The Sporadic Nature of Meridional Heat Transport in the Atmosphere

Gabriele Messori

Thesis Submitted for the Degree of Doctor of Philosophy of Imperial College London and for the Diploma of Imperial College London

December, 2013

Space and Atmospheric Physics Group,

Department of Physics,

Imperial College London,

London SW7 2AZ,

UK

Declaration of Originality

I hereby declare that the work presented in this thesis is my own, and that all else has been appropriately referenced.

Gabriele Messori

September, 2013

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Acknowledgements

This thesis has been a pleasure to write thanks to the support of my family, friends and colleagues.

In the first place I would like to thank my supervisor, Arnaud Czaja, who has spent a good part of his summer reading my drafts. A big thanks also goes to my colleagues, for bearing with me throughout the writing-up period - I know I was rather grumpy! Thanks are also directed to my two assessors, Brian Hoskins and Tim Woollings, who made the *viva* a very stimulating and rewarding experience.

Infine un grandissimo grazie ai miei genitori, amici e alla mia fidanzata, per avermi sempre trascinato via dall'ufficio quando stavo per impazzire.

Abstract

The present study analyses meridional atmospheric heat transport, due to transient eddies, in the European Centre for Medium-Range Weather Forecasts ERA-Interim reanalysis data. Probability density functions of the transport highlight the dominant role played by extreme events. In both hemispheres, events in the top 5 percentiles typically account for over half of the net poleward transport. As a result of this sensitivity to extremes, a large fraction of the heat transport by transient eddies, at a given location and season, is realised through randomly spaced bursts (a few per season), rather than through a continuum of events.

Fast growing atmospheric modes are associated with a large heat transport, suggesting a link between these bursts and growing baroclinic systems (defined here as motions in the 2.5–6 day band). However, wavelet power spectra of the transport extremes suggest that they are driven by very precise phase and coherence relationships, between meridional velocity and moist static energy anomalies, acting over a broad range of frequencies (2-32 days). Motions with periods beyond 6 days play a key role in this framework. Moreover, these longer periods are found to be mainly driven by planetary-scale motions. Notwithstanding this, the heat transport bursts can be matched to specific synoptic-scale patterns. The bursts are therefore interpreted as the signatures of travelling synoptic systems superimposed on larger scale motions.

The dominant role of extreme events can be reproduced in highly idealised simulations. Both a statistical model, where atmospheric motions are assumed to be linear superpositions of sinusoidal curves, and a two-layer model, representing heat transport as a quantised process effected by point vorticity anomalies, are successful in simulating the transport bursts. The fact that two very different idealised models both reproduce the transport's sporadic nature suggests that this must be an intrinsic property of waves in the atmosphere.

Contents

List of Figures

List of Tables

List of Abbreviations

AGCM: Atmospheric General Circulation Model

AOGCM: Atmosphere-Ocean General Circulation Model

APE: Available Potential Energy

CDF: Cumulative Density Function

DJF: December, January and February

ECMWF: European Centre for Medium-Range Weather Forecasts

EVT: Extreme Value Theory

FIR: Finite Impulse Response (filter)

FORTE: Fast Ocean Rapid Troposphere Experiment

JJA: June, July and August

MLV: Most Likely Value (in the context of PDFs)

MSE: Moist Static Energy

MSLP: Mean Sea Level Pressure

NAO: North Atlantic Oscillation

NH: Northern Hemisphere

PDF: Probability Density Function

ppm: parts per million

SH: Southern Hemisphere

SST: Sea Surface Temperature

TOA: Top Of the Atmosphere

WBC: Western Boundary Current

WCB: Warm Conveyor Belt

Part I

Aims of the Thesis

The primary driver of the global climate system is solar energy input. At the top of the atmosphere (TOA), the low latitudes experience a net energy surplus, while the high ones a corresponding net deficit. A direct consequence of this imbalance is a very large meridional heat transport, effected by both the atmosphere and oceans. Figure 1 shows an estimate of the yearly mean transports.

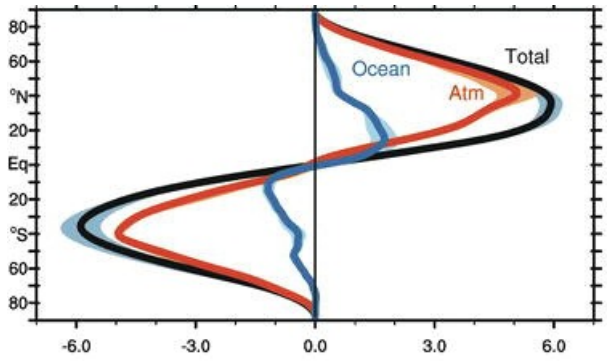

Figure 1: Annual mean meridional heat transport (in PW) by the atmosphere and ocean. The 2σ contours are also shown [from *Fasullo and Trenberth*, 2008].

From the earliest studies of climate dynamics, these have been acknowledged as the key to predicting the time-mean structure of the Earth's climate [*e.g. Budyko*, 1969; *Sellers*, 1969; *Stone*, 1978], and its variability [e.g. *Bjerknes*, 1964]. More recent studies have emphasised the role atmospheric heat transport plays in the response of our climate to anthropogenic forcing, especially at high latitudes [*e.g. Alexeev et al.*, 2005; *Langen and Alexeev*, 2007]. For example, there is evidence that anomalies in atmospheric poleward heat transport might explain the 2007 polar sea-ice minimum [*Graversen et al.*, 2011].

The present thesis focuses on the contribution of time dependent motions to atmospheric poleward heat transport. In mid-to-high latitudes, these motions account for the bulk of the transport. This is especially true in the Southern Hemisphere (SH), where the contribution of stationary waves is small; in the wintertime Northern Hemisphere (NH) stationary and transient eddies provide comparable transports [*Peixoto and Oort*, 1992]. *Lau and Wallace* [1979] have produced a comprehensive analysis of meridional transient-eddy heat transport in terms of rotational and irrotational, and divergent

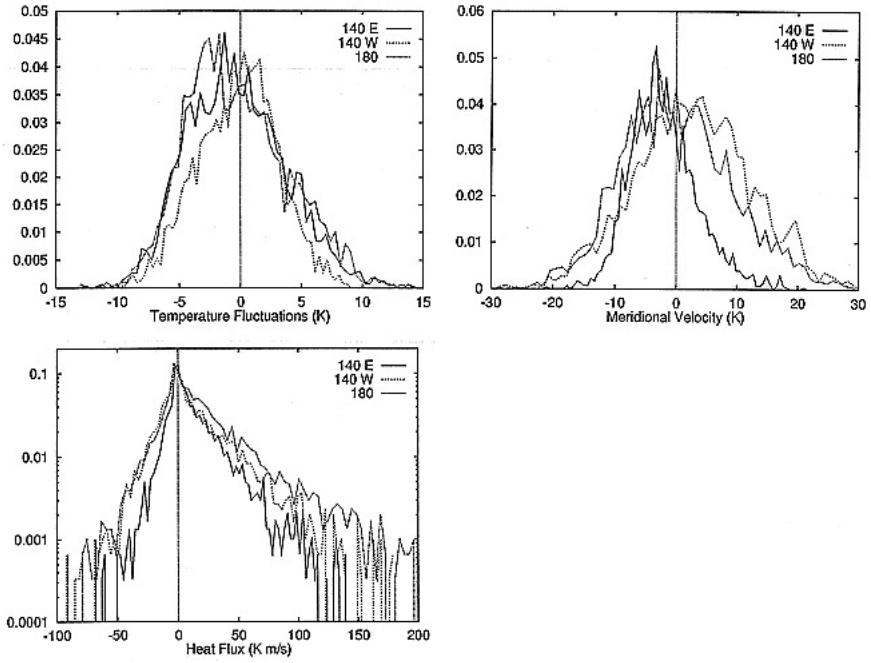

Figure 2: PDFs of potential temperature anomalies (in K, top left), meridional velocity (in ms^{-1}, top right) and meridional heat transport (in Kms^{-1}, bottom left). The data covers the 850 mb fields for November–March 1972–1986. The three locations taken into consideration are 40° N; 140° E, 180°, and 140° W [from *Swanson and Pierrehumbert*, 1997].

and nondivergent, components. This topic has also been the focus of many other studies based on observational *Lau* [e.g. 1978] and theoretical [e.g. *Branscome*, 1983] considerations.

A new perspective on the subject was introduced by *Swanson and Pierrehumbert* [1997], who analysed November–March heat flux probability density functions (PDFs) at three locations in the Pacific storm track. These are shown in figure 2, which also includes the meridional velocity (top left) and potential temperature anomaly (top right) PDFs. The authors noticed that the transport distributions displayed an extended positive tail and found that, surprisingly, the top two percentiles of the distributions typically account for 20% of the heat transport.

This thesis builds on the study by *Swanson and Pierrehumbert* [1997], by extending their results to the Northern and Southern Hemispheres, and to the cold and warm seasons. The key aims can be summarised as follows:

i. Define extreme events in meridional atmospheric heat transport [*Messori and Czaja*, 2013b] (Chap-

ter III);

ii. Demonstrate their key role in setting the net seasonal transport magnitude [*Messori and Czaja*, 2013b] (Chapter III);

iii. Characterise them by looking at their temporal, spatial and spectral structure [*Messori and Czaja*, 2013c] (Chapter IV);

iv. Link them to idealised statistical and minimal physics models (Chapter V).

The extreme events are defined in terms of percentiles of a probability density function (PDF). After robustly establishing the sensitivity of the heat transport to extremes, we demonstrate its sporadic and irregular temporal distribution by showing that a large contribution to the transport arises from a few isolated bursts every season. Next, a detailed analysis of the physical and spectral structure of the extremes, and the associated atmospheric circulations, is performed. The physical interpretation, existence of synoptic analogues and role of large scale motions are discussed in this context. Finally, idealised analogues are used to test and support the findings from points ii) and iii).

Part II

Introduction

1 The Earth's Climate System

This chapter will introduce the basic physical and mathematical framework necessary to understand meridional heat transport in the atmosphere, and will contextualise its role within the global climate system.

1.1 The Global Energy Budget

The basic driver of the Earth's climate system is solar energy input. From the point of view of the energy exchange at the TOA, the equilibrium of the global climate is quite delicate. Indeed, simple energy balance calculations show that a few percent decrease in incoming radiation could initiate a new ice age (e.g. *Sellers*, 1969). Similarly idealised calculations suggest that, if Earth were a perfect black-body, it would be much colder than it currently is. The difference between the measured and theoretical black-body temperatures is due to the greenhouse effect, whereby the constituents of the atmosphere, water vapour *in primis*, absorb and re-emit the outgoing radiation originally emitted at the planetary surface. Another important difference from a black-body is that both the energy input and output are very non-uniform across the globe. This is primarily due to Earth's geometry (which affects incoming TOA radiation) and the surface and atmospheric characteristics (which affect outgoing TOA radiation). The latter generate a complex variability pattern, while the planet's near-spherical geometry gives rise to the underlying trend of decreasing solar input with increasing latitudes. The result is a very inhomogeneous net TOA radiation picture, as shown in figure 3.

As is evident from the figure, the low latitudes experience a net energy gain, while the high latitudes experience a net energy loss; only the mid-latitudes are in approximate balance. A very large meridional heat transport is needed to maintain this apparently imbalanced state. In fact, if both the atmosphere and oceans were to be solid bodies transporting no heat, the temperature

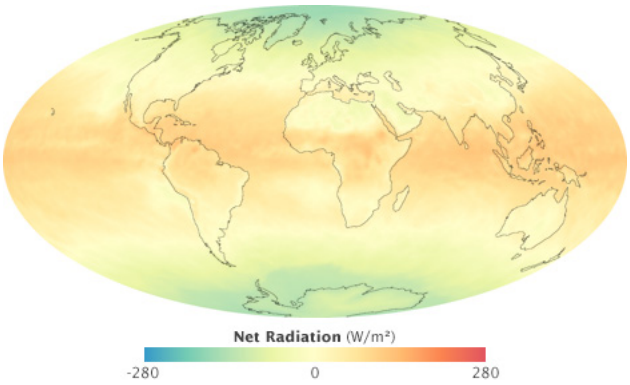

Figure 3: Net TOA radiation during September 2008 (in Wm^{-2}) [by *R. Simmons*, NASA Earth Observatory].

gradient between the equator and the poles would be much larger than what is witnessed [*Salby*, 1996]. The heat transport by the fluid atmosphere and oceans is therefore crucial in maintaining the present day climate.

1.2 The Earth's Atmosphere

The Earth's atmosphere is a very complex and diverse system. Vertically, it is very extended, but approximately 80% of its mass is concentrated within the bottom 10 km. Compared to the radius of the planet, which is roughly 6370 km, the bulk of the atmosphere is a very thin film. This film is made up of permanent gases, notably Nitrogen (78%) and Oxygen (21%), and minor constituents [*Marshall and Plumb*, 1997]. Some of these, even though in very small concentrations, are crucial to the energy balance of the climate system. Water vapour, for example, accounts for approximately 0.5% of the whole atmosphere (and less than 3% of the bottom layer), yet it is at the basis of the greenhouse effect which makes our planet habitable. Similarly, Carbon dioxide (CO_2), at the centre of many Climate Change projections, is present in a concentration so small that it is commonly expressed in parts per million (ppm) rather than a percentage. The current level, which is steadily increasing, is around 400 ppm [*Jones*, 2013].

In the atmospheric sciences, the atmosphere is commonly subdivided into different vertical layers. Here, the interest is in the bottom layer, called the troposphere. This roughly coincides with the 10 km discussed above, and its upper limit is typically defined in terms of the temperature inversion.

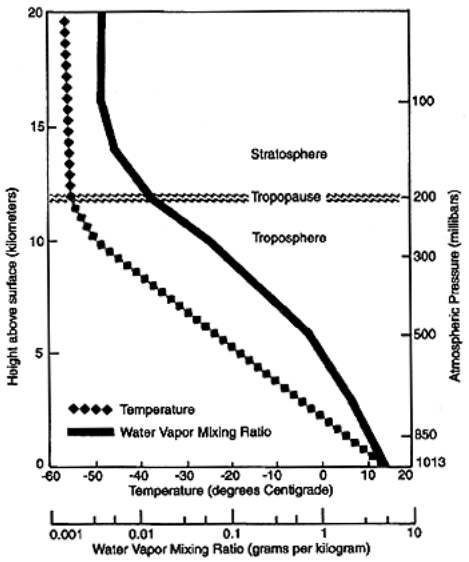

Figure 4: Typical vertical profiles of temperature (in $^{\circ}C$) and water vapor mixing ratio (in gkg^{-1}) in the lower atmosphere [by *D. Gaffen*, Air Resources Laboratory, Maryland].

That is, the height at which the slope of the vertical temperature profile changes abruptly from strongly negative to near-zero [World Meteorological Organisation, 1992]. While this point, called the tropopause, is typically at a height of around 10 km, large local variations are possible. Within the troposphere, the mean vertical temperature and water vapour profiles are reasonably homogeneous, as shown in figure 4.

These two variables are closely interlinked, and are both crucial for atmospheric energetics and dynamics. It is therefore important to define two quantities of interest relating to water vapour: relative and specific humidity. Specific humidity, q, is the ratio of the mass of water vapour to the mass of air per unit volume. It is defined as:

$$q = \frac{\rho_v}{\rho}, \tag{1}$$

where ρ is the total mass of air per unit volume and ρ_v is the mass of water vapour per unit volume. Relative humidity, Q, is then defined as the ratio of specific humidity to saturation specific humidity. The latter is simply the level of specific humidity beyond which condensation occurs. Q is often defined as a percentage:

$$Q = \frac{q}{q_{sat}} \cdot 100, \tag{2}$$

where q_{sat} is saturation specific humidity [*Marshall and Plumb*, 1997]. The temperature dependence mentioned above is contained within q_{sat}.

Another essential quantity needed to describe the atmosphere is potential temperature, θ. Its definition stems from the need for a variable which is conserved under adiabatic displacement, namely a displacement process where there is no heat exchange. Potential temperature is commonly defined as:

$$\theta = T \left(\frac{p_o}{p} \right)^{\frac{R}{c_p}}, \tag{3}$$

where T is temperature, p is pressure, p_0 is a reference pressure relative to which θ is defined (typically chosen as 1000 mb), R is the gas constant and c_p is the specific heat capacity of air at constant pressure. The ratio $\frac{R}{c_p}$ has a value of approximately $\frac{2}{7}$. The vertical gradient of θ is an important measure of atmospheric stability (to dry adiabatic processes), where a positive gradient implies a stable atmosphere. Unlike normal temperature, potential temperature therefore increases with height for a stable atmosphere [*Marshall and Plumb*, 1997]. A typical vertical profile of potential temperature is shown in figure 5.

Thus far, the focus has been on the vertical distribution of the atmosphere. However, the meridional one is equally important. The geometrical effects discussed above (see Section 1.1) play a crucial role in setting this. Figure 6, shows the zonally averaged annual mean temperature distribution, and figure 7 the corresponding potential temperature one. As seen before, T decreases with height and, as expected, also with latitude. The annual mean surface temperature goes below zero beyond approximately 60° N and S. Potential temperature, on the other hand, increases with height, but also decreases with latitude.

As a direct consequence of these temperature gradients, there are also strong meridional pressure gradients. One can define the geopotential height, z_g, of a pressure surface, p, as:

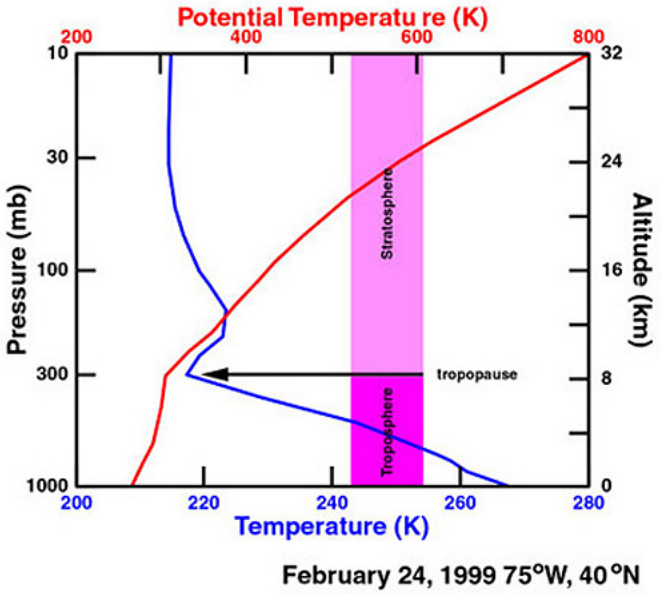

Figure 5: Absolute temperature (blue line) and potential temperature (red line) profiles for a stable atmosphere (in K) [from *Todaro*, 2003].

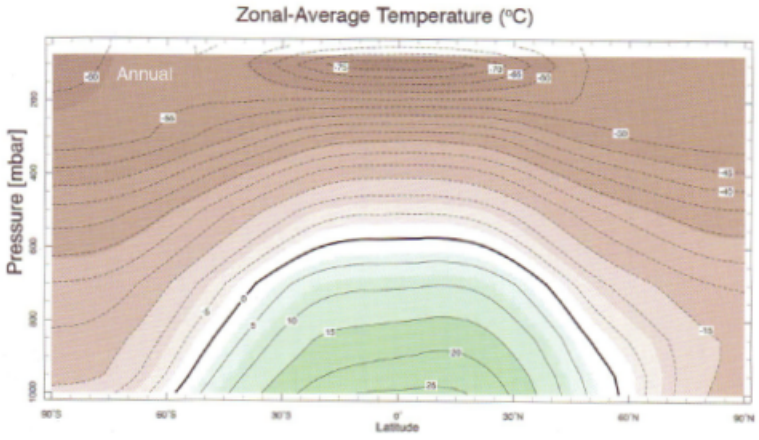

Figure 6: The zonal and annual mean temperature (in $^{\circ}C$) [from *Marshall and Plumb*, 1997].

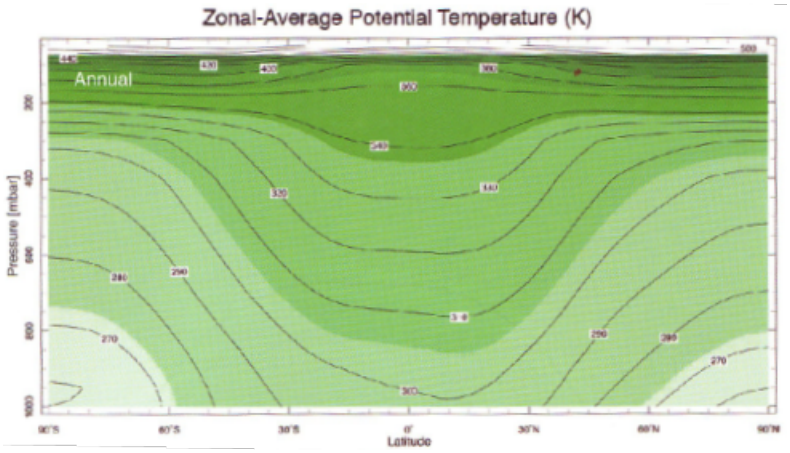

Figure 7: The zonal and annual mean potential temperature (in K) [from *Marshall and Plumb*, 1997].

$$z_g(p) = R \int_p^{p_s} \frac{T}{g} \frac{dp}{p}, \tag{4}$$

where g is gravitational accelleration at mean sea level and p_s is surface pressure. This definition is based on the hydrostatic approximation, which will be discussed in Section 1.3. The temperature dependence in the equation explains the direct link between pressure and temperature gradients. Figure 8 shows the mean 500 mb height in January for the Northern Hemisphere (NH). In the cold polar regions the 500 mb level is significantly lower than in the warm tropics. There are, obviously, other factors affecting the meridional pressure gradient, but the temperature-driven pattern provides a reasonable approximation to what is measured in the real atmosphere. The marked meridional pressure gradient is also a basic driver of atmospheric winds. These will be discussed in the context of large scale circulation patterns in Section 1.3.

1.3 Atmospheric Dynamics

1.3.1 The Equations of Motion

The Earth's atmosphere displays a very rich range of dynamical behaviours. Here, we briefly explore some of the fundamental equations describing them. The starting point is to consider a rotating fluid on a sphere. The equation of motion can be written as:

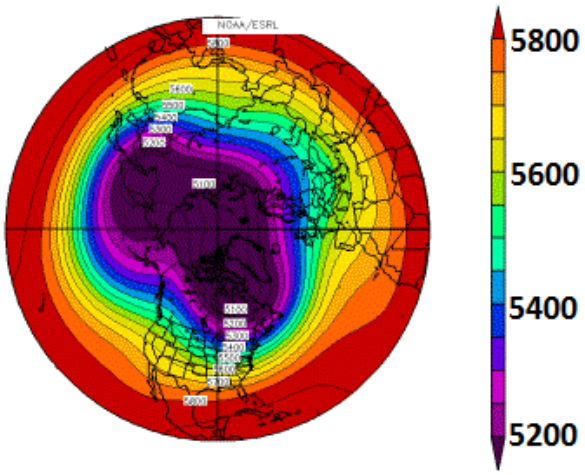

Figure 8: Mean 500 mb height (in m) for NH January 1975-2006 [from NOAA/ESRL Physical Sciences Division].

$$\frac{D\vec{u}}{Dt} + 2\vec{\Omega} \times \vec{u} + \frac{1}{\rho}\nabla p + \nabla\phi = F, \tag{5}$$

where:

$$\phi = gz + \frac{\Omega^2 a^2}{2g} cos^2(\varphi) \tag{6}$$

In the above, $\vec{u}$ is the velocity vector, $\vec{\Omega}$ is the rotation vector of the Earth, ϕ is the modified gravitational potential, F is a lump term representing frictional forces, a is the planetary radius and φ is latitude. Equation 5 can be simplified by making a very simple assumption based on scale analysis, known as hydrostatic balance. This balance describes the variation of pressure with height as a function of density:

$$\frac{\partial p}{\partial z_h} = -g\rho, \tag{7}$$

where z_h is geometrical height. Re-writing equation 5 in local Cartesian coordinates, applying the hydrostatic approximation, and neglecting vertical motions, one obtains:

$$\frac{Du}{Dt} + \frac{1}{\rho}\frac{\partial p}{\partial x} - fv = F_{x,}$$

$$\frac{Dv}{Dt} + \frac{1}{\rho}\frac{\partial p}{\partial y} + fu = F_{y,} \tag{8}$$

$$\frac{1}{\rho}\frac{\partial p}{\partial z} + g = 0,$$

where f, the Coriolis parameter, is given by:

$$f = 2\Omega sin(\varphi) \tag{9}$$

This set of equations provides the basic framework to describe the motion of a thin shell of fluid on a rotating sphere [*e.g.Marshall and Plumb*, 1997].

In fact, these equations account for an incredible diversity of dynamical behaviours, not all of which are relevant to the present work. In terms of the large scale motions, what is crucial is the balance between the different forces acting on the fluid. For the typical large scale atmospheric flow, the Coriolis and pressure gradient terms dominate the left hand side of equation 5, implying that:

$$f\hat{z} \times \vec{u} + \frac{1}{\rho}\nabla p = 0 \tag{10}$$

becomes a good approximation. $\hat{z}$ is simply a unit vector in the vertical direction. This balance between the Coriolis and pressure gradient terms is called geostrophic balance. The resulting horizontal velocity field, called geostrophic wind, is given by:

$$\vec{u}_g = \frac{1}{f\rho}\hat{z} \times \nabla p \tag{11}$$

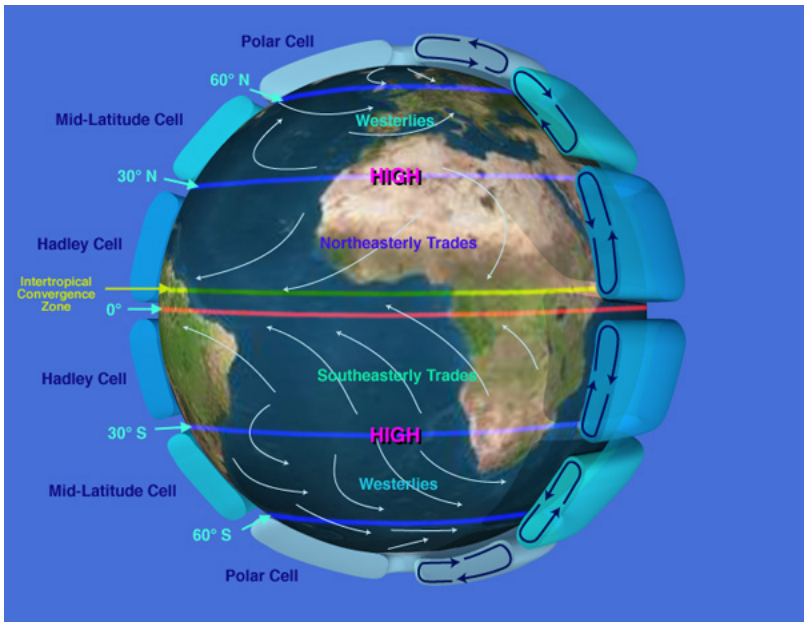

Figure 9: The three circulation cells and the associated surface winds [from NASA Jet Propulsion Laboratories].

1.3.2 The Observed Circulation

The observed large-scale circulation displays three main circulation cells: the Hadley cell (thermally driven circulation), the mid-latitude, or Ferrel, cell (thermally indirect circulation) and the Polar cell (a direct cell). These are schematised in figure 9. The existence of circulation cells is due to the equator-to-pole gradient in net radiation; their number, however, is closely linked to the rate of rotation of the Earth. A very slowly rotating planet would sustain a single equator-to-pole cell, while a very fast rotator, such as Jupiter, will display a higher number of cells [*e.g.* *Peixoto and Oort*, 1992; *Salby*, 1996].

The main surface and high-level winds (again illustrated in figure 9) can be associated with this general circulation structure. In mid-latitudes, the common surface winds are West to East (westerlies) and, in the schematic presented here, would correspond to the bottom branch of the Ferrel cell. The boundary between the Ferrel and Polar cells is called the polar front. Here, cold polar and warm tropical air masses meet, and there is an associated strong zonal wind called the Polar Jet, or often simply the Jet Stream. This jet is found at high levels, typically just below tropopause,

and can attain very high speeds (with winter mean zonal winds at specific locations exceeding 55 ms^{-1}) [*Holton*, 1979]. In the context of the basic equations governing atmospheric dynamics, the Polar Jet is mainly a geostrophic wind. Clearly, this summary is a highly simplified picture of the real circulation, which has a much richer range of features, but correctly reproduces the main observed large scale patterns.

1.3.3 Instabilities in the Atmosphere

The atmospheric dynamics described above all stem from the fact that the atmosphere is subject to instabilities. One of the most fundamental instabilities of geophysical interest is gravitational instability, which is at the base of atmospheric convection. Consider two fluids with the same homogeneous density ρ_2, but different surface heights, separated by a partition. The pressure difference, Δp between two points p_1 and p_2, at equal heights within the fluids, will be given by:

$$\Delta p = \rho_2 g h, \tag{12}$$

where h is the difference in surface heights. Add now a third fluid, ρ_1, placed in equilibrium above the original fluids. The pressure difference between the same two points is now given by:

$$\Delta p = (\rho_2 - \rho_1) g h \tag{13}$$

This can also be expressed as:

$$\Delta p = \rho_2 g' h, \tag{14}$$

where g' is a quantity termed *reduced gravity*, given by:

$$g' = g \frac{\rho_2 - \rho_1}{\rho_2} \tag{15}$$

Figure 10 illustrates this set up. If the partition between the fluids is removed, there will be an adjustment process under gravity. It follows from equation 14 that the driving force behind this adjustement is proportional to $b = \rho_2 g'$. b is termed *buoyancy force* per unit volume [*Gill*, 1982]. If $\rho_1 > \rho_2$, $b < 0$ and the heavier liquid will sink. $Vice-versa$, a lighter liquid on top (*i.e.* $\rho_1 < \rho_2$) will

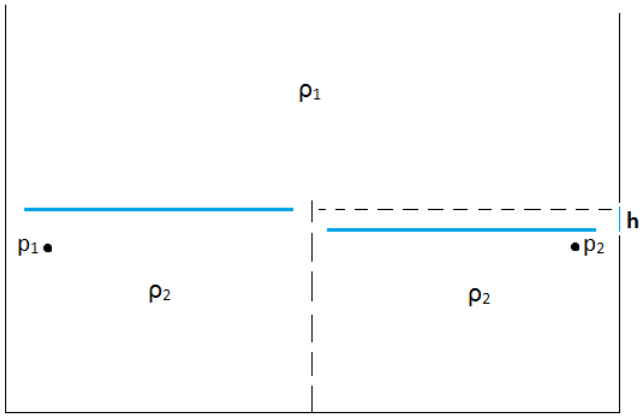

Figure 10: Initial state of the setup illustrating adjustement under gravity. The vertical dashed line is a removable partition between two fluids with density ρ_2, whose upper surfaces are marked by the thick turquoise lines. p_1 and p_2 are two points at equal heights within this fluid. h is the difference in surface heights. Above this there is a third fluid, with density ρ_1.

imply a stable configuration, except for the re-adjustement of the lower fluid to eliminate the height difference h. This basic adjustment process is commonly seen in the atmosphere under the form of convection: if an air parcel becomes lighter (more buoyant) than its surroundings, it will begin to ascend. For an incompressible fluid, and assuming density is only dependent on temperature, an alternative formulation of the stability condition based on temperature in place of density, is:

$$unstable\ if\ \frac{dT}{dz} < 0$$

$$neutral\ if\ \frac{dT}{dz} = 0 \tag{16}$$

$$stable\ if\ \frac{dT}{dz} > 0$$

These conditions do not hold for the real atmosphere, which is a compressible fluid. In order to account for compressibility, potential temperature is usually adopted in place of absolute temperature. It should be noted that humidity and heat exchanges with the environment (diabatic processes) are not considered here. From the point of view of energetics, the air parcel's rising motion will

increase the kinetic energy of the atmosphere but will, at the same time, decrease its potential energy. Returning to the setup illustrated in figure 10, it is clear that, for $b < 0$, the two liquids will re-adjust, with the lighter one on top. Compared to the initial state, the centre of mass of the system has been lowered and its potential energy has decreased. The amount of potential energy which can be converted through this re-adjustement is referred to as *available potential energy* (APE). A more formal definition describes the re-adjustment as an "adiabatic redistribution of mass without phase changes to a statically stable state of rest" [*Gill*, 1982, p.81].

Ignoring the small height difference between the bottom-layer fluids, the mechanism described above is a 1-D process. In the real atmosphere, the conversion of APE can occur through three dimensional mechanisms, and is the energy source for another crucial form of instability, namely baroclinic instability. This is of great interest here, since it can effect heat transport. The crucial component for a baroclinic instability to grow is a horizontal temperature gradient (or, equivalently, a density gradient along an isobar). A baroclinic atmosphere is therefore one where $\rho = \rho(p, T)$. The simplest model of a growing baroclinic instability was developed by *Eady* [1949]. This idealised setup provides an expression for the growth rate and heat transport of the fastest growing mode, and is described in detail in Section 3. In the absence of an isobaric density gradient, there can still be instabilities arising from horizontal variations in the mean velocity. These are called barotropic instabilities, and do not result in a net transport of heat [*Gill*, 1982].

1.3.4 Notable Dynamical Features

One of the main drivers of atmospheric heat transport in the mid to high latitudes are transient eddies [e.g. *Blackmon et al.*, 1977]. These are travelling disturbances resulting from instabilities of the mean flow, and are a direct product of baroclinic instabilities [*Lindzen*, 1990]. Transient eddies are a crucial part of atmospheric dynamics: they are found to drive a range of large-scale patterns, including wave number 6 geopotential anomalies [*Blackmon et al.*, 1984], and are also instrumental in driving synoptic scales.

The term synoptic describes a horizontal length scale, typically taken to be around 1000 km, meant to capture the size of mid-latitude low pressure systems. The most notable synoptic systems are extra-tropical cyclones. It is through these systems that baroclinic instabilities, and the associated eddies, account for a large poleward heat transport [*Eckhardt et al.*, 2004].

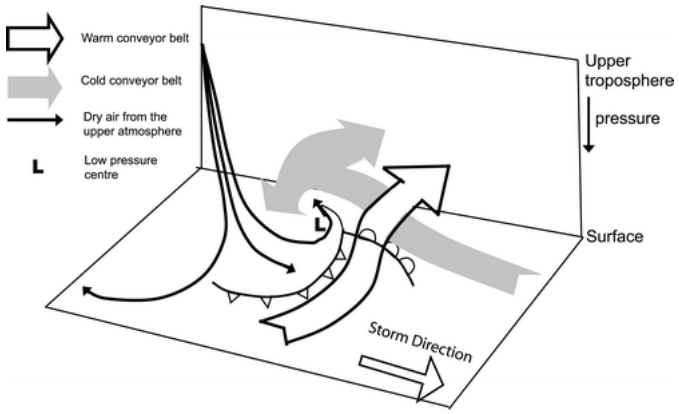

Figure 11: Conveyor belt model schematic [from *Catto et al.*, 2010, adapted from *Browning*, 1997].

If one focuses on heat transport, it is interesting to take a Lagrangian perspective of the mid-latitude dynamics described above, and track air flows. This view is usually applied to cyclones, and has a long history going back to *Bjerknes* [1910]. It was then developed into a quasi-Lagrangian picture, called "conveyor belt" model, by *Harrold* [1973], *Browning et al.* [1973], and *Carlson* [1980]. The basic conveyor belt picture has three air streams: a dry intrusion, a cold belt and a warm belt (WCB). As the name suggests, the WCB is the one driving the largest heat transport, and is characterised by the rapid ascent of warm, moist air. In defining it in precise physical terms, one can impose further constraints in terms of horizontal distance travelled over a given period of time [e.g. *Eckhardt et al.*, 2004]. The conveyor belt model is schematised in figure 11.

The WCBs are largely confined to areas of intense baroclinic activity, such as the storm tracks. These are defined as areas of large synoptic-scale baroclinic wave activity [*Blackmon et al.*, 1977], and are associated with large values of poleward heat transport. One can identify two distinct tracks in the NH (in the Atlantic and Pacific oceans) and a single, continuous band in the SH. Figures 12 and 13 show the geographical position of the NH and SH storm tracks, based on track density of low pressure anomalies. Note that this is only one of several measures that can be used to identify the storm tracks [e.g. *Hoskins and Hodges*, 2002]. It is worth mentioning, in a brief digression, that the oceans also have areas of very intense dynamics, referred to as western boundary currents (WBCs) [*Stommel*, 1948]. These are strong, warm currents found on the western boundaries of ocean basins; the two most important NH ones are the Gulf Stream in the Atlantic and the Kuroshio current in the Pacific. They play important roles in the oceans' meridional heat transport and in ocean-atmosphere

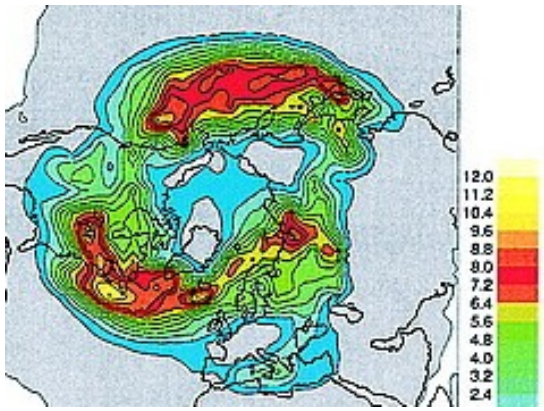

Figure 12: NH storm tracks based on track density of negative mean sea level pressure (MSLP) features [from *Hoskins and Hodges*, 2002].

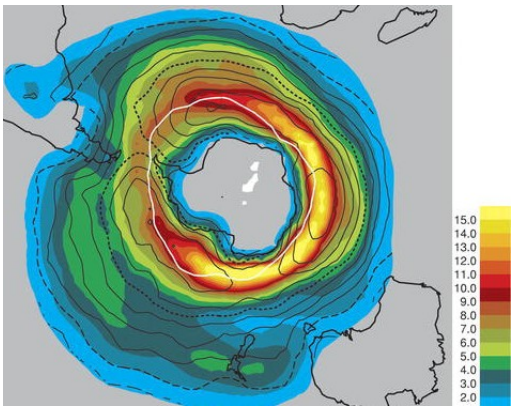

Figure 13: SH storm track based on track density of negative MSLP features [from *Hoskins and Hodges*, 2005].

interaction processes. It is by no means a coincidence that, in the NH, these currents partly match the position of the atmospheric storm tracks [*e.g. Sampe and Xie*, 2007; *Booth et al.*, 2010].

Even though they are traditionally associated with a large portion of the meridional heat transport, synoptic systems are by no means the only scale relevant to the dynamics of the mid-latitude atmosphere. At scales larger than those of individual pressure lows, one can identify fluctuations commonly called planetary waves. These are often influenced by topographical features of the planetary surface, such as mountain ranges, and by the location of the continents and oceans. When they become locked onto these features, they are called stationary waves [*McIlveen*, 1998]. Figure 14 shows

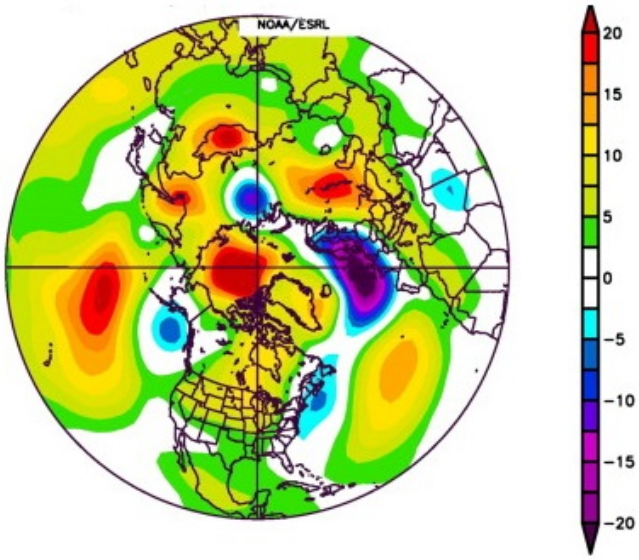

Figure 14: 700 *mb* geopotential height composite anomalies (in m) for NH JJA 1968-1996, showing a clear wavenumber 4 planetary pattern [from NOAA Physical Science Division, in *Wahl et al.*, 2012].

an example of planetary-scale spatial variability, namely 700 *mb* geopotential height anomaly for the NH during June, July and August. These low wave number modes provide the main contribution to long period variability [e.g. *Blackmon and Lau*, 1980] and, crucially, also have an impact on the amplitude and development of the smaller baroclinic waves [*McIlveen*, 1998].

When discussing the role of planetary scales for long periods, *Blackmon and Lau* [1980] define these as $10 < T < 90$ days. There are, however, modes of variability that act on even longer time scales. One of particular interest for the NH mid-latitudes in the North Atlantic Oscillation (NAO). This oscillation is generally defined in terms of the pressure difference between the semi-permanent Icelandic Low and the Azorres High [*Walker and Bliss*, 1932]. The strength of the NAO is quantified by a normalised index based on this pressure difference. For a positive phase of the NAO, the high pressure is higher than usual, and the low pressure lower. A negative phase of the NAO describes the opposite pattern. The sign of the NAO varies with an erratic periodicity, and can remain the same for over a decade [e.g. *Hurrell*, 1995]. The winter average NAO index for the period 1864-1994 is shown in panel a) of figure 15. The sea level pressure difference between strong positive and negative phases of the NAO (NAO > 1 and NAO < -1 respectively), is shown in panel b) of the same figure. The two pressure cores over the Azorres and Iceland are very evident.

The NAO has major large-scale climatic impacts. Positive phases are usually associated with cold weather in Greenland and mild winters in Europe. Negative ones drive the opposite pattern [*Peixoto and Oort*, 1992]. Figure 16 presents a schematic of the two NAO phases and the associated climatic impacts. As could be expected from its wide-ranging climatic influences, the NAO sign can also affect heat transport by altering moisture transports and the intensity, frequency and path of storms [*Hurrell et al.*, 2003].

1.3.5 The Energy View

As noted in Section 1.3.3, energy conversion is a crucial aspect of atmospheric instabilities. More generally, the circulation features described in the previous sections can be understood as a series of physical processes acting to convert energy; this view was first proposed by *Lorenz* [1955].

The uneven solar energy input received by the planet leads to a net heating in the tropical regions and a net cooling at the poles (see Section 1.1). This uneven radiative heating, and the resulting slope of the isentropes, is the primary source of zonal-mean APE, P_M. To obtain the overall APE in the atmosphere, the terms relating to the temporal and zonal variations in temperature must also be considered. These are commonly referred to as transient eddy APE (P_{TE}) and stationary eddy APE (P_{SE}). Here, we will refer to these two terms collectively as eddy APE (P_E). The overall APE, P is then given by:

$$P = P_M + P_{TE} + P_{SE} \tag{17}$$

Similar considerations can be applied to the kinetic energy terms, such that:

$$K = K_M + K_{TE} + K_{SE} \tag{18}$$

These six different forms of energy are then converted and dissipated by the rich range of dynamical behaviours displayed by the atmosphere. The generation of APE, as already mentioned, is largely due to the differential radiative heating between the equator and the poles. In terms of dissipation, a major contribution is provided by the decrease in kinetic energy via frictional and turbulent effects. The conversion processes are extremely varied, and only a brief overview is provided here.

The large-scale circulations cells, such as the Hadley and Ferrel cells, mainly act on the zonal

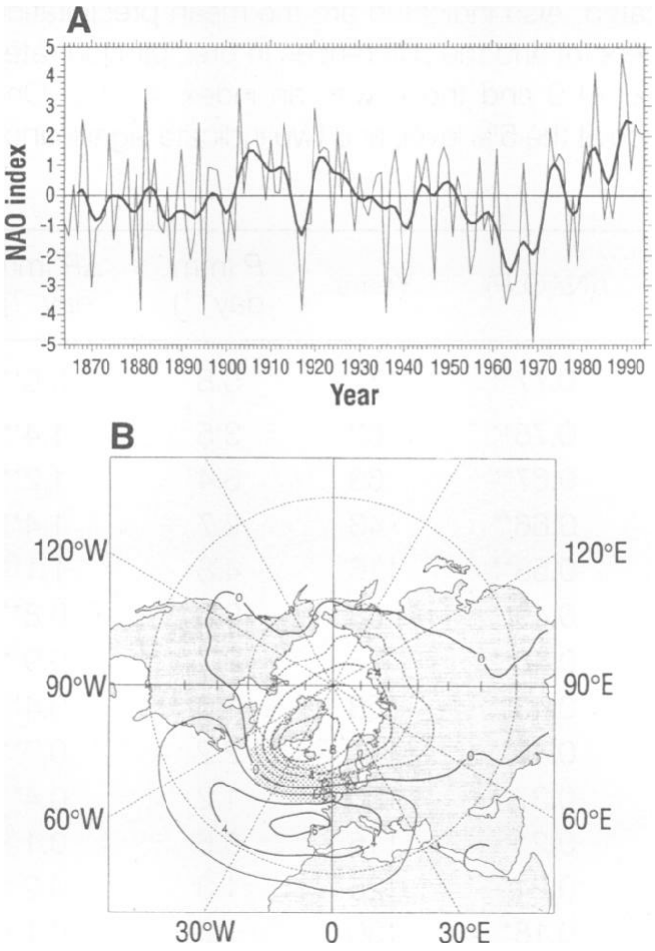

Figure 15: a) December-March NAO index based on the difference of normalized pressures between Lisbon, Portugal, and Stykkisholmur, Iceland, for the period 1864-1994. The heavy solid line represents the meridional pressure gradient smoothed with a low-pass filter which removes fluctuations with periods less than 4 years. b) Difference in SLP between years with an NAO index > 1.0 and those with an index < -1.0 (high index minus low index) since 1899. The contour increment is 2 *mb*; negative values are dashed [from *Hurrell*, 1995].

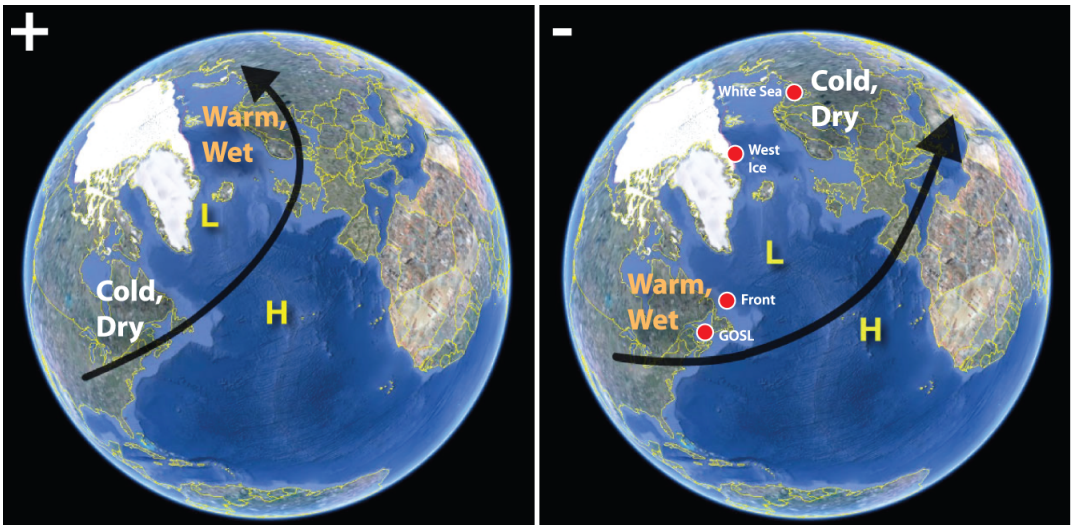

Figure 16: Pressure anomalies and associated climatic impacts for the two phases of the NAO [from *Johnston et al.*, 2012].

mean terms, converting P_M to K_M and *vice-versa*. In the case of the Hadley cell, rising warm air and subsiding cold air lead to a net conversion of APE to kinetic energy, while the opposite occurs in the Ferrel cell. Other motions convert the potential energy between its different forms: eddy motions can distort the isotherms, leading to a loss of zonal symmetry, and hence a decrease in P_M and and increase in P_E. The eddy APE can then be converted to eddy kinetic energy by baroclinic instabilities which, as discussed in Section 1.3.4, are a key driver of the meridional energy transport. Finally, one can consider the conversion of kinetic energy between its different forms. An example of this would be eddies increasing the angular momentum of the high-level jets, and hence converting eddy kinetic energy to zonal-mean one.

This highly simplified overview of the *Lorenz* energy cycle is summarised graphically in figure 17.

1.4 Meridional Heat Transport in the Atmosphere

As already mentioned in Section 1.1, meridional heat transport is a key component of the Earth's climate system and there is an extensive literature on the topic. The transport is the result of the imbalance in solar energy input at the TOA, and has both an atmospheric and an oceanic component. The ocean dominates in the low latitudes, while the atmosphere dominates in the mid and high

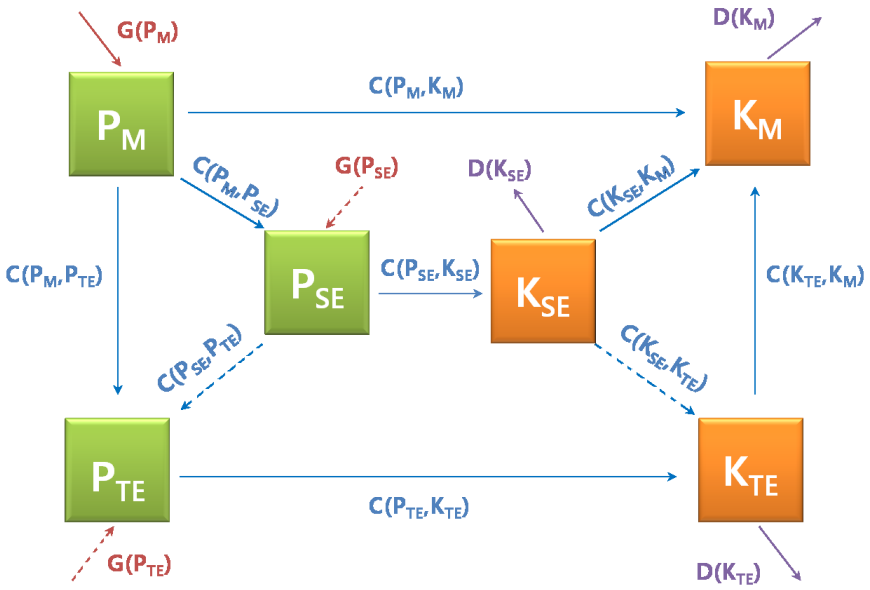

Figure 17: The Lorenz energy cycle in the atmosphere. G stands for generation, D for dissipation and C for conversion. The dashed arrows represent the terms not considered in the text [adapted from *Kim and Kim*, 2012].

latitudes. Figure 18 shows a recent estimate of meridional heat transport by the atmosphere and ocean, as a function of latitude. As can be seen, atmospheric transport peaks in the mid-latitudes at approximately 5 PW, while the corresponding contribution by the ocean is less than 1 PW.

There are indications that, at a very fundamental level, this partitioning of the transport is constrained. The most famous hypothesis in this respect is the Bjerknes compensation concept [*Bjerknes*, 1964], according to which changes in atmospheric and oceanic heat transports balance each other over long timescales. Simple dynamical arguments, however, can provide more detailed indications of the role of the atmosphere relative to the ocean, and suggest that the effect of these compensated changes on the partitioning of the transport is unlikely to be large [*Czaja and Marshall*, 2005].

1.4.1 The Different Components

Meridional atmospheric heat transport can be broken down into three different components. Part of it is effected by the mean background motions, part by the stationary features discussed in Section 1.3.4 and part by the transient eddies. The present thesis is concerned with the temporal variability

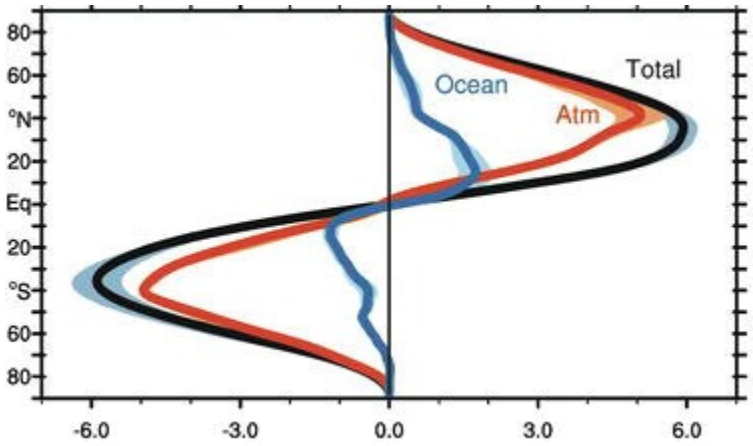

Figure 18: Annual mean meridional heat transport (in PW) by the atmosphere and ocean. The 2σ contours are also shown [from *Fasullo and Trenberth*, 2008].

of heat transport, and will focus on the time-dependent (transient) component. This component is generally associated with the growth and decay of mid-latitude weather disturbances, and is therefore directly linked to baroclinic instability [e.g. *Lau*, 1978]. Typically, the transport driven by the mean circulation dominates in the tropical regions (Hadley Cell), while the baroclinic instabilities dominante in the mid to high latitudes [*Peixoto and Oort*, 1992].

The transient component, τ, is computed as follows:

$$\tau = v'H', \tag{19}$$

where:

$$H = L_v q + c_p T + g z_g \tag{20}$$

H is a thermodynamic variable called moist static energy (MSE), L_v is the latent heat of vaporization of water and c_p is the specific heat capacity of air at constant pressure. The primes denote deviations from the long term average, and will be referred to hereafter as temporal anomalies. To understand where the definition of H originates from, it is instructive to consider dry static energy, s:

$$s = c_p T + g z_g \tag{21}$$

The term *static* is adopted because the above equation does not include a kinetic energy term. One can then describe the energetics of a large volume of air, in pressure coordinates, as:

$$\frac{\partial s}{\partial t} + \nabla \cdot s\vec{u} + \frac{\partial s\omega}{\partial p} = g\frac{\partial (F^R + F^s)}{\partial p} + Q_{LH} \tag{22}$$

Here ω is vertical velocity, p is pressure, F^R and F^S are the upward vertical fluxes of radiative energy and sensible heat and Q_{LH} is latent heating per unit mass [*Neelin and Held*, 1987]. Similarly, moisture conservation can be expressed as:

$$\frac{\partial L_v q}{\partial t} + \nabla \cdot Lq\vec{u} + \frac{\partial \omega L_v q}{\partial p} = -Q_{LH} + g\frac{\partial F^L}{\partial p}, \tag{23}$$

where F^L is the vertical flux of latent heat due to moisture diffusion. Q_{LH} now has a negative sign, to reflect the phase change of water vapour and its subsequent precipitation. Combining equations 22 and 23, one obtains:

$$\frac{\partial H}{\partial t} + \nabla \cdot H\vec{u} + \frac{\partial H\omega}{\partial p} = g\frac{\partial F}{\partial p}, \tag{24}$$

where F is the total energy flux. The large Q_{LH} term has cancelled with itself, and now the total energy flux is entirely in terms of the moist static energy and velocity fields. MSE is only affected by radiation and sensible heat exchanges at the boundaries [*Neelin and Held*, 1987].

1.4.2 Historical Overview

Some of the first works on meridional heat transport in the atmosphere were conducted by *Houghton* [1954]. Early estimates of its magnitude were based on simple parametrisations in energy balance models (EBMs) [*Budyko*, 1969; *Sellers*, 1969]. In addition to numerical estimates, there was also work on *in situ* measurements and variability [e.g. *Oort*, 1971]. This allowed testing the parametrisations commonly used in the models [e.g. *Stone and Miller*, 1980], and to perform detailed analyses on the individual components of the transport. *Peixoto and Oort* [1992], for example, looked at the potential energy, kinetic energy and sensible heat components, while *Lau and Wallace* [1979] separated the divergent, non-divergent, rotational and irrotational fluxes. Often, data and models were combined to obtain optimal estimates, such as in *Trenberth and Caron* [2001]. More recent estimates are based on satellite measurements, which provide a comprehensive global coverage [*e.g. Fasullo and*

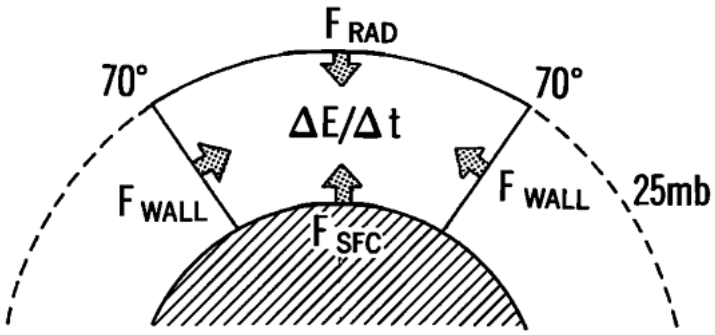

Figure 19: Diagram of the polar energy budget [from *Nakamura and Oort*, 1988].

Trenberth, 2008; *Porter et al.*, 2010].

Most of the above studies focussed on numerical estimates of the transport; others have focussed on its structure and future forecasts [*e.g. Hwang and Frierson*, 2010]. Yet others, such as *Nakamura and Oort* [1988], have focussed on its effect on the polar energy budget. In their 1988 paper, the authors combined GCM output with *in situ* and satellite observations to obtain a full energy budget for the polar caps. The simple schematic used was to consider the polar caps as being delimited by imaginary walls circling the globe at 70° N and 70° S (see figure 19). The energy balance was then defined according to the following equation:

$$\frac{\triangle E}{\triangle t} = F_{rad} + F_{wall} + F_{sfc}, \tag{25}$$

where the left hand side represents the change in stored energy in the atmosphere, F_{rad} is the net incoming TOA radiation, F_{wall} is meridional atmospheric heat transport into the polar cap and F_{sfc} is the energy flux from the surface into the polar cap. It should be noted that this schematic is centered on the atmosphere: both caps extend from the surface to the 25 mb pressure level. The ocean is not considered to be part of the polar caps, and ocean to atmosphere fluxes are incorporated into F_{sfc}. This simple approach is still very relevant, since the very delicate balance of Arctic sea-ice has changed in recent years, and atmospheric heat transport is a crucial energy input for the polar regions [*e.g. L'Heveder and Houssais*, 2001; *Alexeev et al.*, 2005].

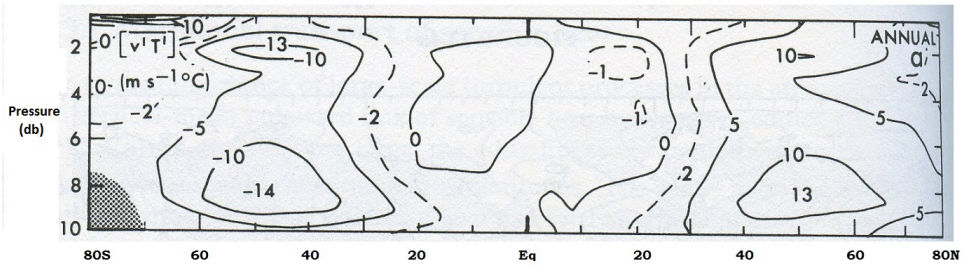

Figure 20: Annual and zonal mean cross section of northward transport of sensible heat by transient motions (in $^{\circ}Cms^{-1}$) [from *Oort and Peixoto*, 1983].

1.4.3 Main Results

As discussed above, simple geometrical constraints can provide powerful insights into the structure of heat transport. *Stone* [1978], for example, assumes that each hemisphere is in radiative equilibrium and that the atmosphere and ocean are dominated by the planetary scale. Applying these two assumptions to one-dimensional EBMs, the author finds that the net meridional energy transport must peak near 35° latitude. Furthermore, the Earth-Sun geometry and the Earth's albedo constraint the peak magnitude of the flux to be around 5 PW. A comparison of these results with figure 18, from a recent work by *Fasullo and Trenberth* [2008], reveals a very good agreement both in terms of position of the maximum flux and peak magnitude. Surprisingly, a one-dimensional EBM and very simple considerations on the geometry of our planet are therefore able to reproduce some key features of the meridional heat flux structure computed with the aid of modern satellite data.

These idealised setups, however, cannot provide physical details of the individual components of the transport in the real atmosphere. The transport by transient eddies is typically found to peak at 850 mb, with a secondary maximum around 200-300 mb [*Lau*, 1978;*Peixoto and Oort*, 1992]. Figure 20 shows the annual and zonal mean cross section of northward heat transport by transient motions. The two peaks discussed above are clearly visible in both hemispheres. It should, however, be noted that the figure only includes the sensible heat component of the transport.

These features are subject to significant temporal and spatial variability. The net atmospheric transport peaks in December in the NH and June in the SH, at values of approximately 6 PW and 5 PW respectively. The summer transport in the SH at 30° S is larger than the corresponding NH one by about 1 PW. These patterns are shown in figure 21, which displays the zonally averaged

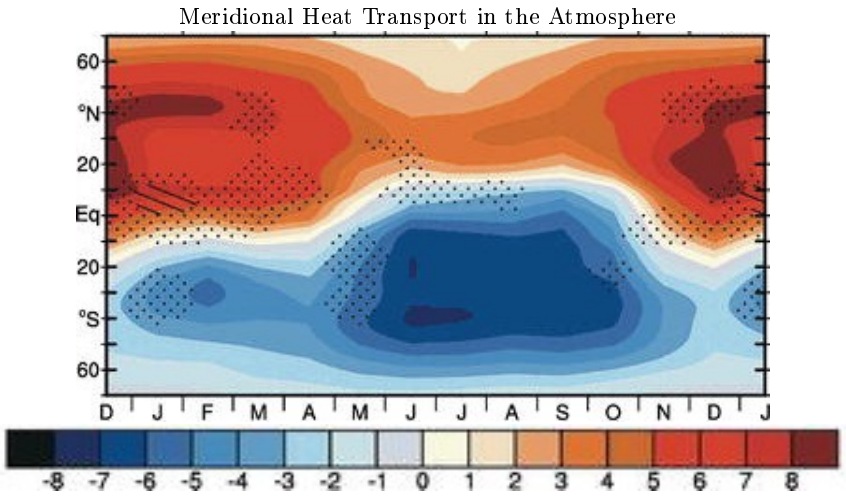

Figure 21: Meridional atmospheric heat transport as a function of latitude and month (in PW) [from *Fasullo and Trenberth*, 2008].

meridional atmospheric heat transport as a function of latitude and month. Notwithstanding the pronounced annual cycle, poleward of about 30° N and S the transient eddies play an important role year round, and their variability matches reasonably closely the one of the net transport [*Peixoto and Oort*, 1992]. Figure 22 shows a comparison of the zonal and vertical mean profiles of meridional heat transport for the three different components and the overall atmosphere. The contribution due to transient motions is displayed in panel b). Again, like in figure 20, only the sensible heat component of the transport is included.

The spatial variability of the transport is as pronounced as the temporal one. Figure 23 shows the global distribution of the vertical and yearly mean northward transport of sensible heat by transient motions. In the SH, there is a uniform band of maxima across the Southern Ocean, with a local peak. In the NH, a more complex pattern emerges. There are maxima over North America, the Gulf Stream and Eastern Asia, associated with the polar front. Other maxima are present over Fennoscandia and the Norwegian sea [*Peixoto and Oort*, 1992].

