate. Equation (6.13) is the beta-plane approximation to the relevant Hough function, and easier to evaluate.

The latitudinal width of the wave varies with phase speed. For $\omega/k = 100$ m/s, $\beta k/2\omega = 2.46 \times 10^{-14}$ m^{-2}. Near the equator, $y \simeq 1256000\phi$, so the wave amplitude is proportional to $\exp(-\phi^2/25)$. Therefore the half-width at half maximum is only $4°$ of latitude. Jupiter's equatorial jet has a HWHM of about $10°$: further evidence for a secondary mechanism, redistributing momentum horizontally.

6.4.3 The mixed Rossby-gravity wave

The analysis of modes with nonzero $\hat{v}$ is slightly more complicated. Here we shall simply quote the results of Andrews et al. (1987). The mixed Rossby-gravity, or MRG wave, is the gravest mode and has dispersion relation

$$m = -\mathrm{sgn}(\omega)\frac{N}{\omega^2}(\beta + \omega k) \tag{6.14}$$

and meridional solution

$$\hat{\Phi} = \omega y \hat{v}_0 \exp\left(\frac{-\beta|m|y^2}{2N}\right) \tag{6.15}$$

The presence of the factor y outside the Gaussian means that the MRG mode is anti-symmetric about the equator. It has a westward phase speed. The mutual orthogonality of the structures of the MRG and Kelvin wave implies that it is possible to force each independently of the other, and with controlled relative amplitudes as well as different frequencies and wavenumbers.

In a zonal shear flow, the MRG wave has propagation properties similar to those of the Kelvin wave. It cannot propagate in a westerly flow faster than its phase speed. The

dissipation becomes a maximum when the phase speed matches the basic wind speed. When this happens, the vertical group velocity becomes small, allowing longer time for dissipation to act and hence greater transfer of momentum.

6.5 Forcing functions for Kelvin waves

We now turn to the practical issues involved with modelling equatorial waves in a GCM, beginning with the choice of forcing functions. Initially we will consider a flow forced by a single Kelvin wave. This has an eastward phase speed, so is expected to excite an eastward (i.e. westerly) equatorial jet, hopefully similar to that observed on Jupiter. The theoretical structure of the Kelvin wave was discussed in Section 6.4 above, using the beta-plane approximation. The zonal part is a travelling sinusoid, while the meridional part is a Hough function (Longuet-Higgins, 1968).

We wish to find a convenient forcing function which closely resembles the Kelvin wave, in order to generate a pure response. In beta-plane geometry the equivalent meridional function is simply a Gaussian, Eq. (6.13). Although our model geometry is spherical, we use this form as a convenient and sufficiently accurate approximation to the Hough functions in equatorial latitudes.

The vertical distribution of heating on Jupiter is largely unknown. We would expect some heating to occur around the level of the water clouds (i.e. a few bars) but we do not know the vertical scale. It may be possible to constrain the latter, by observing the waves in the stratosphere. Chang (1976) found that in a barotropic linear model, the vertical wavelength tended to be around twice the vertical scale of heating. However the presence of vertical shear, and particularly a jet of speed comparable to the wave phase speed, will affect this result. Lindzen (1971) showed that vertical wind shear will reduce the vertical wavelength.

The experiments in this chapter use a vertical heating function proportional to a half-wavelength section of $\cos(z - z_0)$, where z_0 is around the 4 bar level, i.e. near the probable water cloud. The heating function is evident in the model temperature field, shown in Fig. 6.2 for example. The zonal structure is the travelling sinusoid discussed above. Later we shall consider stationary sources, which excite a spectrum of standing waves and bear more resemblance to real physical possibilities.

6.6 Implementation and choice of wave parameters

The travelling Kelvin wave forcing function discussed above may be written

$$\Delta T = A\cos(z - z_0)\exp\left(-\phi/\Gamma\right)^2 \sin\left(kx - \omega t\right) \tag{6.16}$$

where A is the forcing amplitude (typically 0.04), Γ is the meridional width of the forcing, k is the wavenumber and ω the angular frequency. The equatorial Kelvin wave is meridionally trapped, with a width depending on the phase speed (see Eq. 6.13). A wave of phase speed $c = 100$ m/s has $\Gamma = 6$ (degrees^{-1}); one with $c = 145$ m/s has $\Gamma = 7$ (degrees^{-1}). The slowest waves in our models, with $c \simeq 50$ m/s, have widths of $\Gamma \simeq 4$. We choose Γ in the heating function to be a close estimate of the width of the wave, but in practice it is found not to be critical. Tests show that the wave width is not very sensitive to the heating width: reducing the latter has little effect on the results. Nevertheless, we choose an appropriate heating function for each experiment, particularly when forcing more than one phase speed at once. In these superposition experiments (Section 6.10), each component is forced separately, with its natural width and phase speed.

To set up and interpret experiments, it is simplest to think of k in units of wavenumber per channel-length (i.e. the number of waves which can fit into the channel), and of ω in terms of τ, the period in days. To find the phase speed, they must be converted to SI units. We express the forcing at column I, day T as

$$\Delta T = A\sin 2\pi\left(n\frac{I}{N_I} - \frac{1}{\tau}\frac{T}{N_T}\right) \tag{6.17}$$

where n is the zonal wavenumber, τ the period in days, I the gridpoint number, and T the time. N_I and N_T are respectively the number of gridpoints in the channel and the number of seconds per day. This means that $x = I\Delta x$ and $t = T\Delta t$. Then $k = 2\pi n/(N_I\Delta x)$ and $\omega = 2\pi/(\tau N_T)$. The phase speed c is given by

$$c = \frac{\omega}{k} = \frac{N_I\Delta x}{n\tau\Delta t} \tag{6.18}$$

For example, the first experiment has $n = 2$, $\tau = 4$ days, with a channel length of 30° and 60 gridpoints so $\Delta x = 627500$ m. These give $c = 55$ m/s. If we halve n or τ, c becomes 110 m/s. This is the phase speed of the travelling forcing wave: it is not the same as the phase speed of the responding Kelvin wave, which also depends on the vertical wavenumber. The dispersion relation is given by linear theory (Section 6.4.2) as $\omega = -Nk/m$, where m is the vertical wavenumber, and N the Brunt frequency. We can see that, keeping the vertical structure of the forcing constant, there is a linear relation between ω, k, and hence c. This allows us to control c even if we cannot accurately predict m.

We mentioned that the forcing amplitude is typically 0.04 and the frequency $4(\text{days})^{-1}$. To estimate the heating rate, we note that this increment is applied every timestep, i.e. 96 times per day. The temperature anomaly rises from zero to its maximum over a quarter-cycle, or one day. Therefore the mean heating rate is roughly 0.04×96, or almost 4 K per day. This seems quite large but not implausible. We have not conducted formal tests to establish the long-term effect of reducing the heating rate. The model jets certainly spin up rather fast, so we could reduce the heating rate in future experiments if desired.

6.7 Model domain, resolution and drag parameters

The models discussed in Chapter 5 used parameters suitable for the study of waves and vortices. The configuration was a midlatitude channel of around 31° latitudinal width and 30° or 60° longitudinal length. The horizontal resolution was half a degree, with 30 vertical levels covering four decades in pressure from 10 bars. This arrangement has proved suitable for resolving eddies in three dimensions. The domain is sufficiently large that the eddies remain far from the sidewalls, but is not wastefully excessive.

In this section we shall describe and justify a different configuration which is more suited to studies of equatorial waves and jets. We start with the design mentioned above, and progressively modify it to find one more nearly optimal. Our criteria are that the model should have sufficient resolution and domain size, so that the results are consistent and do not depend on the precise values of these parameters. Conflicting with this is the need to minimize computational workload.

The models may be tested for sufficient resolution by examining the gross features of the results, and also by checking the speed of the jet spun up by the waves. One way to do this is by varying the channel length and hence altering the phase speed of the forcing wave. Doubling the channel length while keeping wavenumber and frequency constant will double the phase speed. The effect on the jet speed should be similar to that of doubling the frequency, at constant channel length and wavenumber. This is because the wave will accelerate the jet until it is tuned to the same speed as the wave. The vertical structure will not be quite the same, as the vertical wavelength will differ.

6.7.1 Stratospheric drag coefficient

Before assessing the effect of resolution changes, we must establish a suitable stratospheric drag parameter. Drag of some sort is necessary because in a non-dissipative system, the waves do not (to first order) affect the mean flow. In the present model with simple physics, we do not attempt to pin down the mechanisms responsible. We merely

use a linear drag scheme, tuned to find the smallest drag coefficient capable of causing momentum deposition into the zonal-mean flow on a convenient timescale. The standard model has drag or forcing at the bottom boundary, and a sponge layer in the top three levels. The latter is only there to damp spurious downward reflections from the rigid lid, and we aim to minimize its effect on other features of the model flow, such as upward-travelling waves.

The results of the first experiment are shown in Fig. 6.2. This is a "heavy drag" scenario, with coefficient 0.001 per timestep in the stratosphere and also on the lowest two levels (to simulate an interior at rest). The drag is zero in the troposphere and uniform in the stratosphere for most of our runs, although we sometimes use a smoothly-varying drag parameter. The wave is partly absorbed in the troposphere, accelerating the zonal jet, but partly propagates through to the stratosphere. The meridional structure of the zonal-mean zonal wind has a triple peak. This is unexpected, since the Kelvin wave forcing function (6.16) has a single peak centred on the equator.

It seems unlikely that an equatorial wave could excite such a jet. Previous models (Maxworthy, 1975; Li and Read, 1999) show jets whose meridional structure is close to that of the forcing wave. Therefore we regard the triple-peak jet as probably unphysical. We think it is caused by reflection from the sponge layer, and in Section 6.7.2 below we show how it can be removed by raising the model lid. Note also that the zonal jet does not extend into the stratosphere at all. This is a consequence of the overly heavy damping.

The results of a run with light drag are shown in Fig. 6.3. The stratospheric damping coefficient was 1×10^{-5} per timestep, a factor of 100 smaller than the previous run. The response in stratospheric temperatures is much larger—to the extent that the tropospheric forcing no longer shows up on the plot. The pattern of winds is a series of alternating jets, stacked vertically one on another and extending to the top of the model. The triple-peak pattern with sidelobes is still evident.

The large amplitude of the stratospheric wave in Fig. 6.3 implies that this drag coefficient is probably too small; the wave is propagating freely and depositing little momentum until it reaches the sponge layer at around 30 mb. This is unsatisfactory because the sponge layer is a computational device, with no physical justification beyond damping of reflections. It therefore seems sensible to pick a damping coefficient intermediate between the two values tested. More important is the issue of the jet shape, and the unphysical sidelobes in the jet. We shall now consider how to improve this by altering the model domain and resolution.

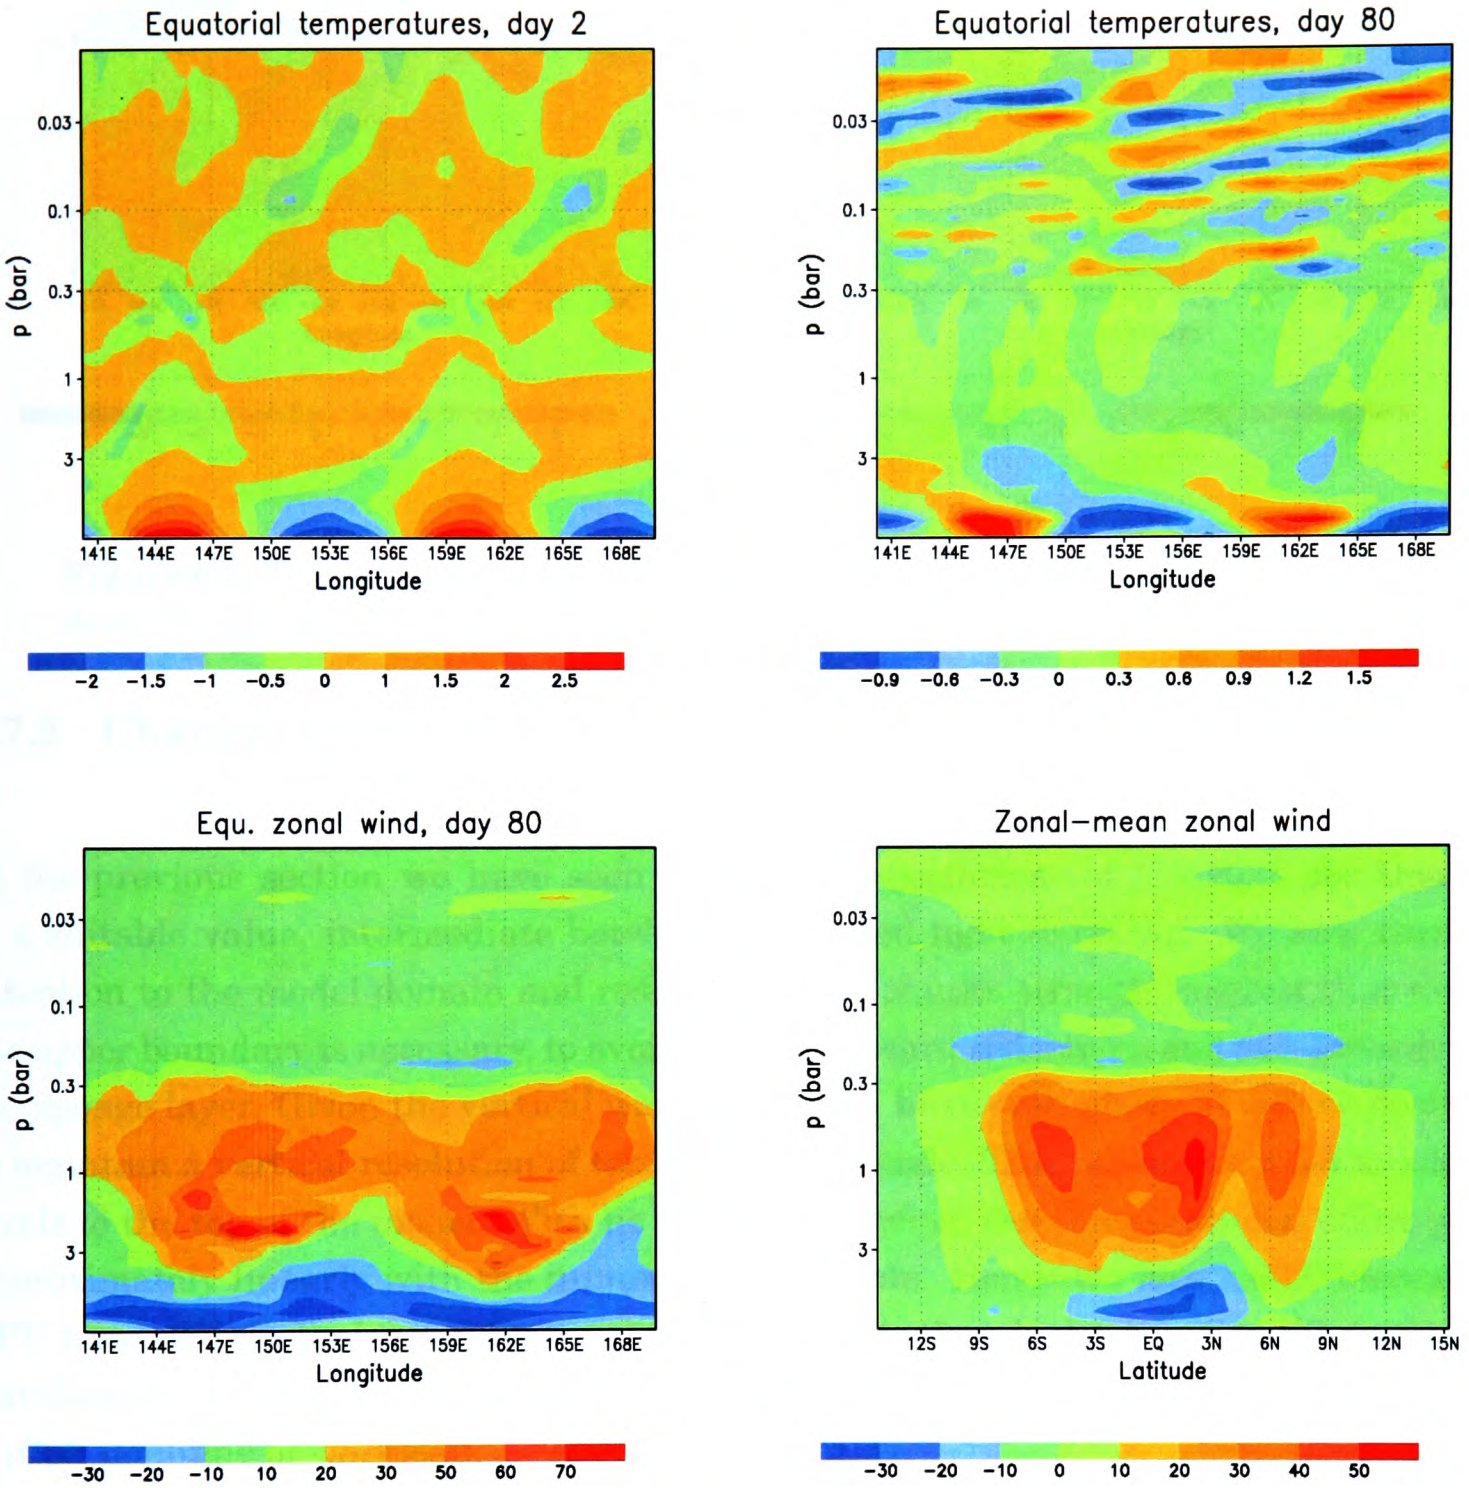


Figure 6.2: Results of the first equatorial wave experiment. The drag coefficient is 0.001 per timestep, or 0.096 per day. The top panels show the temperature anomaly field (zonal-mean subtracted) on the equator. On day 1 the wavenumber 2 forcing is clearly evident, but the response is not yet clear. By day 80 we can clearly see the wave propagating upward and eastward in the stratosphere. The lower left plot shows zonal wind on the equator. The wavenumber 2 structure is still evident but has largely smeared out into a uniform jet. The lower right plot is a cross-section of zonal-mean zonal wind.

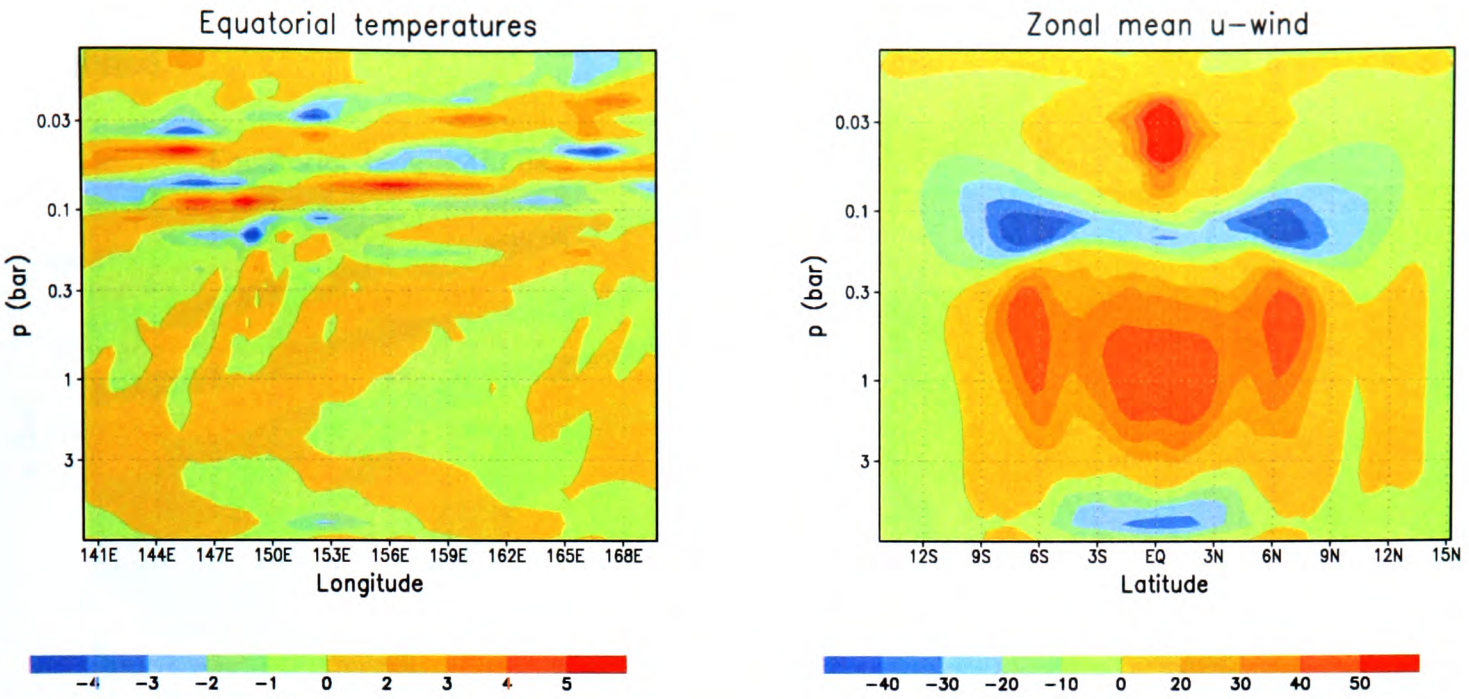


Figure 6.3: Temperatures and winds after 40 days, for a run with light stratospheric drag. The jets now extend into the stratosphere.

6.7.2 Changes to resolution

In the previous section we have seen that a drag coefficient of 1×10^{-4} per timestep is a suitable value, intermediate between heavy and light damping. We now turn our attention to the model domain and resolution. The results strongly suggest that raising the upper boundary is necessary, to avoid both downward reflections and excessive drag in the sponge layer. Given the vertical wavelengths we have seen above, it will be desirable to maintain a vertical resolution of ten levels per decade. This means we need to add ten levels to the top of the model. With no radiation scheme, the computational effort grows approximately linearly with the number of gridpoints. Hence we face a 30% increase in CPU time. We shall also need to use longer channels than hitherto, since the horizontal wavelengths are longer than the size of the eddies studied earlier. This implies at least a further doubling of workload.

In view of these factors, we need to ask whether the horizontal resolution could be degraded without damaging the results. This seems likely to be the case, for the fine structure of the flow is now less important and a one-degree grid should be adequate to resolve features of interest. To test this, we repeat the previous experiment at lower horizontal resolution, but the same channel length of 30°, and the same wavenumber and period of forcing. The damping factor is chosen to be intermediate, 1×10^{-4}.

Figure 6.4 shows the temperature anomaly field above the equator after 80 days. We see a wavenumber-2 signal at the base of the model, and also propagating upward and eastward in the stratosphere. The right-hand panel shows the zonal mean zonal wind as a function of latitude and height at the same date. Again, the most obvious feature is the triple peak around 500 mb. The change in horizontal resolution has not led to any

dramatic change in the results, so we may adopt this one-degree resolution with some confidence.

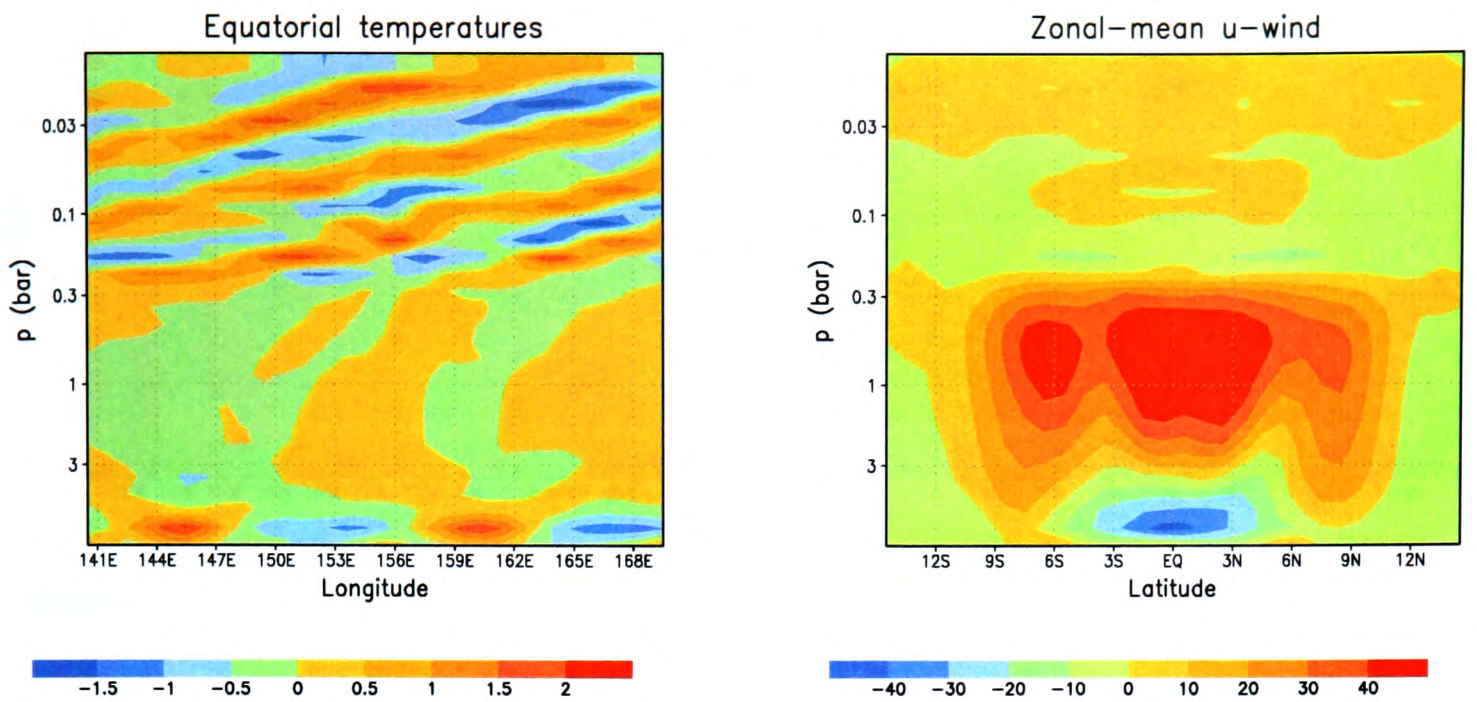


Figure 6.4: Equatorial temperature anomaly and zonal-mean zonal wind fields after 80 days, in a model with lower horizontal resolution. The channel length is 30 degrees.

Figure 6.5 shows, for comparison, the results of a model identical except for the channel length, which is 60 degrees. The temperature anomaly plot shows that the wavenumber is still +2, and the response amplitude is roughly five times larger. The zonal jet is rather stronger, but only by around 50%. There is now clearly a central wave, consisting of easterly and westerly phase alternating with height. The sidelobes are still present, but we will now see how they may be removed, by raising the model lid.

6.7.3 Models extending higher into the stratosphere

We now extend the model to a top level of 1 mb, by adding ten levels to the top and retaining the level spacing. The sponge layer is still the top three levels, so is moved upward by a decade in pressure, or roughly two scale heights. The basic-state temperature profile is extended upward with a small positive lapse rate dT/dz. The resulting profile (Fig. 6.6) approximately follows Voyager IRIS and radio occultation measurements (Hanel et al., 1979a,b; Lindal et al., 1981). The model runs were repeated in the new domain, and the results appear in Fig. 6.7.

We immediately notice that the side lobes have disappeared, and the flow consists of a stack of jets peaking only at the equator. This agrees with the theoretical structure of a Kelvin wave (Matsuno, 1966) and with other models such as that of Li and Read (1999). The maximum jet speed is now much higher, at 120 m/s. The next test of the model is, as before, to double the phase speed and compare the resulting flow. This may be done

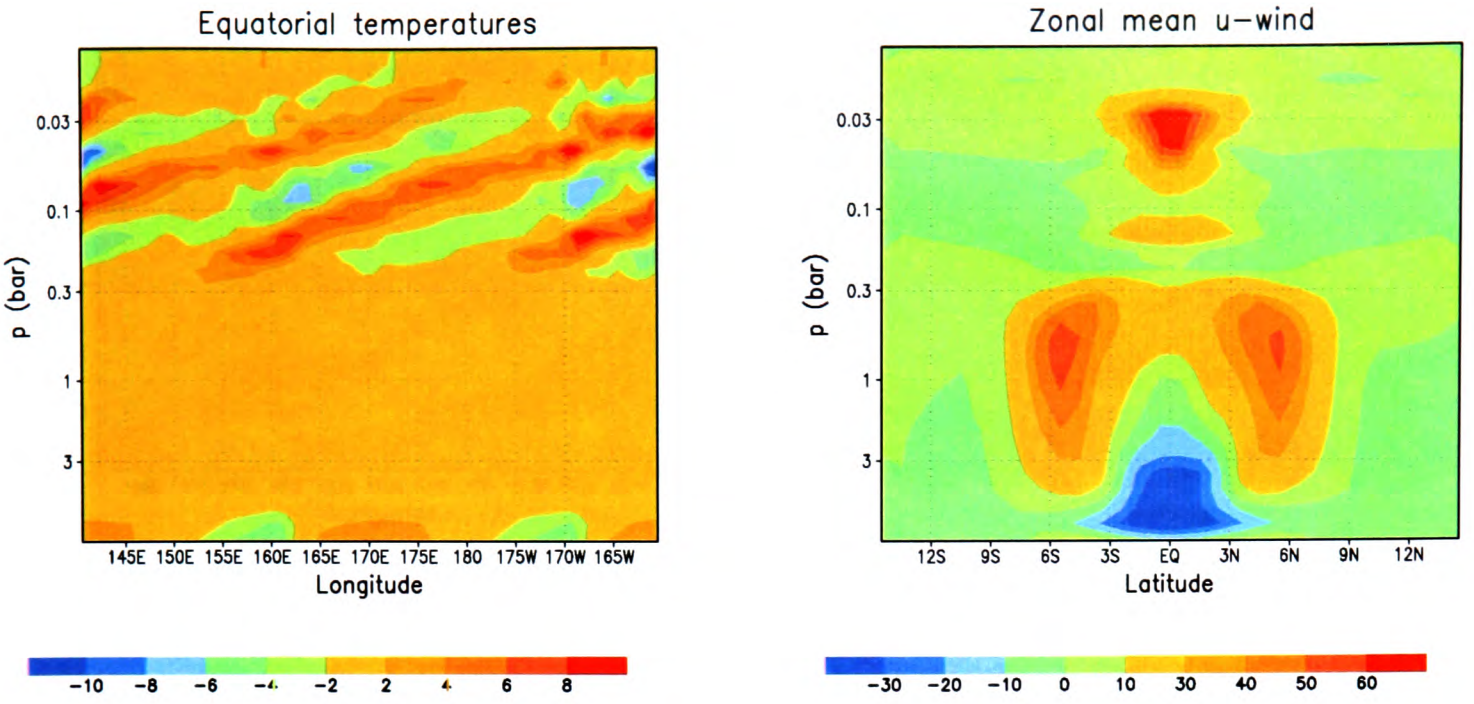


Figure 6.5: Results of the same experiment in a channel of 60 degrees. These differ from the prediction of a jet of similar structure but double speed: the jet speed is almost unchanged.

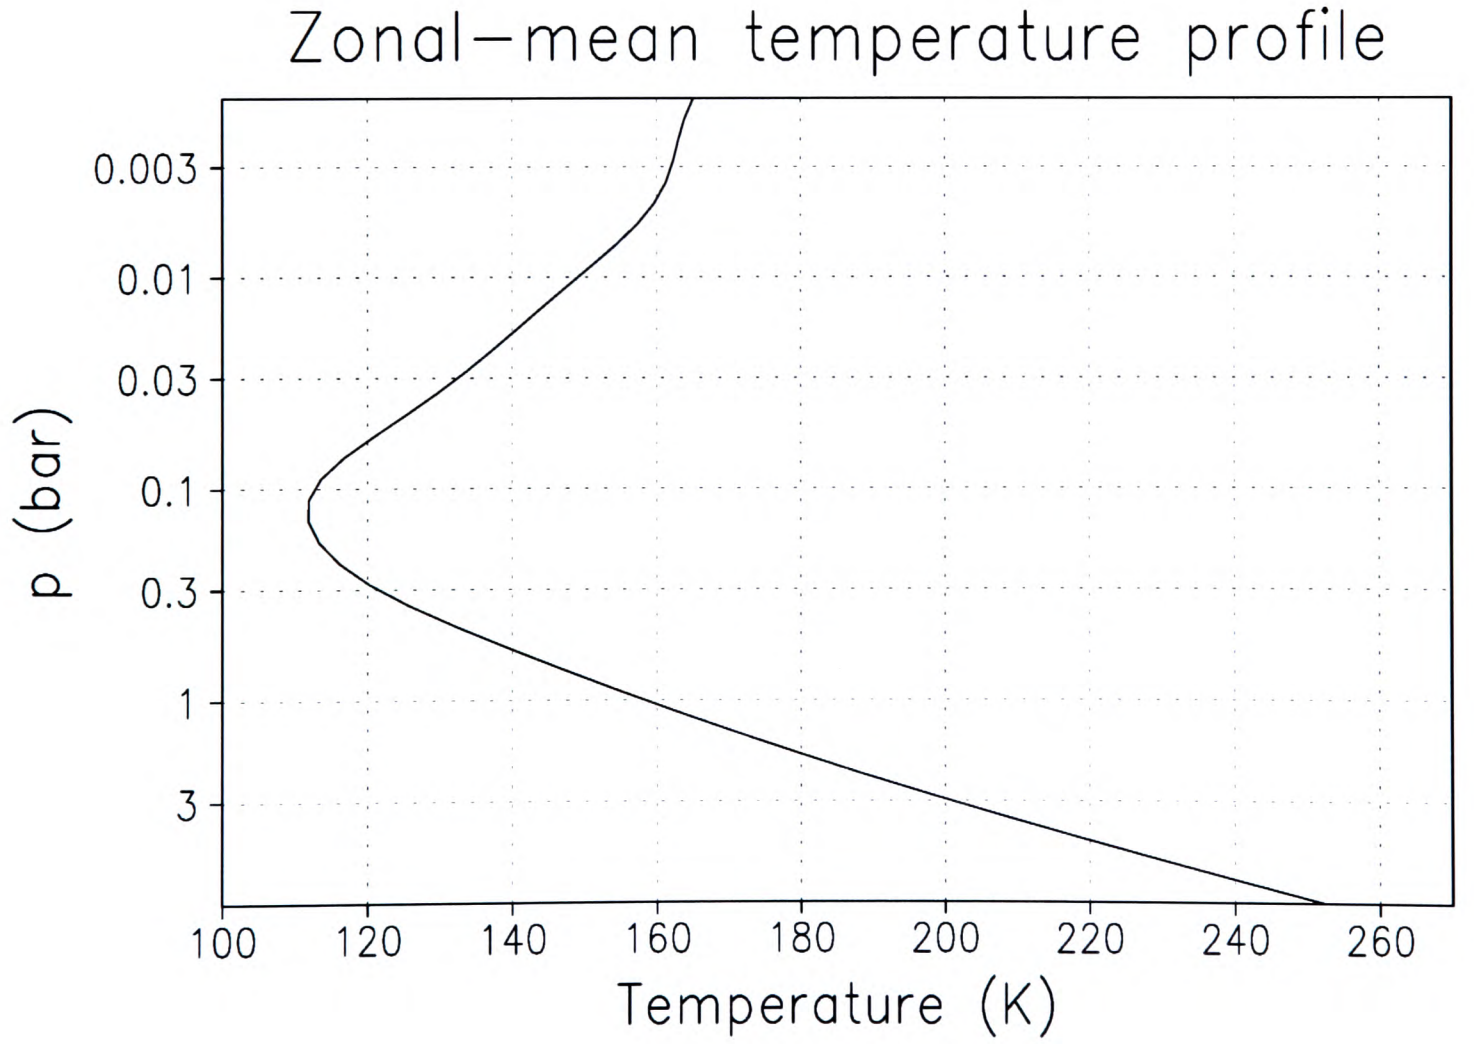


Figure 6.6: Basic state temperature-pressure profile for 40-level models, extending from 10 bars up to 1 mb.

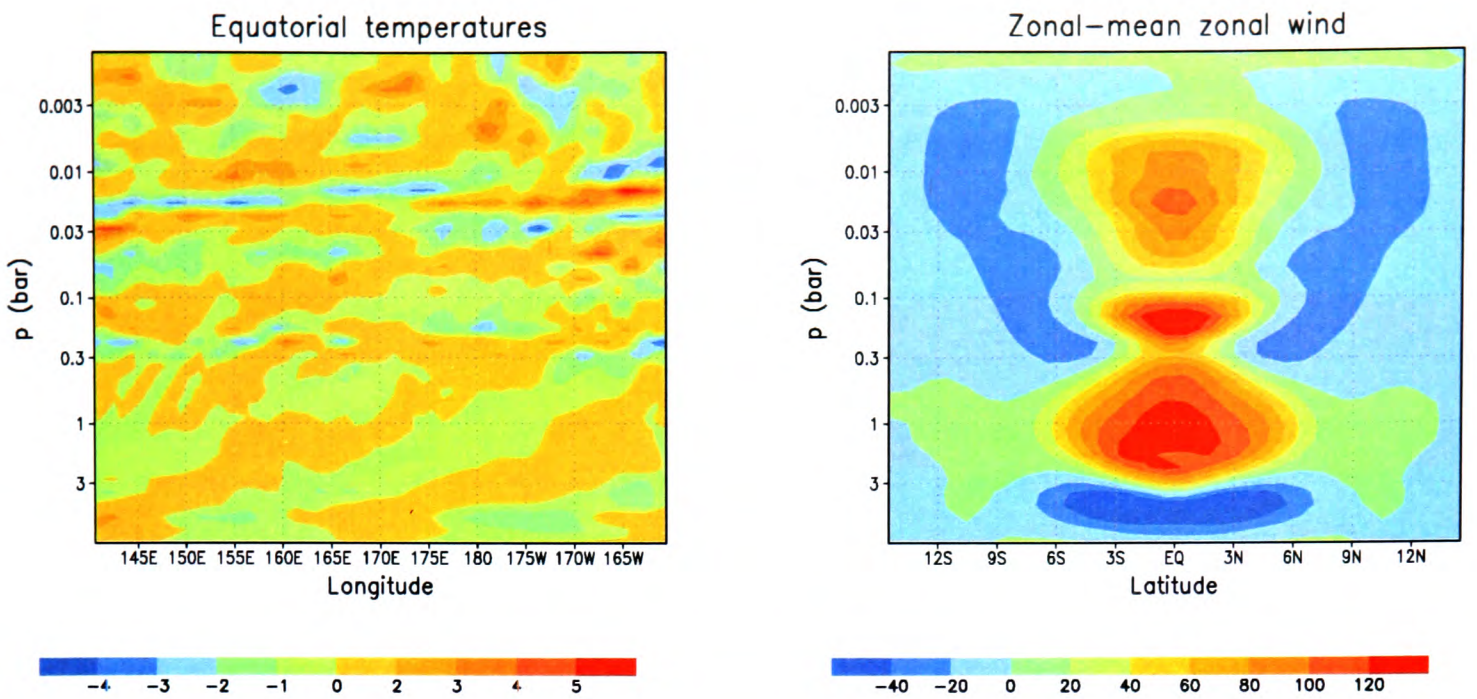


Figure 6.7: Equatorial temperature anomaly and zonal-mean zonal wind fields in a deeper model, after 80 days. The channel length is 60 degrees and the model extends up to the 1 mb level.

either by doubling the channel length to 120° (Fig. 6.8) or by halving the forcing period to 2 days (Fig. 6.9). The former method gives a jet of 240 m/s, i.e. exactly double the previous one. The vertical structure of the flows are different, which we expect because the vertical wavelength will be correlated with the horizontal one. The latter experiment gives a slightly weaker jet, of just over 200 m/s, but the agreement is still reasonably good.