It should be noted that the estimates for atmospheric heat transport are by no means faultless. Older studies suffer from a severe lack of data [*e.g. Oort*, 1971; *Nakamura and Oort*, 1988], and even the most recent estimates are affected by instrument calibration issues [*Porter et al.*, 2010] and the limitations of reanalysis data [*Fasullo and Trenberth*, 2008]. *Porter et al.* [2010] highlight that

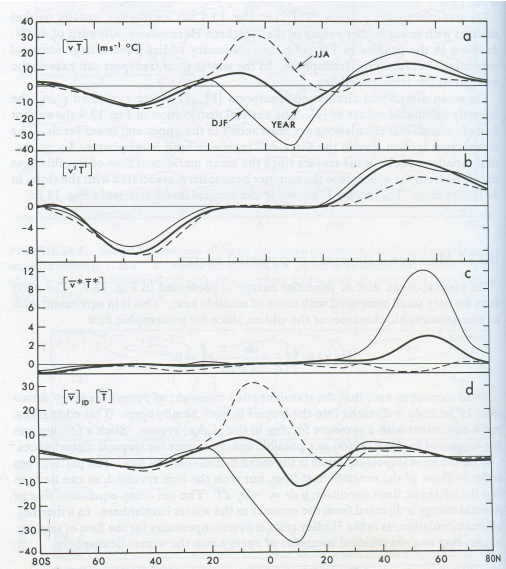

Figure 22: Zonal and vertical mean profiles of meridional transport of sensible heat due to a) all motions; b) transient eddies; c) stationary eddies; d) mean background circulation (in $^{\circ}Cms^{-1}$) [from *Oort and Peixoto*, 1983].

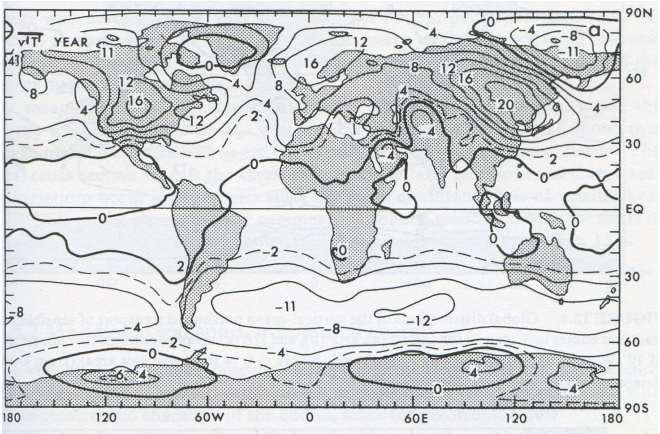

Figure 23: Global distribution of the vertical and yearly mean northward transport of sensible heat by transient motions (in $^{\circ}Cms^{-1}$) [from *Oort and Peixoto*, 1983].

CERES satellite data has shortcomings, mainly arising from instrument calibration, which yield an estimated 6.5 Wm^{-2} global mean net TOA radiative imbalance. This is a surprisingly large value when compared, for example, to the TOA forcing ascribed to anthropogenic influence, which is quantified in 0.6 Wm^{-2} to 2.4 Wm^{-2} [*Forster et al.*, 2000]. *Porter et al.* also find significant discrepancies between reanalysis products from the National Centre for Environmental Prediction/National Centre for Atmospheric Research and from the Japan Metereological Agency. In particular, a large source of error seems to derive from applying a mass correction to these datasets which, in their original form, are not mass-conserving. This translates into differences in estimates of atmospheric meridional heat transport at 70° N of up to 32 Wm^{-2}. Climate models also encounter problems. *Hwang and Frierson* [2010] find that surface albedo and cloud radiative effects, all contribute significantly to model spread in forecasted meridional atmospheric heat transport at 40° N.

The above discussion suggests that, while significant improvements in available data have occurred, this is still far from perfect. There is also ample scope for further research in terms of data analysis and modelling. *Hwang and Frierson* [2010] maintain that a key contribution to better heat flux modelling will come from constraining uncertainties in surface albedo and cloud effects. The problems encountered by *Porter et al.* [2010] in their analysis indicate that better methods for correcting mass imbalance in reanalysis products are needed, since the current ones encounter problems when applied to the high latitudes.

2 Conceptual Models of Heat Transport

2.1 The Eady Model of Baroclinic Instability

The following discussion is largely based on *Gill* [1982], to which the reader is referred for a more detailed illustration of the Eady model.

The Eady model [*Eady*, 1949], is the simplest model of baroclinic instability. It deals only with the development of the initial disturbance, and is not concerned with its full life cycle. The starting point is a free wave in a horizontal temperature gradient. It is assumed that there are no topographical or frictional effects, that the atmosphere is incompressible and that the Coriolis force is constant across the domain (f-plane approximation). Eady's great intuition was to apply the rigid lid approximation to this setup. This approximation, as the name suggests, assumes that the tropopause acts as a

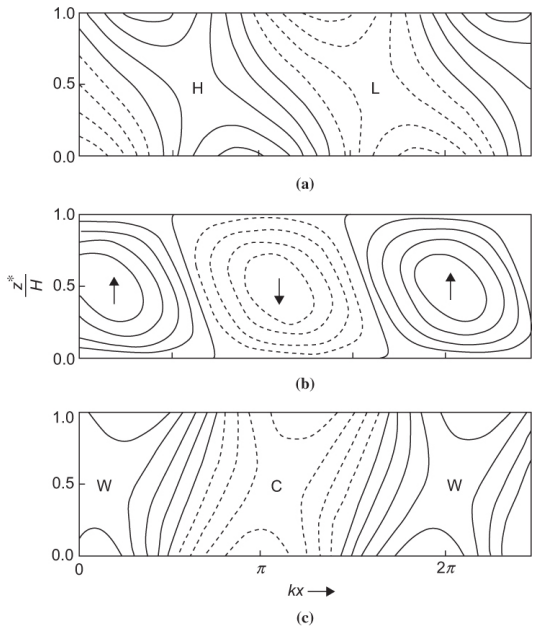

Figure 24: The structure of the most unstable Eady wave in the $x - z$ plane. Panel a) displays geopotential height anomaly contours; "H" and "L" indicate the positions of the ridge and trough, respectively. Panel b) displays contours of vertical velocity; the arrows indicate the positions of strongest ascending and descending motions. Panel c) displays temperature anomaly contours; "W" and "C" indicate the positions of the warmest and coldest temperatures, respectively [from *Holton and Hakim*, 2012].

rigid upper boundary for wave development. Its main physical limitation is that it implies that spatial scales are set by the tropopause height. Nonetheless, it provides a very clear wave picture of phase-driven heat transport by baroclinic instabilities. Figure 24 shows the properties of the most unstable Eady wave. Panel a) displays geopotential height anomaly contours; "H" and "L" indicate the positions of the ridge and trough, respectively. Panel b) displays contours of vertical velocity; the arrows indicate the positions of strongest ascending and descending motions. Panel c) displays temperature anomaly contours; "W" and "C" indicate the positions of the warmest and coldest temperatures, respectively [*Holton and Hakim*, 2012].

The velocity and geopotential contours tilt westward with height, implying that there is a phase shift between the top and bottom boundary waves. The temperature anomaly, on the opposite, tilts eastward. A mathematical analysis of the most unstable (or fastest growing) Eady solution finds that the phase difference between the isobar and isotherm contours at the two boundaries is approximately 21°. Moving westward, temperature leads pressure at the top boundary and lags

pressure at the bottom one. This implies that, at mid-levels, there must be a point where the warmest air is also the fastest moving one, meaning that v and θ' are perfectly in phase. It is further found that the phase shift of the high and low pressure centres between the top and bottom boundaries is of 90°, and that for temperature contours the same shift is of 48°. It is because of the phase difference between temperature and pressure that the Eady wave effects a net heat transport. If the pressure and temperature anomalies were perfectly in phase, there would be no poleward flow where the air is warmest and there would be the strongest flow where the temperature anomaly is zero. For the Eady wave, on the opposite, the poleward air flow corresponds to a warm anomaly and the equatorward flow to a cold one. Indeed, at all levels the poleward motion corresponds to warmer air than the equatorward one. Mass balance is therefore maintained, but the net effect is of a positive poleward heat transport. The importance of the Eady model lies in the fact that it can be easily solved analytically. Further mathematical details are provided in the Appendix.

2.2 The Charney and the Phillips Models

While *Eady* was developing his idealised framework, another model describing baroclinic instability was being studied by *Charney* [1947]. In this model, compressibility effects in the atmosphere are considered, and the rigid lid approximation is discarded in favour of more realistic boundary conditions. The β-effect, namely the linear variation of the Coriolis parameter with latitude, is also taken into account. The final result is surprisingly similar to the one from the Eady model. Figure 25 shows the properties of the most unstable Charney wave. Panel a) displays contours of vertical velocity. Panel b) displays contours of meridional velocity (solid lines) and potential temperature anomalies (dashed lines). “H” and “L” indicate the positions of the geopotential ridge and trough, respectively. They also mark the zero contours of velocity. "W" and "C" indicate the positions of the warmest and coldest airflows at the two boundaries. The small panel on the bottom right hand side displays the variation of poleward heat flux with height. Notwithstanding the many analogies, comparing this result to Eady's one, shown in figure 24, highlights important differences in the phase of the waves. The phase difference bewteen the isobars and isotherms at the surface is now 41°. This then decays to zero at infinity, implying that the meridional heat transport also falls off to zero. It should also be noted that this model is free from the artifical constraint seen in Eady's case, where the spatial scales were set by the tropopause height.

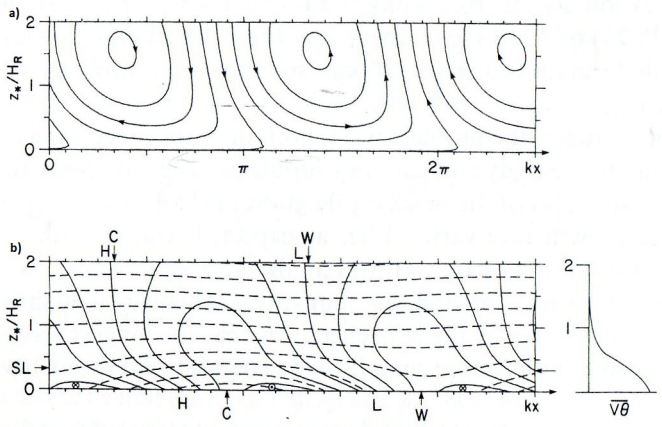

Figure 25: The structure of the most unstable Charney wave in the $x - z$ plane. Panel a) displays contours of vertical velocity. Panel b) displays contours of meridional velocity (solid lines) and potential temperature anomalies (dashed lines). "H" and "L" indicate the positions of the geopotential ridge and trough, respectively. They also mark the zero contours of velocity. "W" and "C" indicate the positions of the warmest and coldest airflows at the two boundaries. The small panel on the bottom right hand side displays the variation of poleward heat flux with height. Note that, unlike in the Eady case, the flow does not end at the upper boundary shown in the figure [from *Gill*, 1982].

An instructive depiction of the baroclinic instability process can also be obtained by considering a simple two layer model. Here, the initial setup consists of two homogeneous fluids, with different densities and mean velocities. The original formulation, illustrated in figure 26, adopts Eady's rigid lid approximation. The baroclinicity, namely the density gradient mentioned in Section 1.3.3, is provided by the deformation of the interface between the two fluids. This model was first studied in detail by *Phillips* [1951], and later extended by other authors to include additional effects, such as the β-effect and bottom slope [e.g. *Gill et al.*, 1974]

2.3 Baroclinic Growth on a Sphere and Life Cycle of a Baroclinic Disturbance

Both the *Eady* and the *Charney* models consider approximations to the full spherical geometry, namely the f-plane and the β-plane. Moreover, in these idealised models, the zonal flow is assumed to be independent of latitude. Studies considering a full spherical geometry and a meridionally varying zonal flow lose the mathematical simplicity of *Eady*'s setup, but have a much closer link to the real dynamics of the atmosphere. *Simmons and Hoskins* [1976], for example, consider the properties of

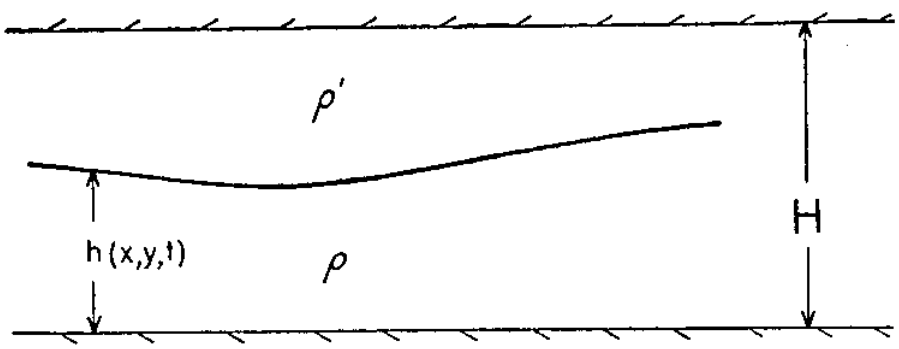

Figure 26: Vertical cross-section of the two-layer Phillips model. The densities of the top and bottom layers are ρ and ρ', respectively. $h(x, y, t)$ is the thickness of the lower layer; H is the constant separation between the top and bottom rigid boundaries [from *Phillips*, 1951].

unstable modes in three zonal flows on a sphere. The jets have the same vertical structure, but differ in their latitudinal profiles. The growth rates of the most unstable modes are found to be similar to those yielded by the more idealised models described above, highlighting the lack of sensitivity of the instability mechanism to the meridional profile of the flow. However, such profile is found to influence the location, meridional scale and horizontal phase characteristics of the disturbance. In order to evaluate the physical significance of the instabilities for the real atmosphere, the additional complexity of the *Simmons and Hoskins* model is therefore crucial[1].

A key limitation of the models described in Sections 2.1 and 2.2 is that they are only concerned with the characteristics of the most unstable mode, but do not consider its temporal evolution. Clearly, an instability cannot grow indefinitely and the details of its growth, interaction with other atmospheric flows, and decay are crucial for understanding its role in the circulation. An instructive example of the life cycle of baroclinic disturbances is discussed in *Simmons and Hoskins* [1978], which considers the evolution of a zonal jet centred at 45° N and 200 *mb*. The zonal flow is initially balanced, and is then perturbed by small-amplitude (maximum 1 *mb*) anomalies. For the case of a disturbance corresponding to the fastest-growing mode of wavenumber 6, the instability develops rapidly by drawing on the mean-flow APE. By the seventh day of the simulation, the disturbance has grown to an amplitude of 32 *mb*, and displays realistic synoptic frontal structures, as shown in figure 27. The upper-level streamfunction displays a growth of the disturbance up to the tenth day of the

[1]The same study also evaluates the accuracy of the *quasi*-geostrophic approximation for the study of baroclinic disturbances. Although this will not be discussed here, it has important implications for the linear theory of baroclinic instability.

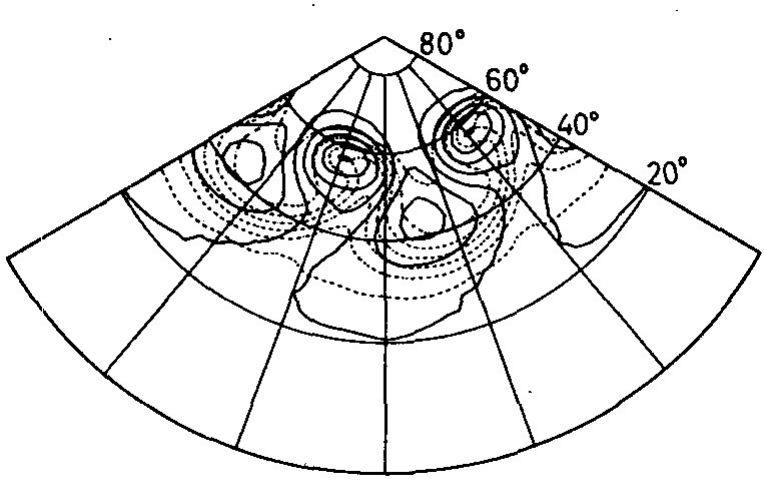

Figure 27: Surface pressure (solid) and low-level temperature (dashed) profiles for a wavenumber 6 perturbation on a zonal jet at 45° N. The picture refers to the perturbation seven days after the initial anomaly was seeded [from *Simmons and Hoskins*, 1978].

simulation, after which there is a rapid decay in wave amplitude and a barotropic strenghtening of the jet. Near the surface, the wave suppresses the meridional temperature gradient between 40° and 60° N, while it enhances it further North and South. In the upper troposphere, on the opposite, very little alteration to the gradient occurs because the effect of the wave is balanced by the strengthening of the upper level jet by momentum fluxes. From the energetics perspective, the maximum baroclinic growth and eddy heat flux occur at upper levels, where the amplitudes are significantly larger than at the surface. The relative difference between upper level and near-surface fluxes is larger than what seen in linear stability analyses. While this might seem like a small difference, it is crucial for a correct parametrization of baroclinic disturbances. The wealth of new information the life cycle provides when compared to *Eady*-type approaches is also evident.

2.4 The Swanson and Pierrehumbert Stochastic Model

In their analysis of lower tropospheric heat transport in the Pacific storm track, *Swanson and Pierrehumbert* [1997] studied the PDFs of meridional atmospheric heat flux at three points in the track. The distributions are shown in figure 2 on page 19. It should be noted that, for these distributions, heat transport τ is computed using a slightly different method to the one adopted here:

$$\tau = v\theta' \tag{26}$$

The prime denotes the anomaly relative to a 30-day running mean. The authors noticed that the distributions displayed an extended positive tail associated with a large skewness, and developed a simple model aimed at reproducing this feature.

As initial step, they constructed a Lagrangian mixing climatology for the 40° N latitude band within the Pacific storm track. This consists of the passive advection of parcels, originating in the track, by the 850 *mb* horizontal wind fields. Such climatology clearly provides no insight into role of thermal interactions with the surface. The purely Lagrangian heat transport, while capturing some of the features of the observed transport, was unable to reproduce the finer structure of the measured values. The authors therefore developed a simple stochastic model, combining both Lagrangian passive advection and thermal interactions, to investigate the role of non-conservative thermal processes.

The model framework is as follows. Velocity is expressed by:

$$v(t + \Delta t) = v(t) - \frac{v(t)\Delta t}{T} + (2V'^2 \frac{\Delta t}{T})^{1/2}\xi, \tag{27}$$

where Δt is the model timestep, T is a characteristic timescale, and ξ is a white noise stochastic forcing. To model the thermal effects, the authors include a temperature equation:

$$\theta(t + \Delta t) = \theta(t) - \Delta t \frac{[\theta - y(t)\partial_y \Theta]}{\Gamma}, \tag{28}$$

where $y(t)$ is the particle position, and Γ and $\partial_y \Theta$ are two free parameters chosen such that the stochastic model correctly reproduces the standard deviation of θ and the mean of τ, as observed within the storm track domain. Γ is interpreted as a damping time for temperature fluctuations.

The model described above successfully reproduces the distributions found in the data (shown in figure 2 on page 19). This implies that both the thermal and dynamical processes are correctly captured, and highlights the important role of thermal damping, which was not included in the initial Lagrangian climatology [*Swanson and Pierrehumbert*, 1997].

2.5 The Heton Model

So far, the traditional wave picture of atmospheric dynamics has beeen presented. However, drawing inspiration from other fields of Physics, where wave phenomena are also prevalent, it is clear that this is not the only paradigm that can be adopted. For example, following the concept of quantisation commonly adopted in Quantum Mechanics, it is very stimulating to think of heat transport as a discrete process effected by particle-like carriers. The Heton model, first developed by *Hogg and Stommel* [1985a], is based on this concept. The model is a two-layer system, seeded with point geostrophic vortices. These give rise to a deformation of the interface between the two layers, whose nature depends on the sign of the vortices. The two layers have different potential temperatures; for a mock atmosphere or ocean, the upper layer would have a higher temperature than the lower one. Deformations in the interface can therefore be likened to thermal anomalies. Figure 28 shows the deformations resulting from having one vortex in each layer. If the two vortices have the same sign (right hand panel), they will correspond to opposite layer deformations. The one in the upper layer (labelled H) will drive a warm anomaly; the one in the bottom layer (labelled C) will drive a cold one. This will lead to a barotropic structure which, as discussed in Section 1.3.3, results in no net heat transport. If, on the contrary, the two vortices have opposite signs (left hand panel), they will drive matching temperature anomalies. For the case shown in the figure this is a warm anomaly (downward deformation), but it could equally well be a cold anomaly, if the signs of the two vortices were to be inverted. This structure is now baroclinic, and can therefore result in a net transport of heat. Such a pair is termed Heton, and is described by *Hogg and Stommel* [1985a] as "two discrete baroclinic geostrophic vortices in a two-layer system". If numerous hetons are seeded with sufficiently small separations, the different vortices will interact and displace one another. The outcomes of simple cases with few vortices can be easily predicted and interpreted. When larger numbers of vortices are considered, however, the behaviour becomes more chaotic and discontinuous. Several vortices can coalesce to form composites, which then propel themselves away from the other hetons. *Hogg and Stommel* [1985b] witnessed this behaviour while performing experiments with clouds of hetons, and termed it *hetonic explosion*. An example of a hetonic explosion is illustrated in figure 29. The hetons are initially seeded in a regular array (top panel). Instabilities start to develop (second panel), and the individual pairs coalesce into larger groups (third panel). These then propel themselves away from the other vortices (fourth panel). The Heton Model will be described in greater detail in

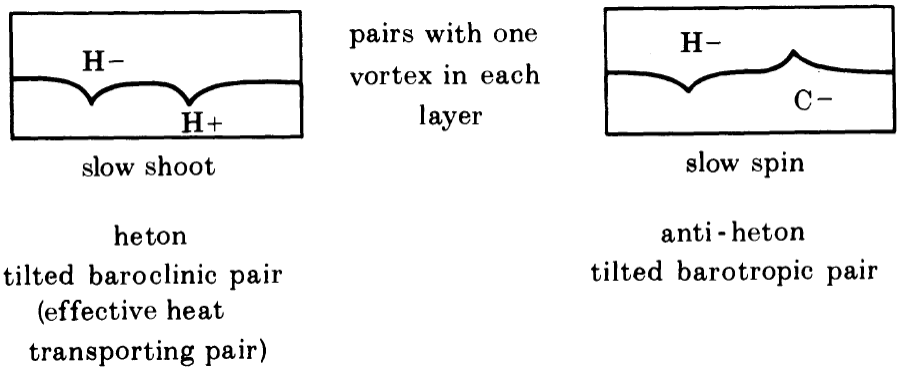

Figure 28: Interface deformations resulting from point geostrophic vortices in a two-layer model. The left hand panel depicts a heat-transporting pair. The right hand one depicts a pair which effects no net transport of heat [from *Hogg and Stommel*, 1985a]

chapter V.

3 Data Overview

3.1 ERA-Interim Reanalysis

ERA-Interim is a reanalysis product developed by the European Centre for Medium-Range Weather Forecasts (ECMWF). The data series starts from 1989 and is currently ongoing. By the term reanalysis, what is intended is a model simulation of the global atmosphere (or ocean, depending on the product), constrained by selected observations. The resulting data set aims to be the best possible estimate of the state of the atmosphere (ocean) at any given time. For ERA-Interim, the observations are assimilated on a 12-hourly basis, and also include satellite data.

The data set used in the present study has a horizontal resolution of approximately $0.75° \times 0.75°$ (more precisely, T255), and 37 vertical levels ranging from 1000 *mb* to 0 *mb*. The data is a daily snapshot at 1200 UTC. It should be noted that different spatial and temporal resolutions of the same data set are available. It should also be pointed out that reanalysis data presents important limitations. Firstly, mass is not conserved [e.g. *Graversen et al.*, 2011]. Furthermore, ERA-Interim assumes a perennial ice cover beyond 82° N; this is a potentially serious problem for more recent years. For a more in-depth discussion of the ERA-Interim data set, the reader is referred to *Simmons et al.* [2007].

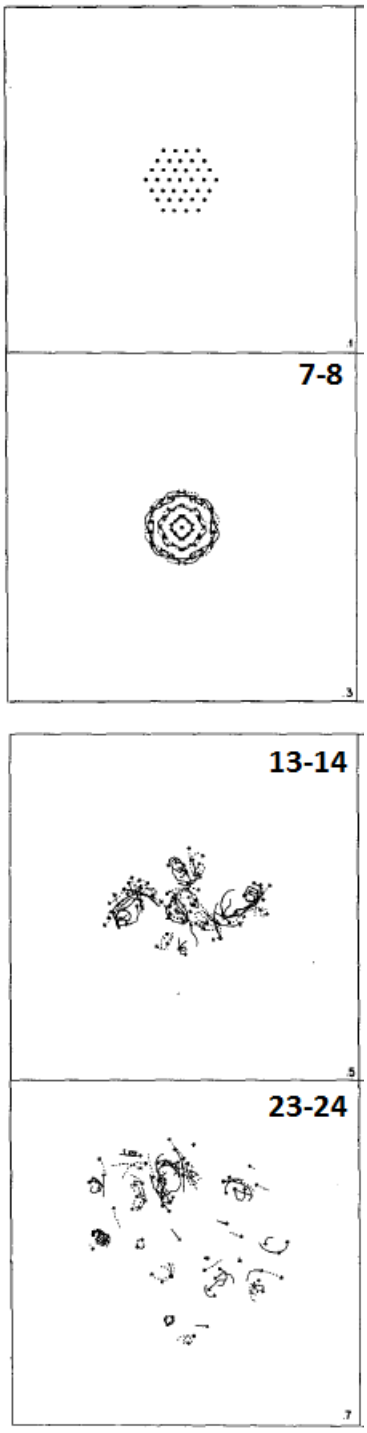

Figure 29: The evolution of a cloud of 37 hetons. The pairs are initially seeded in a regular hexagonal lattice. The lines show the trajectories of the individual vortices. The values in the top right corner of each panel indicate the timesteps depicted therein [from *Hogg and Stommel*, 1985b].

3.2 FORTE Model

FORTE (Fast Ocean Rapid Troposphere Experiment) is a coupled ocean-atmosphere model (*Sinha and Smith*, 2002; *Smith and Gregory*, 2009). It was developed in order to have a coarse resolution, computationally inexpensive model which retains all the critical physical processes. Its key components are the MOMA ocean model (*Webb*, 1996; *Pacanowski et al.*, 1991) and a modified version of the IGCM3 atmosphere model (*Hoskins and Simmons*, 1975; *Forster et al.*, 2000). The two modules are fully coupled on a daily basis. The atmospheric component includes a multi-band radiation scheme, a convective adjustment scheme and a dedicated cloud scheme. The latter, based on a simplified version of *Slingo* [1987], has three cloud levels and can simulate deep convective clouds.

The data set used in the present study is from a simulation with a horizontal resolution of approximately $2.8° \times 2.8°$ (more precisely, T42), and 15 vertical levels. The geographical setup is highly idealised, and is an aquaplanet configuration with a flat ocean bottom and polar islands beyond 78° N and S. These are needed to relax the timestep constraints deriving from grid convergence at the poles. Except for the degraded vertical and enhanced horizontal resolution, the simulation described here matches the one in *Smith et al.* [2006].

In terms of the heat transport, it is interesting to verify whether FORTE can reproduce the observed meridional transport profile. The model profile is shown in figure 30 on the next page. As can be seen, the atmosphere curve is reasonably close the the real world in the mid and high latitudes, while it is unrealistic in the equatorial region. Since the present study will focus on domains beyond 30° N and S, this is not considered to be a limiting factor. From the point of view of dynamics, in an aquaplanet configuration the statistics are only a function of latitude and height, and large-scale waves driven by mid-latitude westerlies dominate the circulation. Mesoscale motions, stationary waves and boundary features of the atmosphere are not captured by the simulation.

3.3 Minobe *et al.* Model

As discussed in section 1.3.4, WBCs are crucial areas for oceanic heat transport and for ocean-atmosphere interactions. In the NH, the Pacific and Atlantic WBCs, namely the Gulf Stream and the Kuroshio, match quite closely the location of the storm tracks, and are also areas of strong meridional atmospheric heat transport. In these regions, it is therefore important to understand to what degree, and by what mechanisms, the ocean can affect the atmosphere. In order to investigate

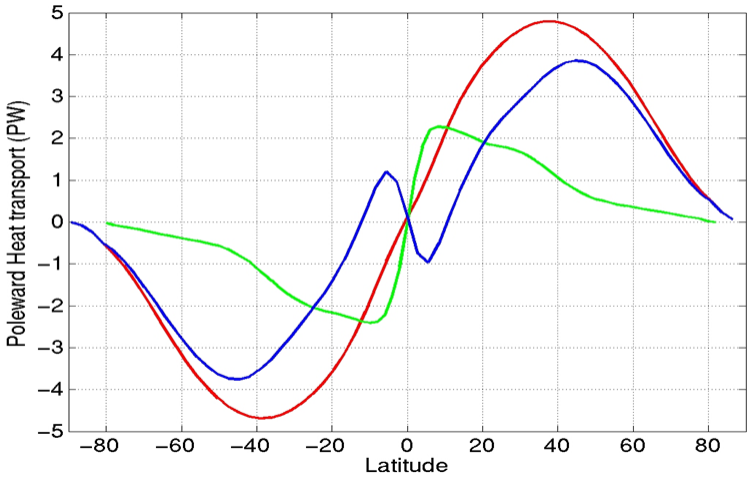

Figure 30: Yearly and zonal mean meridional heat transport by the atmosphere and ocean (in PW) in the FORTE aquaplanet simulation. The blue, green and red curves represent atmospheric, oceanic and total heat transport respectively [by *J. Cheung*, Imperial College, London].

this open question, *Minobe et al.* [2008] developed a model with both realistic and smoothed sea surface temperatures (SSTs). WBCs, in fact, are typically characterised by strong SST gradients.

The starting point is the output from an atmospheric general circulation model (AGCM), with a horizontal resolution of approximately 50×50 km (more precisely, T239) and 48 vertical levels. The period considered is 1989-2001, and 6-hourly data is provided. Two domains are defined: one over Gulf Stream and the other covering the Kuroshio current (30°-50° N and 280°-340° E; 30°-50° N and 130°-200° E respectively). Smoothing is then applied to the SST field in the regions bounded by 280°-330° E and 130°-180° E respectively. The smoothing is performed by applying a running-mean low-pass filter. Further discussions on filters are presented in Chapter IV and in the Appendix.

Part III

Extreme Events in Atmospheric Heat Transport: Definition and Role

1 Aims of the Chapter

The present chapter is largely based on *Messori and Czaja* [2013b]. It aims to:

i) Define extreme events in meridional atmospheric heat transport;

ii) Establish their importance;

iii) Outline their salient features.

A statistical analysis of such events is presented, in order to investigate possible dynamical drivers. A detailed physical interpretation of the heat transport extremes, dealing with atmospheric analogues and placing the events in the context of atmospheric dynamics, will instead be discussed in Chapter IV.

2 Outline

As discussed in Chapter II, meridional heat transport is a crucial component of the climate system. Although there is a vast literature on the transport's temporal variability, very little attention has been dedicated to its statistical distribution. A notable exception is the study by *Swanson and Pierrehumbert* [1997], whose idealised heat transport model was discussed in Chapter II, Section 2.4 on page 51. The latter paper, however, was very limited in scope – both geographically and temporally. The authors considered distributions of data at three individual points in the Pacific storm track, over the November-March season.

In the present chapter similar distributions are presented for broad domains in both the Northern Hemisphere (NH) and the Southern Hemisphere (SH), and for both DJF (December, January and February) and JJA (June, July and August) periods (Section 4.1). After robustly establishing the

sensitivity of the heat transport to extreme events, its sporadic and irregular nature is confirmed by showing that a large contribution to the transport arises from a few isolated bursts every season (Sections 4.2 and 4.3). The focus will be on low levels, at which the heat transport by transient motions is strongest [*Peixoto and Oort*, 1992]. These findings are then tested across a broad range of models, including idealised simulations and fully coupled atmosphere-ocean general circulation models (AOGCMs) (Section 4.4). Having ascertained the importance of the extremes, they are then characterised in terms of their temporal and spatial extents and geographical distribution (Section 5). An idealised framework, based on a wave picture, is also discussed in order to interpret the physical origin of the sporadic nature of heat transport (Section 6). Finally, Sections 7 and 8 present the physical insights and general conclusions that can be gained from the above analysis.

3 Methodology

The analysis in the present chapter is performed on ERA-Interim data, already discussed in Chapter II, Section 3.1 on page 54. Daily outputs (1200 UTC) are considered over a period spanning from December 1993 to August 2005, thereby providing twelve DJF and twelve JJA time series. The latitude and longitude resolution is approximately 0.75°, and the analysis focuses on the 850 mb pressure level. This is the level of peak transport by transient eddies, often used as reference level in the literature [e.g. *Lau*, 1978]. Results for all vertical levels in the data set are also shown, in order to demonstrate that the 850 mb analysis does indeed provide a good indication of the statistics of the full vertical transport. Part of the analysis is also repeated on FORTE output and on the Minobe *et al.* model. Further details of these are provided in Chapter II, Sections 3.2 on page 56 and 3.3 on page 56, respectively.

Transient eddy heat transports are computed as a product of meridional velocity (v) and MSE (H) temporal anomalies. These are defined as departures from the linearly detrended seasonal mean, and are denoted by a prime. They are computed at every grid point for 172 latitude bands between 30° N and 89° N, and 30° S and 89° S. No other time filtering is applied to the data. Part of the analysis was also repeated on non-detrended data, and no significative qualitative differences were found. In the figures described below, H is always given in Kelvin, after division by the specific heat capacity of dry air (taken to be 1005.7 $JK^{-1}kg^{-1}$), and velocity is positive polewards in both hemispheres.

It should be noted that, except where explicitly mentioned in the text, no area or gridbox width weighting is performed in computing the transport, meaning that the relative importance of events at high latitudes might be exaggerated.

The PDF of a continuous random variable is a function describing the relative likelihood that the aforementioned variable has a value within a given interval, termed a *bin*. The values of v', H' and $v'H'$ are binned and the resulting distribution is normalised, so that one can relate the number of events in each bin to a percentage of the overall events. Taking the product $v'H'$ for all data points and repeating the binning process yields the desired PDFs for transient-eddy heat transport. A key element in the analysis of the PDFs is skewness: a measure of the asymmetry of a distribution or, more formally, the distribution's third standardised moment. For a random variable, X, this is defined mathematically as:

$$\varsigma = E\left[\left(\frac{X-\mu}{\sigma}\right)^3\right], \tag{29}$$

where E denotes an expectation value, σ is the variable's standard deviation and μ is the mean. As an example, figure 31 shows the PDFs of four random samples from Pearson distributions with skewnesses of 0, 1, 2 and 4 respectively; all other parameters are held constant. The skewnesses of the samples are -0.09, 0.96, 1.91 and 4.01 respectively. Note that a skewness of zero does not necessarily imply symmetry about the mean. Another oft-used statistical indicator is the most likely value (MLV), which is taken to be the central value of the bin with the highest frequency of events.

4 Transient-Eddy Heat Transport PDFs

4.1 General Features of the PDFs

To investigate the statistical distribution of transient eddy heat transport, we begin by computing composite PDFs, taking into consideration all available NH and SH latitude bands (30°−89° N and 30°−89° S) over twelve DJF and JJA time series (1993/1994-2004/2005), at 850 *mb*. This yields PDFs of almost 5×10^7 data points for each season, which ensures a smooth and representative distribution. The three panels in figures 32, 33, 34 and 35 show the results for v', H' and transport corresponding to NH DJF, NH JJA, SH DJF and SH JJA respectively. The key features of the

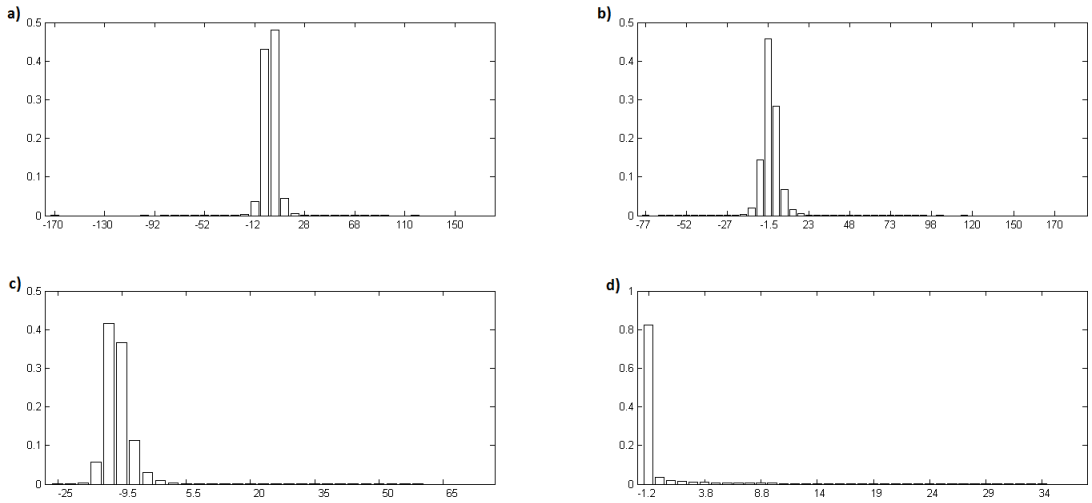

Figure 31: PDFs with skewnesses of a) -0.09, b) 0.96, c) 1.91 and d) 4.01. The data are random samples from Pearson distributions with skewnesses of 0, 1, 2 and 4 respectively. All other parameters are held constant.

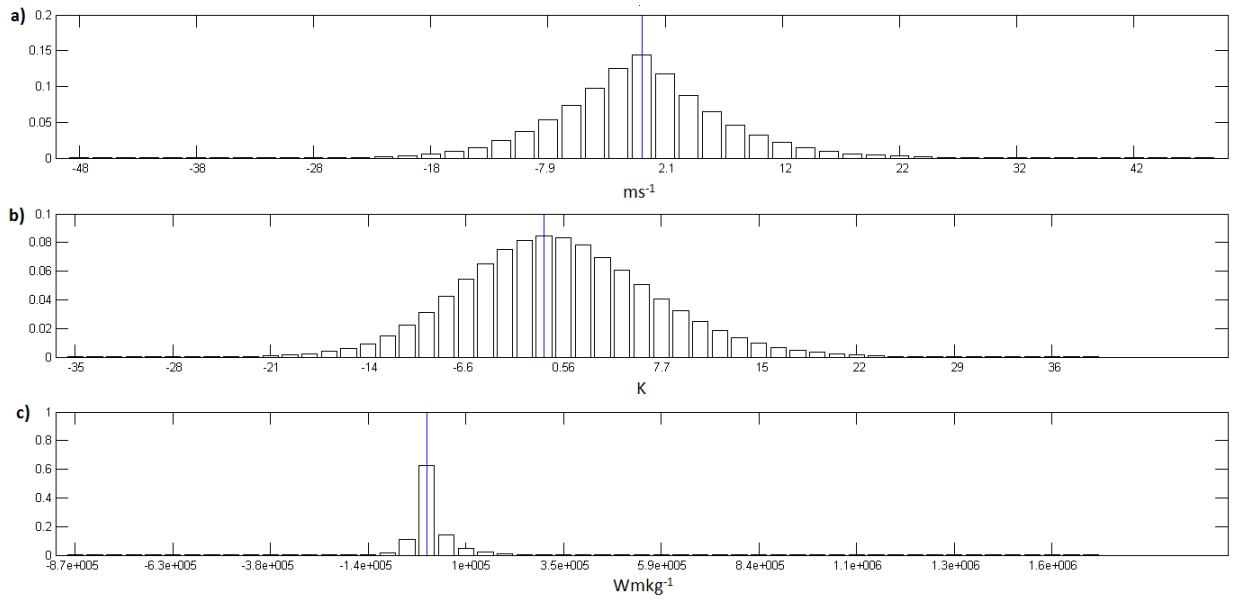

Figure 32: PDFs of a) meridional velocity anomalies, b) moist static energy anomalies and c) atmospheric heat transport due to transient eddies. The data cover the 850 mb fields for NH DJFs from December 1993 to February 2005. All latitude circles between 30° N and 89° N are taken into account. The skewnesses of the PDFs are respectively a) 0.24, b) 0.21 and c) 3.00. The corresponding most likely values are a) 0.1 ms^{-1}, b) -0.87 K and c) 6.8 $\times$ 10^3 $Wmkg^{-1}$. The vertical lines show the bins corresponding to the most likely values [from *Messori and Czaja*, 2013b].

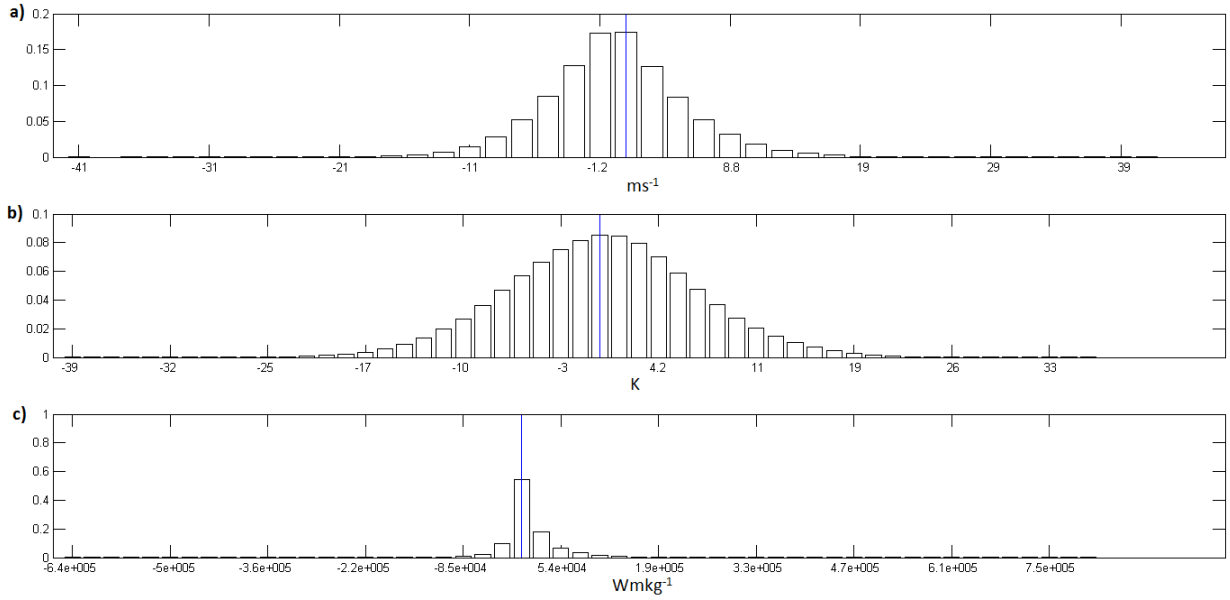

Figure 33: Same as figure 32, but for NH JJAs. The data now cover the period from the period from June 1994 to August 2005. The skewnesses of the PDFs are respectively a) 0.17, b) 0.04 and c) 2.1. The corresponding most likely values are a) 0.82 ms^{-1}, b) -0.16 K and c) -1.2 $\times$ 10^3 $Wmkg^{-1}$.

velocity and MSE PDFs, common to all four figures, are:

i) A very low skewness and a near-symmetrical structure;

ii) Significant positive and negative tails;

iii) A near zero most likely value when compared to the magnitude of the extreme events.

The transport PDFs, on the opposite, have:

i) A very high skewness, associated with a highly asymmetric distribution;

ii) Positive tails which are significantly more extended than the negative ones.

The most likely values of the distributions are typically two orders of magnitude smaller than the extremes, and can therefore be considered near-zero in terms of heat transport.

The symmetry found in velocity and MSE distributions corresponds to comparable realisations of positive and negative anomalies. Since the mean value of the distributions needs, by definition, to

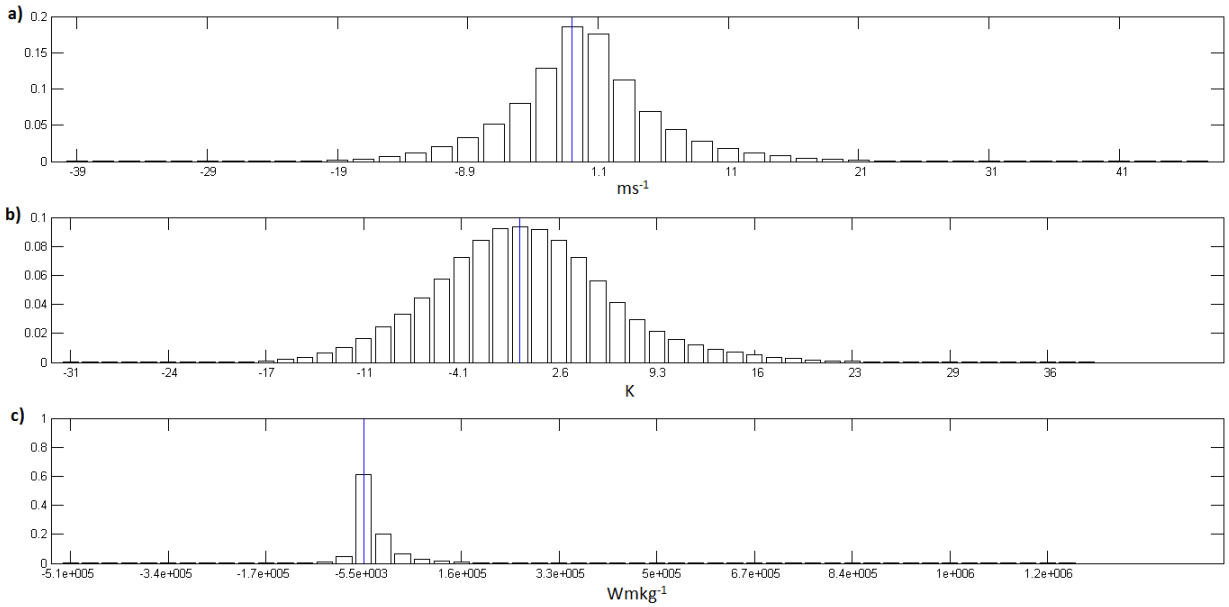

Figure 34: Same as figure 32, but for SH DJFs. The data now cover all latitude circles between 30° S and 89° S. The skewnesses of the PDFs are respectively a) 0.3, b) 0.33 and c) 3.7. The corresponding most likely values are a) -0.92 ms^{-1}, b) -0.09 K and c) -5.5 $\times$ 10^3 $Wmkg^{-1}$.

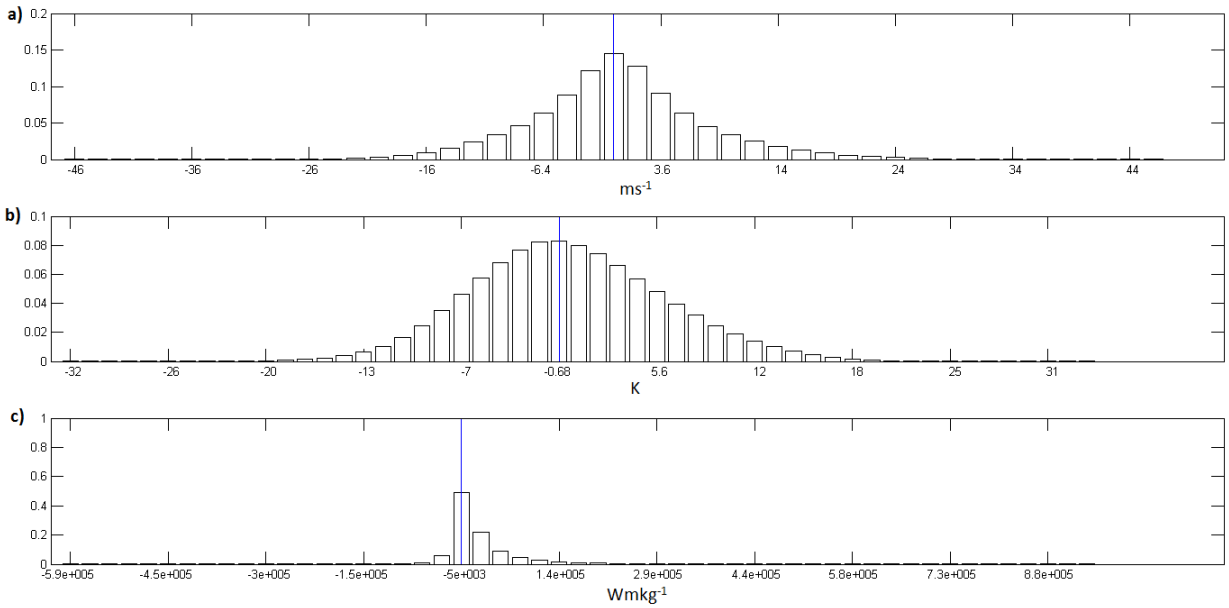

Figure 35: Same as figure 34, but for SH JJAs. The skewnesses of the PDFs are respectively a) 0.31, b) 0.2 and c) 2.7. The corresponding most likely values are a) -0.41 ms^{-1}, b) -0.68 K and c) -5.0 $\times$ 10^3 $Wmkg^{-1}$.

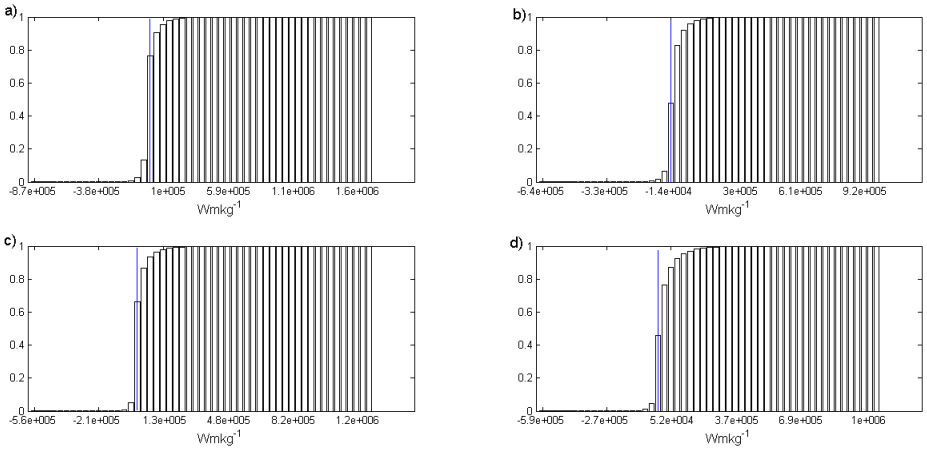

Figure 36: CDFs of atmospheric heat transport due to transient eddies for a) NH DJF, b) NH JJA, c) SH DJF, d) SH JJA. The data range are the same as those in figures 32-35 respectively.

be zero, there is no *a priori* reason to expect one of the two tails to have a radically different profile to the other. The large extreme values appear compatible with strong perturbations, considering that the vast domain analysed here includes the Pacific, Atlantic and Southern Ocean storm tracks. The strong asymmetry found in the transport PDFs, on the other hand, is related to the imbalance inherent to meridional heat transport whereby, as emphasized in Chapter II, there must be a net transport from low to high latitudes. Hence, some measure of asymmetry in a transport PDF is to be expected. What is not obvious is that the asymmetry should be related to very pronounced extreme events and near-zero most likely values.

This feature is illustrated further in figure 36, which shows the cumulative density functions (CDFs) of the transport for the four season and hemisphere combinations. The most likely values correspond to the bulk of the events, and the distribution is almost flat in correspondence with the long positive tail of the PDF which, on the opposite, corresponds to very few events. However, these very few events have a considerable importance. Repeating the analysis on magnitude-weighted CDFs highlights how they account for a significant part of the asymmetry of the transport distribution. A PDF of the modulus of negative transport days is computed, and this is subtracted from the PDF of the positive transport days. The bins are then weighted by the magnitude of the events they represent, and a CDF is computed. Figure 37 shows the resulting distributions for the four season-hemisphere combinations. Where the curve's slope is positive, this indicates that the positive transport events

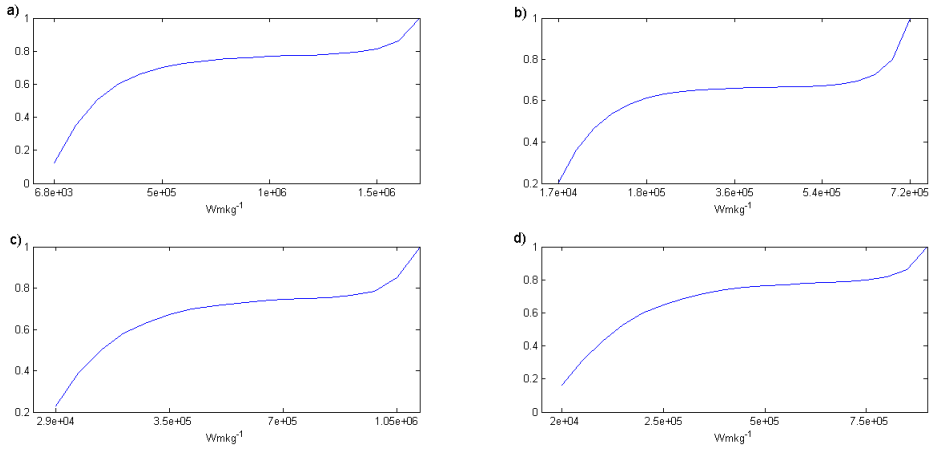

Figure 37: Magnitude-weighted CDFs of atmospheric heat transport due to transient eddies for a) NH DJF, b) NH JJA, c) SH DJF, d) SH JJA. The data range are the same as those in figures 32-35 respectively. A PDF of the modulus of negative transport days is computed, and this is subtracted from the PDF of the positive transport days. The bins are then weighted by the magnitude of the events they represent, and a CDF is computed.

for that given bin outweigh the negative transport ones. The small positive events, clustered around the MLVs, significantly exceed their negative counterparts, owing to their very high frequency. The intermediate events occur with comparable frequencies in both the negative and positive domains, leading to an almost zero slope in the central section of the curves. Finally, the steep section at the end of all plots shows that the very few large events, which make up the long positive tails of the PDFs, outbalance any comparable equatorward transport events. These events therefore account for a significant portion of the skewness seen in the transport PDFs.

Similar results are found for the PDFs comprising all vertical levels of the ERA-Interim data set. Figure 38 shows the vertically integrated transport distribution for NH DJF, where the integral has been performed by weighting each transport value by the corresponding layer thickness and grid box width. Note that the transport is now expressed in W, and that the peak events are of order 0.5 PW. The distribution reproduces all of the features discussed for the 850 mb data.

4.2 The Role of Extreme Events in Transient Eddy Heat Transport

A visual assessment of the heat transport PDFs (figures 32-38) suggests that poleward transport is heavily affected by a small number of very large events. Table 1 displays the contribution of the

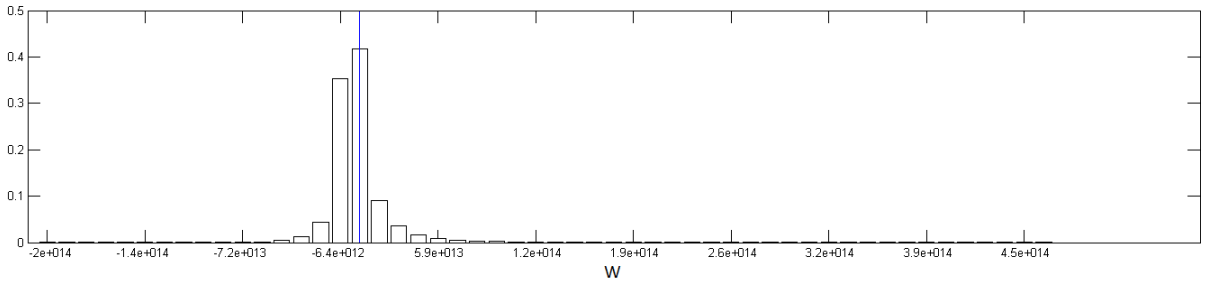

Figure 38: PDF of atmospheric heat transport due to transient eddies. The data cover all ERA-Interim pressure levels for NH DJFs from December 1993 to February 2005. All latitude circles between 30° N and 89° N are taken into account. The skewness of the PDF is 2.7. The corresponding most likely value is 6.7×10^{12} W. The vertical line shows the bin corresponding to the most likely value.

top 2%, 5% and 10% of events to i) the overall and to ii) the poleward-only transports. The values displayed are simply i) the percentage contribution of the selected events to the overall integral of the distribution and ii) the percentage contribution of the selected events to the integral of the positive portion of the distribution. It is immediately clear that, regardless of the percentile used to define extreme events, the higher end of the distribution accounts for a disproportionately large portion of the meridional heat transport. As is shown in table 1, events in the top 5 percentiles typically account for over half of the net transient eddy heat transport, with the exception of SH JJA where the percentage falls just short of this value. The contribution of the top 10 percentiles during NH JJA even approaches 100%, indicating that the transport due to those events is almost as large as the overall net transport. Indeed, the percent contributions systematically peak during the summer months in each hemisphere. These features are found to be independent of the fact that a single pressure level is being analysed. Vertical integrals of $v'H'$ over 1000 mb – 1 mb, whose PDF for NH DJF is shown in figure 38, show even higher contributions to the net transport from the upper percentiles. As is shown in table 2, the extreme events are so preponderant that the contribution of the top 10% during NH JJA even exceeds 100%!

Looking at the numerical values of the 850 mb PDF integrals (not shown), the expected seasonal trend emerges, with the net transport peaking during the winter months of each hemisphere. Although the same cycle is found in the extreme event-only integrals, the variability is attenuated. This explains why the percentage contributions actually peak in the summer months, when the transport is at

a) Hemisphere	Percentile	Overall % weight	Poleward-only weight
N	2	31.8	21.9
	5	56.8	39.2
	10	83.7	57.8
S	2	29.9	23.3
	5	53.3	41.6
	10	77.6	60.6
b) Hemisphere	Percentile	Overall % weight	Poleward-only weight
N	2	33.4	20.8
	5	61.5	38.4
	10	92.2	57.6
S	2	23.2	19.2
	5	43.5	36.1
	10	66.1	54.8

Table 1: Percentage contribution of extreme $v'H'$ events in a) DJF and b) JJA to net and poleward-only meridional atmospheric heat transport due to transient eddies. The data cover all longitudes and latitudes, from 30° N to 89° N and from 30° S to 89° S, over the 24 seasons considered (December 1993–August 2005). The percentile column indicates which percentiles of $v'H'$ events are classed as extreme [from *Messori and Czaja*, 2013b].

a) Hemisphere	Percentile	Overall % weight	Poleward-only % weight
N	2	37.6	23.2
	5	66.5	41.1
	10	96.3	59.5
S	2	37.9	24.7
	5	68.1	44.5
	10	98.8	64.5
b) Hemisphere	Percentile	Overall % weight	Poleward-only % weight
N	2	40.1	23.5
	5	71.0	41.6
	10	102.9	60.3
S	2	27.7	20.0
	5	52.2	37.7
	10	79.2	57.2

Table 2: Same as table 1, but for vertically integrated transport values.

its weakest. The anomalous percentage found for NH JJA, at 850 *mb*, is therefore due to the very pronounced seasonal cycle in net NH transient eddy transport, whose magnitude is not as pronounced in the extreme event-only integral. The same pattern is also witnessed for the vertical integration case.

The amount of heat carried poleward by transient eddies appears to be largely based on very few, very large events. Obviously, the extremely high contributions found could be due to the overall integral of the distribution being close to zero. The poleward-only contributions, however, show that this is not the case: the same 5 percentiles still account for approximately 35% to 45% of the poleward-only transport (see tables 1 and 2). It is interesting to note that the increase in contributions in the vertically integrated transport, compared to the 850 *mb* case, is mainly seen for the overall weight rather than for the poleward-only one. This suggests that the effect is largely due to the integral of the full vertical distribution being even closer to zero, relative to the magnitude of the extreme events, than was the case for the 850 *mb* level. It is important to note that the weight of extreme events is very high regardless of latitude. To make this point, figure 39 displays the contribution of the top 5% of events to overall (blue asterisks) and poleward-only (green crosses) heat transport at selected latitudes, rather than averaged over an ensemble of latitudes, as was done in tables 1 and 2. Panel a) displays the data for NH DJF, panel b) for NH JJA, panel c) for SH DJF and panel d) for SH JJA. The values found are in line with those shown in table 1, and there is no evidence to suggest that the contribution of extreme events is largest at 45° of latitude, where the transient eddy poleward heat transport is almost at its peak [e.g. *Fasullo and Trenberth*, 2008]. Indeed, the highest contributions are found at other latitudes, where the net transient eddy transport is smaller. As was seen in table 1, there is an anomalous contribution approaching 100% for NH JJA (in this case at 30° N). Again, this is due to the full integral of that specific PDF having a smaller value than those of the other distributions. In fact, while the net contribution at 30° N exceeds the second highest data point in the figure by a factor of one and a half, the corresponding poleward-only contribution is not even 20% larger than the next highest value. A similar pattern is seen in panel c) for the high contribution at 75° S.

As seen in figures 32-38, while the magnitude of extreme events and the skewness of the transport PDFs show some seasonality, the key features of the distributions, identified above, are robust characteristics of the data analysed. Distributions for individual latitude bands, individual seasons and

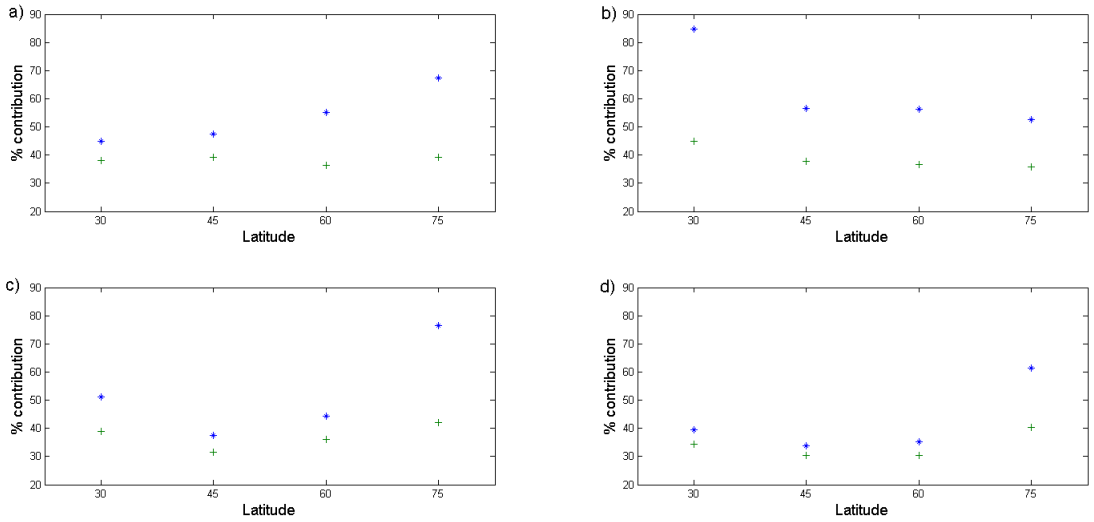

Figure 39: Percentage contribution of the top five percentiles of $v'H'$ events in a) NH DJF, b) NH JJA, c) SH DJF and d) SH JJA to net (blue asterisks) and poleward-only (green crosses) meridional atmospheric heat transport due to transient eddies. The data cover the 850 mb fields for 30°, 45°, 60° and 75° N and 30°, 45°, 60° and 75° S over the 24 seasons considered (December 1993–August 2005).

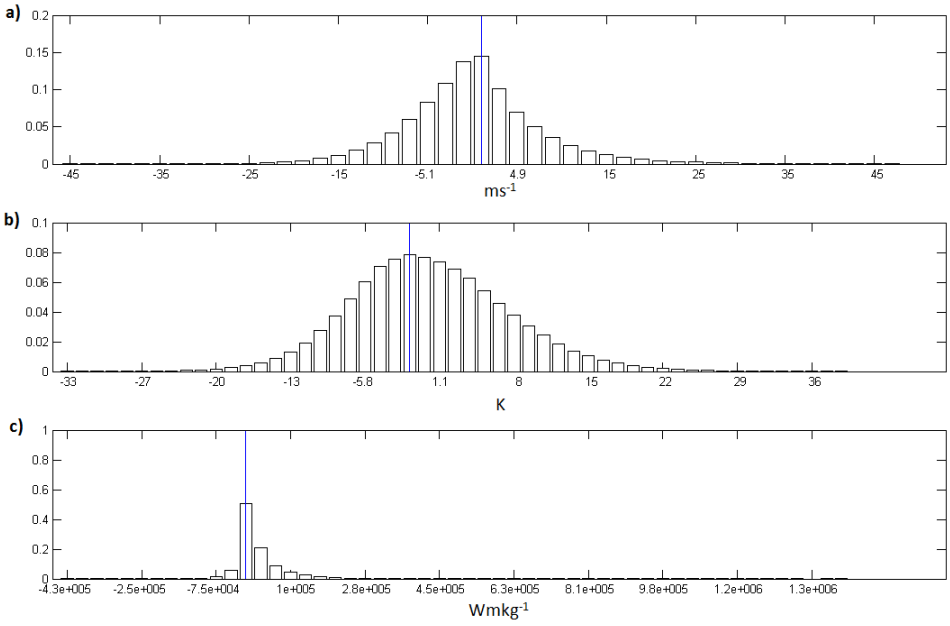

Figure 40: PDFs of a) meridional velocity anomalies, b) moist static energy anomalies and c) atmospheric heat transport due to transient eddies. The data cover the 850 mb fields for 45° N DJFs from December 1993 to February 2005. The skewnesses of the PDFs are respectively a) 0.45, b) 0.27 and c) 4.10. The corresponding most likely values are a) 0.9 ms^{-1}, b) -1.7 K and c) -4.3 $\times$ 10^3 $Wmkg^{-1}$. The vertical lines show the bins corresponding to the most likely values.

different pressure levels were also computed (not shown); significant variability in the magnitudes of skewness and extreme events was found but, again, the aforementioned features of the PDFs were found to be robust. As example, figure 40 shows the v', H' and transport distributions for 45° N, DJF. Regardless of the exact definition of extreme event in terms of percentiles, and regardless of the temporal and spatial domains seletced, very few events each season therefore seem to account for over half of the poleward heat transport by transient eddies. This, and the near zero most likely value seen in the PDFs, are inherent properties of the transport distribution and need to be satisfactorily explained.

4.3 The Zonal Mean View

The analysis presented so far has been performed on heat transport computed at single grid boxes. The possibility of strong return flows balancing the extreme poleward events, and the characteristics of the heat transport from a zonally integrated perspective, have not been addressed. By zonally

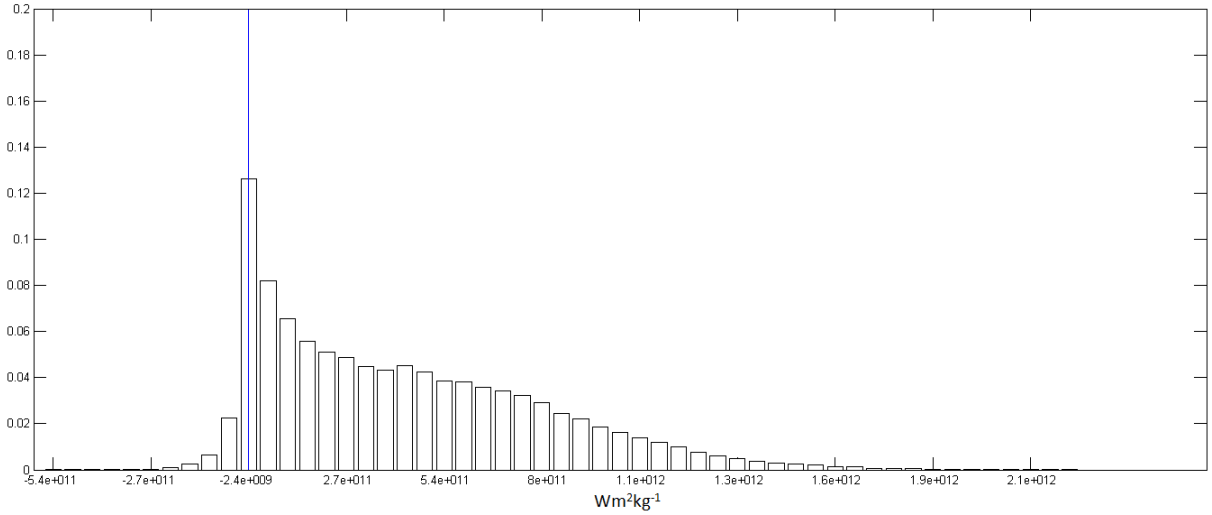

Figure 41: PDF of zonally integrated atmospheric heat transport due to transient eddies. The data cover the same range as in figure 32. The skewness of the PDF is 0.84. The corresponding most likely value is -2.4 $\times$ 10^9 Wm^2kg^{-1}. The vertical line shows the bin corresponding to the most likely value.

integrated transport, what is intended here is simply the sum of all single grid box transport values across a given latitude circle, weighted by grid box width. First, we investigate whether the top percentiles of the zonally integrated transport PDF play any relevant role relative to the distribution's integral. In constructing the PDF, the integrals of heat transport across a latitude circle, on a given day, are treated as single data points. In the interest of conciseness, the resulting distributions are shown only for NH DJF and SH JJA (figures 41 and 42 respectively). The PDFs are significantly different from the single-point ones, both in terms of skewness and of overall shape of the distribution. The PDF for SH JJA even appears to be bi-modal. As would be expected by the strong positive year-round atmospheric heat transport, negative values become almost non-existent. It should be noted, in this respect, that the negative MLV of the PDF for NH DJF lies in the bin containing zero. The value is negative simply because of the convention, adopted here, of expressing the MLV as the central value of the bin with the highest frequency of events. Concerning the bi-modal SH PDF, by splitting the distribution into two latitude bands (30°-60° S and 60°-89° S), it becomes clear that the right-hand side peak is due to the lower latitudes and the left-hand side one to the higher latitudes. This is partly due to the higher frequency of extremes at lower latitudes. As

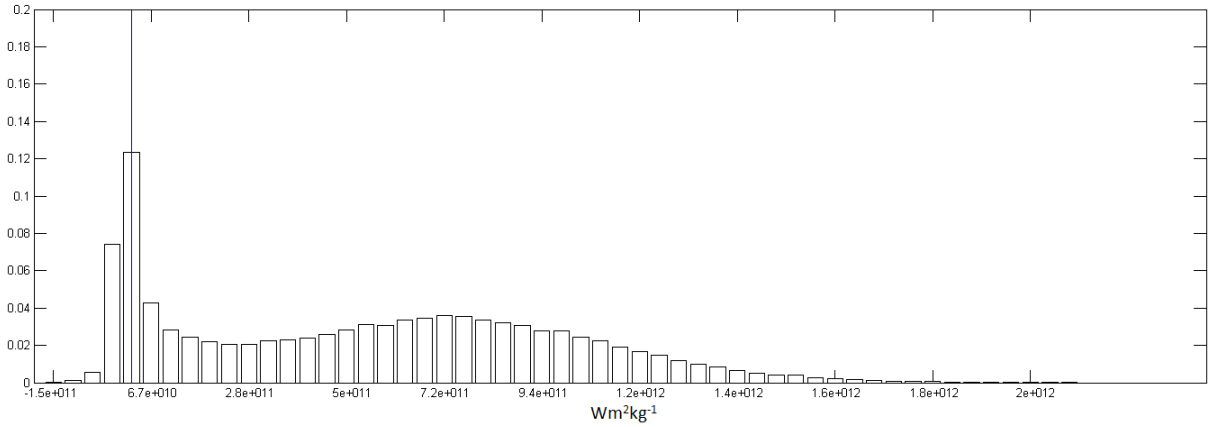

Figure 42: Same as figure 41 but for SH JJA. The data cover the same range as in figure 35. The skewness of the PDF is 0.33. The corresponding most likely value is 2.4 $\times$ 10^{10} Wm^2kg^{-1}. The vertical line shows the bin corresponding to the most likely value.

illustrated in figure 51 on page 87, the very high latitudes in the SH display virtually no extremes; when analysing the 60°-89° S domain, the relative frequency of extremes is therefore significantly lower than for the 30°-60° S band. This leads to a pronounced near-zero MLV. The geographical distribution of extremes is discussed in further detail in Section 5.2. In addition to this, in the zonal integration procedure events at lower latitudes are integrated over wider grid boxes than the ones at higher latitudes. Indeed, reproducing figure 42 without integrating the transport values across grid box width still yields a bimodal distribution, but the two peaks are not as well separated. Care should be taken in comparing these results with the smooth zonally and vertically integrated yearly mean transport curve, shown in figure 18, since here we are looking at the temporal variability of a single pressure level. The PDFs for SH DJF and NH JJA share the same qualitative features as their wintertime counterparts, albeit with some quantitative differences.

As is immediately evident from a visual assessement of the PDFs, the contributions of the top 5 percentiles to the overall integrals of the distributions are significantly lower than those found in Section 4.2. They are presented in table 3. These values can be compared to the weight of the events above the same threshold value in a Gaussian distribution, with the same mean and standard deviation as the zonal transport distributions, shown in the third column of the table. Even though the percentage values for the latitudinal sum are significantly smaller than those found for the single-point events, they are still typically three to five times larger than those found for the

Hemisphere	Season	Overall % weight	Poleward-only % weight	Gaussian % weight
N	DJF	16.1	16.0	2.97
	JJA	15.1	15.0	3.05
S	DJF	16.5	16.4	3.25
	JJA	13.3	13.2	4.70

Table 3: Percentage contribution of the top five percentiles of zonally integrated $v'H'$ events to net and poleward-only meridional atmospheric heat transport due to transient eddies. The fifth column displays the contributions of events above the same threshold in a Gaussian distribution, with the same mean and standard deviation as the zonal transport one. The data cover all latitudes and longitudes, from 30° N to 89° N and from 30° S to 89° S, over the 24 seasons considered (December 1993–August 2005).

corresponding Gaussian distributions. There is therefore some basis for talking about extreme events in heat transport also in the context of zonally integrated values.

Having established this, the next pertinent question to address is how these zonal events might relate to the local extremes discussed previously. We consider the following hypotheses concerning the origin of large values of zonally integrated heat transport:

i) They are due to synchronised local extremes at several gridpoints around a given latitude. Namely, several extreme events occurring on the same day, along the same latitude.

ii) They are due to a larger than average transport across all longitudes, with no significant contribution from the local extremes. That is, to a generalised increase in the transport across large stretches of the latitude circle, without necessarily implying a higher than normal frequency of extreme events. This would result in an increased mean value of single grid box transport around the latitude circle.

iii) A combination of points i) and ii) above.

To test hypothesis i), one can produce PDFs of the number of local extremes on days corresponding to zonal extremes versus normal days. If the PDF for the extreme zonal days is peaked at larger values than the one for all other days, then zonal extremes come about because of local extremes.

Figures 43 and 44 show the resulting PDFs for NH DJF and SH JJA respectively. White bars correspond to data for zonal extremes, while grey bars correspond to data for all other days. For NH DJF, the zonal extremes PDF's MLV and mean are both larger than the non-extreme PDF's ones by a factor of approximately four and two respectively. A very similar pattern is seen for NH JJA

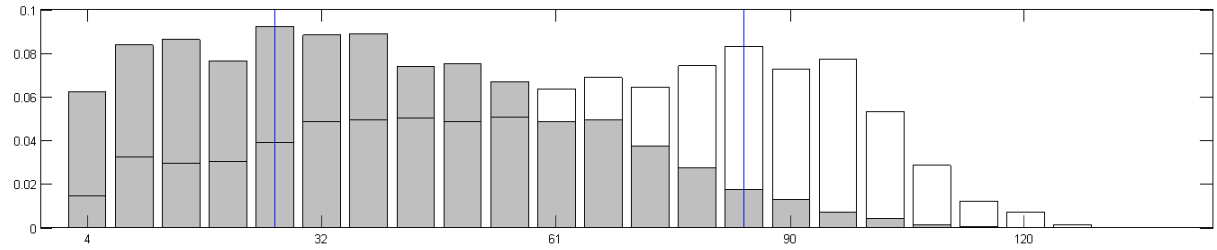

Figure 43: PDFs of the number of local extreme events around a given latitude circle for days which are in the top 5% (white) and days which are in the bottom 95% (grey) of the distribution of the zonally integrated atmospheric heat transport due to transient eddies. The most likely values are respectively a) 88 and b) 22. The corresponding means are a) 69 and b) 39. The data range is the same as in figure 32. The vertical lines show the bins corresponding to the most likely value.

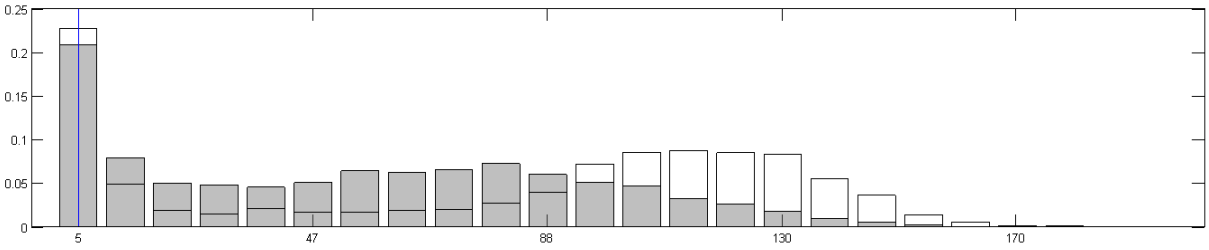

Figure 44: Same as figure 43, but for SH JJA. The data range is the same as in figure 35. The most likely values are respectively 4 (extremes, in white) and 4 (non-extremes, in grey). The corresponding means are 90 and 54.

(not shown). The distributions for SH JJA are again bimodal, as was the case for the magnitude of the zonal extremes. The two PDFs suggest that there are two distinct patterns in SH JJA, both in terms of magnitude of the zonally integrated transport and in terms of the number of local extremes contributing to zonal maxima. One pattern holds for the lower latitudes, where there are larger values of zonally integrated transport resulting from the contribution of a large number of local extremes. The other holds for the higher latitudes, where both the number of local extremes and the magnitude of the zonal transport are smaller. This trend is emphasised by the choice to weight the transport by grid box width in computing the zonal integral. The near-zero MLV of both PDFs is driven by the very high latitudes, where no 850 mb extremes are seen on account of large portions of the Antarctic continent being at a high elevation. As was the case for the NH, the mean of the extreme distribution is larger than its non-extreme counterpart by a factor of approximately two. The shape of the SH DJF distribution (not shown) is very similar to the SH JJA one.

There are two caveats to the methodology adopted here. The first is that the validity of the three hypotheses made above depends on the exact definition of local extremes. This is partly balanced by the fact that, if the quantile defining local extremes changes, so does the one defining the zonal ones. For example, if local extremes were defined as events within the top ten percentiles of the $v'H'$ distribution, the same definition would be applied to the zonally integrated distribution in order to select extreme zonal days. In terms of the distributions shown in figures 43 and 44, a looser definition of local extremes would lead to more overlap between the PDFs for extreme and non-extreme zonal days. At the same time, the effect of applying the same, looser definition to the zonal values would have the opposite effect. This was confirmed by repeating the above analysis for extremes defined as events in the top ten percentiles of the transport distributions. The second caveat is that the zonal integration artificially alters the shape of the distributions, by weighting the transport by grid box width. At the same time, it affects both extreme and non-extreme zonal days in the same fashion, and there is a clear indication that larger zonally integrated values do correspond to larger numbers of local extremes.

To put numbers on this statement, in figures 43 and 44 the overlap between the extreme and non-extreme distributions is 60% and 65% respectively. These are simply the percentages of data points in the distributions which lie in the overlapping portions of the different bins. The percentages provide an indirect measure of how likely it is for an extreme zonal day to have the same number of local extreme

events as a non-extreme zonal day. Hypothesis i) above would imply that the two PDFs have almost no overlap, since local extremes would be the only drivers of zonal ones. On the opposite, hypothesis ii) would correspond to approximately equal PDFs for both extreme and non-extreme zonal days. In physical terms, we therefore interpret the above results as corresponding to a background flow-driven scenario with a contribution from the increased number of extreme events (hypothesis iii) above). The zonal heat transport is therefore characterised by a "weak synchronisation" effect, whereby zonal extremes do, in part, result from synchronised local extremes, but they have a much weaker impact on the overall transport distribution than their single-point counterparts. The extreme zonal days which have the same number of local extremes as non-extreme zonal days display, on average, an enhanced contribution from the transport at non-extreme locations. A very good graphical illustration of this concept comes by producing a barplot of occurrence of extreme events on a given latitude circle over a season. On a given day, the value of the vertical bar is set to the number of local extremes occurring around the latitude circle. The local extremes are selected using a threshold computed relative to events in the whole hemisphere (limited to the $30°-89°$ N band being analysed in the present manuscript). In such a plot, it is found that the vast majority of days have at least one extreme event occurring around the latitude circle. As can be seen in figure 45, computed for $50°$ N JJA, the extreme events tend to happen in "bursts" lasting for a few days, during which very significant numbers of events occur. These are preceded and followed by periods with very little activity. In certain years, more extended periods of activity are present. At the same time, it is clear that there is a continuous background of extremes throughout the season. This visual appraisal of the temporal variability of meridional heat transport is consistent with the previously discussed interpretation of the PDFs in figures 43 and 44.

4.4 Extreme Events in Climate Models

4.4.1 FORTE Model

The above results show that, regardless of the exact definition of extreme events, and regardless of the exact spatial and temporal domains chosen, extreme events play a key role in meridional heat transport in the atmosphere. In order to further test the robustness of these results, it is of interest to verify whether they are reproduced by idealised climate simulations. To this end, the analysis described in Section 4.1 was repeated for the output of a coupled ocean-atmosphere model (FORTE:

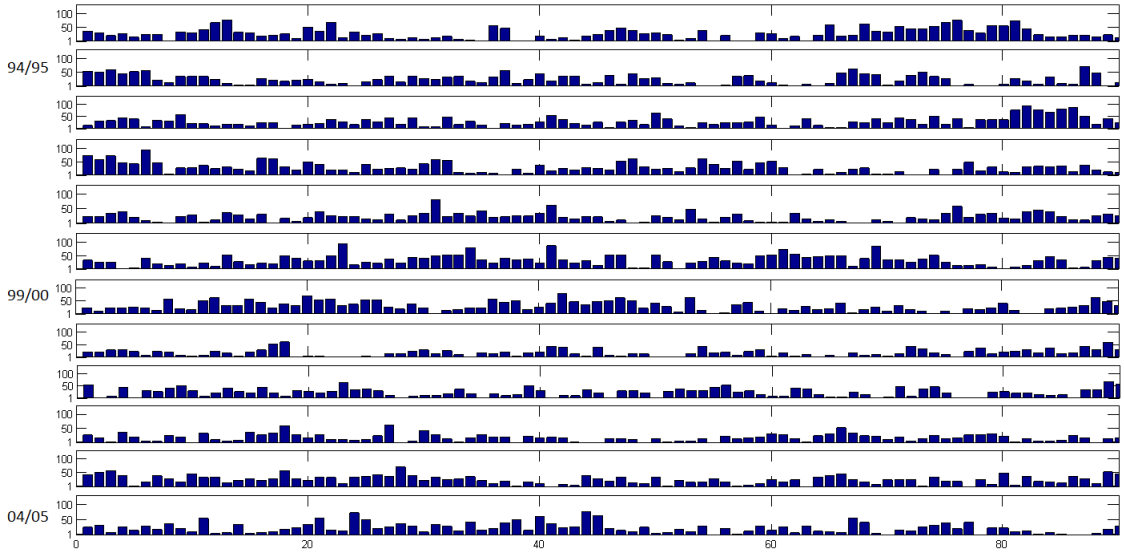

Figure 45: Bar plot of $v'H'$ extreme events along the 50° N latitude circle. Extreme events are defined as events in the top 5 percentiles of the $v'H'$ distribution for the 30°−89° N latitude band. The twelve panels correspond to JJA seasons from JJA 1994 to JJA 2005. The height of the bars corresponds to the number of longitude gridboxes displaying an extreme event on a given day; the abscissa indicate the day of the season.

Sinha and Smith, 2002; *Smith and Gregory*, 2009) run in an aquaplanet geometry (water-covered world, as described in *Smith et al.*, 2006). Analogously to the ERA-Interim data, 12 years of model run were analysed. Further details concerning both the model in general, and the specific run analysed here, are provided in Chapter II, Section 3.2 on page 56.

In an aquaplanet geometry, the statistics are only functions of latitude and height, and the circulation is dominated by large scale waves developing on the mid-latitude westerlies. The results of the PDF analysis at 850 *mb* for FORTE were found to be very similar to the ones for the ERA-Interim data, even though the model's extremes are approximately one order of magnitude smaller. The distribution for NH DJF is shown in figure 46. Even though the skewness of the PDF is smaller than the one of the corresponding ERA-Interim PDF (see figure 32), the near-zero MLV and extended positive tail, identified as the key statistical signatures of the heat transport distribution, are correctly reproduced. It should also be noted that the above distribution is, among the four hemisphere-season combinations in FORTE, the one with the smallest skewness. As reference, the PDF displaying the largest skewness is the one for SH DJF (not shown), which has a value of 2.9. This is further confirmation of the robustness of the results presented above, and will be discussed

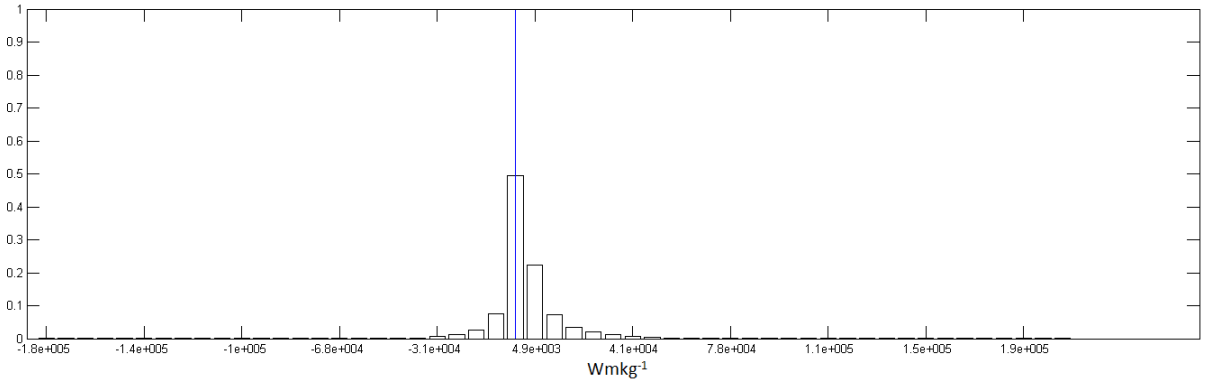

Figure 46: PDF of atmospheric heat transport due to transient eddies. The data cover the 850 mb fields for 12 NH DJFs in the FORTE model. All latitude circles between 30° and 89° N are taken into account. The skewness of the PDF is 1.6. The corresponding most likely value is -2.3 $\times$ 10^3 $Wmkg^{-1}$. The vertical line shows the bin corresponding to the most likely value.

further in Section 7.

4.4.2 Minobe *et al.* Model

To glean additional insights into the nature of the heat transport statistics, the PDF analysis is also performed on another idealised model [*Minobe et al.*, 2008]. To be consistent with the previous analysis, we again focus on twelve DJF and JJA time series. The model, described in detail in Chapter II, Section 3.3 on page 56, is based on a fully coupled AOGCM (control simulation) and applies a smoothing to ocean SSTs over the NH western boundary currents, namely the Gulf Stream and the Kuroshio (smoothed simulation). *Minobe et al.* focus their analysis on the Gulf Stream, and conclude that the SSTs in the region influence the entire troposphere, with important effects on air temperature and pressure fields. In the control simulation, large low pressure systems develop in areas of sharp SST gradients, leading to a band of high precipitation anchored to the SST contours. In the smoothed simulation, the narrow rain band disappears altogether. The annual rainfall rate climatology is shown in figure 47.

As discussed in Chapter II, Section 1.3 on page 26, a large part of the meridional heat transport in the atmosphere is associated with cyclonic activity. If SSTs have a major influence on the atmosphere in a very active region such as the Atlantic storm track, it is natural to expect some influence on the heat transport distribution. The maps shown in figure 47 not only correspond to very different

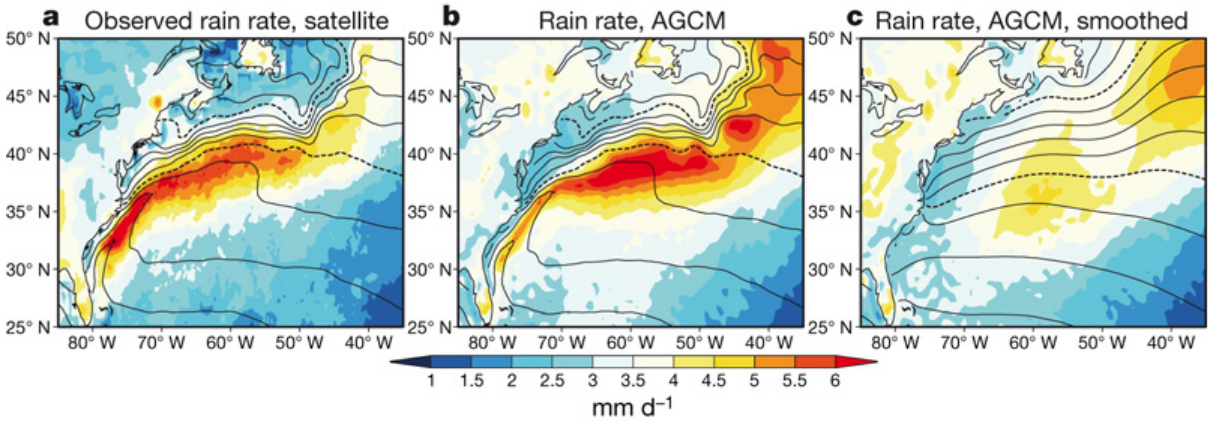

Figure 47: Rainfall rate (in mmd^{-1}) as a) observed by satellites, and as modelled with b) observed and c) smoothed SSTs. The black contours are for SST, with a 2 $°C$ interval and dashed contours for 10 $°C$ and 20 $°C$ [from *Minobe et al.*, 2008].

geographical distributions of rainfall, but also highlight important changes in peak intensity. If the effect on heat transport were to be analogous, the long tails of the PDFs, corresponding to the extreme occurrences, would be significantly shortened by a smoothed SST field. Comparing the results for the smoothed and control cases can therefore:

i) Confirm that fully coupled AOGCMs, in a realistic configuration, reproduce the extremes found both in the ERA-Interim data and in the highly idealised FORTE simulation;

ii) Provide valuable insights into the possible role of ocean-atmosphere interactions in driving the extremes.

In the control dataset, the heat transport PDFs are very similar to those seen for the ERA-Interim data. Since the data only cover a Gulf Stream and a Kuroshio domain (30°−50° N and 280°−330° E; 30°−50° N and 130°−180° E respectively), the skew is more pronounced than for the hemispheric ERA-Interim data[2]. In fact, these are areas of high extreme event activity (see Section 5.2), and the extreme events would therefore be expected to yield a more extended positive tail than seen for the whole hemispheric domain. The distributions for the smoothed data are very similar to the control ones, and display all the major statistical features already discussed. In all of the season-domain combinations, however, a small reduction in the magnitude of extreme events is seen. As example,

[2]Note that the present analysis is limited to the domains where SST smoothing is performed. These are slightly smaller than the whole control domains (see Chapter II, Section 3.3 on page 56).

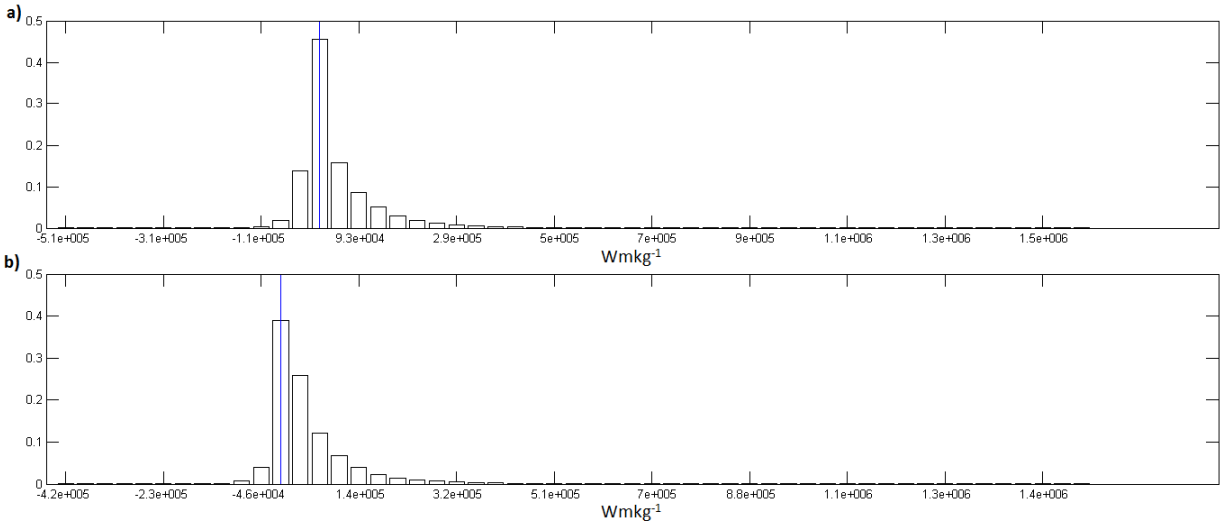

Figure 48: PDFs of atmospheric heat transport due to transient eddies in a) the control and b) the smoothed data from the the Minobe *et al.* simulation. The data cover the 850 mb fields for NH DJFs from December 1989 to February 2001. Only the Gulf Stream domain (30°−50° N and 280°−340° E) is taken into account. The skewnesses of the PDFs are respectively a) 3.5 and b) 3.5. The corresponding most likely values are a) 1.3×10^4 $Wmkg^{-1}$and b) -9.2 $\times$ 10^3 $Wmkg^{-1}$. The vertical lines show the bins corresponding to the most likely values.

figure 48 shows the PDFs for DJF in the Gulf Stream domain, for both control and smoothed data.

In terms of extreme event weights, table 4 shows the contribution of the top 5 percentiles in the control and smoothed cases, and the ratio of the two. The values, although slightly smaller, are of the same order as those found in the ERA-Interim analysis (see table 1 on page 68). The contributions for the control and smoothed data are seen to be extremely similar, and the ratio of the contributions is always close to 1.

The only significant difference between the two model runs arises in the geographical distribution of the extreme events. The smoothed data displays, in general, a broader geographical distribution of events for Atlantic DJF and, to a lesser degree, JJA. This results in areas of high extreme event frequency having less extremes and areas of low frequency having more. For the Pacific, the differences between the controlled and smoothed data are mainly related to a shift in the location of the extremes. Figure 49 shows the map of $v'H'$ events in the top five percentiles of the $v'H'$ distribution, for the Atlantic domain during DJF. Panel a) presents the control data, panel b) the smoothed data and panel c) the difference between the control and the smoothed values. As discussed above, upon smoothing the SSTs there is a decrease in frequency in the region of highest extreme event occurrence, and there

Domain	Season	Control Overall % weight	Smoothed Overall % weight	Control to Smoothed ratio
Gulf Stream	DJF	37.1	38.1	0.97
	JJA	46.7	46.5	1.00
Kuroshio	DJF	41.7	43.6	0.96
	JJA	49.3	47.4	1.04

Table 4: Percentage contribution of the top five percentiles of $v'H'$ events to the net meridional atmospheric heat transport due to transient eddies in the control and smoothed Minobe *et al.* model. The fifth column shows the ratio of the contributions displayed in the previous two. The data cover the 850 *mb* fields for both DJF and JJA from December 1989 to August 2001. The Gulf Stream and Kuroshio domains cover 30°–50° N, 280°–330° E and 30°–50° N, 130°–180° E respectively.

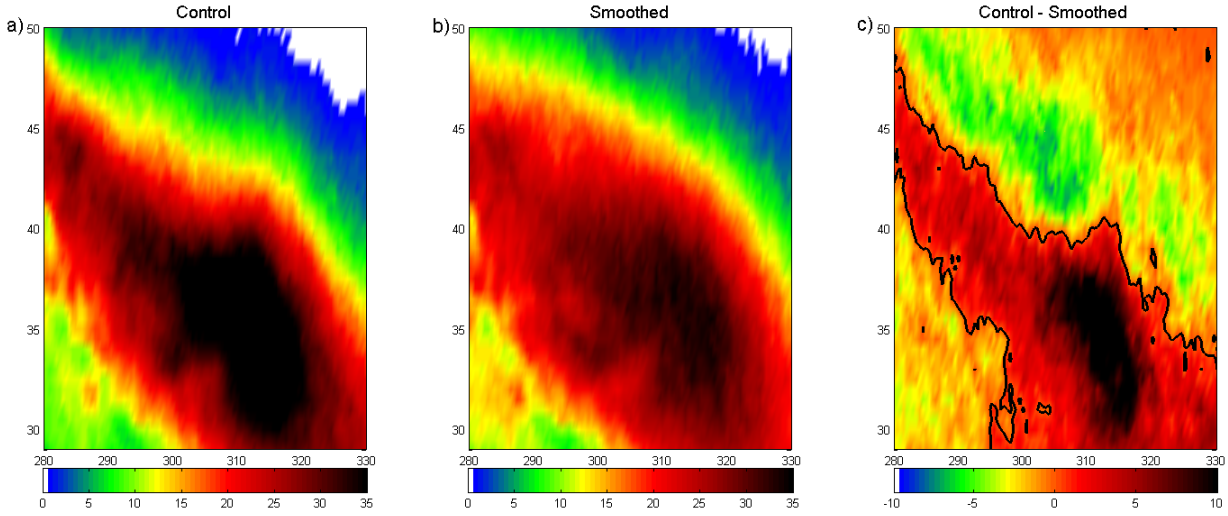

Figure 49: Map of $v'H'$ events in the top five percentiles of the $v'H'$ distribution, for the DJF Atlantic domain in the Minobe *et al.* model. The panels correspond to a) control data, b) smoothed data and c) the difference between the two (control - smoothed). All 12 seasons are taken into consideration. The scale of the colour bar corresponds to the number of data points per season per model grid box. The black contour in panel c) corresponds to zero difference.

is a corresponding increase in regions of lower occurrence. Since, by definition, both the smoothed and control cases include the same number of extreme events, the two effects balance.

The more diffuse distribution resulting from the lack of a sharp SST gradient mirrors the broader spatial distribution of rainfall discussed in *Minobe et al.* [2008]. However, unlike in the latter case, while the geographical location of the heat transport extremes changes, their intensity is almost unaffected by the model's SST configuration. This is mirrored in the contribution of extremes to the net seasonal transport, which is almost identical between the two runs. The implications of these findings will be discussed further in Section 7.

5 Spatial and Temporal Measures of Extremes

5.1 Duration and Spatial Extent

Having ascertained the importance of extreme events in setting the seasonal mean heat transport by transient motions, it is now necessary to investigate their physical structure in the atmosphere. To do so, we use a measure of temporal and spatial scales, whose principle is schematized in figure 50.

To investigate the spatial scales, the $v'H'$ signal on a given day is plotted as a function of longitude over a full latitude circle. Next, extreme events which fulfil given aspect ratio specifications are selected. In particular, the difference between the maximum of an event and the corresponding local minima is required to be larger than the typical magnitude of the near-zero fluctuations in the heat transport distribution. A lower limit of $10^4\ Wmkg^{-1}$ is chosen to represent this constraint. The same value is set as an upper limit for the difference between the two minima of an event. A width (ω) is then measured as the difference in degrees longitude between the points where $v'H'$ reaches half its maximum (dash-dotted line in figure 50). Note that the half maximum point is computed relative to the smaller of the two local minima on either side of the local maximum. From this width, a pseudo-wave number κ is computed as:

$$\kappa = \frac{360}{\omega} \tag{30}$$

Conversely, by plotting the $v'H'$ signal at a given geographical co-ordinate over a season and enforcing the same aspect ratio constraints, the measured width at half maximum is interpreted as a pseudo-

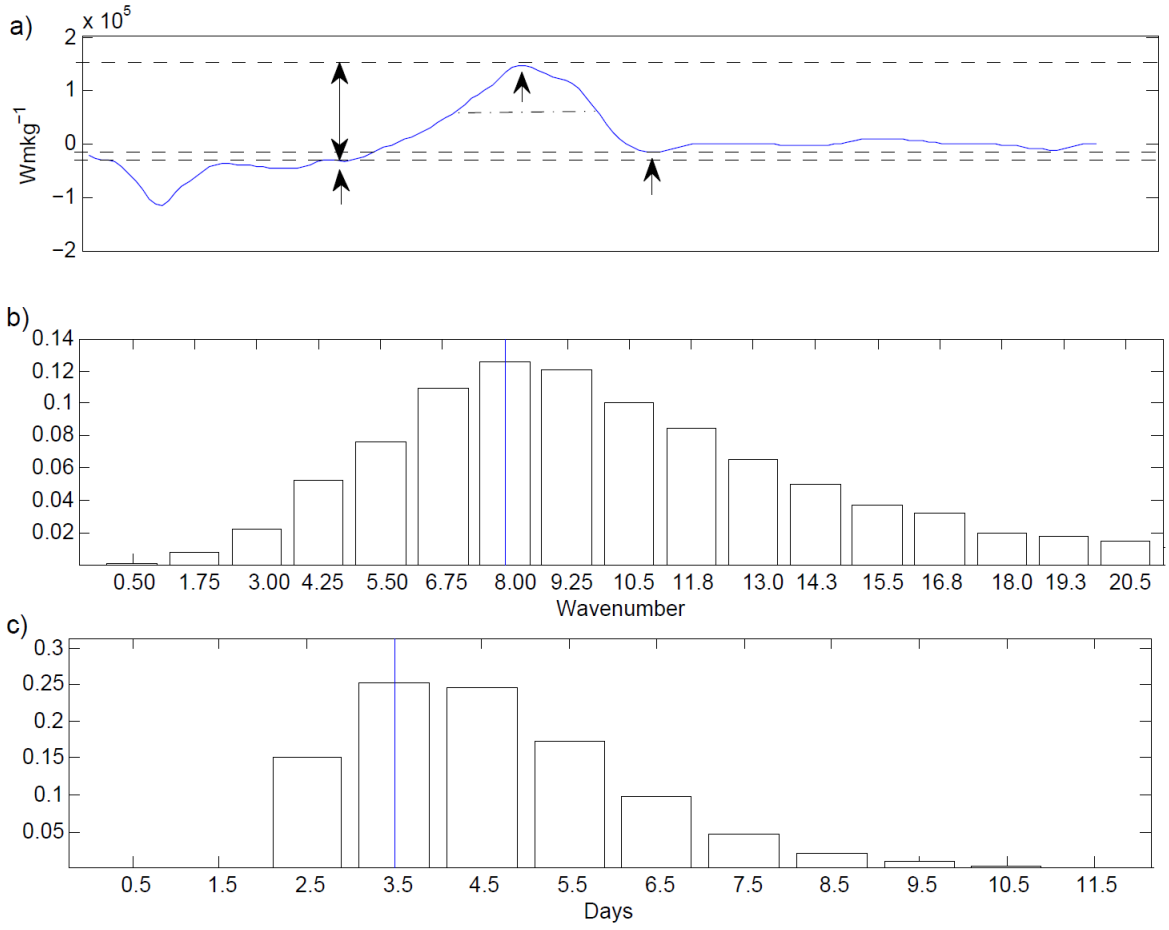

Figure 50: a) Schematic of how the full width/duration at half maximum of the $v'H'$ signal is computed. The three arrows indicate the local maximum corresponding to an extreme event and the two local minima on either side. The double arrow indicates the height of the extreme event relative to the smaller of the two minima. The dash-dot line indicates the measured half-maximum width/duration (ω/τ). In this case, the peak is well defined relative to the local minima and the two minima are close to one another, meaning that the event passes all the aspect ratio tests (see text for details). The two bottom panels are PDFs of b) the wave number K and c) the full duration at half-maximum T of extreme $v'H'$ events. The latter are defined as events in the top five percentiles of the $v'H'$ distribution. The data range is the same as in figure 32. The vertical lines show the bins corresponding to the most likely values [from *Messori and Czaja*, 2013b].

period τ. Repeating the above processes for all latitudes, longitudes and days allows one to obtain PDFs of κ and τ associated with extreme $v'H'$ events. Note that cut-off values for duration and spatial extent are enforced to account for the cases where the half-maximum value is smaller than one of the two minima. In the latter scenario, in fact, the width/duration would be measured as extending all the way to the next point where the $v'H'$ signal went below the half maximum value, hence not yielding the actual extreme event width. The definition of extreme event wavelength in the $v'H'$ signal is equivocal, since the pattern is not necessarily sinusoidal, and extracting the "shape" of the extreme event from the background is nigh on impossible. This is why pseudo-wave number κ and corresponding pseudo-period τ are chosen as measures. However, care should be taken when interpreting them: κ and τ are not the wave number and period corresponding to the wavelength of the extreme event. In fact, in a sinusoidal wave, ω would be half of the wavelength; this 1:2 scaling does not translate perfectly to a non-sinusoidal signal but gives a rough idea of the wave number of the extreme events described here. A similar reasoning applies to τ which, in a sinusoidal wave, would be half of the period of a full wavelength. We therefore define the following:

$$K = \frac{\kappa}{2};\ T = 2 \cdot \tau \tag{31}$$

where K and T are, respectively, the wave number and period corresponding to a full wavelength of an extreme event. Figure 50 shows the probability distribution functions for b) K and c) T for the Northern Hemisphere. The PDFs were computed using all latitude circles between 30° and 89° N and all DJFs and JJAs from December 1993 to August 2005. Events generally have a wave number lying between 4 and 13, with a most likely value in the bin centred on 8. Hence, the distribution in figure 50 compares favourably with the typical Eady wave number of 6 [*Gill*, 1982]. As a point of comparison, *Randel and Held* [1991] find that wintertime heat transport by transients, at 700 mb and 47° N, is dominated by wave numbers 4−7. Concerning duration, the figure suggests that most extreme events have a period of 2 to 7 days, with a most likely value in the bin including timescales between 3 and 4 days. Taking into account the factor of 2, this means that few extreme events persist for longer than 3 days (this will be better illustrated in figure 56 on page 94). At 60° latitude, the midpoint of the domain, wave number 8 corresponds to a wavelength of roughly 2500 km. Taking the period to be $3.0\text{x}10^5$ s (approximately 3.5 days), an order of magnitude calculation yields a phase speed of 8.35

ms^{-1}, which is compatible with a baroclinic system. The result is in good agreement with *Randel and Held* [1991] who, using ECMWF analyses, find phase speeds for 700 mb transient eddies to be of order 5-15 ms^{-1}. The duration, spatial extent and phase speed of extreme events therefore seem to be within the range expected of Eady-type growing systems, as will be further discussed in Section 7. PDFs analogous to those in the figure were also computed for the SH (not shown). Although the spatial extent of the events is slightly larger, the above conclusions are found to hold.

5.2 Geographical Distribution

In order to identify the locations where the extreme events occur most frequently, we compute next the number of $v'H'$ events per season per data grid box that fall in the top 5 percentiles of the distribution of $v'H'$. Figure 51 displays the resulting geographical distribution of extreme events. As expected from the dominant contribution of these events to the seasonal mean heat transport, there are similarities with more traditional measures of storm activity (*e.g.*, maps of time mean transient eddy heat flux, or track density of cyclones – see figures 12 and 13 on page 34). There are also regions of agreement with climatologies of warm conveyor belts, although significant differences exist (*e.g.* figures 3b and 3f in *Eckhardt et al.* 2004, reproduced in figure 52).

In the DJF season, there is an almost uninterrupted band of very high extreme event occurrence spanning the Southern Ocean between 35° S and 55° S, with significant activity over land in the lee of the Andes. This is broadly coherent with SH storm track maps, such as that in *Hoskins and Hodges* [2005] shown in figure 13 on page 34. Similar considerations apply to *Eckhardt et al.*'s WCB climatology. A more discontinuous pattern emerges in the NH, where three areas of high activity can be identified. The first spans from the lee of the Rocky Mountains across continental USA and Canada to the Gulf Stream sector of the North Atlantic. A second one is found over the Greenland and Norwegian Seas. The third area roughly corresponds to the Pacific storm track, with a second local maximum over the Bering Strait. The latter feature, as well as that seen off the North-East coast of Greenland, is distinct from the patterns seen in more traditional storm track diagnostics [e.g. *Blackmon*, 1976]. Since the Bering Strait is at the boundary between cold, dry polar air masses and warmer, moist maritime ones, outflows of cold air over the ocean could explain the local maximum [*e.g. Businger*, 1987; *Overland and Stabeno*, 2004]. An alternative explanation could be to associate these extreme with the low pressure systems often found in the area [*Pfahl and Wernli*, 2012]. In

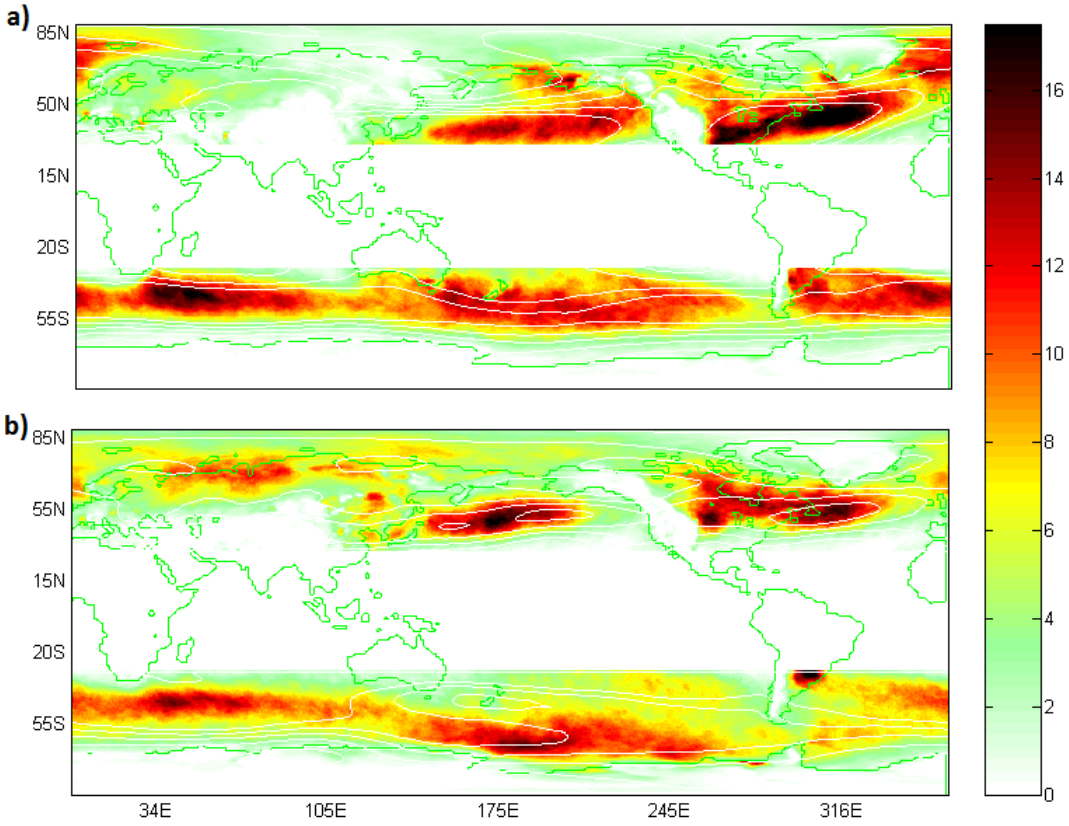

Figure 51: Map of $v'H'$ events in the top five percentiles of the $v'H'$ distribution for a) DJF and b) JJA. All 24 seasons (December 1993–August 2005) are taken into consideration. The scale of the colour bar corresponds to the number of data points per season per $0.75° \times 0.75°$ box. The calculation is not applied equatorward of 30° latitude. The white lines are regularly spaced 5 ms^{-1} contours of 300 mb time-averaged zonal winds. The contours range from 0 to 25 ms^{-1}, with positive wind speeds being eastbound [from *Messori and Czaja*, 2013b].

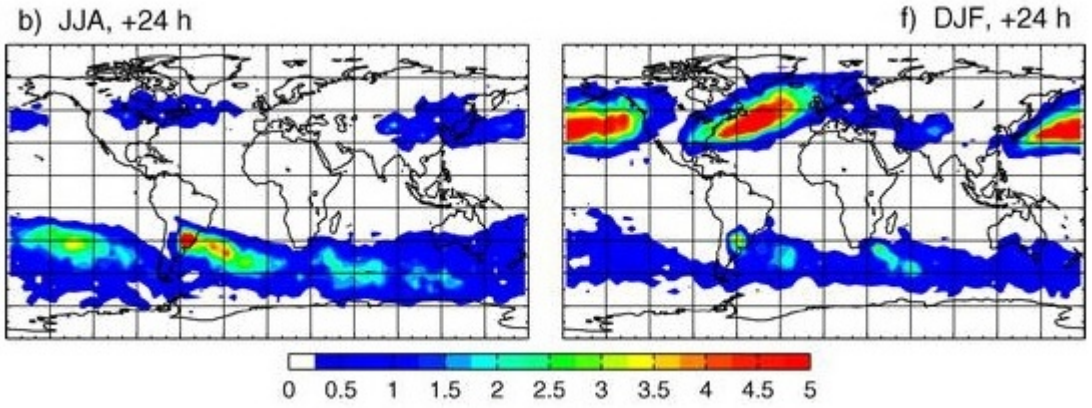

Figure 52: Seasonal mean spatial distribution of WCB trajectory positions 24 h after the beginning of the ascent during b) JJA and f) DJF. Depicted is the fraction (in percent) of all trajectories that fulfill the WCB criteria, averaged over 15 years. For details on the WCB criteria used, the reader is referred to *Eckhardt et al.* [from 2004], from which this figure is taken.

the latter paper, the authors define cyclones as closed 0.5 hPa pressure contours, and find that there is extensive cyclone activity over the Gulf of Alaska/Bering strait area, with a relative frequency of almost 50% during DJF. This is illustrated in figure 53 (originally figure 5 in *Pfahl and Wernli*, 2012).

The very high values found off the North-Eastern coast of Greenland, on the other hand, could be related to mesocyclone genesis occurring in the area (*Hoskins and Hodges*, 2002, see figure 12 on page 34), and match the area of polar low occurrences (see figure 54, originally figure 1 in *Kolstad*, 2011) [3]. Concerning the WCB climatology shown in figure 52, only two of these maxima, namely the storm tracks, are also regions of high WCB activity.

In the JJA season, the picture in the SH remains similar, although the high activity band is shifted poleward, with a well defined maximum over the Ross Sea. This implies that the band of high extreme event frequency occurs slightly further South than the WCBs. In the NH, the main change is that a high activity area appears over the Barents, Kara and Laptev Seas and over parts of Northern Siberia. A weakening of the Greenland and Gulf Stream maxima and the disappearance of the Bering Strait one are also witnessed. The latter feature is reproduced in the measure of cyclonic activity in the ERA-Interim dataset discussed above (see figure 53). The poleward shift of events during the JJA season could be explained by the seasonal shifts in the Jet Stream's position. The

[3] Polar lows are "intense mesoscale low-pressure systems, usually generated by outbreaks of cold, dry polar air over warm water masses" [*Kolstad*, 2006]

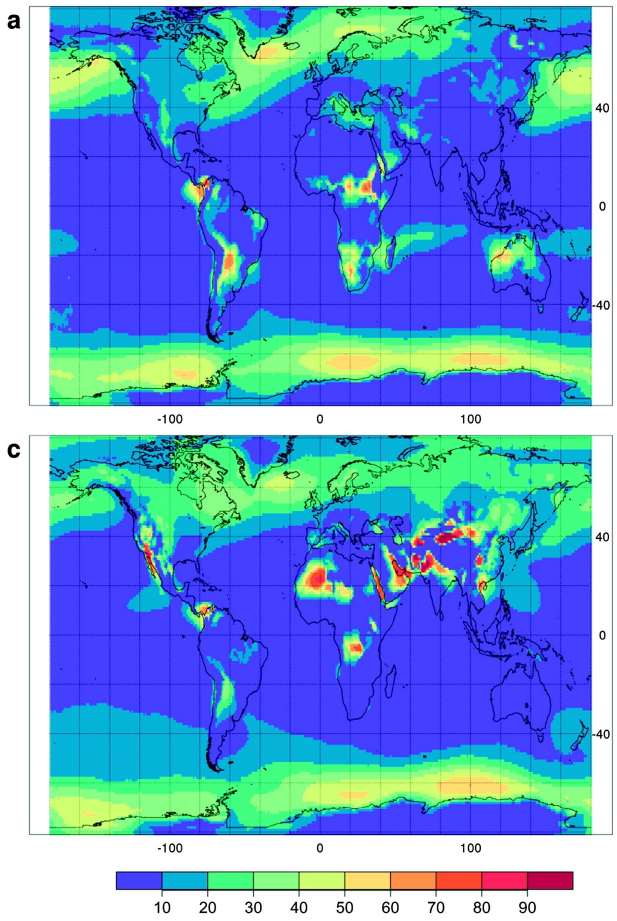

Figure 53: Relative cyclone frequency (%) in a) DJF and c) JJA. The data is 6-hourly ERA-Interim reanalysis and covers the period 1989-2009 [from *Pfahl and Wernli*, 2012].

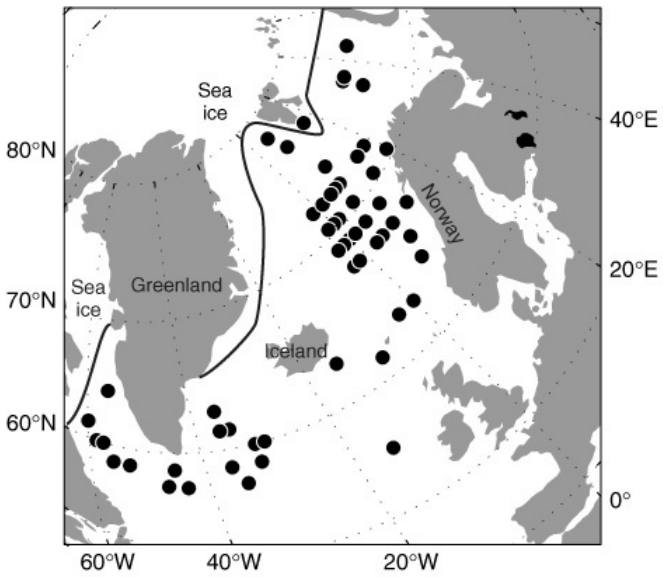

Figure 54: The locations of the polar lows. The data covers the November-March period during 1999-2009. The black curve marks the approximate boundary between the grid points that had open water one-third of the time during the analysis period and the ones that did not. For details of the data and techniques used in diagnosing the lows, the reader is referred to *Kolstad* [2011], from which this figure is taken.

white contours overlaid onto the colour map in figure 51 show mean zonal wind speeds at 300 mb. In the DJF season there is good correspondence between areas of high zonal winds and the Gulf Stream and Pacific maxima. There is also agreement in the SH. During the JJA season, two local wind speed maxima appear in correspondence with the Siberian frequency maximum, and there is very good agreement over the Pacific, Gulf Stream and Ross Sea areas.

6 Extreme Events in Transient-Eddy Heat Transport: Phase vs. Magnitude

As highlighted in table 1 on page 68, extreme events in the $v'H'$ distribution are key contributors to the net poleward heat transport by transient eddies. To identify the origin of these events, it is useful to picture the v' and H' signals as travelling waves, and to reason in terms of phase relationship and magnitude. An anomalously large $v'H'$ event might occur as a result of:

i) Extremely large velocity and MSE anomalies which, regardless of their phase relationship, yield a

very large transport value (magnitude-driven transport).

ii) Non-extreme v' and H' events which occur perfectly in phase (phase-driven transport).

Both mechanisms are possible candidates, since the product of peak v' and H' values typically yields transport values which are almost a factor of two greater than the extreme event threshold in the $v'H'$ distributions. Obviously, a combination of the two, such as near in-phase, large v' and H' events, is entirely possible. To investigate the phase relationship, a valuable tool is a scatter plot of v' versus H'. To obtain a clear picture, it is useful to remove all velocity and MSE anomaly data points which do not correspond to extreme $v'H'$ events. Note that this does not equate to removing all v' and H' data points which are not in the top percentiles of velocity and MSE anomaly distributions, as the selection is performed purely in terms of $v'H'$ percentiles. One can then bin the data points in the scatter plot along the two axes, analogously to a bivariate PDF. Colour coding the resulting distribution yields a discrete colour map, shown in figure 55. In this figure, the top two quadrants show the results averaging over the latitude band $30°-89°$ N in JJA and DJF, while the bottom two show the plots for $30°-89°$ S in the same seasons. The continuous lines in each quadrant show the values of the 5^{th} and 95^{th} percentiles of v' and H'. The squares labelled S contain all extreme $v'H'$ occurrences where neither v' nor H' are in the top/ bottom five percentiles of their respective distributions. The most striking feature of the map is the lack of near-zero v' and H' data points even though, according to figure 32 on page 61, near-zero values are the most likely realisations of the two variables. In a magnitude-dominated scene, extreme events in heat transport would mostly be due to the sheer magnitude of either v' or H'', regardless of their phase relationship. This would imply the presence of events where one of the two anomalies is small while the other one is extremely large. Such events are not seen in any of the panels in figure 55. This therefore immediately excludes the magnitude-driven transport picture. At the same time, a large number of the $v'H'$ events are seen to correspond to at least one of the two variables being in the top 5 percentiles of its distribution. In the figure, this corresponds to all data points not within the squares labelled S. Most extreme transport events therefore seem to be due to (near) in-phase velocity and MSE anomalies where at least one of the two variables qualifies as an extreme event. To put a number on this statement, it is found that only 3% to 8% of the $v'H'$ extreme events correspond to both velocity and MSE anomalies which are not in the top 5 percentiles of their respective distributions. Further proof of the in-phase view

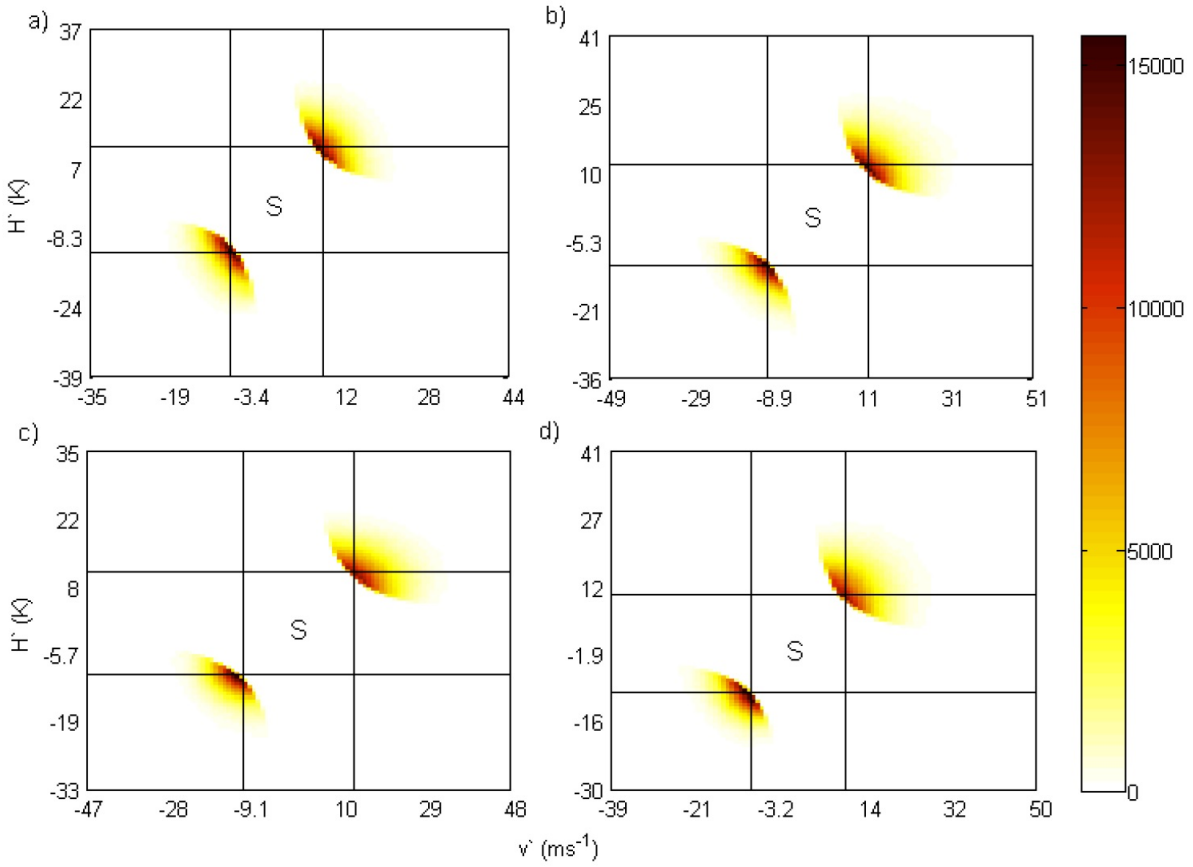

Figure 55: Map of v' and H' data points corresponding to the top five percentiles of $v'H'$ events for a) NH JJA, b) NH DJF, c) SH JJA and d) SH DJF. All 24 seasons (December 1993–August 2005) are taken into consideration. The scale of the colour bar corresponds to the number of data points per discrete bin. The continuous lines correspond to the values of the 5^{th} and 95^{th} percentiles of the v' and H' distributions. The squares labelled S contain all extreme $v'H'$ occurrences where neither v' nor H' are in the top/bottom five percentiles of their respective distributions [from *Messori and Czaja*, 2013b].

can be obtained by normalising velocity and MSE anomalies by their respective standard deviations. Covariance and correlation are related by the following expression:

$$cov(x, y) = corr(x, y)\sigma_x\sigma_y \tag{32}$$

where the correlation component is a measure of phase and the standard deviation one a measure of magnitude. If one normalises both anomalies by their respective standard deviations, and then computes the $v'H'$ PDF as in figure 32 on page 61, only the phase information will be conserved. The results (not shown) still show large skewness, even though it is smaller than that of their unnormalised counterparts. This confirms that the phase relationship between v' and H' is the key mechanism for generating extreme transport events. The simple wave picture suggested here will be discussed further and related to idealised models in Chapter V.