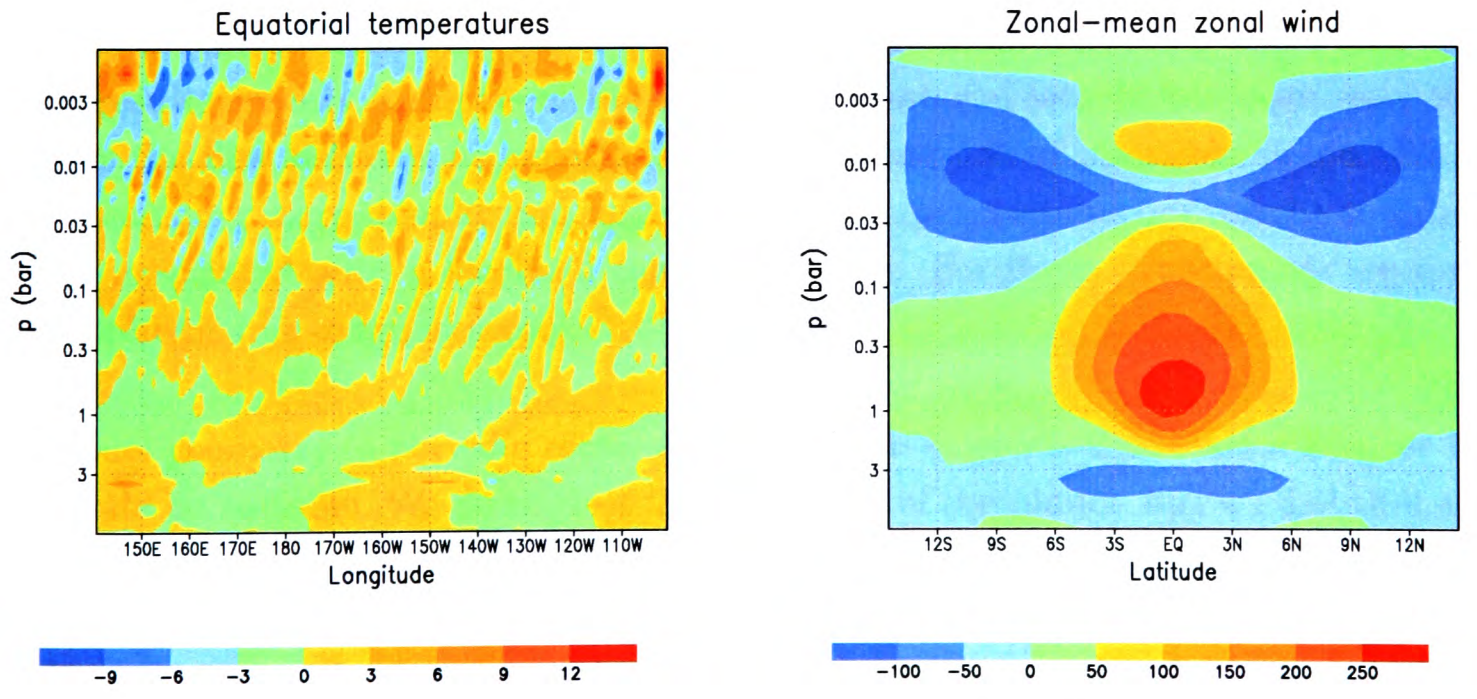


Figure 6.8: Results of the same experiment as in Fig. 6.7, in a channel of 120 degrees. This time the results agree with the prediction. Although the vertical structure is somewhat different, the jet speed has increased linearly with wave phase speed.

This version of the model behaves as expected, in the simple case of motions forced by a pure Kelvin wave. We may now examine more complicated cases including mixed Rossby-gravity waves, wave superpositions, and zonal-mean Hadley cell forcing. Before leaving

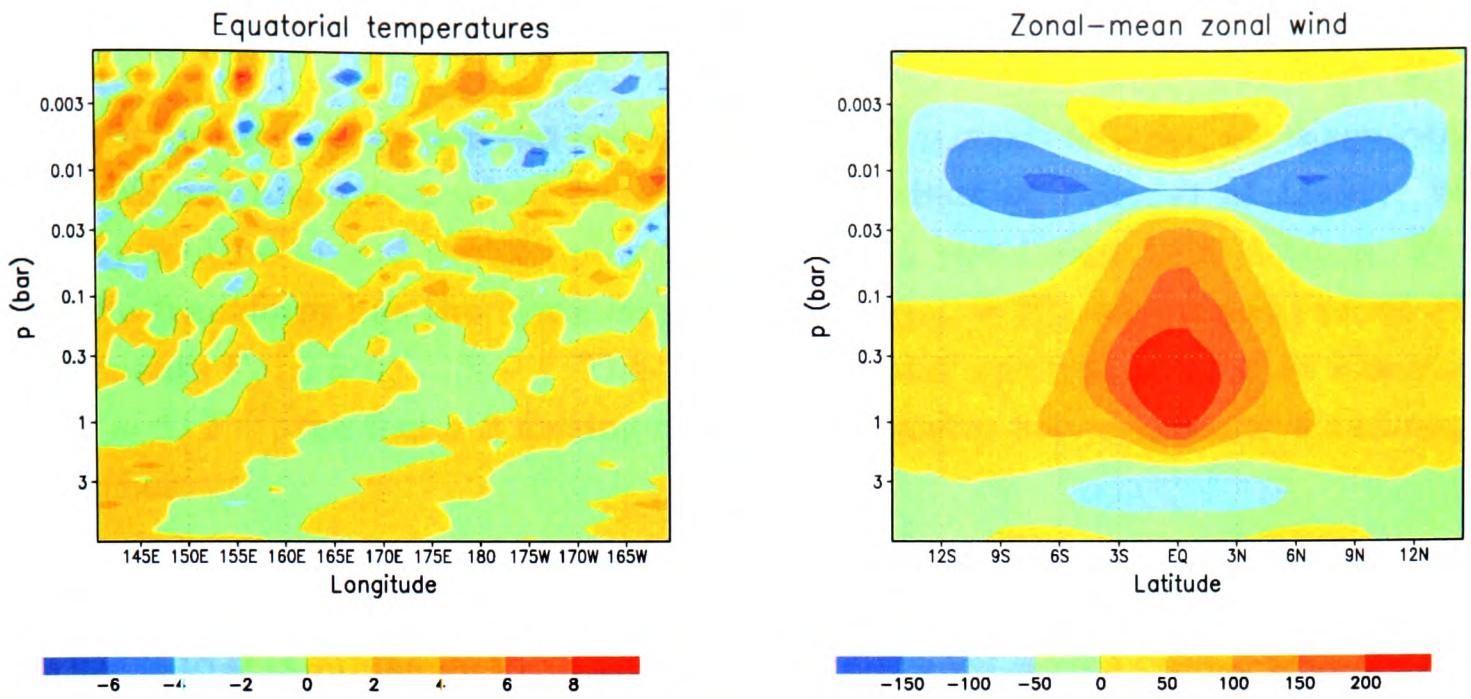


Figure 6.9: Results of the same experiment as in Fig. 6.7, but with a period of 2 days. Again the the vertical structure is different, but the jet speed is almost the same as in Fig. 6.8.

the single Kelvin wave, we will briefly consider a simple Fourier method for analyzing these model results.

6.8 Fourier analysis of waves in the model

Hitherto we have forced the model with a monochromatic travelling wave and looked at the results mainly in terms of the resulting equatorial jet. It has been easy to see whether the flow behaves as expected, i.e. generating jets of plausible shape. We now wish to consider more complicated cases with wave superpositions or even the continuous spectra associated with a localised coherent source. For these experiments we need a more sophisticated diagnostic tool.

The model geometry (a periodic rectangular channel) suggests that simple Fourier analysis would be suitable. We do neglect the curvature of the planet, but we feel that small errors so introduced are outweighed by the greater convenience of sine and cosine functions instead of spherical harmonics. For example, with sine series it is very easy to handle boundaries. Straightforward calculation of the power spectral density is sufficient to detect which signals are present in the model flow.

The frequency spectrum of the model forcing or response is found by extracting a two-dimensional dataset, such as temperature anomaly as a function of longitude and time, or of height and longitude. The fixed parameters in the first case are latitude (typically $\phi = 0$) and height (either around the water clouds to show the forcing, or else a level in the stratosphere to show the response). In the second case we fix latitude and time to

establish the vertical wavelength.

We merely want to estimate the power spectrum of a small dataset, for diagnostic purposes. We use a simple method based on one described by Press et al. (1986). We use a library routine to find the complex transform of a two-dimensional dataset. For real data, the transform has the symmetry property $F_{N-n} = F_n^*$, where any row or column has elements $F_N, n = 0, \ldots N-1$. The "one-sided power spectral density" at wavenumber h is $P_h = 2 \mid F \mid^2$, the factor 2 arising from the symmetry between positive and negative frequencies.

Before giving examples of model spectra we will make a few comments about potential drawbacks to this simple method. None of these have yet arisen but we should be aware of the possible problems and their remedies.

A potential problem in the time dimension is aliasing, i.e. spuriously moving high-frequency signals into low-freqency bins. This cannot arise in the spatial dimension, because the sampling interval Δ is equal to the model grid spacing. Frequencies above the Nyquist frequency $f_c \equiv 1/(2\Delta)$ are absent, because they cannot be supported by the model. However, aliasing may occur in the time dimension if the data are sampled at intervals longer than one timestep, as is almost inevitable. We can detect aliasing by checking that the power spectrum tends to zero as the frequency approaches f_c. If this is not the case, power is being aliased into the sampling interval and we need to increase the sampling rate. In practice this problem does not occur if the dataset is sampled at a frequency at least twice that of the highest component of the forcing.

Another potentially serious problem is that of convergence of the periodogram estimate. Increasing the length of the data sample does not decrease the variance: in fact the standard deviation is always 100%. This is because we would be sampling the data at higher frequency resolution, and hence not increasing the number of samples in the bin of any particular frequency of interest. This could be alleviated by summing (rebinning) the high-resolution estimate to obtain a smoother low-resolution one.

We cannot be sure that the estimate converges onto the true power spectrum, because of possible problems with leakage between adjacent bins. This is usually addressed by the technique of data windowing. The naive method we use here is the special case of a square or top-hat window function. We shall continue to use it so long as we are sure that the model variability modes of interest (i.e. equatorial Kelvin and MRG waves) are restricted to a known frequency band, with periods of a few days. In other words we have a good idea of what to expect and will use the simplest technique until this is shown to be inadequate.

To examine the PSD we first remove half the transformed dataset, which by symmetry contains no extra information. The spectrum is mapped onto positive and negative

wavenumbers following the convention that eastward is positive. Finally we plot the spectrum (or its natural logarithm) over a suitable range of wavenumbers, normally much smaller than the Nyquist interval. This is valid because almost all the power is usually at relatively low frequencies.

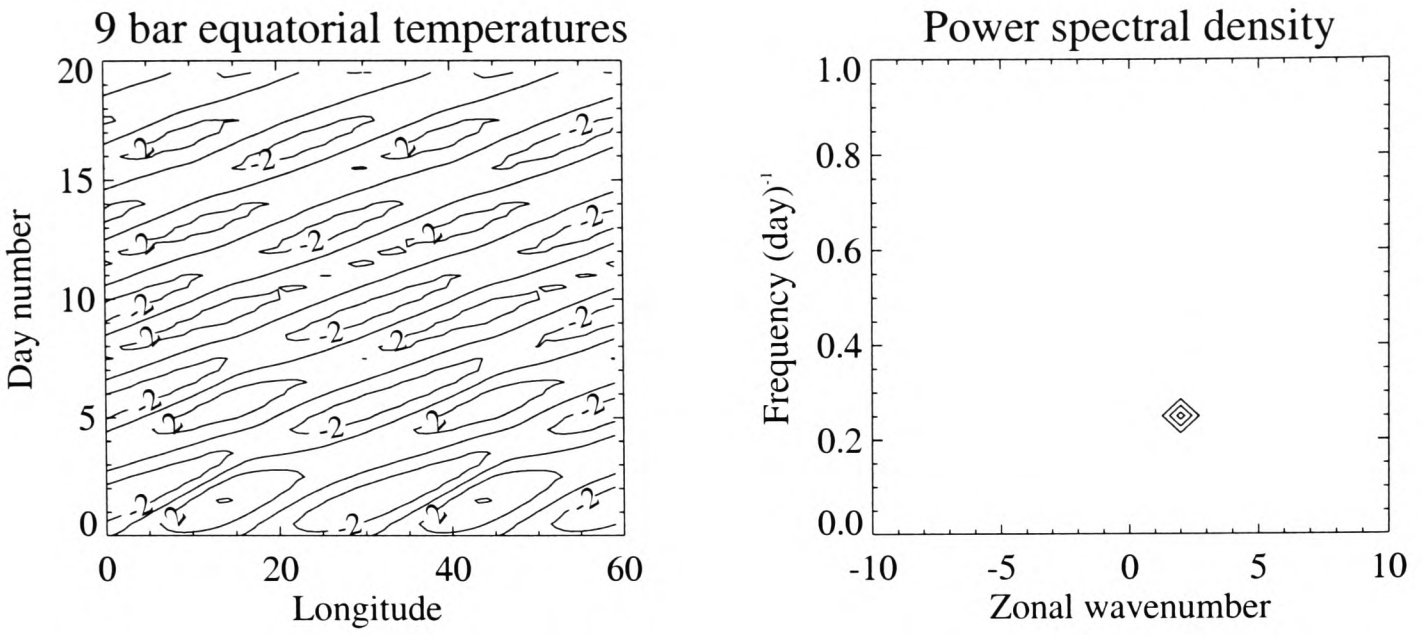


Figure 6.10: Temperature anomaly at 8 bars on the equator, and the power spectrum of this field. The model is the same as shown in Fig. 6.7. These plots show the thermal forcing at the base of the model.

Another example (Fig. 6.10 shows the tropospheric forcing used to generate the Kelvin waves of Fig. 6.7. The wave propagating eastward along the equator is obvious in the temperature plots. The power spectrum has a correspondingly sharp peak at zonal wavenumber +2 (eastward) and frequency 0.25 (day)$^{-1}$— the same as the forcing frequency. This spectrum was obtained by rerunning the model for the first 20 days, sampling the data every 12 hours.

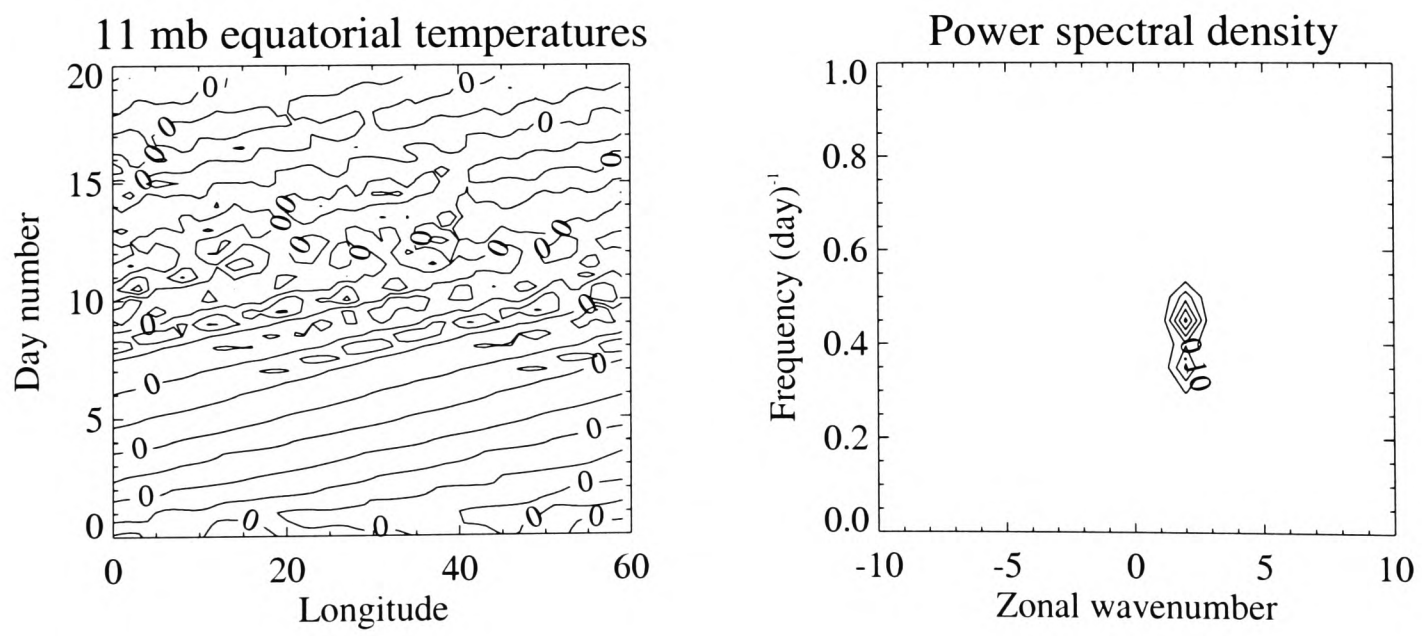


Figure 6.11: Temperature anomaly at 11 mb on the equator, and its power spectrum. This is the response to the forcing of Fig. 6.10.

Figure 6.11 shows the response in stratospheric temperatures. Initially the response is dominated by the zonal wavenumber 2 signal, although after about 6 days, smaller features begin to appear. The spectrum shows that the power is mostly in wavenumber +2

but the response frequency is spread between 0.35 and 0.5 $(\text{days})^{-1}$. The flow is developing from rest, and we are seeing a response at a frequency higher than that of the forcing. Later experiments show that the high-frequency components disappear once the flow has stabilised. We have not analysed these transients, because at present we are more concerned with the equilibrium (or near-equilibrium) model state than the initial response. This could be an area of future development.

6.9 Experiments with other travelling equatorial waves

6.9.1 A mixed Rossby-gravity wave

Next we test how the model handles a monochromatic, travelling MRG wave. Figure 6.12 illustrates the horizontal shape of the forcing function in terms of temperature anomaly. This is directly proportional to the geopotential anomaly given by Eq. 6.15. The forcing signal has wavenumber 2 and period 4 days, travelling westward in a 120° channel. Its vertical form is cosinusoidal in height, centred at level 4 ($p \simeq 4.5$ bars) and vanishing at levels 1 (9 bars) and 7 (2.25 bars). This can be seen in the right-hand plot of Fig. 6.12, which is a latitude-height cross-section of the wave and has a node at the lower boundary.

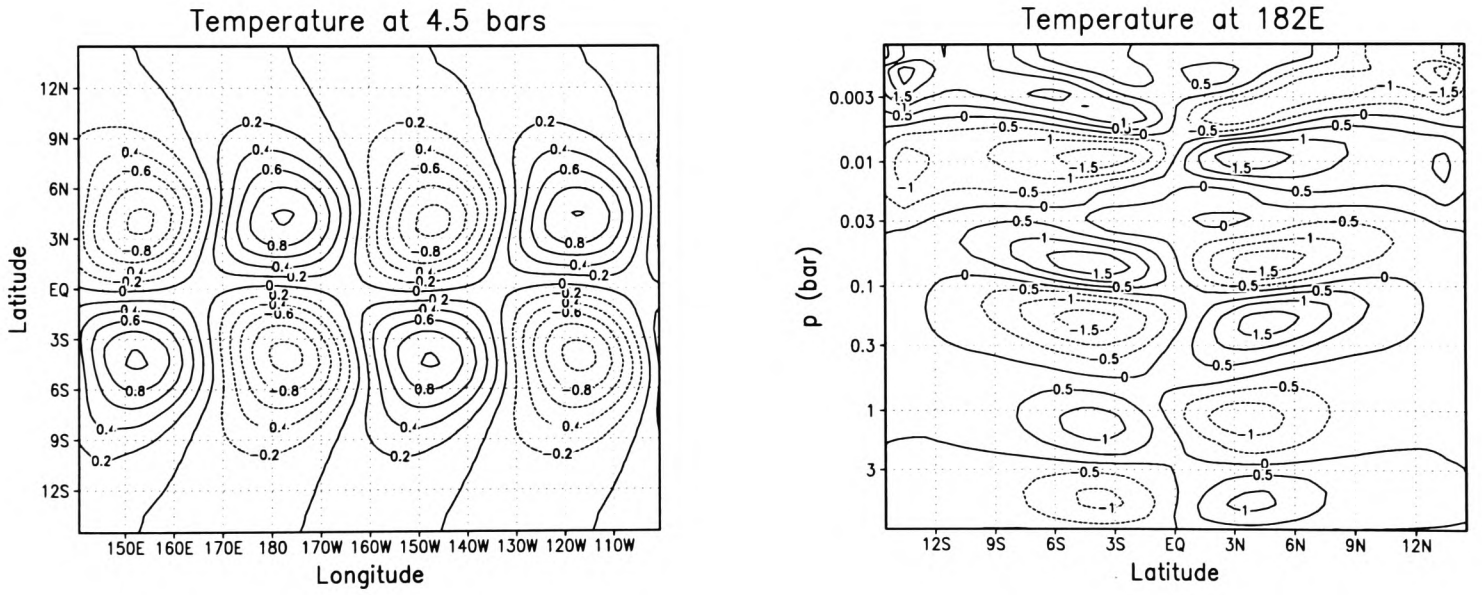


Figure 6.12: Temperatures (zonal-mean subtracted) after 10 days, in an experiment to force a mixed Rossby-gravity wave. Left: a horizontal section through the level of peak heating. Right: a meridional section showing the forcing and the wave propagating upward.

The result (Fig. 6.13) is a westward-propagating wave, seen in both the stratospheric temperature anomaly and (shown here) the meridional wind fields. Meridional wind is a good indicator of an MRG wave, because its value peaks on the equator. An easterly zonal jet appears near the tropopause.

Figures 6.14 and 6.15 show spectral analyses of the v-wind across the equator, at 4.5 bar

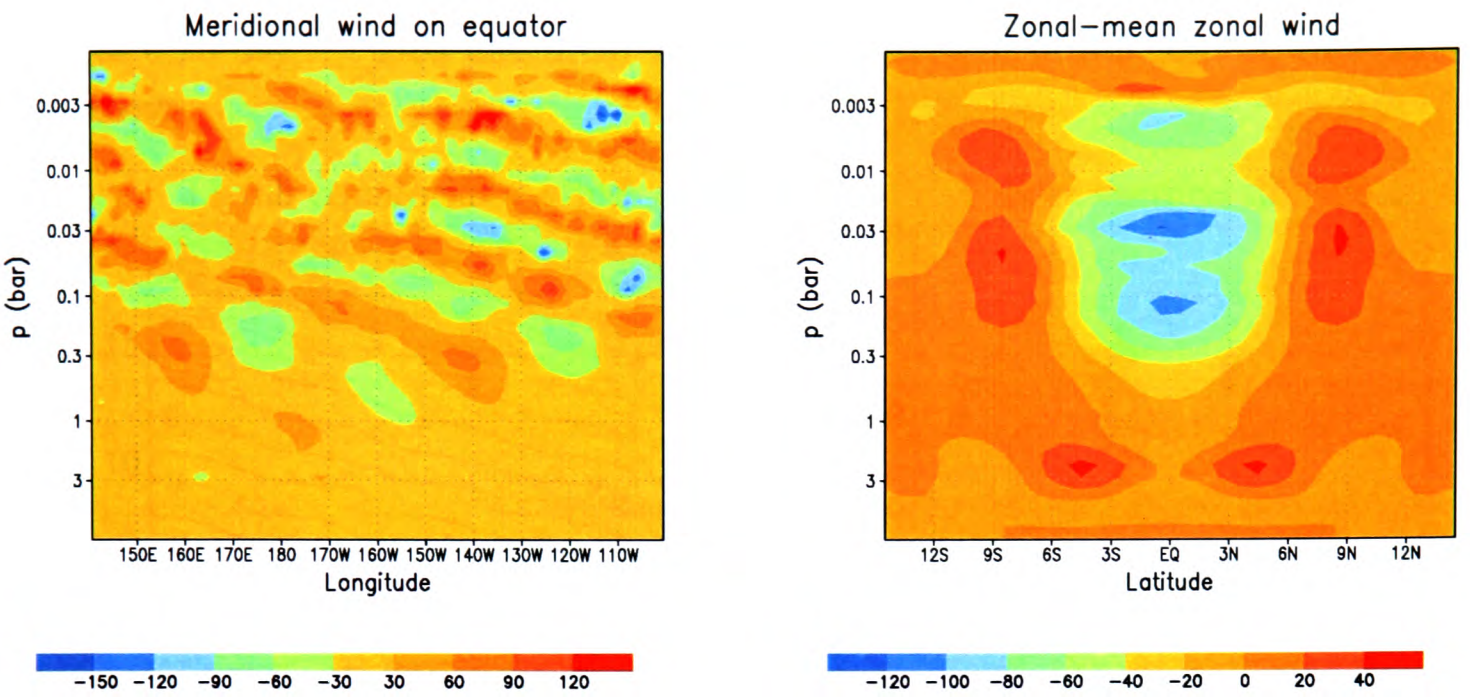


Figure 6.13: Results of an MRG wave model, after 80 days: equatorial section of the meridional wind (note the westward propagation) and the easterly zonal jet forced by the wave.

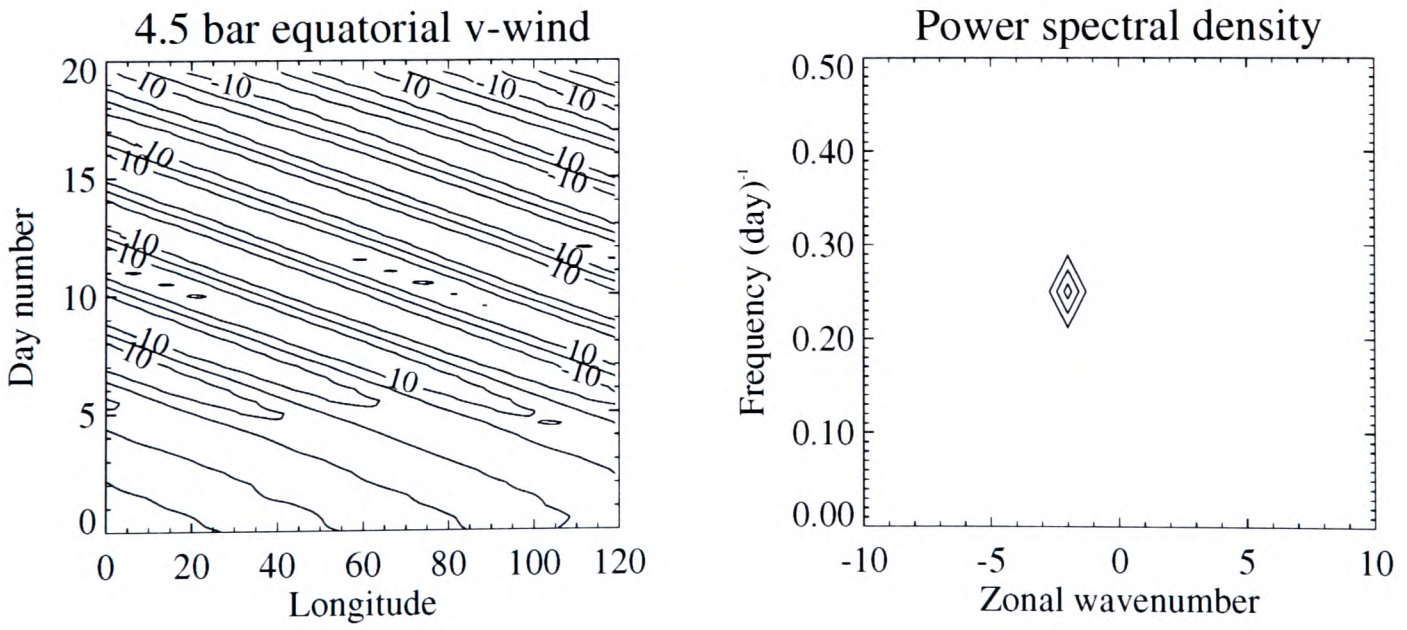


Figure 6.14: Time-longitude plot and power spectrum of meridional wind at 4.5 bar on the equator, forced by a travelling mixed Rossby-gravity mode (Fig. 6.12).

and 11 mb respectively. These spectra are taken from the first 20 days of the run. The tropospheric forcing has wavenumber -2 and a monochromatic frequency of 0.25 day^{-1}, consistent with the heating function. The stratospheric response is also almost purely wavenumber -2. However, the response frequency is spread out above the forcing frequency of 0.25 day^{-1}. This nonlinear effect was also seen in the Kelvin wave example shown above.

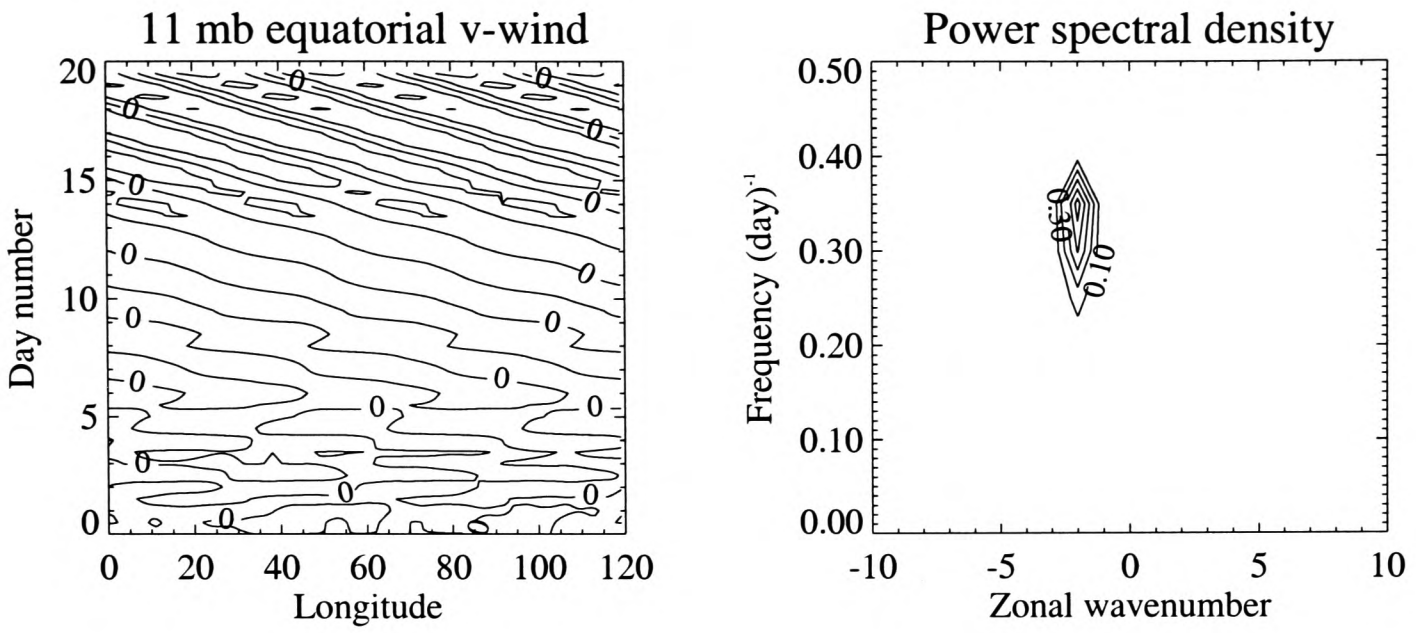


Figure 6.15: Time-longitude plot and power spectrum of meridional wind at 11 mb on the equator, forced by a travelling mixed Rossby-gravity mode (Fig. 6.12). Days 1–20 are shown.

To compare the propagation of the wave once the jet is established, another spectrum was taken over Days 80–100. This is shown in Fig. 6.16. In comparison with Fig. 6.15 it demonstrates a general broadening of the zonal wavenumber spectrum, together with a shift towards lower temporal frequencies. The peak is centred on wavenumber -2, frequency 0.25 day^{-1}. These are precisely the values of the forcing (Fig. 6.14). The higher frequency components have therefore died out and the response is now dominated by the forcing frequency, broadened almost isotropically.

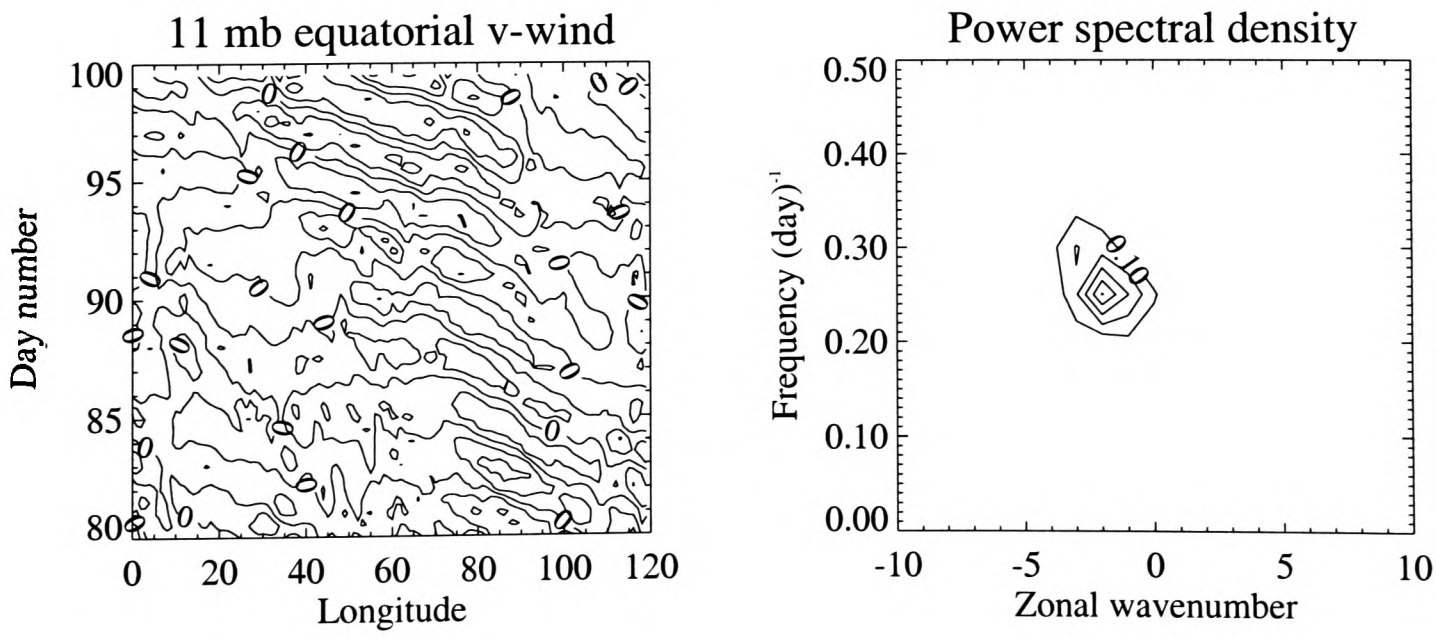


Figure 6.16: As Fig. 6.15, but for days 80–100.

6.9.2 A superposition of Kelvin and MRG waves

The next step towards a more realistic model of Jupiter's equatorial zone is to study the interaction of two modes. Later we shall consider models forced by extended, coherent sources analogous to the Earth's ITCZ. Such sources excite a red-noise spectrum of waves on all scales larger than the source iself. Before we can understand such sources, we need to look at the simpler case of two superimposed pure travelling modes. Another reason is that we are laying the foundations for a possible future study of stratospheric oscillations using this model. We shall not treat this subject in any detail in this thesis. A serious attempt to model the QQO would require very long model integrations, which is time-consuming. Our longest experiment (Fig. 6.21) lasts about 8 years, and shows behaviour which may or may not be the beginnings of an oscillation.

The experiment of Fig. 6.17 was forced with an eastward Kelvin mode of wavenumber 4 (in a 120° channel) and period 4 days, and a superposed westward MRG mode of the same period but wavenumber -4. The heating distributions are the same as in the experiments with individual waves.

We enseure that approximately equal thermal power is injected into each mode, by adjusting the amplitudes to normalize each forcing function. This is a simplistic assumption based on the linearity of the heating and cooling for small temperature differences, but we expect it to be broadly accurate. (To check this, the experiment was repeated with a different amplitude ratio—see below).

The results (Fig. 6.17) are interpreted in the same way as before, i.e. the signature of the Kelvin mode is evident in the equatorial temperature anomaly, and that of the MRG mode in the cross-equatorial meridional wind. The Kelvin wave is clearly evident above 10 mb, and the MRG wave between 100 mb and 10 mb. The transition between the two regions is marked by a strong easterly zonal jet, just above 10 mb. This suggests that the MRG mode is being absorbed by this easterly jet while the Kelvin mode propagates through. This is consistent with the theoretical prediction that westward-propagating modes are absorbed in westward (i.e. easterly) jets.