7 Discussion

The present study robustly demonstrates a new feature of meridional heat transport: namely that the majority of the transport is carried out by only a few occurrences every season, in easily identifiable bursts. This property can be clearly seen in figure 56 which illustrates, in a binary format, the sporadic nature of the heat transport process at a single point (50° N 0° E). Similar plots are obtained for other grid points (not shown). On a given day, a value of unity (vertical bar) is set to the curve if $v'H'$ falls in the top 5 percentiles of the distribution for the 50° N latitude circle, and a value of zero is used otherwise (no vertical bar). By definition, there are only a few extreme events every season, yet these account for a very large portion of the overall poleward heat transport at this location, sometimes exceeding 50% (the numbers for each winter are indicated in the top right corner of each panel)! Repeating this analysis for the number of extreme events per day along a full latitude circle, as shown in figure 45 on page 78, yields a similar picture, with bursts of extreme events lasting for a few days and typically involving less than 20% of the gridpoints at that latitude. This further illustrates the sensitivity of the atmosphere to very few, temporally and spatially localised features.

The analysis of the FORTE data, described in Section 4.4, implies that the sporadic nature of heat transport is not explained by the detailed features of the atmosphere's lower boundary (coastline,

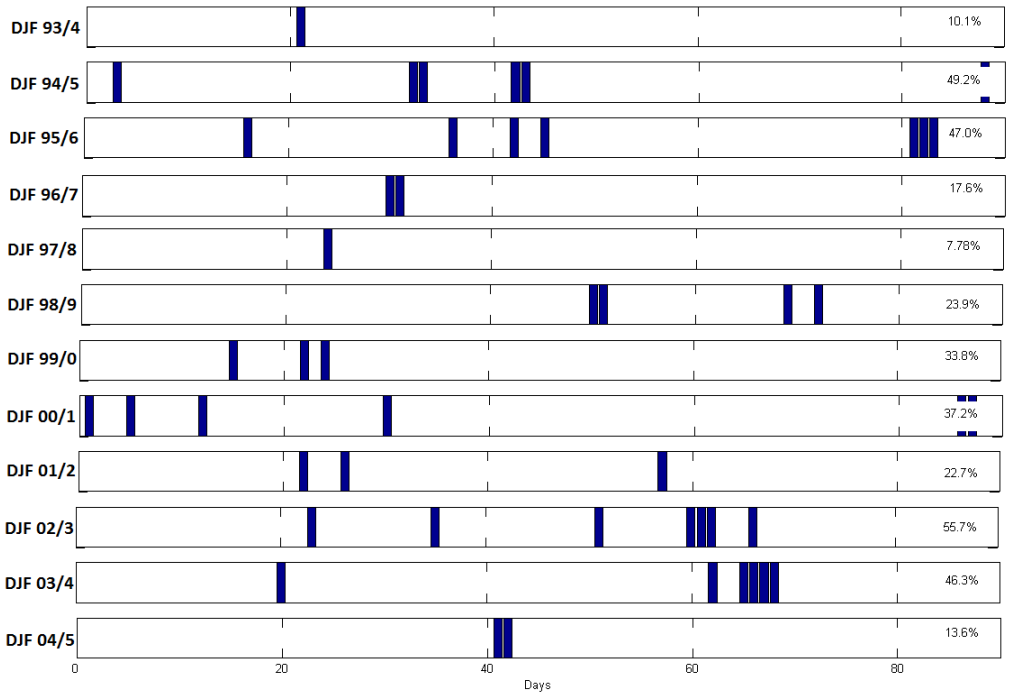

Figure 56: Bar plot of $v'H'$ extreme events at 50° N 0° E. Extreme events are defined as events in the top five percentiles of the $v'H'$ distribution for the 50° N latitude circle. The 12 panels correspond to DJF seasons from DJF 1993/1994 to DJF 2004/2005. Bars correspond to an extreme event occurring on a given day; the abscissa indicate the day of the season. The percentages in each panel indicate the contribution of the selected events to the net seasonal meridional transient-eddy heat transport. See text for further details [from *Messori and Czaja*, 2013b].

ocean fronts, sea ice, orography etc.), stationary waves and mesoscale features. Indeed, these are excluded by the coarse resolution of FORTE, but present in the reanalysis through data assimilation. Rather, this must be an intrinsic property of waves in the atmosphere.

Concerning the physical drivers of the extremes, the results found in this study are compatible with the traditional picture of growing systems having the minimal phase shift between v' and H', and accounting for the bulk of the heat transport [*Eady*, 1949]. At the same time, there is no attempt to discuss the specific motions in the v' and H' signals contributing to the extreme transport events. Furthermore, the considerations made above are based on a very idealised wave picture of heat transport. These results do not exclude that extreme events are simply growing baroclinic systems linked to local mesoscale phenomena, but it is not claimed that they demonstrate this either. This seems to conflict with the previous considerations deriving from the FORTE analysis. However, the fact that the statistical distribution of the transport can be reproduced without mesoscale features does not exclude that, in the real atmosphere, these might play some role in setting the location and initial development of large heat transport events. For example, the output from the the Minobe *et al.* model suggests that, while ocean fronts are not important in setting the magnitude of extreme events in heat transport, they might have an influence on their location. The strength of the extreme events was found to be almost unaffected when oceanic SSTs were smoothed, but the location of the extremes appears to be linked to oceanic fronts, analogously to what *Minobe et al.* found for rainfall. This consideration leads to possible speculations concerning the role of larger spatial scales, in addition to shorter baroclinic-scale motions, in driving the extremes. In fact, one would expect events entirely driven by atmospheric mesoscales to be more closely linked to the local state of the ocean, especially in the domains analysed by the Minobe *et al.* model, where the SST patterns have been shown to have a significant impact throughout the troposphere [*Minobe et al.*, 2008].

While no systematic tests were conducted, no clear link to a single physical scale or phenomenon emerges from the statistical analysis performed here. The only possible picture coherent with the speculations made seems to be one where local effects combine with the larger scales to drive large heat transport events. This idea is further analysed and discussed in Chapter IV.

A final note should be made on the different components of the meridional heat transport by transient motions. The moist static energy anomalies, H', are the sum of a humidity, a temperature and a geopotential term. These are taken collectively when computing the transport, but their

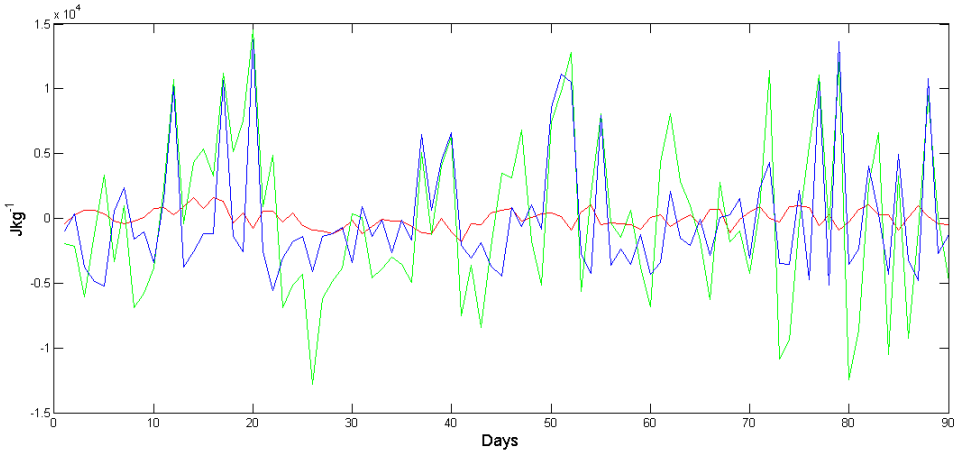

Figure 57: Time series of the humidity (blue), temperature (green) and geopotential (red) components of H'. The data considers fields at 40° N, 65° E during DJF 2000.

respective magnitudes differ significantly. Taking as reference anomalies[4] $q' = 5\ gkg^{-1}$, $T' = 10\ K$ and $z'_g = 100\ m$, an order of magnitude calculation yields:

$$H' = L_v q' + c_p T' + g z'_g \sim 10^4 Jkg^{-1} + 10^4 Jkg^{-1} + 10^3 Jkg^{-1} \tag{33}$$

The magnitudes of the humidity and temperature terms are comparable, while the geopotential term is smaller. An example of anomalies from the ERA-Interim data is shown in figure 57. The plot presents the time series of the humidity (blue), temperature (green) and geopotential (red) terms, at 40° N, 65° E during DJF 2000. The location was chosen to match an area of high extreme event occurrence, while the year was chosen randomly. The peak magnitudes of the anomalies match very well the order of magnitude calculation in equation 33, and clearly show the relative importance of the three terms. It should be noted that some days are dominated by temperature fluctuations, while others are dominated by humidity ones. The strongest H' values, which presumably correspond to large heat transports, will occur when peaks in the two anomalies are perfectly in phase. Figure 33 displays several instances of this behaviour. The figure focusses on a single point and is for illustrative purposes only - the magnitudes of the anomalies will change significantly across the seasons and the broad spatial domains considered in the present thesis. It should also be noted that, in a low-pressure

[4]The specific humidity value is taken from *Ahrens* [2007], the temperature and height are taken from reference plots in *Marshall and Plumb* [1997].

system, the geopotential term has a very different phase to the temperature and humidity ones, being $\pi/2$ out of phase. There is therefore an argument to exclude it all together from the calculations.

While not performed here, an analysis separating the temperature and humidity terms would be very interesting for investigating potential links between extreme heat transport events and moisture transport to the polar regions. Moisture is not a passive tracer; however, the *Swanson and Pierrehumbert* [1997] model discussed in Section 2.4 suggests that a large part of the remarkable statistical features of the heat transport PDFs (*e.g.* figures 32-35) might be expected to emerge also in tracer transports, as long as a damping term is present. In fact, the authors find that a damped advective model manages to reproduce most of the salient features of the observed heat transport distribution. This analogy will obviously not apply to undamped passive tracers, for which the considerations made in this chapter are largely unapplicable.

8 Conclusions

The present chapter studies meridional atmospheric heat transport due to transient eddies, focussing on low levels in the mid-latitudes. The analysis is in terms of the probability distribution functions of meridional velocity anomalies v', moist static energy anomalies H' and their product $v'H'$. Two outstanding features of the distribution of $v'H'$ are the near-zero most likely value and the very pronounced positive skewness. These appear to be robust features of the distribution and are only marginally affected by season, hemisphere or latitude. Outputs from an intermediate complexity climate model, run in an aquaplanet configuration, present similar characteristics. The magnitude of the extremes of the distribution also seems to be independent of oceanic SST gradients. These findings suggests that the shape of the $v'H'$ PDF is not solely due to mesoscale phenomena, stationary waves and complexities associated with surface boundary conditions. As a direct consequence of the distribution's skewness, the top 5% of $v'H'$ events accounts for over half of the net poleward heat transport by transient eddies. This large sensitivity to extremes was noted by *Swanson and Pierrehumbert* [1997] at three locations in the Northwest Pacific during winter. Here, it is being established robustly for all extra-tropical regions as well as winter and summer seasons. A zonal perspective is also proposed, and it is shown that there are extremes in zonally integrated heat transport. These are partly due to synchronised local extremes occurring around a given latitude circle, but also have

a significant contribution from increased mean transport at non-extreme locations.

The results presented suggests that the extreme events result from a near in-phase relationship between v' and H' anomalies of sizeable amplitude, as is typically expected of growing Eady waves. Indeed, further analysis shows that extreme events in heat transport have wave numbers and timescales compatible with Eady-type systems, although this finding is by no means conclusive. Regions where extreme events occur most frequently partly coincide with the storm track and WCB regions singled out in previous studies via Eady growth rate diagnostics [e.g. *Hoskins and Valdes*, 1990], or tracking algorithms [*e.g. Hoskins and Hodges*, 2002; *Hoskins and Hodges*, 2005; *Eckhardt et al.*, 2004; *Pfahl and Wernli*, 2012], although important differences emerge. It should further be noticed that this partial agreement does not exclude the role of broader spatial scales. An important consequence of the sensitivity of heat transport to extreme events is that a very large fraction of the transport occurs in a few discrete bursts, each lasting for only a couple of days. The transient heat transport process in mid-latitudes is therefore fundamentally sporadic in the temporal domain.

Part IV

Extreme Events in Atmospheric Heat Transport: Physical and Spectral Features

1 Aims of the Chapter

In the previous chapter we have established the importance of extreme events for setting the mean seasonal transport, and described the basic statistical features that characterise them. This allowed making a number of speculations about the physical drivers of the extremes, but few firm conclusions. The present chapter aims to extend the analysis of the physical origin of extremes by:

i) Providing insights into their relationship to local atmospheric dynamics and synoptic-scale analogues;

ii) Discussing the role of baroclinic time scales (here defined as 2.5-6 days) versus motions with longer periods.

In the spirit of the statistical analysis presented in Chapter III, the analysis of the local atmospheric dynamics will not be performed on individual case studies, but will be kept as general as possible. Points i) and ii) above, combined, will provide an answer to the very fundamental question of: *what are the heat transport extremes?*

2 Outline

Messori and Czaja [2013b] robustly established the sporadic nature of atmospheric heat transport by transient motions, and described the role of extremes in setting the mean seasonal transport. Furthermore, the authors found that the extreme occurrences displayed temporal and spatial scales compatible with phase-driven, baroclinic growing systems [*Eady*, 1949]. The exact role of baroclinic time scales in driving these extremes, however, was not quantified. Here, the term *baroclinic* will be applied to periods between 2.5 and 6 days, following the terminology introduced by *Blackmon et al.*

[1977].

The present chapter will analyse in more detail the atmospheric motions driving extreme event occurrences. Synoptic analogues will be proposed, and it will be shown that the v and H anomalies[5] leading to extreme events are not predominantly linked to baroclinic time scales. Moreover, the fractional contribution of these timescales to the heat transport occurring on extreme days is found to be essentially indistinguishable from that seen on other days. After robustly establishing the role of longer time scales, we will confirm that the key drivers of extreme heat transport are the phase and coherence relationships between the v and H anomalies.

First, an analysis of the circulation associated with extreme events is presented, for both extreme (Section 3.2) and median (Section 3.3) events, and for both global and regional (Section 3.4) composites. This includes an analysis of the different components of the heat transport and of the underlying wind fields. A physical interpretation is provided in Section 5. Next, the spectral features of the transport are analysed, using wavelet transforms (Section 4). A wave number decomposition is presented in Section 5, so as to relate the temporal scales to the spatial ones. The implications of the inferences made by analysing both local dynamics and the spectral decomposition are discussed in the same section. Last, the main conclusions reached in the present chapter will be summarised in Section 6.

3 Vertical and Zonal Structure of Heat Transport Extremes

3.1 Methodology

The data range and the computation of the heat transport by transient motions are analogous to those described in Chapter III, Section 3 on page 59 and in *Messori and Czaja* [2013b]. The extreme events are also defined as before, namely as the top 5 percentiles of the distribution of heat transport at 850 mb, for a given season and hemisphere. *Messori and Czaja* [2013b] have shown that the exact percentile chosen as threshold does not affect the characteristics of the extremes. Here, however, only the extremes corresponding to local maxima are selected, so that only a single point per event is retained. In fact, in choosing a percentile threshold to define extreme events, an extensive region of strong heat transport could contribute with multiple data points to the statistics. This is

[5] Details of how these anomalies are computed are provided in Section 3.

desirable when computing, for example, a climatology of heat transport bursts. However, it becomes problematic when analysing atmospheric circulation, since it would imply replicating several times the same circulation system associated with a single extreme event peak. Using the wave number values found in *Messori and Czaja* [2013b], 50 model grid boxes (corresponding to approximately 35° longitude), are retained on either side of the selected local maxima. This ensures that the full extent of the extreme events, including the surrounding circulation features, is captured. To avoid double-counting data points, if two successive extremes on the same day and latitude are less than 100 grid boxes apart, half of the data points in the interval are assigned to one of the extremes and half to the other. The heat transport is then computed across all the selected longitude data points, at all pressure levels present in the ERA-Interim dataset, namely 37 levels ranging from 1000 mb to 0 mb. Masking is applied where given pressure levels are below the topographical surface. This provides a pressure-longitude transport cross-section of each extreme. All the extremes thus analysed are then composited, and the values found are normalised by the number of data points being composited. Note that, since some extremes are less than 100 grid boxes apart, the normalisation factor will not be a constant across the composite. A similar procedure is applied in order to obtain cross-sectional plots of MSE anomalies, meridional velocity anomalies, and net values of the wind fields corresponding to extreme events. Analogous plots are also produced for the temporal dimension. Here, the two days before and after each local extreme are selected, and only the single latitude-longitude coordinate corresponding to the maximum is retained. The transport values of all local extremes are then composited and normalised, yielding the corresponding pressure-time composite colour maps.

Note that no vertical integration is performed over the quantities analysed here. This means that heat transport cannot be expressed in W, but at the same time avoids problems associated with mass non-conservation in the ERA-Interim data set. Even though the mass imbalance should not affect a calculation based on anomalies, the reanalysis estimates might not be accurate for the whole vertical extent between two adjacent pressure levels. Indeed, as noted by *Trenberth* [1991], the values archived in the ECMWF reanalyses should be interpreted as "the most accurate values available at those levels, but ... not representative of layers". Such issue has vastly improved in the passage from ERA-40 to ERA-Interim, but is still present in the latter data set and should not be ignored [e.g. *Graversen et al.*, 2011].

Since the atmospheric circulation is analysed in composite plots, one needs to ensure that the

Domain Name	Abbreviation	Boundaries	Seasons
Gulf Stream	GS	30° N−55° N; 265° E−335° E;	DJF, JJA
Pacific Storm Track	PS	30° N−50° N; 150° E−230° E;	DJF, JJA
Bering Strait/Gulf of Alaska	BS	55° N−70° N; 180° E−200° E;	DJF
Nordic Seas	NS	55° N−80° N; 335° E−15° E;	DJF
Siberian Plateau	SP	60° N−75° N; 40° E−90° E;	JJA
Southern Ocean DJF	SO	35° S−55° S; 295° E−275° E;	DJF
Souther Ocean JJA	SO	40° S−70° S; 295° E−275° E;	JJA

Table 5: Names, abbreviations and boundaries of the domains used in the analysis of extreme event regional composites. Where a domain is only considered for one season, this is specified in the fourth column of the table.

mean picture represents individual events well. Standard deviation maps have been computed to verify that this is indeed the case. The transport standard deviation is very similar in all four season/hemisphere combinations, and is more than an order of magnitude smaller than the transport values (not shown). The variability peaks at high levels ($\sim$300 mb), which correspond to the typical location of the Jet Stream. The standard deviations of the velocity fields, though larger in proportion, are also significantly smaller than the fields themselves. The meridional velocity anomaly standard deviation is, instead, comparable in magnitude to the velocity anomalies themselves, indicating a very large variability from one event to the next. Again, the peak variability is at altitude. The MSE standard deviation is also of the same order of magnitude as the MSE anomalies, but peaks at lower levels. We therefore have a very large variability in the individual components of the transport, while the overall variability in the transport itself is very small when compared to the latter's magnitude. This suggests that, while the transport and velocity composites can be interpreted confidently, care should be taken when looking at the velocity and MSE anomalies. It further implies that extreme events occurring in different parts of the globe might be similar in terms of heat transport values, but might be driven by very different local wind, temperature and moisture anomalies. This motivates performing a regional analysis, so as to identify the circulation associated with extreme events in different areas of the globe.

The regional domains are chosen based on the areas of high extreme event occurrence identified in the climatology of extremes, shown in figure 51 on page 87. Seven domains are selected; their names, abbreviations and geographical boundaries are listed in table 5. Note that, due to seasonal differences in the distribution of extreme events, three of the domains (BS, NS and SP), are only examined during one season (DJF, DJF and JJA respectively). Even though there are seasonal

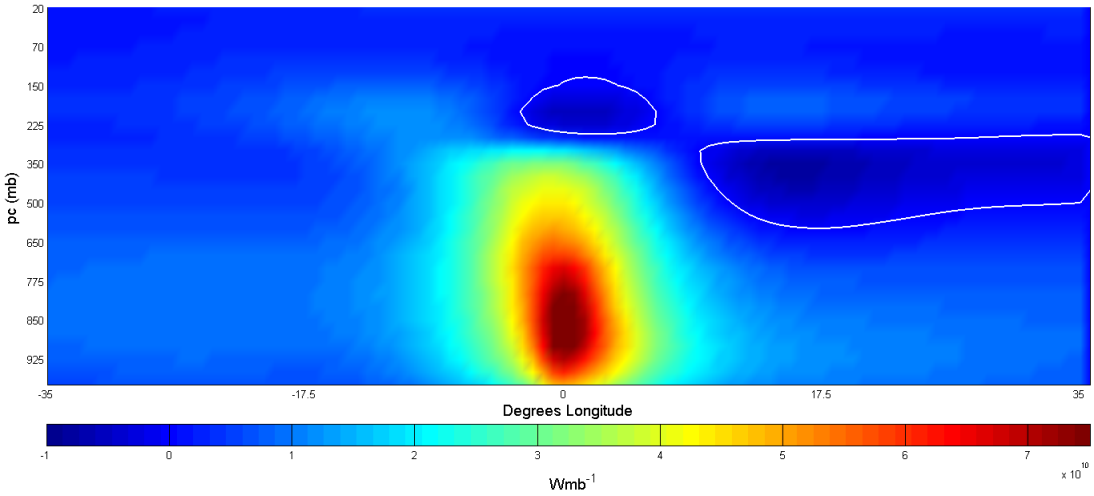

Figure 58: Composite pressure vs. longitude meridional heat transport colour map for extreme events during NH DJFs, in Wmb^{-1} . The data cover the period from December 1993 to February 2005. All latitude circles between 30° N and 89° N are taken into account. The white contours correspond to zero meridional heat transport. Note that the colourbar is not symmetric about zero.

variations in the distribution of extremes across all domains, the boundaries are defined so as to be kept constant for both seasons. The only case for which this isn't possible is the SO domain, where the marked poleward shift of the high frequency belt during the austral winter implies that a shift in the domain boundaries is also needed. At the same time, it should be noted that the exact position of domain boundaries is not so important. Indeed, the data analysed only includes extreme events, so if a domain also includes areas with a low event frequency, these will count very little towards the regional composites. The standard deviations of the velocity and MSE anomalies were generally found to be smaller than for the hemispheric composites, although some domains present significant variability in specific variables. Where this is the case, it is highlighted in the analysis.

3.2 Zonal Cross-Sections of Extreme Events

To investigate the vertical and zonal structure of transient-eddy heat transport extremes, we begin by computing composite transport colour maps for all season/hemisphere combinations. These take into account all available latitude bands (30°–89°) and time series (1993/1994–2004/2005). Figure 58 shows the composite maps for events in NH DJF. The other season/hemisphere combinations (not shown) yield almost identical maps. As would be expected from the definition of extreme events used

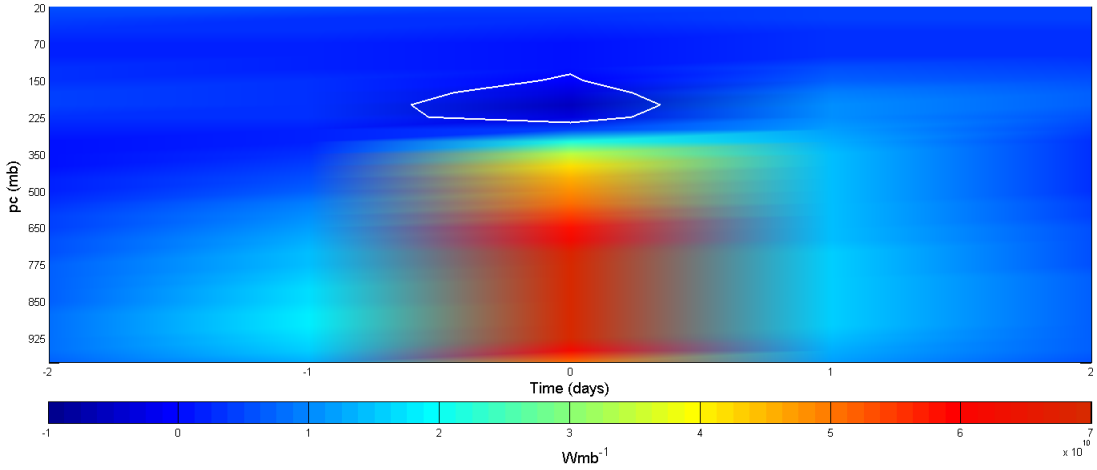

Figure 59: Composite pressure vs. time transport colour map for extreme events during NH DJFs. The data cover the same range as in figure 58. The white contours correspond to zero meridional heat transport. Note that the colourbar is not symmetric about zero.

here, the peak transport is found around 850 mb. The general spatial structure of the extreme events seems to be that of a deep vertical column of poleward transport flanked by weaker equatorward transport regions to the east and at high levels. The whole pattern displays a small westward tilt. Even though the area of poleward heat transport is very extensive, the core of the extreme event is quite narrow, covering on average only a few degrees longitude. The spatial scales of the transport are generally consistent with the conclusions drawn in Section 5, Chapter III, which found that the full extent of an extreme event, including the possible recirculation features, typically corresponds to wavenumber 8 (or, equivalently, 45 degrees longitude).

Looking at the corresponding pressure-time plots, no striking features emerge. All four season/hemisphere combinations display a similar pattern, with an increase in transport during the day before the extreme event, followed by a rapid decrease thereafter. The structure is roughly symmetric about the maximum, and the peak at $t = 0$ dominates strongly over the other days. While there is no evidence of a return transport in time, the upper level return seen in the zonal cross-sections is instead captured, and is found to be limited to the day at which the poleward transport peaks. As example, the NH DJF temporal composite is shown in figure 59.

Due to the large standard deviations in v' and H', discussed in Section 3.1, the composites for these variables are not analysed on a hemispheric scale. Rather, their discussion is deferred to Section

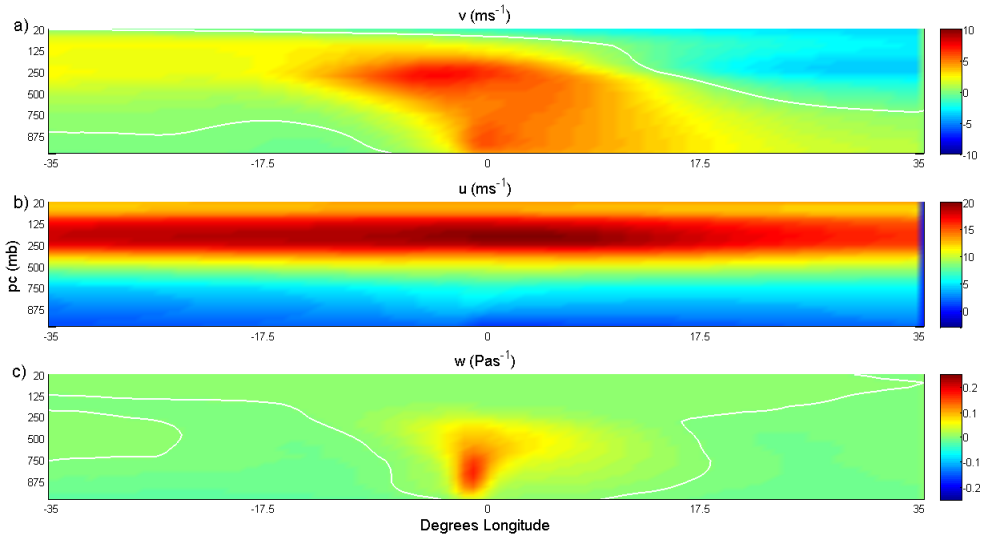

Figure 60: Composite pressure vs. longitude wind field map for extreme events during NH DJFs. u and v are in ms^{-1}; w is in Pas^{-1}. The data cover the same range as in figure 58. The top panel depicts meridional velocity, the middle one zonal velocity and the bottom one vertical velocity. Meridional, zonal and vertical velocities are positive in the polewards, eastwards and upwards directions, respectively. The white lines in all panels are zero velocity contours. Note that the zonal velocity colourbar is not symmetric about zero.

3.4, dealing with regional domains, where their implications can be better assessed.

Next, maps analogous to that in figure 58 are produced for individual velocity components, in order to reconstruct the circulation corresponding to the extreme events. Figures 60, 61, 62 and 63 show the u, v and w wind fields in the longitude-pressure plane, for events in NH DJF and JJA, and SH DJF and JJA, respectively. The top panels represent meridional velocity, the middle ones zonal velocity and the bottom ones vertical velocity. Meridional, zonal and vertical velocities are positive in the polewards, eastwards and upwards directions, respectively. The white lines in all panels are zero velocity contours.

Looking at the figures, it is immediately clear that, regardless of hemisphere and season, extreme events are characterised by a strongly ascending air stream corresponding to the core of the event, flanked by two regions of subsiding air. As for what concerns the zonal wind, extremes systematically correspond to eastward flow, in agreement with the prevailing mid-latitude westerlies. The meridional velocity pattern is very consistent throughout all seasons and hemispheres, and is characterised by a core of strong poleward velocity, corresponding to the location of the extreme, surrounded by a strong equatorward flow on the eastern flank and a weaker one on the western flank. There are,

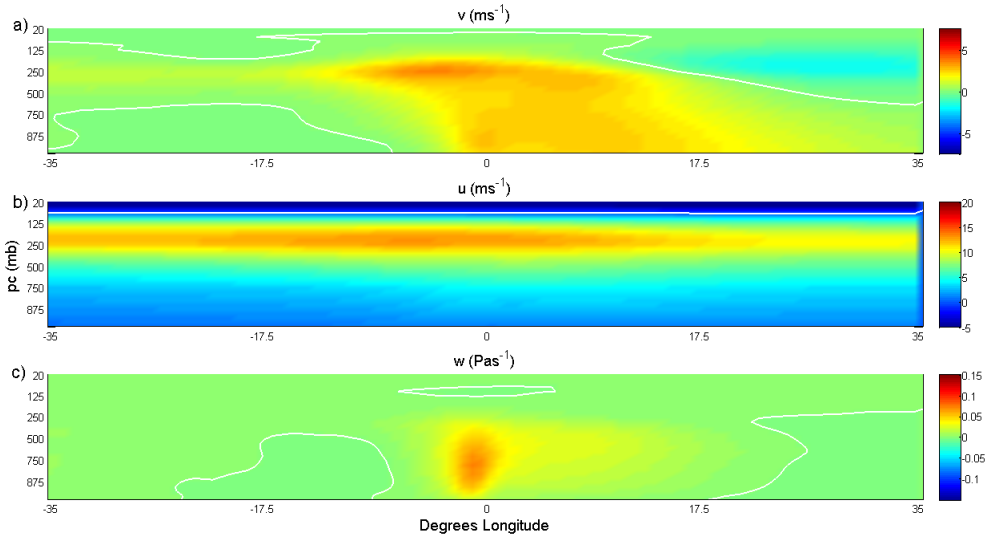

Figure 61: Same as figure 60 but for NH JJAs. The data now cover the period from June 1994 to August 2005.

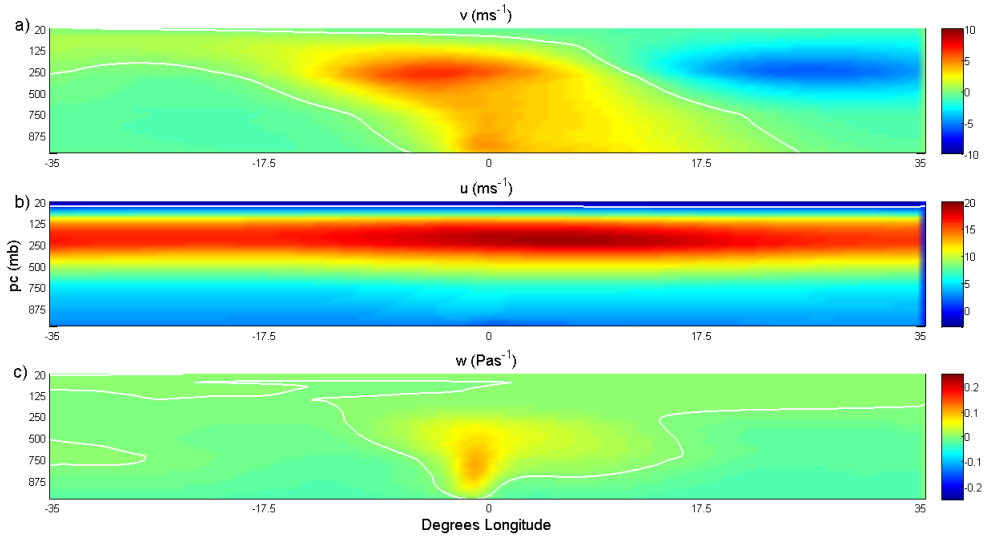

Figure 62: Same as figure 60 but for SH DJFs. The data now cover all latitude circles between 30° S and 89° S.

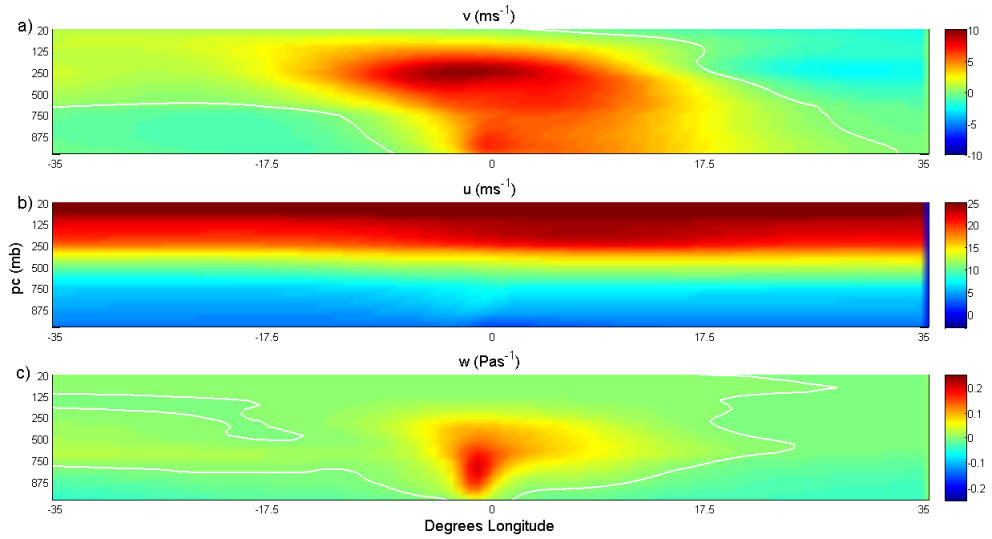

Figure 63: Same as figure 62, but for SH JJAs.

obviously, seasonal and hemispheric variations in the strength and positions of the patterns described above. It is interesting to note that a persistent feature of the maps is a high-altitude meridional jet at $\sim$200 mb, which is slightly higher than the typical position of the Jet Stream. It should further be noted that not all features of ERA-Interim's stratospheric circulation are realistic [*Dee et al.*, 2011].

It is immediately evident that there are a number of shared features between the v and transport colour map shown in figure 58. In particular, the poleward transport seems to be associated with strong poleward velocities, and the opposite for the equatorward one. However, these similarities should be interpreted with care because the transport is driven by velocity anomalies, as opposed to absolute values. Indeed, a poleward velocity does not automatically translate into a positive v', just as an equatorward flow should not be equated with a negative velocity anomaly. In addition to this, the equatorward heat transport seen at upper levels is not matched by a corresponding negative meridional velocity.

Further information concerning the role of the velocity and MSE anomalies, and their relation to the circulation features described above, can be obtained by analysing composite covariance colour maps. The composite covariance (which should not be confused with the temporal covariance) at every grid box of the zonal cross-section is computed as:

$$cov(v',\, H') = \overline{v'H'} - \overline{v'}\,\overline{H'}, \tag{34}$$

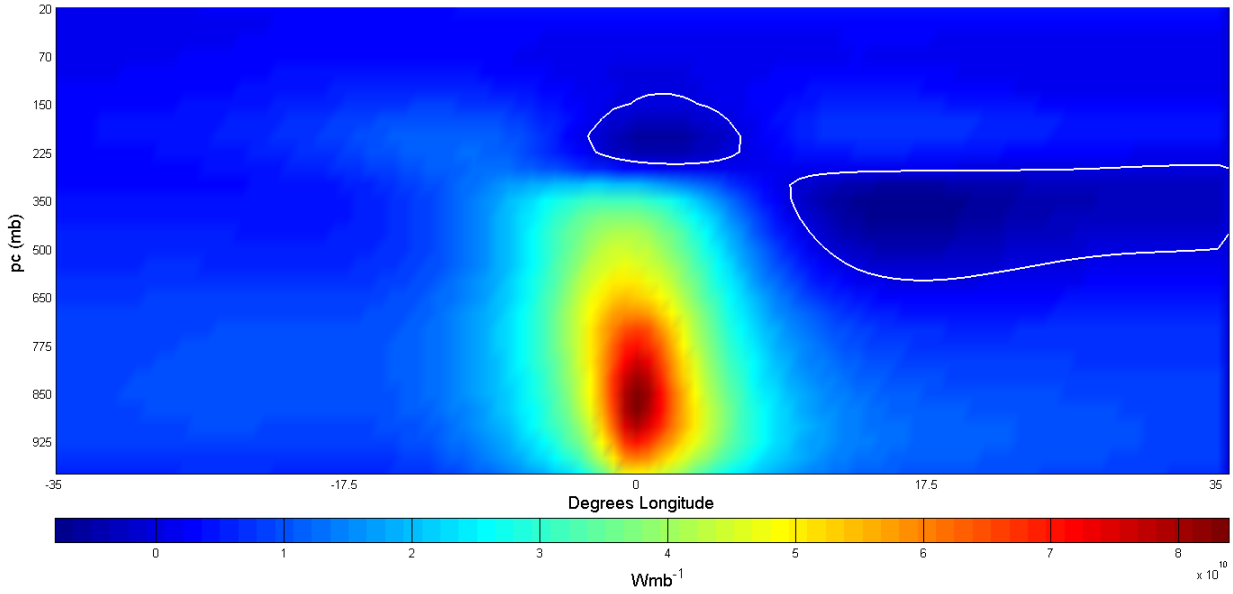

Figure 64: Composite pressure vs. longitude composite covariance map for extreme events during NH DJFs. The data cover the same range as in figure 58. The white lines correspond to zero composite covariance contours. Note that the colourbar is not symmetric about zero.

where the overbars denote an average over all extreme events. When large values of one variable correspond to large values of the other variable, and the same holds for the smaller values (*i.e.* the variables tend to show a related behavior), the composite covariance is positive. In the opposite case, when large values of one variable mainly correspond to small values of the other, the composite covariance is negative. The sign of the composite covariance therefore provides information on the linear relationship between the variables. The extreme event composite covariance maps are very similar for all four season/hemisphere combinations. As an example, figure 64 shows the one for NH DJF. The white lines mark contours of zero composite covariance. Positive composite covariance corresponds to poleward transport and negative composite covariance (albeit much smaller than the positive one) corresponds to equatorward transport. This indicates that extreme events in poleward transport come about as a phase-driven effect, because the composite covariance is highest (most positive) where the poleward transport peaks and *vice-versa* for the equatorward transport. Note that, since composite covariance is not necessarily related to the sign of each variable, this result is not obvious *a priori*. Indeed, it is very easy to construct an artificial dataset where poleward transport is generated by v' and H' arrays displaying negative composite covariance. This result also suggests

Value	v'	H'
1	+	+
2	+	-
3	-	+
4	-	-

Table 6: The four combinations of v' and H' signs and the corresponding values, as assigned in figure 65.

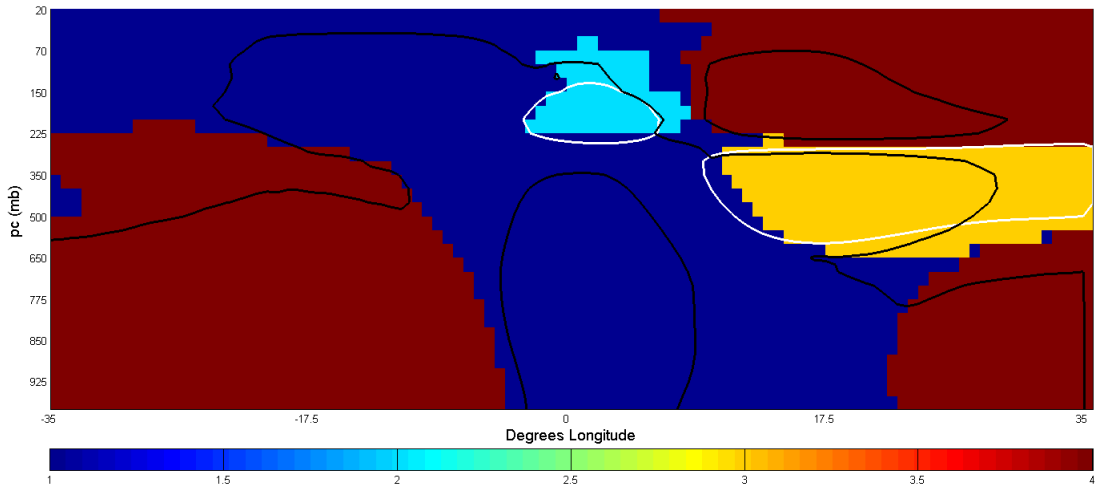

Figure 65: Composite v' and H' sign combination colour map for extreme events during NH DJFs. The data cover the same range as in figure 58. The black contours correspond to 30% and 50% of extreme events sharing the same sign combination. The white contours correspond to zero meridional heat transport.

that the return equatorward transport associated with extreme events does not follow a phase-driven (in this case anti-phase) mechanism as closely as the poleward one, otherwise the negative composite covariance would be as large as the positive one.

To further investigate the $v'H'$ combinations driving poleward and equatorward transports, one can plot a sign covariance matrix. Here, the sign of v' and H' at each grid box of the cross section is recorded. A value between 1 and 4 is then assigned to each grid box depending on the sign combination of the two variables. The sign combinations, and the corresponding values, are summarised in table 6. The value which occurs most frequently at a given grid box, for all extreme events in a given season and hemisphere, is assigned to that grid box. In all seasons and hemispheres, the poleward transport corresponds to a ++ case, with the core corresponding to more than 50% of occurrences having this sign combination. The - - combination, on the other hand, accounts for the weaker background flow

areas, but with a lower frequency of occurrence. An equally consistent pattern is seen for the two areas of equatorward flow (the one aloft and the one to the east of the main transport maximum), which are due to different sign combinations. The upper level flow is due to a +- case, with cold, dry air being advected polewards, while the eastern flow is due to a -+ case, where warm, moist air is being advected equatorward. Except for the ++ combination, these patterns are typically indicative of no more than half of the extremes, as could be guessed from the high standard deviation of the velocity and MSE anomalies, discussed in Section 3.1. As example, figure 65 shows the sign combinations for NH DJF. The black contours show areas where the most frequent values correspond to more than 30% and 50% of extreme days. We attempt to give a physical interpretation of these results in Section 5, but it should be kept in mind that they only represent a limited number of extremes.

Due to the significant contribution of both ++ and - - cases to the extremes, it is of interest to investigate whether the two sign combinations correspond to radically different transport profiles. Figures 66 and 67 show the composite transport maps for ++ and - - extreme events respectively, during NH DJFs. Compared to the overall composite, shown in figure 58, there are some immediate differences. The ++ combinations display a stronger transport at the core of the event, and do not have any equatorward transport regions. The - - events, on the opposite, have a much weaker poleward transport maximum and display significant regions of equatorward transport. In terms of the structure of the poleward maximum, namely a well defined central core and a westward tilt with height, no major variations emerge between the three plots. The only difference worth noting is that the - - core is slightly broader than its ++ counterpart, which is consistent with an extended region of equatorward flow of cold air.

The single-sign analysis can also be repeated for the individual velocity components associated with the extremes. Figures 68 and 69 show the composite velocity maps for ++ and - - extreme events respectively, during NH DJFs. Panels a), b) and c) show the u, v and w wind fields in the longitude-pressure plane, analogously to figure 60. As was the case for the transport, the single-sign plots display some immediate differences when compared to the overall composite. The ++ events have a purely positive meridional velocity field, suggesting that positive v anomalies generally correspond to poleward absolute velocities. The vertical velocity field is also almost entirely positive, unlike what seen in the overall composite. The features of the velocity fields at the core of the poleward heat

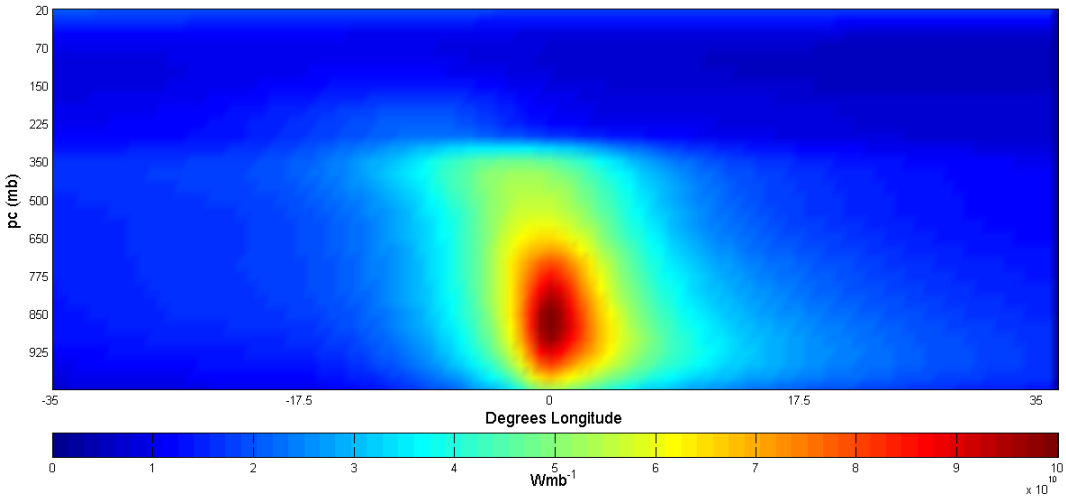

Figure 66: Same as figure 58, but for ++ extreme events only. Note that the colourbar is positive only.

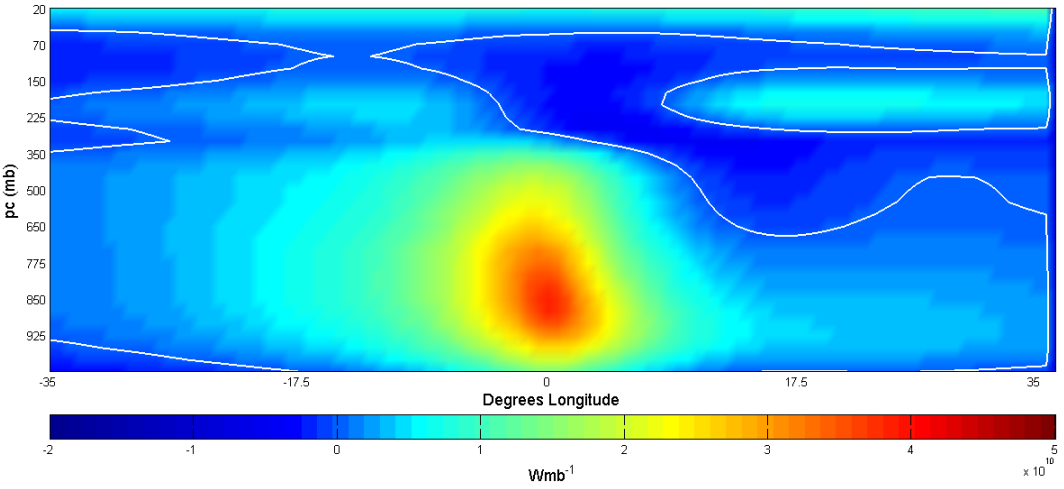

Figure 67: Same as figure 58, but for - - extreme events only. The white contours correspond to zero meridional heat transport. Note that the colourbar is not symmetric about zero.

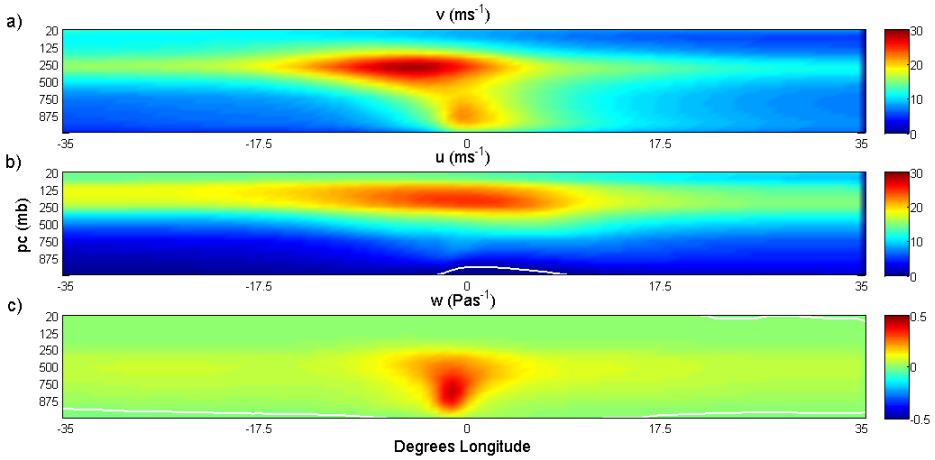

Figure 68: Same as figure 60, but for ++ extreme events only during NH DJFs. The white contours correspond to zero velocity. Note that the meridional and zonal velocity colourbars are not symmetric about zero.

transport are, however, very similar to those seen in figure 60. Overall, the picture is consistent with warm, rising air parcels being advected polewards and eastwards, much like in the conveyor belt schematic discussed in Section 1.3.4, Chapter II.

The - - composites display weaker fields than their ++ counterparts, and also differ more from the overall composites. Again, the meridional velocity sign matches that of the corresponding anomalies. Instead of having a central core of strong velocity, the peak velocities are found on the two flanks of the transport extreme. The same applies to the zonal wind, which has two peaks in correspondence with the meridional velocity ones. The vertical velocity corresponds to a subsiding air mass, again in stark contrast to the overall composite seen in figure 60. This is consistent with cold, dry, subsiding air parcels being advected equatorwards, as could be expected in the lee of a low pressure system. In this picture, the strong upper-level jet would correspond to the back side of a large trough in the jet stream. What is more surprising is that a second equatorward jet is seen to the east of the location of extreme heat transport. This could perhaps correspond to the circulation induced by the low pressure system preceding the one associated with the large heat transport. A second unexpected feature is the lack of poleward velocities inbetween the two equatorward jets. Even though the events have specifically been selected to have negative v', a weak positive v associated with the low pressure circulation would still be expected. A possible explanation could be that the process of compositing different events smears the weak positive velocities, which do not show up in the final average.

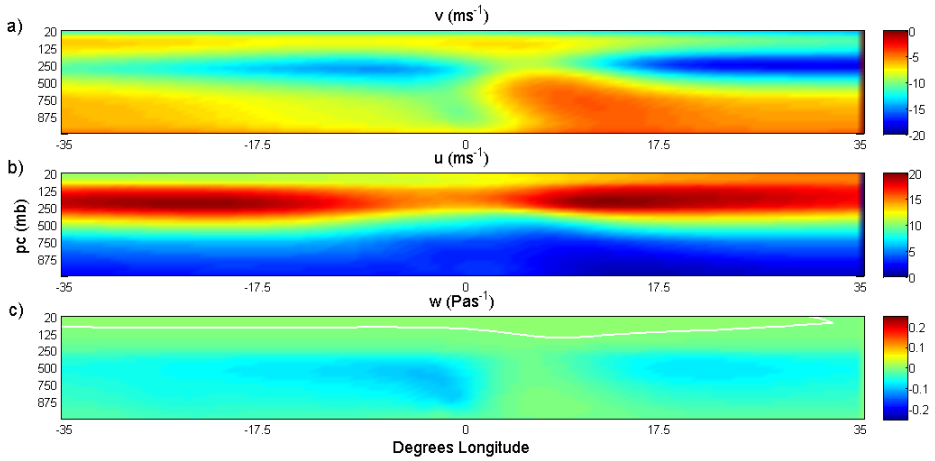

Figure 69: Same as figure 68, but for - - extreme events only during NH DJFs.

The other season-hemisphere combinations, although they differ in the fine features, present a similar overall picture. The single-sign composites highlight how the structure of the transport extremes is very similar, regardless of the underlying anomaly combination driving them. The overall transport composite therefore gives a very good representation of the typical extreme event. The velocity fields, on the opposite, present significantly different profiles. In part, this is to be expected because meridional velocity has a direct link to the velocity anomalies driving the extremes. However, it also suggests that the full composite plots give a good representation of the ++ events, which are stronger and more frequent, while they largely overlook the - - ones. The analysis in the following sections will focus on the overall composites, except for specific domains where the - - events are particularly frequent.

3.3 Zonal Cross-Sections of Median Events

The extreme event composites, discussed above, can be compared to the corresponding plots computed for median events. These are defined as events within 5 percentiles of the median of the poleward-only transport distribution. This will enable to verify the extent to which the patterns discussed above are specific to the extremes. The composite for median events during NH DJF is shown in figure 70, as indication of the typical aspect of these events. Due to the nature of the transport PDFs (see figures 32 - 38 on page 67) , it is clear that median events will be associated with a very small heat transport. In fact, figure 70 displays low transport values at the selected median events, flanked by

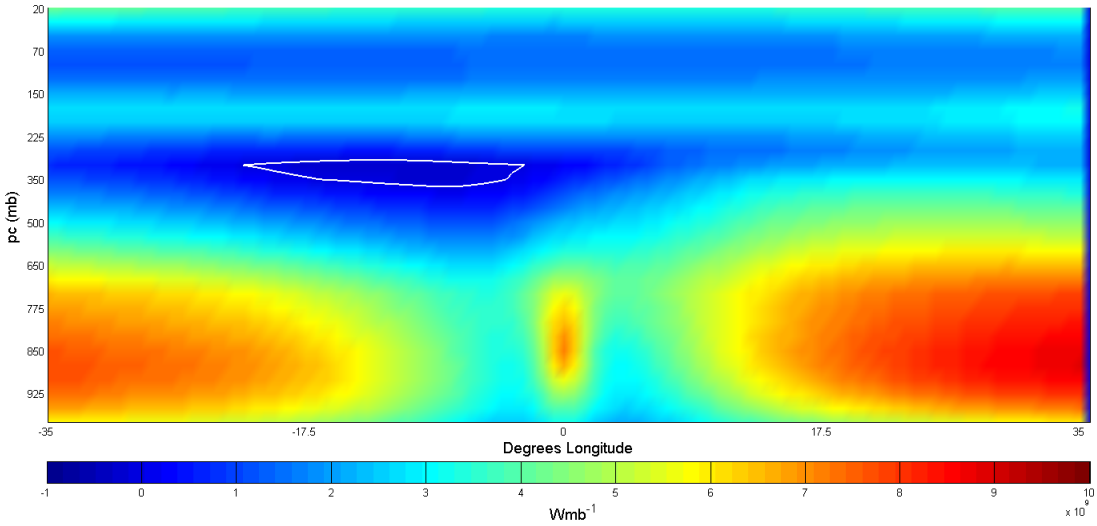

Figure 70: Same as figure 58, but for median events during NH DJFs.

two narrow bands of near-zero transport. These, in turn, are flanked by higher transport values. In all seasons and hemispheres, except for NH JJA, a weak equatorward transport nucleus is seen above the median event, at around 350 mb. The weak transport zones that flank the median events are simply due to the fact that, just as was done for the extremes, the median events are also selected as local maxima in the $v'H'$ signal. It is therefore inbuilt in their definition that the transport values immediately adjacent will be smaller than the local maximum. However, moving further away from the selected events the transport will, on average, return to its mean value. This is clearly of a similar magnitude to the median events.

Next, wind fields corresponding to the median events are analysed. Again, these are computed using exactly the same procedure adopted for the extremes. In contrast to the extended column of ascending air which characterised the latter, the median events correspond to weak, positive w near the surface and near-zero velocities further aloft. The meridional velocity fields display a dipole structure, with weak equatorward motion matching the central area of poleward heat transport. Figure 71 shows the wind composite for NH DJF, as indication of the typical flow pattern associated with median events.

Having found that the structure of the extremes is very different from that of the median days, it is interesting to verify whether the extremes are unique even relative to below-threshold strong poleward transport events. To test this, the transport and wind field composites shown above are

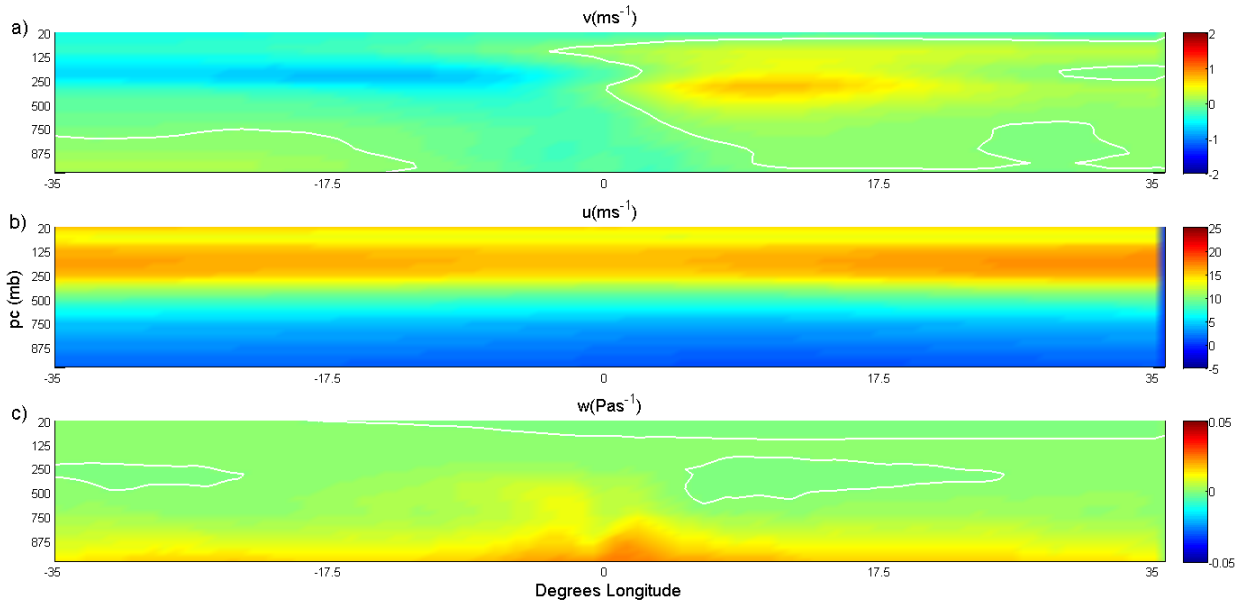

Figure 71: Same as figure 60, but for median events during NH DJFs.

replicated for events corresponding to different percentiles of the poleward transport distribution. Intervals of 5 percentiles are again selected, so as to compare composites representing equal numbers of events. The 60^{th}-65^{th} percentile events look very similar to the median events, both in terms of structure of the transport and of wind fields (not shown). The 80^{th}-85^{th} percentile events for NH DJF start looking more similar to the transport extremes, with a clear poleward transport core and an ascending airstream roughly matching the position of the largest heat transport (not shown). For the other season/hemisphere composites, no strong ascent is seen in correspondence with the transport. Moreover, there is no evidence of equatorward transport regions, and the meridional velocity sign patterns look extremely different to those seen in figure 65. The 85^{th}-90^{th} percentiles look similar to the extreme events for all seasons and hemispheres, but still miss key features such as the equatorward transport and the negative v' anomaly on the transport's eastern flank, seen in figure 65. As illustration, the heat transport composite for NH DJF is shown in figure 72. The 90^{th}-95^{th} percentiles, although displaying some minor differences (especially in terms of v'), finally match most of the features of the extreme event maps. The structure of the extreme events, discussed in Section 3.2, is therefore specific to only the highest percentiles of the heat transport distribution.

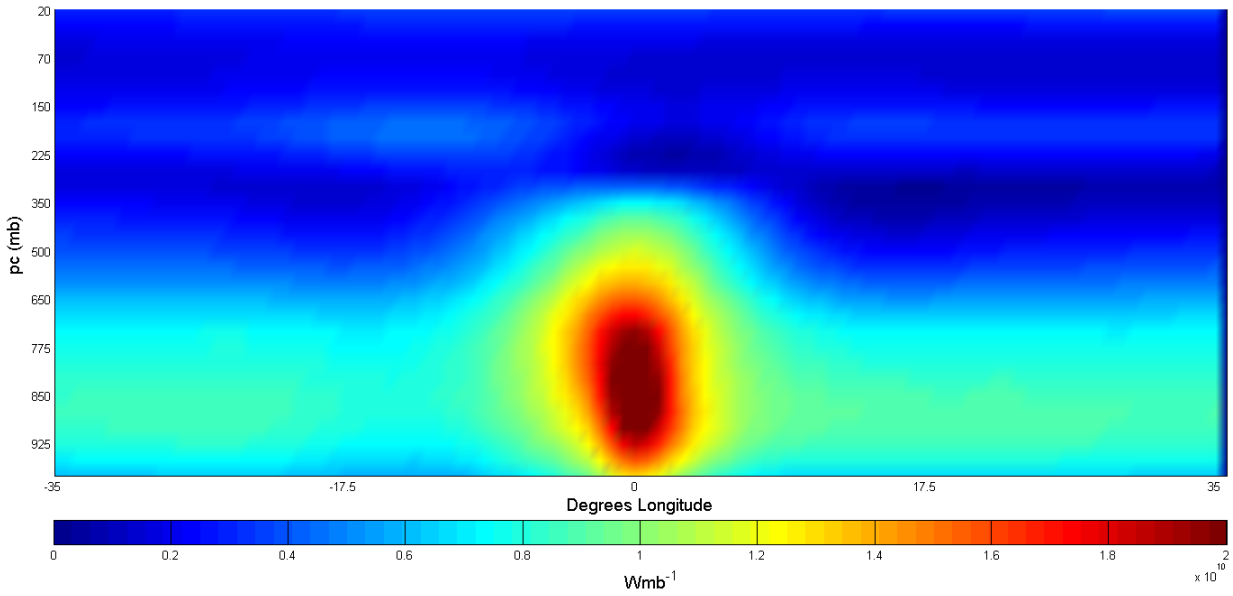

Figure 72: Same as figure 58, but for events in the 85^{th}-90^{th} percentiles of the heat transport distribution, during NH DJFs.

3.4 Regional Domains

In the interest of conciseness, we do not present the cross-sectional plots for all domains, but only highlight the features that will be later relevant to the physical interpretation of the events, discussed in Section 5.

As expected from the low standard deviation of the hemispheric transport composites, discussed in section 3.1, all regional domains display a very similar heat transport picture. There are, however, significant variations in the strength of the extremes, with the ones in the storm tracks being more intense than the ones at higher latitudes. Another minor difference between regions is that three display, in addition to the eastward flank return flow, a small return flow area on the western flank (BS and NS DJF, PS JJA). This feature is not seen in any of the hemispheric composites. As an example, the transport colour map for NS DJF is shown in figure 73.

Concerning the winds, in all the domains the extremes correspond to a rapidly ascending airstream. The GS and PS domains display similar circulations during DJF, although the PS one has a more extended poleward velocity field; the composite plots for the latter domain are shown in figure 74. This is very similar to the full NH composite. The same pattern, albeit noisier and with smaller velocities, is seen for PS JJA. In the GS domain, on the other hand, the summer months are

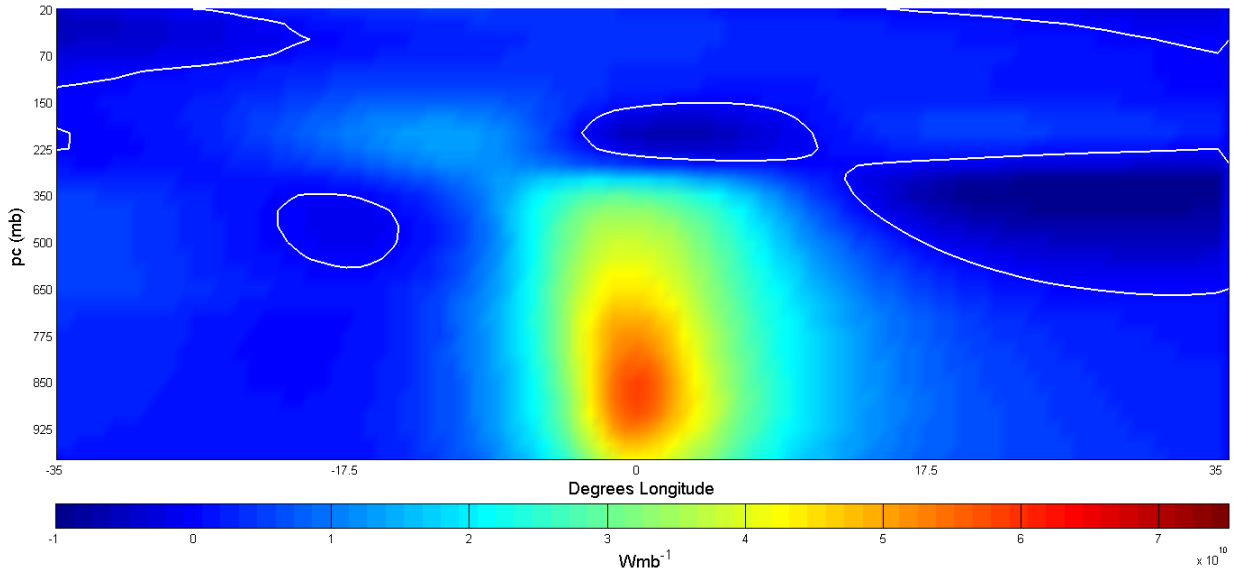

Figure 73: Same as figure 58, but for extreme events during DJFs in the the Nordic Seas (NS) domain.

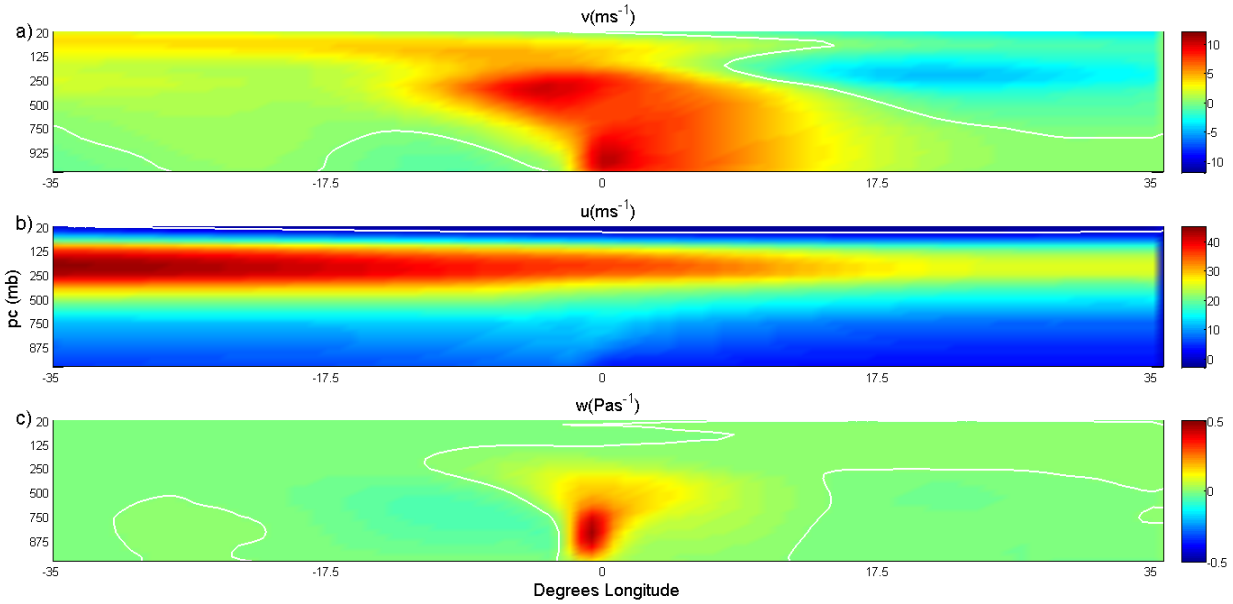

Figure 74: Same as figure 60 but for extreme events during DJFs in the Pacific storm track (PS) domain.

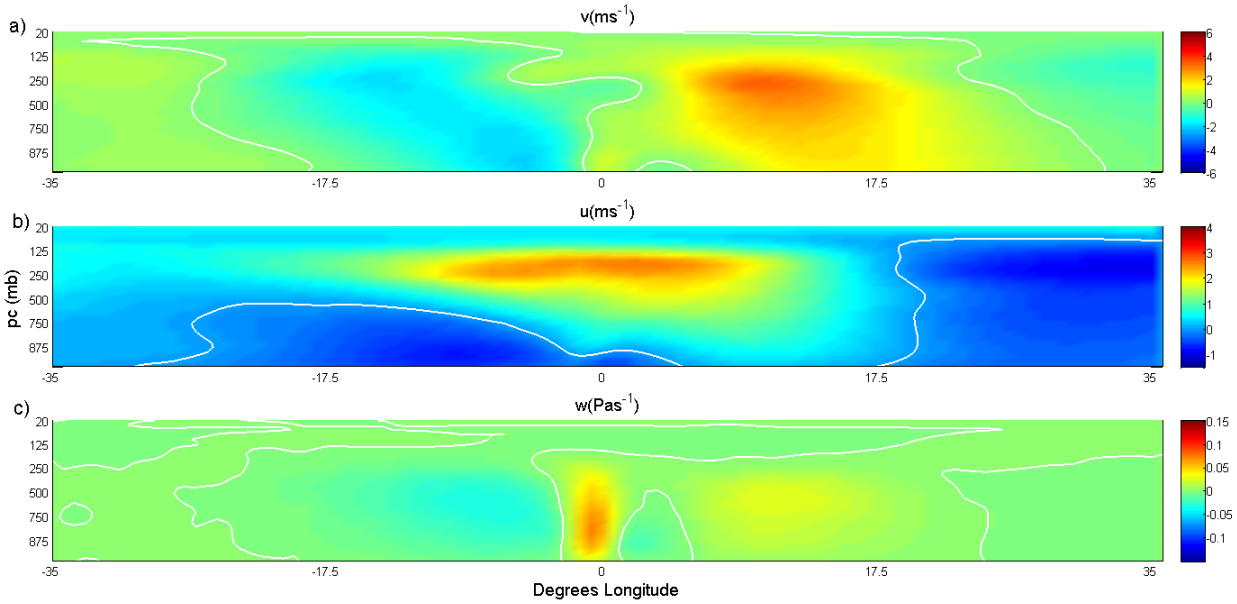

Figure 75: Same as figure 60 but for extreme events during JJAs in the Gulf Stream (GS) domain.

characterised by a large area of north-westerly flow. The u, v and w fields for GS JJA are shown in figure 75. The meridional velocity pattern is strongly reminiscent of the median event composite shown in figure 71, although here the location of the event corresponds to a poleward velocity. It should be noticed that the standard deviation of the latter map is large relative to the magnitude of the velocities. The SO domain is very similar to the PS one, although the JJA zonal velocity picture is less uniform due to the poleward shift of the extreme event band, which places it beyond the latitudes at which the low-level westerlies peak [*Risien and Chelton*, 2008]. NS DJF displays a very longitudinally extended high level poleward flow west of the location of the extremes, but otherwise again presents a picture similar to the wintertime storm track domains, as does SP JJA. The third seasonal domain, BS DJF, displays a low-level north-easterly wind associated with extended upwelling.

Next, the regional v' and H' composites are presented. Theses were not analysed for the hemispheric domains due to the high standard deviations found. The v' composites are found to match closely the v maps discussed above, and will not be analysed in detail. Figure 76 shows the v' composite for PS DJF, which can be compared to the v map shown in figure 74a. The v' values may seem very small compared to those in the tails of the PDFs shown in figures 32 - 35 on page 64.

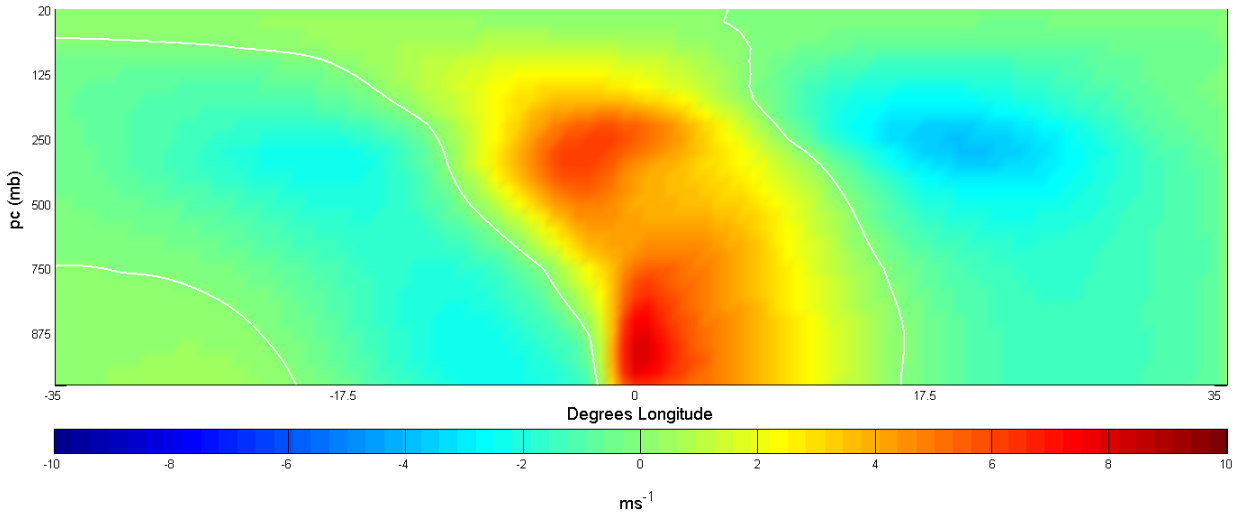

Figure 76: Composite pressure vs. longitude v' colour map for extreme events during DJFs in the Pacific storm track (PS) domain. The white contours correspond to zero v'.

However, it should be kept in mind that the typical v's associated with heat transport extremes do not necessarily correspond to the extremes of the v' PDFs. Taking as example NH DJF, the bulk of the extreme heat transport events corresponds to v anomalies of order 10 ms^{-1}, as illustrated in figure 55 on page 92. The anomalies found in the composite are therefore seen to be compatible with the results discussed in Chapter III.

Concerning the MSE anomalies, during DJF both GS and PS show a clear structure, with a high H' core at the location of the extremes and a negative anomaly intruding at altitude from the east. The western flank is characterised by an area of significant negative anomaly. The composite for PS DJF is shown in figure 77. The JJA composites display, in both domains, a greater prevalence of negative anomalies, especially for GS where the low-level H' is negative even at the extreme event location. This composite, however, is characterised by a high standard deviation. For the SO domain, the pattern in both seasons is very similar to the one shown in figure 77. The same applies to the NS and SP maps. The BS map shows a very weak but diffuse positive MSE anomaly, and a large standard deviation.

The sign combination maps (analogous to that shown in figure 65 on page 109) for the GS and PS regions, during DJF, show an almost identical pattern to the full NH hemisphere one. The same applies, during both seasons, to the SO domain and the full SH composites. The GS summer

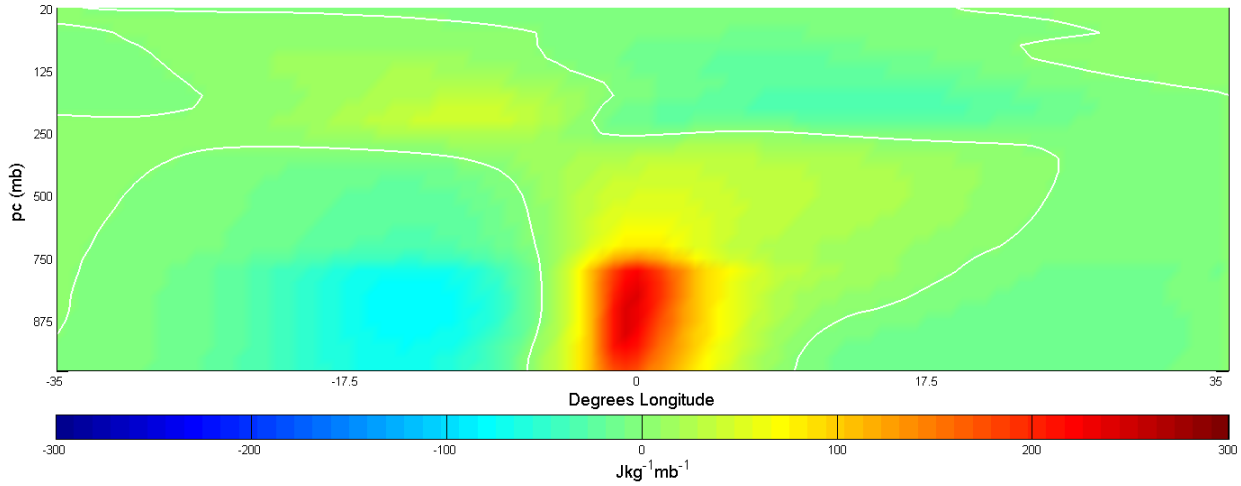

Figure 77: Composite pressure vs. longitude H' colour map for extreme events during DJFs in the Pacific storm track (PS) domain. The units are $Jkg^{-1}mb^{-1}$. The white contours correspond to zero H'.

pattern, on the opposite, shows a large contribution from the negative v', negative H' combination. The core of the extreme event is still dominated by the $++$ combination, but the confidence level is low and the - - one has a comparable frequency of occurrence. None of the other domains show any significant differences in sign combinations. The percentage of events sharing the same sign combination is generally higher than seen for the full hemispheric composites, although all domains are still below 70% at the core of the extreme events. It should be noted that the BS domain displays lower percentages at the extreme event location than the other domains, except for GS JJA.

4 Wavelet Spectra

4.1 Inferences from Basic Principles

As discussed in Chapter III and in *Messori and Czaja* [2013b, hereafter *MC13*], extreme heat transport occurrences display temporal and spatial scales compatible with phase-driven, baroclinic growing systems. In terms of the spatial scales, the cross-sectional composites described above support this conclusion. However, they do not provide a precise indication of the time scales involved, and therefore motivate a more systematic analysis of the latter.

While baroclinic motions might not be the single dominant component when considering the whole

of the transport [e.g. *Blackmon et al.*, 1977], they would be expected to be the main drivers of the extreme occurrences, both in view of what found in *Messori and Czaja*, and of the intuitive link between large heat transport and growing systems. A priori, there is therefore no reason to doubt that the traditional view of growing systems as developing baroclinic disturbances should not apply to these extreme events. Considering the transport as a product of anomalies in meridional velocity (v) and moist static energy (H), one can apply the bandwidth theorem to test the plausibility of the argument. The theorem applies to any wave phenomenon, and is commonly applied to wave packets. It relates the spread in frequencies ($\Delta\omega$) of said phenomenon to its duration (Δt) via the following inequality:

$$\Delta t \Delta\omega \geq \frac{\pi}{2} \tag{35}$$

This relation then suggests that baroclinic timescales in v and H would constrain the transport signal duration according to:

$$\Delta t \geq \frac{\pi}{2\Delta\omega}, \tag{36}$$

$$\Delta\omega = \frac{2\pi}{2.5} - \frac{2\pi}{6}, \tag{37}$$

where the denominators in equation 37 come from the definition of baroclinic time scales as having periods between 2.5 and 6 days (see Section 2). This corresponds to $\Delta t \geq 1$ day. MC13 have shown that heat transport extremes typically last for a few days. This lower bound therefore does not exclude baroclinic timescales from driving extremes, but it is clearly inconclusive and implies that further analysis is needed.

4.2 Methodology

The present analysis utilises the ECMWF ERA-Interim reanalysis data, previously described in Chapter II, Section 3.1 on page 54. Due to a more restrictive selection of events when compared to the previous sections of the present thesis, the number of seasons considered has been extended in

order to obtain a sufficient number of data points. The analysis now includes 22 DJF and 22 JJA time series, from June 1989 to February 2011. Part of the analysis presented in Chapter III has been replicated over this longer period in *Messori and Czaja* [2013a], and the qualitative conclusions were found to remain essentially unaltered. The transient-eddy heat transport is computed following the procedure described in Section 3, Chapter III. The extreme events are also defined as before (namely as the top 5 percentiles of the distribution of heat transport at 850 mb, for a given season/hemisphere combination). In order to avoid double counting, in the case of several consecutive days at a given location exceeding the threshold, only the local maximum of the $v'H'$ signal is retained. A Morlet wavelet transform is then applied to the selected $v'H'$ and the corresponding v' and H'. Note that the transform of the transport is computed as the transform of a single variable $v'H'$, not as the product of the transforms of v' and H'.

Wavelet transforms are a technique whereby a two dimensional spectral picture (here in period-time space) is obtained from a one-dimensional time series [*Daubechies*, 1990]. Wavelets are therefore the optimal tool to analyse the velocity, MSE and heat transport data in the time-frequency domain, allowing a detailed analysis of the signals' power throughout each extreme event. The wavelet transforms of two signals can also be combined into a single picture providing an indication of the frequencies and times at which the two signals share the most power. This is termed cross-wavelet transform. Furthermore, wavelets can provide a time-frequency picture of phase and coherence of separate signals. These two features provide a complete overview of the frequency and phase relationships driving the extreme events, making wavelet transforms perfectly suited to the present spectral analysis. For further details on the Morlet wavelet and on the wavelet formalism, the reader is referred to *Goupillaud et al.* [1984], *Torrence and Compo* [1998] and *Grinsted et al.* [2004]. An overview of the technical aspects of the transform is also provided in the Appendix.

Similarly to more traditional Fourier transforms, wavelets do not yield significant results if applied to time series which are short compared to the time scales considered. In order to minimise this problem while avoiding adding data external to the seasons being considered, only the extreme events on the central day of every season are retained (*i.e.* 15^{th} January and 16^{th} July). While this criterion might seem excessively restrictive, thanks to the extended time period now being considered it still provides around 4×10^4 events per season/hemisphere combination. As discussed below, this number is found to be sufficient to provide significant results. The wavelet power spectra for the

selected events are then composited, so that a single spectrum for each signal is obtained for every season/hemisphere combination. *Torrence and Compo* [1998] have suggested that it is not good practice to composite wavelet spectra because the power maxima will be smeared out. Here, we feel justified in adopting this approach because we are co-locating the extremes, where we expect the power to peak.

Part of the analysis is also performed on time-filtered data. The filter used is a 21-point high pass Finite Impulse Response (FIR) filter, with a half-power cut-off at 8 days, and is designed to capture the full breadth of baroclinic timescales. Even though our definition of these timescales considers a more restrictive interval, *Chang* [1993] suggested that filters with a 6 day cut-off loose a key part of the baroclinic variance. Here, we therefore follow *Nakamura et al.* [2002] in choosing an 8 day cut-off so as to retain the full variability of the 2.5–6 day band. The filter can be tested by computing the spectral power of the filtered and unfiltered v' and H' signals in the 2–6 day range. It is found that typically less than 5% of the power is lost when the filter is applied. Further details concerning the filter are provided in the Appendix.

The present chapter also analyses wave number power spectra. To obtain them, the v', H' and $v'H'$ signals around a given latitude circle, for all days of a given season, are decomposed into single wave number contributions using a Fourier Transform algorithm. This yields one coefficient per day per wave number. A Fourier amplitude spectrum is then computed in the time domain for each wave number, yielding spectral power versus period (or frequency). This procedure is repeated for all latitudes and years considered. Finally, the spectra thus obtained are composited, yielding a single spectrum for every season/hemisphere combination.

To verify that none of the results presented below depend on artefacts of the statistical analysis performed, a number of verifications were carried out. First, the wavelet analysis was repeated in logarithmic space, as described in detail in the Appendix. This robustly establishes that variations in the composite's power, in correspondence with the extreme events, are representative of the individual events. It also indicates that a sufficient number of extremes have been selected. The whole analysis was also repeated using non-detrended data. Next, the wavelet transforms were repeated using a Paul wavelet. To further test the robustness of the wavelet analysis, the period-averaged wavelet spectra were compared to the traditional Fourier spectra. Finally, the filtering was repeated using the band-pass filter from *Blackmon and Lau* [1980]. None of these changes affected the qualitative

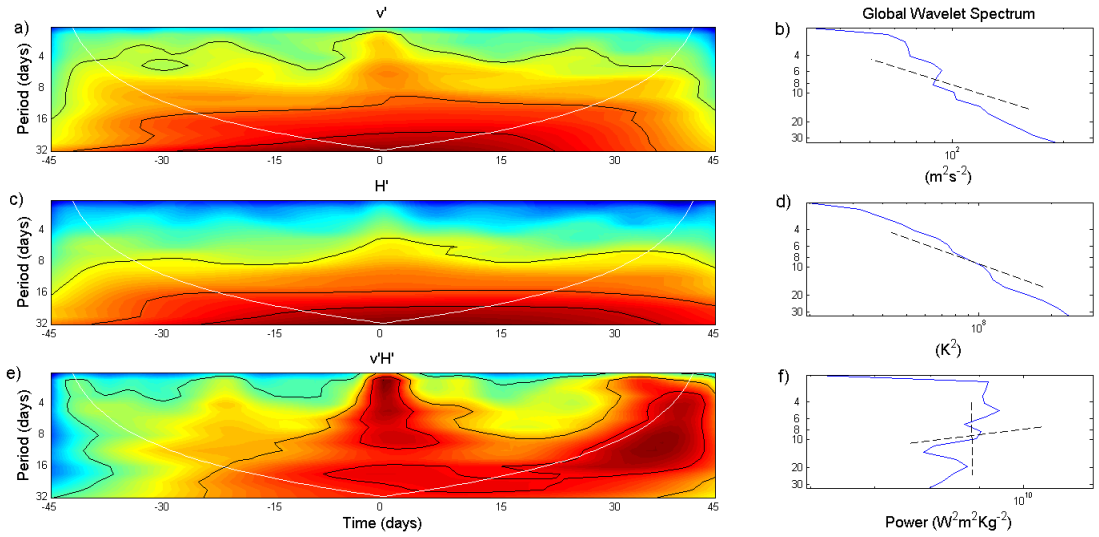

Figure 78: Composite wavelet power spectra of a) meridional velocity anomalies (in m^2s^{-2}), c) moist static energy anomalies (in K^2) and e) atmospheric heat transport due to transient eddies (in $W^2m^2kg^{-2}$). The spectra are centred on local maxima of transport extreme events. Darker colours indicate higher values. The white contours represent the cones of influence. Panels b), d), and f) display the corresponding time-averaged spectra, for the five days centred on the extreme, on a log-log scale. The dashed lines in panels b) and d) are positive unit slopes. Those in panel f) have slopes 0 and -1 respectively. The data cover NH DJFs from December 1989 to February 2011. All latitude circles between 30° N and 89° N are taken into account [from *Messori and Czaja*, 2013c].

aspects of the results presented in this study. Further technical details on the significance of the results presented here can be found in the Appendix.

4.3 General Features of the Power Spectra

To investigate the spectra of transient-eddy heat transport extremes, we begin by computing composite power spectra, taking into consideration all available NH and SH latitude bands (30°–89° N and 30°–89° S) over twenty-two DJF and JJA time series (1989/1990-2010/2011). Panels a, c and e in figures 78, 79, 80 and 81 show the results for v', H' and transport for NH DJF, NH JJA, SH DJF and SH JJA respectively. It should be noted that the latter spectrum was computed by treating the transport as a single variable, as described in Section 4.2. The white contours indicate the cones of influence, which represent the limits beyond which edge effects become important. Only the data above these lines should therefore be considered. Darker shades indicate higher spectral power. Also displayed are the time averaged spectra for the 5 days centred on the extreme event, and

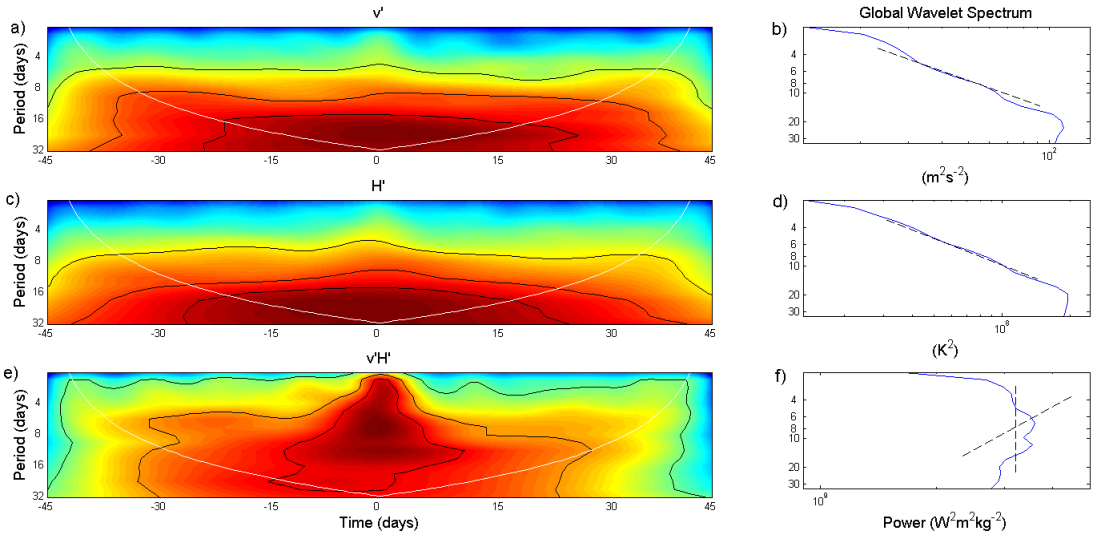

Figure 79: Same as figure 78, but for NH JJAs. The data now cover the period from June 1989 to August 2010.

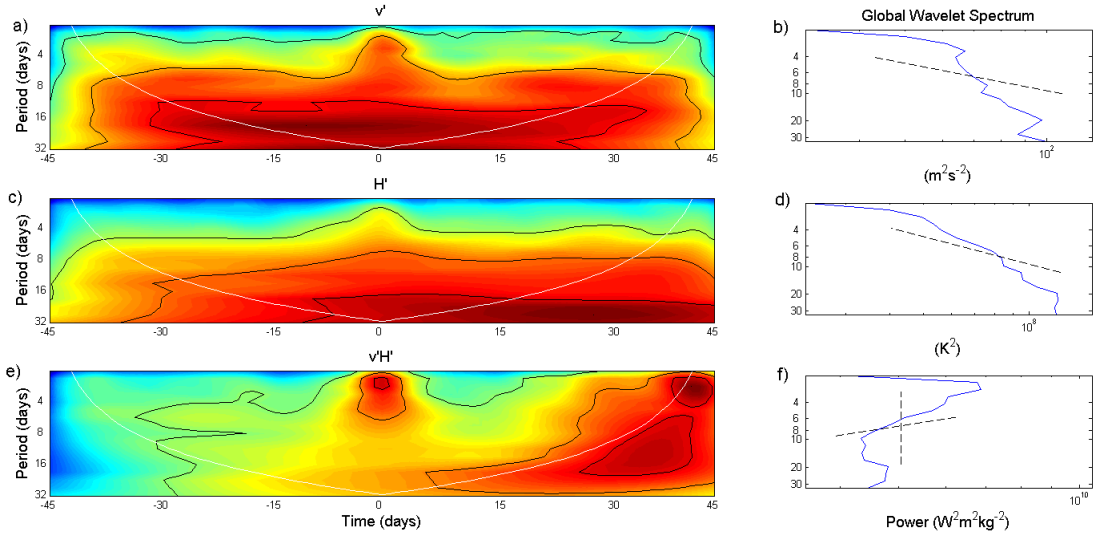

Figure 80: Same as figure 78, but for SH DJFs. The data now cover all latitude circles between 30° S and 89° S.

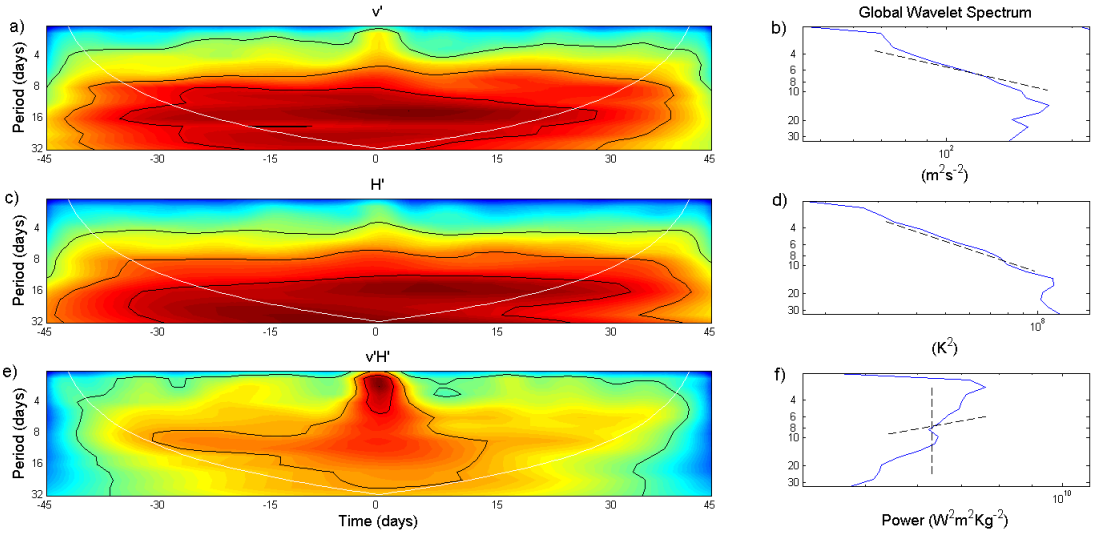

Figure 81: Same as figure 80, but for SH JJAs [from *Messori and Czaja*, 2013c].

reference slopes (panels b, d and f respectively). The 5 day window was chosen to ensure that the full breadth of the extremes was captured, following the durations found by MC13. The key features of the NH spectra are: i) a clear increase in power with increasing period in v' and H' and, ii) an approximately flat power spectrum in $v'H'$, with a modest peak at periods around 3-5 days. The latter feature is consistent with the typical extreme event duration of 2-7 days reported in MC13. The differences between the spectra are illustrated very clearly in panels b, d and f of figures 78 and 79. The dashed lines in panels b and d both have positive unit slope, while the ones in panels f have zero and negative unit slopes respectively, highlighting the relative flatness of the transport spectra. Similar observations apply to the corresponding SH time series, although here the $v'H'$ spectra tend to be more tilted than their NH counterparts. Two of the $v'H'$ spectra, namely NH and SH DJF, also display a marked asymmetry in the power before and after the event. While a large part of said asymmetry occurs within the cone of influence of the wavelet, the shorter periods also clearly display it. This could be linked to the spatial structure of the extreme events, and will be discussed further in Section 5.

The discrepancy between the $v'H'$ power spectra and those of v' and H', suggests that there must be a systematic interaction pattern between v' and H' driving the "flattening" of the transport spectrum. Furthermore, since transport is the result of a product of two signals, and the wavelet transform is computed on the product, a scaling factor must be taken into account. So, for example,

Event	Hemisphere	Season	Ratio
a) Extreme	N	DJF	0.31
		JJA	0.22
	S	DJF	0.35
		JJA	0.31
b) Median	N	DJF	0.24
		JJA	0.18
	S	DJF	0.36
		JJA	0.33

Table 7: Fractional contribution of baroclinic v' and H' fluctuations to the period-integrated $v'H'$ power spectrum over a 5 day window centred on events in a) the top 5 percentiles of the $v'H'$ distribution and b) within 5 percentiles of the median value of the poleward-only distribution. The data cover all the 44 seasons considered (June 1989-February 2011) [from *Messori and Czaja*, 2013c].

periods of 10 days in v' and H' could potentially contribute to the 5 day band in $v'H'$. This suggests that a broad range of frequencies, beyond the 2.5-6 day range, might be contributing to this effect.

To investigate these contributions in greater detail, the v' and H' data is high-pass filtered (periods < 8 days) using the filter described in Section 4.2. The filter cut-off period is well beyond that of our range of interest (2.5-6 days), to ensure that power loss due to filter design is minimised. A composite power spectrum is produced following the same procedure as for the unfiltered data (see Section 4.2). Next, the power in the wavelet spectrum in the 5 days centred on the extreme event is summed over all periods within the cone of influence. Last, a ratio between the summed power of the filtered and unfiltered data is computed. In taking the ratio, the same range of periods is included in both the filtered and unfiltered integrals, even where these fall well outside the baroclinic range. This ensures that the full breadth of contributions by the baroclinic-scale v' and H' motions is captured. Depending on the hemisphere/season combination, it is found that the power spectrum resulting from the filtered data captures between 22% and 35% of the power found in the original spectrum (see table 7). If one replicates the analysis selecting events within 5 percentiles of the median of poleward heat transport, rather than for extreme events, the ratios are very similar, with only a modest difference emerging during the NH winter (see table 7). The contribution from baroclinic periods in v' and H' is therefore comparable for extreme and median events. In both cases, these timescales do not drive the bulk of the transport, which must therefore come from longer periods.

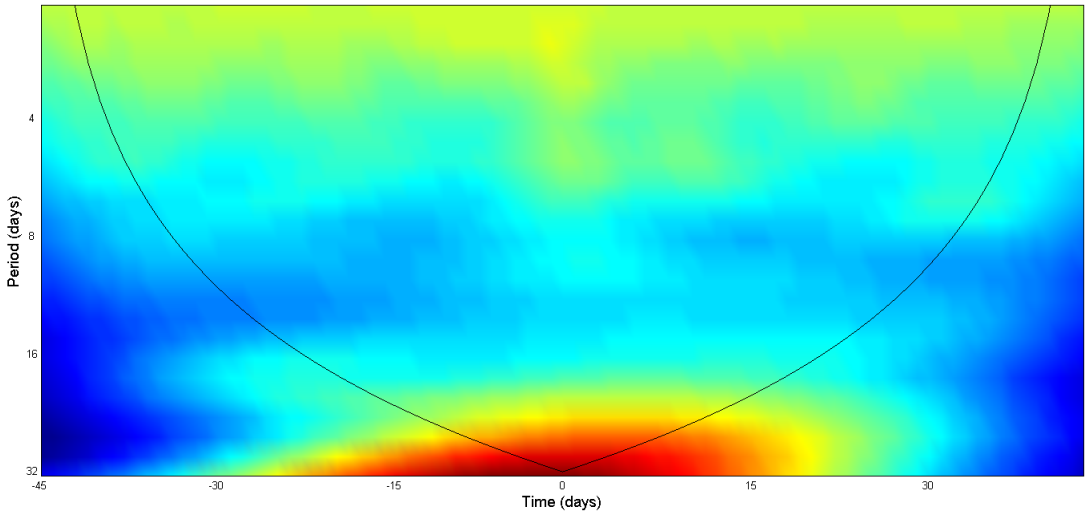

Figure 82: Wavelet cross-spectral power for v' and H' during NH DJFs. Darker colours (red) indicate higher values. The black contour represents the cone of influence. The data cover the same range as in figure 78.

4.4 The Role of Phase and Coherence

The velocity and MSE fluctuations do not display a clear peak at baroclinic scales, as seen in the spectra in figures 78-81, and cross-wavelet spectra of v' and H' confirm that most of the power the two signals share resides at long periods, well above baroclinic timescales. As illustration, figure 82 shows the v' versus H' cross-wavelet power for NH DJF. The other seasons and hemisphere display comparable characteristics (not shown).

The shift towards higher frequencies seen in the $v'H'$ spectra can only be explained by systematic phase and coherence relationships between v' and H' perturbations. Figures 83 and 84 show the composite coherence-phase plots for NH DJF and SH JJA transport extremes, respectively. Panel a) in both figures depicts coherence as colours and phase as arrows. The black contour indicates the cone of influence, which represents the limit beyond which edge effects become important. Again, only the data above this line should be considered. Since phase for incoherent signals has no meaning, phase arrows are only retained in regions where coherence is greater than 0.5. When the arrows point right, the v' and H' signals are perfectly in phase; when they point upwards, the two signals are in quadrature with v' leading H'. The converse holds for leftwards and downwards pointing arrows. Panel b) depicts phase angle in colours for the 10 days centred on the extreme events. Again, positive values indicate that velocity leads MSE; the black contours bound the regions where coherence exceeds

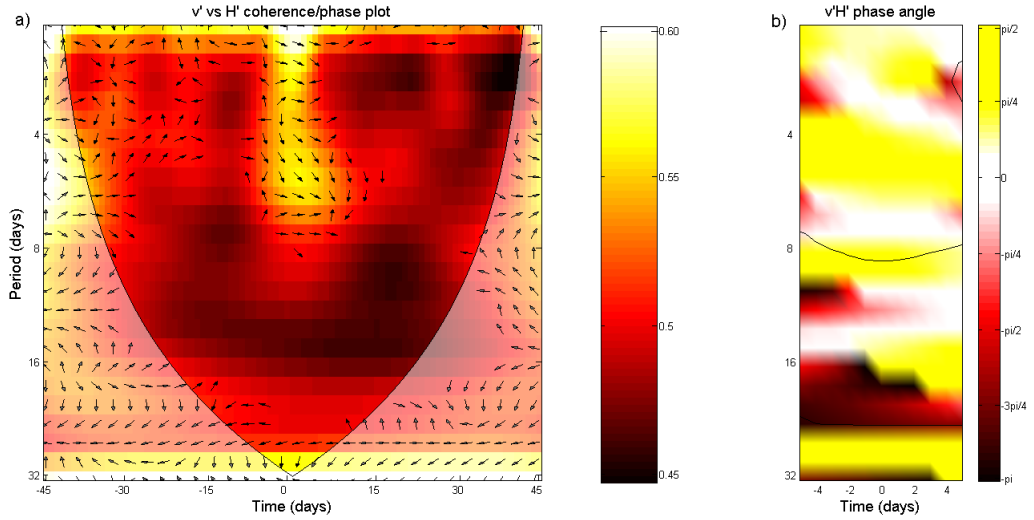

Figure 83: a) Composite wavelet coherence/phase spectrum for meridional velocity and moist static energy anomalies. The spectra are centred on local maxima of extreme transport events. The colour map represents coherence values, the arrows the phase relationship. When the arrows point right, the v' and H' signals are perfectly in phase; when they point upwards, the two signals are in quadrature with v' leading H'. The converse holds for leftwards and downwards pointing arrows. The black contour represents the cone of influence. b) Depicts phase angle in colours for the 10 days centred on the extreme events. The scale is in radians. Again, positive values indicate that velocity leads MSE; the black contours bound the regions where coherence exceeds 0.5. The data range covers NH DJFs and is the same as in figure 78 [from *Messori and Czaja*, 2013c].

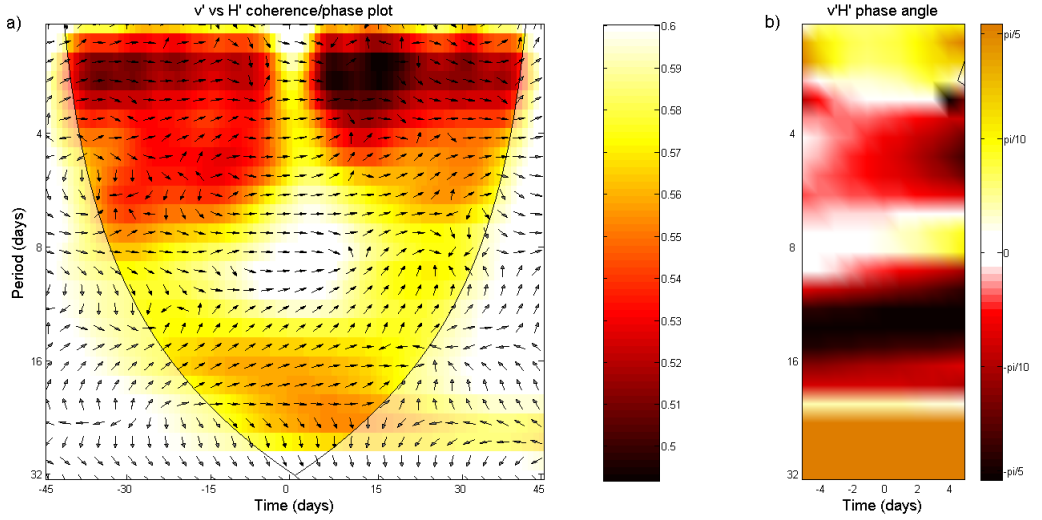

Figure 84: Same as figure 83 but for SH JJAs [from *Messori and Czaja*, 2013c].

0.5.

For the NH, three features emerge, which explain the patterns seen in figure 78: i) the coherence is highest at periods of approximately 2-8 days, ii) the phase angle is smallest at approximately 3 and 7 days, and iii) at longer periods v' and H' approach either quadrature or anti-phase. Similar considerations, albeit with slightly different numbers, can be made for SH DJF (not shown), and SH JJA (figure 84). In the case of NH JJA, the coherence values are, perhaps, too small to allow us to make definitive statements using phase. For SH JJA, at all but the lowest frequencies the phase angle stays relatively small, while the coherence values are systematically higher than those seen in the NH. Nonetheless, there is still a clear peak in coherence, now shifted more towards the 6-12 day range, and the phase angle is smallest at 2-3 and 7-8 days. The flattening in the $v'H'$ spectrum is therefore driven by coherent, in phase v' and H' disturbances, while the unfavourable phase relationship at higher periods masks the high power of the two signals there. Note that, since heat transport results from the product of v' and H', a scaling factor in frequency (and therefore in period) must be taken into account here. If not for this "braking" role of phase, there would therefore be significantly more power at long periods in the $v'H'$ spectrum. This effect is evident in both hemispheres, while the differences in coherence between the SH and NH lead to more subtle considerations, and will be discussed below.

It is interesting to note that, in both figures, the data seems to display phase angles smaller than the Eady angle (21°). The latter is the theoretical lower bound on the phase angle between the velocity and temperature components of baroclinic disturbances near the surface. The Eady model, however, considers a highly idealised setup of monochromatic, coherent waves, so it is not obvious that the numerical limits it sets should be exactly obeyed by real-world instabilities. To further illustrate the importance of phase and coherence, composite wavelet spectra for a set of random events were computed. These are data selected at random locations, on the same days and over the same range of latitudes and longitudes as the extreme events. A number of data points similar to the number of extreme events were chosen and the spectra were then composited, exactly as for the extremes. These composites are shown in figure 85 for NH DJF. The v' and H' spectra are comparable to those seen in figures 78a and 78c , while the $v'H'$ spectrum is extremely different from figure 78e and is, in fact, very similar to the velocity and MSE ones. The distinctive feature of extreme events therefore does not lie in the power spectra of the v' and H' fluctuations but, rather, in their phase

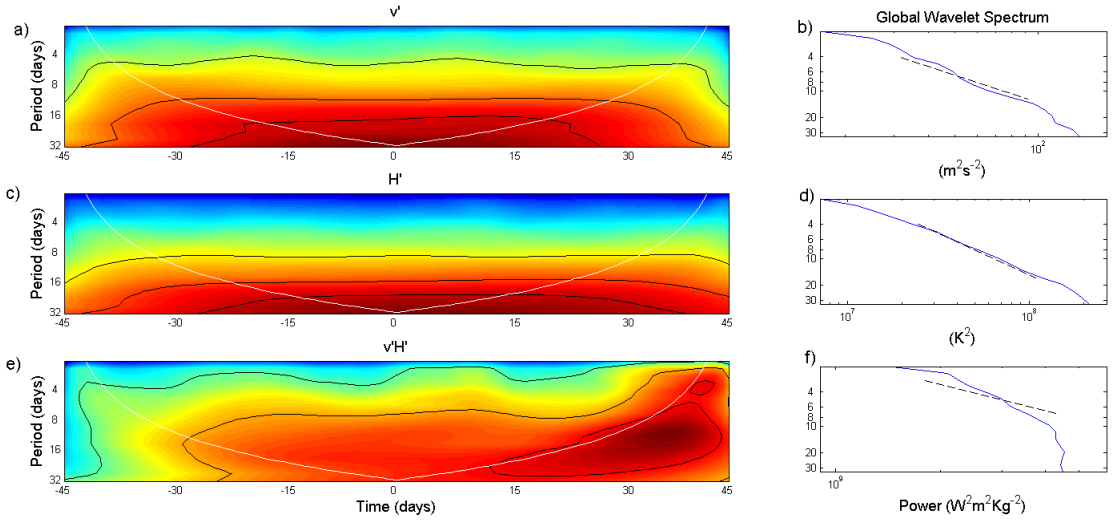

Figure 85: Same as figure 78 but for heat transport events selected at random locations, on the same days and over the same range of latitudes and longitudes as the extreme events. The data cover NH DJFs. Unlike in figure 78, the dashed line in panel f) also has positive unit slope [from *Messori and Czaja*, 2013c].

and coherence relationship. As might be intuitively expected, both aspects are essential. Indeed, an artificial $v'H'$ spectrum was reconstructed using the actual coherence of the v' and H' signals, but applying a random phase at all periods, and none of the patterns discussed above were reproduced. Therefore neither coherence nor, obviously, phase alone can explain the $v'H'$ spectra.

The spectra of the different seasons and hemispheres present a coherent picture and share the same qualitative results. As noted above, however, the coherence values in NH JJA are, perhaps, too small to allow us to make definitive statements using phase, while coherence values in the SH are generally larger. This is evident by comparing figures 83 and 84, and a similar pattern emerges when comparing NH JJA to SH DJF. A possible explanation could be related to the different storm track configurations in the two hemispheres, with the SH storm track being more extensive than its NH counterpart. Since a high coherence can be expected in these areas of strong growth, this difference could drive the higher coherence values found in the SH. To the author's knowledge, however, no references or conclusive explanations for this difference exist in the literature.

The differences in coherence between the two hemispheres drive some interesting mechanisms. Systematically higher coherence values mean that phase angles at all periods have a strong impact

on the resulting $v'H'$ spectrum. In the SH, coherence therefore has a comparatively less important role than in the NH. As thought experiment or *gedankenexperiment*, consider the situation in which velocity and MSE have, throughout their full range of frequencies, the same phase angle. Assume further that this angle happens to be the optimal one for transporting heat; in the spirit of this idealised setup, this is taken to be the angle found in the Eady model of baroclinic instability (21°). If this is imposed on the actual coherence found, it will result in a flattening of the $v'H'$ spectrum in the NH, but not in the SH. In the NH, where coherence is generally low, the phase angle only has an effect at periods where coherence is sufficiently higher than average. These happen to be within the baroclinic range, as discussed above. In the SH, where coherence is high at all periods, phase becomes instrumental in directing the power in the $v'H'$ spectra, meaning that imposing an Eady phase at all periods leads to the loss of the flattening effect seen in the real data.

5 Physical Interpretation

The cross-sections discussed in Section 3 and the spectral features analysed in Section 4 present a comprehensive overview of the space-time structure of extreme events in heat transport. At the same time, the two analyses do not necessarily blend into a single, coherent picture. The cross-sectional composites emphasise synoptic-scale features of the circulation, but do not afford any insight into the interaction of local dynamics with larger scale variability. The wavelet plots, on the other hand, provide a wealth of information regarding the timescales driving the atmospheric motions, but no direct link to the spatial scales these correspond to. In order to offer a physical interpretation of both the spatial and temporal analyses, while at the same time reconciling some apparently contrasting aspects of the two, the following points will be discussed:

i) Wave number spectra, which link the temporal wavelet information to precise spatial scales;

ii) The vertical and zonal structure of the heat transport in relation to the idealised Eady model, and possible signatures of this structure in the wavelet spectra;

iii) The local circulation and synoptic analogues;

iv) Alternative interpretations to the wave picture adopted here.

Points i) and ii) suggest that, to some degree, a joint interpretation of the spatial and temporal scales is indeed possible.

Wave number Spectra

According to both physical [*Messori and Czaja*, 2013b] and purely mathematical considerations (*e.g.* the bandwidth theorem, discussed in Section 4.1), the observed extremes in $v'H'$ could plausibly be solely due to baroclinic disturbances. Here, the latter are defined as motions with periods between 2.5 and 6 days. Traditional measures find that these timescales and motions with periods longer than 10 days typically account for comparable portions of the 850 mb NH transient-eddy heat transport [*Blackmon et al.*, 1977]. Intuitively, extremes must correspond to growing systems, which are associated with baroclinic perturbations. Therefore, when focusing specifically on extremes, one would expect the contribution of baroclinic scales to be significantly higher than normal. Here, it is found that this is not the case: the v' and H' power spectra are found to peak at long periods and, as discussed above, these long periods provide the greatest contribution to the power in the $v'H'$ spectrum. Even though the coherence and phase characteristics of the v' and H' signals favour short period fluctuations, this is not enough for the latter to account for the majority of the power in the $v'H'$ spectrum (see table 7 on page 127). It should also be noted that, as mentioned in Section 4.3, periods beyond the baroclinic range in v' and H' can contribute to baroclinic motions in the $v'H'$ spectrum. Regardless of hemisphere and season, v' and H' fluctuations therefore combine to produce phase and coherence-driven extreme events in heat transport which, as discussed in detail in Chapter III, account for over half of the poleward heat transport by transient eddies.

The secondary contribution of baroclinic timescales to the $v'H'$ power spectrum, highlights the key role played by the longer periods. These can be related to physical analogues in the atmosphere by investigating the spatial scales driving the transport. In order to do this, the power in the v', H' and $v'H'$ signals is decomposed into single wave number contributions (hereafter also referred to as k), following the procedure described in Section 4.2. A caveat of this technique is that, since it takes into account full latitude circles, it is not limited to the locations of extreme events. Figures 86, 87, 88 and 89 show the composite power spectra for wave number as a function of period for NH DJF, NH JJA, SH DJF and SH JJA, respectively. Panels a, c and e correspond to v', H' and $v'H'$ spectra. Darker shades indicate higher powers. Also displayed are the period-averaged spectra for the 10-32

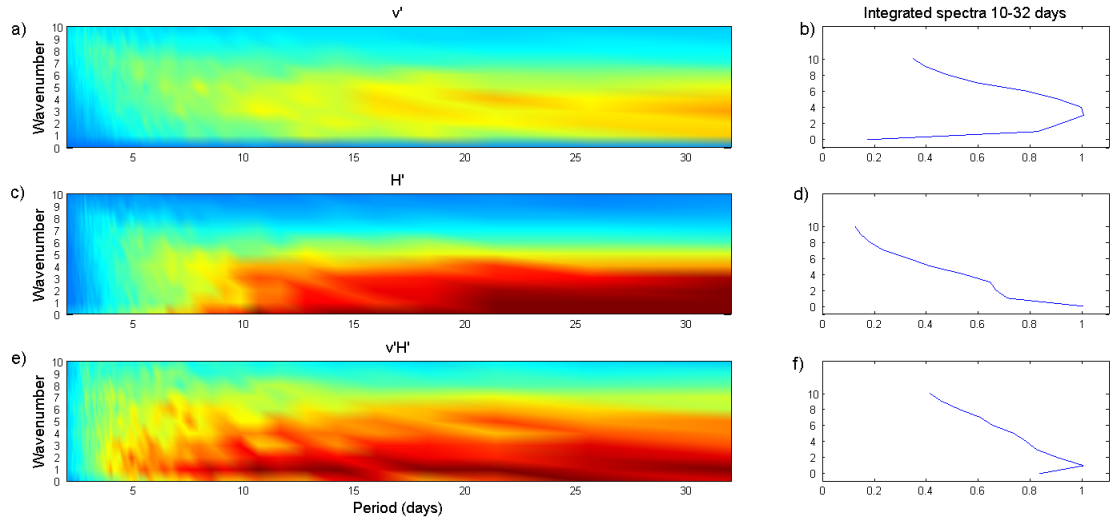

Figure 86: Composite Fourier power spectra of a) meridional velocity anomalies, c) moist static energy anomalies and e) atmospheric heat transport due to transient eddies. The spectra are a function of wave number versus period. Darker colours indicate higher values. Panels b), d), and f) display the corresponding spectra averaged over the range 10-32 days. The data range covers NH DJFs and is the same as in figure 78 [from *Messori and Czaja*, 2013c].

day range (panels b, d and f respectively). These are normalised relative to the spectral peak, so that the contribution of each wave number is expressed as a dimensionless fraction of the contribution from the dominating wave number. The lower bound of the integration range was set to 10 days to ensure that it was well beyond the range of the filter used in Section 4.3, which has its half-power cut-off at 8 days. The upper bound was chosen to coincide with the longest periods resolved by the wavelet spectra discussed above. From the figures, it is immediately clear that planetary-scale waves ($k \leq 4$) dominate the spectra. During NH DJF, the peak contributions in the v', H' and $v'H'$ spectra come from $k = 3 - 4$, $k = 0$ and $k = 1$ respectively. Similar considerations apply to the NH JJA season. The corresponding plot for SH JJA, again shows similar results, with the peaks for the v', H' and $v'H'$ spectra now being respectively at $k = 4$, $k = 0, 3$ and $k = 0$. The SH DJF season differs slightly from the above, in that the transport spectrum peaks at $k = 1$ and $k = 4$. The latter spectrum reproduces quite closely the corresponding v' spectrum, and could be related to the typical $k = 4 - 7$ patterns seen in the SH summertime circulation [e.g. *Hamilton*, 1983].

As already mentioned, the results discussed above refer to an analysis of full latitude circles, and are not limited to the locations of extreme events. Nevertheless, they can provide important insights

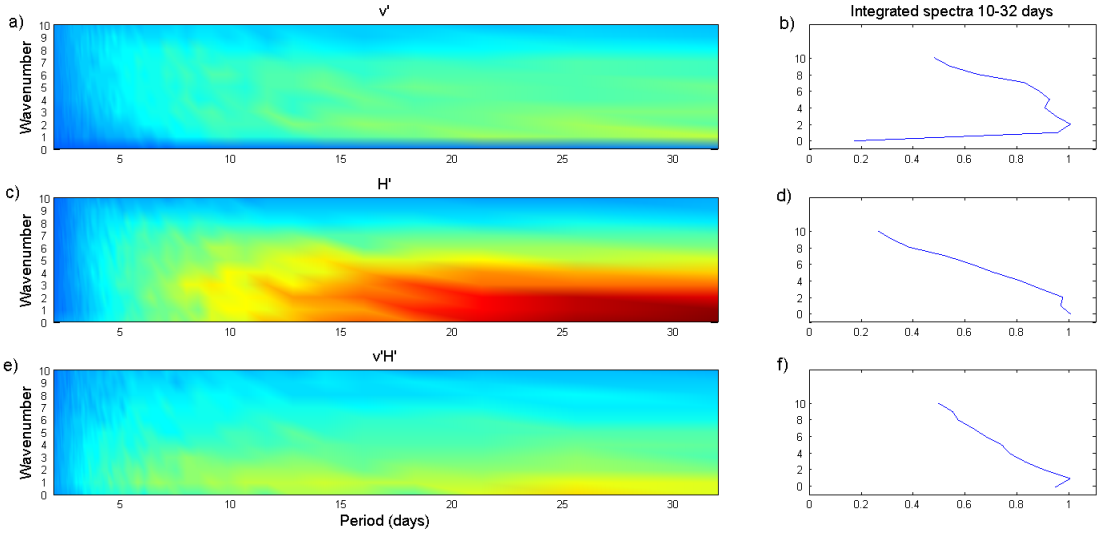

Figure 87: Same as figure 86, but for NH JJAs. The data cover the same range as in figure 79.

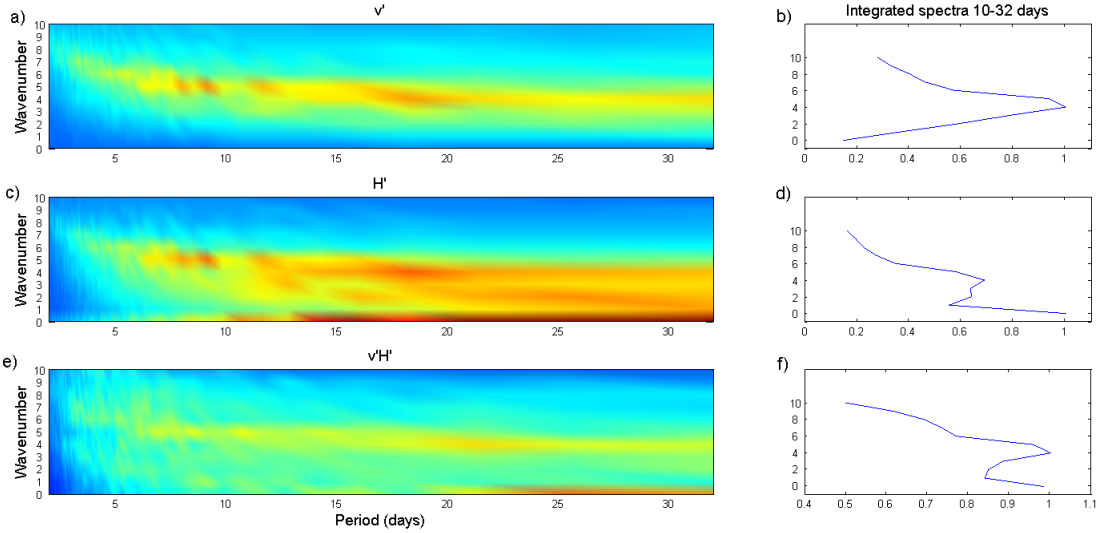

Figure 88: Same as figure 86, but for SH DJFs. The data cover the same range as in figure 80.

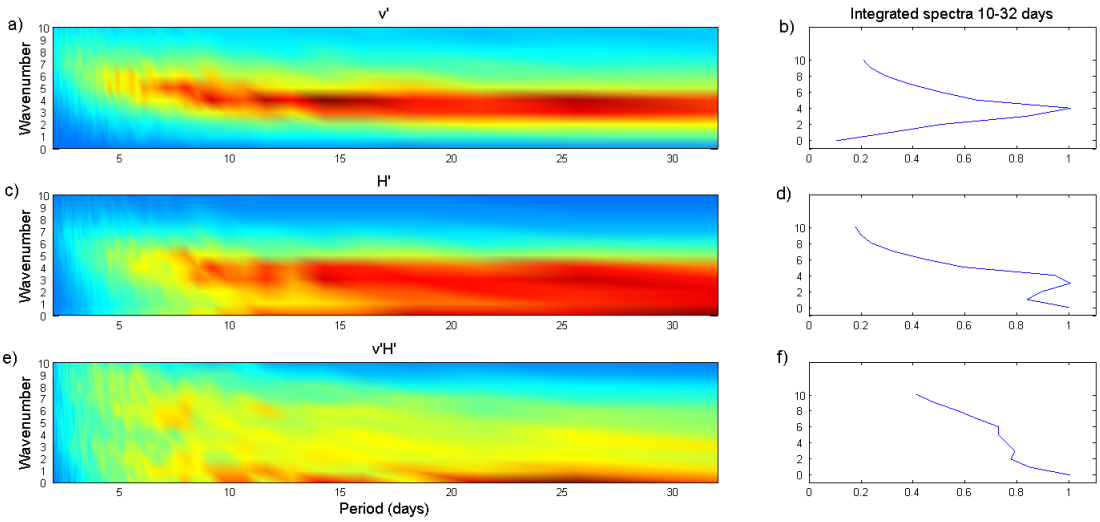

Figure 89: Same as figure 88, but for SH JJAs [from *Messori and Czaja*, 2013c].

into the origin of the low frequencies in the extremes' power spectra. Indeed, there is no reason to suspect that the motions accounting for low frequencies in the extreme events are distinct from those driving the same frequencies in the full latitudinal data. The power lost in applying the baroclinic filter to v' and H', namely that due to motions with periods beyond 8 days, therefore comes primarily from very low wave numbers, in the range $0 \leq k \leq 4$. Similarly low wave numbers account for the majority of the transport's spectral power at long periods. Allowing for the significant seasonal and inter-hemispheric differences found, we interpret these results as indicating that the power seen in the $v'H'$ wavelet transforms at long periods is driven by planetary-scale waves.

At the same time, the long periods do not account for the entirety of the power of the extreme events. It is difficult to relate directly percentiles of the heat transport distribution, used to define the extremes, to percentiles of the time-integrated power spectrum[6]. However, keeping in mind that periods below 6 days typically account for just under one-third of the total power of an extreme, it is reasonable to assume that these timescales will play an instrumental role in determining whether an event is extreme or not. That is, if a given day is in the top percentiles in terms of power at periods beyond 6 days, it might not necessarily be in the top percentiles in terms of power over all periods. Imagine, as example, two events with identical, very high power, at long periods. Since we single out

[6]Remember that the baroclinic contribution to the power of extreme events was computed by integrating the power spectrum over a five-day window centred on the extremes themselves.

events in the top 5 percentiles of the heat transport distribution, a factor of two variation in power at periods below 6 days might well imply that only one of the two classifies as an extreme.

The above suggests two considerations:

i) Extreme events in meridional heat transport by transient eddies are not merely the signatures of passing synoptic systems, but are driven by larger scale modes of atmospheric variability.

ii) Even though the long periods account for the majority of the extreme events' spectral power, they are not necessarily sufficient, alone, to determine whether an event classifies as extreme or not.

Point i) is a novel result in the context of extreme events, and is consistent with the known influence of large scale modes, such as the North Atlantic Oscillation, on meridional eddy heat transport [*e.g.* *Carleton*, 1988; *Hurrell et al.*, 2003]. The second consideration suggests that extremes might be associated with precise atmospheric features at short periods. One might therefore expect to see clear synoptic patterns, when studying the local atmospheric circulation corresponding to extreme events.

The Vertical and Zonal Structure

The extreme event cross-sectional composites, presented in this chapter, do indeed display some common features across both hemispheric domains. The most striking one is the systematic westward tilt of the narrow poleward heat transport core and its associated return flow. This is very similar to the tilt expected of an idealised Eady wave (see Chapter II, Section 2.1 on page 46). Figure 90 shows the meridional heat transport associated with such a wave, here computed as $v'T'$. Since the magnitude of the transport is entirely arbitrary, and depends on the ambient parameters chosen for the simulation, the transport has been normalised such that the poleward maximum has unit value. The continuous contours correspond to values of 0.25, 0.5 and 0.75 respectively. The dashed contours correspond to values of 0 and -0.25 respectively. The periodicity of the figure is simply due to the periodic nature of the Eady wave itself, and should not be taken into account when comparing the plot to the ERA-Interim data.

The vertical tilt of the heat transport is very similar to the one seen in figure 58 on page 103. In the Eady model, such tilt is given by the small phase difference between velocity and temperature anomalies. Indeed, for perfectly in-phase perturbations, there is no zonal asymmetry in the system.