The timeseries in the lower-right panel of Fig. 6.17 shows that the main tropospheric westerly jet is rapidly excited by the Kelvin mode, and remains present throughout this (admittedly brief) integration. However, there is a noticeable downward trend in the jet height with time, with a suspicion of this trend in the stratosphere too. Around the 10 mb level a westerly is excited first (as with all other levels), but this subsequently decays and reverses, becoming easterly after Day 50. One explanation for this could be that the Kelvin mode responsible for exciting it initially can no longer reach this height because it is being absorbed lower down, in the westerly jets. The mixed Rossby-gravity mode, unaffected by the westerlies, is then able to deposit westward momentum and reverse

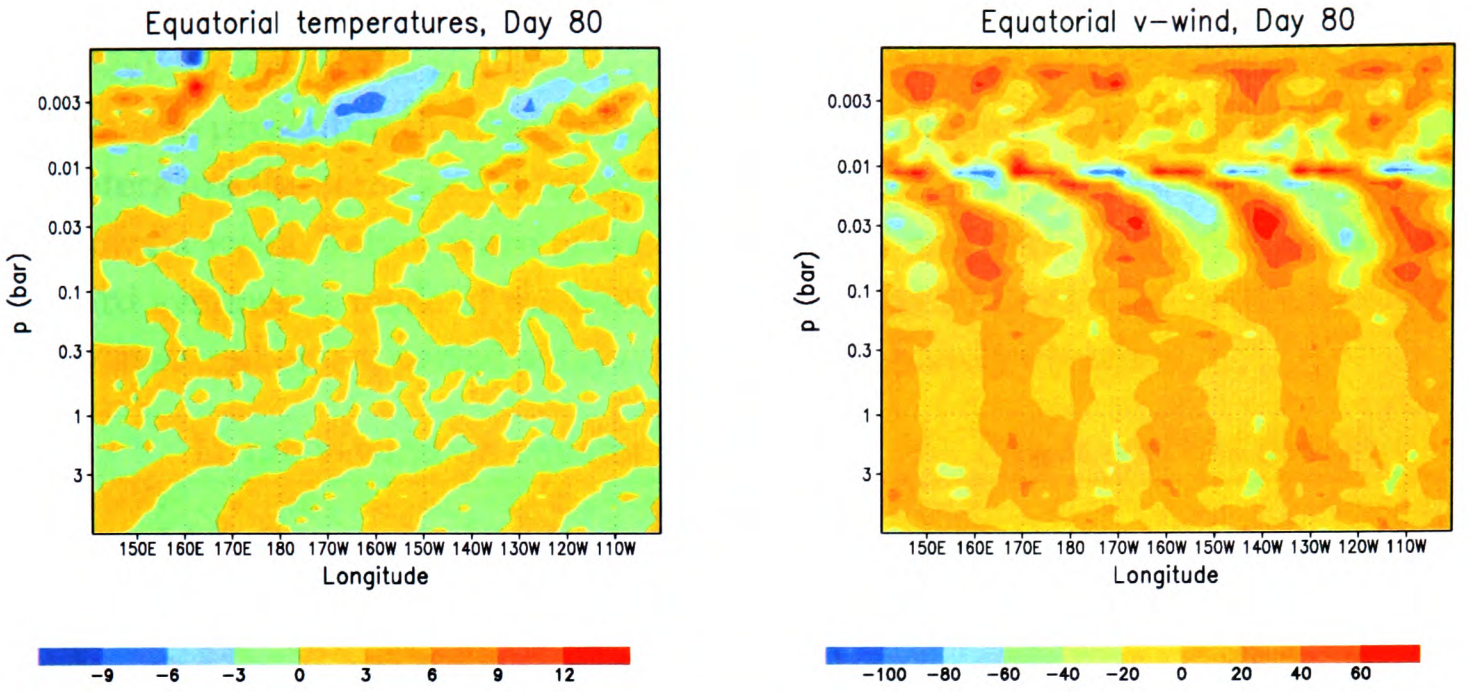


Figure 6.17: Results of a model forced by Kelvin and MRG waves of equal amplitude and wavenumber, and opposite phase speed. The top two panels show the stratospheric equatorial waves present after 80 days: on the left a Kelvin wave seen in the temperatures at the top of the model domain, and on the right an MRG seen just below this. The lower panels show the zonal-mean zonal wind, on the left as a cross-section in latitude, and on the right as a timeseries.

the jet. Clearly a much longer run is required to assess the long-term effects of this slow evolution.

The interpretation that the Kelvin wave is absorbed while the MRG is not, is supported by spectral analysis. We cannot directly compare the power present in each wave, for their signatures are seen in two different fields (temperature and meridional wind respectively). However, the timeseries at 11 mb (Figs.6.18 and 6.19) demonstrate that the MRG wave is present in the stratosphere, with wavenumber -4 and period 4 days. These parameters are the same as the forcing, so it is undistorted. By contrast, the temperature plot shows a weak Kelvin wave of wavenumber $+1$ and period 10 days, and a hint of a westward symmetric mode of wavenumber -1, period 20 days. Clearly the forced Kelvin wave (wavenumber $+4$, period 4 days) has been completely absorbed.

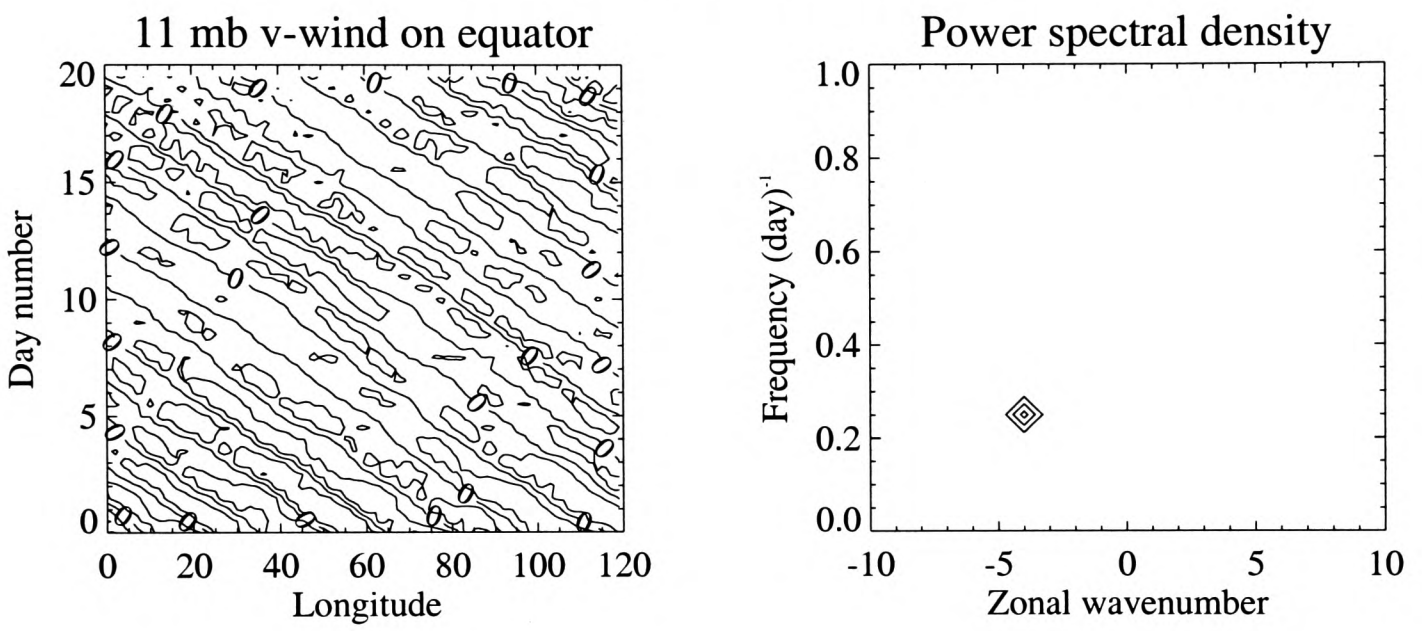


Figure 6.18: Time-longitude plot of stratospheric meridional wind, showing that the mixed Rossby-gravity wave reaches the 11 mb level undistorted.

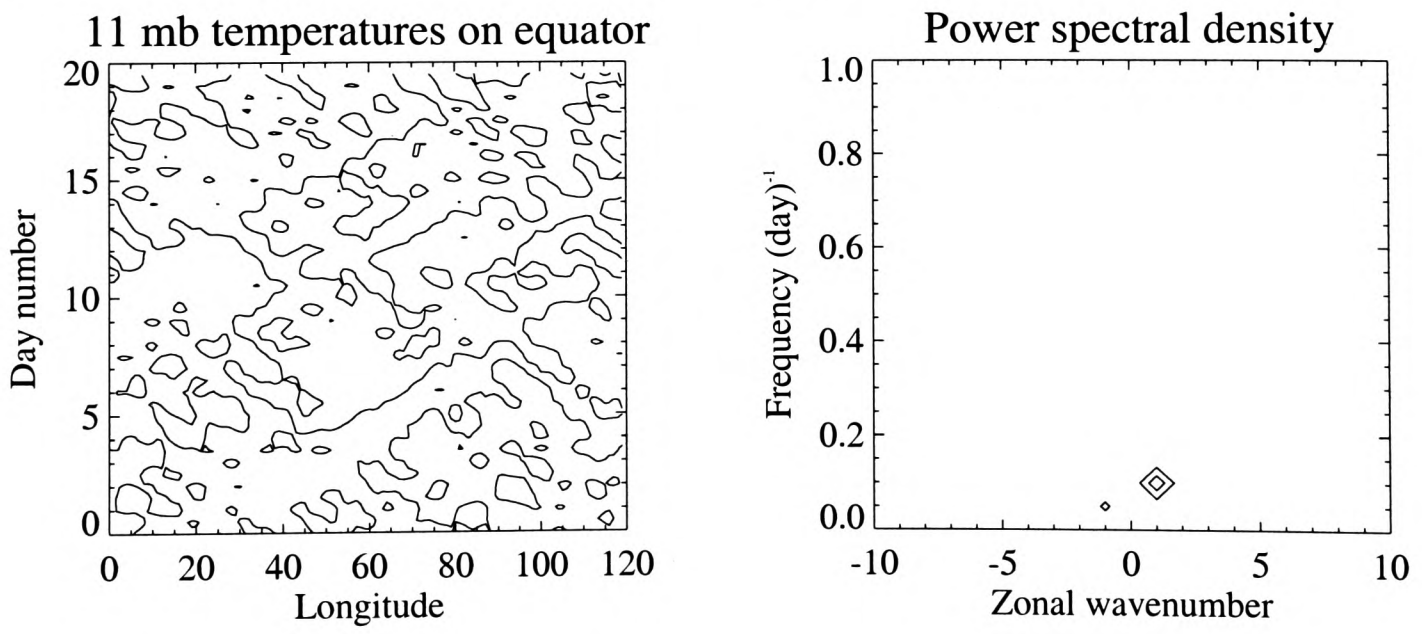


Figure 6.19: Time-longitude plot of stratospheric temperatures. The structure is complicated and there is no trace of the forced Kelvin wave. Instead, most of the power is in a wavenumber 1 wave with period 10 days. The Kelvin wave has been absorbed by the tropospheric jet.

Our final conclusion from this experiment is that, despite the Kelvin and MRG modes

being forced with equal amplitude, the Kelvin wave is dominant because the main jet is westerly. The presence of the MRG mode has hardly affected the growth of the westerly jet. It is important to establish the robustness of this result. If it turns out that Kelvin modes are preferentially selected within the troposphere, this would help to explain the observed westerly equatorial jet. Conversely, if Kelvin and MRG modes are selected equally, we would have no reason to prefer either. Depending on the filtering effect of the troposphere, we might need to make arbitrary assumptions about the forcing sources in order to choose waves which force a westerly. We will now test this mode selection by increasing the relative amplitude of the MRG wave and comparing the results.

6.9.3 Kelvin and MRG waves with different amplitude ratios

In this section we present the results of much longer integrations of the model. The reason for this, as stated above, is to study the interaction between the different wave modes over long timescales. For comparison we first show (Fig. 6.20) a plot of the winds up to Day 80, in a model with MRG forcing double that of the Kelvin forcing. Comparison with the equal-amplitude results (Fig. 6.17) shows that the flow is similar. There is still a 90 m/s westerly near the cloud-tops, although the easterly above has gained speed. The increase in MRG amplitude has had little impact on the main westerly jet.

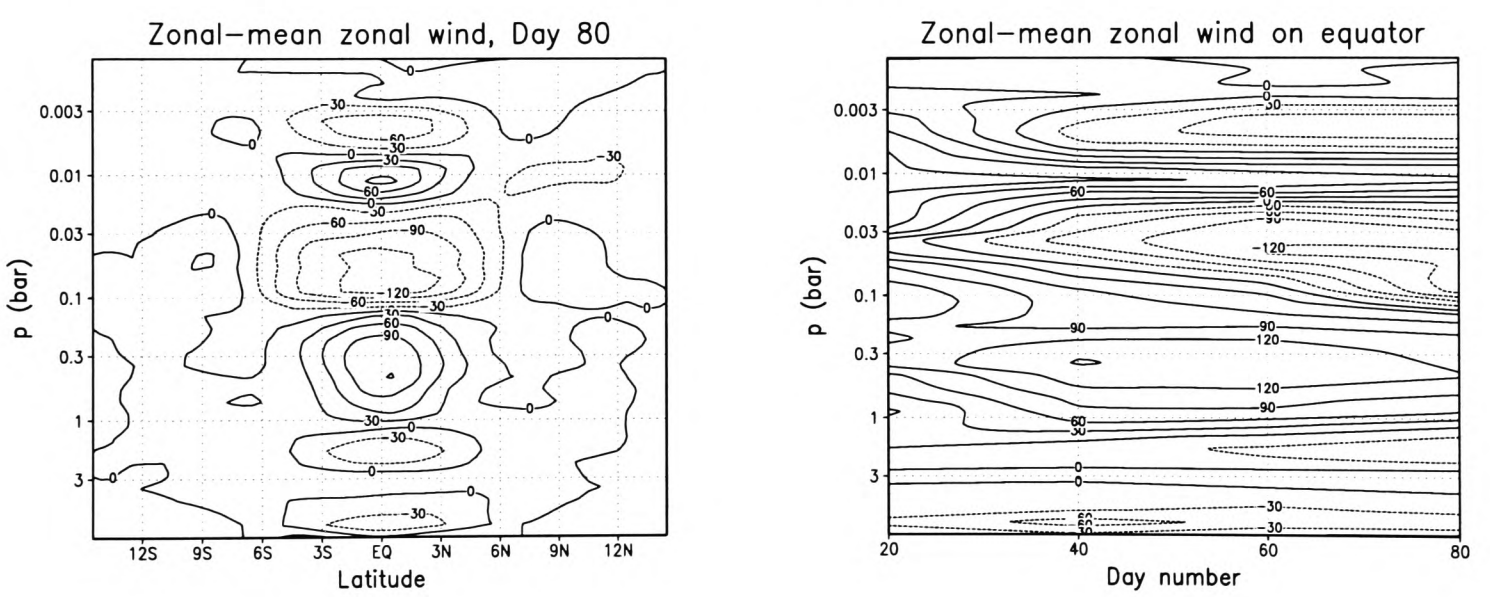


Figure 6.20: Results of a model where the MRG forcing is double the amplitude of the Kelvin forcing. Note that this model does not write output until Day 20, so the time axis of the timeseries is not the same as in Fig. 6.17.

Figure 6.21 illustrates the changes in speed and direction of the equatorial winds in this model, as a function of time over a period of around eight (Earth) years. There is a steep downward progression for the first few months, followed by a period of slower change in which the easterly jet dies out and the cloud-top westerly becomes deeper. During year 6 there is a downward phase shift in the stratosphere, and then in year 7 the main westerly jet shuts off abruptly and reverses to become easterly. This suggests that the MRG mode is becoming dominant.

This tropospheric behaviour is clearly not consistent with that observed on Jupiter, but in the stratosphere if we ignore the rapid downward shifts during year 6, we do see a period of westerly winds (red) for the first two years, followed by easterly winds (green) for around four years, and then westerly again for the final two or so years. Although this model run is too short to be certain, the flow in some ways resembles the observed quasi-quadrennial stratospheric oscillation (Leovy et al., 1991; Friedson, 1998; Li and Read, 1999) . This is worthy of further study, probably with a weaker MRG forcing which would avoid the destruction of the main westerly jet.

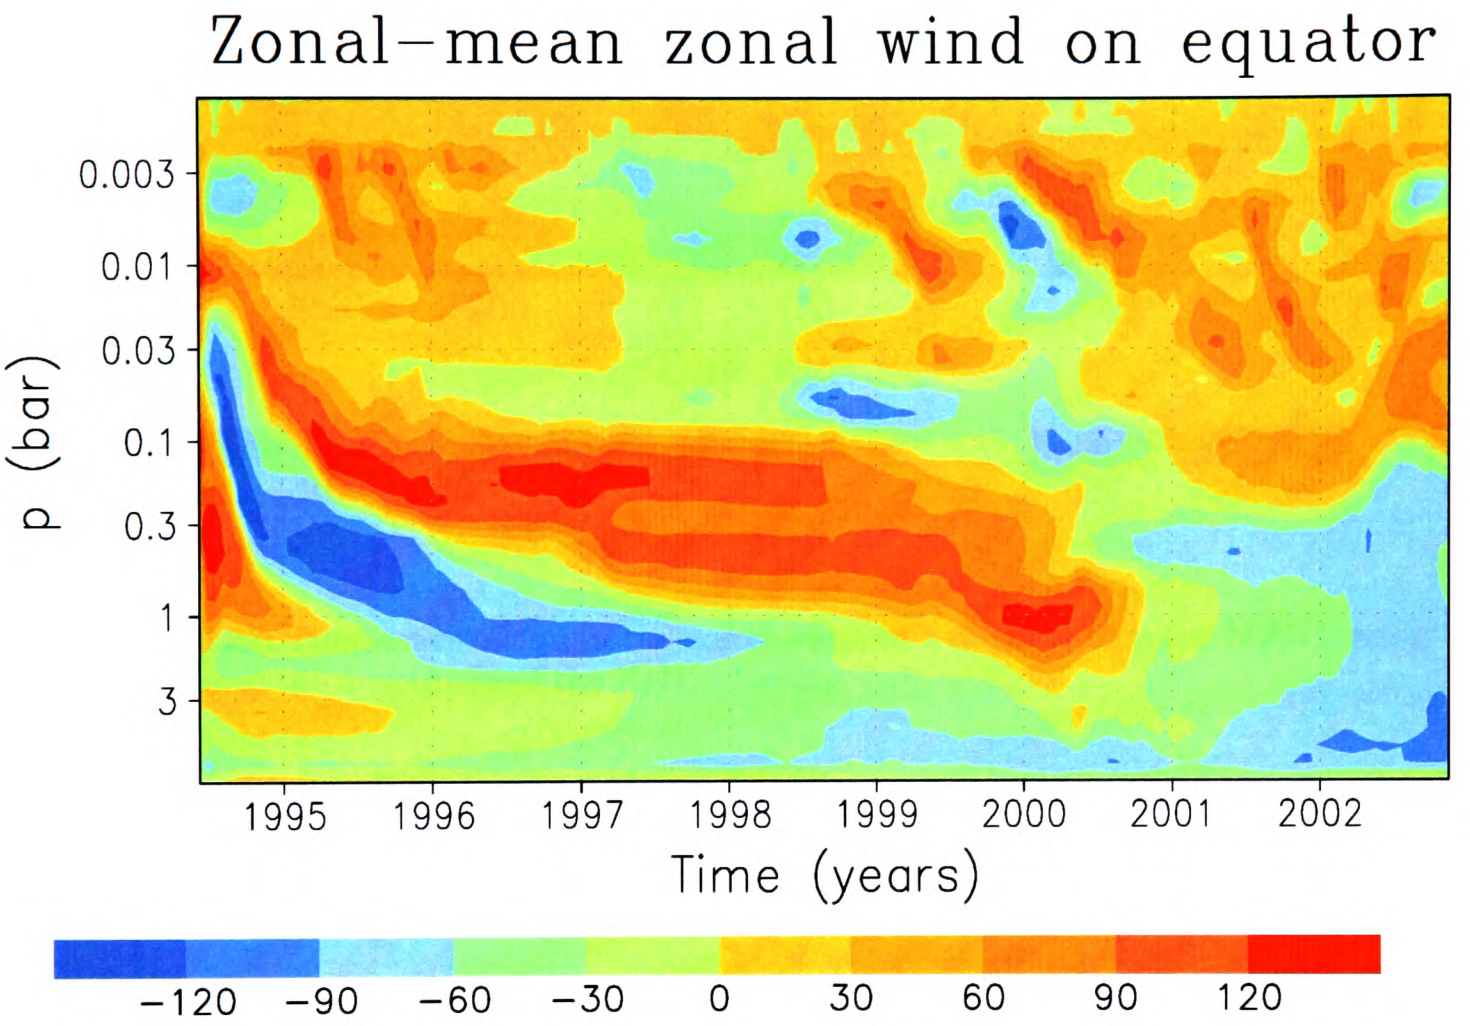


Figure 6.21: Timeseries of equatorial zonal wind, in a model where the MRG forcing is double the amplitude of the Kelvin forcing.

As the effect of doubling the MRG relative amplitude is not particularly marked, at least in the early stages of the integration, we increase the MRG amplitude to a value four times that of the Kelvin wave (Fig. 6.22). Some features remain the same—for example the downward-trending jets during the first 100 days—but this time the tropospheric jet is a strong easterly, so the MRG wave does now dominate the forcing.

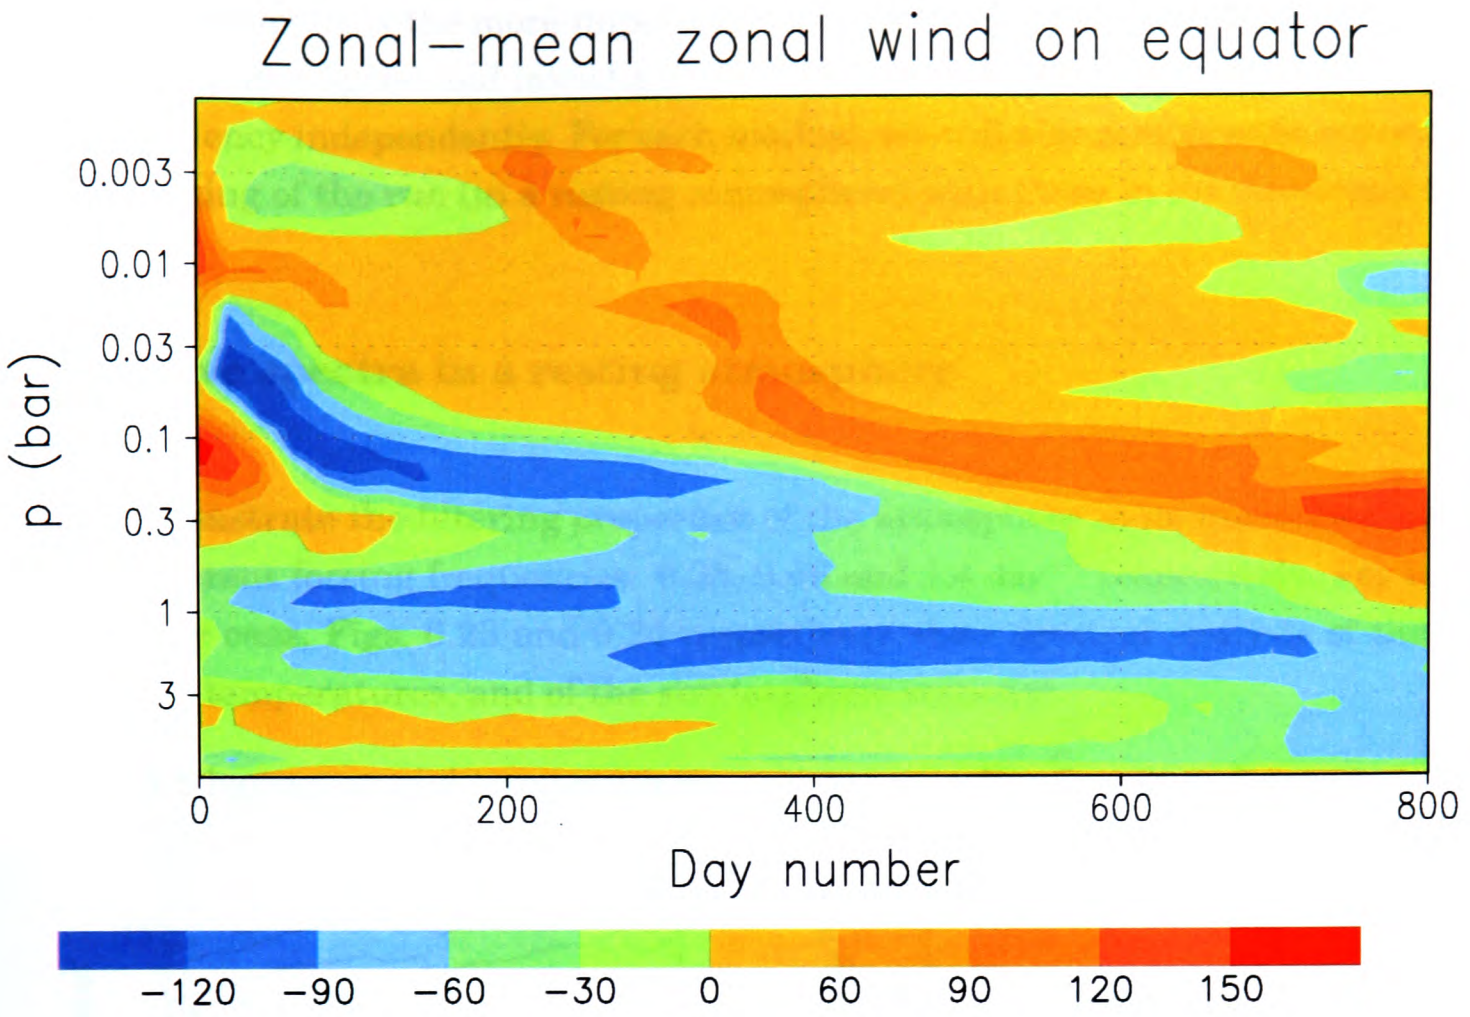


Figure 6.22: Timeseries of equatorial zonal wind, in a model where the MRG forcing is four times the amplitude of the Kelvin forcing.

6.10 Wavenumber selection for travelling waves: superpositions of equal frequency

Having studied the flows generated by monochromatic Kelvin and MRG waves, the next logical step is to try to infer the filtering properties of the atmosphere. Knowledge of which spectral regions are preferentially selected will allow us to understand the effects of more complicated sources. The filter properties can be explored as a function of wavenumber or of frequency, and will in general also vary between different modes. Here we will concentrate on wavenumber selection in Kelvin waves. The reason for this choice is that the Kelvin wave is a good candidate for the mechanism driving the superrotating jet. Therefore the propagation of Kelvin waves through the atmosphere, both with and without westerly jets, is of particular interest.

The technique is to force the model with a superposition of (typically eight but sometimes twelve) travelling Kelvin waves of different horizontal wavenumber k. Given that we are interested in propagation as a function of k, we have the choice of fixing either the frequency or phase speed in each experiment.

We show results of both methods, starting with waves of equal frequency and hence differing phase speeds. This is the approach taken by Chang (1976). However, our results sug-

gest that phase speed is the more important variable in Jupiter's atmosphere, so we then fix a suitable phase speed and investigate the secondary effects of selection by wavenumber and frequency independently. For each method, we will also compare the wave spectra at the beginning of the run (in a resting atmosphere) with those in the developed flow.

6.10.1 Wave spectra in a resting atmosphere

We will demonstrate the filtering properties of the atmosphere in its initial resting state, at three different forcing frequencies: 0.33, 0.66 and 1.4 day^{-1} respectively. For the lowest frequency case, Figs. 6.23 and 6.24 respectively show spectral analysis of the forced tropospheric temperatures, and of the stratospheric response.

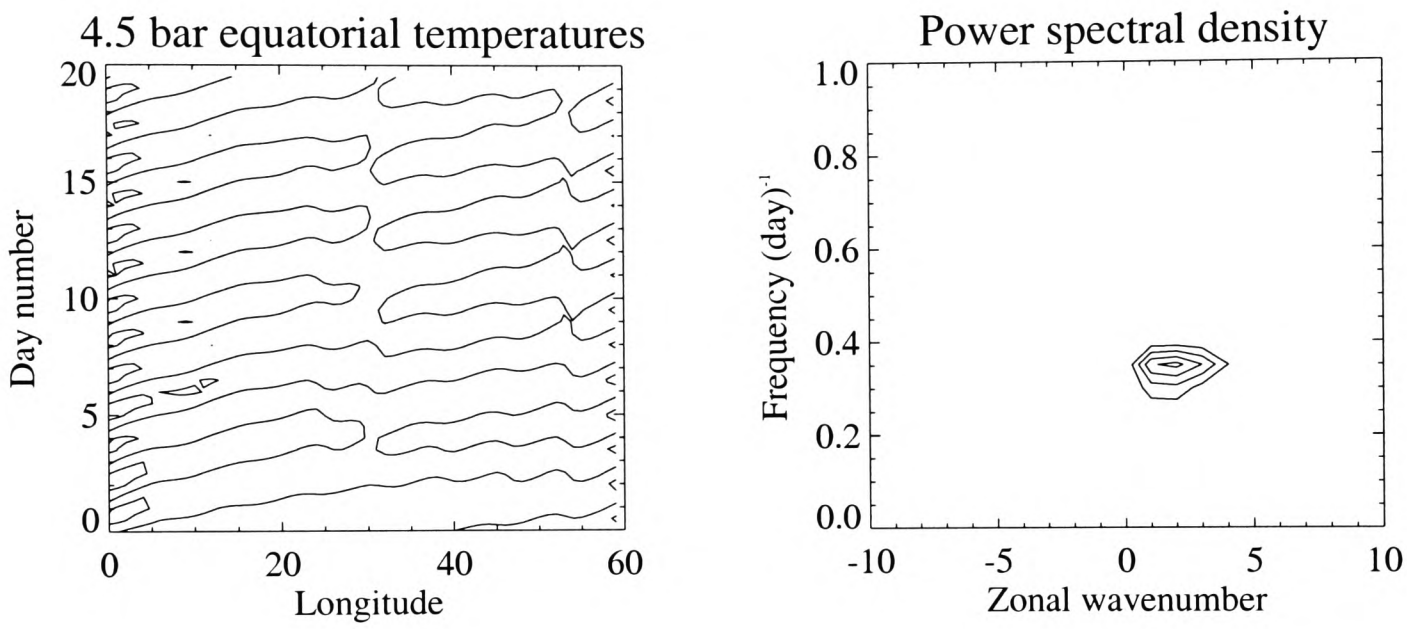


Figure 6.23: Showing the travelling-wave forcing function. Time-longitude plot and power spectrum of temperatures at 4.5 bars on the equator, for 20 days starting from rest. There is a broad spread of wavenumbers 1–4, with frequency 0.33 day^{-1}. However, wavenumbers greater than 4 are not seen.

The tropospheric signal has quite a broad peak, from wavenumber +1 to +4. This is not as broad as the heating spectrum, which includes waves up to +8. A plot of log(PSD) (not shown) reveals a gradual decline in power with increasing wavenumber. The reason is unclear but may be to do with horizontal diffusion in the model, which is much more effective on shorter length scales. Also, the technique of direct heating with a wavelike structure may be more effective on longer scales. This is not a major difficulty because throughout this chapter we shall only compare the waves in the troposphere with those in the stratosphere. We do not directly compare with the forcing signal. We shall see that the stratospheric signal is invariably spectrally narrower than the tropospheric one. We always ensure that this is so, by broadening the forcing band where necessary.

Although the tropospheric signal is broadly spread over wavenumbers 1–4, the stratospheric response has a sharper peak at wavenumbers +2 and +3 only. We need to discover why the wavenumber 1 has been removed as well as the shorter wavelengths. Other work

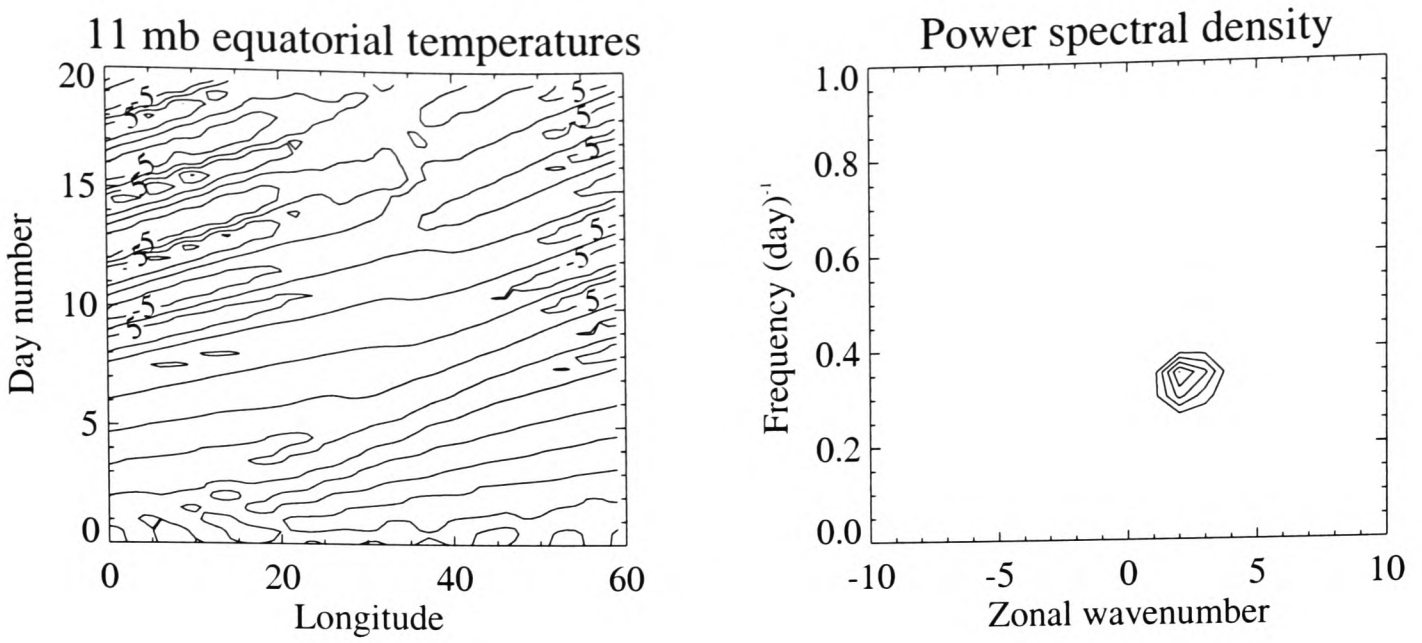


Figure 6.24: Time-longitude plot and power spectrum of equatorial temperatures at 11 mb. Wavenumber 2 dominates, and the temporal frequencies are spread around 0.35 day^{-1}

.

(e.g. Holton, 1972) has suggested that the Earth's atmosphere acts as a low-pass filter for equatorial waves, and it is possible that Jupiter acts similarly. This does not however explain the removal of wavenumber 1.

For Kelvin waves the vertical and horizontal wavelengths are inversely proportional, because the dispersion relation is $\omega = \pm Nk/m$. It is possible that wavenumber 1 has not been excited because the vertical scale of the forcing is too small. Like all waves, Kelvin modes are most efficiently excited when the forcing is of scale comparable to the wavelength (Holton, 1972). This was checked by repeating the experiment with a much thicker heating function, peaking at 2.25 bars and vanishing at 9 bars and 450 mb. The zonal phase speed was kept constant. The result was that no wavenumber 1 signal was seen, and no jet was formed except for a weak one immediately below the sponge layer.

Therefore we still need to find a reason for the removal of wavenumber 1. We do this by repeating the experiment with different forcing frequencies. When the forcing frequency is doubled (Fig. 6.25) the stratospheric response becomes mostly wavenumbers 4, 5 and 6 at the new frequency. Wavenumber 2 is absent, which suggests that the atmosphere does not filter by wavenumber but rather by phase speed. A wavenumber 2 at 0.35 day^{-1} has the same phase speed as wavenumber 4 at 0.66 day^{-1}. This result is supported by the third case, forced with frequency 1.4 day^{-1}, four times the initial frequency. Here (Fig. 6.26) we have wavenumbers 6 and 7 at the forcing frequency (i.e. approximately the same phase speed as the previous two cases). Interestingly there is also a weaker signal at wavenumber 4 and 0.66 day^{-1}. This component is not present in the forcing (it has only half the temporal frequency) but does have the same phase speed as the other wave.

These experiments have shown that the signal reaching the 10 mb level (and other stratospheric levels) in the initial, resting atmosphere is selected only by phase speed. The forc-

ing spectrum comprises eight discrete spatial wavenumbers, all at the same frequency. (The highest-frequency experiment was repeated with twelve components, to check that wavenumber 8 really was the highest which could be transmitted at 1.4 day^{-1}; the results were unchanged). In each case, only a few (up to four at most) of these wavenumbers are transmitted to the stratosphere. Varying the frequency allowed the selection of different wavenumbers. In frequency terms, the filter function has quite a wide band—for example, wavenumber 5 is transmitted at both 0.66 and 1.4 day^{-1}. However, it seems that the frequency and spatial wavenumber individually are each less important than the phase speed.

The phase speeds in every case were in a range between 100 and 150 m/s. It may be significant that this is close to the observed speed of Jupiter's equatorial jet. In the concluding section (6.15) of this chapter we shall discuss filtering by phase speed, and cite some evidence from other recent work (Li and Read, 1999) with similar results. We also note that the logical next step in this work would be to repeat experiments with different static-stability profiles. If our reasoning is correct, N^2 is the main variable determining the filtering and hence the equatorial jet speed. We shall return to this point in Section 7.2 (Future Work).

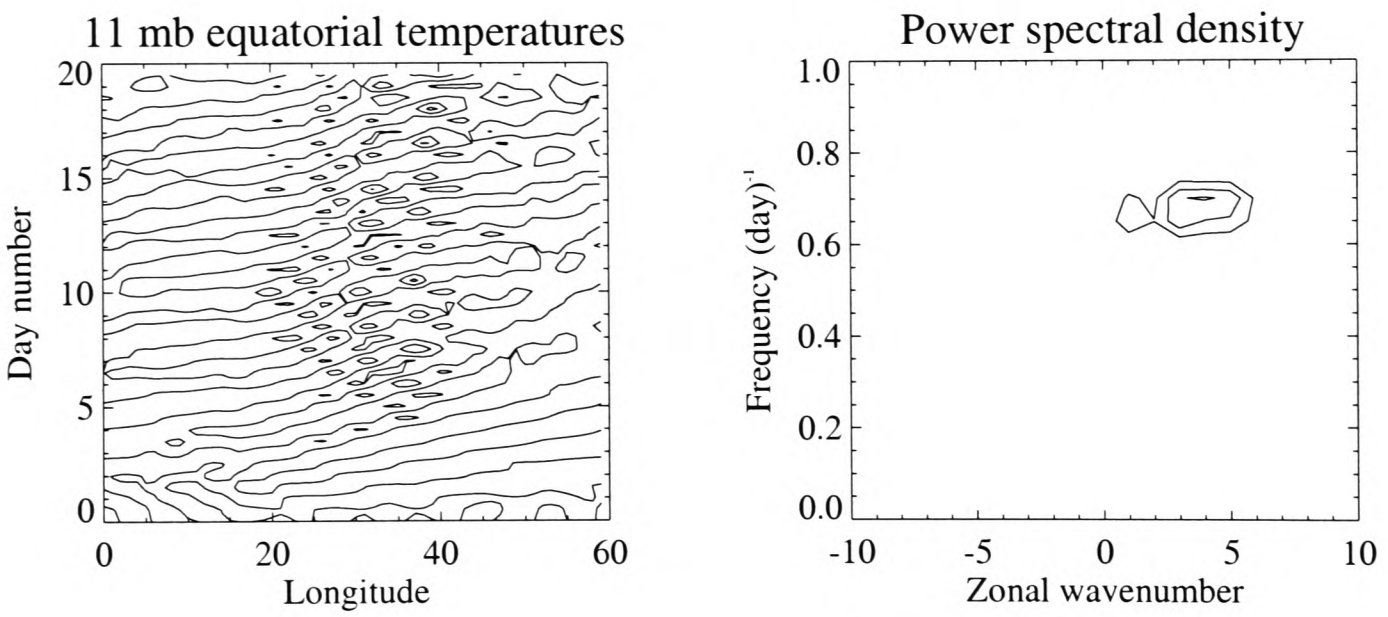


Figure 6.25: As Fig. 6.23 but with forcing of frequency 0.66 day^{-1}.

6.10.2 Wave-mean flow interaction

Fig. 6.27 shows the results of a 400-day integration of the model, forced by a superposition of Kelvin waves with frequency 0.33 day^{-1}. The jet speed reaches 140 m/s after 400 days. This velocity is close to the wavenumber 2 phase speed, which is consistent with the fact that Fig. 6.24 shows wavenumber 2 as the most powerful component reaching the stratosphere. Although the jet appears still to be developing slowly after 400 days, a continuation run for another 400 days shows that it does not grow further, but actually declines slightly.

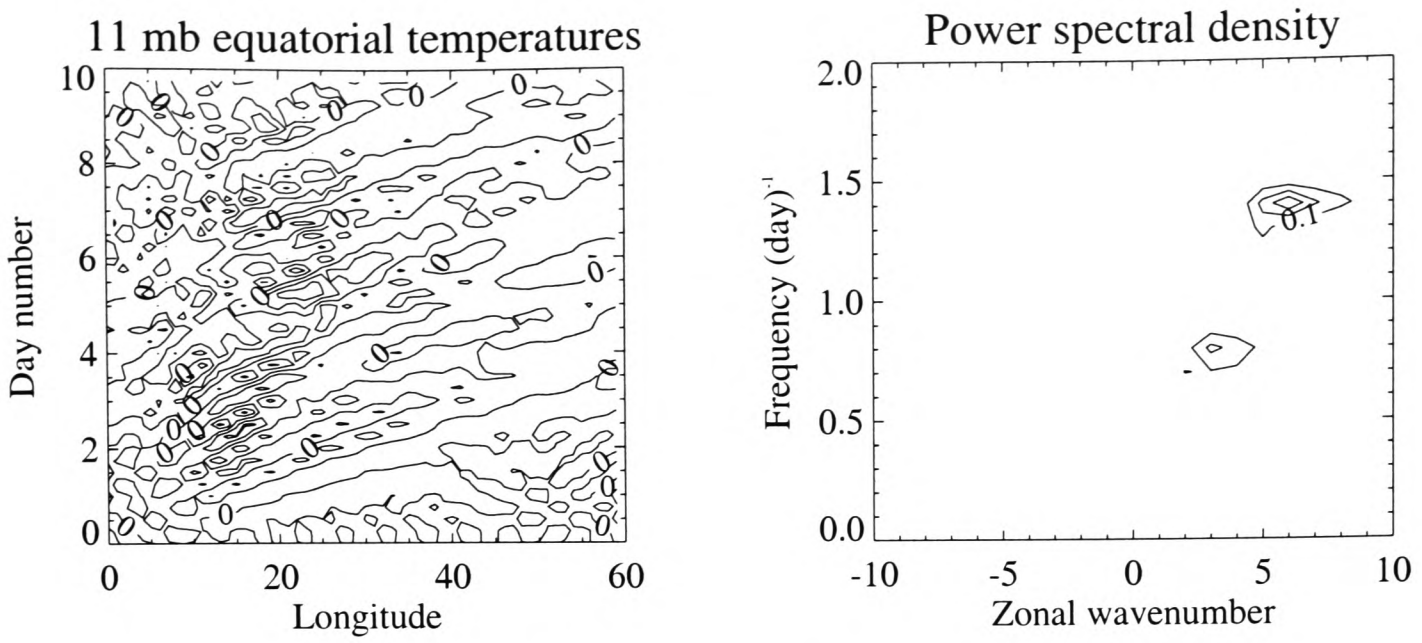


Figure 6.26: As Fig. 6.23 but with forcing of frequency 1.4 day^{-1}. Note the appearance of a low-frequency component not present in the forcing, but with equal phase speed to the higher-freqency waves present.

Note that although the jet in the upper troposphere is westerly, there is an easterly below approximately 3 bars. This is a common feature of our results. It is probably an unrealistic one, in view of the Galileo Probe Doppler Wind Experiment (Atkinson et al., 1996, 1997), which showed westerlies down to at least 22 bars. Ideally we would like to place the model sources deeper in the atmosphere, but cannot do this without extending the model down below 10 bars. This development has so far been beset by technical difficulties, but ought in principle to be possible in future (see Section 7.2).

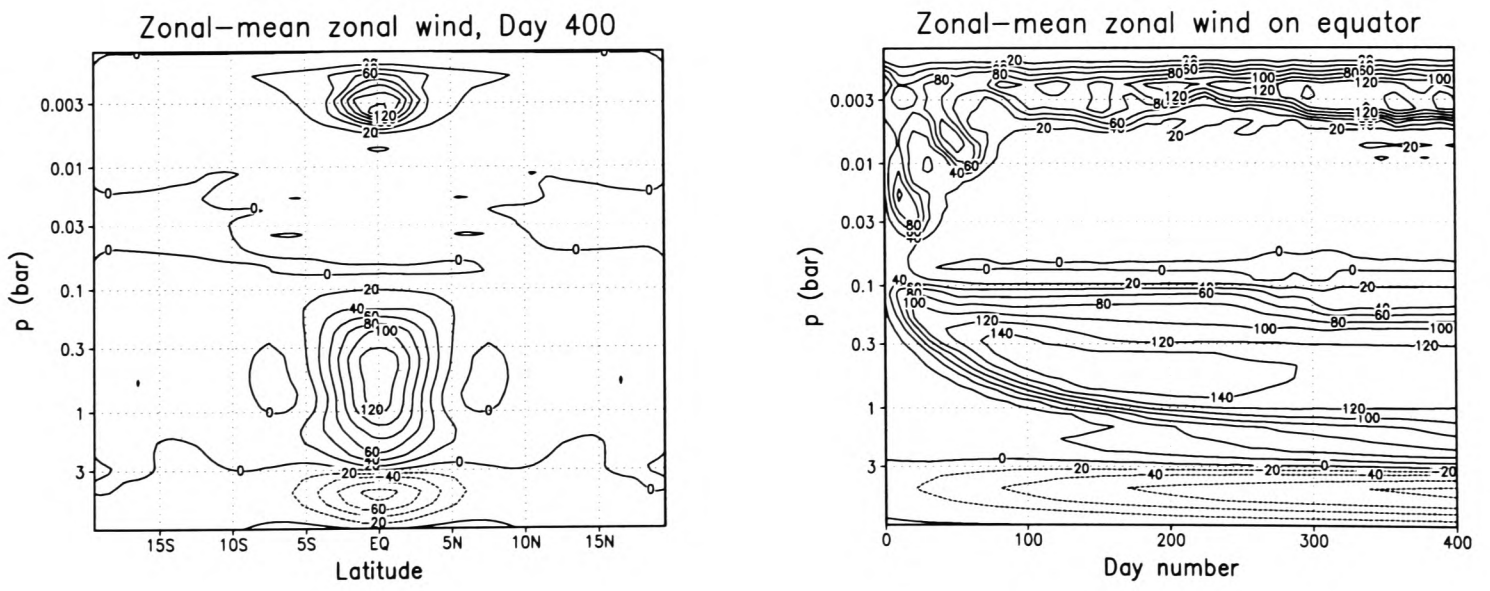


Figure 6.27: Results of a model forced with an equal-frequency superposition of Kelvin waves. Left: cross-section of zonal mean zonal wind, also at Day 400. Right: timeseries of zonal wind on the equator, showing growth and saturation of the tropospheric jet. The core jet speed is 120 m/s.

The spectrum of waves passing through the developed jet is rather different from those obtained at the start of the runs, in zero background flow. In all cases the tropospheric jet acts as a low-pass filter, since wavenumber 1 is always seen. The phase speeds are the same as in the resting atmosphere. The most interesting example is the highest frequency one, where initially we saw (Fig. 6.26) waves 3, 4, 6, and 7, in two groups with

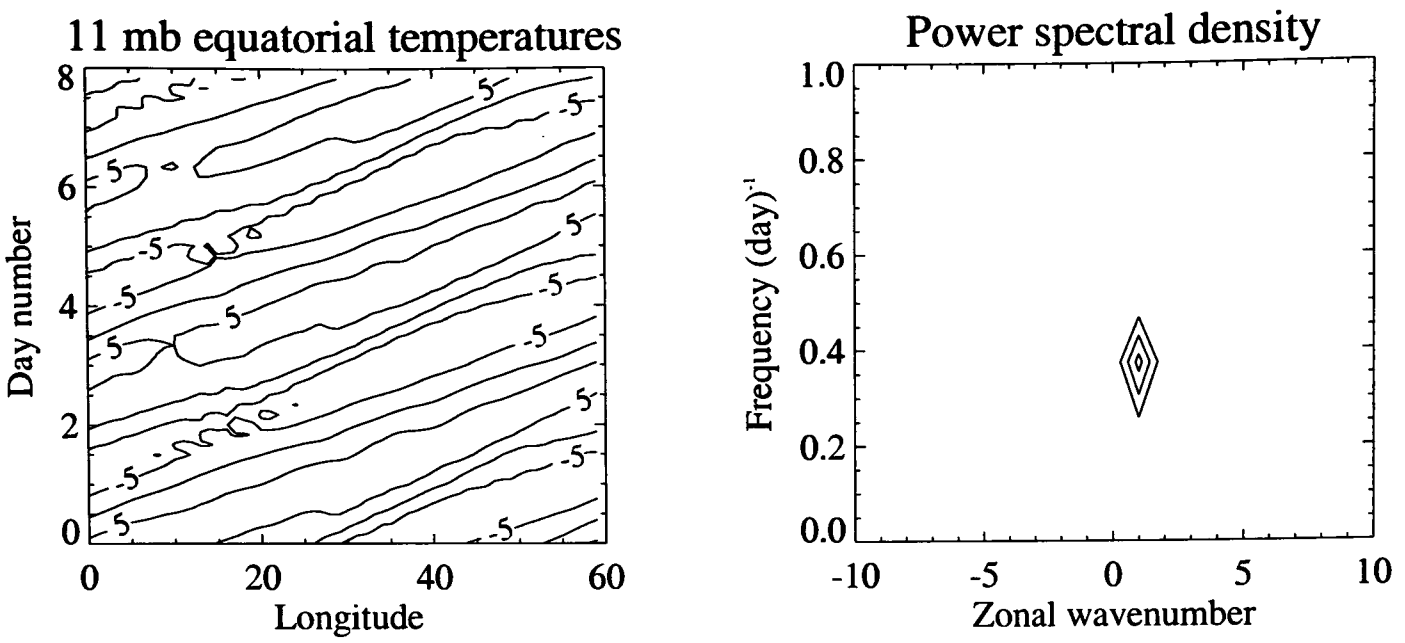


Figure 6.28: Time-longitude plot and power spectrum of temperatures at 11 mb on the equator, for the developed flow shown in Fig. 6.27 (i.e. after 400 days). In contrast with the initial flow (Fig. 6.23), the only signal is a pure wavenumber +1 with frequency 0.3 day^{-1}.

equal phase speeds. As the flow developed, both groups were filtered out, to be replaced with wavenumber 1 but at the same phase speed.

These results indicate firstly that phase speed is the most important factor determining the propagation of a Kelvin wave in the Jovian atmosphere, and secondly that the presence of a westerly jet acts as a low-pass filter in terms of wavenumber, so only the largest scales propagate through. In the next section we will try to confirm this conclusion by fixing the phase speed, and forcing the flow with a superposition of different frequencies and wavenumbers.

6.10.3 Superpositions with equal phase speed

In this section we examine the alternative to the previous experiment. This is a superposition of travelling waves all with the same phase speed, different wavenumbers 1–8, and frequencies determined by the dispersion relation. Again we keep the vertical shape of the forcing fixed, allowing the vertical wavenumber to be determined by this and the model stratification.

We will show examples with three different zonal phase speeds: 290, 145, and 72.5 m/s. The channel length is 60° and the period of the wavenumber 1 components are respectively 3, 6, and 12 days. Periods of the higher components are scaled appropriately to match the phase speeds.

For the highest phase speed, we can see from the equatorial flow after 400 days (Fig. 6.33) that this is an unsuitable choice, because the waves propagate easily through the whole

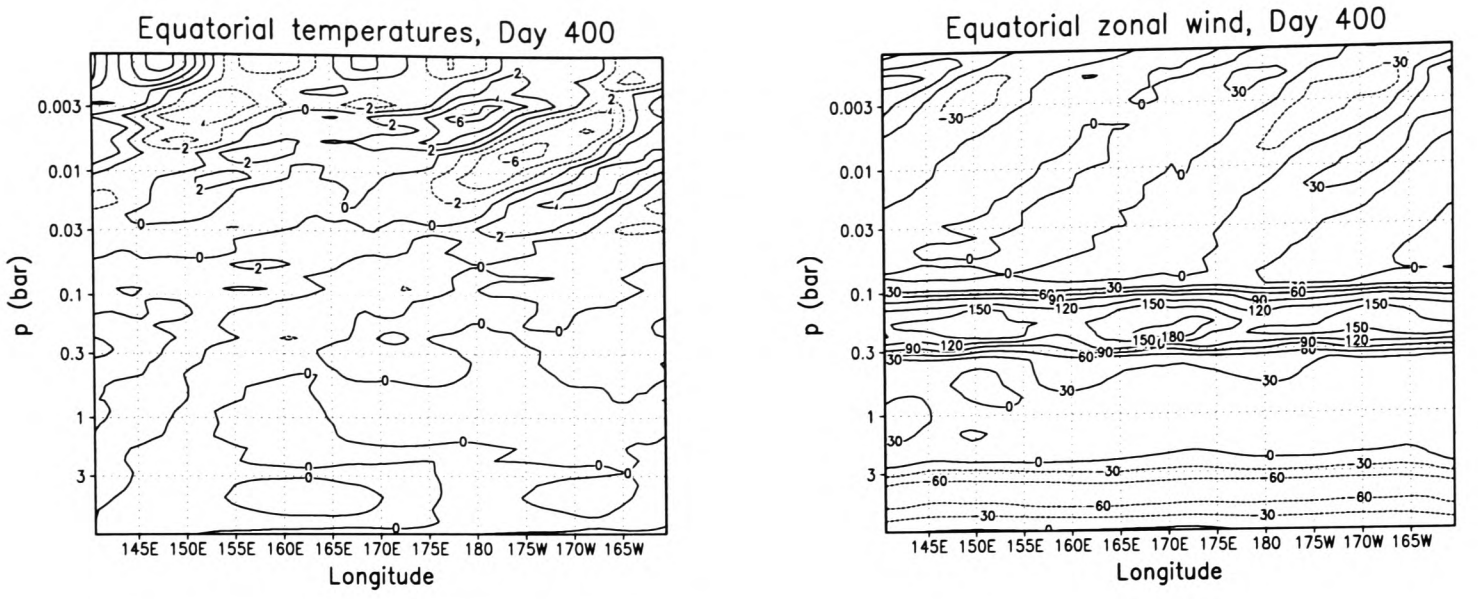


Figure 6.29: As Fig. 6.27, but forced at twice the frequency. Above: temperature anomaly and zonal wind at the equator, after 400 days. Lower left: cross-section of zonal mean zonal wind, also at Day 400. Lower right: timeseries of zonal wind on the equator. The jet reaches 140 m/s, i.e. 17% faster than the previous case.

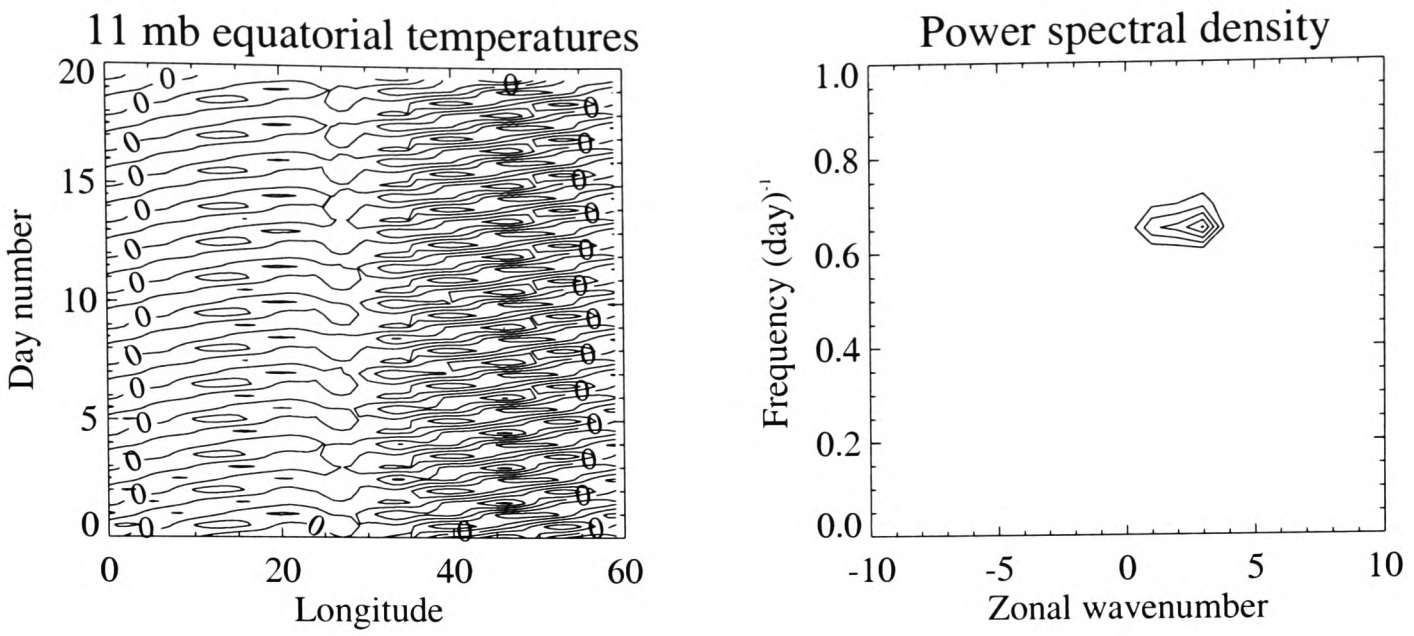


Figure 6.30: Time-longitude plot and power spectrum of temperatures at 11 mb on the equator, for the flow shown in Fig. 6.29. The dominant frequency is now 0.65 day^{-1}, equal to the forcing frequency. Wavenumbers 1, 2, and 3 are present.

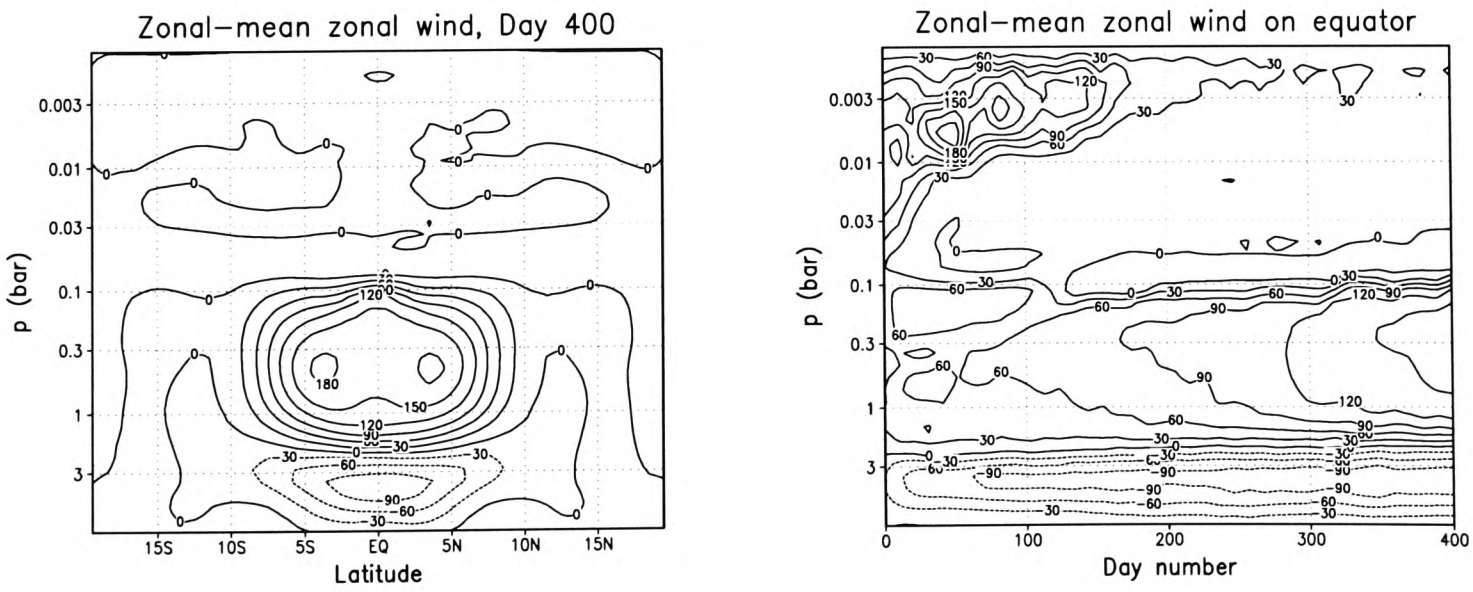


Figure 6.31: As Fig. 6.27, but forced at four times the frequency. The jet peaks at 180 m/s, again rather faster than the previous case. It also continues to increase in speed.

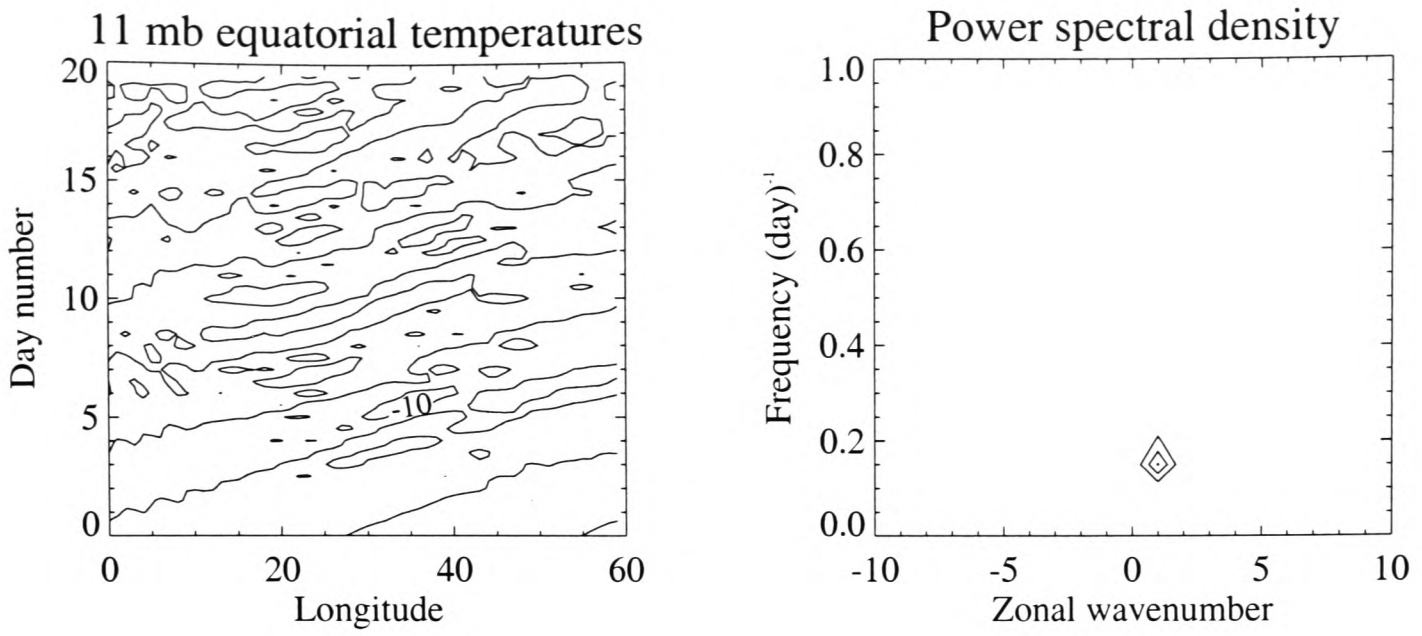


Figure 6.32: Time-longitude plot and power spectrum of temperatures at 11 mb on the equator, for the flow shown in Fig. 6.31. The only mode present is wavenumber 1 with frequency 0.33 day^{-1}. Comparison with Fig. 6.26 reveals that although this wave has much lower frequency than the forcing, it has the same phase speed as all the waves present in the initial spectrum.

atmosphere before being damped in the sponge layer. Even after 400 days, very little zonal momentum has been deposited anywhere except in the sponge layer. This is consistent with the results in the previous section, where the highest phase speed seen was around 150 m/s.

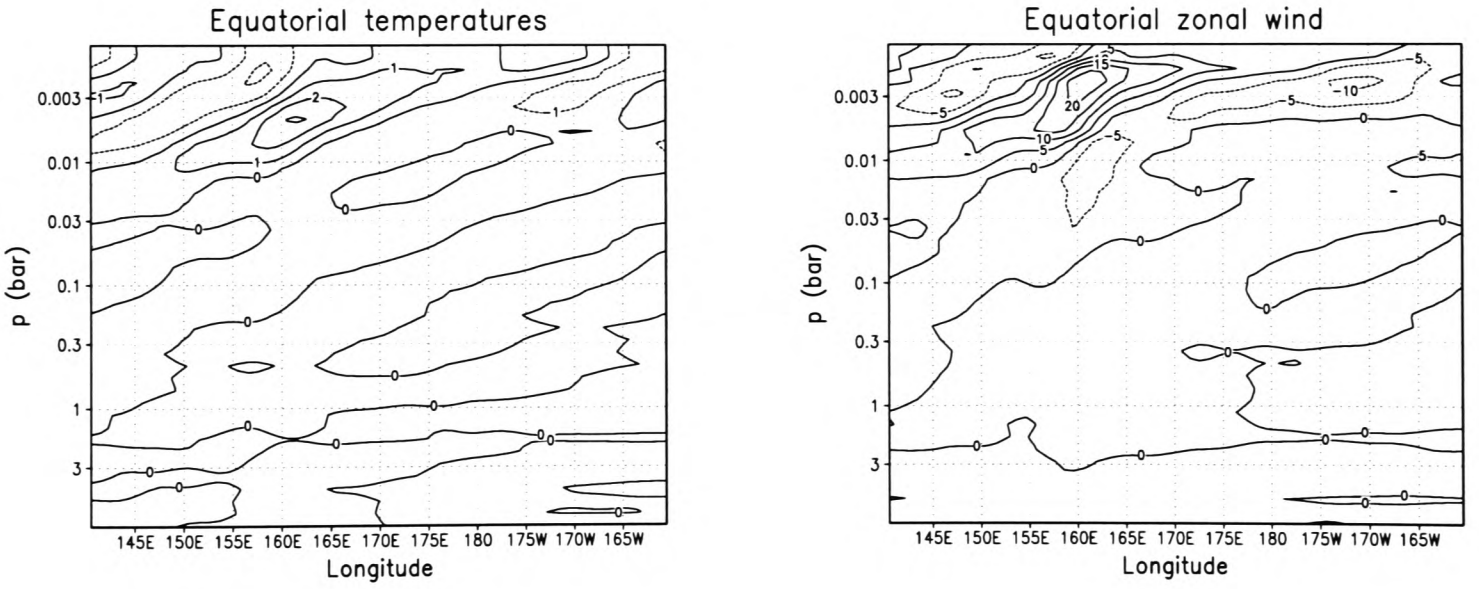


Figure 6.33: Temperature anomaly and zonal wind on the equator after 400 days, forced by a broad band of waves of equal phase speed. The waves have not excited any significant flow anywhere except in the sponge layer; we conclude that this phase speed is unrealistic.

The second case, $c = 145$ m/s, is more interesting (see Fig. 6.34). The waves are now able to deposit momentum more efficiently. A very large jet is generated, filling the upper troposphere and extending into the lower stratosphere. Its speed reaches 250 m/s at the 2 bar levels, and is faster than 150 m/s up to the tropopause. The large meridional shear of the stratospheric wind near the sidewalls suggests that the finite channel width is affecting the flow.

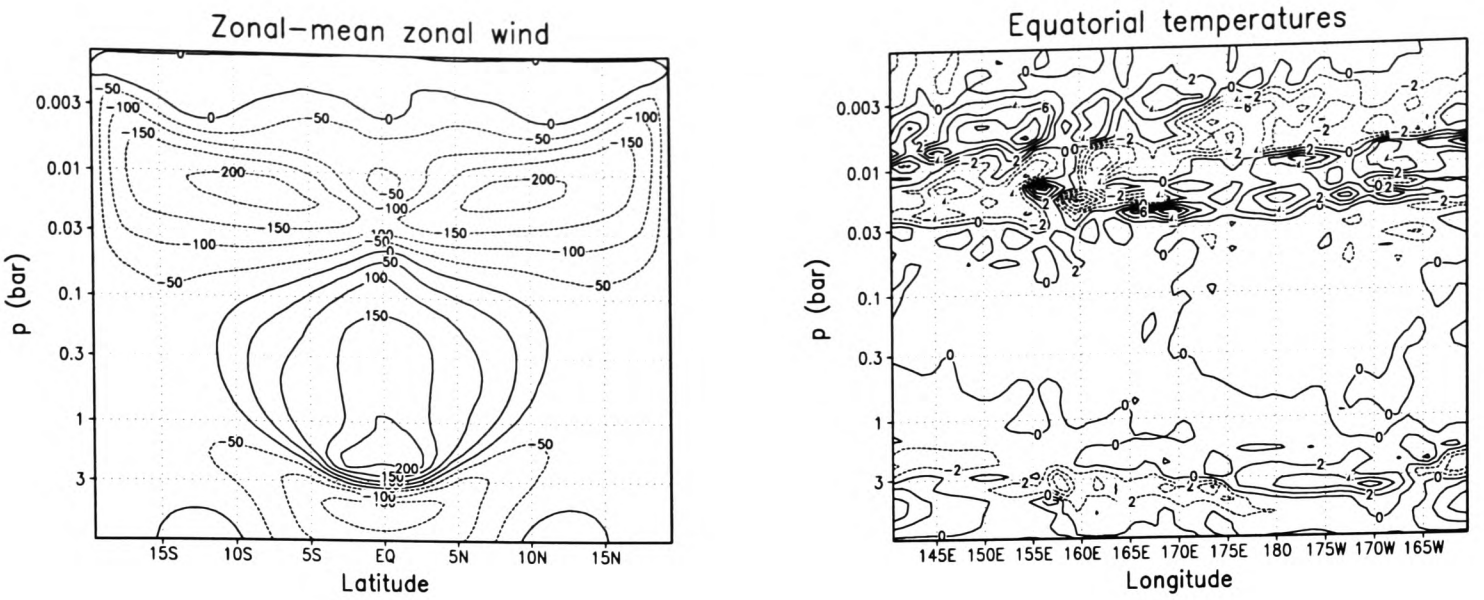


Figure 6.34: Zonal-mean zonal wind, and equatorial temperature anomaly after 300 days, forced by a broad band of waves of equal phase speed.

Spectral analysis of the early part of the model run (Fig. 6.35) shows that the forced waves initially propagate up to the 11 mb level undistorted. The spectrum is linear in frequency and wavenumber, so all modes have equal c as required. A weak low-pass filtering effect seems to be acting, because the power decreases slowly with wavenumber.

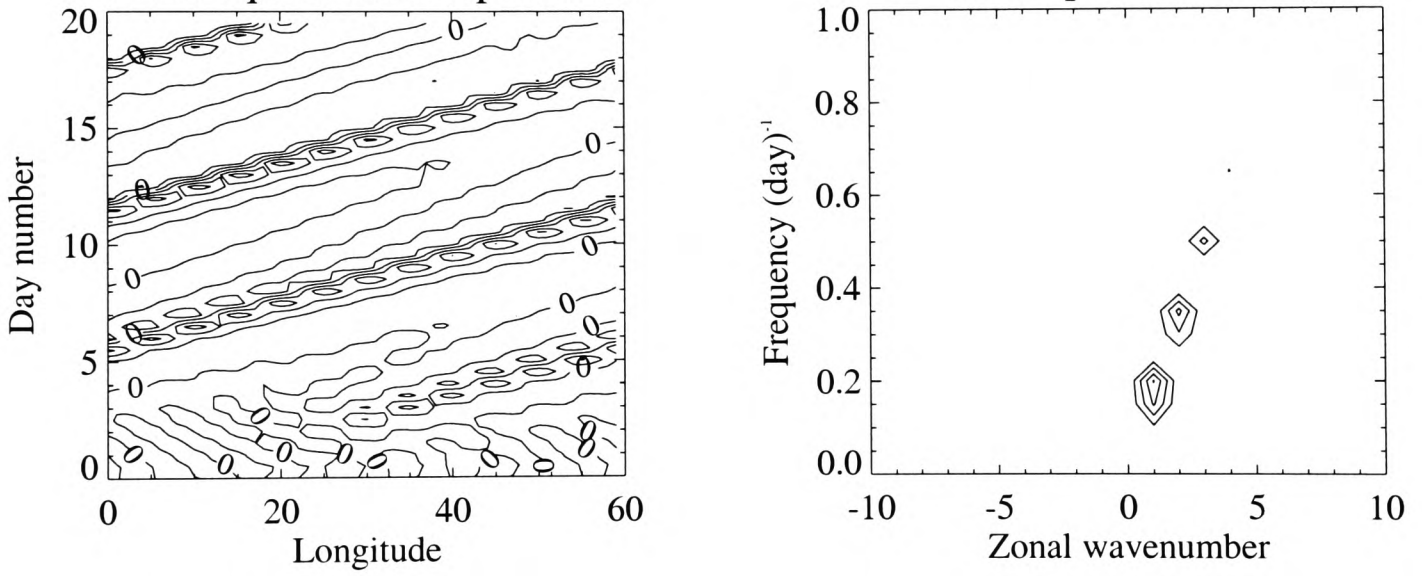


Figure 6.35: Time-longitude plot and power spectrum of temperatures at 11 mb on the equator, for the initial stage of the flow shown in Fig. 6.34.

However, later in the run, the equatorial temperature wave spectrum of Fig. 6.36 has a peak centred on wavenumber +1 with period 5 days. In this case it appears the model atmosphere (presumably because of the Doppler effect in the jet) is low-pass filtering the signal. Comparison of the early and late spectra suggests that this 5-day wave is just the gravest mode present initially; the others have been removed. However, note that this jet is perhaps rather faster and wider than observed (or normally generated in the model).

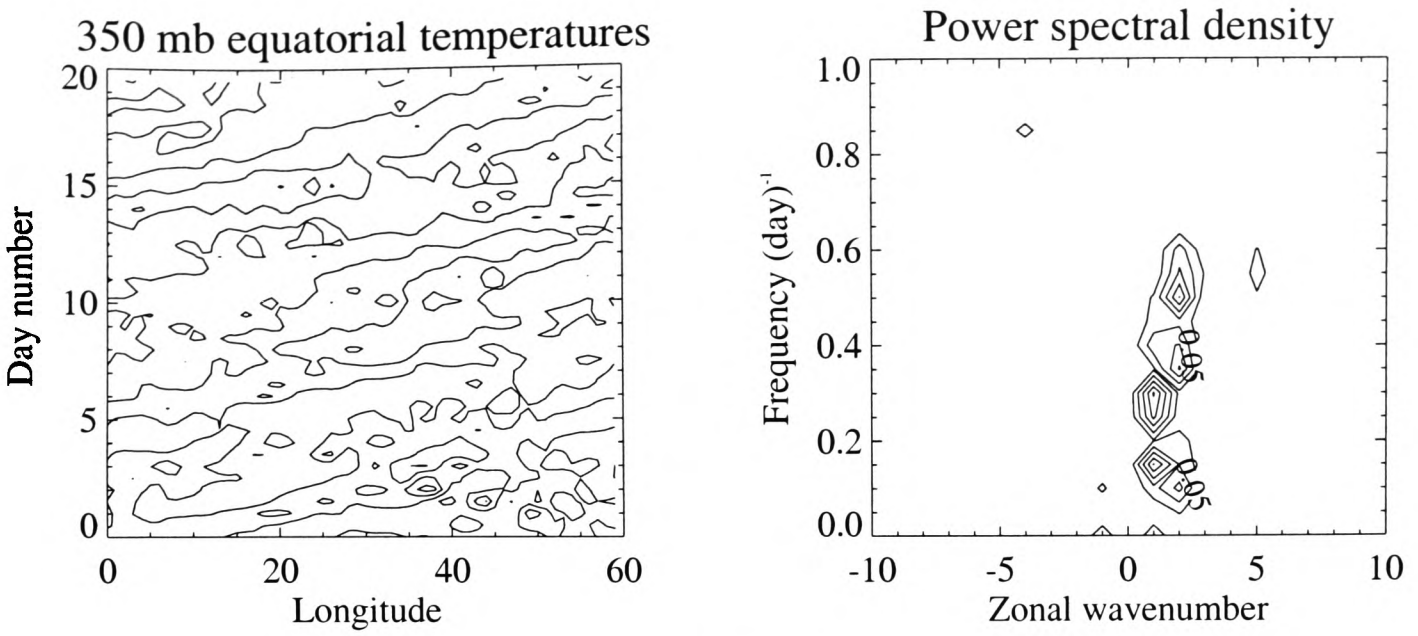


Figure 6.36: As Fig. 6.35, but for the developed stage of the flow shown in Fig. 6.34.

6.11 Waves forced by a single stationary source

In this section we shall consider the atmosphere's response to a stationary source on the lower boundary, oscillating coherently in time. We investigate the effects of various sources, including those chosen by Holton (1972). The situation is more complicated than before, because the sources are not points, but have a finite longitudinal extent. This means they potentially excite all modes with wavelengths greater than the length of the source but shorter than the channel length. Also, unlike the case of a travelling wave, we cannot restrict the response to eastward or westward propagation, but must allow both.