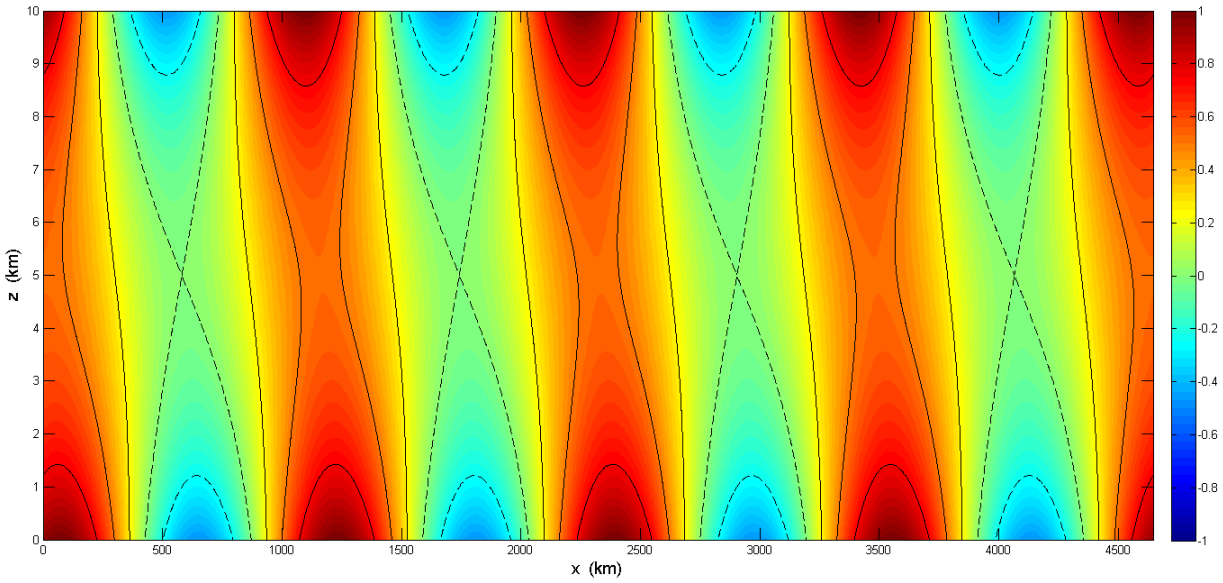

Figure 90: Composite pressure vs. time transport colour map for an idealised Eady wave. The transport values have been normalised such that the poleward maximum has unit value. The continuous contours correspond to values of 0.25, 0.5 and 0.75 respectively. The dashed contours correspond to values of 0 and -0.25 respectively.

In view of the discussion on phase presented in Chapter III, a similar tilt is also expected in the real atmosphere. A second similarity between the Eady model and reanalysis is the breadth of the poleward heat transport core, although in the model the horizontal scale is somewhat arbitrarily set by the tropopause height. Furthermore, in both data and model, the return flow on the eastern flank of the poleward transport is separated into two distinct areas. Due to the periodic pattern, however, in the Eady wave the same return flow is also witnessed on the western flank, unlike what is seen in the data. Another difference is given by the boundary intensification of the transport seen in figure 90 which, in the case of the upper bound, is a direct consequence of the rigid lid approximation. Nonetheless, considering the extreme idealisation of the model, the similarities with figure 58 are striking.

Concerning the vertical tilt, it is interesting to note the effect it might have on the transport's wavelet spectra. In the context of large scale eastward advection during the NH winter months (DJF) and in the SH, one would expect to see this zonal asymmetry mirrored in the wavelet spectrum, as is indeed the case in figures 78 and 80 on page 125. The asymmetry in the power spectrum is almost absent during SH JJA (see figure 81 on page 126), possibly because of the poleward shift in the location of the extreme events, which places them beyond the latitudes of peak zonal winds [*e.g. Risien and Chelton*, 2008].

The vertical asymmetry of the heat transport is mirrored in the v' and H' sign combinations driving it (see figure 65 on page 109). The area of largest poleward transport is seen to correspond to the combination of positive anomalies in both variables. This points towards the traditional WCB-type view of poleward heat transport corresponding to the advection of warm, moist air towards the high latitudes. As for what concerns the equatorward transport, the top and bottom parts of the tilted column correspond to different sign combinations. The positive v', negative H' pattern characterising the higher levels could be linked to a tropopause bending, driven by high-level eastward jets. These were found to match the locations of extreme occurrences quite closely (see figure 51 on page 87), and would imply a strong negative MSE anomaly. The opposite sign combination, seen at lower levels, is likely to be a temperature-driven effect, related to re-circulating air parcels which have come from the low latitudes and have not lost all of their heat at the time of the equatorward return. This is consistent with the southward flow east of the extremes, seen in the meridional velocity composites (see figures 60- 63 on page 107) . Ultimately, it should be noted that these sign combinations are not

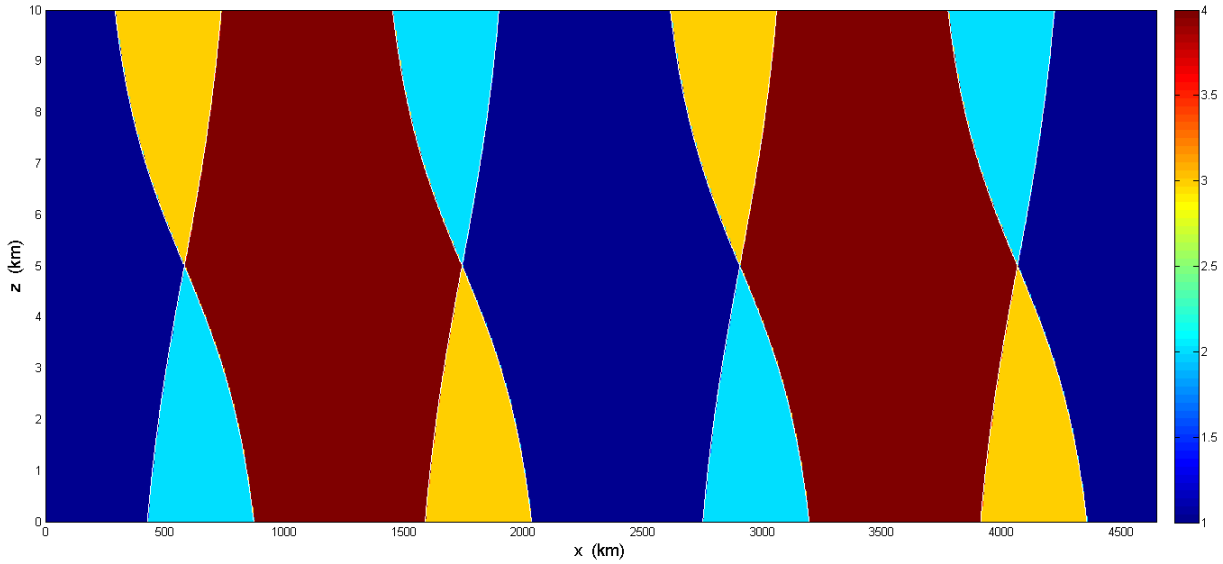

Figure 91: Composite v' and H' sign combination colour map for the meridional heat transport associated with an Eady wave. The white contours correspond to zero meridional heat transport. The sign convention is analogous to the one adopted in figure 65.

very robust, since at most grid points they represent between 30% and 50% of the extremes forming the composite. They are therefore a good basis on which to formulate hypotheses, but should not be taken as the motivation for a detailed investigation.

As was done for the transport, these sign combinations can be compared to the ones in the Eady model. Figure 91 shows the sign combinations driving the Eady transport, using the same convention as figure 65 on page 109. Unit values correspond to a positive v', positive H' pattern, values of 2 to the +- case, 3 to the -+ case and 4 to the - - one. Again, the pattern is periodic. The first striking aspect is that the transport maxima alternate between ++ and - - sign combinations. This could partially explain why the confidence levels in figure 65 are so low. Since all extreme events have a positive transport at their core, if just over 50% of events correspond to a ++ combination, it means that just under 50% of events must correspond to a - - sign pattern. It should also be noted that, if one takes a ++ maximum as reference, the sign combination of the equatorward transport is the opposite of what seen in the ERA-Interim data. What is, however, consistent with the data, is that the two areas of equatorward transport always display different sign combinations. As was the case for figure 90, there are therefore both important differences and striking similarities between the Eady

model and the reanalysis.

Synoptic Analogues

The large domains taken into consideration for the hemispheric composites make it difficult to associate precise synoptic analogues to the extremes. By looking at the regional composites, however, a number of common elements and robust features emerge. Here, we present possible synoptic analogues for the composites discussed in Section 3.4, although the suggested matches are by no means conclusive. It should be noted that only a brief analysis was performed regarding the differences between ++ and - - sign combinations in extreme events. As seen in figures 65 and 66, the ++ events tend to be stronger and more frequent than the - - cases, and therefore dominate the composites. The synoptic analogues discussed below are therefore mostly relevant to ++ combinations. The - - cases are only considered for the domains in which they play a major role, namely: GS JJA and BS DJF.

The three storm tracks domains (GS, PS and SO) present similar pictures, suggestive of a direct correspondence with WCBs. Referring back to the schematic shown in figure 11 on page 33, the rapid ascent at the location of the extremes matches the warm conveyor itself (see figure 74 on page 117). The descending motions seen in the velocity composites and the negative H' seen in the MSE ones would correspond to the dry air from the upper atmosphere on the western flank of the WCB. Similarly, the high level negative MSE intrusion coincides with the cold conveyor (see figure 77 on page 120). What the WCB schematic does not necessarily explain is the fact that the return heat transport is generally seen only on the eastern flank of the extremes (with the exception of the PS JJA case). In fact, idealised simulations of midlatitude cyclones predict two re-circulating branches of the WCB, one to the east and one to the west of the location of rapid ascent [*e.g. Boutle et al.*, 2010].

The good match between WCBs and extreme heat transport events over the storm track domains does not come as a surprise since, as observed in Chapter III, Section 5.2 on page 86 these are, broadly speaking, the areas where the extreme event distribution matches the WCB climatology. The latter is shown in figure 52 on page 88. The values shown in the figure, which refer to the percentage frequency of WCB occurrence, are systematically lower than 5% except for NH DJF. The extreme transport events, however, are identified as being the top 5 percentiles of the heat transport

distribution. Moreover, extreme events are computed relative to the full hemispheric distribution, meaning that they can correspond to more than 5% of events in very active regions such as the WBCs. However, the WCB climatology shown in figure 52, considers very restrictive criteria for identifying the conveyors. Other studies find similar geographical distributions, but slightly higher percentage occurrences. Moreover, the WCB algorithms are usually designed to avoid counting multiple trajectories associated with the same rising air mass, while no such filter is applied when selecting extreme events [*Madonna et al.*, Submitted].

The picture described above does not fit the GS JJA data (see figure 75 on page 118). In the latter composites, there is a strong signature of northerly cold, dry air flows, which is mirrored in the domain's sign combination map. This could be due to the fact that, especially during JJA, a large number of the extremes within the domain occur over the eastern portion of the North American continent. This interpretation agrees very well with the cold air outbreaks often witnessed in the region [e.g. *Konrad and Colucci*, 1989].

The NS domain, analysed here only during the winter season, presents a picture similar to the storm track domains discussed above. Even though the WCB frequency in the area is significantly lower than in the storm tracks, it is nonetheless a region of high cyclonic activity (see figure 53 on page 89). Moreover, the area of strong negative H' just east of the extreme event location, which is associated with an equatorward flow, matches very well the typical precursor conditions for polar lows, as discussed in Chapter III, Section 5.2 on page 86. The general distribution of extremes within the domain also matches polar lows very well (see figure 54 on page 90). It should be noted that the polar lows have warm cores and typically travel in a southward direction. In most cases, this will correspond to a positive MSE anomaly and a negative velocity one. The integrated heat transport effect would therefore be negative. However, the composites discussed here look at zonal snapshots at fixed times, and therefore tend to suppress the effect of the polar low drift, since the typical drift velocity would be dominated, in calculating v', by the vigorous circulation associated with the low. Suitable phase differences between the velocity and MSE fields, such as those highlighted in Section 4.4, can then drive a strong poleward heat transport. This is subject to the caveat of how well the reanalysis data captures polar lows, which is a problem which will not be discussed here. For further information on this topic, the reader is referred to *Condron et al.* [2006].

The SP domain presents an overall picture similar to the NS one, and its composites are consistent

with a cyclonic circulation. Even though it is not an area of polar low development, the domain corresponds to an area of significant cyclonic activity in the ERA-Interim data (see figure 53 on page 89). These local low pressure systems are therefore the most fitting synoptic analogues for this domain. It should be noted that the definition of cyclonic activity (see section 5.2) is based on much less stringent criteria than those adopted to define WCBs. This is reflected in the different schematic chosen in figure 92 to represent the SP domain.

The BS domain has a very weak, positive MSE signal. This excludes cold air outbreaks from the Arctic basin from being the only drivers of the poleward heat transport, as hypothesised in Section 5.2, Chapter III. As for the SP domain during JJA, the Bering Strait and Gulf of Alaska region during DJF is associated with significant cyclonic activity in the ERA-Interim data. Indeed, the Gulf of Alaska is a region of polar low genesis [*Businger*, 1987]. However, compared to the former domain, there is the additional complication of the presence of sea-ice. Studies have found evidence for ice-cyclone feedback mechanisms [*Carleton*, 1984], and the presence of sea-ice also heavily affects ocean-atmosphere fluxes. Upon closer analysis, the sign of the MSE at the extreme event location is seen to vary frequently between individual events, which is consistent with the low confidence of the sign combination map. There are, indeed, a number of events displaying southward airflows with negative MSE anomalies at the extreme event location, consistent with the prevailing northeasterly winds blowing across the frozen ocean [*Overland and Pease*, 1982]. Part of the events therefore correspond to cold air outflows from the Arctic basin, while the rest match the more traditional interpretation of cyclonic activity driving poleward heat transport.

The above synoptic analogues are summarised in figure 92.

The combined space-time picture

The cross-sectional plots of extreme events in heat transport suggest that the events across both hemispheres are localised in space and time, and share a number of common features, such as ascending airflows in correspondence with the core of the extreme and a clear westward tilt in the transport column. As could be expected by characteristics applying to such a broad domain, these can be matched to fundamental characteristics of baroclinc disturbances, which are reproduced in idealised simulations such as the Eady model. If one focusses on regional domains, it is also possible to suggest potential synoptic analogues for the extreme events. At the same time, the wave number and wavelet

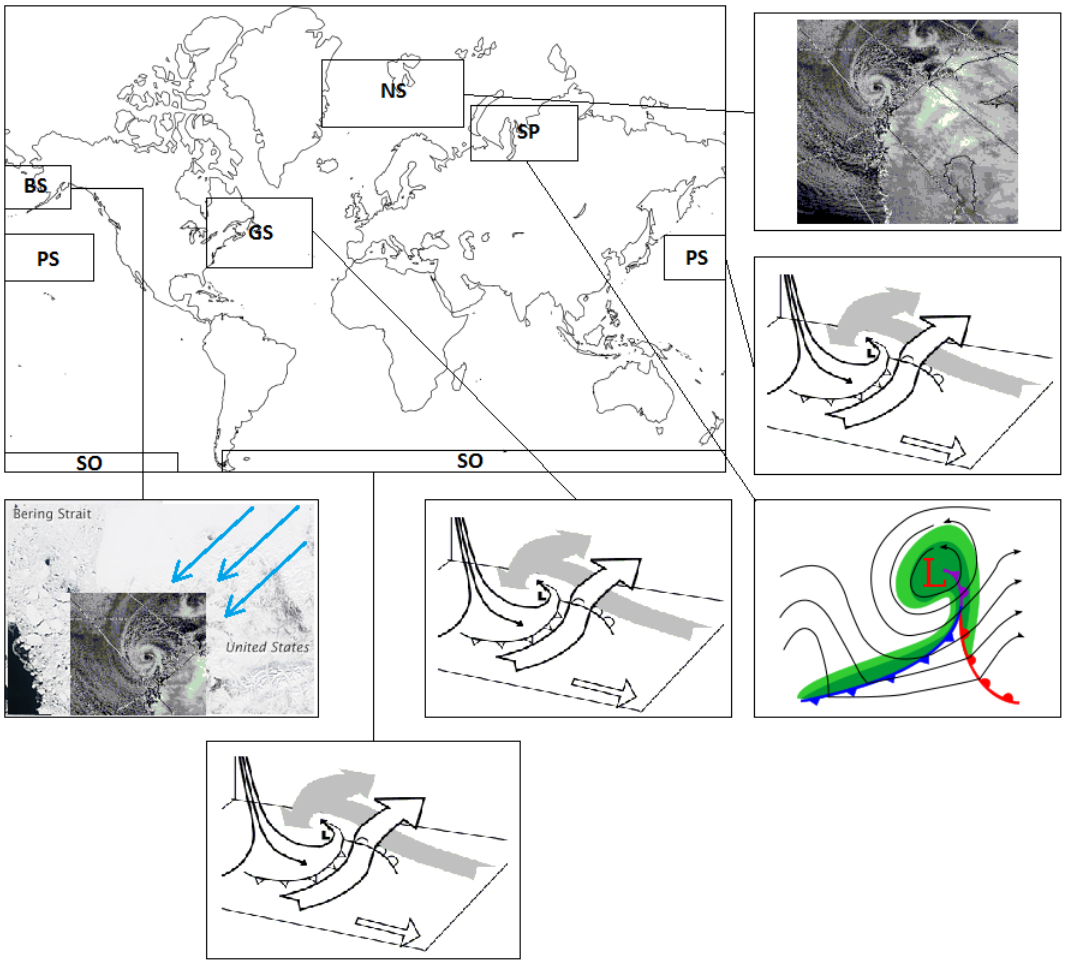

Figure 92: Synoptic analogues for extreme events in heat transport. See text for further details.

spectra shown above (see both the present section and Section 4.3) suggest that a large portion of the power of extreme heat transport events comes from low wave numbers and long periods. These considerations pose two questions:

i) Are the long temporal and spatial scales compatible with a direct correspondence between extreme events and atmospheric synoptic and mesoscale features?

ii) Why are the long spatial and temporal scales evident in the spectral analysis but not in the cross-sectional composites?

Concerning the first question, the wavelet and wavenumber spectra do not preclude a direct

correspondence between the extremes and synoptic systems. In fact, large variations in power at periods below 6 days and at high wavenumbers could still be sufficient to determine whether events with high power at longer periods and broader spatial scales classify as extremes or not. One can therefore reconcile the spectral and synoptic analyses, and interpret the extreme events as regionally coherent synoptic features superimposed on planetary-scale variability. Physically, the extremes are therefore linked to synoptic systems, but these are slaves to the larger scales. Adopting this view, the atmosphere will effect a heat transport "burst" only when a vigorous local circulation is in phase with a larger-scale anomaly.

The above interpretation is based on a synoptic and planetary wave view of the atmosphere. It should be noted that there are alternatives to this framework, such as that suggested by *Larichev and Held* [1995]. In the latter model, heat flux (or, more precisely, potential vorticity flux in the two-layer model framework) is dominated by the largest eddies excited by the inverse barotropic energy cascade, rather than by the deformation-scale ones. In this perspective, the low wave number systems discussed in this section would be interpreted as equivalent barotropic patterns, such as the annular modes. Note that, while in the net equivalent barotropic modes account for no heat transport, they can drive strong local fluxes. A caveat of this interpretation is that it is unlikely to apply to the SH, where the high coherence values discussed in Section 4.4 are more reminiscent of a linear wave regime.

Regarding the second question, there are two distinct aspects which require an answer: a first pertaining to the temporal domain and a second to the spatial one. For what concerns the temporal scales, there are indeed extremes whose duration is beyond baroclinic timescales. This is clearly illustrated in figure 50, where a significant number of extremes has periods beyond 6 days. In addition to this, the importance of the longer periods is possibly over-emphasised by the localised nature of the extremes. As discussed in Section 4.1, a local phenomenon requires a minimum spread in the frequencies (or spatial scales) generating it. Very localised events will therefore yield low-frequency components in a spectral analysis. Because of these two factors, it is not entirely surprising that a broad range of temporal scales play a role in the spectra of the extreme events.

These longer periods are not so evident in the cross-sectional composites, such as figure 59. The figure clearly shows a peak in transport, lasting for approximately 1 day and then rapidly dying away. However, two days before and after the event, the transport is still greater than 1×10^{10}

Wmb^{-1}, which places it around the 90th percentile of the transport distribution for NH DJFs. These values are around seven times smaller than the central peak, but are still in the top percentiles of their distribution, and indicate that the strong transport associated with the extreme has all but died away. The impression that the long timescales are not shown by the cross-sectional composite is therefore due to the predominance of the central peak of the extreme which, in the colourmaps, misleadingly suggests that the extreme events effectively last for a single day.

For what concerns the spatial dimension, the integrated spectra (panels b, d and f of figures 86-89) only refer to the long time periods (10-32 days). It is not necessarily unexpected that these longer periods should correspond to large spatial scales. For example, this would be the case for non-dispersive Rossby waves. For periods below six days, the low wavenumbers are much less prominent (see panels a, c and e of the same figures). The cross-sectional composites of the transport (*e.g.* figure 58) obviously include motions at all periods, and this explains why the long spatial scales are more evident in the integrated spectra than in the colourmaps. Moreover, as was the case for the temporal composites, the central peak of the extremes is very prominent, masking transport values which are nevertheless in the top percentiles of their respective distributions. Finally, the localised nature of the extremes might artificially enhance the importance of the long spatial scales in the spectral analysis.

6 Conclusions

The present chapter focuses on the physical interpretation of the extreme events. A detailed analysis of the circulation characterising the extremes suggests that there are features common to both hemispheres. These include a westward tilt of the heat transport column and ascending airflows in correspondence with the core of the extreme. Many details of the local circulation, however, vary on a regional scale. Seven geographical domains, corresponding to areas of high extreme event occurrence, were selected and analysed. The storm track ones (Atlantic, Pacific and Southern Ocean), were found to have WCB-type extremes, characterised by rapidly ascending warm, moist airflows and cold intrusions (akin to cold conveyors). The North Seas domain was seen to be characterised by cold air outbreaks and warmer air return flows, which recall the precursors to the development of polar lows. The Siberian Plateau and Bering Strait domains, finally, display features consistent with the local low-pressure, warm-core systems commonly found in those regions. These are not captured

by the WCB maps, which adopt very stringent selection criteria. However, they can still effect large heat transports, either by poleward advection of the warm core or by small phase differences between the velocity and MSE fields. It should also be noted that the Bering Strait domain is characterised by a significant variability, and some of the events correspond to cold air outbursts from the Arctic basin.

A detailed investigation of the timescales involved in driving the extreme events was also performed, based on wavelet spectra. The analysis is in terms of v', H' and $v'H'$ power spectra and v' and H' cross-power and coherence-phase plots. The outstanding feature of the v' and H' spectra is the increase in power with increasing period, which is not reproduced in the $v'H'$ spectra. As a result of this feature, baroclinic timescales in v' and H' fluctuations account for a limited part of the power in the extreme $v'H'$ days, comparable to that seen for average transport days. It is also found that power at periods beyond baroclinic timescales is mainly driven by planetary scale waves, confirming that the extreme events are not simply signatures of passing synoptic systems.

As discussed in Section 4.1, the importance of the longer periods is not obvious from a simple evaluation of the bandwidth theorem. However, the theorem only sets a lower bound on the required range of frequencies. Indeed, in quantum mechanics, wave packets are commonly expressed as an integral over an infinite spread of wave numbers (and hence of angular frequencies). Analogously to the quantum wave packet, the present study therefore finds that the localised extreme occurrences in $v'H'$ require a much broader range of frequencies than the lower limit set by the bandwidth theorem. This, combined with the conclusions drawn from the regional domains, motivates the interpretation of the extremes as being characterised by the synoptic features discussed above, superimposed on larger-scale motions. It also highlights the roles of phase and coherence, across a broad range of frequencies, in generating extreme $v'H'$ occurrences.

The importance of phase in generating extremes had already been hypothesised in *Messori and Czaja* [2013b], in the context of baroclinic timescales. Here, the result is established robustly and placed in the context of the time-frequency characteristics of the spectra. Furthermore, the role of coherence is added to the picture. Indeed, it is shown that the key to generating an extreme lies in the coincidence of ideal coherence and phase relationships, and that only one of the two is a necessary but not sufficient condition. An important consequence of this finding is that extreme event days might not be immediately distinguishable from random days in terms of v' and H' spectral characteristics,

meaning that, potentially, moderate v' and H' anomalies at periods beyond baroclinic timescales could drive very large meridional heat transport events.

Part V

Extreme Events in Atmospheric Heat Transport: Idealised Physical and Statistical Analogues

1 Aims of the Chapter

The previous chapters have established the importance of extreme events in heat transport, in both reanalysis and model data. Furthermore, idealised GCM simulations suggest that the sporadic nature of heat transport might be a basic property of the atmosphere. The physical interpretation of the extremes has also been discussed. The present chapter aims to:

i) Relate these extremes and their interpretation to minimal physics models;

ii) Confirm that they are a fundamental component of the dynamics of the atmosphere.

In order to satisfy both points, we will present the simplest possible setups which still capture the salient features of heat transport extremes.

2 Outline

Messori and Czaja [2013, hereafter MC13] robustly established the sporadic nature of heat transport in the atmosphere. The present thesis has shown that this feature is reproduced in both idealised experiments in intermediate complexity, coarse-resolution AOGCMs, and realistic simulations with smoothed SSTs (Chapter III). This suggests that the transport's sporadicity has nothing to do with mesoscale signals and characteristics of the atmosphere's lower boundary but, rather, must be an intrinsic property of waves in the atmosphere. Here, we bring this reasoning a step further, and try to relate the presence and nature of heat transport extremes to basic statistical considerations and minimal physics models.

First, the classic wave framework of atmospheric dynamics is adopted. A sinusoidal picture of v and H anomalies is presented, and it is shown that such simplified statistical view succesfully reproduces the skewed transport distribution seen in the reanalysis data (Section 3). Next a radically different viewpoint is taken, and a quantised picture of heat transport is adopted, basing the approach on the heton model developed by *Hogg and Stommel* [1985a; 1985b] (Section 4). The implication of the model results for the considerations made in the previous chapters will be discussed in Section 5. The main inferences and conclusions reached in this chapter will then be summarised in Section 6.

3 An Idealised Statistical Model

Messori and Czaja [2013b; 2013c] have shown that phase relationships between v' and H' play a key role in driving meridional heat transport extremes. The simplest test of this picture is to adopt a wave view of atmospheric motions, and liken v and H perturbations to plain sinusoidal curves. Figure 93 shows the PDFs resulting from the product of two sine curves of equal amplitude (here unit amplitude) and frequency (here unit angular frequency). Panel a) considers the case where the two waves are in quadrature; panel b) where they are in phase. For the quadrature case, the PDF of the product is symmetric about its median value, namely zero. As can be easily verified, the minimum of the distribution is then found at the median, the most likely value coincides with the two extreme values, and the skewness is null. When the same waves are in phase, the PDF simply shifts towards positive values of $v'H'$, with one of the two most likely values now being at zero, but the shape of the distribution remains unchanged. Both distributions are therefore far from resembling that shown in panel c of figure 32 on page 61.

Next, consider v' and H' being represented by a superposition of sine waves over a broad spectrum of frequencies and amplitudes. The product of a high (say for v') and a low frequency (say for H') wave yields a wave packet-like pattern, with frequent oscillations around zero. This is shown in figure 93c. Hence, this interaction is able to reproduce the near-zero most likely value of the $v'H'$ distribution. The extended positive tail of the PDF, on the other hand, originates from the range of amplitudes considered. As discussed above, the product of two sinusoidal curves will yield most likely values at the two extremes of the distribution. For near in-phase waves, such a distribution will mostly cover positive values, and the magnitude of the extremes will depend on the amplitude

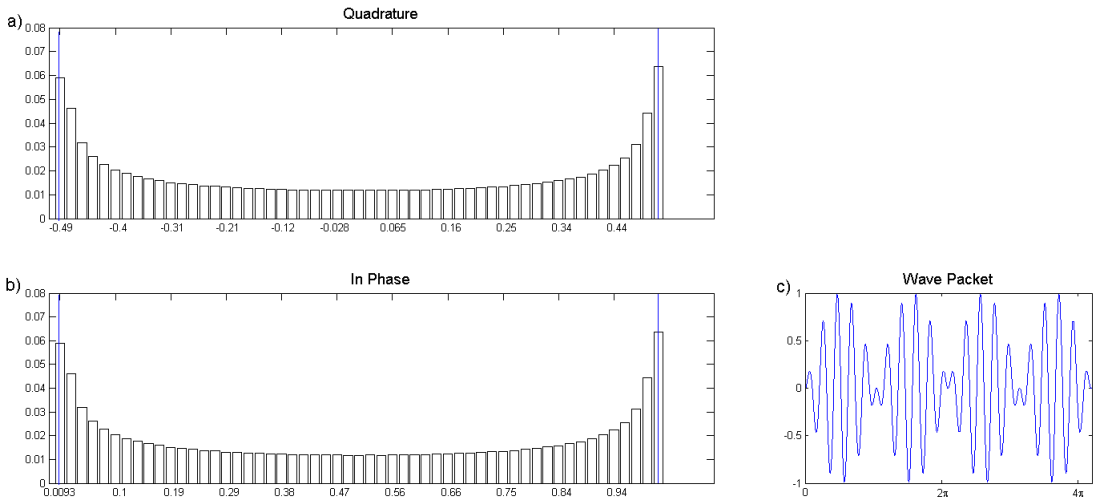

Figure 93: The left hand side panels display the PDFs resulting from the product of two sine curves of unit amplitude and angular frequency. The relative phase is a) $\pi/2$ and b) 0. Panel c) shows the product of two sine curves with unit amplitude and angular frequencies of $\omega = 1$ and $\omega = 10$ respectively.

of the two waves. The effect of having a distribution of random amplitudes for near in-phase waves is, therefore, to spread out these most likely values. Those corresponding to the maxima of the individual distributions will yield a long positive tail. The result is a PDF with a near-zero most likely value and a long positive tail, analogously to that shown in figure 32c. A simple example is given in figure 94. In order to produce the PDF shown in the figure, the following equations were adopted:

$$v' = \sum_i A_i sin(ix + \varphi_i)$$

$$H' = \sum_i B_i sin(ix + \varphi_i + \Delta\varphi) \tag{38}$$

$$v'H' = \sum_i \sum_j A_i B_j sin(ix + \varphi_i) sin(jx + \varphi_j + \Delta\varphi)$$

Here i and j represent summation over a range of frequencies (corresponding to periods ranging from

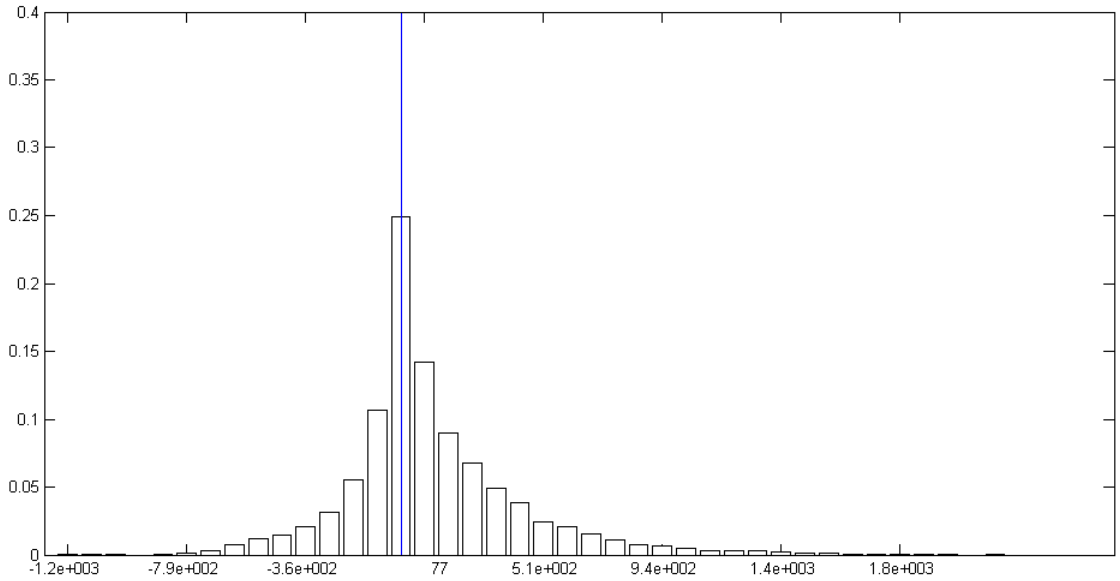

Figure 94: PDF resulting from the product of two linear sums of sinusoidal waves, taken as simple models for v' and H'. The distribution reproduces both the long positive tail and the near-zero most likely value of the $v'H'$ PDF seen in figure 32c [from *Messori and Czaja*, 2013b].

seasonal to daily scales), each with corresponding amplitudes A_i and B_j and random phases φ_i and φ_j. $\Delta\varphi$ is a fixed phase relationship between all v' and H' waves, here set to $\pi/10$ as an indicative near in-phase scenario. This reinforces our interpretation of the distributions of v', H' and their product as resulting from a broad spectrum of travelling waves.

4 The Heton Model

4.1 Model Description

The description of the heton model provided in this section is largely based on *Hogg and Stommel* [1985a; 1985b].

We now adopt a very different picture, abandoning altogether the wave view of the atmosphere. Since the key conclusion of the present thesis is that heat transport is sporadic, as illustrated graphically by the bar plots shown in figure 45 on page 78 and figure 56 on page 94, we try to relate the extremes to a quantised view of heat transport. Indeed, by drawing on ideas from other fields of Physics (*e.g.* Quantum Mechanics) one can think of heat as a discrete, particle-like process. This

idea is not entirely new, and has been applied to the atmospheric and oceanic sciences by *Hogg and Stommel* in their heton model [*Hogg and Stommel*, 1985a; 1985b]. Such model was outlined in Chapter II, Section 2.5 on page 53, and is described here in greater detail. It will be used to help interpret the results presented in this thesis from a very different viewpoint compared to that adopted thus far.

The heton model is a two-layer system, seeded with point geostrophic vortices. These are expressed as local potential vorticity anomalies, and give rise to deviations in the interface between the two layers, hence inducing flow anomalies. The nature of the interface deviations depends on the sign of the vortices, and on the layer in which these are placed. The two layers have different potential temperatures; for a mock atmosphere or ocean, the upper layer would have a higher temperature than the lower one. Deformations in the interface can therefore be likened to thermal anomalies. The terminology adopted here will call *positive* vortex one that drives a cyclonic circulation, and *negative* vortex one associated with an anti-cyclonic one. A positive vortex will draw the interface towards itself, regardless of whether it is in the top or bottom layer of the model. Conversely, a negative one will repel it. Depending on the interface deviation associated with it, a vortex may be further termed *hot* or *cold*. Any vortex leading to a downward displacement of the interface (namely a negative vortex in the upper layer or a positive vortex in the lower one) will be hot. Conversely, any vortex leading to an upward displacement of the interface (namely a positive vortex in the upper layer or a negative vortex in the lower one) will be cold. One can further consider combinations of two vortices. Where both share the same sign, they will form a barotropic pair. Where their signs differ, they form a baroclinic pair. If in different layers, the two vortices in a baroclinic pair will both displace the interface in the same direction; one can therefore have hot or cold baroclinic pairs. Because they have a well defined "temperature", when such pairs move they will transport heat in the direction of movement. The heat-transporting baroclinic pairs are termed hetons. Figure 95 provides a graphical summary of the terminology described here, relating the different vortices to the associated interface deformations. The top four panels show examples of hot and cold positive and negative vortices, while the bottom panels show vortex pairs. The left hand side one depicts a hot heton, while the right hand side one illustrates an example of barotropic pair effecting no net heat transport.

An example of a two-layer model and the anomalies associated with a single cold heton are illustrated in figure 96. Panel a) shows the two layers of the model. L_H is the total thickness of the

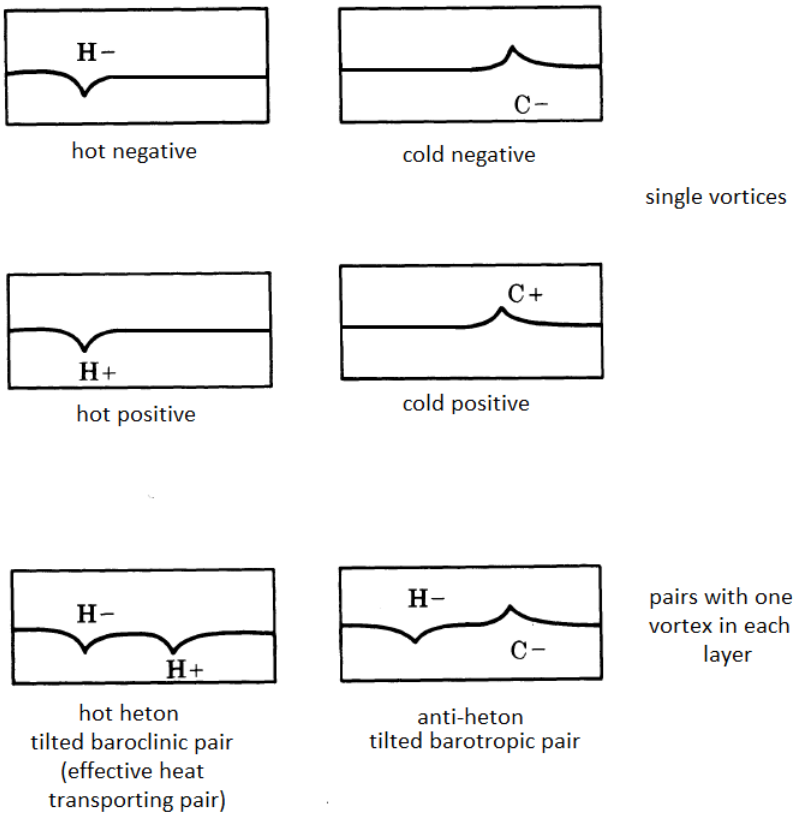

Figure 95: Interface deformations resulting from point geostrophic vortices in a two-layer model. The top four panels show examples of hot and cold positive and negative vortices, while the bottom panels show vortex pairs. The bottom left hand panel depicts a heat-transporting pair. The bottom right hand one depicts a pair which effects no net transport of heat [adapted from *Hogg and Stommel*, 1985a].

layers, while h is the thickness of a single one. Each layer has a potential vorticity anomaly q associated with the vortex, and a corresponding streamfunction ψ. Here, we limit ourselves to a qualitative description of the setup; the mathematical relations governing the model are discussed below. Panel b) illustrates the upper-layer potential vorticity profile of the heton. Panels c) and d) show the induced interface deviation and upper-layer azimuthal velocity. r is simply the radial distance from the heton core. Note that, since the vortex is a point anomaly, mathematically represented by a *delta*-function, the velocity at the location of the vortex tends to infinity (or, more precisely, varies as $1/r$). The heton in the example has a cold core, and the interface domes upwards (compare this to the hot heton depicted in figure 95). This induces a cold temperature anomaly in correspondence with the doming interface, and a cyclonic circulation in the top layer of the model. The azimuthal velocity of the bottom layer (not shown) would be anti-cyclonic. A more complex example of hetons in a two-layer model is illustrated in figure 97. This shows the upper-layer temperature anomalies associated with a cluster of randomly seeded hetons (marked by the red dots). In this specific example the hetons are again cold, meaning that they cause an upward deformation of the layer interface. The red points therefore correspond to the largest negative temperature anomalies. The large-scale cyclonic circulation around the heton cluster is evident but, as can be seen, the hetons also give rise to finer anomaly patterns in the interior of the cluster.

The mathematical formulation of the heton model is reasonably simple, and is based on the *quasi*-geostrophic formulation of potential vorticity. Here *quasi*-geostrophy refers to a flow where the advection terms in the momentum equation are too large to be ignored but still one order of magnitude smaller than the Coriolis and pressure gradient terms. The latter are the terms which define geostrophic balance, as discussed in Chapter II, Section 1.3 on page 26. Vorticity is simply the curl, or rotational part, of the velocity field. The vorticity of an air parcel can be decomposed into a planetary component, due to the Earth's rotation, and a relative component, due to the parcel's motion relative to the planetary frame. The sum of the two gives the absolute vorticity, namely the vorticity for an inertial frame. *Quasi*-geostrophic potential vorticity is then the sum of the absolute vorticity and a stretching term related to the vertical displacement of isentropes. It is a useful quantity because it is conserved in adiabatic, frictionless flows.

In a two-layer model seeded with point potential vorticity anomalies, the streamfunctions[7] ψ_1 and

[7] The streamfunction ψ of a two-dimensional flow $\vec{u}$ is commonly defined as: $\vec{u} = -\nabla \times \psi$.

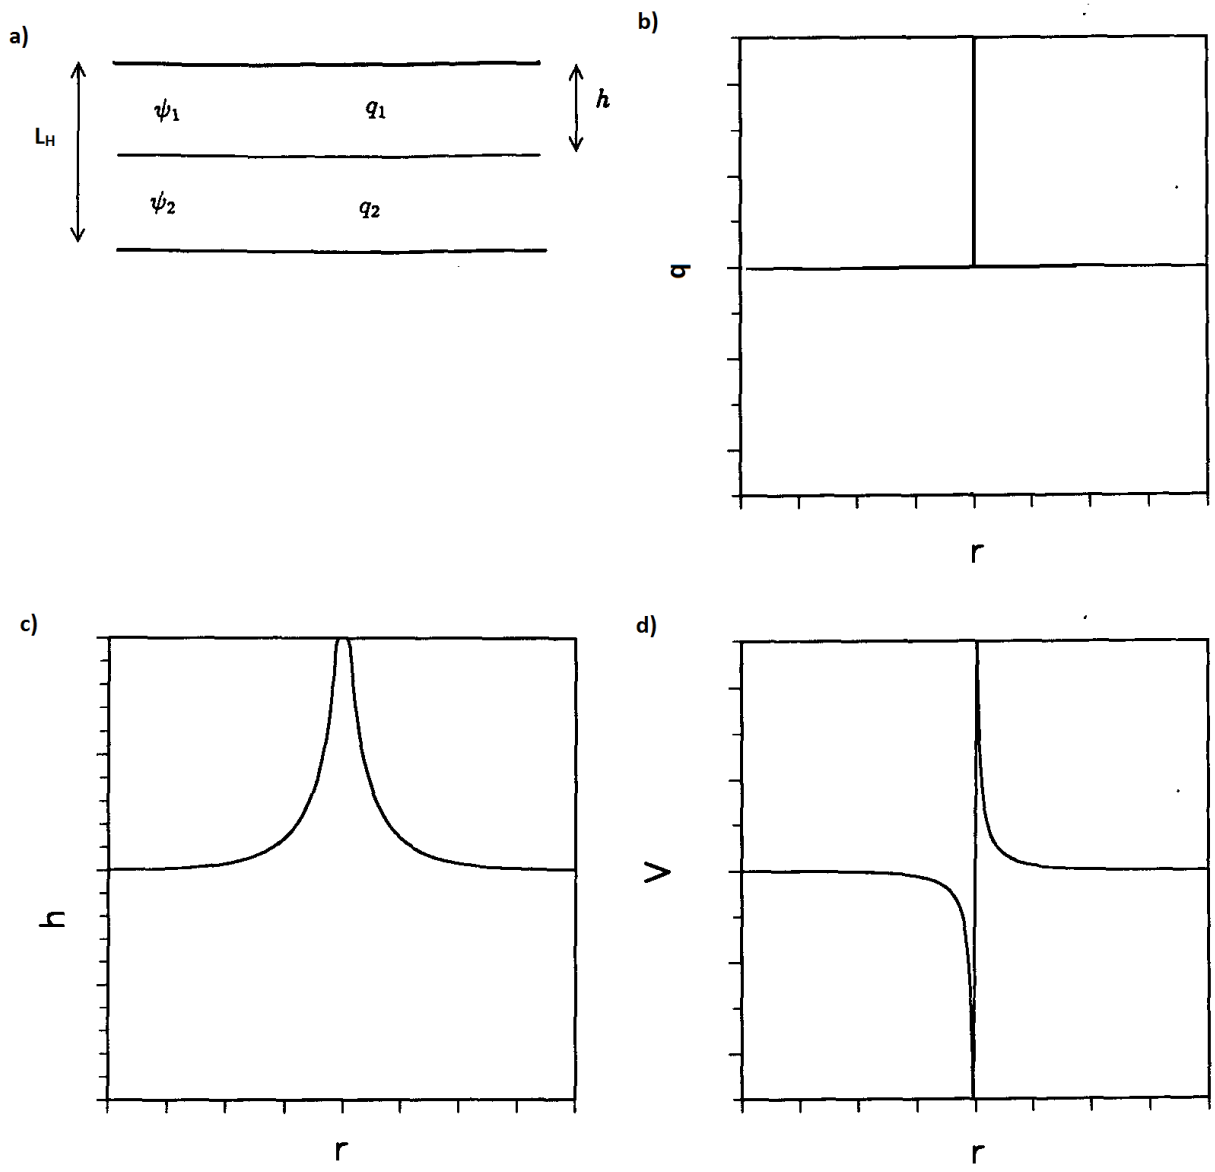

Figure 96: Panel a) shows the two-layer model cross-section. L_H is the total thickness of the model, h is the thickness of a single layer, and q and ψ are the potential vorticity anomalies and streamfunctions associated with each layer. Panel b) shows the upper-layer vorticity profile of a single, cold heton. Panels c) and d) show the induced interface deviation and upper-layer azimuthal velocity. r is simply the radial distance from the heton core [adapted from *Legg and Marshall*, 1993].

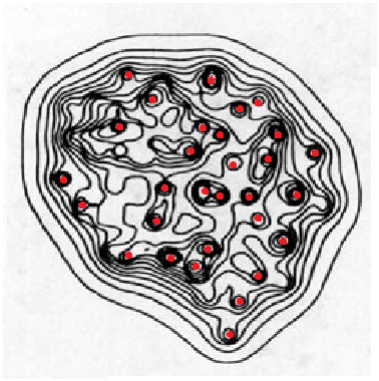

Figure 97: Temperature anomalies associated with a cluster of randomly seeded hetons (marked by the red dots). The hetons are cold, meaning that the red points correspond to the largest negative temperature anomalies. The contour interval is 0.005 K [adapted from *Legg and Marshall*, 1993].

ψ_2 in the two layers must satisfy the following equations:

$$\nabla^2\psi_1 - \frac{1}{2}\lambda^{-2}(\psi_1 - \psi_2) = \delta_1(r) \tag{39}$$

$$\underbrace{\nabla^2\psi_2}_{A} + \frac{1}{2}\lambda^{-2}\underbrace{(\psi_1 - \psi_2)}_{B} = \delta_2(r) \tag{40}$$

Here, the subscripts identify the layer and $\delta(r)$ is a delta function. λ is the Rossby radius of deformation[8], commonly expressed as:

$$\lambda = \frac{N_{BV}L_H}{f}, \tag{41}$$

where N_{BV} is the Brunt-Väisälä frequency, L_H is a vertical length scale and f is the Coriolis parameter. Here, λ can be interpreted as an interaction length scale for baroclinic pairs, as discussed below. A is a relative vorticity term, while B is a stretching term accounting for deformations in the interface between the two layers, as discussed above. There is no planetary vorticity factor because the heton model adopts an f-plane formulation, meaning that the Coriolis parameter is set to a constant. Equations 39 and 40 are the standard expressions for the streamfunctions of a two-layer

[8]The Rossby radius of deformation is the length scale at which rotational and buoyancy effects become comparable. Consequently, it sets the scale for atmospheric cyclonic and anti-cyclonic motions.

baroclinic model. Detailed derivations can be found, among others, in *Holton* [1979] and in *Legg and Marshall* [1993]. The barotropic and baroclinic streamfunctions, $\psi = \psi_1 + \psi_2$ and $\psi' = \psi_1 - \psi_2$ respectively, are then given by:

$$\nabla^2 \psi = \delta_1 + \delta_2 \tag{42}$$

$$\nabla^2 \psi' - \lambda^{-2} \psi' = \delta_1 - \delta_2 \tag{43}$$

From the streamfunctions ψ_1 and ψ_2, one can determine the motion of any given vortex as a result of its interactions with the other vortices in the domain. By the definition of the streamfunction, the horizontal velocity components are given by: $u_i = -\frac{\partial \psi_i}{\partial y}$ and $v_i = \frac{\partial \psi_i}{\partial x}$, $i = 1, 2$. The streamfunctions for each layer can then be re-written in terms of their barotropic and baroclinic counterparts:

$$\psi_1 = \frac{1}{2}(\psi + \psi') = \frac{1}{2}[(s_1 + s_2) log(r) - (s_1 - s_2) K_0(r/\lambda)] \tag{44}$$

$$\psi_2 = \frac{1}{2}(\psi - \psi') = \frac{1}{2}[(s_1 + s_2) log(r) + (s_1 - s_2) K_0(r/\lambda)], \tag{45}$$

where s_i are the nondimensional strengths of the point vortices, r is the radial distance from the vortex and K_0 is a modified Bessel function of the second kind, namely a type of exponentially decaying function. It is the $K_0(r/\lambda)$ term which motivates the interpretation of λ as an interaction scale for baroclinic vortices.

In a cloud of vortices, each vortex will be affected by the presence of all the others. Using the above framework, the equations of motion for a cloud of N vortices can then be expressed, for each vortex ν_i, in terms of the effect of all other vortices ν_j, $j \neq i$:

$$u_i = -\sum_{j \neq 1}^{j=1...N} \frac{1}{2} s_j \left(\frac{y_i - y_j}{r_{ij}^2} \right) \left[1 + \beta_{ij} \frac{r_{ij}}{\lambda} K_1 \left(\frac{r_{ij}}{\lambda} \right) \right] \tag{46}$$

$$v_i = -\sum_{j\neq 1}^{j=1\ldots N} \frac{1}{2} s_j \left(\frac{x_i - x_j}{r_{ij}^2}\right) \left[1 + \beta_{ij} \frac{r_{ij}}{\lambda} K_1 \left(\frac{r_{ij}}{\lambda}\right)\right], \tag{47}$$

where (x_i,y_i) is the position of vortex ν_i, r_{ij} is the separation between vortex ν_i and vortex ν_j and β_{ij} is a binary quantity, with positive unit value if ν_i and ν_j are in the same layer and negative unit value otherwise. K_1 is a different modified Bessel function of the second kind. If, in the same model domain, one seeds a number of different pairs, these will therefore interact. The sign of β_{ij} in equations 46 and 47, implies that vortices in same layer will have a stronger interaction than those in different layers. When discussing vortex pairs, the pairing is based on interaction strength, which is therefore not inversely correlated to spatial distance.

Baroclinic pairs are seen to have a limited interaction range, given by λ; within this range, they act to produce a shear between the two layers, and therefore split neighbouring pairs. It is instructive, in this regard, to consider a very simple interaction between two hetons, one hot and one cold. This example, discussed in detail in *Hogg and Stommel* [1985a], can be interpreted physically as representing thermal processes in a two-layer ocean. The vortices are initially perfectly aligned with their companions, and are therefore only affected by the other pair. As soon as the vertical alignment is altered, and the heton has split, companion vortices also interact among themselves. What is seen to happen is that both hetons initially move towards each other, while the split between the individual vortices increases. During this phase, a meridional heat transport occurs, which can be shown mathematically to correspond to a decrease in the system's APE. As the four vortices approach, new pairings are formed between vortices in the same levels. These two new pairs interact very strongly, and rapidly move in opposite directions. This motion effects no heat transport because the new pairs are in same layer. The trajectories of the vortices in this simple setup are shown in figure 98. Numerous experiments of a similar type can be performed, involving more hetons or adding barotropic pairs to the picture. The key point of these simple setups is that hetons with the same temperature (*i.e.* both hot or cold), tend to repel if further away than one deformation radius, and split if closer. On the other hand, hetons with opposite temperatures will attract at far range and split when closer.

Due to the easily tractable mathematical framework of the model, it is possible to go beyond the

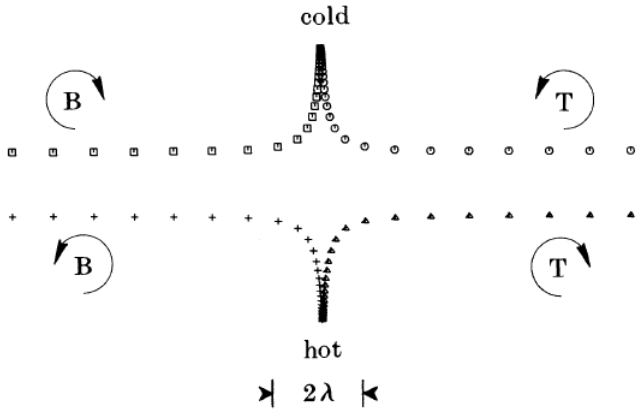

Figure 98: The trajectories of individual vortices in a two-heton experiment. The hetons are unit strength, initially aligned, and of opposite temperatures [from *Hogg and Stommel*, 1985a].

simple case presented above and consider the interactions of large groups of hetons. In a uniformly dense array of same temperature hetons, the baroclinic fields of the individual vortices cancel out. At the edges of the array, however, the field is non-zero and a rim of baroclinic velocity forms. Individual vortex pairs which attempt to leave the array are split by this rim current and cannot escape. Since, in the near field, hetons with same temperature do not repel, clusters of these vortices can form. These effectively behave like a composite heton. Figure 99 shows how the boundaries of the initial heton cloud deform, and how composite hetons can form at the rim. If a sufficiently large composite forms, this can then "break" the rim current. This behaviour is termed *hetonic explosion*, an example of which is shown in figure 29 on page 55. The composites, however, often do not have the same number of top and bottom vortices, and therefore do not explode outwards until interactions with other pairs balance the numbers. The spatial details of the composite hetons depend on the deformation radii of the individual vortices. From a physical standpoint, composite hetons can be likened to the varying amplitudes of atmospheric waves. In the context of heat transport, hetonic explosions can be taken as analogues for extreme transport events, since they correspond to very large, rapidly moving thermal anomalies.

4.2 Heat Transport Calculation

A key part of the framework described above is that it allows to compute the heat transport as a result of heton motions. The technique applied here is an extension of *Legg and Marshall* [1993], and

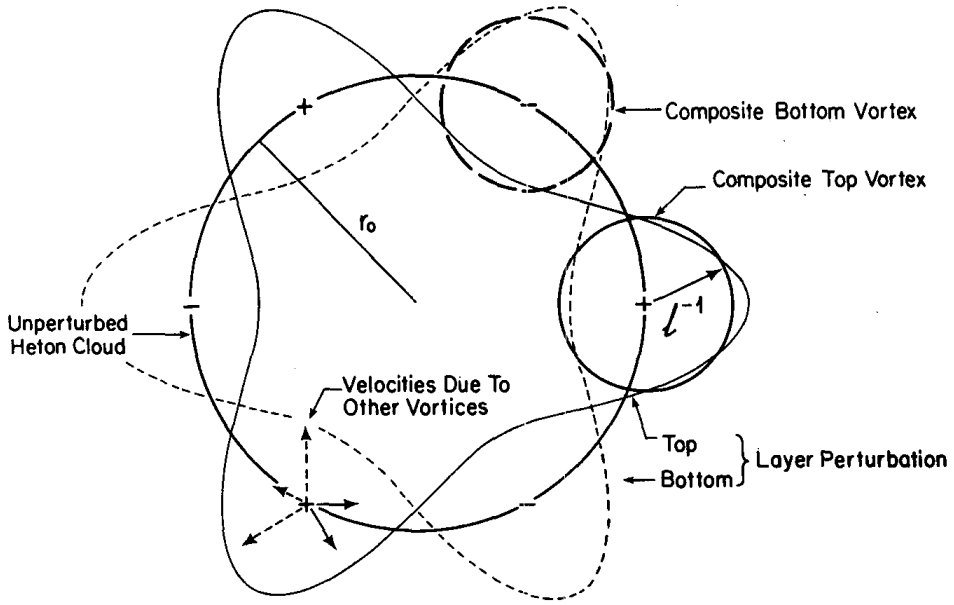

Figure 99: Schematic of the deformation of a heton cloud and formation of composite hetons at the rim. The continuous line corresponds to the top layer, the dashed line to the bottom one [from *Hogg and Stommel*, 1985b].

was originally applied to an ocean model. It considers hetons of strength e_0 seeded at a rate N_r per unit time, within a given domain. The hetons are free to exit and re-enter the domain. Here, we try to relate this framework to an atmospheric domain. If seeding cold hetons, this setup would mimic atmospheric radiative heat loss at high latitudes.

For N^{in} being the number of hetons within the domain, and N^{out} being the number without, the following will be true at all times:

$$\frac{\partial}{\partial t}\left(\sum_{i=1}^{N^{in}} e_i^{in} + \sum_{i=1}^{N^{out}} e_i^{out}\right) = \frac{\partial}{\partial t}\left(Q_r + Q_d\right) \tag{48}$$

Here, e is the strength of the vortices, expressed as an energy (thermal) anomaly in J. In order to make our computation widely applicable, this is left as a variable. $\frac{\partial}{\partial t}Q_r = e_0 N_r$ is the energy change due to the new hetons being seeded (for N_r expressed in *per unit time*). Provided the new seeds are cold hetons, in atmospheric terms this would correspond to the radiative cooling term mentioned above. $\frac{\partial}{\partial t}Q_d$ is the energy change due to changes in heton strength over time. If heton strength decays, this term could be interpreted as a damping factor for thermal anomalies in the atmosphere.

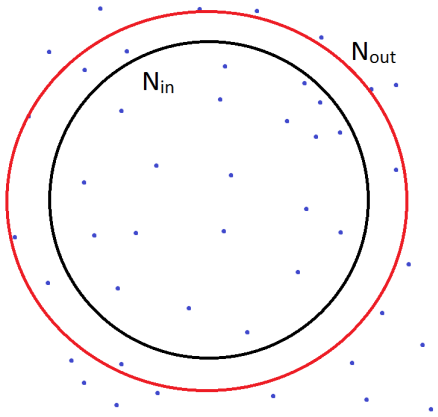

Figure 100: Schematic showing a seeding (black circle) and a heat transport (red circle) domain. The heat transport described by equation 50 is the integrated heat flux across the red domain's boundary. N_{in} are the hetons within the red circle and N_{out} those without.

Note that, in the above discussion, the seeding domain and the domain relative to which the heat transport is computed are taken to be the same. It is, however, entirely possible to adopt the same formulation while computing the transport across a domain boundary that does not coincide with the seeding domain. This setup is illustrated in figure 100. The black contour shows the seeding domain, and the red contour a hypothetical, larger domain relative to which the heat transport is computed. N^{in} and N^{out} are then the numbers of hetons within and without the red curve.

For the case where the hetons decay over time at a rate α_d, the corresponding energy change satisfies:

$$\frac{\partial}{\partial t} Q_d = -\alpha_d \left(\sum_{i=1}^{N^{in}} e_i^{in} + \sum_{i=1}^{N^{out}} e_i^{out} \right) \tag{49}$$

Combining the above with equation 48, the following can be derived:

$$\frac{\partial}{\partial t} \left(\sum_{i=1}^{N^{in}} e_i^{in} - Q_r \right) = \frac{\partial}{\partial t} \left(Q_d - \sum_{i=1}^{N^{out}} e_i^{out} \right)$$

$$\Rightarrow \underbrace{\frac{\partial}{\partial t}\left(\sum_{i=1}^{N^{in}} e_i^{in}\right)}_{A} + \underbrace{\alpha_d \left(\sum_{i=1}^{N^{in}} e_i^{in}\right) - \frac{\partial}{\partial t} Q_r}_{B} = -\frac{\partial}{\partial t}\left(\sum_{i=1}^{N^{out}} e_i^{out}(1+\alpha_d)\right) = H_T, \tag{50}$$

where H_T is the heat flux across the domain boundaries. The two terms on the left hand side represent (A) the local change in heat content and (B) the diabatic processes within the domain. The heat transport therefore depends on the number of vortices within and without the domain. Composite hetons crossing the domain boundary will clearly have a large effect on it.

4.3 Experiment Design

Two different experiment setups are designed. Both consider a circular domain, meant to represent an idealised polar region. The domain size is chosen to roughly coincide with a spherical cap bounded by the 60° N latitude circle. The hope is to reproduce the sporadic nature of heat transport witnessed in the real atmosphere in terms of hetonic heat transport into the cap. Because of this, all atmospheric parameters, listed in table 9 a), were chosen to match realistic values measured in the atmosphere and, where possible, were set to round values in the spirit of the idealised simulation performed here.

We term the first experiment *initial value experiment*. Cold hetons are initialised at random locations within the domain, and are then left to evolve freely over time. Their strength at time t, e_t varies according to:

$$\frac{de_t}{dt} = -\alpha_d e_t \tag{51}$$

No additional vortices are seeded at later timesteps. The heat transport is computed across the domain boundary and at a number of larger radii. While this setup does not relate directly to the type of continuous radiative cooling one might envisage in the atmosphere, it has the advantage of being very computationally inexpensive. It is adopted mainly to study the heton dynamics and verify the existence of composite hetons. To get meaningful statistics, the experiment is replicated 50 times.

The second experiment, termed *continuous seeding experiment*, considers the same initial domain. Now, however, cold hetons are seeded within the domain at every timestep, always at random locations. There is an initial spin-up phase, where the number of hetons grows with time. Equilibrium

is reached once the initial seeds decay to zero, and the new seeds balance the decay but do not increase heton numbers further. This setup mimics a constant radiative cooling, and presents a more immediate physical link to the real atmosphere than the initial value setup. Again, heat transport is computed at the domain boundary and at various other radii. In order to get meaningful outputs, each setup is run at equilibrium for 160 model timesteps, corresponding to just under one month. This number is heavily limited by the computational facilities available.

Both the above experiments present a large number of parameters which need to be set. The two main ones are heton strength, e, and number density, n. The latter quantity is closely related to N_r, and is simply the spatial density of hetons within the model domain. Obviously, once e is set so is n, and *vice-versa*. For the initial value experiment, heton strength can be related to an initial heat content anomaly in the domain. In the real atmosphere, such anomaly is a volume integral of a temperature perturbation, given by:

$$E = \iiint dxdydp \left(\frac{c_p T'}{g} \right) \approx c_p T' \frac{\Delta p}{g} L_x L_y, \tag{52}$$

where E is the heat content anomaly, c_p is the specific heat capacity of dry air at constant pressure and L_x and L_y are typical horizontal length scales. If we take E to represent the heat content anomaly of the whole domain (which will henceforth be called E_d), these represent the size of the domain itself. For the anomaly associated with an individual heton, e, these would be related to the latter's interaction length scale, λ. Using the domain parameters specified in table 9, and setting $\Delta p = 7.75 \times 10^4 \ Pa$, the approximate thickness of the troposphere, and $T' = 5 \ K$, we obtain $E_d \approx 1.18 \times 10^{21} \ J$. There is certainly some degree of arbitrariness in selecting the exact values for these quantities, but they are generally chosen to be within realistic ranges for the real atmosphere in the mid-to-high latitudes.

Since there is no continuous seeding, $N_r = 0$, except for the first time step, where $N_r = N_{tot} = N_0$, the initial number of hetons within the domain. The strength of an individual heton can then be expressed as:

$$e = \frac{E_d}{N_0} \tag{53}$$

For the continuous seeding experiment, heton strength can be related to the rate of radiative cooling in the atmosphere in the mid-to-high latitudes, which is of order $H = 200\ Wm^{-2}$ [*Campbell and Vonder Haar*, 1980]. This value refers to the net radiative cooling of an air column; a rough conversion to a rate in K/day can be performed by assuming a constant cooling throughout the column. By taking the values for specific heat capacity (c_p), density (ρ) and vertical length scale (L_H) provided in table 9 a), H can be expressed as:

$$H\{K/day\} = \frac{H\{Wm^{-2}\} \times 3600 \times 24}{c_p \rho L_H} \approx 2.87 K/day$$

Such value, together with the optional heton decay rate, effectively sets the mean heat transport in the model's steady state.