6.11.1 Review of some relevant terrestrial results

Nearly thirty years ago, satellite observations of tropospheric clouds above the western Pacific (Wallace, 1971) revealed westward waves of wavelength $\sim$ 2000–4000 km and period 4–5 days, and also a stationary oscillation in cloud brightness of a similar period. The latter is longitudinally confined, and latitudinally antisymmetric about the equator. Holton (1972) modelled equatorial Kelvin waves forced by this ITCZ standing-wave heat source. Holton represented the ITCZ by a function with meridional part

$$Q(y) = (y/y_0)\exp\left[-y^2/\left(2y_0^2\right) + 1/2\right] \tag{6.19}$$

where $y = \sin\phi$, $y_0 = \phi_0$ and ϕ_0 is the latitude of maximum heating.

The longitudinal part was expressed as another Gaussian:

$$Q = Q(y,z)\exp\left[-(\lambda/\lambda_0)^2\right]\cos\omega t \tag{6.20}$$

with $\lambda_0 = 36°$. This function was expanded as a Fourier series of five zonal wavenumber components, each with the same frequency but different phase speed. The wavenumber

distribution of Eq. 6.20 is red, i.e. the Fourier amplitudes of smaller wavenumber k are greater than those of larger k.

Holton found that the stratospheric response (in a sheared zonal wind) was dominated by the $k = 1$ modes. For the westward-propagating symmetric Rossby modes, only this mode has high enough phase speed c to penetrate the easterly jet. Among the eastward propagating Kelvin modes, the $k = 1$ wave also dominated, but this was because vertical wavelength is proportional to zonal wavelength, for fixed ω. Holton's forcing function was quite thick vertically, so large vertical scales were excited more efficiently. Hence long zonal wavelengths ($k = 1$) were preferred. In contrast, antisymmetric forcing excited all five zonal components.

Chang (1976) later showed in a simple linear model that the assumption of a red-noise spectral distribution is unnecessary. For a given vertical scale, the upward energy flux across the tropopause was proportional to the zonal wavelength even for a white-noise distribution, so long as the heating rate is independent of the zonal wavenumber.

6.11.2 Stationary sources on Jupiter

We can only guess at the details of the tropospheric forcing on Jupiter, so our approach will be to try various possibilities suggested by this terrestrial work, and to look for robust features of the response. The water cloud is probably concealed beneath at least two other cloud decks, plus tropospheric haze above, so we cannot expect to observe directly an ITCZ-style oscillation on Jupiter. However we can explore the possibilities it presents. We do this by considering monotonic and bitonic symmetric sources, and also a bitonic antisymmetric source.

6.11.3 A monotonic, symmetric source on the equator

This is the simplest case to consider, for it has only a single peak centred on the equator. The horizontal form of the source is the Gaussian

$$\Delta T = A \exp\left(-\phi^2/16\right) \exp\left[-(\lambda - \lambda_0)^2/16\right] \tag{6.21}$$

while the vertical form is the usual $\cos(z - z_0)$, and the temporal variation is sinusoidal. The period for this experiment is 10 days.

Figure 6.37 shows the results of this forcing. There is evidence of waves in the stratospheric temperatures but the spectral composition is not clear (see below). The time-height plot (lower right) shows the initial development of a series of stacked jets by day 50. The higher-level ones decay with time however, leaving only the tropospheric one

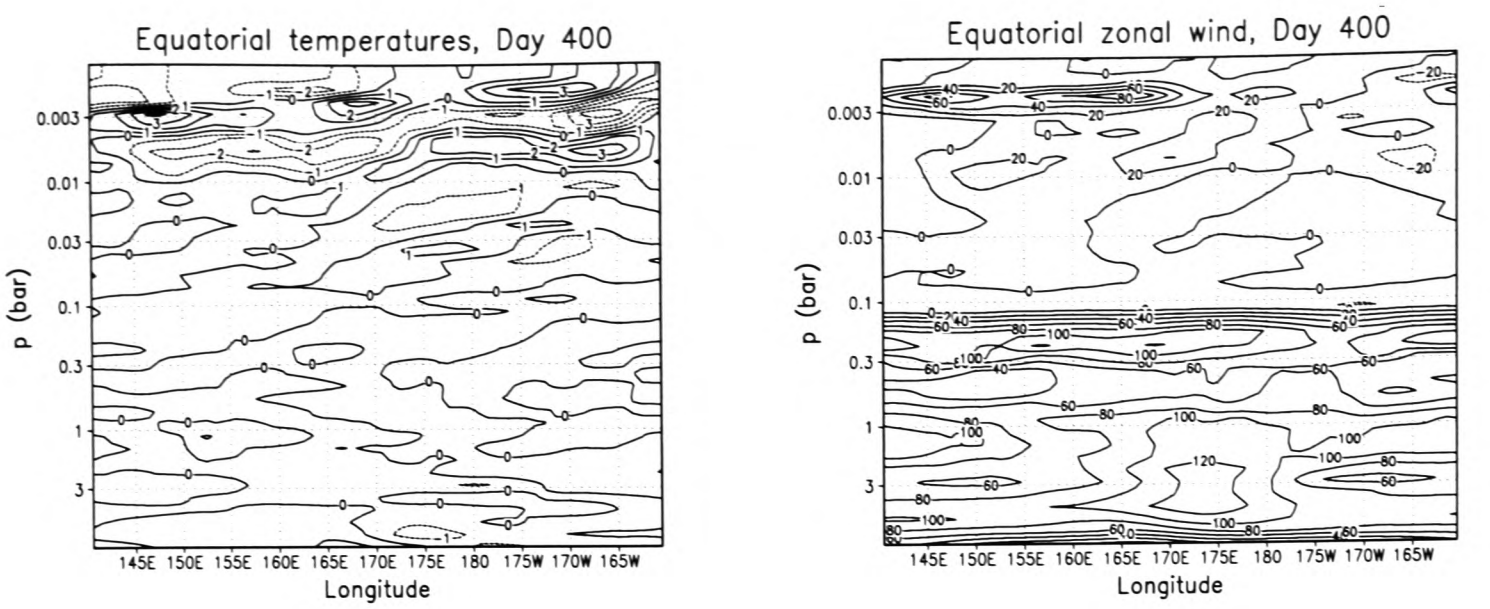


Figure 6.37: The flow forced by a monotonic symmetric source. The top plots show the final state on the equator: temperatures on the left, winds on the right. The zonal wind plot shows that a westerly jet has been spun up throughout the equatorial troposphere. The lower-left panel is the latitude-height cross section of the zonal mean zonal wind at the end of the run. This shows the westerly centred on the equator, flanked by weaker easterlies. On the lower right is a time-height series of zonal-mean zonal wind on the equator. Initial stratospheric jets disappear whilst the main tropospheric jet is stable.

which persists until the end of the integration. The stacked pattern arises because initially the waves can propagate (through the resting atmosphere) up to the top of the model. The subsequent decay is probably due to the damping in the troposphere. The high level jets cannot later reform, because the wave is strongly absorbed in the tropospheric jet. It therefore does not penetrate into the stratosphere.

The spectrum of waves was studied with a shorter run of the early part of this experiment, with output every 12 hours (Fig. 6.38) and 6.39). The tropospheric temperatures (at the level of the forcing) show both eastward and westward long-wavelength modes, with the eastward stronger. In the stratosphere, only the eastward mode is present. This is consistent with the later behaviour of the model, i.e. the appearance of an eastward jet.

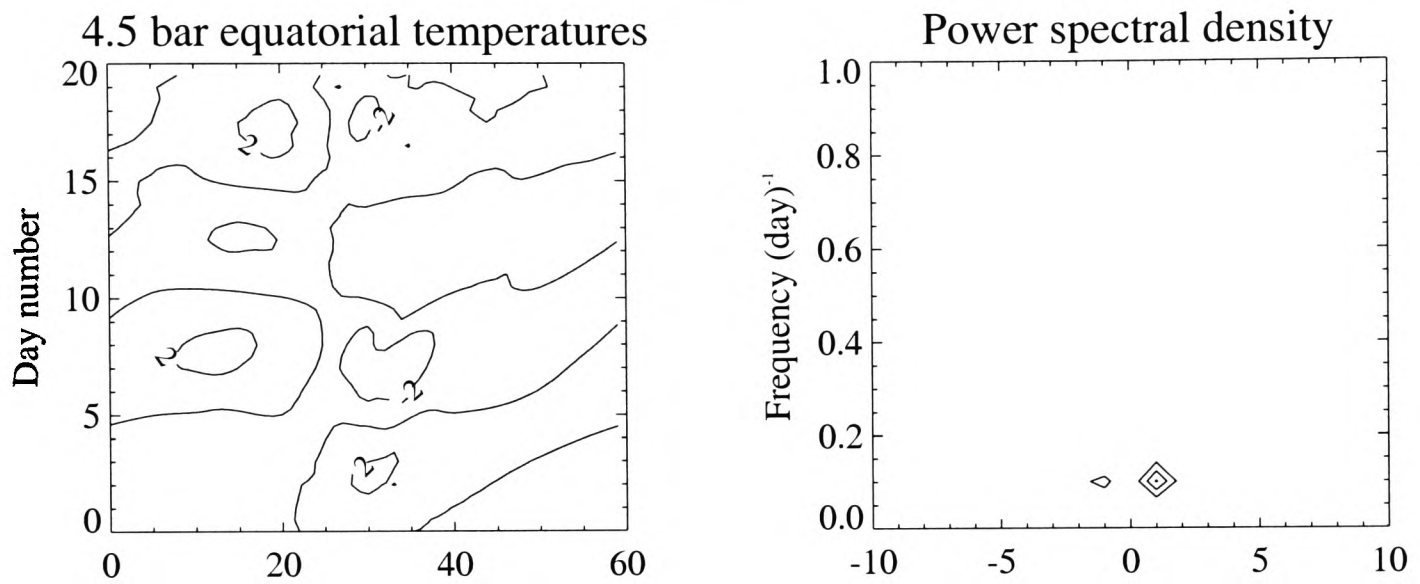


Figure 6.38: Tropospheric temperature anomaly on the equator, and its power spectrum. The model is the same as shown in Fig. 6.37. Wavenumbers ±1 are dominant, with the +1 mode the stronger.

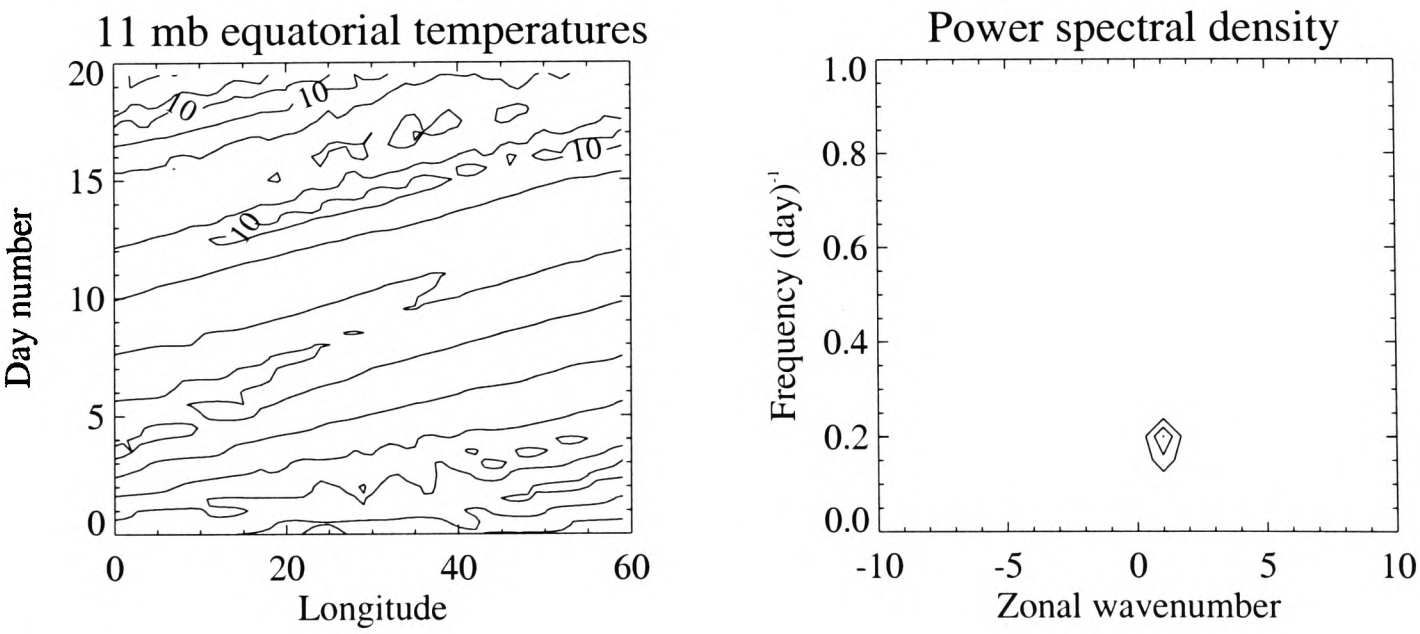


Figure 6.39: Stratospheric temperature anomaly on the equator, and its power spectrum. The model is the same as shown in Fig. 6.37. Wavenumbers ±1 are dominant, with the +1 mode stronger.

The jet produced by this source is perhaps a reasonable first model of that observed on Jupiter. It has the correct direction (unlike the jets produced by other sources we shall

consider later), and peaks at roughly the cloud-top height.

Other features are less realistic when compared to the observations: the peak wind of 80 m/s is rather slow, and the vertical shear above the jet is too great. The latitudinal width is much too narrow, with tropical easterlies at 7° instead of around 20°. However, as pointed out in Section 6.4.2, this is precisely what we would expect from a Kelvin wave alone.

In principle all these deficiencies could be addressed by modifying the shape and frequency of the source or by tuning the stratospheric drag. We shall very briefly examine the sensitivity to these parameters.

Frequency dependence

To explore the sensitivity of the previous result, the model was re-run with the source frequency increased by 10%. Now if the source were emitting one dominant mode responsible for the jet, we would expect the jet speed to increase by the same factor. However, this is not quite what happens (Fig. 6.40). The 80 m/s jet is replaced by one of 100 m/s, with a slightly different vertical structure. A 10% increase in source frequency has led to a 20% increase in jet speed. This sensitivity shows that if a convective source like this were to force an equatorial jet of the speed observed on Jupiter, its oscillation period should be around (9 ± 1) days. Further experiments would be required to assess the effect of multiple sources, however.

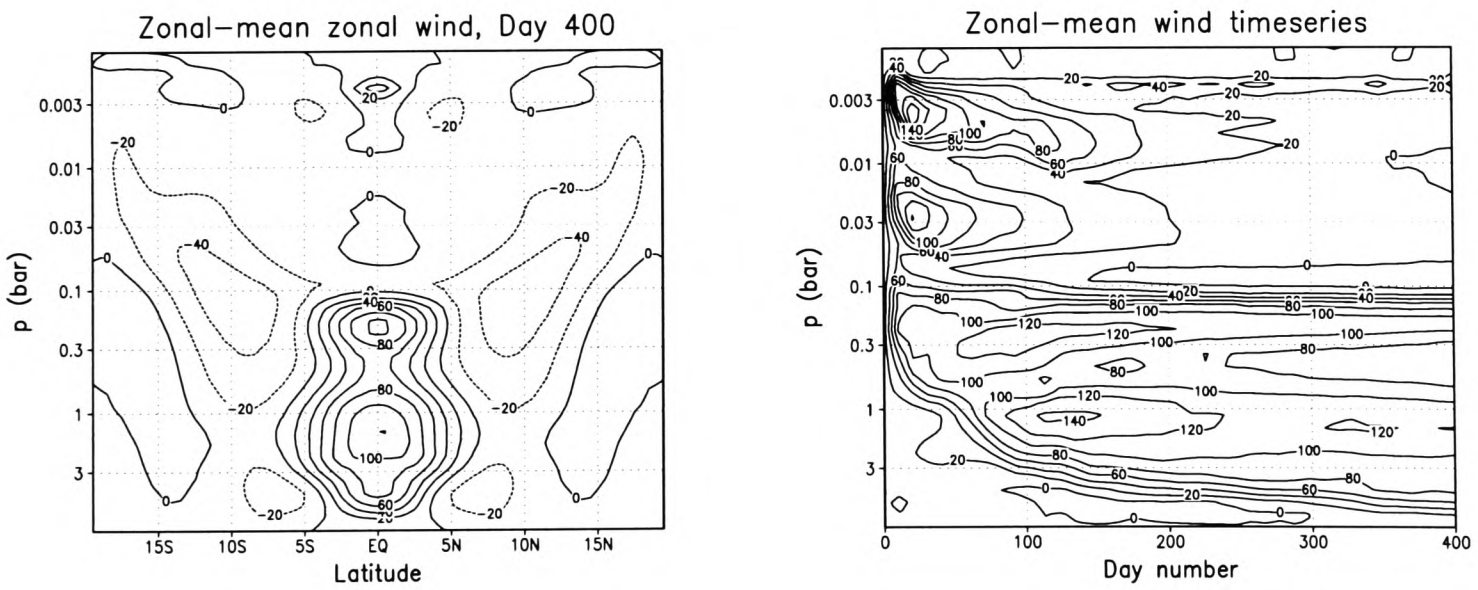


Figure 6.40: Results with a source of 10% higher frequency; compare with Fig. 6.37. The flow structure is similar but the peak speed has increased to 100 m/s (left panel). The right panel shows the time evolution of the zonal-mean wind at the equator.

Effect of stratospheric drag profile

As an extreme example of a different drag profile, the run was repeated with uniform drag in the troposphere as well as the stratosphere, that is throughout the model atmosphere (except in the sponge layer which was left unchanged). The flow is shown in Fig. 6.41. There are now westerly jets in the stratosphere as well as the troposphere, and strong easterlies in the tropical stratosphere. However, the flow in the troposphere is not very different to that seen in the other model, so we can hope that the details of the drag parametrization are not too important.

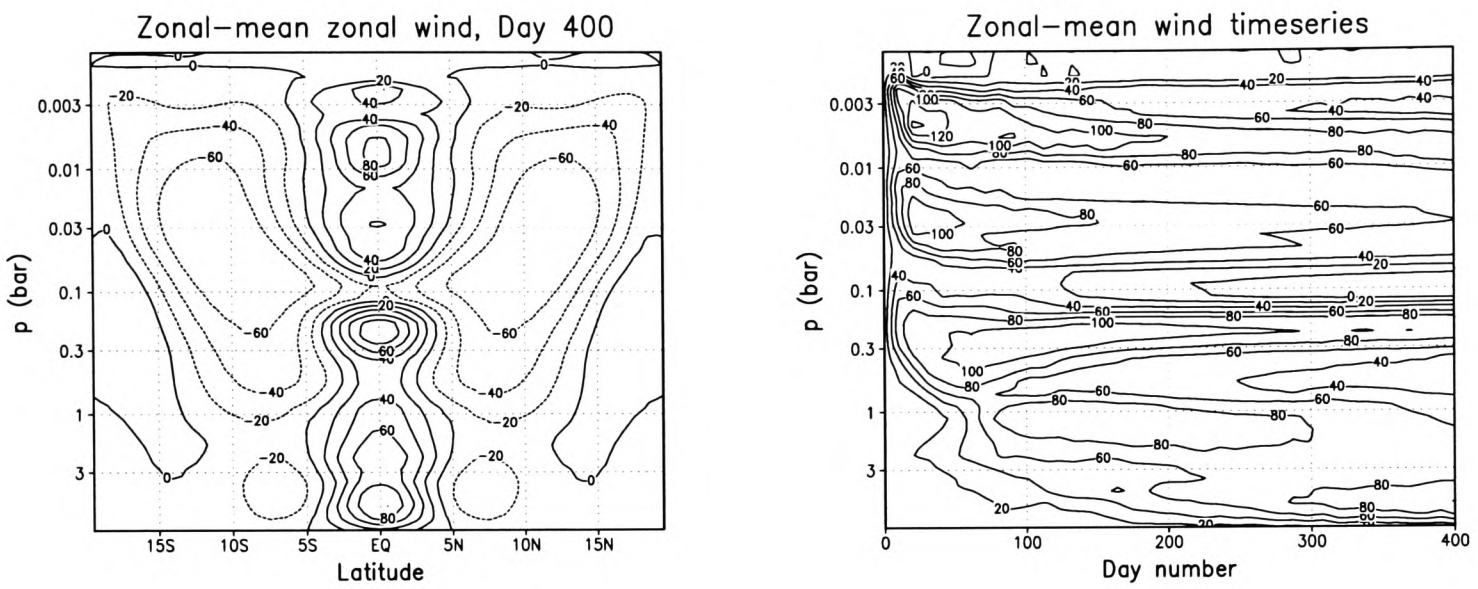


Figure 6.41: Results with a source of 10% higher frequency; compare with Fig. 6.37. The flow structure is similar but the peak speed has increased to 100 m/s (left panel). The right panel shows the time evolution of the zonal-mean wind at the equator.

6.11.4 A bitonic symmetric source

Following Holton, we look at the flow forced by a source of the form given by Eqs. (6.19) and (6.20). This is bipolar and symmetric in ϕ about the equator, with a maximum heating effect at latitudes $\phi = \pm 6°$. Figure 6.42 shows the results. The equatorial temperature plot shows evidence of eastward modes in the stratosphere, while the plot of meridional wind on the equator shows westward (probably Rossby) modes. It may be surprising that this source generates eastward Kelvin waves, since its meridional structure does not match that of the waves. However, because the source is symmetric, it is able to excite the Kelvin wave although not as strongly. An ITCZ-style symmetric source produced Kelvin waves in Holton's model too.

The westward mode is not seen between 300 mb and 100 mb, which is consistent with the zonal wind plot (lower left) showing an easterly jet at this level. Clearly the easterly (westward) jet is absorbing the westward wave mode. Some energy remains to propagate higher though. The timeseries of zonal-mean zonal wind against height indicates that the flow has not reached equilibrium after 400 days. Unfortunately, computing resources

did not allow long integrations of all cases, so we restrict ourselves to discussion of only a few examples of particular interest.

The presence of an equatorial easterly jet in the model suggests that this case is not relevant to Jupiter. It may be more appropriate to Uranus or Neptune, since both these planets have easterly jets centred near the equator. However, there are many other considerations which prevent us from drawing a simplistic conclusion. For example, the very large orbital inclination (97.9°) of Uranus means that the annual-mean solar heating at its poles is probably greater than at the equator. This means that something akin to a reverse Hadley cell may be operating, and the overall picture is complicated.

To check that the result is not strongly dependent on our arbitrary choices, we repeat the experiment firstly in a channel of half the length (Figure 6.43) and secondly in the longer channel but with a source of shorter length. For the latter, we set $\lambda_0 = 18°$ in Eq. (6.20), to replace the previous value of 36°.

A similar flow is obtained by repeating the experiment in a channel of half the length, i.e. 60 degrees in longitude. As shown in Fig. 6.43, the central easterly jet is almost identical in position, shape and magnitude. The westerly above this is somewhat stronger in the shorter channel.

Finally (Fig. 6.44), changing the longitudinal size of the source can be seen to have little effect on the results. The flow is weaker, because the source is smaller but has unchanged amplitude. Therefore the total power input is smaller and the response correspondingly so.

It seems that this result—bitonic symmetric heating in the tropics causes easterly equatorial jets—is insensitive to the details of the forcing. On the Earth, where the source represents the ITCZ, weak easterlies are indeed observed on the equator. We therefore conclude that convective sources on Jupiter are more likely to be monotonic and centred near the equator than near the tropics, because we have seen that a source on the equator can generate superrotation similar to that observed.

6.11.5 An antisymmetric source

For completeness we include the results of an antisymmetric, bitonic source. This has the same structure as the symmetric one described above, except that the two extrema of the function have opposite signs of heating. The period was 6 days. The results (Fig. 6.45) are rather different, in that no zonal jet is spun up over 400 days of model integration. Instead we see a mixed Rossby-gravity wave, propagating almost unaffected into the stratosphere. Not surprisingly, the wave looks very like that in Holton's model. The temper-

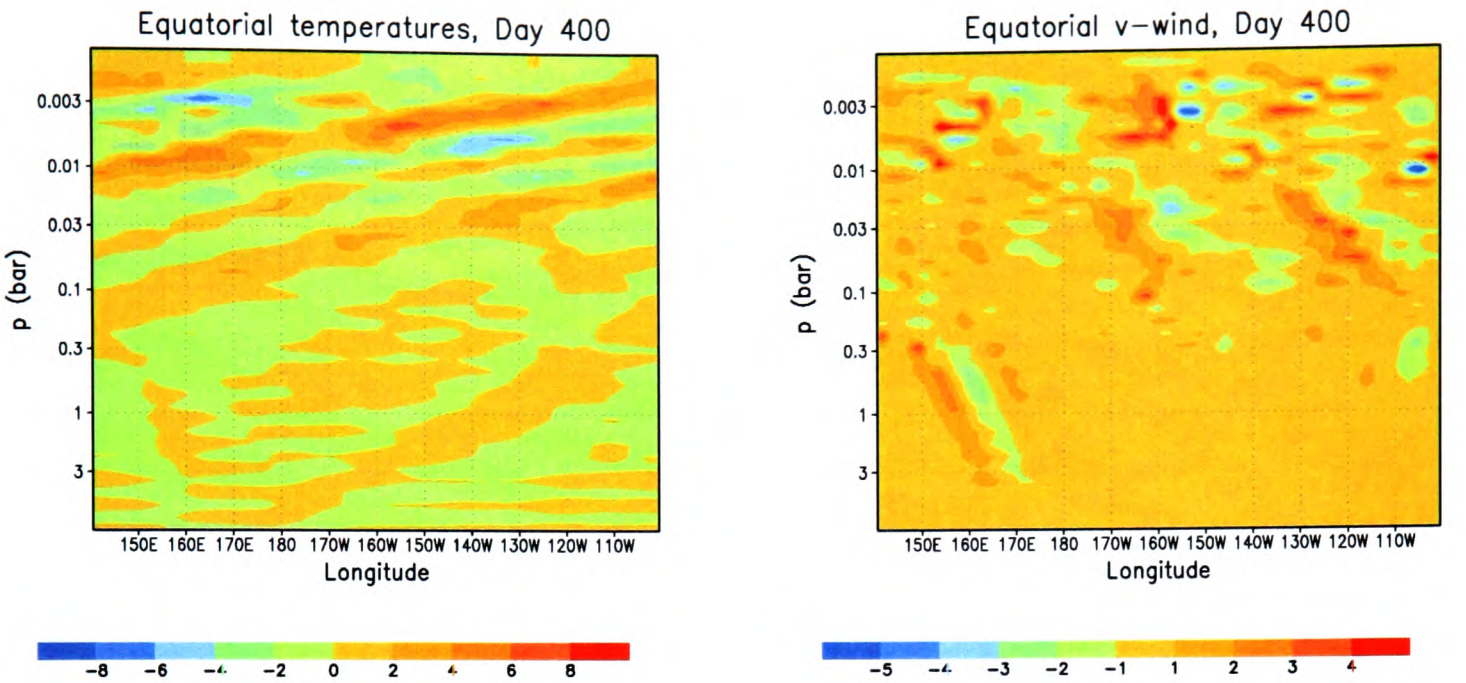


Figure 6.42: Flow forced by an ITCZ-style symmetric source. Above: temperatures and meridional winds at the equator, showing westward and eastward modes. Below: zonal-mean zonal wind, as a latitude-height section and a timeseries.

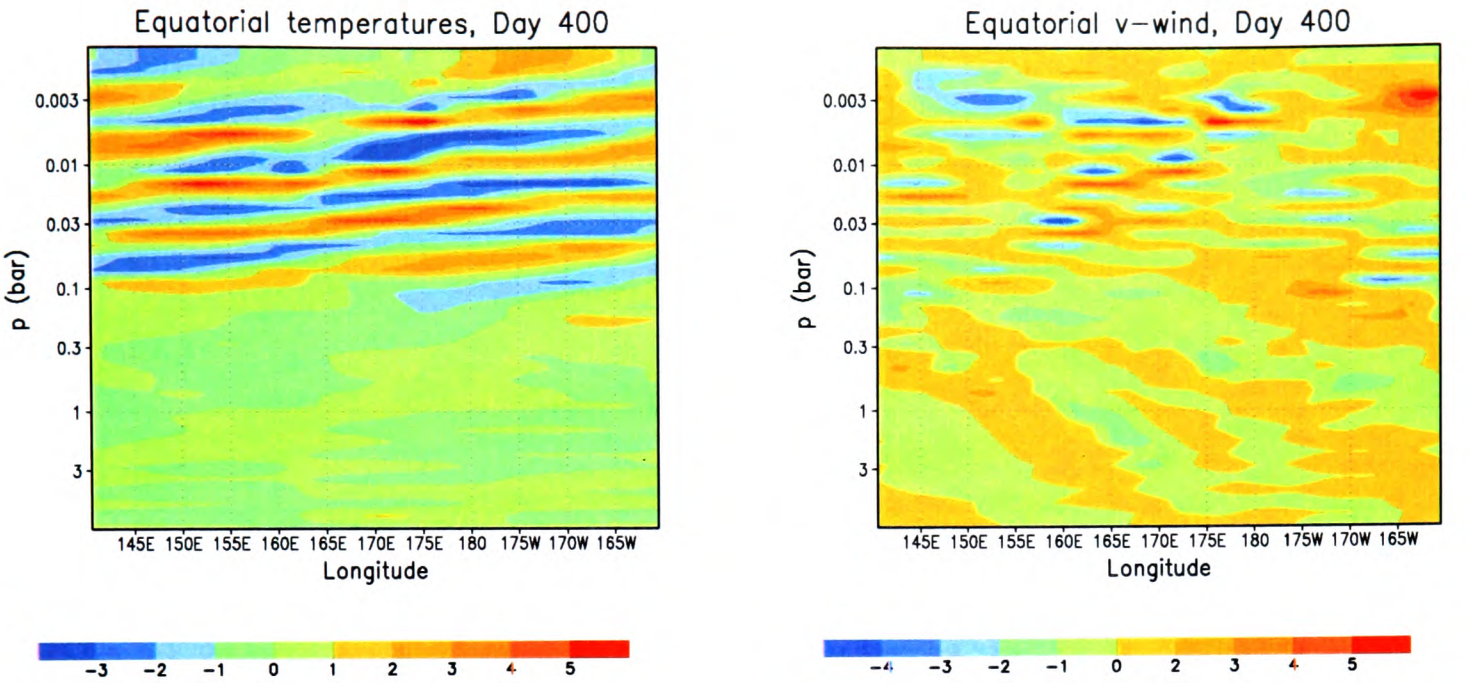


Figure 6.43: Results of a similar experiment in a channel 60 degrees long. Above: temperatures and meridional winds at the equator, showing westward and eastward modes. Below: zonal-mean zonal wind, as a latitude-height section and a timeseries.

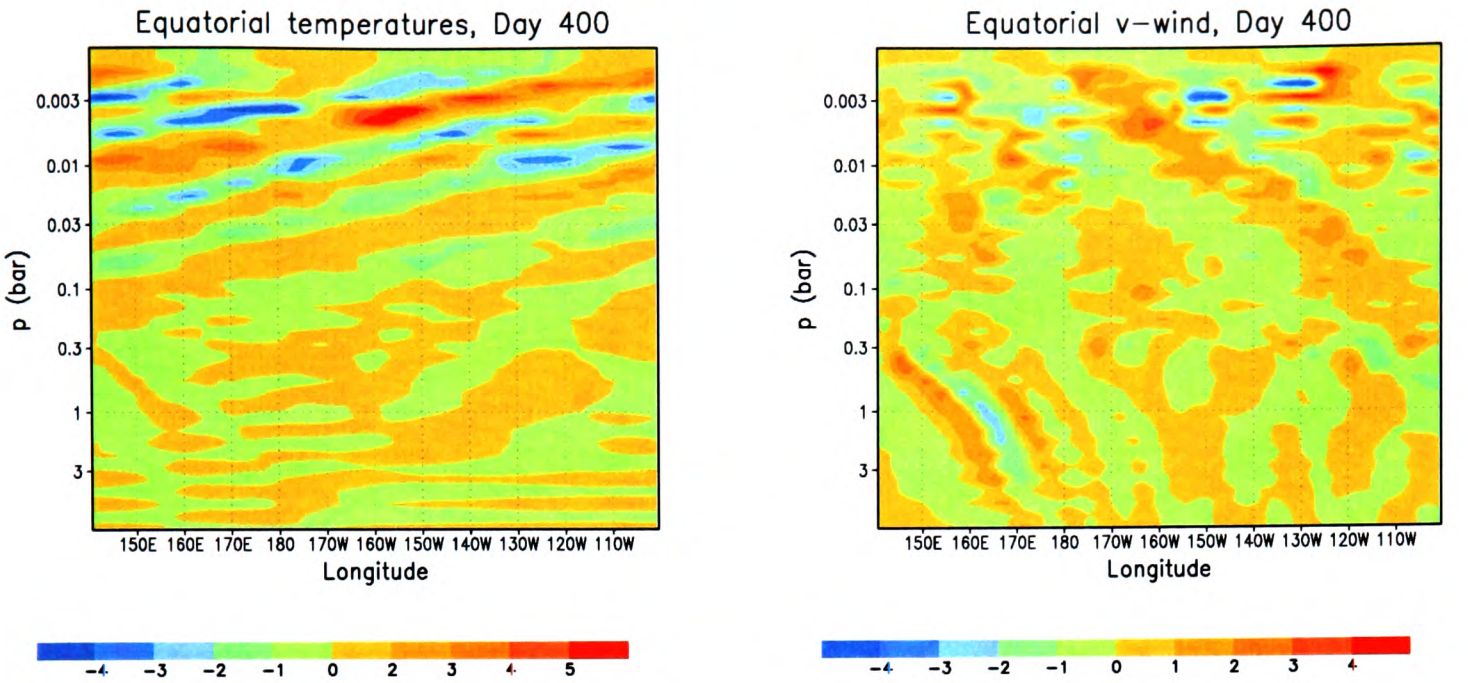


Figure 6.44: Results of a similar experiment in a channel 120 degrees long but with a smaller source. Above: temperatures and meridional winds at the equator, showing westward and eastward modes. Below: zonal-mean zonal wind, as a latitude-height section and a timeseries.

ature signal seems to propagate purely vertically in the troposphere and then eastward in the stratosphere. The strong signal in the meridional wind propagates eastward and upward. As Holton points out, the *group* velocity of an MRG wave is eastward, although the phase velocity is westward. Hence there is no conflict with our previous results. The lower two plots confirm the diagnosis of an antisymmetric, mixed Rossby-gravity wave: the temperature field is antisymmetric while the v-wind is symmetric. There seems to be no evidence of a Kelvin wave—and we would expect none, as the forcing has no symmetric component.

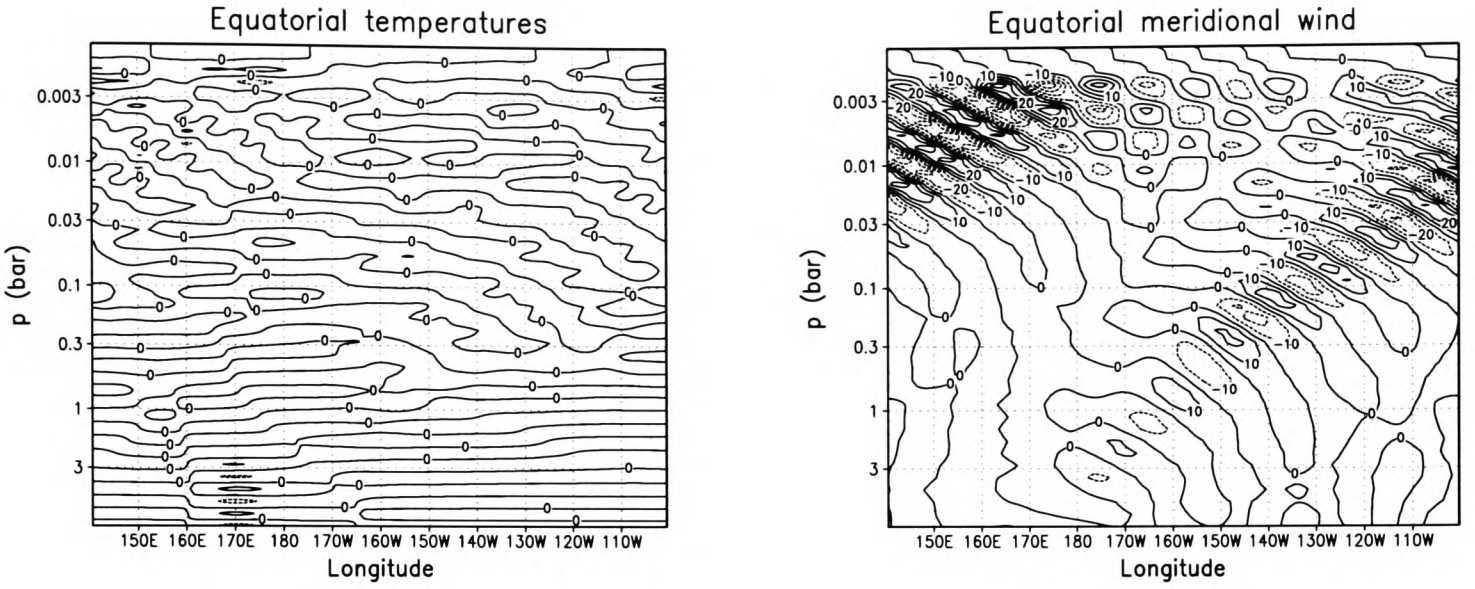


Figure 6.45: Waves generated by an antisymmetric stationary source, shown after 400 days. The source is located at $170°E$. The wave in the temperature anomaly propagates straight upward, as there is no significant zonal flow generated. The v-wind plots show an eastward-propagating antisymmetric mode. The lower two panels are pressure sections in the stratosphere, showing that the temperature is antisymmetric and the v-wind symmetric, i.e. an MRG mode.

6.12 Linear models of the Hadley circulation

An axisymmetric, thermally direct overturning motion was suggested by Halley (1686), and then Hadley (1735), to explain the Earth's trade winds and the associated low-pressure zone over the equator. Such a circulation will also generate westerly jets at the poleward sides of the cell. The double-peaked structure of Jupiter's equatorial jet (Fig. 6.1), with its local minimum on the equator, suggests that angular momentum transport by a Hadley cell may be partly responsible for generating the jet, or at least for modifying its profile from a single peak to a double one.

6.12.1 An overturning circulation, conserving angular momentum

Held and Hou (1980) developed a simple analytic model of the Hadley cell, a recent description of which may be found in James (1994). The essential features are as follows. Air near the surface of the planet is assumed to be at rest. Solar heating drives convective ascent in the tropics, with associated divergence above. The meridional flow must conserve angular momentum, hence giving rise to a pair of westerly jets. Descent in the tropics and convergence to the equator at low levels complete the cycle.

We attempt to represent some aspects of Jupiter's equatorial jet analytically using a modified Held-Hou model, in which the winds are a maximum at the base of the model (the cloud-tops) and decay with height. We also introduce a nonzero zonal wind $u = u_0(z)$ at the equator. The following argument is formulated in spherical coordinates and log-pressure pseudoheight (cf. Chapter 2). We assume zonal symmetry and use absolute temperature, not potential temperature as used by Held and Hou. The argument is the same but our units are consistent with the equations elsewhere in this thesis.

The specific angular momentum of the atmosphere at latitude ϕ is

$$M(\phi) = (\Omega a \cos\phi + u)\, a \cos\phi \tag{6.22}$$

which at the equator is

$$M(0) = \Omega a^2 + u_0 a \tag{6.23}$$

Conservation of angular momentum implies that

$$\Omega a^2 \cos^2\phi + ua\cos\phi = \Omega a^2 + u_0 a \tag{6.24}$$

so the zonal wind is given by

$$u(\phi, z) = \frac{\Omega a \sin^2\phi + u_0(z)}{\cos\phi} \tag{6.25}$$

For a narrow cell of halfwidth $7^\circ = 0.12$ rad, $\sin\phi \simeq y/a$ and $\cos\phi \simeq 1$ and Eq. (6.25) simplifies to

$$u(\phi, z) = \frac{\Omega y^2}{a} + u_0(z) \tag{6.26}$$

The temperature field may be calculated if we assume thermal wind balance:

$$2\Omega \sin\phi \frac{\partial u}{\partial z} = -\frac{R}{H_0}\frac{\partial T}{\partial y} \tag{6.27}$$

and also assume linear wind shear from u_0 at the cloud-tops to zero at ($z = Z, p = 10$ mb), i.e.

$$\frac{\partial u}{\partial z} = -\left(\frac{\Omega y^2}{aZ} + \frac{u_0}{Z}\right) \tag{6.28}$$

where $z = 0$ is the cloud-top level and Z the height of the cell above this. The combination of Eqs. (6.27) and (6.28), using the small-angle approximations, yields the temperature gradient:

$$\frac{\partial T}{\partial y} = \frac{2H_0\Omega y}{RZa}\left(\frac{\Omega y^2}{a} + u_0\right) \tag{6.29}$$

which has the solution

$$T(y) = T_0 + \frac{H_0\Omega}{aRZ}\left(\frac{\Omega}{2a}y^4 + u_0 y^2\right) \tag{6.30}$$

The thermal forcing is taken to be Newtonian cooling towards an equilibrium temperature distribution

$$T_E(y) = T_{E0} - \frac{\Delta T}{a^2}y^2 \tag{6.31}$$

where T_{E0} is the global-mean temperature and ΔT the pole-to-equator difference. (The small angle approximation has been incorporated into this expression.)

Held and Hou extended their analysis to give an estimate of the cell halfwidth Y. This is done by equating the integral of (6.30) within the cell to that of the radiative equilibrium temperature, thereby requiring zero net heating:

$$T_0 + \frac{H_0\Omega}{aRZ}\left(\frac{\Omega}{10a}Y^4 + u_0 Y^2\right) = T_{E0} - \frac{\Delta T}{3a^2}y^2 \tag{6.32}$$

Also, at the edge of the cell $T = T_E$ so

$$T_0 + \frac{H_0\Omega}{aRZ}\left(\frac{\Omega}{2a}Y^4 + u_0 Y^2\right) = T_{E0} - \frac{\Delta T}{a^2}Y^2 \tag{6.33}$$

Equations (6.32) and (6.33) may be solved to give

$$Y = \left(\frac{5\Delta T RZ}{3H_0\Omega^2}\right)^{1/2} \tag{6.34}$$

6.12.2 Results of the modified Held-Hou model

Firstly we can predict the cell width from Eq. (6.34). The values of the parameters required are not particularly well known, but we can estimate $\Delta T \simeq 20$K (this is generous) and $Z = 2H_0$. Other quantities are $H_0 = 21.5$ km, $\Omega = 1.774 \times 10^{-4}$ rad/s, $R = 3745$ J/kg, and $a = 71995$ km. These give $Y \simeq 4000$ km $\simeq 3°$.

We insert the Jovian values into Eqs. (6.26) and (6.30). Figure 6.46 shows the results. We have taken $u_0 = 100$ m/s as a reasonable value for the speed of the core of the jet, and $4°$ as the latitudinal half-width. This has led to predicted zonal winds of 150 m/s at the poleward edges of the cell. The observed winds do reach 150 m/s but not until about $10° N$. In other words, the observed increase of wind speed away from the equator is less rapid than would be the case with uniform angular momentum. Our estimate of ΔT is large; a more conservative value, such as 10K, would give an even narrower cell. Therefore the mismatch between model and observations is not really sensitive to ΔT, as no reasonable value could lead to agreement.

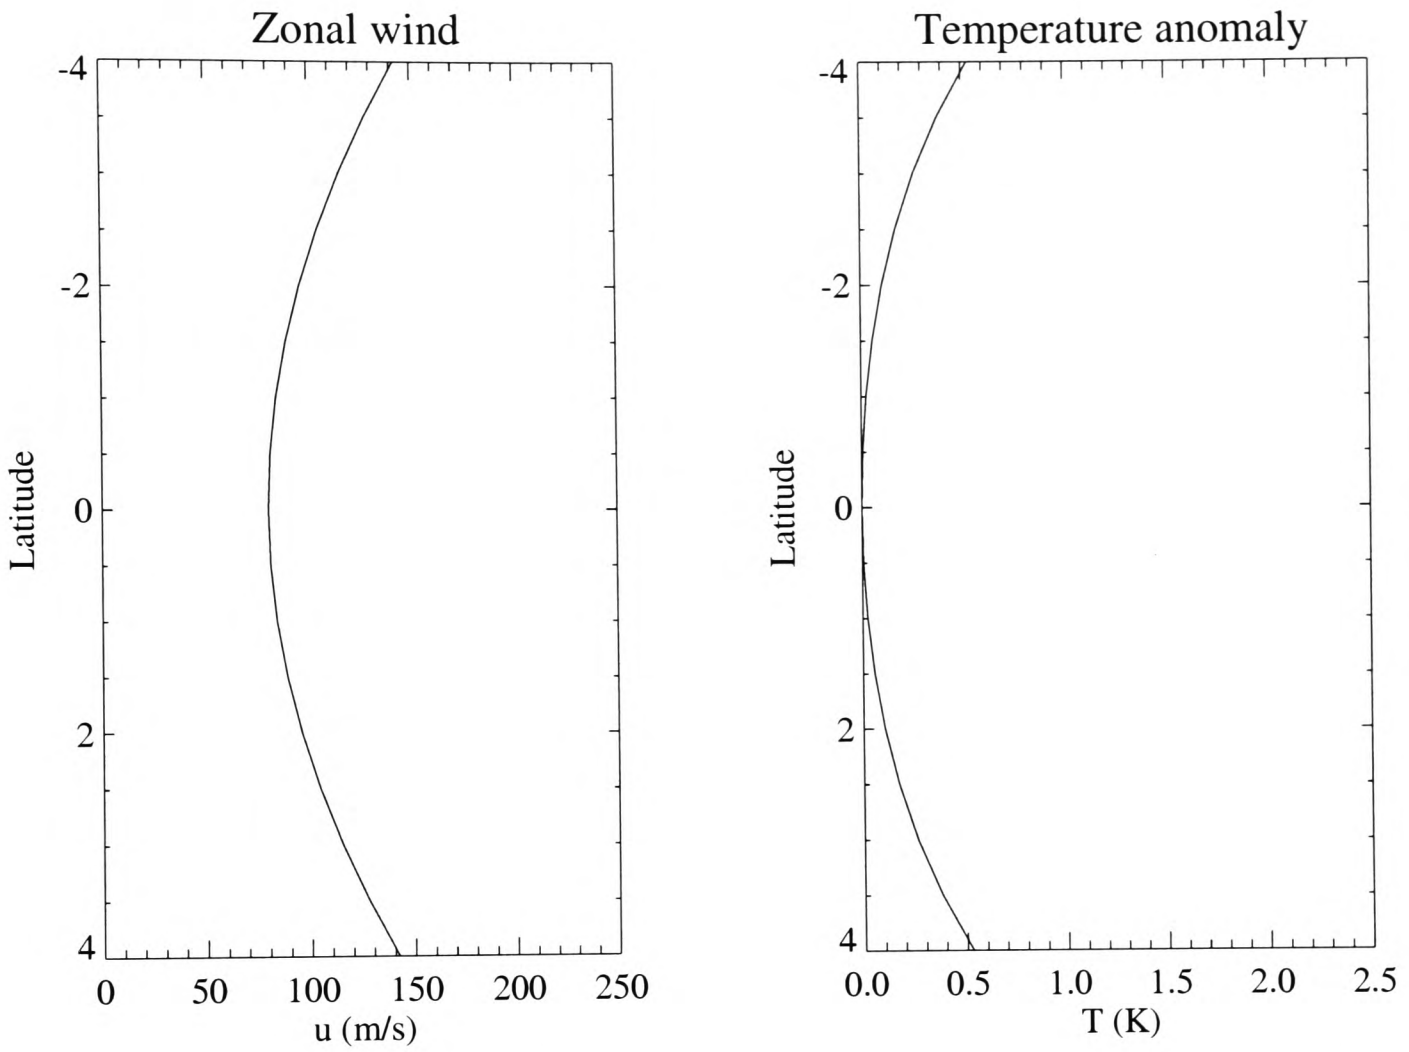


Figure 6.46: Cloud-top winds and temperatures predicted by the inverted Held-Hou analytic model. The range shown is from 4°N to $4° S$, slightly wider than the predicted width of the cell.

We conclude that the Hadley circulation may be a plausible explanation for some features of the equatorial jet, although it cannot account for the strong westerlies at the equator, whose presence we have assumed *a priori*. Nevertheless, the wind profile for $-4° < \phi < 4°$ does agree qualitatively with the observed increase in u away from the equator. The model wind shear is much greater than observed, indicating that on Jupiter, other processes are probably damping the winds of the Hadley cell.

6.13 Symmetric forcing of a Hadley cell in the GCM

Having established the shortcomings of the linear model of a Hadley cell, we next investigate the behaviour of the Hadley cell in the simple general circulation model. This is done initially in isolation and then in conjunction with wave forcing. We need to choose a suitable forcing function which is is zonally symmetric, incorporating heating near the equator and compensatory cooling in the tropics. It is convenient, when looking for a model equilibrium state, to use a function whose net heating is zero.

On Jupiter the heating is presumably linked to the cloud distribution, since local heating generates ascending motion and therefore condensation in moist air. Clouds also affect the transfer of radiation, reflecting short-wave solar flux but absorbing and emitting long-wave thermal flux. Without a detailed cloud model we cannot determine the details of the vertical or meridional extent of a suitable function. We do not even know if we are justified in assuming that it should integrate to zero, because of the effects of the internal heat source and of transport. However, for the present study we shall assume that latent-heat release in the zones is balanced by radiative cooling in the belts. We may then choose a function whose net heating is zero, and this should not force any long-term trend in the model.

In the absence of compelling evidence for any particular distribution we conduct a sensitivity study, comparing the flows generated by a variety of heating functions. The first example is chosen to generate a flow in some ways resembling Jupiter's equatorial jet. The temperature anomaly function has the vertical form

$$\Delta T \propto \cos((l - 10) \times 2\pi/16) \tag{6.35}$$

where l is the model pressure full-level, and level 10 is at $p = 1.13$ bar. The heating is applied between levels $l = 6$ and 14, the function vanishing at these two extrema. This corresponds to a pressure range from 2.8 bars to 450 mb. In latitude it has the form $\Delta T \propto \cos(A\phi) + B$ where $A = 17.3$ and $B = 0.214$. These values give a zero integral over the channel width and zero amplitude at the channel walls. The heating is uniform in longitude.

This scheme relaxes temperatures towards the anomaly distribution, following Newton's Law as usual. A good estimate of a suitable coefficient would be 2.7×10^{-9}, the reciprocal of the orbital period (see Section 3.3.1 and Conrath et al. (1990)). However, we multiply the relaxation coefficient by 400 so that the flow equilibrates after about 50 days. This is necessary for reasons of computational economy.

This approach, i.e. regarding the clouds as modifying the equilibrium temperature profile, quickly leads to an equilibrium state. The alternative (continuous constant heating) does not readily converge to equilibrium, even after hundreds of model days. This is probably

because there is little meridional or vertical heat flux in the model except in the lower troposphere. Therefore the extra heat cannot easily be dissipated and the jets continue to accelerate almost linearly. The former approach, as well as working better, is more physically justifiable than constant heating with no feedback from the model state.

Figure 6.47 shows the results of the first experiment. The heating, around the 1 bar level, spins up westerly jets centred between the maxima and minima of temperature. In the vertical the winds peak at around 450 mb (just above the heating), decaying both above and below this level. The circulation induces a subsidiary temperature anomaly above the main one, associated with the vertical shear of the jets. The timeseries at lower right shows that the jets, beginning at rest, reach their equilibrium speeds after about 60 days. After this there is little change except for decay of the weak easterlies at the edges.

At the equator, we see (surprisingly at first) westerly winds of up to 30 m/s, while the 120 m/s jets are quite sharply confined poleward of $\pm 4°$. The model has quite heavy horizontal diffusion, and it is probably due to this that the winds are small but nonzero at the equator. This supports the view that horizontal diffusion could drive equatorial superrotation, but suggests that the required diffusion would be very large.

The twin jets have approximately the correct amplitude, direction, and combined width, but on the equator itself the winds are only of around 20 m/s. Therefore, as expected from the linear models of Section 6.12, this cannot of itself explain the equatorial superrotation. One possibility is strong meridional diffusion (Williams and Robinson, 1973), which draws momentum from the tropical jets to the equator. However, with a nearly pure zonal flow and few observed eddies in the equatorial region (Vasavada et al., 1998), such a large eddy diffusivity seems unlikely. We continue to pursue the hypothesis that the Hadley cell merely reinforces and modifies an equatorial jet driven by wave momentum flux convergence.

6.14 Hadley and Kelvin-wave forcing superposed

In an effort to produce a jet of realistic Jovian properties but with the simplest possible mechanisms, we will add a travelling Kelvin wave forcing to the previous model. We know that the Kelvin mode will give a westerly jet centred on the equator; together with the symmetric forcing this is hoped to produce a flow resembling that observed.

The parameters of the Kelvin wave chosen for convenience. They are a period of three days and a wavenumber of $+3$, in a channel of length $60°$. By reference to the model described in Section 6.7.3, we would expect this to give roughly the correct phase speed. We know from the broad-band forcing experiments that this will be within the resonant part of the spectrum. Of course this forcing is a physical idealization, but we have seen

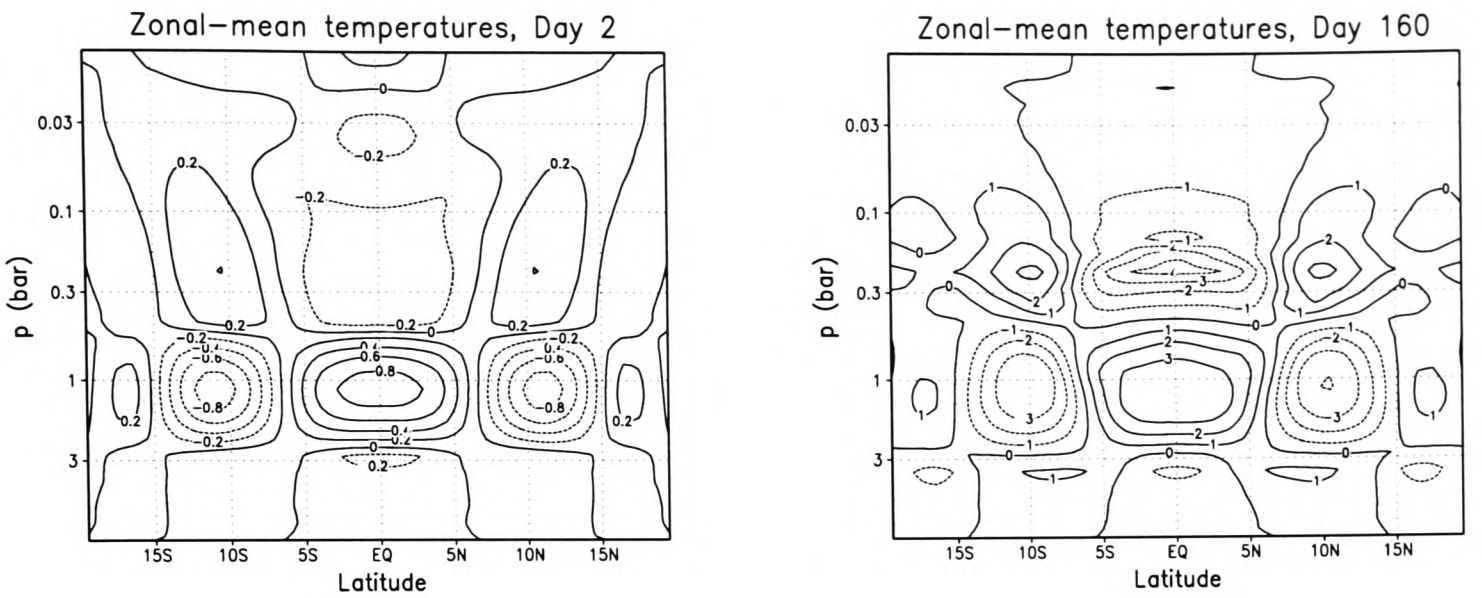


Figure 6.47: Results of a model Hadley cell. Top left: the zonally-symmetric heating function is illustrated by a cross-section of temperature anomaly early in the run. Top right: the equilibrium temperature distribution. The flow has induced anomalies of the opposite sign, above the forcing region. Lower left: the final zonal wind, showing a pair of westerly tropical jets. The equatorial winds are very light. Lower right: timeseries showing adjustment to equilibrium.

(Section 6.11.3) that a similar wave could be generated by a feature such as a stationary coherent source on the equator. The Hadley forcing is the same as in the model detailed above.

Two sets of results are presented, one with intermediate stratospheric drag and another with light drag. The models are alike in all other respects, including the sponge-layer drag. Figure 6.48 illustrates the flow in the case with intermediate drag.

The second case, with light drag, shows that the wave forcing generates strong westerlies in the stratosphere as well as in the upper troposphere. The upper jet appears very rapidly at the start of the integration, and remains stable after this. This suggests that perhaps the wave forcing amplitude could be reduced, which might mean the lower jet would take longer to form, but less energy would be available to penetrate through to the stratosphere. In fact the effect of reducing the amplitude is to *delay* the appearance of the jet until around Day 20. It still grows just as rapidly and reaches a very similar final state.

It is possible to reproduce the asymmetric nature of Jupiter's observed jet but adjusting the Hadley cell heating distribution. However, this is merely introducing more adjustable parameters and it is probable that eddy transport in the tropics, ignored here, has an effect. We therefore end the study at this point. Figure 6.50 shows an example profile of cloud-top zonal mean zonal wind, for comparison with the cloud-top observations (Limaye, 1986). As we have already noted, the easterly jet beneath the main westerly (centred at 4 bar) is refuted by the Galileo Probe observations. To remove this unrealistic feature, we would probably need to make the model channel deeper, and apply the forcing at a deeper level.

6.15 Conclusions

In this chapter we set out to explain the existence of a superrotating equatorial jet, by showing that a simple mechanism (dissipation of upward propagating Kelvin waves) can generate a plausibly realistic jet profile. This is not a new idea—in fact it dates back at least to Maxworthy (1975)—but we believe it has not previously been explored using a modern three-dimensional GCM. The theory has stood up well to this more stringent test, generating jets of realistic speed, but drawbacks have been demonstrated. One example is the single-peaked meridional profile of the model winds. We have extended the model, to show how solar heating at the cloud-tops could generate a small Hadley-cell like meridional circulation. This circulation, conserving angular momentum, can modify this profile to produce one like that observed. This has a double peak and a local minimum at the equator.

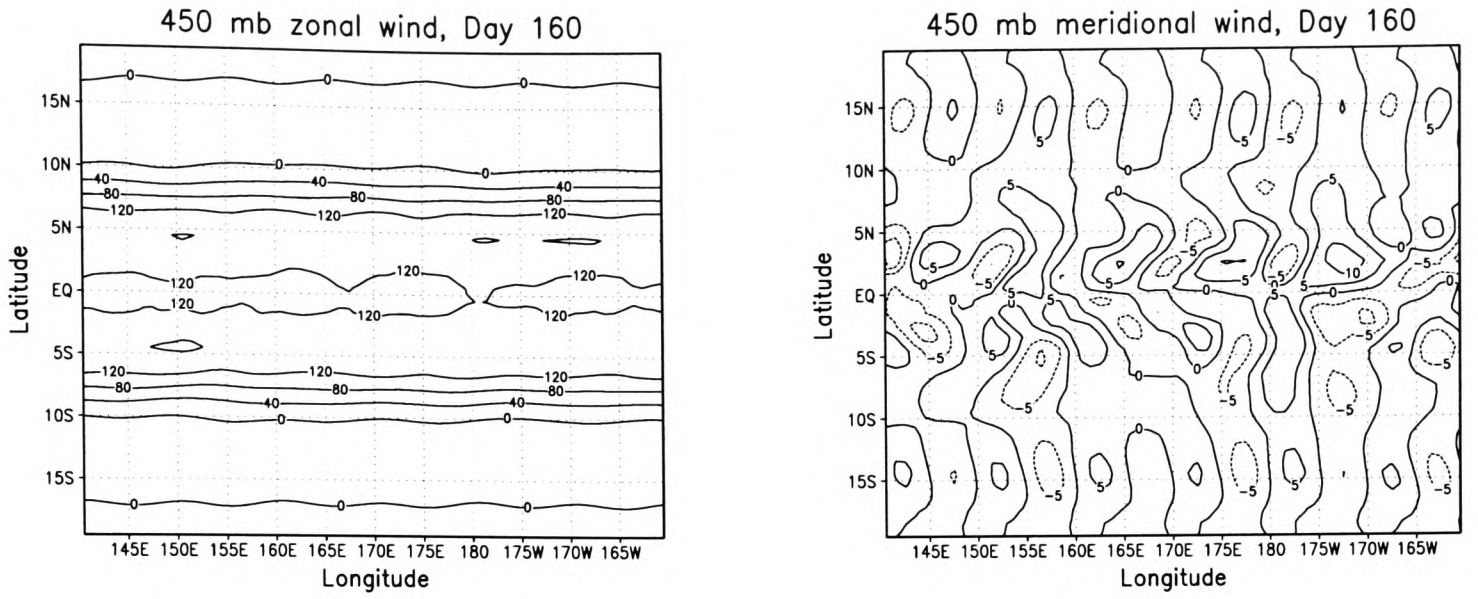


Figure 6.48: Results with superposed Hadley and wave forcing. Top panels: wind components at the cloud-tops. Lower left: cross-section of zonal wind. Lower right: timeseries showing development of cloud-top zonal mean zonal wind.

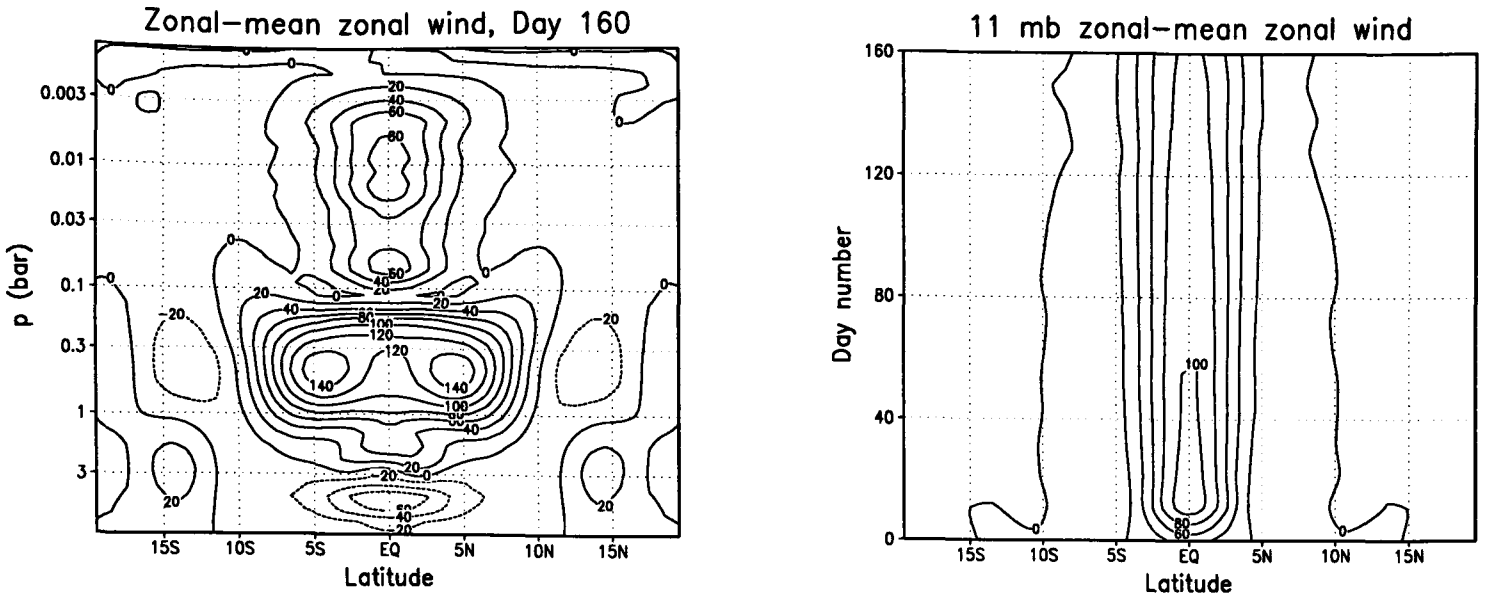


Figure 6.49: Results with superposed Hadley and wave forcing, and lighter drag in the stratosphere. Left: cross-section of zonal wind after 160 days. Right: timeseries showing zonal mean zonal wind in the stratosphere. Note the rapid acceleration towards a stable state.

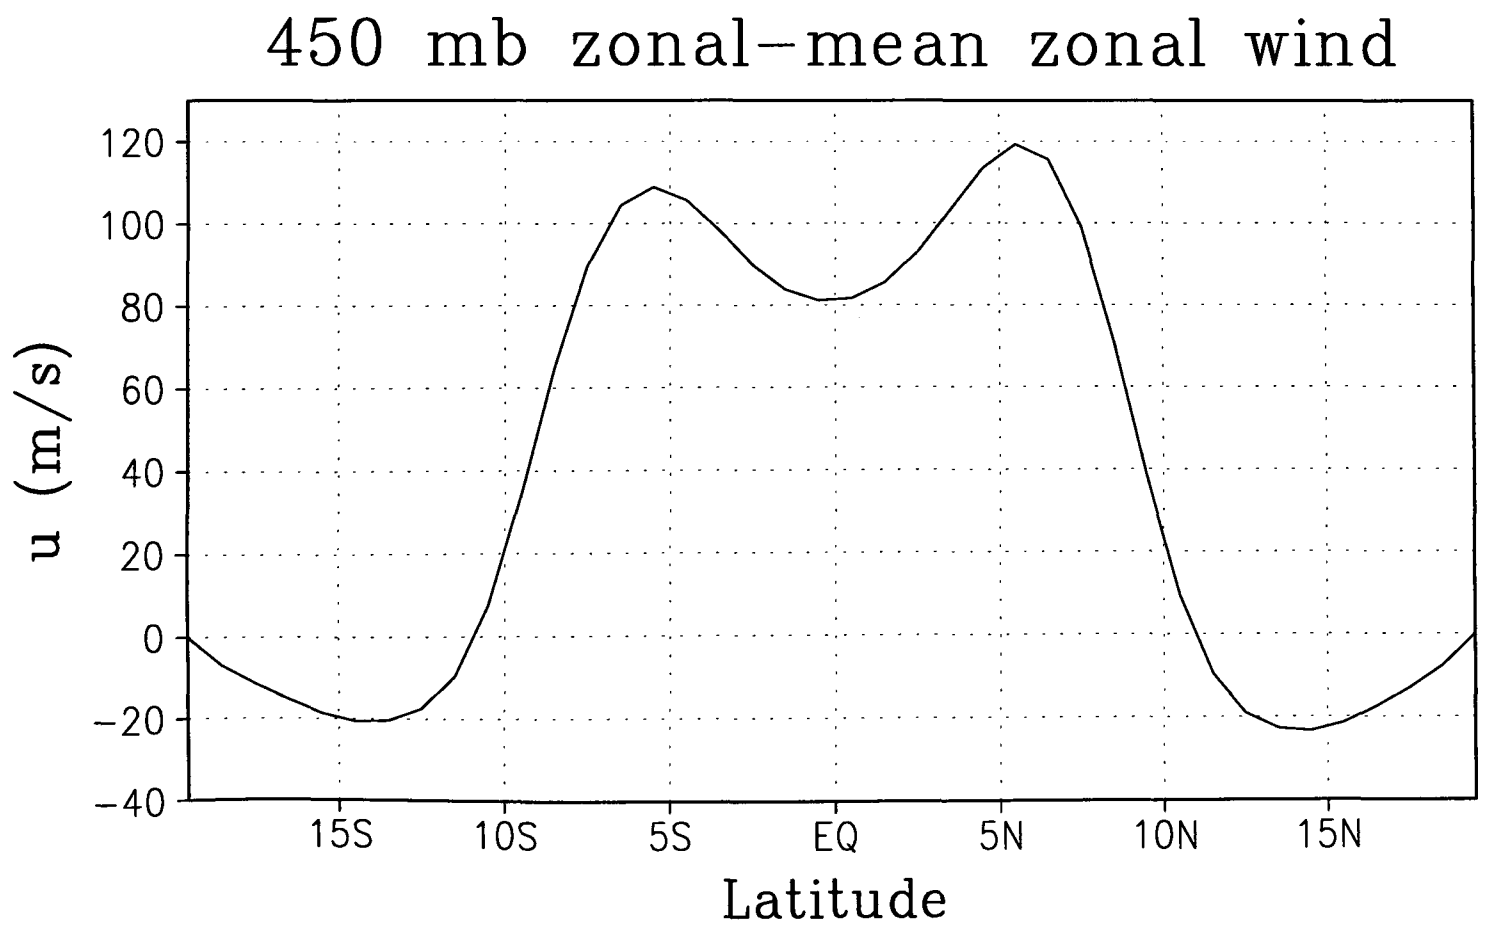


Figure 6.50: Cloud-top wind profile from an experiment with slightly asymmetric Hadley-cell forcing.

There are clear deficiencies in our model as it stands, one being the idealized nature of our wave sources. This limitation is perhaps defensible in terms of the lack of direct evidence of what lies below the cloud tops. A highly detailed model would make specific but unverifiable predictions. We restrict ourselves to arguing that if a wave source or sources are present, and they have some of the spectral and spatial properties of those in the model, then this could explain the direction, speed and profile of Jupiter's equatorial jet. This is perhaps a rather limited conclusion, and certainly a more detailed model could serve by showing that sources like ours are really possible in Jupiter's lower troposphere.

We have already mentioned the unrealistic easterly jets at the forcing level, below the main westerly. To address this, we propose to look at the effect of changing the depth of the sources, which could be done in a deeper version of the model. Another possibility is a moving source. We do not know if Jupiter's interior is in solid body rotation, and our assumption of a source stationary in System III is no more than a working hypothesis. It would be relatively straightforward to model a moving source.

Other more subtle failings have been glossed over. One of these is the question of dissipation of the waves. We have parametrized this as simple linear drag in the stratosphere, with an arbitrary amplitude and vertical profile of drag coefficient. This is a crude first approximation, with some scope for refinement. It may be that results from the study of equatorial waves in the Earth's stratosphere will be of use. Plumb and Bell (1982) discussed the question of dissipation in some detail. They found the use of thermal dissipation (Newtonian cooling) alone, with no mechanical drag, required very high model resolution. Therefore in most experiments they used mechanical drag for convenience, as we have done. They point out that the use of scale-dependent damping would be preferable, because it would be more effective where the vertical wavelength is short, i.e. in the dynamically important regions. It would also have the desirable side-effect of allowing coarser vertical resolution.

This issue of parametrization of sub-grid scale flow has wider implications. One of the difficulties of modelling Jupiter is that in many ways it resembles the Earth's oceans more closely than the atmosphere. A common feature is the importance of unresolvable, sub grid-scale eddies. The jovian situation is actually more clear-cut than the oceanic one, since even a superficial glance at an image of Jupiter shows very fine structure in filaments, eddies and waves. On the other hand, the small-scale flow in the ocean is quite difficult to measure. We note in passing that another feature common to Jupiter and the oceans is the paucity of vertical profile information, compared to the Earth's atmosphere.

In the Earth's atmosphere, such sub-grid scale features do not significantly affect the synoptic scale flow, although they are of importance in tracer transport, atmospheric chemistry, and mesoscale forecasting. Conversely, both in the oceans and on Jupiter, a significant fraction of the energy resides on short length scales which we cannot hope to resolve. There is much scope for improvement of schemes such as those dealing with

wave-drag, but it is not straightforward.

To support the claim that Kelvin-wave forcing drives the equatorial jet, we have studied the filtering properties of the atmosphere in response to a range of wavelengths and phase speeds. The results suggest that the phase speed of the forcing waves must lie reasonably close to 100 m/s or so, and also that forcing seems more efficient at small wavenumbers. The latter result agrees that of Holton (1972), whilst the former is a new result which offers an explanation for the speed of the observed equatorial jet. This is because the jet will be accelerated to equal the phase speed of the wave.

The model predicts that the jet speed may be determined more by the atmosphere's filtering properties than by the wave-sources themselves. Even if the sources produce a range of modes with varied phase speeds, only a relatively narrow band of phase speeds can reach the upper troposphere to excite a superrotating jet. We do not know the sensitivity of the jet velocity (100 m/s or so) to the vertical temperature profile. Limited tests have shown that adjusting the vertical distribution of the drag coefficient has little effect, so perhaps the filter band value depends on static stability, and hence on temperature. Other possibilities include model parameters such as horizontal diffusion. An interesting future experiment might be to repeat the runs with different (unrealistic) $T(p)$ profiles. This may demonstrate a connection between stratification and jet speed, which would confirm the idea that the details of the wave sources are of less importance than the atmospheric filter. Conversely if the filter depends more on numerical parameters of our particular model, this would cast doubt on our conclusions.

The latitudinal width of the jet is probably governed by the balance between horizontal advection of vorticity and the β-effect (Hide, 1969), which means that the width will be given approximately by the Rhines length-scale $(UR/2\Omega)^{\frac{1}{2}}$. As Hide points out, this does agree with observations. The Rhines length also appears in the expression for the width of the Kelvin wave, so these are connected. The present work allows us to take a tentative step further, with a theoretical reason for the observed speed U. Model experiments have shown that the wave and jet widths do not vary strongly with the width of the forcing. This agrees with the view that the jet width is determined by the Rhines length rather than by the forcing.

We have also looked at the effects of different coherent, stationary sources. It has been shown that a symmetric coherent source centred on the equator can generate the required modes, although antisymmetric sources (cf. the Earth's ITCZ) cannot. This is logical because the main eastward-propagating mode (the Kelvin wave) is symmetric about the equator. Therefore any mechanism to transport eastward (westerly) zonal momentum should feature symmetric sources. Antisymmetric sources tend to generate mixed Rossby-gravity waves, which propagate westward. These probably have a part to play in driving the quasi-quadrennial oscillation (Li and Read, 1999), so we expect such sources to be present.

The effect of the vertical scale of the sources has not been assessed, partly because the model channel is relatively shallow. It is possible that future observations of stratospheric waves could be used to validate these model results and infer the properties of the deep sources. The simple view is that the vertical wavelength of waves in the stratosphere is equal to the thickness of the source, i.e. the height of the convective cells.

Chapter 7

Conclusions and future work

7.1 Conclusions

O mihi praeteritos referat si Iuppiter annos. (Virgil: *Aeneid* 8:560)

The aim of this modelling study has been to investigate (and attempt to explain) some features of the observed flow above Jupiter's clouds, in terms of possible forcing mechanisms, and of the effect of various plausible flows in the unobserved region below the clouds. This has been made possible by the technical advance of a general circulation model, with realistic stratification and 30–40 vertical levels. Unlike most previous models, the GCM was run in a limited area or channel mode, to allow high resolution with modest computational resources. Being more realistic than previous models, ours allows direct comparison with observations as a function of height.

The model has several inherent limitations and drawbacks, evident from the outset. A critical review of these is perhaps in order, before we summarize the main results. Then, in the future work section (7.2) we will outline plans for developing the model to overcome some of these limitations. We shall cover the issue of radiative transfer in more detail (Section 7.2.3) than the other aspects, because we have already begun to develop a radiation scheme for the model. This part could be considered as work in progress, rather than future plans.

7.1.1 Limitations of our model

The first restriction stems from our use of a limited area configuration. This is the necessity for impermeable channel sidewalls, which potentially have a noticeable (and unrealistic) effect on the flow. On the other hand, flow on Jupiter is mostly zonal, and we

are most interested in features which do not move meridionally—for example, giant eddies confined in latitude. Limited-area modelling is unlikely to be replaced completely by global models, since any available increases in computational power can easily (and profitably) be absorbed by increasing the resolution and complexity of the model. Expanding the channel size is perfectly possible however.

The current simplified approach to radiation is a limitation which could readily be remedied. A linear-relaxation radiation scheme could perhaps be justified on grounds of efficiency, the lack of observational data, and by arguing that a detailed transfer scheme might not change the model results very significantly. Definitely desirable, though, is a proper calibration of the Newtonian scheme against a more sophisticated one, perhaps incorporating new observational data more explicitly.

Perhaps a more serious omission is that of realistic thermodynamics. We have included dry convective adjustment in the model, because without it there is a risk of statically unstable regions generating spurious large-amplitude waves. Such waves were observed in the early development of the model, and were thought to be due to the combination of an unsuitable sponge layer, and lack of convective adjustment.

Dry convection is probably a rather crude approximation to jovian conditions. To improve on it we suggest a scheme incorporating a parametrization of moist convection (preferably with several condensibles, as on the real planet), and also ortho-para hydrogen conversion. We will give details and references in Section 7.2.