For $N^{tot} = N^{in} + N^{out}$, ignoring decay, we have:

$$\frac{\partial N^{tot}}{\partial t} = \frac{HA}{e}, \tag{54}$$

where A is the area of the seeding domain. For a seeding timestep t_s, this provides an equation for e:

$$e = \frac{HAt_s}{N^{tot}} \tag{55}$$

The next step is to relate the heton strength, e, to its nondimensional counterpart, s, in terms of which the model is formulated (see equations 44 and 45). Starting from equation 52, the temperature anomaly can be re-formulated in terms of heton strength according to:

$$e \approx c_p T' \frac{\Delta p}{g} \lambda^2 = \frac{1}{4} \frac{\rho c_p f}{g\alpha} \lambda^2 q_s s, \tag{56}$$

where α is the coefficient of thermal expansion for dry air and q_s is a dimensionalisation factor for s in $m^2 s^{-1}$ [*Legg and Marshall*, 1993].

In the idealised setups presented here, it is difficult to relate e to a physical quantity that could suggest the most appropriate value. Combinations of e and n, however, can be readily related to either a heat content anomaly or a radiative cooling rate (see equations 53 and 55). The approach adopted here is therefore to test a range of combinations of these two parameters, such that their product is always constant and corresponds to physical values of E_d and H. The different combinations adopted

a)	s	N_0	b)	s	N_r
	100	10		50	2
	50	21		25	4
	25	41			

Table 8: The different heton strength-heton number combinations for a) initial value and b) continuous seeding experiments. N_0 is the initial number of hetons within the domain; N_r is the number of hetons seeded at every timestep.

are shown in table 8. The number of hetons is presented in terms of N_0 and N_r, rather than n, to provide values that can be more readily related to the dynamics of the experimental setups described above. Since only whole numbers of hetons can be seeded, the E_d and H values corresponding to the combinations shown below vary slightly from experiment to experiment. For the initial seeding case, the number of hetons was simply rounded to the closest integer. For the continuous seeding case, however, since the number of hetons is much smaller, it was chosen to keep the same linear relationship between s and N_r for both experiments. The number of combinations for the continuous seeding experiment is severely limited by the computational cost of running the simulation.

Table 9 provides a summary of the quantities relevant to the two experiments. All atmospheric parameters, described in table a), were chosen to match realistic values measured in the atmosphere and, where possible, were set to round values in the spirit of the idealised simulation performed here. The chosen values for the atmospheric parameters are then used to compute a domain-integrated heat content anomaly, E_d, for the initial value experiment, according to equation 52. The corresponding radiative cooling rate for the initial value experiment is set to a realistic atmospheric value. The atmospheric parameters can also be used to compute the value of the Rossby radius of deformation λ:

$$\lambda = \frac{N_{BV} L_H}{f} \sim 770\, km$$

The domain size is selected so as to mimic a polar cap bounded by the 60° latitude circle, and its radius is set to 4λ. The heton lifetime, from time of seeding to time of disappearance, is set to approximately 1 week, to match a realistic thermal relaxation timescale in the atmosphere. Since the model is formulated in terms of non-dimensional variables, the dimensional parameters listed in the table are simply used to rescale model output to SI units.

a)

Parameter Name	Value	Notes
Combined layer height (L_H)	10 km	Approximate height of tropopause
Combined layer pressure thickness (Δp)	$p = 7.75 \times 10^4$ Pa	Approximate thickness of troposphere
Reference velocity (v_{ref})	5 ms^{-1}	
Reference tempertaure anomaly (T')	5 K	
Brunt-Vaisala frequency (N_{BV})	10^{-2} S^{-1}	e.g. *Peixoto and Oort* [1992]
Coriolis parameter (f)	1.3×10^{-4}	Value at 65° N
Coefficient of thermal expansion for dry air (α)	3.9×10^{-3} K^{-1}	From *Dixon* [2007]
Air density at mid-levels (ρ)	0.6 kgm^{-3}	From *ICAO* [1952]
Specific heat capacity for dry air (c_p)	1005 $Jkg^{-1}K^{-1}$	e.g. *Peixoto and Oort* [1992]

b)

Parameter Name	Initial Value	Continuous Seeding	Notes
Heat content anomaly (E_d)	1.22×10^{21} J	-	See equation 52
Radiative cooling rate (H)	-	200 Wm^{-2}	From *Campbell and Vonder Haar* [1980]
Domain radius	3080 km	3080 km	$\sim 60°$ latitude circle
Domain area (A)	2.97×10^7 km^2	2.97×10^7 km^2	
Model timestep	1.54×10^4 s	1.54×10^4 s	
Seeding timestep (t_s)	1.54×10^4 s	1.54×10^4 s	Same as model timestep
Dimensional heton strength (q_s)	3.85×10^6 m^2s^{-1}	3.85×10^6 m^2s^{-1}	
Dimensional length scale (λ)	770 km	770 km	Rossby radius of deformation
Dimensional time scale	1.54×10^5 s	1.54×10^5 s	
Heton lifetime (T_H)	$4 \times 1.54 \times 10^5$ s	$4 \times 1.54 \times 10^5$ s	

Table 9: a) Atmospheric and b) model parameters for both initial value and continuous seeding experiments.

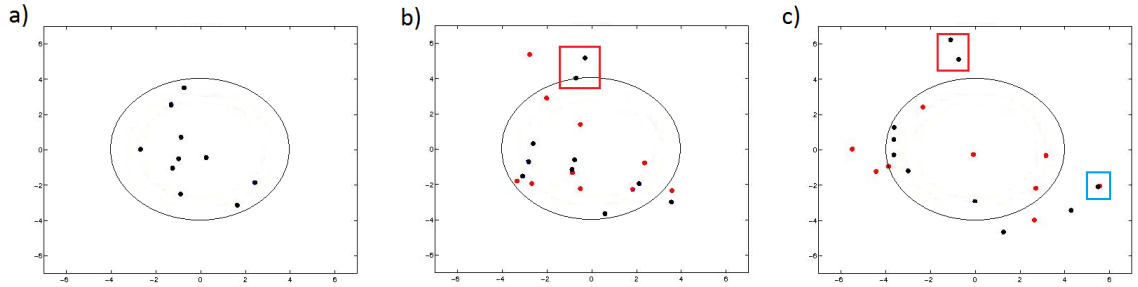

Figure 101: Snapshots of vortex positions during the first three timesteps of the initial value experiment with $N_0 = 10$ and $s = 100$. The black dots are vortices in the upper layer; the red dots ones in the lower layer. The red rectangles in panels b) and c) highlight a rotational trajectory. The blue rectangle in panel c) highlights a linear trajectory. The black contours mark the seeding domain with radius of 4λ.

4.4 Initial Value Experiments

We begin the analysis of the heton model by focussing on the output of the three initial value simulations, whose reference parameters are given in table 8. As discussed above, this experiment considers an expanding heton cloud and is mainly aimed at investigating heton dynamics. Because there is no heton seeding after the initial timestep, the experiment has a limited physical significance in relation to the real atmosphere. As could be naively guessed from the dynamics described in *Hogg and Stommel* [1985a; 1985b], for low number densities no composite hetons form, and the individual vortices simply propel outwards from the seeding domain. This is the case of the setup with 10 hetons of non-dimensional strength 100. There are, obviously, numerous interactions between the hetons, but these do not generally involve large numbers of vortices. Similarly to what seen in the *Hogg and Stommel* experiments, the top and bottom vortices forming a heton often detach themselves from their original partners. When two detached vortices in the same level meet, they are seen to exhibit circular trajectories around a common centre. The typical evolution of the heton cloud for this experiment is shown in figure 101. Panels a), b) and c) show the three initial time steps of one of the iterations of the experiment. The red rectangles in panels b) and c) highlight an example of the rotational behaviour described above. The blue rectangle in panel c) highlights a pair which is tightly coupled and proceeds on a linear trajectory away from the seeding domain. The latter domain is marked by the circular contours, and has a radius of 4λ.

For higher number densities, some of the vortices get close enough to form clusters, which then

propel outwards at greater speed than the individual hetons. These events are comparable to the hetonic explosions described in *Hogg and Stommel* [1985b], where groups of vortices detach themselves from the areas of denser population. Due to the fact that many vortices lose their original partners, these vortex clusters are often imbalanced, with more vortices in one layer than in the other. Figure 102, taken from the experiment with $N_0 = 21$ and $s = 50$, displays such composite behaviour. Panels a) - f) represent the first six timesteps of a specific run of the experiment. The black and red dots are top and bottom layer vortices respectively. The circular contours mark the seeding domain with radius of 4λ. Starting from panel c), three clusters of vortices (marked by the black rectangles) are clearly visible across the domain. Two of the clusters in the figure contain similar numbers of upper and lower level vortices, and propel quickly outwards (panel d)). The third contains mostly bottom level vortices, and this initially hinders its radial accelleration. The fastest moving of the three clusters leaves the domain, crossing the bottom right hand boundary (panels d) - e)). The remaining two clusters interact, and one of them sheds top-layer vortices which go to balance numbers in the other one. It should be noted that, while the cluster formation begins within the seeding domain, the clusters tend to become well formed only beyond the latter domain's boundary. If, as hypothesised here, hetonic explosions correspond to large heat transport events, these will be measured only if the heat transport is computed across a domain with radius larger than the seeding domain's one.

Since the hetons are not continuously seeded, and their strength decays over time (see equation 51 on page 163), the heat transport across the domain boundary falls off very rapidly. Figure 103 presents the heat transport time series for the three initial value simulations. The transport is computed across the boundary of the seeding domain ($r = 4\lambda$). The continuous curves are the values averaged across all 50 iterations performed for each setup. The blue curve coresponds to the $N_0 = 10$, $s = 100$ case, the red curve to the $N_0 = 21$, $s = 50$ case and the black curve to the $N_0 = 41$, $s = 25$ case. The dashed lines are the 1σ limits for the red curve. The transport values during the initial timesteps are extremely large, approximately two to three times larger than what one would expect in the real atmosphere. This is mostly due to the very quick expansion of the heton cloud, leading to term A in equation 50 on page 163, becoming very large. The decrease in transport is very rapid, and after approximately 2 days nearly all the vortices have left the seeding domain. In the experiments with more numerous, weaker hetons, the transport decays more slowly. This is consistent with the fact that weaker hetons interact less strongly with one another, and therefore take longer to cross

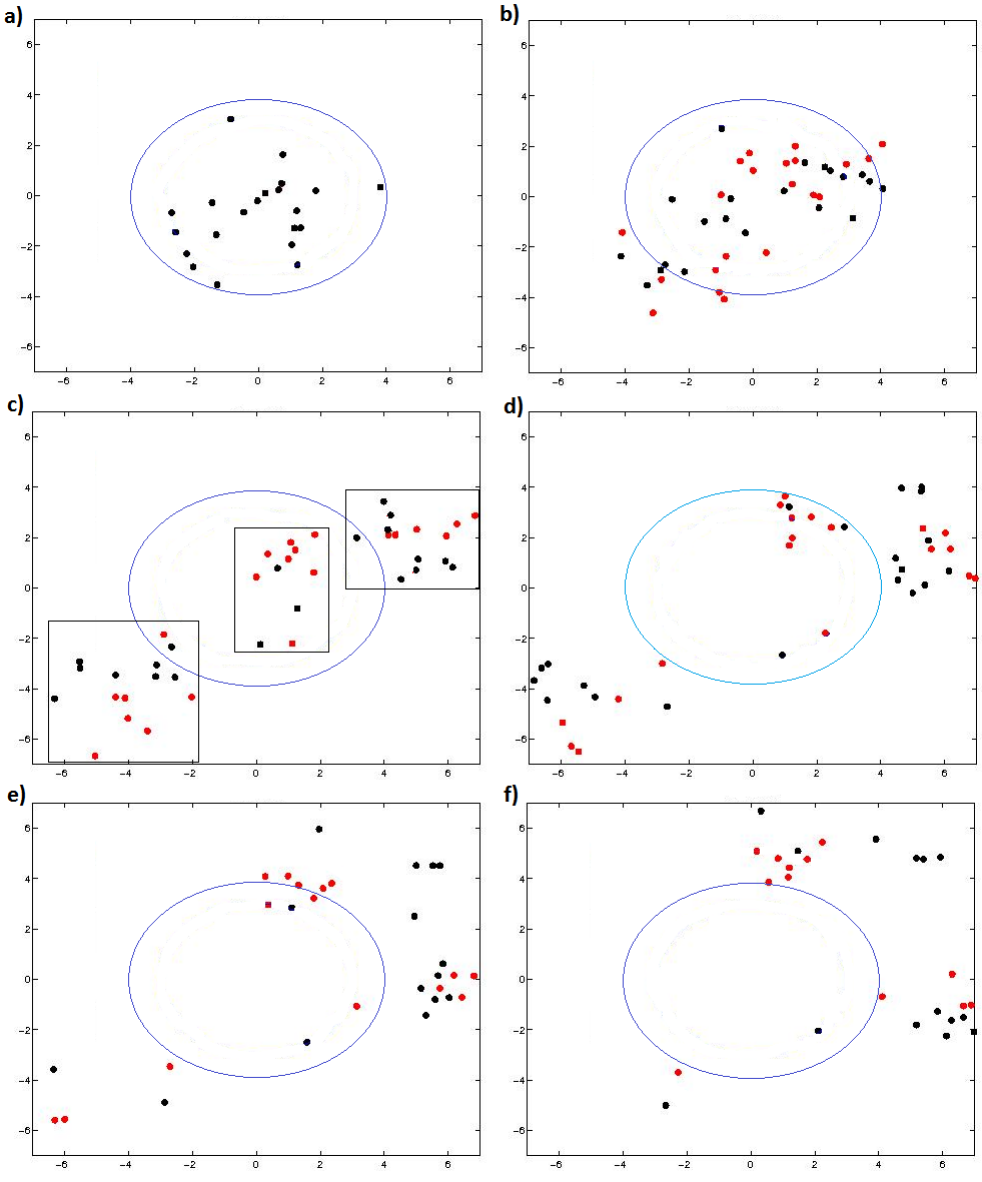

Figure 102: Snapshots of vortex positions during the first six timesteps of the initial value experiment with $N_0 = 21$ and $s = 50$. The black dots are vortices in the upper layer; the red dots ones in the lower layer. The three rectangles in panel c) mark three vortex clusters. The blue contours mark the seeding domain with radius of 4λ.

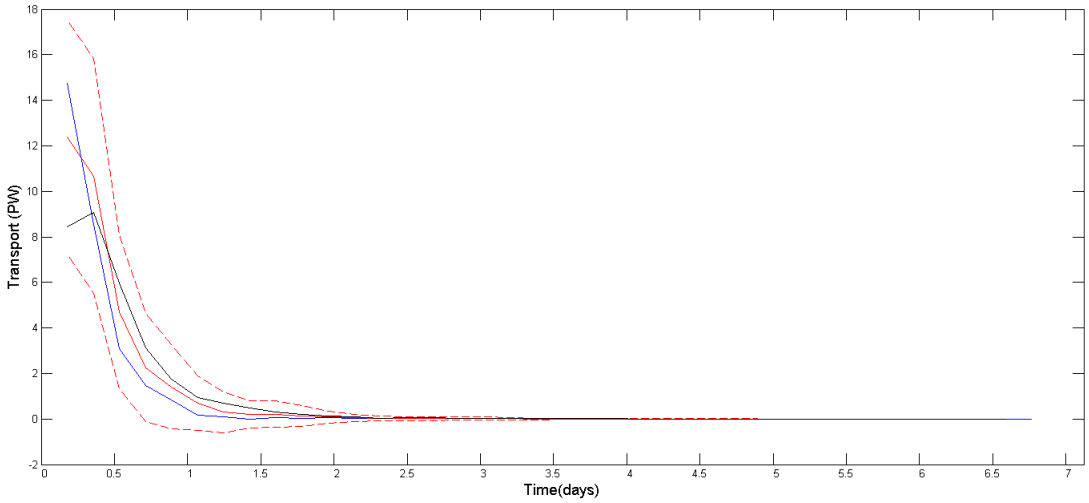

Figure 103: Heat transport time-series, in PW, for the three initial value setups. The transport is computed across the boundary of the seeding domain ($r = 4\lambda$). The continuous curves are the values averaged across all 50 iterations performed for each setup. The blue curve coresponds to the $N_0 = 10$, $s = 100$ case, the red curve to the $N_0 = 21$, $s = 50$ case and the black curve to the $N_0 = 41$, $s = 25$ case. The dashed lines are the 1σ limits for the red curve.

the domain boundary when they experience a repulsive interaction. The lower dashed line suggests that some runs display negative transport values. This occurs because of the circular trajectories characteristic of same-level vortex interactions. When two vortices start to gyrate around a common point and their trajectories straddle the domain boundary, they provide a negative contribution to term A in equation 50 whenever they re-enter the domain. Since the looping trajectories can persist for several timesteps, this behaviour can lead to prolongued negative transport values.

Heat transport PDFs for the initial value experiment are not necessarily indicative of the characteristics of the hetons' transport. Indeed, one could argue that the PDFs of such experiment are skewed, and with a near-zero MLV, almost by design. The large heat transport values resulting from the bulk of the vortices leaving the seeding domain would provide the PDF's long tail, while the near-zero MLV would result from the timetseps when the hetons are beyond the domain's boundary and have already decayed significantly in strength. These could be likened to neutral days in the ERA-Interim data, when the atmosphere does not carry much heat. Figure 104 presents the composite heat transport distribution for all 50 iterations of the $N_0 = 10$, $s = 100$ experiment. The transport is computed across the boundary of the seeding domain. In order to limit the influence

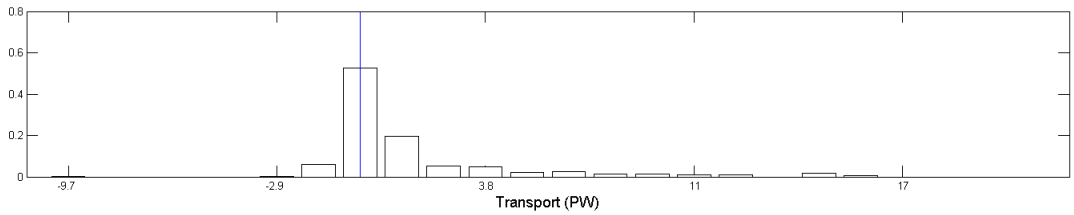

Figure 104: PDF of heat transport, in PW, across the boundary of the seeding domain ($r = 4\lambda$) for the initial value experiment with $N_0 = 10$, $s = 100$. The skewness is 2.4, the MLV is -2.4×10^{12} W. The vertical line shows the bin corresponding to the most likely value.

of the timesteps where all the hetons are very weak and well outside the seeding domain, only the first five timesteps of every iteration are included. Notwithstanding this, the PDF reproduces the predicted features, and looks very similar to the distributions found for the ERA-Interim data (see figure 41 on page 72). The main caveat is that the heat transport extremes are not related to hetonic explosions, but are a direct consequence of the experiment's set-up. The MLV is negative simply because it is defined as the central value of the bin with the highest frequency of events. The bin it corresponds to is the bin containing zero.

4.5 Continuous Seeding Experiment

We now analyse the output of the continuous seeding experiments. Compared to the initial value case, a more limited range of parameters is explored; however, the continuous seeding setup has a more immediate link to the real atmosphere. The simulation, as outlined in Section 4.3, aims to represent atmospheric heat transport into the polar cap. As was the case for the initial value experiment, the number density of hetons across the domain is very uneven. Occasionally, vortices coalesce in dense clusters and hetonic explosions occur. Figure 105, taken from the experiment with $N_r = 2$ and $s = 50$, displays such composite behaviour. Three dense clusters of vortices, rapidly moving away from the domain, are marked by black rectangles. As was seen in the initial seeding experiment, the clusters tend to form well outside the seeding domain (marked by the circular contour), suggesting that heat transport extremes might be best captured at large radii.

The very complex dynamics which govern large numbers of hetons lead to a highly variable heat transport, very different from the monotonic decay witnessed in the initial value setups. As example, Figure 106 shows the time series of heat transport for the two continuous seeding simulations.

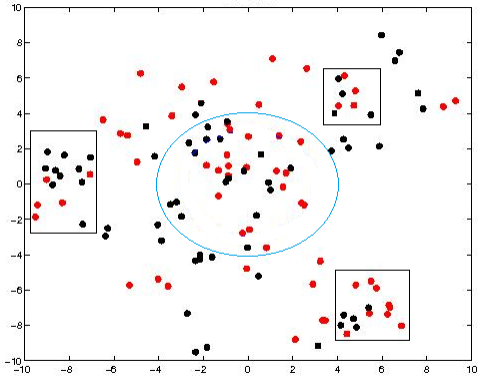

Figure 105: Snapshot of vortex positions during the continuous seeding experiment with $N_r = 2$ and $s = 50$. The black dots are vortices in the upper layer; the red dots ones in the lower layer. The three rectangles mark three vortex clusters rapidly moving away from the bulk of the hetons. The blue contour marks the seeding domain with radius of 4λ.

The transport is computed across the boundary of the seeding domain ($r = 4\lambda$). The blue curve corresponds to the $N_r = 2$ and $s = 50$ case; the red curve corresponds to the $N_r = 4$ and $s = 25$ case. The horizontal dashed lines mark the 95^{th} percentiles of the two curves. Following the definition adopted throughout the present thesis, the values above such threshold are termed extreme events. As expected from the experiment's design, the mean transport of the two simulations is the same (~ 3.5 PW). There is a very small difference between the two due to computational rounding errors. The blue curve has larger peaks, but its baseline value, which would correspond to the MLV of its PDF, appears to be slightly lower than that of the red curve. This is consistent with the more quantised nature of the $N_r = 2$ experiment, where there are fewer, stronger hetons. To understand this better imagine the most extreme quantisation possible: a binary system where a single heton is either inside or outside the domain. When the heton crosses the boundary, it will effect the whole transport of the simulation. The MLV, or mode, of the time-series, on the opposite, will be zero. Now assume that there are 100 hetons within the domain. Unless they cross the boundary all together, the peaks in transport will be smaller than in the binary case. However, the MLV will be larger since, on average, a few hetons will cross the boundary at every timestep.

The differences between the two continuous seeding simulations are reflected in the variability of the two model runs, illustrated by the box and whiskers plots shown in figure 107. Plot a) corresponds to the $N_r = 2$ and $s = 50$ case; plot b) corresponds to the $N_r = 4$ and $s = 25$ case.

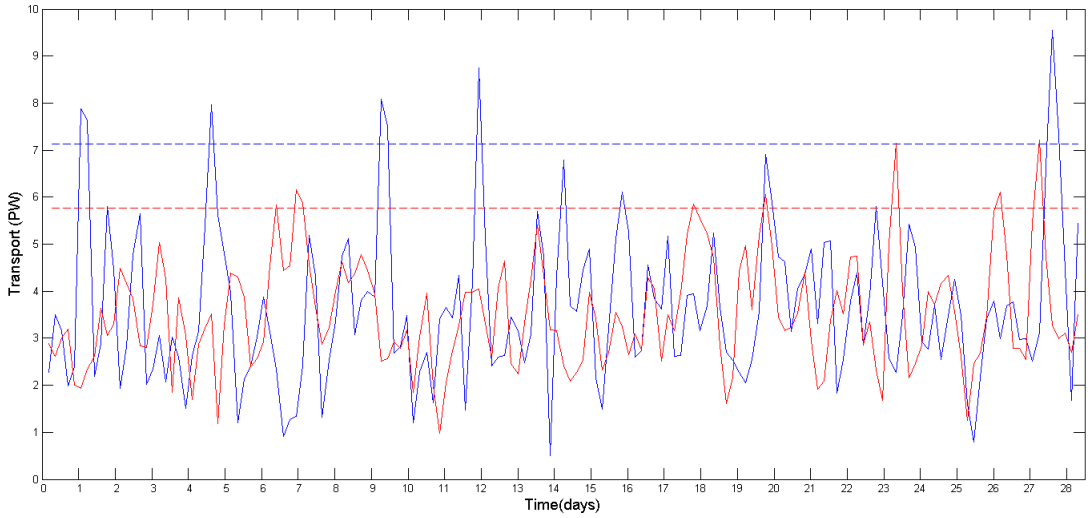

Figure 106: Heat transport time series, in PW, for the two continuous seeding setups. The transport is computed across the boundary of the seeding domain ($r = 4\lambda$). The blue curve corresponds to the $N_r = 2$ and $s = 50$ case; the red curve corresponds to the $N_r = 4$ and $s = 25$ case. The horizontal dashed lines mark the 95^{th} percentiles of the two curves.

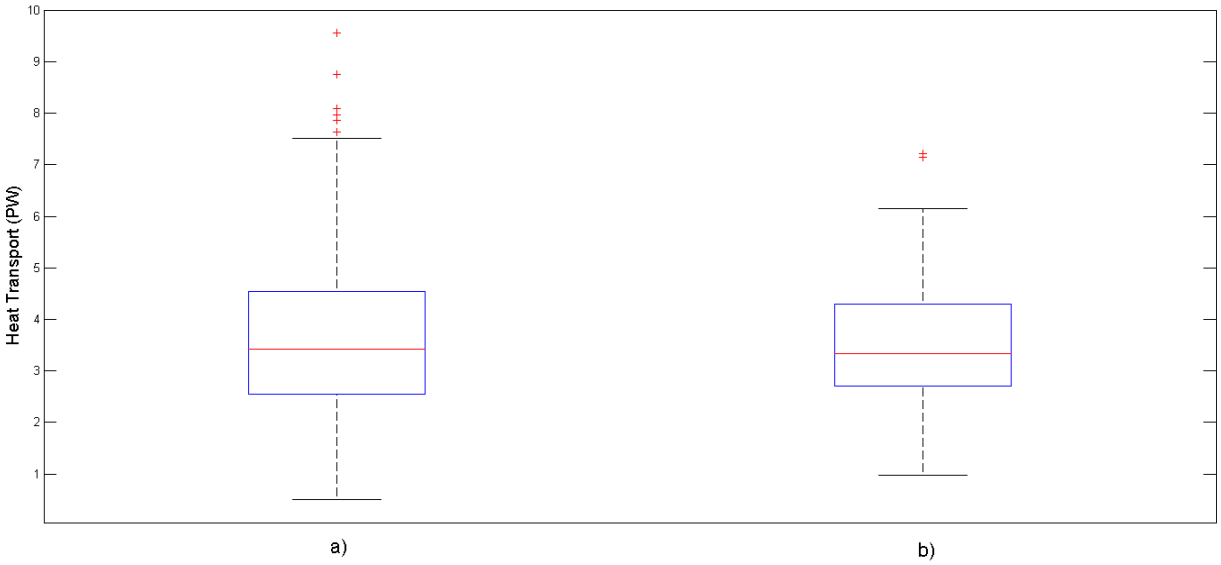

Figure 107: Box and whiskers plots for heat transport, in PW, in the continuous seeding experiments. The transport is computed across the boundary of the seeding domain ($r = 4\lambda$). Plot a) corresponds to the $N_r = 2$ and $s = 50$ case; plot b) corresponds to the $N_r = 4$ and $s = 25$ case. The horizontal red lines mark the 50^{th} percentiles of the two heat transport distributions. The horizontal blue lines mark the 25^{th} and 75^{th} percentiles. The whiskers' length corresponds to 1.5 times the interquartile range. Values beyond these limits are marked by the red crosses.

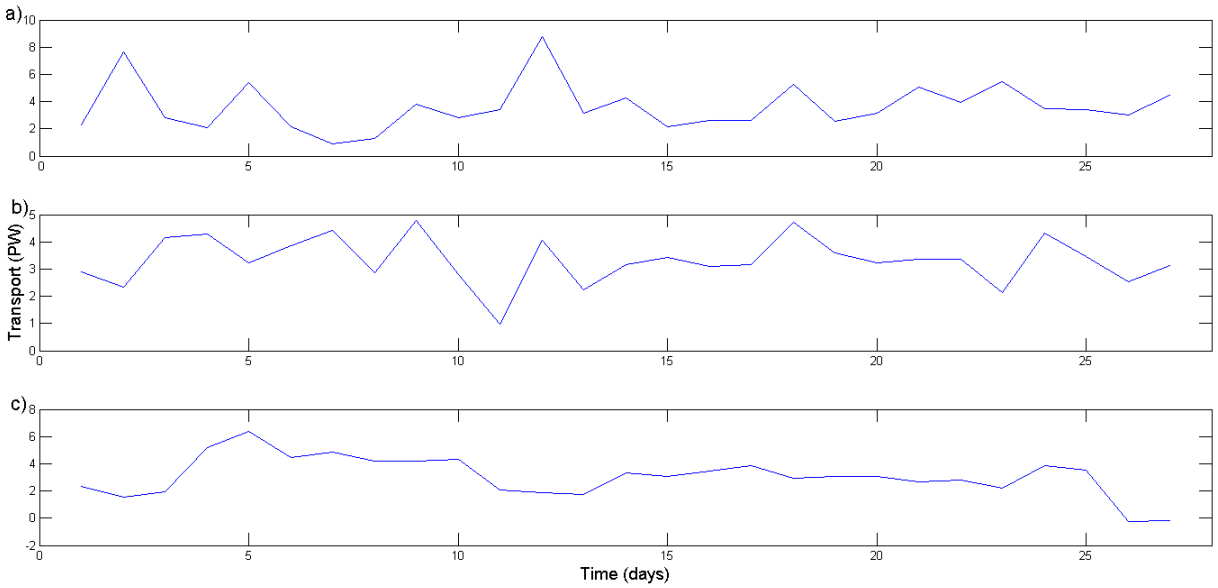

Figure 108: Heat transport time-series, in PW, for a) the continuous seeding experiment with $N_r = 2$ and $s = 50$; b) the continuous seeding experiment with $N_r = 4$ and $s = 25$; c) ERA-Interim data at 60° N, during DJF. The continuous seeding transport is computed across the boundary of the seeding domain ($r = 4\lambda$), and has been downsampled to one point per day to match the sampling frequency of the reanalysis.

The horizontal red lines mark the 50^{th} percentiles (medians) of the two heat transport distributions. The horizontal blue lines mark the 25^{th} and 75^{th} percentiles. The whiskers' lengths corresponds to 1.5 times the interquartile ranges. Values beyond these limits are marked by the red crosses. The simulation with fewer hetons (plot a)) is indeed seen to have a larger variability; the medians are almost identical and the lower mode of a) results in a broader inter-quartile range. As expected, the higher variability of a) corresponds to longer whiskers and more numerous outliers.

The magnitude and variability of the heton transport can be compared to the values seen in the daily ERA-Interim reanalysis. Figure 108 displays both datasets. The reanalysis data covers a randomly selected 28-day interval at 60° N, during DJF. The transport was zonally integrated over the full latitude circle and weighed by pressure level thickness in the vertical to obtain a value in W. The heton model data has been downsampled to one point per day to match the sampling frequency of the reanalysis. Panel a) presents data from the continuous seeding experiment with $N_r = 2$ and $s = 50$; panel b) from the experiment with $N_r = 4$ and $s = 25$; and panel c) from the ERA-Interim reanalysis. The continuous seeding transport is computed across the boundary of the seeding domain

N_r	s	$r=4$	5	6	7	8
2	50	11.1	11.6	13.5	13.4	14.1
4	25	8.95	10.1	10.5	10.8	10.0

Table 10: Percentage contribution of extreme heat transport events to the net meridional atmospheric heat transport in the continuous seeding heton experiment.

($r = 4\lambda$). Unlike for the initial seeding experiment, the heton transport values are comparable to the ones seen in the real atmosphere. As reference, the mean value of plotted the ERA-Interim curve is $\sim 3.05\ PW$. It should be noted that this value refers to the component of heat transport effected by transient motions only, over a random 28 day interval during an NH winter. It is therefore a purely indicative value, but confirms that the heton simulation captures the correct scale. Concerning the variability of the signal, a visual estimate shows that this is highest in the $N_r = 2$ heton experiment, while the $N_r = 4$ one looks very similar to the reanalysis curve. The standard deviations of the three curves are found to be $1.64 \times 10^{15}\ W$, $1.16 \times 10^{15}\ W$ and $1.49 \times 10^{15}\ W$, respectively.

Concerning the extreme events in heat transport in the heton model, there are two questions which need to be answered:

i) Do the extremes account for very high portions of the overall integral of the heat transport distribution, as seen for the ERA-Interim data?

ii) If so, do these extremes correspond to hetonic explosions?

An important note to make is that the heton model should be compared to the zonal integrals in the ERA-Interim data. Indeed, the method adopted here for computing the heton's heat transport does not provide single-point values, but only a value integrated across the whole domain boundary. In the model's idealised setup, this corresponds to integrals around latitude circles.

In order to answer the first question, table 10 displays the percentage contribution of the top 5 percentiles of the distribution to the net heat transport. The different rows and columns display the contributions for different experimental parameters and radii. These values can be compared to those for the latitudinal composites of the ERA-Interim data, shown in table 3 on page 74. For $r = 4$, the percentages found for the heton experiments are significantly lower than the corresponding values seen in the reanalysis. At larger radii, however, the two become comparable, especially for the $N_r = 2$ experiment. The fact that the percentage contributions are systematically lower for the

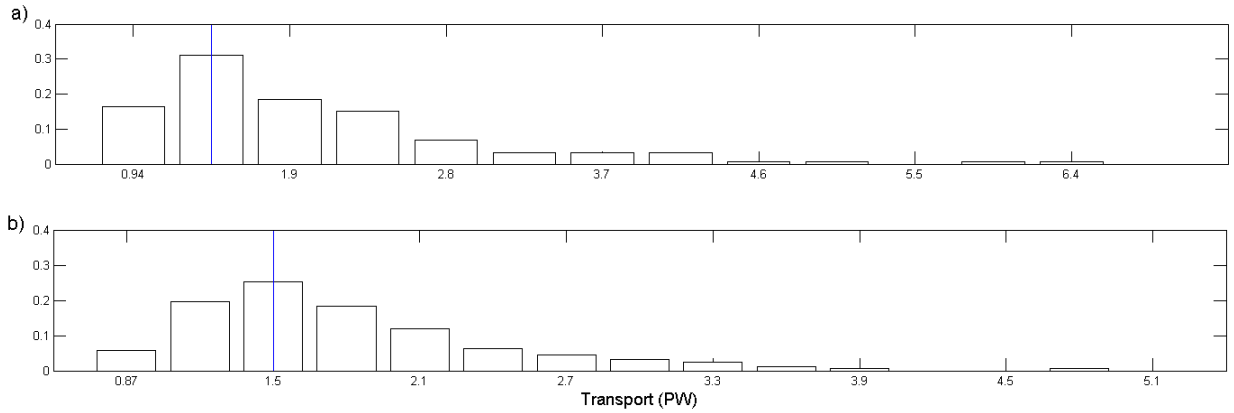

Figure 109: PDF of heat transport, in PW, across the boundary of the domain with $r = 7\lambda$, for the continuous seeding experiments with a) $N_r = 2$, $s = 50$, and b) $N_r = 4$, $s = 25$. The skewnesses are a) 2.2 and b) 1.8 respectively. The corresponding MLVs are a) $1.4 \times 10^{15} W$ and b) 1.5×10^{15} W. The vertical lines show the bins corresponding to the most likely values.

simulation with more hetons is in agreement with the time series shown in figure 106. Additionally, the transport accounted for by the extremes is generally seen to increase with increasing radius. This supports the link between hetonic bursts and transport extremes since, as was noted above, vortex clusters often form beyond the seeding domain's bounds.

Next, it is interesting to verify whether the distributions associated with the heton transport are similar to the ERA-Interim ones. We focus on the $r = 7$ case, where both continuous seeding simulations present large contributions from extremes. The heat transport PDFs are shown in figure 109. Comparing these distributions to those for the ERA-Interim zonal integrals, shown in figures 41 and 42 on page 73, it is immediately clear that there are some similarities with the NH ones, while the SH ones look extremely different. The general shapes of the NH and hetonic distributions are comparable: both display a clear peak, a long positive tail and a very short tail to the left of the maximum. There are, however, some differences in the shape of the positive tail. In the heton case the frequencies decrease gradually, while in the ERA-Interim case the initial bins of the tail have approximately constant frequencies. Concerning the SH distribution, there is no evidence of bimodality in the hetonic case and the two distributions are, visually, extremely different. Last, it should be noted that the skewnesses of the heton distributions are higher than those of the ERA-Interim ones.

Having ascertained that, at least for the $N_r = 2$ experiment, there are extreme heat transport

events comparable to those seen in the real atmosphere, the next question that needs to be answered concerns the role of hetonic explosions in driving these extremes. The heton model provides a visual overview of the heat transport process, and allows for an easy break-down of its different components. It can be immediately verified that the local change in heat content (term A in equation 50 on page 163) drives the extremes. Indeed, the Q_r factor in term B of the same equation is a constant for the constant seeding rate adopted in the current experiment. Therefore, only term A and the α_d factor in term B can affect the transport's variability. Of the two, the former term accounts for the largest variability, and has a correlation with the total transport of $R^2 = 0.95$ for $N_r = 2$ and $R^2 = 0.94$ for $N_r = 4$. The variability of the heat transport is therefore largely determined by the number of vortices crossing the domain boundary; this further strengthens the link between hetonic explosions and heat transport extremes.

To demonstrate conclusively that hetonic explosions correspond to heat transport extremes, it is possible to look at snapshots of the model domain at the timesteps in which the heat transport peaks. Results for the transport at $r = 7\lambda$ are presented for both simulations. For the $N_r = 2$ case, there are five distinct peaks, for a total of 8 data points. The transport time series is shown in figure 110. The horizontal dash-dotted line marks the 95^{th} percentile threshold. The numbers mark the locations of the five extreme peaks. The locations of the individual vortices in correspondence with these peaks are shown in figure 111. Rows a)-e) correspond, respectively, to the maxima labelled 1-5 in figure 110. The panels along each row display consecutive timesteps centred on the maxima. The blue ellipsoids mark vortex clusters crossing the domain boundary. The outer circular contours marks the domain with radius of 7λ. All maxima, and especially the strongest one (row e)), clearly show clusters of hetons crossing the domain boundary. Maximum number 3, in row c), is perhaps the least clear one, with two separate locations which appear to be displaying weak clustering behaviour. For the $N_r = 4$ case, there are again five distinct maxima. Three of these show very clear clustering behaviour, while the other two are comparable to panel c) in figure 111. As illustration, one of the three clustered cases is shown in figure 112. The blue ellipsoids again mark a vortex cluster crossing the domain boundary. The panels display consecutive timesteps centred on the extreme. In this case, since the density of hetons is higher than for the $N_r = 2$ experiment, it is obviously harder to visually identify the heton clusters.

As a final piece of evidence on the role of heton clusters in driving extremes, it is possible to

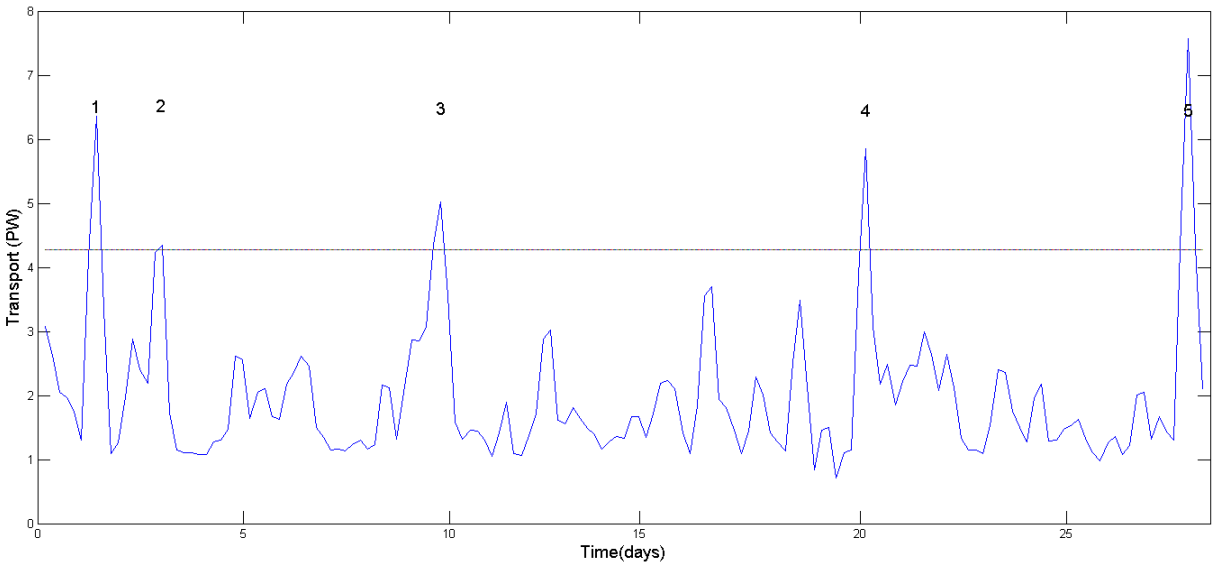

Figure 110: Heat transport time series, in PW, for the $N_r = 2$ continuous seeding setup. The transport is computed across the boundary of the domain with $r = 7\lambda$. The horizontal dash-dotted line marks the 95^{th} percentile of the curve. The numbers mark the locations of the five extreme peaks.

compute the heat transport directly associated with one of these explosions. One can then estimate the percentage of the extreme transport accounted for by the heton cluster crossing the domain boundary. As case study, we select extreme event number 5 for the $N_r = 2$ simulation (see figure 110). Row e), in figure 111, allows estimating the number of vortices crossing the border as part of the hetonic explosion. Focussing on the transition from the first to the second timestep in the plot, such number is seen to be around 15. Each vortex initially corresponds to a set heat anomaly, but decays with time. For this specific case, the mean dimensionless vortex strength s of the cluster is ~ 10.9. Converting this value to W, the hetonic explosion is found to drive a transport of roughly 6.0 PW, while the local maximum of the heat transport extreme is ~ 7.5 PW. The hetonic explosion therefore accounts for approximately 86% of the extreme event's heat transport!

5 Discussion

In their study of the North Pacific storm track, *Swanson and Pierrehumbert* [1997] noted the large sensitivity of heat transport to extreme events. They were able to reproduce the salient features of the resulting PDFs using a model in which temperature anomalies (akin to the H' discussed in

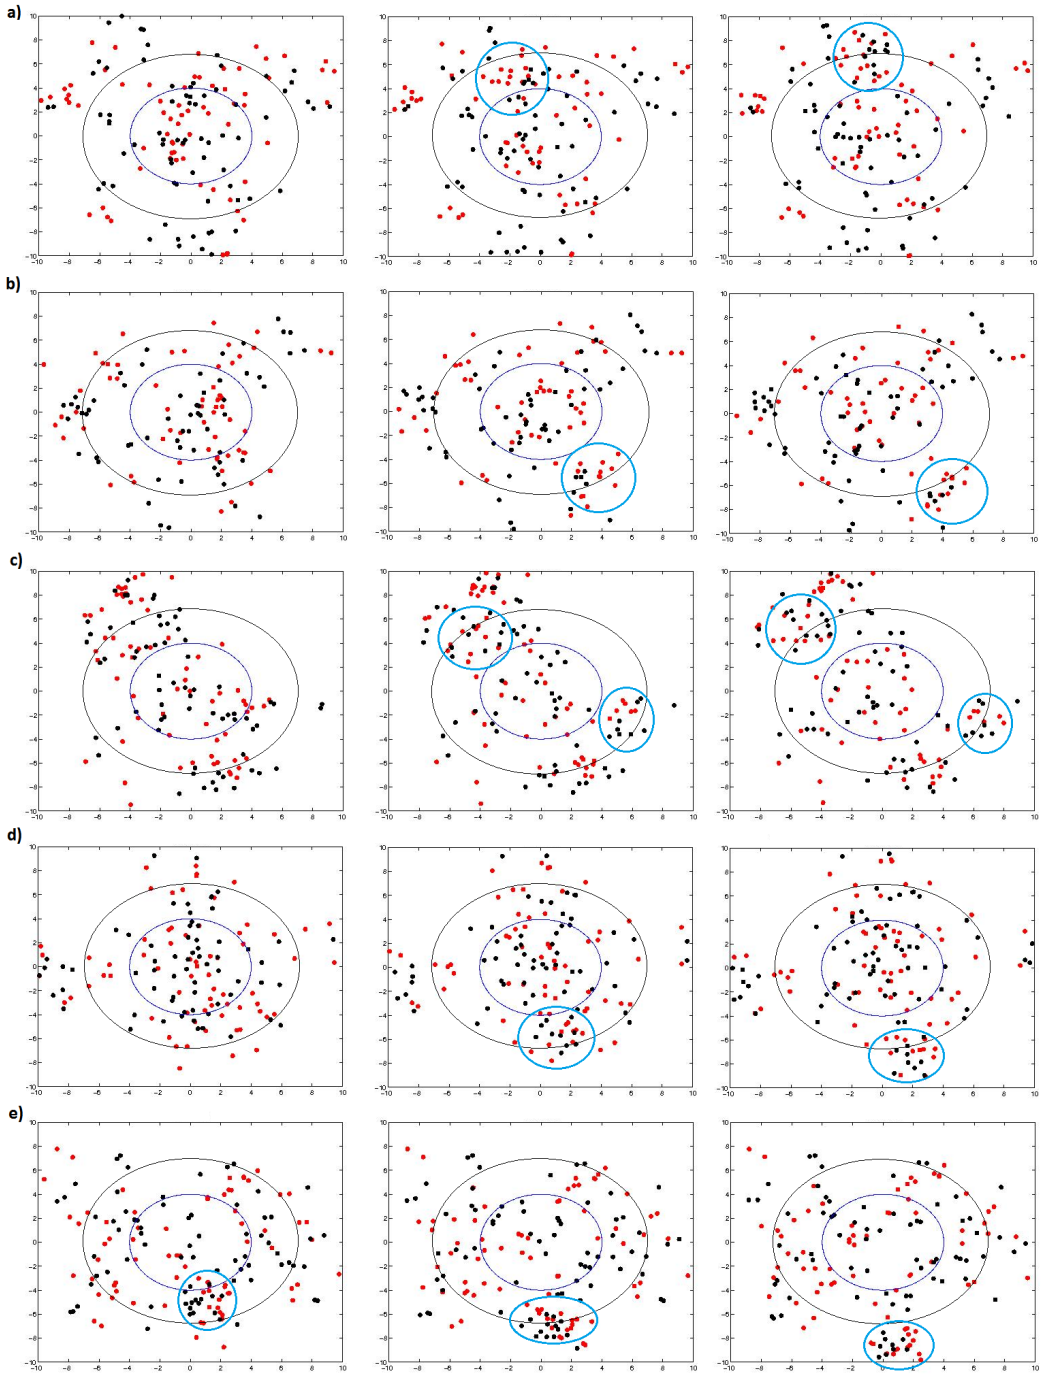

Figure 111: Snapshots of vortex positions during extreme heat transport events for the continuous seeding experiment with $N_r = 2$ and $s = 50$. The panels along each row display consecutive timesteps centred on the extreme. The transport is computed across the boundary of the domain with $r = 7\lambda$. The black dots are vortices in the upper layer; the red dots ones in the lower layer. Rows a)-e) correspond, respectively, to the maxima labelled 1-5 in figure 110. The blue ellipsoids mark vortex clusters crossing the domain boundary. The purple and black contours mark domains with radii of 4λ (seeding domain) and 7λ, respectively.

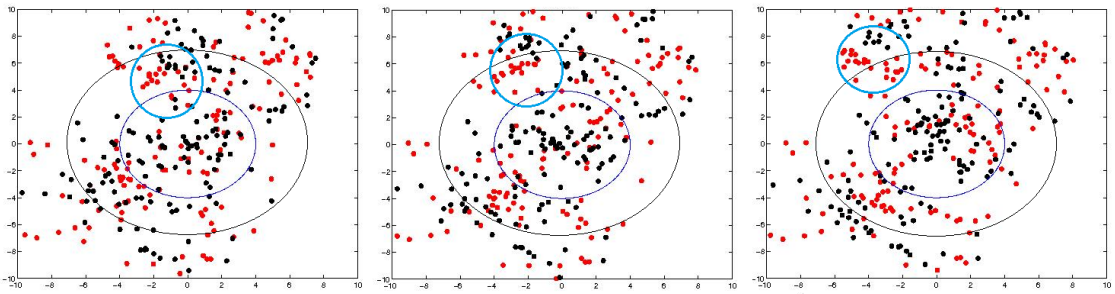

Figure 112: Snapshots of vortex positions during an extreme heat transport event for the continuous seeding experiment with $N_r = 4$ and $s = 25$. The panels display consecutive timesteps centred on the extremes. The transport is computed across the boundary of the domain with $r = 7\lambda$. The black dots are vortices in the upper layer; the red dots ones in the lower layer. The blue ellipsoids mark a vortex cluster crossing the domain boundary. The purple and black contours mark domains with radii of 4λ (seeding domain) and 7λ, respectively.

the present thesis) were created by anomalous meridional advection (v'), and damped through heat exchange with the underlying ocean (see Chapter II, Section 2.4 on page 51 for further details). This picture is somewhat different from the wave model introduced in Section 3, in that *Swanson and Pierrehumbert* treated temperature as a passive scalar, rather than as coupled to the velocity field like in a growing Eady wave. In addition, the authors provided a mechanism limiting the temperature variance (thermal damping), while the mechanisms leading to the equilibrium distribution of baroclinic waves were not addressed in the present chapter. Nevertheless, and more importantly, the simplicity of both views suggests that:

i) One should not be surprised by the large influence of extreme events on the mean poleward heat transport by transient eddies;

ii) One should expect to see this statistical signature in very idealized models of the atmosphere and not solely in nature.

The results from the aquaplanet experiment run in an intermediate complexity AOGCM, described in Chapter III, Section 4.4 on page 77, only strengthen this theory.

The heton model brings these considerations to the extreme, since it provides an atmospheric model essentially based on a single variable – potential vorticity – from which a heat transport can be inferred. Notwithstanding its simplicity, it can still reproduce important aspects of the sporadic

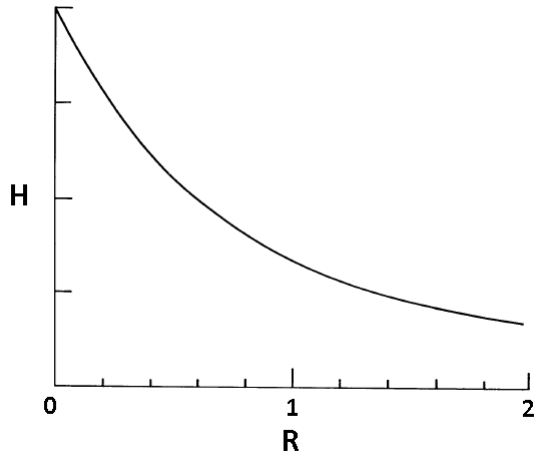

Figure 113: Hetonic heat transport as a function of vortex separation. R is the separation of the two vortices in units of deformation radii. H is the corresponding meridional heat transport [from *Hogg and Stommel*, 1985a].

nature of the heat transport seen in the ERA-Interim data. The percentage contributions of extreme events to the net transport are slightly lower than for the reanalysis but, at least at large radii, of comparable magnitude. As for what concerns the heat transport distributions, these are reasonably similar for the NH, while the two data sets present some major differences for the SH. Most importantly, however, the heton model enables to associate the extreme events to a precise physical mechanism, namely hetonic explosions. These are large clusters of hetons which, effectively, behave like a composite vortex, and propel themselves away from the bulk of the other vortices.

The synchronised displacement of a large number of hetons corresponds to the advection of a large thermal anomaly. This interpretation is the most physically immediate if one views a single vortex as a local temperature anomaly. However, it overlooks the essential role played by the the tilt, or separation, of vortex pairs. A baroclinic pair, such as the heton, will act to shear the fluid between the two layers, and will effect a larger heat transport than the one solely due to the advection of the vortices themselves. Indeed, a single heton will drive a larger heat transport if the separation between the two vortices is small, even though this implies that the pair moves more slowly. Figure 113 illustrates this behaviour. R is the separation of the two vortices in units of deformation radii, and H is the corresponding meridional heat transport. For R approaching zero, the vortex pair will be almost stationary, but the heat transport reaches a maximum.

The overall transport effected by a heton is therefore a combination of the transport due to the displacement of the vortices and the transport due to the heton's tilt. The displacement component vanishes for vanishing vortex separation, because vortices which are perfectly aligned in the vertical are stationary. It also vanishes when R tends to infinity, and the two vortices decouple. It reaches a maximum for intermediate separations, when the vortices are far apart enough to move rapidly but not so far that the distance weakens their interaction. On the opposite, the tilt, or phase, component peaks at vanishing separation. It rapidly decays with increasing distance, and becomes negligible for separations beyond the radius of deformation. This decay is more rapid than the one of the advective component, which therefore dominates for well separated vortices [*A. Czaja, personal communication*].

The heat transport calculation performed in Section 4.2 focusses on the advective component of the transport. It is deemed to be a suitable approximation because the heton clusters analysed in the present chapter, generally display well separated top and bottom layer vortices (*e.g.* figures 111 and 112). The term *well separated* is used here to indicate vortex pairs which have separations comparable to or larger than one deformation radius. Note that this refers heton-type pairings between vortices in different levels, and does not apply to same-level pairs.

The tilt of the hetonic pairs has a direct link to more classical models, such as the Eady wave. A tilted heton will drive tilted anomaly fields, and a westward tilt against the shear of the background flow is common in perturbations leading to a large transient growth. The heton simulation obviously has no constant background shear but, rather, each heton feels the influence of all other vortices (see equations 46 and 47). *Badger and Hoskins* [2001] studied the role of the vertical tilt in Eady-type states, in a PV framework. They found that perturbations with an initial against-shear tilt display an augmented growth when compared to vertically aligned cases. Such growth exceeds the one of the fastest growing Eady-type normal mode. The essential prerequisite for rapid growth, however, is not the vertical tilt but, rather, a varied vertical PV structure. This is certainly present in the heton model, where the two layers have opposite PV signs. Dense clusters of hetons can be approximated as uniform PV patches, implying that the vertical structure of a heton cluster is that of a PV dipole.

The typical size of the clusters involved in hetonic explosions is a few deformation radii (for *e.g.* see figure 111). In the atmosphere, these would be analogous to large thermal anomalies with very high advection speeds, that is large $v'T'$ values. However, because their size is a multiple of the

deformation radius, they cannot be equated to single low pressure systems. In physical terms, the initial state of the experiments described above corresponds to a cyclonic circulation in the upper layer of the model, much like the polar vortex seen in the real atmosphere. This would drive a westerly flow, which then breaks up once the vortices are left free to evolve, cluster and then abandon the seeding domain. This is reminiscent of standing wave patterns developing in the jet stream. This behaviour is illustrated in figure 114, which shows the streamfunction of the top layer of the model for the continuous seeding experiment with $N_r = 2$ and $s = 50$. The six panels display consecutive model timesteps, preceeding and following extreme event number 5 in figure 110 on page 179. Blue shades correspond to low values, and red shades to high values. The continuous black contours are streamlines. In the first panel the streamfunction corrresponds to a cyclonic circulation with a pronounced wavenumber 3 disturbance. One of the troughs begins to grow, extends and finally detaches itself from the main circulation pattern and exits the domain.

The wavenumber of the instability can be compared to the results of a linear stability analysis of a circular baroclinic piecewise constant vortex [*Pedlosky*, 1985]. Following the example of *Legg and Marshall* [1993], and approximating the heton cluster to such a vortex, the rim current should be unstable to perturbations with wavenumber k if:

$$k \leq \frac{r_0}{\lambda}, \tag{57}$$

where r_0 is the radius of the domain. For the radius chosen in the experiments discussed here, we therefore have that $k \leq 4$, in agreement with the wavenumber 3 instability seen in figure 114. Such wavenumber implies a wavelength of approximately 6500 km; long but still consistent with the scale of planetary waves.

The heton setup could therefore be likened to a very pronounced planetary wave pattern, where the composite heton would play the role of a strong trough. Adopting the PV view, the initial heton patch sets up a meridional PV gradient. In *Badger and Hoskins* [2001], an Eady-type initial state was found to support the propagation of Rossby waves once such a gradient was imposed, even in the absence of a β-effect. This provides an alternative interpretation of the instabilities developing on the edge of the hetonic polar vortex.

An important caveat of the model is that the hetonic troughs are seen to propagate at a much faster

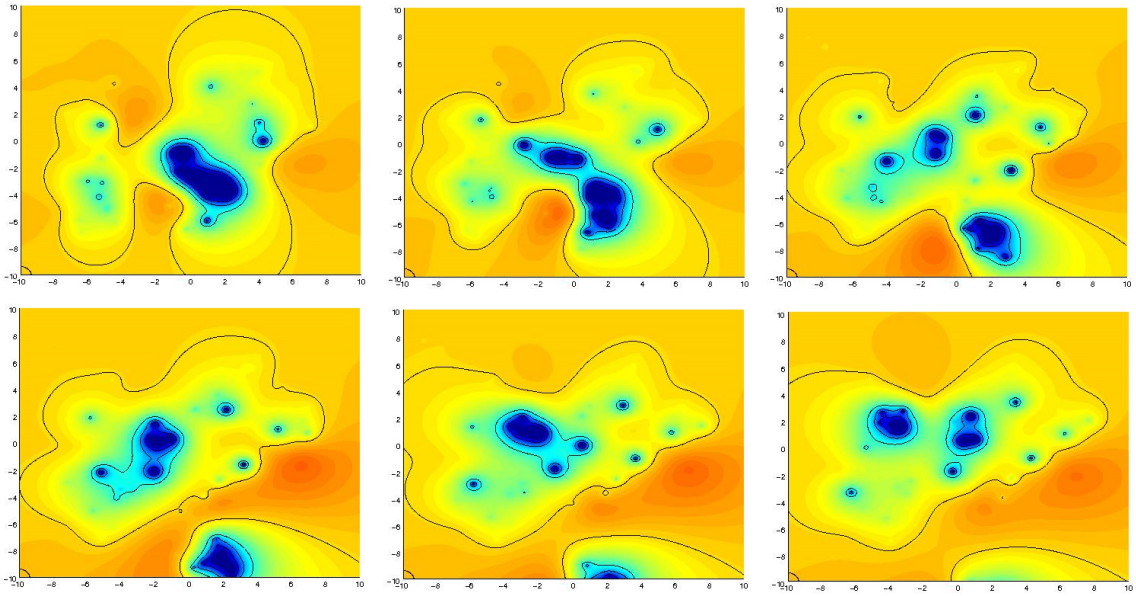

Figure 114: Top layer streamfunction for the continuous seeding experiment with $N_r = 2$ and $s = 50$. Blue shades correspond to low values, and red shades to high values. The continuous black contours are streamlines. The six panels display consecutive model timesteps, preceeding and following extreme event number 5 in figure 110. One of the troughs in the original wavenumber 3 disturbance begins to grow, extends and finally detaches itself from the main circulation pattern and exits the domain.

rate than what would be expected in the real atmosphere. Keeping in mind that a model timestep corresponds to roughly 4.5 hours, and that a cluster moves a distance of order λ over one timestep, a very rough calculation suggests that the velocities attained follow: $v \sim \frac{\lambda}{1.5\times10^4} \approx 50ms^{-1}$. This unrealistic velocity is an inborn characteristic of the model. In fact, the strength of the interactions between the vortices, and hence the resulting velocities, are largely governed by the strength of the vortices themselves. Large clusters behave as composite vortices, where the strength of the individual members is summed, enabling them to attain a much swifter displacement than what would otherwise be seen. A second caveat is that the large $v'T'$ anomalies associated with hetonic explosions correspond, in physical terms, to equatorward advection of cold air. In fact, the heton simulations described in this chapter are all based on cold hetons, namely vortices corresponding to a negative thermal anomaly. Since the circular domains considered are meant to reproduce a polar cap, the hetons abandoning the domain would be flowing equatorwards. The large, positive transport values associated with hetonic explosions would therefore correspond to a combination of negative thermal and velocity anomalies. This is the opposite of the WCB concept discussed in Chapter IV, Section 5 on page 132, which corresponds to combinations of positive thermal and velocity anomalies.

The very simple heton model adopted here is only able to reproduce cold thermal anomalies, since hot hetons would be needed for warm ones. The sign combination driving the heat transport is therefore the direct result of a design feature of the model, and should not be viewed as a shortcoming of the heton dynamics. Indeed, the main aim of the model setup was to reproduce the real atmosphere's heat transport variability, rather than match specific features of the circulation.

The fact both an aquaplanet AOGCM and the heton model present heat transport extremes suggests that they must be a very basic physical property of atmospheric dynamics. The fact that something as simple as the statistical wave model discussed in the present chapter can also reproduce them highlights that they are also a very basic statistical property. It is evident that the high level of abstraction of the latter two models makes them amenable to arbitrary choices. For example, the heton model was seen to display different dynamics depending on vortex strength and heton number density, and these in turn were found to affect the importance of heat transport extremes. Nonetheless, the immediate links that can be drawn between highly idealised models and complex analyses of real-world data are very stimulating.

6 Conclusions

The role of extreme events in setting the mean seasonal value of meridional heat transport has been found, throughout the present thesis, to a be very robust characteristics of the atmosphere. This chapter proposes two highly idealised models which successfully reproduce some salient features of the extreme events in heat transport. The first model is based on a wave view of atmospheric motions, and simulates meridional velocity and MSE anomalies as superpositions of sinusoidal waves. These superpositions result in a heat transport distribution which is very similar to the one seen in the ERA-Interim data. The model succesfully reproduces, among other features, the long positive tail of the distribution, which is the statistical signature of the predominant role played by extreme events.

The second model discussed in this chapter, the heton model, abandons altogether the traditional wave view of atmospheric dynamics, and adopts a quantised view of heat transport. The model's atmosphere is a simple two-layer system seeded with point baroclinic vortices. The complex dynamics driving vortex displacement imply that this toy atmosphere displays large heat transport events, much like the real one. The value of the model lies in its ability to provide a direct analogue for these events, namely hetonic explosions. The explosions are clusters of vortices which act as a composite vortex, and have the potential to rapidly advect large thermal anomalies across the model domain. In physical terms, they can be likened to vigorous outflows of cold air from the polar domain to the lower latitudes.

The two models presented in this chapter are simple examples of the wide range of possible minimal physics simulations which can be used to mimic the variability of atmospheric heat transport. Other authors have proposed different, but equally valid, idealised approaches. For example *Ambaum and Novak* [2013] discuss an oscillator model for storm-track variability, where a temperature gradient gradually builds up and is then rapidly relaxed by transient eddy activity. Unlike the wave model presented here, this mechanism explicitly accounts for the non-linear growth of the instability, albeit in a very simplified fashion. More interestingly, however, the non-linearity results in a very discontinuous heat transport, with well-defined peaks occurring at regular intervals. These could be a direct analogue for the heat transport extremes, much like the hetonic explosions in the heton model. Yet a different idealised model reproducing heat transport extremes is discussed in *Swanson and*

Pierrehumbert [1997]. In the latter study, the variability is driven by a combination of Lagrangian passive advection and thermal interactions (see Chapter II, Section 2.4 on page 51).