Let us now move on to look at the results. Two particular cases have been studied in detail: midlatitude giant eddies and their interaction with jetstreams, and equatorial waves and their role in forcing the equatorial jet. The GCM studies are complemented by two simpler models. One uses a potential vorticity inversion technique, to solve for the balanced time-independent flow in a circularly symmetric eddy. This model shows that both deep and shallow balanced eddies can exist, with almost indistinguishable structures above the cloud-tops.

On the other hand, the validity of the PV model is rather limited, since it takes no account of the environment of the vortex. On Jupiter this would normally be a shear zone, breaking the axial symmetry. Furthermore it solves for balanced states, but does not guarantee that such solutions will be stable. It would be quite possible to incorporate time dependence into the code, which would test stability. However, we would not wish to invest disproportionate effort into pushing the model too far, given its other intrinsic limits. For more detailed, time dependent experiments we use the Unified Model.

We have also developed programs which perform a one-dimensional analysis of the vertical structure of the flows, in terms of a superposition of a few normal modes of the atmosphere. This has proved illuminating in allowing interpretation of the flow structures

from a Gaussian shape (as generated purely by a wave) and with suitable tuning, a jet closely resembling Jupiter can be achieved.

7.1.4 Implications for Jupiter

We wil end by attempting to summarize what has been learned about Jupiter itself as a result of this work. Model jets can spontaneously generate eddies extending barotropically down to the base of the model. This may suggest that the Great Red Spot need not be a shallow vortex confined to a thin weather layer, but could extend far below the clouds. The GCM can produce both shallow and deep vortices in the same environment, and perhaps this is true of Jupiter too. Observations only above the clouds may not be able to discriminate between shallow and deep flows. Could it be that the differences between the GRS (with its quiescent interior) and White Ovals (without one) are due to different vertical structures? Of course, 10 bars is not a very deep level in the jovian context, but our PV inversion results imply that our conclusion may even be valid down to 100 bars.

We may have learned something about the jets too. Although we have not tried to explain the midlatitude jets, we have confirmed that the equatorial jet may be explicable in terms of wave forcing. The jet shape is not particularly sensitive to the details of the wave source, which is fortunate since we know nothing directly about any real wave sources on Jupiter. We predict that the waves are filtered by phase speed, and that this filtering largely determines the jet speed. Future observations of stratospheric equatorial waves may help to constrain our model.

Finally, we are in a strong position to refine these conclusions by means of future improvements to the model. These improvements will include enlarging the domain, and including much more realistic physics parametrizations. We will outline the possibilities in the next section.

7.2 Future work

A good plot, good friends, and full of expectation. (Shakespeare: *Henry IV Part 1*)

The major part of this project has been the design and development of a new model of Jupiter's atmosphere, starting from an existing general circulation model of the Earth. We focused initially on the less-than-glamorous work of tuning the model resolution, diffusion parameters, etc. Once the model was working consistently, we were able to obtain some interesting results, even with our rather simple physics parametrizations. Along

generated in the GCM. Spontaneous instability of forced jets in the GCM sets up eddies whose vertical structure consists of only one or two normal modes. The major failing of the normal-mode calculations is the assumption that the ambient atmosphere is at rest. This is patently not so in most of our GCM simulations, although he modes are directly applicable to the single-vortex experiments in the first half of Chapter 5. We assume that they are also reasonable approximations to the modes of the basic zonal flow in the later experiments. In any case, the two eddy structures we generate (shallow and deep) are so dramatically different that the details of the normal modes are not vitally important.

7.1.2 Midlatitude vortices and jets

In Chapter 5 we presented a systematic study of the properties of midlatitude vortices on Jupiter. This followed on from the time-independent models of Chapter 2, which suggested that deep and shallow vortices might behave in similar ways when viewed from above the cloud tops. We asked whether the atmosphere can generate and maintain both shallow and deep eddies, whether these two types have different migration or stability properties, and whether they could be distinguished with available observations.

As the jovian environment is very complicated, we began with the simple case of a Gaussian geostrophic vortex initialised in an atmosphere at rest. Even this example had more complex behaviour than expected. Most vortices showed a tendency to shrink and intensify, steepening the PV gradient at the edges. This steepening seems to be an ageostrophic process similar to frontogenesis. All anticyclones migrate westward and equatorward, the migration rate increasing with vortex intensity (measured by peak wind speed or amplitude of the PV anomaly).

The steepening process was opposed by numerical diffusion, and the adoption of fourth-order hyperdiffusion, and a shorter timestep than originally adopted, proved capable of maintaining an isolated vortex at its original intensity. Hyperdiffusion is justifiable on the grounds that our gridpoint spacing of half a degree fails to resolve many fine-scale motions which are visible in images. These sub-grid scale processes are represented by scale selective diffusive mixing in the model. Because the diffusion has such a high order, it is very effective on these small scales but insignificant on the properly-resolved synoptic scale.

We found that the normal modes predicted by the one-dimensional model of Chapter 4 turned out to match the GCM results rather well. The model vortices often break up, especially under weak diffusion. When this happens, new vortices are formed and these invariably have a vertical structure closely resembling the first baroclinic normal mode. The model seems to cascade energy into the gravest mode compatible with the boundary conditions (drag at the top and bottom). Vortices initialized with arbitrary vertical

structures tend to adjust until they look like a normal mode.

Having gained some understanding of vortices in a resting atmosphere, we moved on to the more realistic case of vortices embedded in zonal jet streams. These zonal jets are driven by direct mechanical forcing at the bottom of the model, and weak thermal relaxation throughout the domain towards geostrophically-balanced temperatures. This arrangement works well for midlatitude jets. The vertical structure is barotropic below about 300 mb, then a linear decrease to zero at the top of the model.

Vortices initialized in the shear zone between the jets behaved as expected: they became elongated into ellipses, and their equatorward drift was initially prevented, by the PV barrier associated with the jet. This was true for both shallow baroclinic eddies (similar to those studied in the rest case) and for deep eddies whose vertical structure was the same as that of the jets. Eddies grow by baroclinic instability until they become intense enough to overcome the PV barrier, at which point they migrate rapidly, and are normally destroyed on reaching the model sidewall.

There is considerable natural variability in the model, even when run with a simple case of a single strong easterly jet. This is baroclinically unstable, and when an instability is seeded (deliberately or by model noise), we observe a range of waves growing and rolling up into vortices. These are seen with both the shallow, confined first-baroclinic structure as before, but also a deep structure similar to that of the jets.

We gain further insight by means of energy diagnostics. These are based on QG theory and have other limitations, so cannot be taken as quantitatively precise. Nevertheless, the study of model flows together with timeseries of energy quantites and conversion rates has shown that most eddies are created by mixed barotropic-baroclinic instability. Their subsequent destruction is a baroclinic process, as they are sheared apart on reaching the northern sidewall.

Thus we have seen the relatively realistic GCM can spontaneously generate vortices which are very similar to the idealised cases investigated with the PV inversion model. It is very difficult to establish the vertical structure indirectly by looking at one level only, because deep vortices form, migrate and disperse in much the same ways as do shallow ones. This implies (unfortunately) that existing and forseeable observations are insufficient to answer the question of how deep Jupiter's weather extends.

Another question raised by the results is that winds on Jupiter seem to peak at or below the cloud tops: thermal winds above the clouds show a decrease with height. In contrast to this, the shallow model vortices are trapped rather higher, in the lower stratosphere. Of course the inclusion of realistic cloud physics and radiation may alter this result, but there may be scope for use of the 1D normal mode model to see whether it is possible to construct an atmosphere in which the gravest baroclinic mode peaks at a lower level

than in this one. It may be that the present GCM, with a slightly different reference atmospheric profile, could produce circulations peaking rather lower in the atmosphere.

Did this part of the research succeed? It must be said that the point reached to date is rather immature, in that the answer to most of the questions is "maybe". We have seen that the model can often simulate several of the alternative theoretical mechanisms proposed to explain observations. The aim of science is to refute hypotheses, and in this we have often been unsuccessful. In answering one question we usually raise another, or show that several alternative answers are possible.

However, there are two grounds for cautious optimism. One is that Jupiter models are still very simple, by comparison with those of the Earth and Mars. It would be very satisfactory if a particular feature—e.g. very deep vortices—were present in the simple models, but not in more realistic future models. We would then have precise physical grounds for rejecting the hypothesis. The second reason not to despair is that Jupiter's atmosphere may well be much more complicated than we think, with a wide range of different processes and features all present together. In other words, our questions are not specific enough.

It is not difficult to refine the scope of the problem. For example, instead of asking "what possible vertical structures can vortices adopt?" we could ask "what is the structure of the Great Red Spot?". The difficulty then is that by moving from the generic to the specific, we suffer seriously from the lack of observational data. The greatest deficiency is in our knowledge of winds and temperatures below the cloud-tops. In Chapter 2 we saw that our simple PV-inversion model, making use of published data for the flow at and above the clouds, cannot convincingly discriminate between a deep and shallow case for the GRS and White Ovals. There seems to be little option but to attack the general case from a simplified viewpoint and to hope that increasingly complicated models will be able to explain more phenomena in the future. The Galileo mission has taken a vast amount of new data, and we can expect that this will be analysed and published for many years to come. Unfortunately, there will be no more data on the flow beneath the clouds, at least for many years.

Apart from the proposed developments listed in Section 7.2, there are many possibilities for further work with the model as it stands. An obvious next step is to study the interaction between two or more vortices. We have performed experiments in which two vortices are initialized at the same latitude. Animations show them approaching each other and being mutually repelled. We might expect large eddies to absorb smaller ones, as is observed. Detailed studies of this have not been done, because it was thought better to draw a line at one vortex plus jets, rather than include a superficial study of the more general case. The model is ready for this work to be done in the future, if required.

7.1.3 Waves and the equatorial jet

Chapter 6 contains a detailed study of equatorial waves in the GCM. The emphasis here was on trying to pin down the possible nature of wave sources in the deep atmosphere, by means of modelling a variety of sources and then comparing the resulting waves with observations. After reviewing the theory of linear waves, we described the tuning of model parameters necessary to achieve a suitable model configuration. The considerable complexity of the Unified Model (despite its great flexibility) meant that this was quite a lengthy task. We showed how the resulting model correctly supports pure travelling Kelvin and mixed Rossby-gravity (MRG) waves in an atmosphere at rest. Furthermore, the numerical dissipation in the model is sufficient for the wave forcing to spin-up a Jupiter-style equatorial jet. For many runs, weak linear drag was added in the stratosphere. This is convenient because it shortens the spin-up timescale, but we showed that the drag coefficient is not critical to the long-term development of the flow.

Next we took advantage of the high vertical resolution and realistic stratification in the model, using it to identify the filtering properties of a resting jovian model atmosphere. We found that low wavenumbers propagate more easily than higher ones. More interesting was the result that the the model atmosphere filters Kelvin waves primarily by phase speed. The pass band lies near to a zonal velocity of 100 m/s, which agrees rather well with that observed at the centre of the equatoral jet. The preferred speed seems to be determined by some static property of the atmosphere, which is initially at rest. We speculate that it is the static stability which is important, but have not yet tried to verify this by repeating the experiments with different N^2 profiles.

This result suggests that, so long as wave sources are present and can generate a wide spectrum of waves, the atmosphere will select the observed phase speed. We therefore do not need to constrain the sources to emit waves near this speed: instead we would expect them to generate a broad band, most of which is not transmitted through the troposphere.

We looked at the effects of several different coherent, stationary sources, following the work of e.g. Holton (1972) on organized convection in the Earth's atmosphere. We found that symmetric sources generate Kelvin waves and hence prograde jets, while antisymmetric sources generate mixed Rossby-gravity waves (as expected) and retrograde jets. Reasonable values of the amplitude, frequency and spatial size of the sources led to superrotating jets of realistic speed.

The final part of Chapter 6 was concerned with zonally symmetric heating and a Hadley circulation. This was motivated by the wish to explain the latitudinal structure of the observed jet, with its subsidiary minimum of u at the equator. Model results showed that weak "solar heating", balanced by enhanced cooling and localized near the cloud-tops is able to induce a suitable meridional circulation. This circulation modifies the jet structure

the way, we have seen which aspects of the model are likely to benefit most from improvements. Aspects of our proposed programme have been granted financial support by PPARC, and a postdoctoral researcher appointed. The outline given here is based on our research proposal. It includes estimates of the likely difficulty of each item, and mentions specific problems of which we are already aware.

The proposals are grouped into three categories: extension to deeper levels, inclusion of detailed thermodynamics, and explicit modelling of radiative transfer. We shall briefly describe plans for the first two, and then cover the third in more detail. This is because work has already begun on the implementation of a simple radiation scheme. It is included in the "future work" section because the scheme has not yet been fully tested, nor coupled to the dynamics model.

7.2.1 Depth of model domain

Perhaps the most immediate need is for the downward extension of the model domain, to a new lower boundary at 100 bars or even deeper. This would have many advantages. We could look at the effect of a possible stable radiative layer (Guillot et al., 1994b). We could place equatorial wave-sources at deeper levels, and/or increase their thickness if desired. This might allow us to simulate an equatorial westerly extending down to at least 20 bars, to determine whether or not our model is compatible with the Galileo Probe observations. The underlying easterly jet seems to be the least realistic feature of our present model flows.

A feasibility study has shown that, although the implementation is trivial, jets in the deeper model tend inexplicably to tilt in the meridional plane. Limited attempts to solve this problem, by means of vertical diffusion, direct forcing, or very strong heating, have so far failed. Further investigation is required to establish the reason for this tilting.

Another important possibility is to repeat the wave-selection experiments of Chapter 6 but in a deeper domain and with a different static stability profile. This would test our conclusion that the filtered phase speed (and hence the equatorial jet speed) depend mostly upon the stratification in the lower troposphere. If a range of unrealistic stablity profiles gave unrealistic model jet speeds, this would confirm our result. On the other hand, it may be perfectly possible to generate model jets of 100–150 m/s almost independent of the static stability. The latter result would reopen the question.

7.2.2 Thermodynamics: moist convection and hydrogen conversion

Molecular hydrogen, which is by far the most abundant constituent of Jupiter's atmosphere, exists in two forms, known as *ortho* and *para*. These refer to the relative orientations of the nuclear spin of the two protons in the molecule. The significance here is that the conversion between the two forms is slow, temperature-dependent, and involves a change in c_p and release of latent heat. Thus hydrogen conversion may affect the thermodynamics of convection. We propose to adapt the scheme of Smith and Gierasch (1995) for use in the Unified Model. The present dry convective adjustment scheme was designed and coded with this in mind, so generalizing it should not be too difficult.

If we are not merely to assume some idealized distribution of para-fraction hydrogen, we would need to treat this quantity as an active tracer, advected by the flow and mixed by convection. The UM has powerful capabilities for handling tracer advection, let down only by difficulties with initialization. See Appendix A for technical information. We plan also to treat water as a tracer, using a cumulus-convection scheme along the lines of Del Genio and McGrattan (1990). We would be able to include ammonia and perhaps ammonium hydrosulphide using the same tracer methods.

7.2.3 Radiative transfer

Radiative equilibrium models of Jupiter can predict a realistic vertical temperature profile, even with only a simple parametrization of solar heating and IR cooling. An early example (Cess and Khetan, 1973) considered just the strongest methane bands (absorption at 3.3 μm and emission at 7.7 μm), along with hydrogen emission. Hydrogen also has a vibration-rotation band at 2.5 μm which is a potentially significant source of solar absorption. However, Cess and Khetan showed that the only effect of including this band was to raise the equilibrium temperature of Jupiter by 0.6 K; modern schemes neglect solar absorption by hydrogen.

Wallace et al. (1974) developed a treatment of the weakly absorbing methane bands at $\lambda <$ 1.7 μm. These are more significant than the strong bands at higher pressures, because the strong bands are saturated. Heating by aerosols was also included. Aerosol cooling was ignored, as the intention was to find an upper limit on the effects of aerosols.

Cess and Chen (1975) combined these two approaches, and added the effects of two other strong methane bands, at 1.7 and 2.3 μm. They also considered the influence of ethane and acetylene. There are important emission bands, of ethane at 12.2 μm and of acetylene at 13.7 μm. These bands are in a much stronger region of the infrared Planck function than is the CH_4 7.7 μm band, which compensates for the low abundances of C_2H_6 and C_2H_2. Cess and Chen showed that absorption by C_2H_6 and C_2H_2 are a minor source of

heating; the two strongest bands contribute only about 5% of the solar heating, because the abundances are so small. Most subsequent models neglect these and the weaker bands.

Time-dependent models followed, incorporating the physics established by the equilibrium models. The model of Cess and Caldwell (1979) predicted zonal-mean temperatures for Saturn above 100 mb. Cess and Caldwell assumed the infrared heat flux from the interior to be independent of latitude. This approximation was refined by Bézard et al. (1984), who assumed the emission temperature of the interior to be linear in sin(latitude). This later model also covered the upper troposphere, down to 600 mb. Hydrocarbon emission was neglected, as being small at the levels of interest ($p > 100$ mb). Hydrogen cooling was handled with the direct cooling-to-space approximation (Harshvardhan and Cess, 1976).

The methods of Bézard et al. and of Wallace et al. were applied to radiative-dynamical modelling by Conrath et al. (1990). They assumed the lower boundary temperature to be latitudinally independent, and found the solar heating to be only weakly dependent on latitude. This surprising result was explained by the fact that the model heating was dominated by the weak methane bands, which are optically thin at most levels. The dynamical model used was the linearized zonal-mean scheme of Leovy (1964). The effect of aerosols was neglected, on the basis that latitudinal temperature variations were of more interest than the absolute values. The dependence of aerosol abundance on latitude is not well known; it was assumed to be weak. Hydrogen cooling was handled by an exchange formulation with 64 levels, which is more realistic than the assumption of direct cooling to space.

7.2.4 The proposed radiation scheme

The ideas described above could be applied to parametrize the radiative forcing of a GCM. One possible scheme, following Bézard et al. (1984) and Conrath et al. (1990), is set out below. We have implemented this in a radiative-convective column model. Preliminary results are discussed in Section 7.2.5. These are reasonably realistic at levels above the clouds, but not near the lower boundary.

Net flux deposition and heating rate

The heating rate at height z is given by

$$\frac{\partial T}{\partial t} = \frac{-1}{\rho c_p} \frac{\partial F}{\partial z} \tag{7.1}$$

where T is the temperature, ρ the local density, F the net upward flux, and c_p the heat capacity per unit mass. For a hydrostatic atmosphere, $dp/dz = -\rho g$, so

$$\frac{\partial T}{\partial t} = \frac{g}{c_p}\frac{\partial F}{\partial p} \tag{7.2}$$

The flux may be split into contributions from hydrocarbon opacities, from aerosol heating, and from hydrogen cooling:

$$\frac{dF}{dp} = \Sigma K_i + K_a - K_{H_2} \tag{7.3}$$

Heating by strong methane bands

Conrath et al. used a Curtis-Godson approximation in their expression for the instantaneous flux deposition K_i. For the case of Jupiter, however, they assumed constant hydrocarbon mixing ratios. Under these circumstances the expression simplifies to that of Bézard et al., namely

$$K_i = \pi\alpha B_{\sigma_i}(T_s)\frac{dA_i}{dp}\cos\theta \tag{7.4}$$

Here α is the planetary angle factor, defined as $\alpha = (R_s/L)^2$ where R_s is the Sun's radius and L is the (seasonally varying) orbital radius. Values are $\alpha = 0.81 \times 10^{-6}$ for Jupiter, and 0.225×10^{-6} for Saturn. $B_{\sigma_i}(T_s)$ is the Planck function at wavenumber σ_i, for the effective temperature of the Sun (which is close to 5785°K in this spectral region). θ is the solar elevation angle. Combining this with the expression for the total band absorbance A_i yields

$$K_i = 2\pi\alpha B_{\sigma_i}(T_s)\frac{A_i^\circ}{p}\left(\frac{\cos\theta}{1+(A_i^\circ d_i \cos\theta/2\gamma_i p S_i N)^{1/2}}\right) \tag{7.5}$$

N is the column abundance of the absorbing gas, measured in cm-amagat. (An abundance of 1 cm-am is the amount of gas such that a column of 1 cm^2 cross-section at stp has a height of 1 cm. Thus 1 cm-am $\simeq 4.46 \times 10^{-5}$ moles.) The quantities S_i, A_i, γ_i, and d_i are band parameters: respectively, the band intensity, bandwidth parameter, half-width per unit pressure, and mean line spacing. Their values are taken from Tables I and II of Cess and Chen (1975).

Heating by weak methane bands

Wallace et al. (1974) treated the bands shortward of 1.5 μm by Goody's exact exp $(-S/\sigma)$

line intensity distribution function (Goody and Yung, 1989). Doppler broadening was ignored; the exponential distribution function is integrable for a Lorentz profile, and for a constant mixing ratio the opacity reduces to $C_i N$, where N is the column abundance and C_i a constant for each band i. Wallace et al. calculated values of C for 18 bands between $0.74\ \mu\text{m} < \lambda < 1.50\ \mu\text{m}$.

Conrath et al. (1990) simplified this approach, averaging the 18 parameters C_i into just two, corresponding to the wavelength ranges $0.74\ \mu\text{m} < \lambda < 1.02\ \mu\text{m}$ and $1.06\ \mu\text{m} < \lambda < 1.50\ \mu\text{m}$. The solar flux was also averaged over each spectral interval. Weak lines between 3.3 and 1.5 μm were neglected, although they were included by Wallace et al. The expression for weak band absorption is

$$\frac{dF}{dp} = \pi\alpha(B_a \Delta\nu_a C_a e^{-C_a N/\mu} + B_b \Delta\nu_b C_b e^{-C_b N/\mu}) \tag{7.6}$$

where $\Delta\nu$ is the spectral interval in cm^{-1} and μ is the cosine of the solar elevation angle. The two band constants have the values $C_a = 1.36 \times 10^{-4}$ and $C_b = 4.8 \times 10^{-5}$ (cm-am)$^{-1}$.

Cooling by methane, ethane and acetylene

Conrath et al. assumed this contribution cools directly to space. With constant mixing ratios, their expression becomes

$$\frac{dF}{dp} = -2\pi\alpha B_i \frac{A_i}{p}\left(\frac{\mu}{1 + (A_i^\circ d_i \mu / 2\gamma_i p S_i N)^{1/2}}\right) \tag{7.7}$$

Once again the band parameters are taken from Cess and Chen (1975).

Cooling by hydrogen and helium

Cess and Khetan (1973) considered the effects of H_2 and He opacities. In the far infrared, hydrogen opacity is caused by rotational-translational and pure translational transitions. At stratospheric pressures, the opacity is induced by binary collisions, so the absorption coefficient varies with the pressure squared. Helium increases the pressure induced hydrogen opacity, although it has no opacity itself. In the direct cooling-to-space approximation (Harshvardhan and Cess, 1976), the flux emitted to space by the slab between pressures p and $p + dp$, is $dF = A\sigma T^4 dp$. The emissivity A is given by

$$A = \frac{2\pi}{\sigma T^4}\int_0^\infty B_\sigma(T)\frac{d\tau_\sigma}{dp}E_2(\tau_\sigma)d\sigma \tag{7.8}$$

where τ_σ is the optical depth at wavenumber σ. E_2 is the second exponential integral (Goody and Yung, 1989):

$$E_2(x) = \int_1^\infty e^{-\omega x}/\omega^2 d\omega \tag{7.9}$$

An emissivity model is required to find τ_σ and evaluate this expression. The simple analytic model of Cess and Khetan (1973) was rejected by Bézard et al. as inappropriate for the troposphere; newer laboratory results were used instead. Conrath et al. used an exchange formulation, although they have published no details. It appears that emissivity calculations may be the most difficult part of the scheme to implement efficiently.

7.2.5 Preliminary results of the radiative-transfer model

We implemented this scheme in a one-dimensional column model. In the hydrogen cooling section we currently assume cooling-to-space, rather than the more rigorous exchange formulation we would eventually adopt. There is no time dependence, the solar zenith angle being fixed. The lower boundary condition is forced toward a fixed temperature of 300 K, representing Jupiter's interior. Unfortunately the convective adjustment scheme does not keep the lower boundary temperature fixed, so the boundary is always slightly colder than intended. Our usual model atmosphere has $T = 273$ K at 10 bars, and the radiation model does come close to this. This boundary condition is unsatisfactory and we intend to modify it, preferring to specify the upwelling thermal flux rather than the temperature. The convection scheme would probably not then need to be changed, since the lower boundary temperature would not be fixed.

The initial condition is an isothermal atmosphere at 165 K. The model relaxes from this toward a radiative-convective equilibrium profile. Figure 7.1 shows the different contributions to the heating and cooling rates: strong- and weak-band methane absorption, methane cooling and hydrogen cooling. The net flux deposition (which is proportional to the heating rate) is shown in the lower right panel.

The equilibrium profile (Fig. 7.2) has some realistic features but also some unrealistic ones. Above 1 bar, the profile is rather close to observations: the temperature at 1 bar is around 170 K, falling to 110 K at the tropopause, which is correctly situated at about 100 mb. The stratospheric temperatures are reasonable too, at least up to the 1 mb level, which is the lid level of our GCM.

Below the 1 bar level the model profile is a little less satisfactory. It appears that, between 10 bars and about 5 bars, the lapse rate is adiabatic, i.e. governed by convection. This is reasonable enough, but there is a small turning point at around 4 bars, where the lapse rate becomes governed by radiation rather than convection.

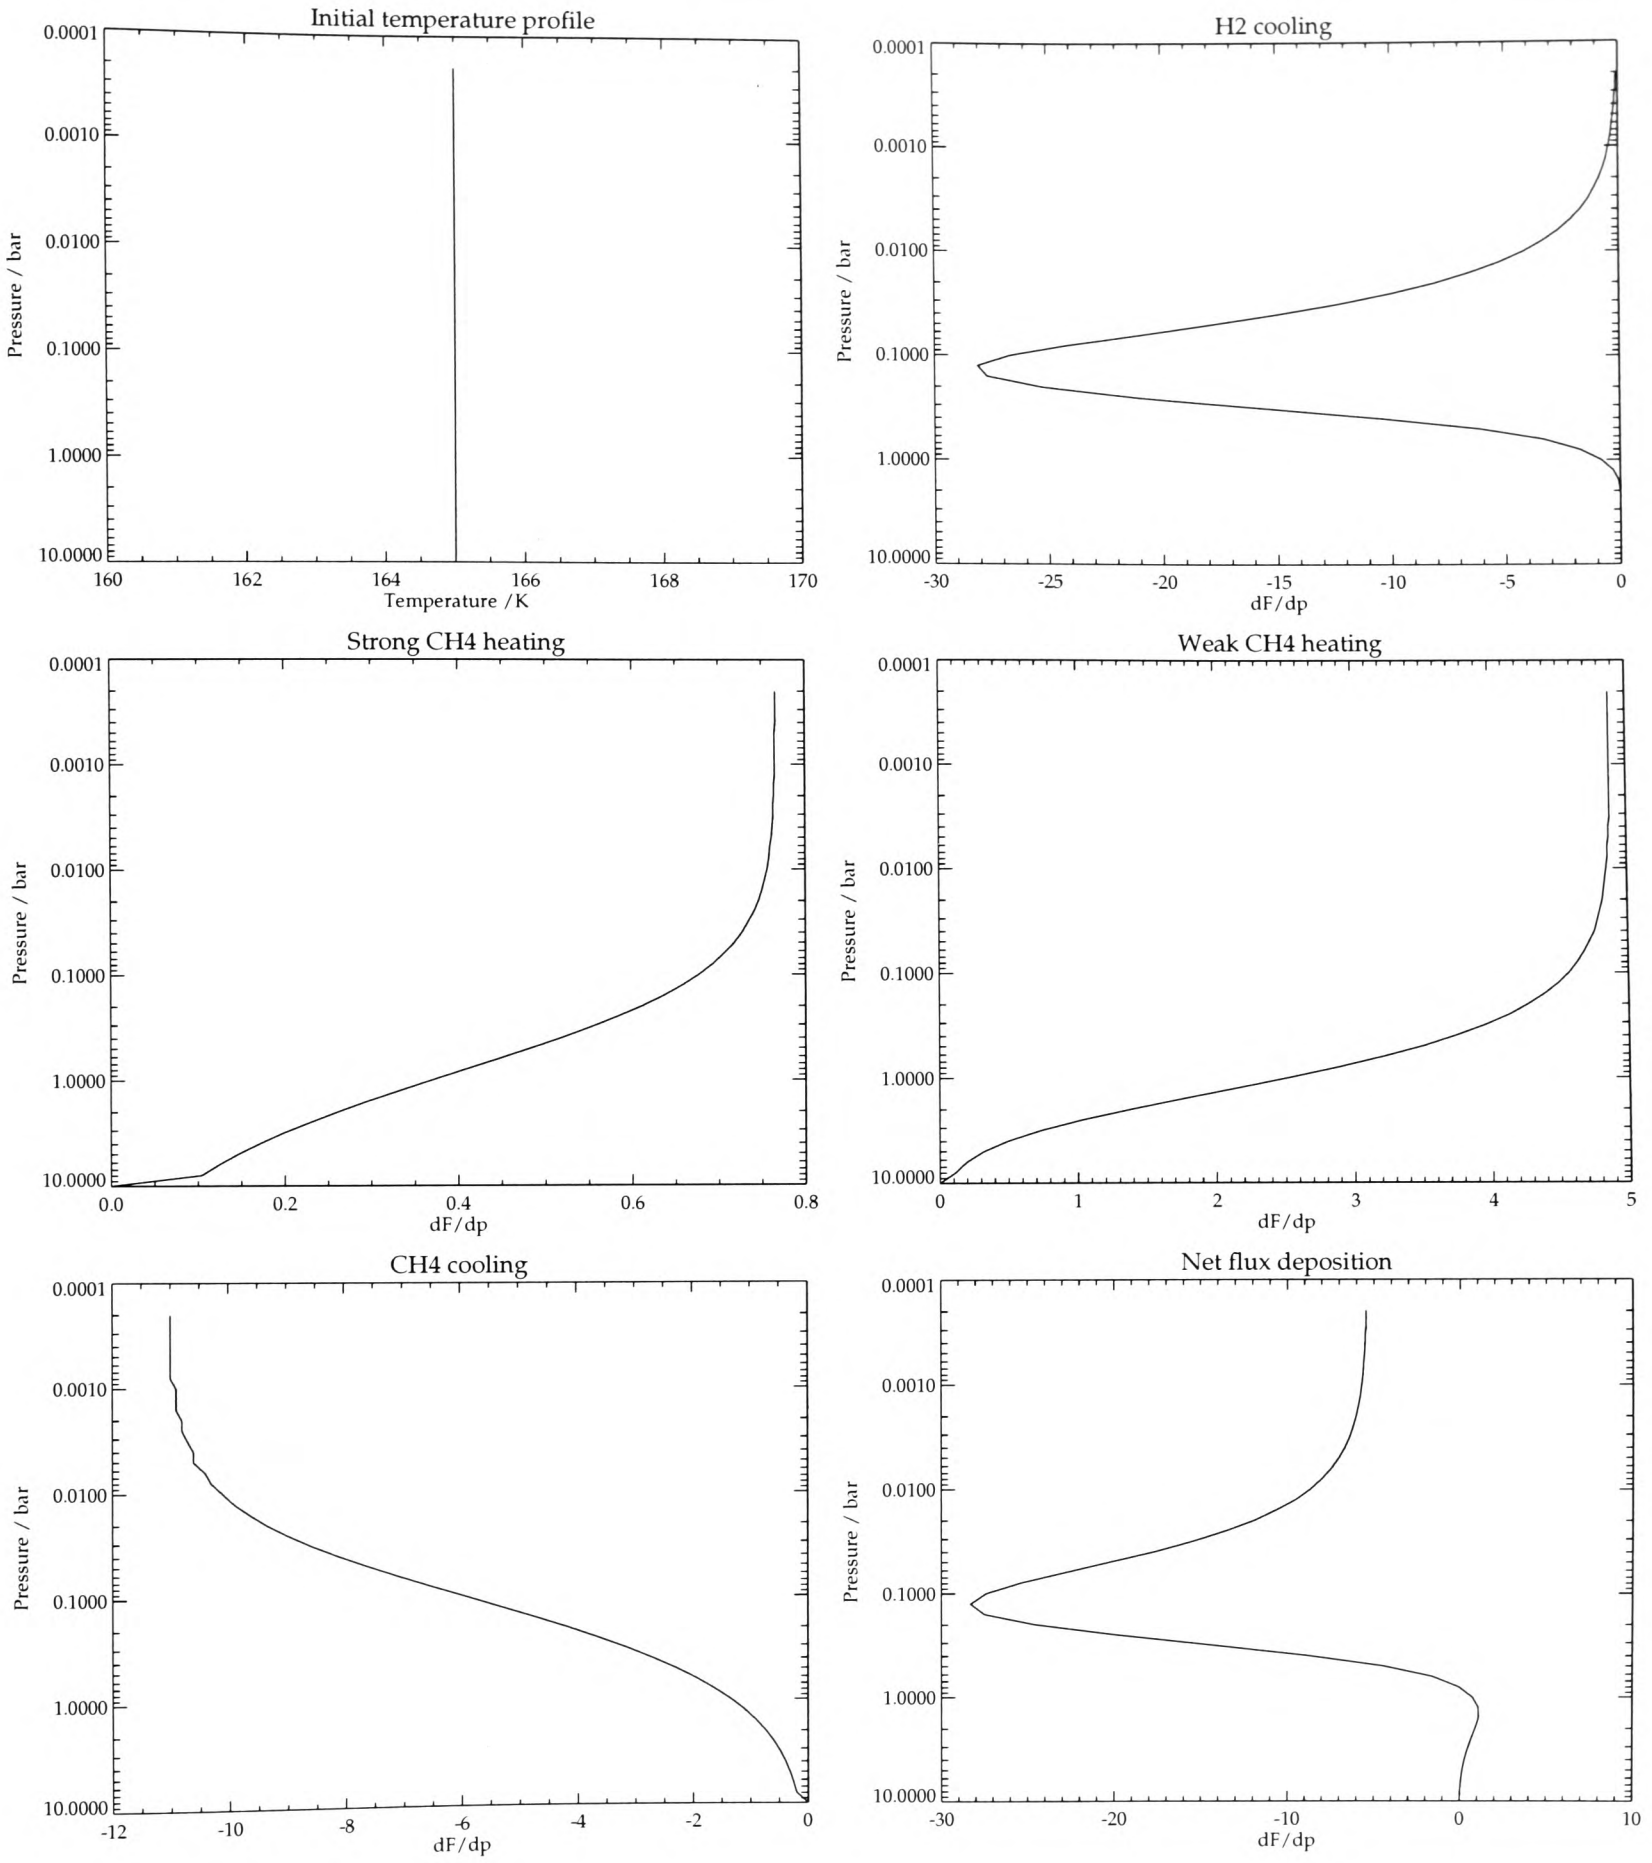


Figure 7.1: The initial temperature profile, and contributions to the net flux deposition from different processes.