All these idealised models, including the two presented in this chapter, succesfully reproduce extreme events in heat transport, notwithstanding their radically different conceptual approaches. This confirms the extremes' fundamental nature, which had already been hypothesised in the previous chapters. The ease in establishing links between basic physical and statistical concepts and the outcome of much more complex analyses performed on reanalysis data suggest that, perhaps, one should not be surprised by the sporadic nature of meridional heat transport in the atmosphere.

Part VI

Conclusions

The present thesis focusses on atmospheric meridional heat transport by transient eddies. The three main research Chapters are parts III, IV and V. The analysis is structured as follows:

i. Chapter III is concerned with the temporal and spatial variability of meridional heat transport, and the role played by the upper percentiles of the transport distribution. It is demonstrated that these percentiles, termed extreme events, are instrumental in setting the magnitude of the net seasonal transport [*Messori and Czaja*, 2013b].

iii. Chapter IV characterises the extreme events by looking at their temporal, spatial and spectral structures. Synoptic analogues are proposed and the role of the different timescales is investigated [*Messori and Czaja*, 2013c].

iv. Chapter V links the extremes to idealised simulations. It is shown that both a statistical model and a non-conventional approach, based on *Hogg and Stommel*'s heton model, succesfully capture the sporadic nature of the transport.

Data from ECMWF's ERA-Interim reanalysis is used throughout the thesis. Specific sections use output from the FORTE model [*Sinha and Smith*, 2002; *Smith and Gregory*, 2009] and *Minobe et al.*'s GCM [*Minobe et al.*, 2008]. The domains considered include the entire extra-tropical regions of the Northern and Southern hemispheres.

1 Summary

Meridional heat transport, effected by both the atmosphere and oceans, is a key component of the climate system. Outside the tropics, the atmosphere accounts for the majority of this transport. The present thesis focuses on a specific component of atmospheric heat transport, namely that driven by transient motions. This is computed as the product of meridional velocity (v') and moist static energy (H') anomalies. The analysis aims to explore the transport's temporal variability, focussing on the role of the highest percentiles of the transport distribution.

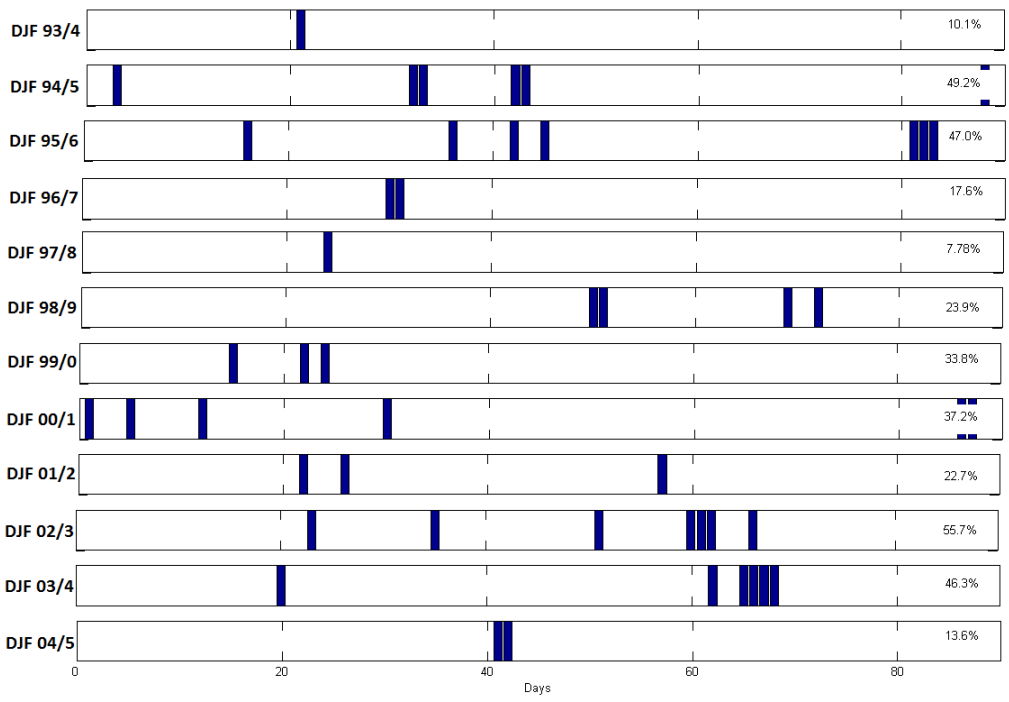

Figure 115: Bar plot of $v'H'$ extreme events at 50° N 0° E. Extreme events are defined as events in the top five percentiles of the $v'H'$ distribution for the 50° N latitude circle. The 12 panels correspond to DJF seasons from DJF 1993/1994 to DJF 2004/2005. Bars correspond to an extreme event occurring on a given day; the abscissa indicate the day of the season. The percentages in each panel indicate the contribution of the selected events to the net seasonal meridional transient-eddy heat transport. See page 94 for further details [from *Messori and Czaja*, 2013b].

Messori and Czaja [2013b] found that, at any given location beyond the tropics, very few large events effectively set the mean seasonal transport value. These events, termed *extremes*, have been defined throughout the present thesis as the top five percentiles of the heat transport PDF. The role of extremes is clearly illustrated in figure 115 which presents, in a binary format, the variability of the heat transport process at a single point (50° N 0° E). Similar plots are obtained for other grid points (not shown). On a given day, a value of unity (vertical bar) is set to the curve if $v'H'$ falls in the top 5 percentiles of the distribution for the 50° N latitude circle, and a value of zero is used otherwise (no vertical bar). By definition, there are only a few extreme events every season, yet these account for a very large portion of the overall poleward heat transport at this location, sometimes exceeding 50% (the numbers for each winter are indicated in the top right corner of each panel)! Extending the analysis to the whole of the extra-tropical regions it is found that, on average, the top five days

per season contribute to over half of the transport at any given point. The atmosphere is therefore very sensitive to a few, temporally and spatially localised features, implying that meridional heat transport by transient motions has a fundamentally sporadic nature.

These extreme events in meridional heat transport are found to be compatible with the traditional picture of growing systems having the minimal phase shift between v' and H', and accounting for the bulk of the heat transport [*Eady*, 1949]. An idealised aquaplanet GCM simulation also reproduces the transport extremes, implying that the sporadic nature of heat transport is not explained by the detailed features of the atmosphere's lower boundary (coastline, ocean fronts, sea ice, orography etc.), stationary waves and mesoscale features. This suggests that it might be a fundamental property of atmospheric waves.

A more detailed analysis, focussing on the local atmospheric dynamics, finds that the synoptic scale motions associated with heat transport extremes depend on the region of the globe considered. In some regions, the extremes are associated with rapidly ascending warm, moist air streams, called warm conveyor belts. This is the case in the Pacific, Atlantic and Southern Ocean storm tracks. At higher latitudes, some events are linked to cold air outflows from the polar basin (*e.g.* in the Bering Strait). Other events are associated with local low-pressure systems, which do not necessarily fit the criteria used to define warm conveyors (*e.g.* over the Nordic Seas and the Siberian Plateau). However, they can still effect significant heat transports, either by poleward advection of warm core lows or by small phase differences between the velocity and MSE fields.

A study of the spectral features of the v', H' and $v'H'$ signals confirms that phase plays an important role in explaining extreme events. However, it also suggests that the baroclinic timescales (here defined as motions with periods between 2.5 and 6 days) can only explain part of the story [*Messori and Czaja*, 2013c]. Indeed, it is found that the extreme events correspond to very precise phase relationships between v' and H' over a broad range of periods (2-32 days), and that low wavenumbers ($k \leq 4$) play an important role in shaping the transport's power spectrum at long timescales. Panels a), b) and c) in figure 116 show the wavenumber spectra for v', H' and $v'H'$ during NH DJF. The spectra are averaged over the 10-32 day range, taken as representative of atmospheric motions with long periods. They are then normalised relative to the spectral peak, so that the contribution of each wave number is expressed as a dimensionless fraction of the contribution from the leading wave number. The dominance of low wavenumbers in panel c) is immediately evident.

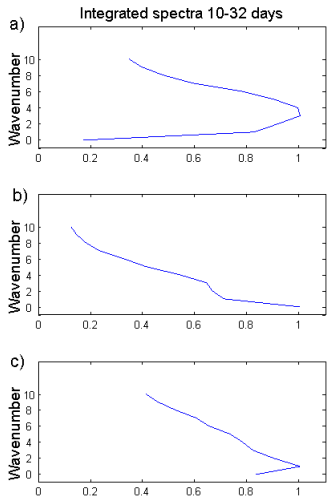

Figure 116: Composite Fourier power spectra of a) meridional velocity anomalies, b) moist static energy anomalies and c) atmospheric heat transport due to transient eddies. The spectra are averaged over the range 10-32 days, and are normalised relative to the spectral peak, so that the contribution of each wave number is expressed as a dimensionless fraction of the contribution from the dominating wave number. The data cover NH DJFs from December 1989 to February 2011. All latitude circles between 30° N and 89° N are taken into account [from *Messori and Czaja*, 2013c].

The fact that planetary-scale perturbations play a significant role in the heat transport's power spectrum might seem at odds with the clear synoptic-scale structures found to be associated with the extremes. The two results can be succesfully reconciled by interpreting the extremes as synoptic features superimposed on larger-scale modes of variability.

Notwithstanding their complex physical analogues in the real atmosphere, extreme events in heat transport can be succesfully reproduced by highly idealised models. A good example is provided by a statistical model, which simulates the v' and H' signals as linear superpositions of sinusoidal waves. Starting from this simple assumption it is possible to capture the salient features of the heat transport PDF. The model and reanalysis distributions are illustrated in panels a) and b) of figure 117. While there is an evident discrepancy in terms of orders of magnitude, the two PDFs share some key elements. Just like the PDF obtained from the reanalysis data, the model distribution displays a highly asymmetric structure with a clear near-zero peak and a very extended positive tail. The latter is the statistical signature of the extreme events, which constitute the upper end of the tail.

A second idealised model of the atmosphere, the heton model, is also discussed. The heton model is a two-layer model, which represents heat transport as a quantised process effected by point vorticity

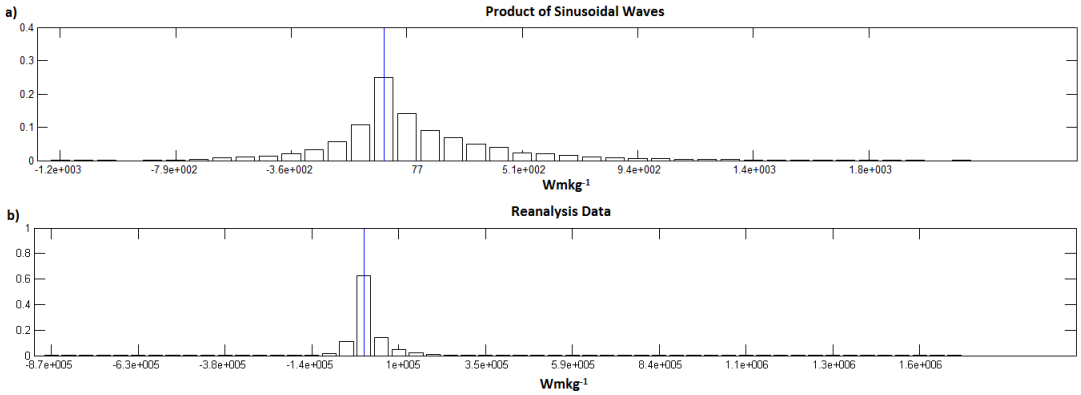

Figure 117: PDFs of atmospheric heat transport due to transient eddies (in $Wmkg^{-1}$) for a) the statistical model and b) reanalysis data. The data cover the 850 mb fields for NH DJFs from December 1993 to February 2005. All latitude circles between 30° N and 89° N are taken into account. The vertical lines show the bins corresponding to the most likely values [from *Messori and Czaja*, 2013b].

anomalies [*Hogg and Stommel*, 1985a]. Its quantised nature naturally lends itself to reproducing the sporadic aspect of the transport. Most importantly, the model provides a precise mechanism generating the heat transport extremes, namely hetonic explosions. These occur when a number of individual point vortices cluster together, acting as a single, stronger vortex. By virtue of its increased strength, such vortex moves very rapidly, separating itself from the bulk of the other vortices and *exploding* outwards. In physical terms, these explosions could be likened to large $v'T'$ values, much like the real-world extremes. The scale of these explosions is typically well in excess of the deformation radius, making them reminiscent of a very pronounced planetary wave pattern. The composite heton would then play the role of a strong trough. An example of hetonic explosion is illustrated in figure 118. The three panels display snapshots of vortex positions during three consecutive model timesteps, centred on the extreme heat transport occurrence. The scale of the plot is in units of deformation radii (λ), where $\lambda \approx 770\ km$. The transport is computed across the boundary of the domain with $r = 7\lambda$, marked by the outer circular contour. The blue ellipsoids mark the vortex cluster crossing the domain boundary.

Two other idealised models relating to heat transport extremes have also been described in the present thesis: *Swanson and Pierrehumbert*'s mixing model (Section 2.4, Chapter II) and *Ambaum and Novak*'s oscillator model (Section 7, Chapter V). The four models seem to contain radically different mechanisms for explaining the nature of meridional heat transport. These are, respectively:

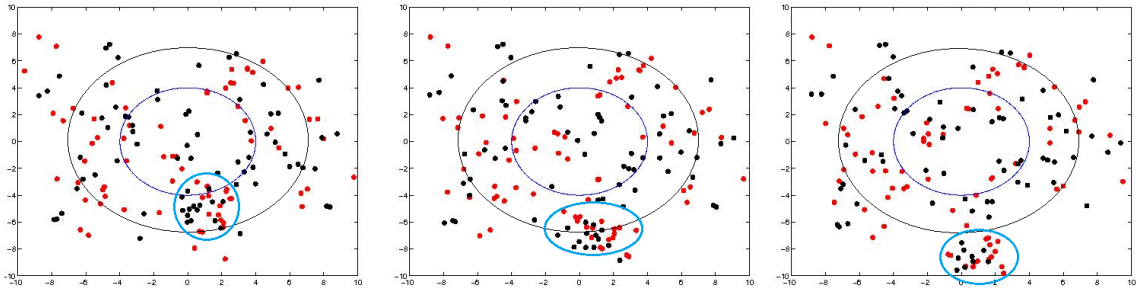

Figure 118: Snapshots of vortex positions during an extreme heat transport event in the heton model. The panels display consecutive timesteps centred on the extreme. The transport is computed across the boundary of the domain with $r = 7\lambda$. The dots are vortices in the model layers. The blue ellipsoids mark a vortex cluster crossing the domain boundary. The outer circular contour marks the domain with radius of 7λ.

i) Lagrangian passive advection and non-conservative thermal processes for the *Swanson and Pierrehumbert* model;

ii) A non-linear oscillator for the *Ambaum and Novak* model;

iii) Phase differences between sinusoidal anomalies in the wave model;

iv) Vortex movements and interactions in the Heton model.

All these approaches are aimed at representing a rapidly growing instability. Diverse as the mathematical and conceptual formulations may be, they are all designed to allow for a rapid growth of the heat transport. In the case of Lagrangian advection, the model is based on the statistics of the observed wind field. In the case of the *Ambaum and Novak* oscillator, the transport bursts occur at regular intervals, determined by the oscillator's frequency. Physically, this frequency corresponds to a build up of baroclinicity which is then released by the eddies. For the wave model, peak transport is attained at points where the different components of the meridional velocity and MSE anomalies are closest to all being in phase. For the heton model, peak transport occurs when different vortices coalesce into a single cluster and then move very rapidly. The *Swanson and Pierrehumbert* model is the one with the closest link to observations in the real atmosphere, and aims to provide a statistically accurate representation of low-level heat transport. The relationship between baroclinicity and eddy heat transport proposed by the oscillator model is seen in real data, but the regular periodicity of the

phenomenon is not realistic. The heton model, and even more so the wave model, have very abstract formulations. Due care should therefore be taken when linking results from these two models with specific features of the real atmosphere. However, all four models successfully reproduce the sporadic nature of heat transport. This is because they highlight the same, fundamental aspect of baroclinic instabilities, namely their rapid growth. The different approaches should not be seen as competing effects but, rather, as different interpretations of a same mechanism observed in the real atmosphere.

The fact that idealised GCM simulations reproduce extreme events in heat transport suggests that the transport's sporadicity has nothing to do with mesoscale signals and characteristics of the atmosphere's lower boundary but, rather, must be an intrinsic property of waves in the atmosphere. The fact that four highly idealised, yet extremely different, models both succesfully capture the transport's sporadic essence confirms this.

2 Future Work

The main consequence of the sensitivity of the heat transport to extreme events is that a very large fraction of the transport occurs in a few discrete bursts, each lasting for only a couple of days. This new perspective has an intriguing application to the climate change debate. Polar amplification describes the phenomenon by which the two polar regions, and especially the Arctic, are warming at a significantly faster pace than the rest of the globe. This has traditionally been ascribed to surface-albedo feedback [e.g. *Hall*, 2004], but a number of studies have also highlighted the important contribution of atmospheric heat transport to the phenomenon [*e.g. Alexeev et al.*, 2005; *Graversen*, 2006; *Lee et al.*, 2011]. This suggests that there could be a strong link between extreme transport events and the climate of the Arctic. The fact that the dominance of extremes in the transient eddy heat transport is present throughout the summer months, which form a consistent part of the polar cap's melt season, also suggests a connection to sea-ice coverage. Even more strikingly, the fraction of the NH poleward heat transport due to extremes is actually higher during the summer months than during the winter ones, as illustrated in tables 1 on page 68, and 2 on page 68. While it is beyond the scope of the present thesis, an analysis relating extreme heat transport events to summertime NH sea-ice extent, along the lines of that performed by *Graversen et al.* [2011, hereafter G11], might uncover interesting connections. G11 suggest that the 2007 sea-ice minimum in the Arctic could be linked

to atmospheric heat transport. In particular, the study finds anomalous atmospheric poleward heat transport and anomalous atmospheric transport convergence during that year in the region interested by the greatest areal loss of sea-ice. G11 then conclude that the additional downwelling long-wave radiation forcing generated by these atmospheric features had a significant role in initiating the melt process. A natural extension of this work would involve looking for anomalies in the frequency of extreme meridional transient eddy heat transport events in areas of the Arctic subject to enhanced sea-ice melting. Analysing in detail the area of high extreme event frequency found in JJA over Siberia and part of the Arctic Ocean would offer the perfect starting point.

In more general terms, the sporadic nature of heat transport in the atmosphere was seen to be driven by a large range of spatial scales, with planetary waves playing an important role. This suggests that specific global regimes, namely quasi-stationary states of the atmosphere [for *e.g.* see *Cheng and Wallace*, 1993], could be associated with an enhancement of extreme transport features. Such a link can be investigated very effectively by using extreme value theory (EVT). EVT is a statistical technique aimed at describing the tails of the distributions of random variables, by characterising them in terms of a small number of key parameters [*Gumbel*, 1958]. These parameters can then be related to a relevant external variable, which in this case could be, for example, geopotential height. Such technique has been succesfully applied to weather extremes over France by *Yiou et al.* [2008], and could provide a robust physical interpretation of the larger scales driving extreme heat transport events.

On the modelling side, the idealised heton setup described in Chapter V could be developed further, by designing a model domain where both the low and high latitudes are represented. One of the main caveats of the model presented here is that the exclusive presence of cold hetons only enables heat transport to be effected as a result of cold anomalies being advected outside the model domain. In the real atmosphere, the majority of the extreme heat transport events were actually found to correspond to poleward advection of warm anomalies. Including the lower latitudes in the model domain would allow both hot and cold hetons to be seeded, overcoming this limitation. Moreover, a greater range of model parameters needs to be tested in order to fully understand their effects on heton dynamics and the resulting heat transport. The results discussed here suggest that the strength and number density of the vortices are extremely important in determining the temporal variability of the heat transport, but the limited computational facilities available did not allow for a systematic

exploration of the model's parameter space.

3 Implications and Learnings

Identifying and describing specific atmospheric features and their driving mechanisms is a never-ending undertaking. The analysis of the features themselves can always be refined, and the links to specific drivers can always be verified with new, more robust techniques. In the specific case of the extreme events described in the present thesis, a more rigorous statistical approach, such as the EVT analysis suggested above, would perhaps have been appropriate. With hindsight, this would have prepared the ground for a more insightful and systematic physical interpretation of the extreme events, which in turn would have made the modelling efforts more effective.

Nonetheless, the analysis described in the present work does result in some important findings and implications, which can be summarised in three main points:

i. Meridional atmospheric heat transport by transient motions is fundamentally sporadic in nature. Very few days every season, termed extremes, effectively set the mean seasonal transport.

ii. These very few days come about because of a complex interplay between synoptic and planetary scale motions. They are further characterised by very precise phase relationships between temporal anomalies in the meridional velocity and moist static energy fields, over a broad range of periods.

iii. The existence of these events is a basic property of the atmosphere. Simple statistical considerations, idealised experiments with realistic atmospheric dynamics, and minimal physics two-layer setups all successfully reproduce the main features of the heat transport distribution.

These points suggest that the present thesis has explored and characterised a largely unknown, yet fundamental, aspect of the variability of heat transport in our atmosphere. Ultimately, very few days every season could hold the key to explaining some large scale features of our climate system.

Appendix

1 The Eady Model

Here, we present the mathematical framework of the Eady Model, which complements the discussion presented in Chapter II, Section 2.1 on page 46. The intention is not to provide a step-by-step derivation but, rather, to highlight the key formulae describing Eady's idealised heat-carrying wave. The analysis is mostly based on *Gill* [1982].

The Eady model will be treated in log-pressure co-ordinates, namely a co-ordinate system which uses the logarithm of pressure as third dimension. The vertical co-ordinate, z_{lp}, is defined as:

$$z_{lp} = -H_s ln\left(\frac{p}{p_r}\right), \tag{58}$$

where ln is the natural logarithm, p_r is a reference pressure level and H_s is a reference scale height[9]. The initial setup considers an incompressible fluid on an f-plane, with a uniform temperature gradient in y and z_{lp}. The presence of this gradient implies that the sytem has APE. Using a relation termed *thermal wind*, it is possible to infer the vertical gradient of the zonal velocity from the horizontal temperature gradient:

$$f\frac{du}{dz_{lp}} = -\alpha g\frac{\partial T}{\partial y}, \tag{59}$$

where α is the thermal expansion coefficient for dry air. Small perturbations to this system satisfy simple wave-like equations of the form:

$$\Phi' = \Phi_0 sin(ly) sin(k(x-ct)) e^{-z_{lp}/H_R} \tag{60}$$

Here Φ is geopotential, namely geopotential height multiplied by gravitational acceleration, the prime denotes anomalies, and the subscript $_0$ denotes a constant amplitude. k and l are the horizontal wave

[9] The term scale height indicates the vertical length scale over which a given quantity decreases by a factor of e.

numbers, and c is the zonal phase speed of the disturbance. H_r is a vertical scale for the decay of the perturbation. The potential temperature perturbation corresponding to equation 60 is given by:

$$\theta' = \frac{1}{\alpha g}\frac{\partial \Phi'}{\partial z_{lp}} = \frac{-\Phi'}{H_R} \tag{61}$$

Negative temperature anomalies are therefore associated with positive geopotential ones, and *vice-versa*.

Next, a surface boundary condition is imposed, namely that at $z_{lp} = 0$, $w = 0$. Applying this to solutions of the form shown in equation 60, one can then obtain an expression for the meridional heat transport in terms of the geopotential anomaly. The mathematical details of the intermediate steps will not be presented here, but the final result is:

$$\overline{v\theta} = \frac{1}{f}\overline{\frac{\partial \Phi'}{\partial x}\theta'} \propto \overline{\Phi'\frac{\partial \Phi'}{\partial x}} = 0, \tag{62}$$

where equations 60 and 61 have been used to obtain the final result. The overbar denotes the zonal mean over a full wavelength. The net heat transport of the perturbations is therefore zero.

Eady, in his seminal 1949 paper, considered this same system, but applied a rigid upper boundary to the setup. If one takes a frame of reference where the vertical mid point of the system is at $z_{lp} = 0$, the two boundaries at $z_{lp} = \pm H$ and the horizontal flow is zero at this mid point, the velocity can then be expressed as:

$$\vec{u} = z_{lp}\frac{d\vec{u}}{dz}, \tag{63}$$

where $\vec{u}$ is the horizontal velocity vector. Equation 60 can now be re-written in terms of the hyperbolic functions as:

$$\Phi' = A(x,y,t)sinh\left(\frac{z_{lp}}{H_R}\right) + B(x,y,t)cosh\left(\frac{z_{lp}}{H_R}\right), \tag{64}$$

where the boundary condition of no vertical velocity applies to both $z_{lp} = \pm H$. Amplitudes of the type:

$$A, B \propto cos(ly)e^{ik(x-ct)} \tag{65}$$

satisfy this new wave solution. There will be a range of possible wave numbers yielding valid solutions to this problem, but only one will correspond to the fastest growing instability. This will soon dominate because of the exponential term in the amplitudes. The growth rate can be expressed as $\sigma = kc_i$, where c_i is the imaginary component of the zonal phase speed. Its maximum value is given by:

$$\sigma_{max} = 0.3098 \frac{f}{(k^2 + l^2)^{1/2} H_R} \frac{d\vec{u}}{dz_{lp}} \tag{66}$$

The numerical pre-factor results from maximising a hyperbolic function of scale height. The corresponding solution for the fastest growing mode is:

$$\Phi' = \left[cos(kx) \frac{sinh(z_{lp}/H_R)}{sinh(H/H_R)} + sin(kx) \frac{cosh(z_{lp}/H_R)}{cosh(H/H_R)} \right] e^{\sigma t}, \tag{67}$$

which can be also expressed as a potential temperature perturbation:

$$\theta' = \frac{1}{\alpha g H_R} \left[cos(kx) \frac{cosh(z_{lp}/H_R)}{sinh(H/H_R)} + sin(kx) \frac{sinh(z_{lp}/H_R)}{cosh(H/H_R)} \right] e^{\sigma t}, \tag{68}$$

and results in a heat transport given by:

$$\overline{v\theta} = \frac{1}{\alpha f g H_R} \overline{A \frac{\partial B}{\partial x}} \tag{69}$$

This solution is positive and independent of height. It can be shown that it remains positive for all growing waves.

2 Wavelet Transforms

2.1 Mathematical Formulation, Cross-Wavelet Transforms and Coherence-Phase Spectra

The discussion of wavelets and continuous wavelet transforms in this section is largely based on *Torrence and Compo* [1998]. The one of cross-wavelet and coherence-phase spectra is based on *Grinsted et al.* [2004].

Wavelets transforms have been discussed in Chapter IV of the present thesis. They are a technique whereby a two dimensional spectral picture (here period-time) is obtained from a one-dimensional time series (here v', H' and $v'H'$) [*Daubechies*, 1990]. The key component of the transform is the wavelet itself, ψ. There are a number of such wavelets available; in the present thesis, we have adopted the Morlet wavelet, defined by:

$$\psi_M(\eta) = \pi^{-1/4} e^{i\omega_0 \eta} e^{-\eta^2/2} \tag{70}$$

This is the product of a Gaussian and a plane wave, where ω_0 is a non-dimensional frequency and η is a non-dimensional time parameter. A different type of wavelet, the Paul wavelet, was used to test the robustness of the analysis performed. These two types of wavelets are depicted in figure 119.

Wavelets can be applied to a data sequence, x_n, via the continuous wavelet transform. This is defined as:

$$W_n(s) = \sum_{n'=0}^{N-1} x_{n'} \psi^* \left[\frac{(n'-n)\delta t}{s} \right], \tag{71}$$

where the asterisk indicates a complex conjugate, s is a scale factor and n is a time index. W_n is essentially the convolution of the original time series with the wavelet, mediated by a scaling and a shifting factor. It is by varying these factors that one can obtain a two-dimensional spectral view of the original signal. The convolution process should be repeated N times for each scale, N being the length of the time series analysed. The wavelet power is then the modulus squared of the above expression.

A number of arbitrary choices need to be made when performing a wavelet transform. The first is the choice of wavelet type, and is based on an inspection of the original signal. The wavelet is

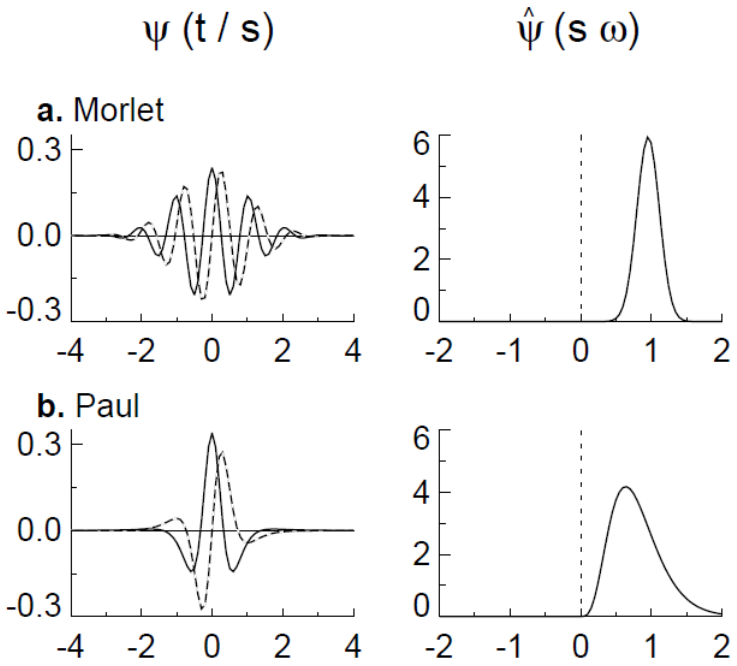

Figure 119: Real (solid lines) and imaginary (dashed lines) parts of the Morlet and Paul wavelets. The panels on the left hand side depict the wavelets in the time domain. Those on the right hand side in the frequency domain [from *Torrence and Compo*, 1998].

chosen so as to match the general "shape" and features of the signal. Next, a set of scales s needs to be selected. These are commonly expressed as an exponential series:

$$s = s_0 2^{j\delta s},\ j = 0, 1, 2... J, \tag{72}$$

where s_0 is the smallest resolvable scale and J corresponds to the largest scale. The smallest scale should clearly be chosen in relation to the sampling frequency of the original time series. Here, following *Torrence and Compo* [1998], it is chosen such that the corresponding Fourier frequency is $2dt$, where dt is the sampling period. The algorithm relating Fourier frequency to scale will not be discussed here; a full description can be found in *Meyers et al.* [1993]. Conveniently, for the Morlet wavelet the Fourier frequency is almost equal to the wavelet scale. A suitable incremental step for scale must also be selected. Here, again following *Torrence and Compo* [1998], a step $\delta s = 0.25$ was chosen. The scale used in the present analysis is therefore given by:

$$s = 2\delta t\, 2^{0.25j},\ j = 0, 1, 2... J \tag{73}$$

In interpreting the wavelet spectrum, one must be careful of edge effects, since the v', H' and $v'H'$ time series analysed here are not periodic. In fact, similarly to more traditional Fourier transforms, wavelets do not yield significant results if applied to time series which are short compared to the time scales considered. The cone of influence, namely the limit beyond which edge effects become important, is defined at each scale as the point where the wavelet power for a discontinuity at the edge of the time series is attenuated by a factor of $1/e^2$.

If two different signals are analysed (such as v' and H'), continuous wavelet transforms can be used to construct cross-wavelet and coherence-phase spectra. The first show the times and periods at which the two signals have the highest common power. The second show the coherence and relative phase of the two signals. The cross-wavelet transform of two time series, x and y, is defined as:

$$W^{xy} = W^x W^{y*}, \tag{74}$$

where W^x and W^y are the individual continuous wavelet transforms, as defined in equation 71.

The complex argument $arg(W^{xy})$ can be interpreted as the local relative phase between x and y in time-frequency space. It is further possible to compute a measure of wavelet coherence, always in the time-frequency domain. For the mathematical details of this procedure the reader is referred to *Torrence and Webster* [1999]. Combining the phase and coherence information, one can produce plots such as those shown in Chapter IV, Section 4.4 on page 128, termed coherence-phase spectra.

2.2 Statistical Tests

The analysis in Chapter IV was largely based on composites of wavelet spectra. It is important to verify that these composite spectra provide a good representation of the typical individual spectrum, particularly at the location of the extreme events. As test, the mean of the logarithm of wavelet power is computed. If the mean in logarithmic space (assuming power follows a log-normal distribution) increases in coincidence with the extreme events, then this suggests that the increase in power is due to a shift in the distribution, rather than a change in its shape. Next, the mean's standard error is computed, in order to verify that the increase is statistically significant. Last, the power's standard deviation, always in logarithmic space, is calculated. An increase in standard deviation in correspondence with the spectral peak would indicate that few extremes are increasing the mean power. If the standard deviation is constant, but the mean increases, this implies that the whole distribution shifts to higher values, and that the bulk of the events contributes to the spectral peak. This procedure needs to be repeated for each scale over which the wavelet transform has been performed. Here, we show as an example the plots for the $v'H'$ signal during NH DJF (see figure 120). For clarity only three scales are presented, all beyond the cone of influence for the times included in the plot, and selected such that they correspond to periods at which the extremes show high power. Panel a) shows that there is a clear increase in the mean in correspondence with the extreme events. Such increase is large compared to the standard error (dashed red lines). The standard deviations (panel b)) show no increase but, on the opposite, decrease at the extreme event location. The two panels therefore indicate that the wavelet composites gives a good representation of the typical extreme event. Similar plots for all scales, hemispheres and seasons were analysed in order to verify the robustness of the results presented here.

Some authors, such as *Torrence and Compo* [1998], suggest techniques to compute significance levels for the wavelet spectra. These techniques are by no means universally agreed upon, and there

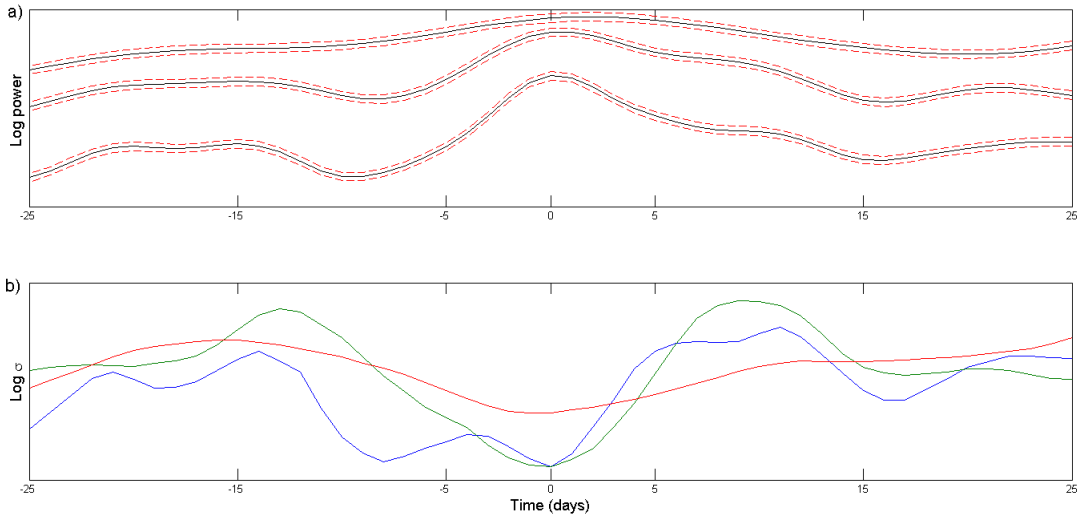

Figure 120: Panel a) illustrates the mean of the logarithm of wavelet power for $v'H'$. The three scales presented are beyond the cone of influence over the time period considered. The dashed red lines indicate the standard errors. Panel b) illustrates the standard deviation of the logarithm of wavelet power for the same three scales. Both panels are centred on the location of extreme events (0 days). The data cover the 850 *mb* fields for NH DJFs from December 1989 to February 2011. All latitude circles between 30° N and 89° N are taken into account.

are different view points concerning whether stationary statistical tests are suitable to evaluate wavelet confidence intervals [*e.g. Lau and Weng*, 1995; *Torrence and Compo*, 1998]. As stated in *Lau and Weng*: "...one of the shortcomings of the wavelet transform analysis is the lack of a proper statistical significance test... Almost all traditional significance tests are derived from the assumption of identical repeated cycles for nonstationary processes and are therefore inappropriate for wavelet transfroms.... Monte Carlo methods are also unsatisfactory in the context." Furthermore, the significance levels presented in *Torrence and Compo* [1998] are suitable for an underlying Gaussian distribution, while the present analysis is based on a distribution that is distinctly non-Gaussian. Even though far from ideal, the only possible test is therefore a Monte Carlo type approach. A loose Monte Carlo test is indeed discussed in Chapter IV, when the composite spectra for random days are compared to those for the extreme days.

3 Filters

The filter used as part of the wavelet analysis was briefly described in Chapter IV, Section 4.2 on page 121. This is a 21-point high pass FIR filter, with a half-power cut-off at 8 days. FIR filters are filters whose response to any input of finite length is itself finite, meaning that it decays to zero at $t < T$, where T is finite. The filter used here was designed using a *MatLab* algorithm based on *McClellan et al.* [1973]. Figure 121 shows the filter's amplitude response as a function of the fractional Nyquist frequency (green curve). The Nyquist frequency is simply half the sampling

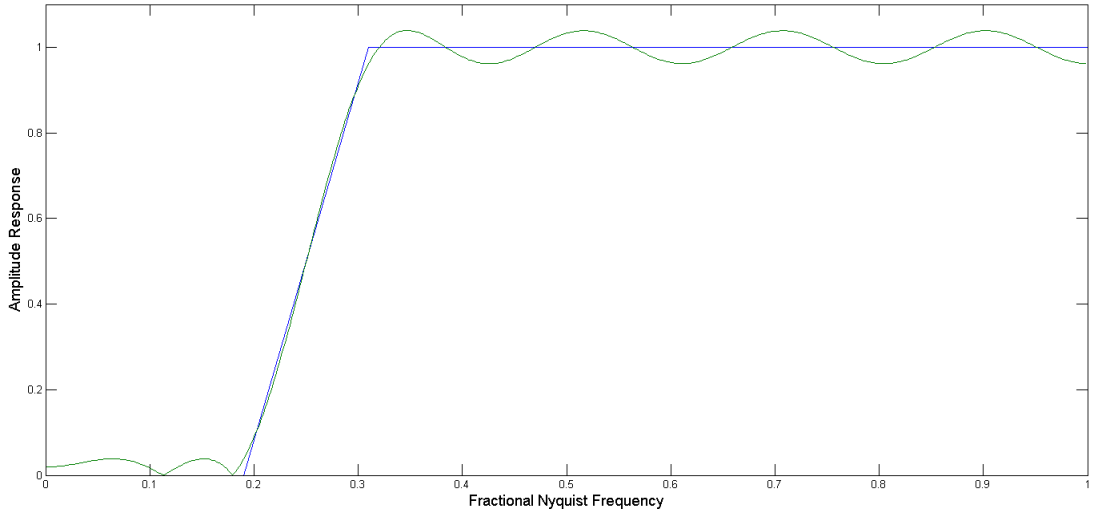

Figure 121: The amplitude response of the filter used in Chapter IV, as a function of fractional Nyquist frequency (green curve). This is a 21-point high pass FIR filter, with a half-power cut-off at 8 days. The response of the ideal filter is shown in blue.

frequency; here it is normalised to the interval $[0, 1]$. The 8-day cut-off corresponds to 0.25 in terms of this dimensionless scale. The 21 filter coefficients are:

[-0,0276636710903229; -0,0126869494554124; 0,00255569651900081; 0,0267124072125839; 0,0456647784637818; 0,0415326955178977; 0,00236913784968678; -0,0691491262509086; -0,153644545402426; -0,221977864735121; 0,751809257520219; -0,221977864735121; -0,153644545402426; -0,0691491262509086; 0,00236913784968678; 0,0415326955178977; 0,0456647784637818; 0,0267124072125839; 0,00255569651900081; -0,0126869494554124; -0,0276636710903229].

The effect of the filter on the v' and H' signals, analysed in the present thesis, is illustrated

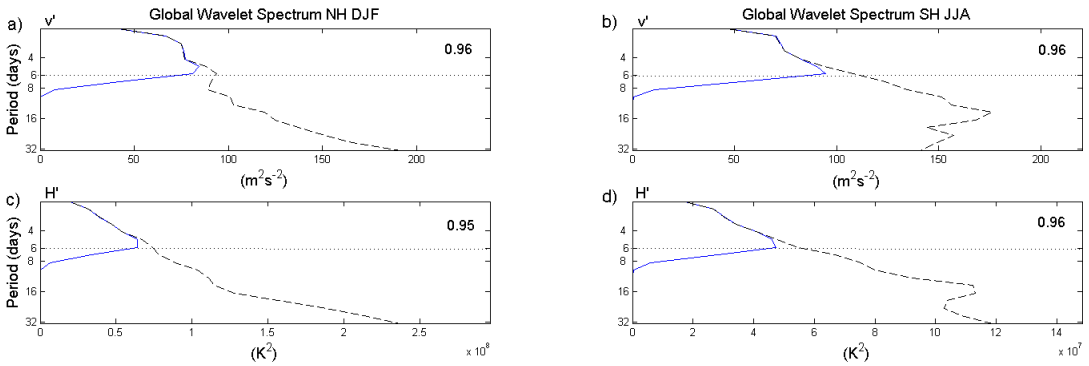

Figure 122: Global (time-averaged) wavelet spectra for v' (in m^2s^{-2}) and H' (in K^2) during NH DJFs (panels a and c respectively) and SH JJAs (panels b and d). The time span considered covers the five days centred on the extremes. The continuous lines are the spectra of the filtered v' and H'. The dashed lines are the spectra of the unfiltered signals. The dotted lines mark the limit of baroclinic timescales, defined here as motions with periods of 2-6 days. The numbers in the top left hand corner of the panels are the ratios between the period-integrated powers of the filtered and unfiltered signals in the baroclinic range only. The data cover NH DJFs from December 1989 to February 2011 and SH JJAs from June 1989 to August 2010 [from *Messori and Czaja*, 2013c].

in figure 122. The figure displays the time averaged wavelet spectra for the 5 days centred on the extreme event, for both the filtered (continuous line) and the unfiltered (dashed line) signals. Panels a and c display the spectra for v' and H' respectively, during NH DJF. Panels b and d display the corresponding spectra for SH JJA. The plots for NH JJA and SH DJF (not shown) present a very similar picture. The dotted lines mark the limit of baroclinic timescales, defined here as motions with periods of 2-6 days. The numbers in the top right hand corner of the panels are the ratios between the period-integrated powers of the filtered and unfiltered signals in the 2–6 day range only. The key feature of the figure is the very small loss in power of the v' and H' signals at periods below 6 days. This, combined with the ratios shown in the panels, demonstrates that the filter's design has a reduced impact on the v' and H' signals at baroclinic timescales.

References

International Meteorological Vocabulary, Secretariat of the World Meteorological Organization, 1992.

Ahrens, C., *Meteorology today: an introduction to weather, climate, and the environment*, Cengage Learning, 2007.

Alexeev, V., P. L. Langen, and J. Bates, Polar amplification of surface warming on an aquaplanet in "ghost forcing" experiments without sea ice feedbacks, *Clim. Dynam.*, *24*, 655 – 666, 2005.

Ambaum, M., and L. Novak, A nonlinear oscillator describing storm track variability, *Quarterly Journal of the Royal Meteorological Society, Submitted*, 2013.

Badger, J., and B. Hoskins, Simple initial value problems and mechanisms for baroclinic growth, *Journal of the atmospheric sciences*, *58*(1), 38–49, 2001.

Bjerknes, J., Atlantic air-sea interaction, *Advances in geophysics*, *10*(1), 82, 1964.

Bjerknes, V., Synoptical representation of atmospheric motions, *Quarterly Journal of the Royal Meteorological Society*, *36*(155), 267–286, 1910.

Blackmon, M., A climatological spectral study of the 500 mb geopotential height of the Northern Hemisphere, *Journal of the Atmospheric Sciences*, *33*(8), 1607–1623, 1976.

Blackmon, M., and N.-C. Lau, Regional characteristics of the Northern Hemisphere wintertime circulation- A comparison of the simulation of a GFDL general circulation model with observations, *Journal of the Atmospheric Sciences*, *37*(3), 497–514, 1980.

Blackmon, M., J. Wallace, N.-C. Lau, and S. Mullen, An observational study of the Northern Hemisphere wintertime circulation, *Journal of the Atmospheric Sciences*, *34*(7), 1040–1053, 1977.

Blackmon, M., Y. Lee, and J. Wallace, Horizontal structure of 500 mb height fluctuations with long, intermediate and short time scales, *Journal of the atmospheric sciences*, *41*(6), 961–980, 1984.

Booth, J., L. Thompson, J. Patoux, K. Kelly, and S. Dickinson, The signature of the midlatitude tropospheric storm tracks in the surface winds, *Journal of Climate*, *23*(5), 1160–1174, 2010.

Boutle, I., S. Belcher, and R. Plant, Moisture transport in mid-latitude cyclones, *Quarterly Journal of the Royal Meteorological Society*, *136*, 1–14, 2010.

Branscome, L., A parameterization of transient eddy heat flux on a beta-plane, *Journal of the atmospheric sciences*, *40*(10), 2508–2521, 1983.

Browning, K., The dry intrusion perspective of extra-tropical cyclone development, *Meteorological Applications*, *4*(4), 317–324, 1997.

Browning, K., M. Hardman, T. Harrold, and C. Pardoe, The structure of rainbands within a mid-latitude depression, *Quarterly Journal of the Royal Meteorological Society*, *99*(420), 215–231, 1973.

Budyko, M., The effect of solar radiation variations on the climate of the earth, *Tellus*, *21*(5), 611–619, 1969.

Businger, S., The synoptic climatology of polar-low outbreaks over the Gulf of Alaska and the Bering Sea, *Tellus A*, *39*(4), 307–325, 1987.

Campbell, G., and T. Vonder Haar, *Climatology of radiation budget measurements from satellites*, Department of Atmospheric Science, Colorado State University, 1980.

Carleton, A., Cloud-cryosphere interactions, in *Satellite Sensing of a Cloudy Atmosphere: observing the Third Planet*, edited by A. Henderson-Sellers, chap. 8, pp. 289–325, Taylor & Francis, Philadelphia, 1984.

Carleton, A., Meridional transport of eddy sensible heat in winters marked by extremes of the North Atlantic Oscillation, 1948/49-1979/80., *Journal of climate*, *1*, 212–224, 1988.

Carlson, T., Airflow through midlatitude cyclones and the comma cloud pattern, *Monthly Weather Review*, *108*(10), 1498–1509, 1980.

Catto, J., L. Shaffrey, and K. Hodges, Can climate models capture the structure of extratropical cyclones?, *Journal of Climate*, *23*(7), 1621–1635, 2010.

Chang, E., Downstream development of baroclinic waves as inferred from regression analysis, *Journal of the atmospheric sciences*, *50*(13), 2038–2053, 1993.

Charney, J., The dynamics of long waves in a baroclinic westerly current, *Journal of Meteorology*, *4*(5), 136–162, 1947.

Cheng, X., and J. Wallace, Cluster analysis of the Northern Hemisphere wintertime 500-hPa height field: Spatial patterns, *Journal of the Atmospheric Sciences*, *50*(16), 2674–2696, 1993.

Condron, A., G. Bigg, and I. Renfrew, Polar mesoscale cyclones in the northeast Atlantic: Comparing climatologies from ERA-40 and satellite imagery, *Monthly weather review*, *134*(5), 1518–1533, 2006.

Czaja, A., and J. Marshall, The Partitioning of Poleward Heat Transport between the Atmosphere and Ocean, *J. Atmos. Sci.*, *63*, 1498 – 1511, 2005.

Daubechies, I., The wavelet transform, time-frequency localization and signal analysis, *Information Theory, IEEE Transactions on*, *36*(5), 961–1005, 1990.

Dee, D., et al., The ERA-Interim reanalysis: Configuration and performance of the data assimilation system, *Quarterly Journal of the Royal Meteorological Society*, *137*(656), 553–597, 2011.

Dixon, J., *Appendix B: Properties of Air*, pp. 375–378, John Wiley & Sons, Ltd, doi: 10.1002/9780470516430.app2, 2007.

Eady, E., Long waves and cyclone waves, *Tellus*, *1*(3), 33–52, 1949.

Eckhardt, S., A. Stohl, H. Wernli, P. James, C. Forster, and N. Spichtinger, A 15-year climatology of warm conveyor belts, *Journal of climate*, *17*(1), 218–237, 2004.

Fasullo, J., and K. Trenberth, The annual cycle of the energy budget. Part II: Meridional structures and poleward transports, *Journal of Climate*, *21*(10), 2313–2325, 2008.

Forster, P. d. F., M. Blackburn, R. Glover, and K. Shine, An examination of climate sensitivity for idealised climate change experiments in an intermediate general circulation model, *Climate Dynamics*, *16*(10-11), 833–849, 2000.

Gill, A., *Atmosphere-Ocean Dynamics*, vol. 30, Academic press, 1982.

Gill, A., J. Green, and A. Simmons, Energy partition in the large-scale ocean circulation and the production of mid-ocean eddies, *Deep Sea Research*, *21*, 499–528, 1974.

Goupillaud, P., A. Grossmann, and J. Morlet, Cycle-octave and related transforms in seismic signal analysis, *Geoexploration*, *23*(1), 85–102, 1984.

Graversen, R., Do changes in the midlatitude circulation have any impact on the Arctic surface air temperature trend?, *Journal of Climate*, *19*(20), 5422–5438, 2006.

Graversen, R., T. Mauritsen, S. Drijfhout, M. Tjernström, and S. Mårtensson, Warm winds from the Pacific caused extensive Arctic sea-ice melt in summer 2007, *Climate dynamics*, *36*(11-12), 2103–2112, 2011.

Grinsted, A., J. Moore, and S. Jevrejeva, Application of the cross wavelet transform and wavelet coherence to geophysical time series, *Nonlinear processes in geophysics*, *11*(5-6), 561–566, 2004.

Gumbel, E., *Statistics of extremes*, Columbia University Press, 1958.

Hall, A., The role of surface albedo feedback in climate, *Journal of climate*, *17*(7), 1550–1568, 2004.

Hamilton, K., Aspects of wave behaviour in the mid-and upper troposphere of the southern Hemisphere, *Atmosphere-Ocean*, *21*(1), 40–54, 1983.

Harrold, T., Mechanisms influencing the distribution of precipitation within baroclinic disturbances, *Quarterly Journal of the Royal Meteorological Society*, *99*(420), 232–251, 1973.

Hogg, N., and H. Stommel, The heton, an elementary interaction between discrete baroclinic geostrophic vortices, and its implications concerning eddy heat-flow, *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, *397*(1812), 1–20, 1985a.

Hogg, N. G., and H. M. Stommel, Hetonic explosions: the breakup and spread of warm pools as explained by baroclinic point vortices, *Journal of the atmospheric sciences*, *42*(14), 1465–1476, 1985b.

Holton, J., *An introduction to dynamic meteorology*, Academic press, 1979.

Holton, J., and G. Hakim, *An introduction to dynamic meteorology*, Academic press, 2012.

Hoskins, B., and K. Hodges, New perspectives on the Northern Hemisphere winter storm tracks, *Journal of the Atmospheric Sciences*, *59*(6), 1041–1061, 2002.

Hoskins, B., and K. Hodges, A new perspective on Southern Hemisphere storm tracks, *Journal of Climate*, *18*(20), 4108–4129, 2005.

Hoskins, B., and A. Simmons, A multi-layer spectral model and the semi-implicit method, *Quarterly Journal of the Royal Meteorological Society*, *101*(429), 637–655, 1975.

Hoskins, B., and P. Valdes, On the existence of storm-tracks, *Journal of the atmospheric sciences*, *47*(15), 1854–1864, 1990.

Houghton, H., On the annual heat balance of the northern hemisphere, *Journal of meteorology*, *11*(1), 1–9, 1954.

Hurrell, J., Decadal trends in the North Atlantic oscillation, *Science*, *269*, 676–679, 1995.

Hurrell, J., Y. Kushnir, G. Ottersen, and M. Visbeck (Eds.), *The North Atlantic Oscillation: climatic significance and environmental impact*, vol. 134, American Geophysical Union, 2003.

Hwang, Y.-T., and D. Frierson, Increasing atmospheric poleward energy transport with global warming, *Geophysical Research Letters*, *37*, L24,807, 2010.

ICAO, *Manual of the ICAO Standard Atmosphere*, International Civil Aviation Organization, 1952.

Johnston, D., M. Bowers, A. Friedlaender, and D. Lavigne, The effects of climate change on harp seals (Pagophilus groenlandicus), *PloS one*, *7*(1), e29,158, 2012.

Jones, N., Troubling milestone for CO2, *Nature Geoscience*, *6*(8), 589–589, 2013.

Kim, M.-K., and Y.-H. Kim, Examination of the global Lorenz energy cycle and preliminary applications for regime shift, in *APCC Seminar*, 2012.

Kolstad, E., A new climatology of favourable conditions for reverse-shear polar lows, *Tellus A*, *58*(3), 344–354, 2006.

Kolstad, E., A global climatology of favourable conditions for polar lows, *Quarterly Journal of the Royal Meteorological Society*, *137*(660), 1749–1761, 2011.

Konrad, C., and S. Colucci, An examination of extreme cold air outbreaks over eastern North America, *Monthly Weather Review*, *117*(12), 2687–2700, 1989.

Langen, P., and V. Alexeev, Polar amplification as a preferred response in an idealized aquaplanet GCM, *Climate dynamics*, *29*(2-3), 305–317, 2007.

Larichev, V., and I. Held, Eddy amplitudes and fluxes in a homogeneous model of fully developed baroclinic instability, *Journal of physical oceanography*, *25*(10), 2285–2297, 1995.

Lau, K., and H. Weng, Climate signal detection using wavelet transform: How to make a time series sing, *Bulletin of the American Meteorological Society*, *76*(12), 2391–2402, 1995.

Lau, N.-C., On the three-dimensional structure of the observed transient eddy statistics of the Northern Hemisphere wintertime circulation, *Journal of the Atmospheric Sciences*, *35*(10), 1900–1923, 1978.

Lau, N.-C., and J. Wallace, On the Distribution of Horizontal Transports by Transient Eddies in the Northern Hemisphere Wintertime Circulation', *Journal of the Atmospheric Sciences*, *36*(10), 1844–1861, 1979.

Lee, S., T. Gong, N. Johnson, S. Feldstein, and D. Pollard, On the possible link between tropical convection and the Northern Hemisphere Arctic surface air temperature change between 1958 and 2001, *Journal of Climate*, *24*(16), 4350–4367, 2011.

Legg, S., and J. Marshall, A heton model of the spreading phase of open-ocean deep convection, *Journal of physical oceanography*, *23*(6), 1040–1056, 1993.

L'Heveder, B., and M.-N. Houssais, Investigating the varaibility of the Arctic sea ice thickness in response to a stochastic thermodynamic atmospheric forcing, *Clim. Dynam.*, *17*, 107 – 125, 2001.

Lindzen, R., *Dynamics in Atmospheric Physics*, 310 pp., Cambridge University Press, 1990.

Lorenz, E. N., Available potential energy and the maintenance of the general circulation, *Tellus*, *7*(2), 157–167, 1955.

Madonna, E., H. Wernli, H. Joos, and O. Martius, Warm conveyor belts in the ERA-Interim data set (1979-2010). Part I: Climatology and potential vorticity evolution., *Journal of Climate*, –, –, Submitted.

Marshall, J., and R. Plumb, *Atmosphere, ocean and climate dynamics: an introductory text*, Academic Press, 1997.

McClellan, J., T. Parks, and L. Rabiner, A computer program for designing optimum FIR linear phase digital filters, *Audio and Electroacoustics, IEEE Transactions on*, *21*(6), 506–526, 1973.

McIlveen, R., *Fundamentals of weather and climate*, Stanley Thornes Publishers, 1998.

Messori, G., and A. Czaja, The temporal variability and associated spectral features of transient-eddy meridional heat transport, in *Proceedings of the First Annual Conference of the Italian Society for Climate Sciences*, 2013a.

Messori, G., and A. Czaja, On the sporadic nature of meridional heat transport by transient eddies, *Quarterly Journal of the Royal Meteorological Society*, *139*, 999–1008, 2013b.

Messori, G., and A. Czaja, Some considerations on the spectral features of meridional heat transport by transient eddies, *Quarterly Journal of the Royal Meteorological Society*, *Accepted*, –, 2013c.

Meyers, S., B. Kelly, and J. O'Brien, An introduction to wavelet analysis in oceanography and meteorology: With application to the dispersion of Yanai waves, *Monthly Weather Review*, *121*(10), 2858–2866, 1993.

Minobe, S., A. Kuwano-Yoshida, N. Komori, S.-P. Xie, and R. Small, Influence of the Gulf Stream on the troposphere, *Nature*, *452*(7184), 206–209, 2008.

Nakamura, H., T. Izumi, and T. Sampe, Interannual and decadal modulations recently observed in the Pacific storm track activity and East Asian winter monsoon, *Journal of Climate*, *15*(14), 1855–1874, 2002.

Nakamura, N., and A. Oort, Atmospheric Heat Budgets of the Polar Regions, *J. Geophys. Res.*, *93*, 9510 – 9524, 1988.

Neelin, J., and I. Held, Modeling tropical convergence based on the moist static energy budget, *Monthly Weather Review*, *115*(1), 3–12, 1987.

Oort, A., The Observed Annual Cycle in the Meridional Transport of Atmospheric Energy, *J. Atmos. Sci.*, *28*, 325 – 339, 1971.

Oort, A., and J. Peixoto, Global angular momentum and energy balance requirements from observations, *Adv. Geophys*, *25*, 355–490, 1983.

Overland, J., and C. Pease, Cyclone climatology of the Bering Sea and its relation to sea ice extent, *Monthly Weather Review*, *110*(1), 5–13, 1982.

Overland, J., and P. Stabeno, Is the climate of the Bering Sea warming and affecting the ecosystem?, *Eos, Transactions American Geophysical Union*, *85*(33), 309–312, 2004.

Pacanowski, R., K. Dixon, and A. Rosati, The GFDL modular ocean model users guide, version 1.0, *GFDL Ocean Group Tech. Rep*, *2*(46), 08,542–0308, 1991.

Pedlosky, J., The instability of continuous heton clouds, *Journal of the atmospheric sciences*, *42*(14), 1477–1486, 1985.

Peixoto, J., and A. Oort, *Physics of Climate*, 520 pp., American Institute of Physics, 1992.

Pfahl, S., and H. Wernli, Quantifying the relevance of cyclones for precipitation extremes, *Journal of Climate*, *25*(19), 6770–6780, 2012.

Phillips, N., A simple three-dimensional model for the study of large-scale extratropical flow patterns, Ph.D. thesis, University of Chicago, Department of Meteorology., 1951.

Porter, D., J. Cassano, M. Serreze, and D. Kindig, New estimates of the large-scale Arctic atmospheric energy budget, *Journal of Geophysical Research*, *115*(D8), 2010.

Randel, W., and I. Held, Phase speed spectra of transient eddy fluxes and critical layer absorption, *Journal of the atmospheric sciences*, *48*(5), 688–697, 1991.

Risien, C., and D. Chelton, A global climatology of surface wind and wind stress fields from eight years of QuikSCAT scatterometer data, *Journal of Physical Oceanography*, *38*(11), 2379–2413, 2008.

Salby, M., *Fundamentals of atmospheric physics*, vol. 61, Academic press, 1996.

Sampe, T., and S.-P. Xie, Mapping high sea winds from space: A global climatology, *Bulletin of the American Meteorological Society*, *88*(12), 1965–1978, 2007.

Sellers, W., A Global Climatic Model Based on the Energy Balance of the Earth-Atmosphere System, *J. Applied Meteorology*, *8*, 392–400, 1969.

Simmons, A., and B. Hoskins, Baroclinic instability on the sphere: Normal modes of the primitive and quasi-geostrophic equations, *Journal of the Atmospheric Sciences*, *33*(8), 1454–1477, 1976.

Simmons, A., and B. Hoskins, The life cycles of some nonlinear baroclinic waves, *Journal of the Atmospheric Sciences*, *35*(3), 414–432, 1978.

Simmons, A., S. Uppala, D. Dee, and S. Kobayashi, ERA-Interim: New ECMWF reanalysis products from 1989 onwards, *ECMWF newsletter*, *110*, 25–35, 2007.

Sinha, B., and R. Smith, Development of a fast coupled general circulation model (FORTE) for climate studies, implemented using the OASIS coupler, *Southampton Oceanography Centre Internal Document*, *81*, 67pp, 2002.

Slingo, J., The development and verification of a cloud prediction scheme for the ECMWF model, *Quarterly Journal of the Royal Meteorological Society*, *113*(477), 899–927, 1987.

Smith, R., and J. Gregory, A study of the sensitivity of ocean overturning circulation and climate to freshwater input in different regions of the North Atlantic, *Geophysical Research Letters*, *36*(15), L15,701, 2009.

Smith, R., C. Dubois, and J. Marotzke, Global climate and ocean circulation on an aquaplanet ocean-atmosphere general circulation model, *Journal of climate*, *19*(18), 4719–4737, 2006.

Stommel, H., The westward intensification of wind-driven ocean currents, *Trans. Amer. Geophys. Union*, *29*(2), 202–206, 1948.

Stone, P., Constraints on dynamical transports of energy on a spherical planet, *Dynamics of Atmospheres and Oceans*, *2*(2), 123–139, 1978.

Stone, P., and D. Miller, Empirical Relations Between Seasonal Changes in Meridional Temperature Gradients and Meridional Fluxes of Heat, *J. Atmos. Sci.*, *37*, 1708 – 1721, 1980.

Swanson, K., and R. Pierrehumbert, Lower-tropospheric heat transport in the Pacific storm track, *Journal of the atmospheric sciences*, *54*(11), 1533–1543, 1997.

Todaro, R. (Ed.), *The Stratospheric Ozone Electronic Textbook*, Goddard Space Flight Center Atmospheric Chemistry and Dynamics Branch, 2003.

Torrence, C., and G. Compo, A practical guide to wavelet analysis, *Bulletin of the American Meteorological society*, *79*(1), 61–78, 1998.

Torrence, C., and P. Webster, Interdecadal changes in the ENSO-monsoon system, *Journal of Climate*, *12*(8), 2679–2690, 1999.

Trenberth, K., Climate diagnostics from global analyses: Conservation of mass in ECMWF analyses, *Journal of climate*, *4*(7), 707–722, 1991.

Trenberth, K., and J. Caron, Estimates of meridional atmosphere and ocean heat transports, *Journal of Climate*, *14*(16), 3433–3443, 2001.

Wahl, E., H. Diaz, and C. Ohlwein, A pollen-based reconstruction of summer temperature in central North America and implications for circulation patterns during medieval times, *Global and Planetary Change*, *84*, 66–74, 2012.

Walker, G., and E. Bliss, World Weather. V., *Mem. R. Meteorol. Soc.*, *4*, 53–84, 1932.

Webb, D., An ocean model code for array processor computers, *Computers & Geosciences*, *22*(5), 569–578, 1996.

Yiou, P., K. Goubanova, Z. Li, and M. Nogaj, Weather regime dependence of extreme value statistics for summer temperature and precipitation, *Nonlinear Processes in Geophysics*, *15*(3), 365–378, 2008.