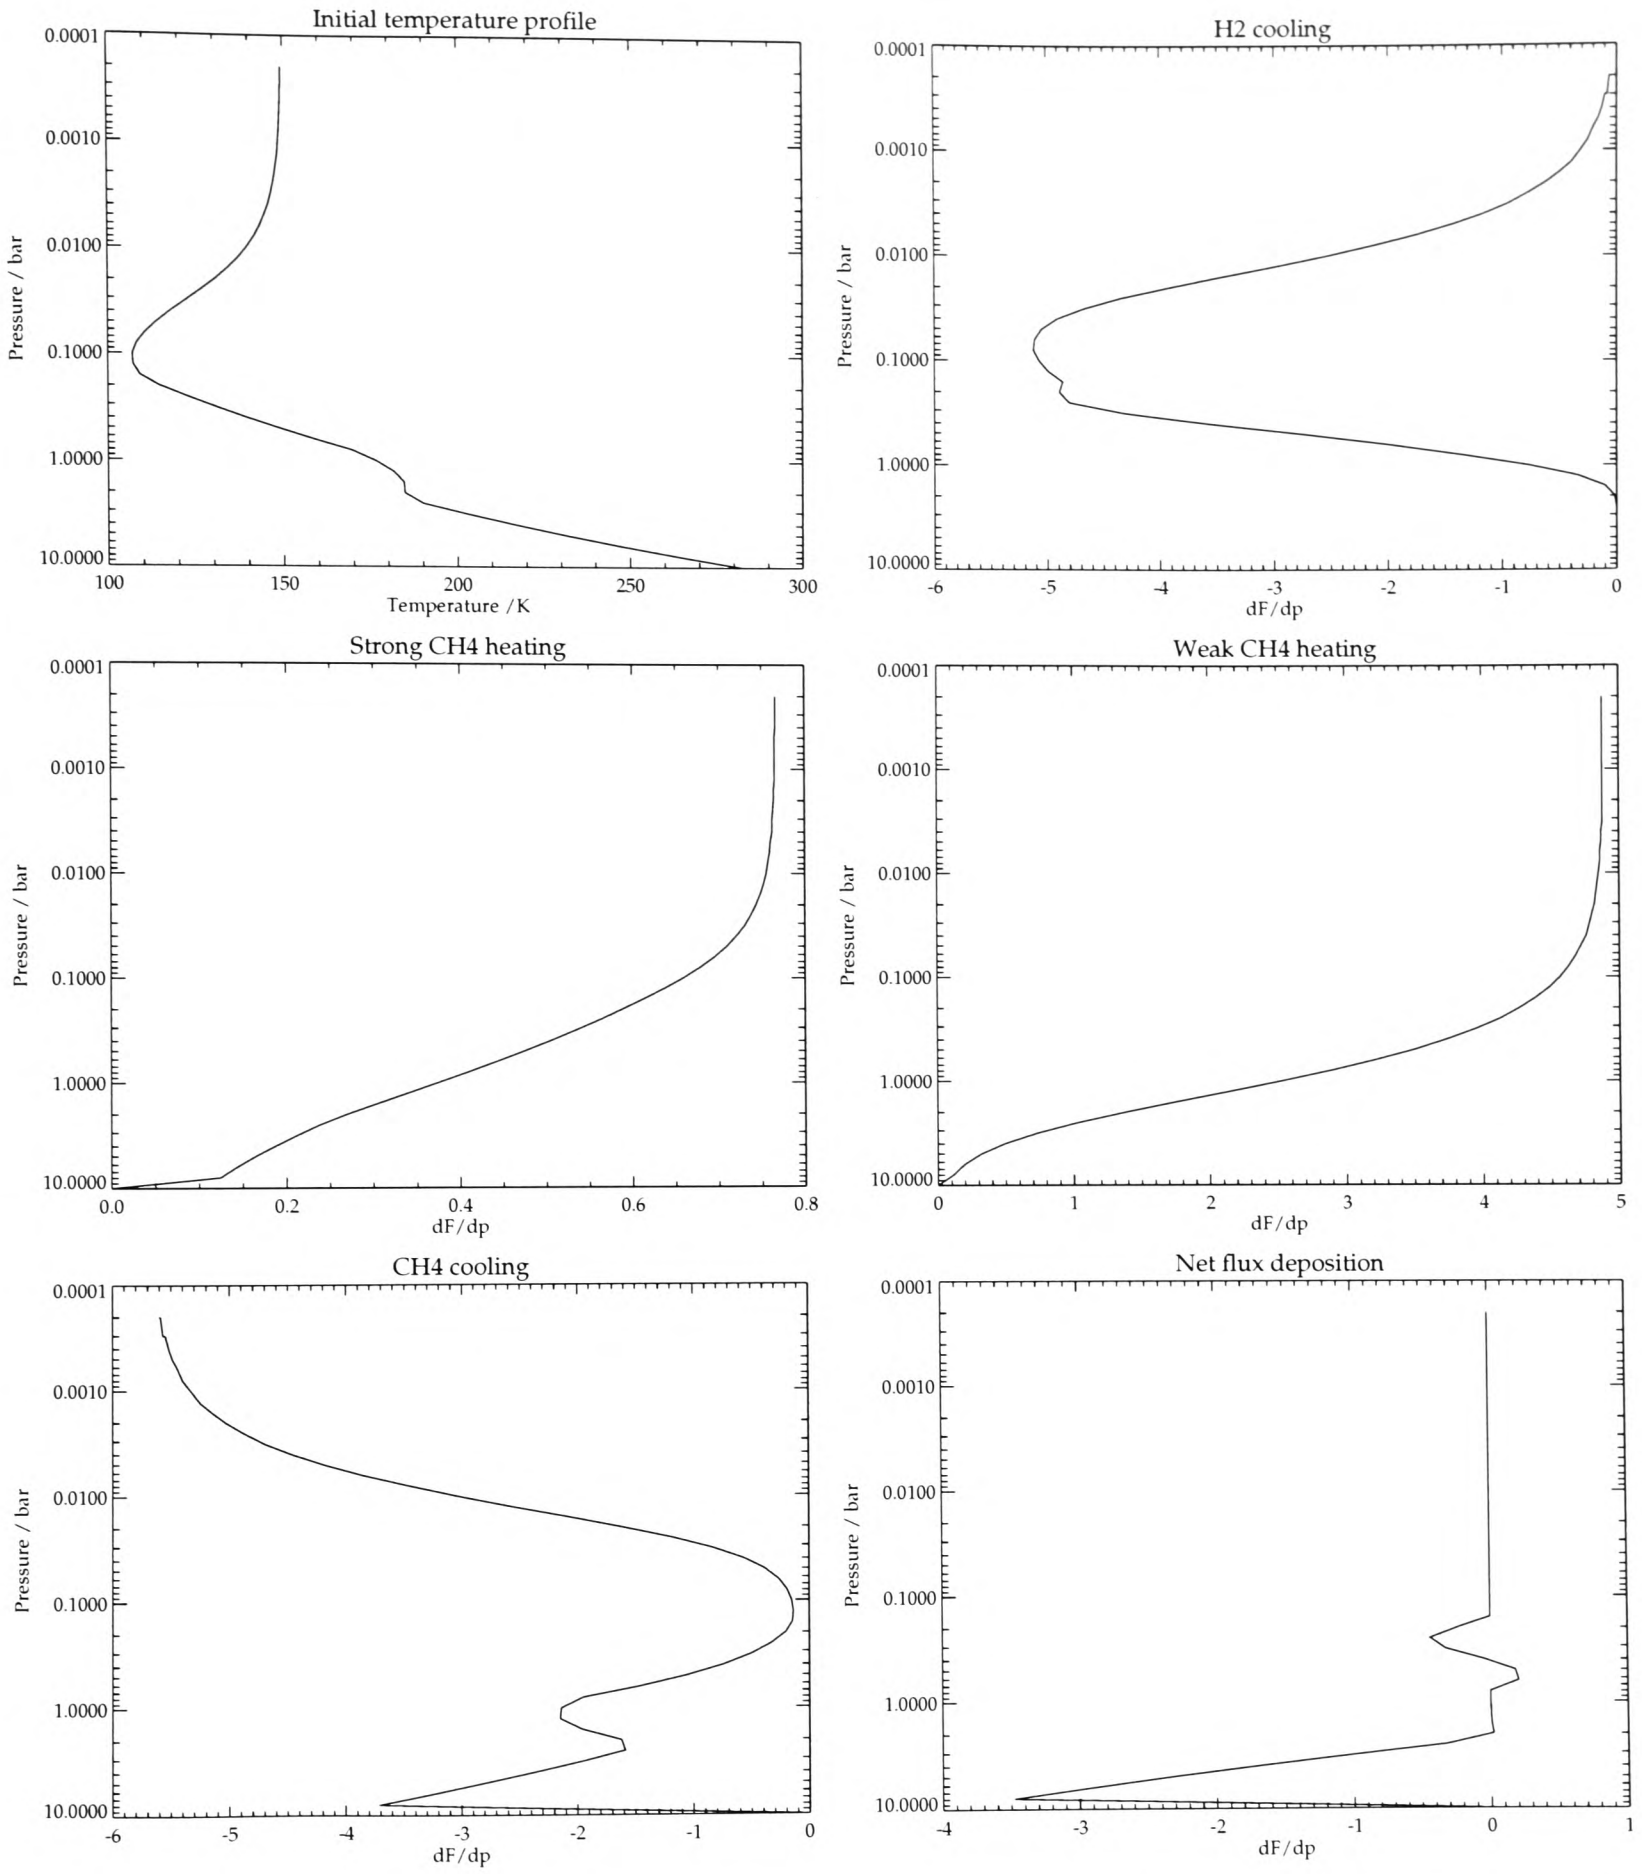


Figure 7.2: The final temperature profile, and the various heating-rate contributions.

In conclusion, the work of Conrath et al. (1990) indicates that this radiation scheme should be suitable for our model. Our implementation, based only on published descriptions, does not yet work perfectly. Results above the 1 bar level are encouraging, but below this there are small problems, we think largely due to the lower boundary condition, but perhaps also to the use of the cooling-to-space approximation. This is increasingly invalid as the levels become deeper. We note that the experiments done by Conrath et al. did not extend below 1 bar. Finally, it would be useful to obtain other workers' cooling/heating rates for each individual contribution. These are not always included in published papers but are important for validating our own model. It is difficult to debug the code when we do not know which aspects are incorrect, and mistakes may offset each other.

Appendix A

Running the Portable UM

The Portable Unified Model, as distributed for use outside the UK Met. Office, is a complicated package and is not entirely straightforward to install or use. The model lacks any graphical facilities for data visualisation and analysis. There are several bugs and limitations, some of which are documented and others not. For these reasons, it is worthwhile describing the operation of the model in some detail. The intention is to provide enough technical information to allow the continuation or repeat of parts of this project. The "Unified Model User Guide", supplied with the software, provides an introduction to the model and its user interface. Here we try to summarize the information, and fill in some gaps in the User Guide. We also explain the methods used to view and analyse the model output.

Installation

The Portable UM is designed for easy installation on Cray supercomputers, or Hewlett-Packard and DEC Alpha workstations running under suitable Unix flavours—UNICOS, HP-UX and Digital Unix respectively. Automated build scripts are provided for these systems. It is possible, but time-consuming, to modify the scripts for other environments.

The UM may only be run under the Unix Korn shell; this must be the user's default login shell. The bugs in the model build scripts are well documented. The model will fail unless its shell has adequate stack space, so the line `ulimit -s 32767` should be included in the login script. Other system resources may also be set to their maximum values, to prevent problems.

The Unified Model User Interface (UMUI) is also simple to install. For this project, Version 4.0 of the UMUI is used, rather than the improved version 4.0.1. The reason for this is that, for the specialized Jupiter version, reconfiguration fails if the model is run using

version 4.0.1. Therefore the older version must be used instead.

This has two known bugs, one of which may be fixed. The required fix (for both versions 4.0 and 4.0.1) is to edit the file `umui-src-4.0.09/vn4.0/src/database_tcl.c`. The declaration on line 128 should be changed from `char rule[64]` to `char rule[256]`. The second bug, found only in version 4.0, concerns the selection of the horizontal domain for model output. This may be avoided, by requesting output for the entire horizontal domain.

Anatomy of a model run

A model job is begun by using the User Interface to define the job parameters, and to submit the job to a suitable machine. Each model run consists of a number of distinct phases, not all of which are required for all runs. These are called in sequence by a master script. The first major step is the reconfiguration. This is optional (selected with the User Interface) and involves interpolating the model start dump on to the required domain and resolution. This provides the model with its initial data. Reconfiguration is not required if the start dump has been created by another run with identical domain and resolution.

After reconfiguration, the model is compiled and linked, if necessary. The UM is very flexible, and many changes may be made using the UMUI, which do not require recompilation. These include changes of resolution, diffusion parameters, and model domain. However, other changes (principally to the model physics) involve modifying the source code, which requires recompilation. Modifications are introduced by means of the `nupdate` source code preprocessor.

`nupdate` is a portable version of a Cray proprietary source code revision management system `update`. This utility extracts the required code from the FORTRAN and C libraries. It executes a file of instructions to insert, delete and replace specified lines of code. This method prevents corruption of the main program files, and it provides a neat way of collecting and recording the changes made.

After `nupdate`, the next step is another source code preprocessor, `lexcon`. This implements an extension to standard FORTRAN, allowing it to allocate array storage dynamically. FORTRAN-77, by contrast, permits only static arrays, for which space is reserved throughout the life of the program. The use of static arrays is faster but wasteful of memory. `lexcon` works by inserting calls to C memory allocation functions into the FORTRAN. Modern compilers such as Digital Fortran do allow dynamic arrays, theoretically removing the need for `lexcon`. However, the Digital Alpha version of the UM has not yet been tested without it, so it was thought wiser to retain `lexcon` and accept the conse-

quent processing overheads.

After the two preprocessor passes, the FORTRAN and C code is compiled. Linking is done at runtime, so a large number of error messages are generated, regarding unresolved symbols. These may be ignored, provided the model runs successfully. Finally, the main executable is called and the model is run. Once the run is complete, a few housekeeping tasks are performed by the control scripts, and processing ends.

Diagnostic output

During a model run, the control scripts and the model itself produce a large amount of textual output. This assists the user both in finding problems, and in monitoring the progress of a successful run. The files are identified by the job identification name, which is a five-letter string chosen within the User Interface. This name is stored in the Unix environment variable `$RUNID`. For external users, the first letter is always 'x'. The executable usually has the same name as the job ID: for example, `xaaaa`. The main output log file would be called `xaaaa.out`, the reconfigured start dump `xaaaa.start` and the numerical output files `xaaaa.pp0`, `xaaaa.pp1`, `xaaaa.pp2`, etc. See below for a description of pp files.

The contents of the log file are explained in the User Guide, Section 3.21 "Troubleshooting". The output from each stage of the job is found here, in chronological order. The progress of a successful job may be monitored by examining the last line of the `.out` file: it will usually be "ATMOS TIMESTEP n", where n is the timestep number reached so far. The log file is deleted at the end of each job.

A second, permanent output file is contained in the directory `umui_out`. This has a long, unique filename ending in `.leave` and contains all text output from the model scripts and program. At the end of the job, the contents of the `.out` file are inserted into this file, to preserve a complete record of the job.

Scientific output

The Unified Model contains a sophisticated output processing system, known as STASH. This derives the values of many different meteorological quantities, and writes the requested fields to a choice of up to four files. Spatial and temporal averaging may be performed, and the time and space domain for the output altered as required. The most significant limitation to this system is that some variables are output on the p, θ grid and others on the u, v grid. Some quantities, such as temperature, are available interpolated onto the u, v grid, but others (including PV) are not.

The User Interface STASH windows perform validation checks on the output requests. This causes problems when pressure levels greater than 1.1 bar are specified. The valid input range may be changed by editing line 430 of the `tcl` program `umui/vn4.0/tcl/UM/chk_profile_vars.tcl`.

STASH output is converted to so-called PP (post-processing) format by a UKMO utility, `convpp`. This produces an unformatted flat file consisting of interleaved headers and data records. The format is documented in UM Documentation Paper F3, "Format of atmospheric and oceanic dumps, fieldsfiles,...".

Data visualization

The package used to visualize the data is the Grid Analysis and Display System (GrADS 1.5) by Brian Doty. This is a freely available tool which is designed for display of geophysical data. A dataset is considered as a four-dimensional array, which may be "sliced" along any plane for viewing. For example, with height held constant, GrADS will display an animated contour plot of the data as a function of latitude, longitude, and time. With longitude and time held constant, we obtain a meridional cross-section. With time the only variable, the result is a time-series at a fixed point, and so on.

GrADS also includes many mathematical functions, including averaging, centred differencing and integration. Operations may be automated by use of the GrADS scripting language, and we use this extensively to calculate energy diagnostics, potential-vorticity gradients, and other derived quantities. GrADS can handle several files simultaneously, which is necessary because our UM fields are held on two staggered grids, and therefore in two files. The graphical output may be converted to PostScript for printing or inclusion in documents such as this thesis.

The output is converted from PP format to GrADS format by a set of programs written in the Interactive Data Language (IDL). These read the header records to establish the grid size and resolution, write this information into a GrADS control file, and write the data into a GrADS dataset file. The latter is a raw binary file containing only the data array with no headers.

Tracers

The use of tracers is relatively straightforward in the Unified Model, which is designed to allow up to 29 different tracers at once. Although diagnostic output of tracer concentrations is simple to arrange, initialization of a tracer field is difficult. The UM performs very thorough validation of input data, and ancillary files cannot yet be reconfigured (interpo-

lated) so the tracer ancillary file must precisely match the start dump. The main reason seems to be that few users wish to set up tracers with arbitrary distribution; most use ancillary files interpolated from operational analyses. If software to reconfigure tracer files cannot be obtained, it would be necessary to write a program to generate files which can pass the validity checks. Ironically, the actual data do not matter—they can easily be overwritten at the start of the model run.

Many of these remarks apply equally to the treatment of water and moist convection in the model. The standard UM does of course explicitly include water in all three phases, with calculation of virtual temperature, etc. We caution against the use of these intrinsic capabilities. At least one bug has been identified, which corrupts the data when a model restart-dump is written, effectively filling the atmosphere with water-ice clouds! In our Jupiter version, the humidity and condensate fields are explicitly set to zero, as are the latent heats of water. This treats the symptoms but the cause is still mysterious. We suggest the implementation of two or three tracers, representing water, ammonia, para-hydrogen, and perhaps extra tracers for the other phases of water. These would be passive as far as the main model is concerned; their physical effects would all be parametrized within the new convection scheme.

References

Achterberg, R. K. and Ingersoll, A. P., 1989: A normal-mode approach to Jovian atmospheric dynamics. *J. Atmos. Sci.*, **46**, 2448–2462.

Allison, M., 1990: Planetary waves in Jupiter's equatorial atmosphere. *Icarus*, **83**, 282–307.

Allison, M., 2000: A similarity model for the windy Jovian thermocline. *Plan. Space Sci.*

Andrews, D. G., Holton, J. R. and Leovy, C. B., 1987: *Middle Atmosphere Dynamics*. Academic Press.

Arakawa, A. and Lamb, V. R., 1977: Computational design of the basic dynamical processes of the UCLA general circulation model. *Methods in Computational Physics*, **Vol. 17**.

Atkinson, D. H., Ingersoll, A. P. and Seiff, A., 1997: Deep winds on Jupiter as measured by the Gailieo probe. *Nature*, **388**, 649–650.

Atkinson, D. H., Pollack, J. B. and Seiff, A., 1998: The Galileo Probe Doppler Wind Experiment: Measurement of the deep zonal winds on Jupiter. *JGR Planets*, **103**(E10), 22911–22928.

Atkinson, D. M., Pollack, J. B. and Seiff, A., 1996: Galileo Doppler measurements of the deep zonal winds at Jupiter. *Science*, **272**, 842–843.

Banfield, D., Gierasch, P. J., Bell, M., Ustinov, E., Ingersoll, A. P., Vasavada, A. R., West, R. A. and Belton, M. J. S., 1998: Jupiter's cloud structure from Galileo imaging data. *Icarus*, **135**, 230–250.

Bézard, B., Gautier, D. and Conrath, B. J., 1984: A seasonal model of the Saturnian upper troposphere: comparison with Voyager infrared measurements. *Icarus*, **60**, 274–288.

Cess, R. D. and Caldwell, J., 1979: A Saturnian stratospheric seasonal climate model. *Icarus*, **38**, 349–357.

Cess, R. D. and Chen, S. C., 1975: The influence of ethane and acetylene upon the thermal structure of the Jovian atmosphere. *Icarus*, **26**, 444–450.

Cess, R. D. and Khetan, S., 1973: Radiative transfer within the atmospheres of the giant planets. *J. Quant. Spectrosc. Radiat. Transfer*, **13**, 995–1009.

Chang, C.-C., 1976: Forcing of stratospheric Kelvin waves by tropospheric heat sources. *J. Atmos. Sci.*, **33**, 740–744.

Conrath, B. J and Gierasch, P. J., 1984: Global variation of the para fraction in Jupiter's atmosphere and implications for dynamics on the outer planets. *Icarus*, **57**, 184–204.

Conrath, B. J., Gierasch, P. J. and Leroy, S. S., 1990: Temperature and circulation in the stratospheres of the outer planets. *Icarus*, **83**, 255–281.

Cullen, M. J. P. and Davies, T., 1991: A conservative split-explicit integration scheme with fourth-order horizontal advection. *Q. J. R. Meteorol. Soc.*, **117**, 993–1002.

Cullen, M. J. P., Davies, T. and Mawson, M. H., 1993: *UMDP 10: Conservative finite difference schemes for a unified forecast/climate model*. Tech. rept. UK Met. Office.

Del Genio, A. D. and McGrattan, K. B., 1990: Moist convection and the vertical structure and water abundance of Jupiter's atmosphere. *Icarus*, **84**, 29–53.

Deming, D., Mumma, M. J., Espenak, F., Jennings, D. E., Kostiuk, T., Wiedemann, G., Loewenstein, R. and Piscitelli, J., 1989: A search for p-mode oscillations of Jupiter: Serendipitous observations of non-acoustic thermal wave structure. *Astrophys. J.*, **343**, 456–467.

Deming, D., Reuter, D., Jennings, D. and Bjoraker, G., 1997: Observations and analysis of longitudinal thermal waves on Jupiter. *Icarus*, **126**, 301–312.

Denning, W. F., 1899: Early history of the Great Red Spot on Jupiter. *Monthly Notices RAS*, **59**, 574–584.

Dowling, T. E., 1995: Estimate of Jupiter's deep zonal wind profile from Shoemaker-Levy 9 data and Arnol'd's second stability criterion. *Icarus*, **117**, 439–442.

Dowling, T. E. and Ingersoll, A. P., 1988: Potential vorticity and layer thickness variations in the flow around Jupiter's Great Red Spot and White Oval BC. *J. Atmos. Sci.*, **45**, 1380–1396.

Dowling, T. E. and Ingersoll, A. P., 1989: Jupiter's Great Red Spot as a shallow water system. *J. Atmos. Sci.*, **46**, 3256–3278.

Dowling, T. E., Fischer, A. S., Gierasch, P. J., Harrington, J., LeBeau, Jr., R. P. and Santori, C. M., 1998: The Explicit Planetary Isentropic-Coordinate (EPIC) Atmospheric Model. *Icarus*, **132**, 221–238.

Flasar, F. M., 1986: Global dynamics and thermal structure of Jupiter's atmosphere. *Icarus*, **65**, 280–303.

Flasar, F. M., Conrath, B. J. and Pirraglia, J. A., 1981: Thermal structure and dynamics of the Jovian atmosphere. 1. The Great Red Spot. *J. Geophys. Res.*, **86**, 8759–8767.

Flierl, G. R., 1978: Models of vertical structure and the calibration of two-layer models. *Dynamics of Atmospheres and Oceans*, **2**, 341–381.

Friedson, A. J., 1998: New observations and modelling of a QBO-like oscillation in Jupiter's stratosphere. *Icarus*. in submission.

Gehrels, T. (Ed.)., 1976: *Jupiter*. Univ. of Arizona Press.

Gierasch, P. J. and Conrath, B. J., 1992: Dynamics of the Atmospheres of the Outer Planets.

Gierasch, P. J., Conrath, B. J. and Magalhaes, J. A., 1986: Zonal mean properties of Jupiter's upper troposphere from Voyager infrared observations. *Icarus*, **67**, 456–483.

Gill, A.E., 1982: *Atmosphere-Ocean Dynamics*. Academic Press.

Golitsyn, G. S., 1970: A similarity approach to the general circulation of planetary atmospheres. *Icarus*, **13**, 1–24.

Goody, R. M. and Yung, Y. L., 1989: *Atmospheric Radiation: Theoretical Basis*. Oxford University Press.

Green, N. E., 1887: On the belts and markings of Jupiter. *Memoirs RAS*, **49**(2), 259–270.

Guillot, T., Gautier, D., Chabrier, G. and Mosser, B., 1994a: Are the Giant Planets Fully Convective? *Icarus*, **112**, 337,353.

Guillot, T., Chabrier, G., Morel, P. and Gautier, D., 1994b: Nonadiabatic Models of Jupiter and Saturn. *Icarus*, **112**, 354,367.

Hadley, G., 1735: Concerning the cause of the general trade-winds. *Phil. Trans. Roy. Soc.*, **39**, 58–62.

Halley, E., 1686: An historical account of the Trade Winds, and Monsoons, observable in the Seas between and near the Tropicks, with an attempt to assign the Phisical cause of the said Winds. *Phil. Trans. Roy. Soc.*, **16**, 153–168.

Hanel, R., Conrath, B., Gautier, D., Gierasch, P., Kumar, S., Kunde, V., Lowman, P., Maguire, W., Pearl, J., Pirraglia, J., Ponnamperuma, C. and Samuelson, R., 1977: The Voyager spectroscopy and radiometry investigation. *Space Sci. Rev.*, **21**, 129–157.

Hanel, R., Conrath, B., Flasar, M., Kunde, V., Lowman, P., Maguire, W., Pearl, J., Pirraglia, J., Samuelson, R., Gautier, D., Gierasch, P., Kumar, S. and Ponnoamperuma, C., 1979a: Infrared observations of the Jovian system from Voyager 1. *Science*, **204**, 972–976.

Hanel, R., Conrath, B., Flasar, M., Herath, L., Kunde, V., Lowman, P., Maguire, W., Pearl, J., Pirraglia, J., Samuelson, R., Gautier, D., Gierasch, P., Horn, L., Kumar, S. and Ponnoamperuma, C., 1979b: Infrared observations of the Jovian system from Voyager 2. *Science*, **206**, 952–956.

Hanel, R., Conrath, B. J., Herath, L. W., Kunde, V. G. and Pirraglia, J. A., 1981: Albedo, internal heat and energy balance of Jupiter: Preliminary results of the Voyager infrared experiment. *J. Geophys. Res.*, **86**, 8705–8712.

Harshvardhan and Cess, R. D., 1976: Stratospheric aerosols: effect upon atmospheric temperature and global climate. *Tellus*, **28**, 1–10.

Held, I. M. and Hou, A. Y., 1980: Nonlinear axially symmetric circulations in a nearly inviscid atmosphere. *J. Atmos. Sci.*, **37**, 515–33.

Held, I. M. and Suarez, M. J., 1994: A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. *Bull. AMS*, **75**(10), 1825–1830.

Hide, R., 1961: Origin of Jupiter's Great Red Spot. *Nature*, **190**, 895–896.

Hide, R., 1969: Dynamics of the atmospheres of the major planets with an appendix on the viscous boundary layer at the rigid bounding surface of an electrically-conducting rotating fluid in the presence of a magnetic field. *J. Atmos. Sci.*, **26**, 841–853.

Hide, R., 1980: High vorticity regions in rotating thermally-driven flows. *Met. Mag.*, **110**, 335–344.

Hignett, P., White, A. A., Carter, R. D., Jackson, W. D. N. and Small, R. M., 1985: A comparison of laboratory measurements and numerical simulations of baroclinic wave flows in a rotating cylindrical annulus. *Quart. J. Roy. Met. Soc.*, **111**, 131–154.

Holton, J. R., 1972: Waves in the equatorial stratosphere generated by tropospheric heat sources. *J. Atmos. Sci.*, **29**, 368–375.

Holton, J. R., 1992: *An Introduction to Dynamic Meteorology*. 3 edn. Academic Press.

Hooke, R., 1665: A spot in one of the belts of Jupiter. *Phil. Trans. Roy. Soc.*, **1**, 3.

Hoskins, B. J. and Bretherton, F. P., 1972: Atmospheric frontogenesis models: mathematical formulation and solution. *J. Atmos. Sci.*, **29**, 11–37.

Hoskins, B. J. and Pedder, M. A., 1980: The diagnosis of middle latitude synoptic development. *Quart. J. R. Met. Soc.*, **106**, 707–719.

Hoskins, B. J., MacIntyre, M. E. and Robertson, A. W., 1985: On the use and significance of isentropic potential vorticity maps. *Q. J. R. Meteorol. Soc.*, **111**, 877–946.

Hough, G. W., 1905: Our present knowledge of the condition of the surface of the planet Jupiter. *Popular Astronomy*, **13**(121), 19–30.

Hubbard, W. B., 1968: Thermal structure of Jupiter. *Astrophys. J.*, **152**, 745–754.

Hunt, G. E., Conrath, B. J. and Pirraglia, J. A., 1981: Visible and infrared observations of Jovian plumes during the Voyager encounter. *J. Geophys. Res.*, **86**(A10), 8777–8781.

Hunten, D. M., Colin, L. and Hansen, J. E., 1986: Atmospheric science on the Galileo mission. *Space. Sci. Rev.*, **44**, 191–240.

Ingersoll, A. P., 1969: Inertial Taylor columns and Jupiter's Great Red Spot. *J. Atmos. Sci.*, **26**, 744–752.

Ingersoll, A. P., 1973: Jupiter's Great Red Spot: a free atmospheric vortex? *Science*, **182**, 1346–1348.

Ingersoll, A. P. and Cuong, P. G., 1981: Numerical model of long-lived Jovian vortices. *J. Atmos. Sci.*, **38**, 2067–2076.

Ingersoll, A. P. and Cuzzi, J. N., 1969: Dynamics of Jupiter's cloud bands. *J. Atmos. Sci.*, **26**, 981–985.

Ingersoll, A. P., Kanamori, H. and Dowling, T. E., 1994: Atmospheric gravity waves from the impact of comet Shoemaker-Levy 9 with Jupiter. *Geophys. Res. Lett.*, **21**(11), 1083–1086.

Irwin, P. G. J., Weir, A. L., Smith, S. E., Taylor, F. W., Lambert, A. L, Calcutt, S. B., Cameron-Smith, P. J., Carlson, R. W., Baines, K., Orton, G. S., Drossart, P., Encrenaz, T. and Roos-Serote, M., 1998: Cloud structure and atmospheric composition of Jupiter retrieved from Galileo near-infrared mapping spectrometer real-time spectra. *J. Geophys. Res.*, **103**(E10), 23001–23021.

James, I., 1994: *Introduction to Circulating Atmospheres*. Cambridge University Press.

Kálnay de Rivas, E., 1972: On the use of nonuniform grids in finite-difference equations. *J. Comp. Phys.*, **10**, 202–210.

Kritzinger, H. H., 1914: On the physical constitution of Jupiter. *J. British Ast. Assoc.*, **24**, 452–463.

LeBeau, R. P. and Dowling, T. E., 1998: EPIC simluations of time-dependent, three-dimensional vortices with application to Neptune's Great Dark Spot. *Icarus*, **132**, 239–265.

Leovy, C. B., 1964: Simple models of thermally driven mesosphere circulation. *J. Atmos. Sci.*, **21**, 327–341.

Leovy, C. B., Friedson, A. J. and Orton, G. S., 1991: The quasiquadrennial oscillation of Jupiter's equatorial stratosphere. *Nature*, **354**, 380–382.

Li, X. and Read, P. L., 1999: A mechanistic model of the quasi-quadrennial oscillation in Jupiter's atmosphere. *Plan. Space. Sci.*

Limaye, S. S., 1986: New estimates of the mean zonal flow at the cloud level. *Icarus*, **65**, 335–352.

Lindal, G. F., Wood, G. E., Anderson, J. D., Sweetnam, D. N., Hotz, H. B., Buckles, B. J., Holmes, D. P. and Doms, P. E., 1981: The atmosphere of Jupiter: an analysis of the Voyager radio occultation measurements. *J. Geophys. Res.*, **86**, 8721–8727.

Lindzen, R. S., 1971: Equatorial planetary waves in shear. *J. Atmos. Sci.*, **28**, 609–622.

Lindzen, R. S., 1974: Wave CISK in the tropics. *J. Atmos. Sci.*, **31**, 156–179.

Longuet-Higgins, M. S., 1968: The eigenfunctions of Laplace's tidal equations over a sphere. *Phil. Trans. Roy. Soc.*, **262**, 511–607.

Lunine, J. I. and Hunten, D. M., 1987: Moist convection and the abundance of water in the troposphere of Jupiter. *Icarus*, **69**, 556–570.

Mac Low, M. and Ingersoll, A. P., 1986: Merging of vortices in the atmosphere of Jupiter: an analysis of Voyager images. *Icarus*, **65**, 353–369.

Marcus, P. S., 1993: Jupiter's Great Red Spot and other vortices. *Annu. Rev. Astron. Astrophys.*, **31**, 523–573.

Massie, S. T. and Hunten, D. M., 1982: Conversion of para and ortho hydrogen in the jovian planets. *Icarus*, **49**, 213–226.

Matsuno, T., 1966: Quasi-geostrophic motions in the equatorial area. *J. Meteor. Soc. Japan*, **44**, 25,33.

Maxworthy, T., 1975: A wave driven model of the Jovian equatorial jet. *Planet. Space Sci.*, **23**, 1223–1233.

Maxworthy, T. and Redekopp, L. G., 1976: A solitary wave theory of the Great Red Spot and other observed features in the Jovian atmosphere. *Icarus*, **29**, 261–271.

Maxworthy, T., Redekopp, L. G. and Weidman, P. D., 1977: On the production and interaction of planetary solitary waves: applications to the Jovian atmosphere. *Icarus*, 388–409.

Mayr, H. G. and Harris, I., 1983: Quasi-axisymmetric circulation and superrotation in planetary atmospheres. *Astron. Astrophys.*, **121**, 124–136.

Mesinger, F., 1981: Horizontal advection schemes of a staggered grid—an enstrophy and energy conserving model. *Mon. Weather Rev.*, **109**, 758–766.

Mitchell, J. L., Beebe, R. F., Ingersoll, A. P. and Garneau, G. W., 1981: Flow fields within Jupiter's Great Red Spot and White Oval BC. *J. Geophys. Res.*, **86**, 8751–8757.

Niemann, H. B., Atreya, S. K., Carignan, G. R., Donahue, T. M., Habermann, J. A., Harpold, D. N., Hartle, R. E., Hunten, D. M., Kasprzak, W. T, Mahaffy, P. R., Owen, T. C. and Way, S. H., 1998: The composition of the Jovian atmosphere as determined by the Galileo probe mass spectrometer. *J. Geophys. Res.*, **103**(E10), 22831–22845.

Norton, W. A., 1988: *Balance and potential vorticity inversion in atmospheric dynamics*. Ph.D. thesis, University of Cambridge.

Ortiz, J. L, Orton, G. S., Stewart, S. T. and M., Fisher. B., 1998: Evolution and persistence of 5-μm hot spots at the Galileo Probe entry latitude. *J. Geophys. Res.*, **103**(E10), 23051–23069.

Orton, G., Ortiz, J. L., Baines, K., Bjoraker, G., Carsenty, U., Colas, F., Dayal, A., Deming, D., Drossart, P., Frappa, E., Friedson, J., Goguen, J., Golisch, W., Griep, D., Hernandex, C., Hoffmann, W., Jennings, D., Kaminski, C., Kuhn, J. Laques, P., Limaye, S., Lin, H., Lecacheux, J., Martin, T., McCabe, G., Momary, T., Parker, D., Puetter, R., Ressler, M., Reyes, G., Sada, P., Spencer, J., Spitale, J., Stewart, S., Varsik, J., Warell, J., Wild, W., Yanamandra-Fisher, P., Fazio, G., Hora, J. and Deutsch, L., 1996: Earth-based observations of the Galileo Probe entry site. *Science*, **272**, 839–840.

Pedlosky, J., 1979: *Geophysical Fluid Dynamics*. Springer-Verlag.

Pedlosky, J., 1982: *Geophysical Fluid Dynamics*. Springer-Verlag.

Peek, B. M., 1958: *The Planet Jupiter*. Faber and Faber, London.

Pirraglia, J. A., 1984: Meridional energy balance of Jupiter. *Icarus*, **59**, 169–176.

Plumb, R. A. and Bell, R. C., 1982: Equatorial waves in steady zonal shear flow. *Quart. J. R. Met. Soc.*, **108**, 313–334.

Press, W. H., Flannery, B. P., Teukolsky, S. A. and Vetterling, W. T., 1986: *Numerical Recipies: the Art of Scientific Computing*. Cambridge University Press.

Ragent, B., et al., 1998: The clouds of Jupiter: Results of the Galileo Jupiter Mission Probe Nephelometer Experiment. *JGR Planets*, **103**(E10), 22891–22909.

Read, P. L., 1986: Stable, baroclinic eddies on Jupiter and Saturn: a laboratory analogue and some observational tests. *Icarus*, **65**, 304–334.

Read, P. L. and Hide, R., 1984: An isolated baroclinic eddy as a laboratory analogue of the Great Red Spot on Jupiter. *Nature*, **308**, 45–48.

Read, P. L., Lewis, S. R. and Hide, R., 1997: Laboratory and Numerical Studies of Baroclinic Waves in an Internally Heated, Rotating Fluid Annulus: A Case of Wave/Vortex Duality? *Journal of Fluid Mechanics*, **337**, 155–191.

Reese, E. J. and Smith, B. A., 1968: Evidence of vorticity in the Great Red Spot of Jupiter. *Icarus*, **9**, 474–486.

Rogers, J. H., 1995: *The Giant Planet Jupiter*. Cambridge University Press.

Roos-Serote, M. and authors, 12., 1998: Analysis of Jupiter North Equatorial Belt hot spots in the 4–5 μm range from Galileo/near-infrared mapping spectrometer observations: Measurements of cloud opacity, water and ammonia. *J. Geophys. Res.*, **103**(E10), 23023–23041.

Sada, P. V., Beebe, R. F. and Conrath, B. G., 1996: Comparison of the structure and dynamics of Jupiter's Great Red Spot between the Voyager 1 and 2 encounters. *Icarus*, **119**, 311–335.

Seiff, A., Kirk, D. B., Knight, T. C. D., Mihalov, J. D., Blanchard, R. C., Young, R. E., Schubert, G., von Zahn, U., Lehmacher, G., Milos, F. and Wang, J., 1996: Structure of the atmosphere of Jupiter: Galileo Probe measurements. *Science*, **272**, 844–845.

Seiff, A., et al., 1998: Thermal structure of Jupiter's atmosphere near the edge of a 5-μm hot spot in the north equatorial belt. *JGR Planets*, **103**(E10), 22857–22889.

Simon, A. A., Beebe, R. F., Gierasch, P. J., Vasavada, A. R. and Belton, M. J. S., 1998: Global context of the Galileo-E6 observations of Jupiter's White Ovals. *Icarus*, **135**, 220–229.

Smith, B. A., Briggs, G. A., Danielson, G. E., Cook, A. F., Davies, M. E., Hunt, G. E., Masursky, H., Soderblom, L. A., Owen, T. C., Sagan, C. and Suomi, V. E., 1977: Voyager imaging experiment. *Space Sci. Rev.*, **21**, 103–127.

Smith, M. D. and Gierasch, P. J., 1995: Convection in the outer planet atmospheres including *ortho-para* hydrogen conversion. *Icarus*, **116**, 159–179.

Stevenson, D. J., 1982: Interiors of the Giant Planets. *Ann. Rev. Earth Plan. Sci.*, **10**, 257–295.

Sundqvist, H. and Veronis, G., 1970: A simple finite-difference grid with non-constant intervals. *Tellus*, **22**, 26–31.

Vasavada, A. R., Ingersoll, A. P., Banfield, D., Bell, M., Gierasch, P. J., Belton, M. S, Orton, G. S., Klassen, K. P, DeJong, E., Breneman, H. H., Jones, T. J., Kaufman, J. M., Magee, K. P and Senske, D. A., 1998: Galileo imaging of Jupiter's atmosphere: the Great Red Spot, Equatorial Region, and White Ovals. *Icarus*, **135**, 265–275.

von Zahn, U. and Hunten, D. M., 1996: The helium mass fraction in Jupiter's atmosphere. *Science*, **272**, 849–851.

Wallace, J. M., 1971: Spectral studies of tropospheric wave disturbances in the tropical western Pacific. *Rev. Geophys.*, **9**, 557–612.

Wallace, L., Prather, M. and Belton, M. J. S., 1974: The thermal structure of the atmosphere of Jupiter. *Astrophys. J.*, **193**, 481–493.

West, R. A., Strobel, D. F. and Tomasko, M. G., 1986: Clouds, aerosols, and photochemistry in the Jovian atmosphere. *Icarus*, **65**, 161–217.

White, A. A. and Bromley, R. A., 1995: Dynamically consistent, quasi-hydrostatic equations for global models with a complete representation of the Coriolis force. *Q. J. R. Meteorol. Soc.*, **121**, 399–418.

Williams, G. P., 1975: Jupiter's atmospheric circulation. *Nature*, **257**, 778.

Williams, G. P., 1978: Planetary circulations: 1. Barotropic representation of Jovian and terrestrial turbulence. *J. Atmos. Sci.*, **35**, 1399–1426.

Williams, G. P., 1979: Planetary circulations: 2. The Jovian quasi-geostrophic regime. *J. Atmos. Sci.*, **36**, 932–968.

Williams, G. P., 1996: Jovian dynamics: part I: vortex stability, structure, and genesis. *J. Atmos. Sci.*, **53**(18), 2685–2734.

Williams, G. P., 1997: Planetary vortices and Jupiter's vertical structure. *J. Geophys. Res.*, **102**(E4), 9303–9308.

Williams, G. P. and Robinson, J. H., 1973: Dynamics of a convectively unstable atmosphere: Jupiter? *J. Atmos. Sci.*, **30**, 684–717.

Williams, G. P. and Wilson, R. J., 1988: The stability and genesis of Rossby vortices. *J. Atmos. Sci.*, **45**, 207–241.

Young, R. E., 1998: The Galileo probe mission to Jupiter: Science overview. *JGR Planets*, **103**(E10), 22778–22790.

Zhang, K. and Schubert, G., 1996: *Science*, **273**, 941–943.