Department of Physics and Astronomy, University of Canterbury,
Private Bag 4800, Christchurch 8140, New Zealand

Application of Complexity Measures to Stratospheric Dynamics

A thesis
submitted in partial fulfilment
of the requirements for
a
Masters Degree in Physics
at the
University of Canterbury

by

Nikolai Christian Krützmann

University of Canterbury

2008

Supervisors: Dr. A. McDonald
Dr. S. George

Abstract

This thesis examines the utility of mathematical complexity measures for the analysis of stratospheric dynamics. Through theoretical considerations and tests with artificial data sets, e.g., the iteration of the logistic map, suitable parameters are determined for the application of the statistical entropy measures *sample entropy* (*SE*) and *Rényi entropy* (*RE*) to methane (a long-lived stratospheric tracer) data from simulations of the SOCOL chemistry-climate model.

The *SE* is shown to be useful for quantifying the variability of recurring patterns in a time series and is able to identify tropical patterns similar to those reported by previous studies of the "tropical pipe" region. However, the *SE* is found to be unsuitable for use in polar regions, due to the non-stationarity of the methane data at extra-tropical latitudes. It is concluded that the *SE* cannot be used to analyse climate complexity on a global scale.

The focus is turned to the *RE*, which is a complexity measure of probability distribution functions (PDFs). Using the second order *RE* and a normalisation factor, zonal PDFs of ten consecutive days of methane data are created with a Bayesian optimal binning technique. From these, the *RE* is calculated for every day (moving 10-day window). The results indicate that the *RE* is a promising tool for identifying stratospheric mixing barriers. In Southern Hemisphere winter and early spring, *RE* produces patterns similar to those found in other studies of stratospheric mixing. High values of *RE* are found to be indicative of the strong fluctuations in tracer distributions associated with relatively unmixed air in general, and with gradients in the vicinity of mixing barriers, in particular. Lower values suggest more thoroughly mixed air masses.

The analysis is extended to eleven years of model data. Realistic inter-annual variability of some of the *RE* structures is observed, particularly in the Southern Hemisphere. By calculating a climatological mean of the *RE* for this period, additional mixing patterns are identified in the Northern Hemisphere. The validity of the *RE* analysis and its interpretation is underlined by showing that qualitatively similar patterns can be seen when using observational satellite data of a different tracer. Compared to previous techniques, the *RE* has the advantage that it requires significantly less computational effort, as it can be used to derive dynamical information from model or measurement tracer data without relying on any additional input such as wind fields.

The results presented in this thesis strongly suggest that the *RE* is a useful new metric for analysing stratospheric mixing and its variability from climate model data. Furthermore, it is shown that the *RE* measure is very robust with respect to data gaps, which makes it ideal for application to observations. Hence, using the *RE* for comparing observations of tracer distributions with those from model simulations potentially presents a novel approach for analysing mixing in the stratosphere.

Contents

List of Figures

Chapter 1

Introduction

> "*The fluttering of a butterfly's wing in Rio de Janeiro, amplified by atmospheric currents, could cause a tornado in Texas two weeks later.*" - Edward N. Lorenz

1.1 Motivation

The above quote constitutes what today is known as the "butterfly effect". It is arguably the most well-known result of Chaos Theory and illustrates a very important consequence: the atmosphere is a very complex system and accurate long-term weather predictions are almost always impossible.

In 1961, E. N. Lorenz discovered that many non-linear systems, under certain conditions, may exhibit high sensitivity to initial conditions. This means that a very small uncertainty or perturbation in the initial state of a system, can eventually result in entirely different dynamics on a much larger scale. Accordingly, the flap of a butterfly's wings can have an effect on the Earth's weather system. Most natural systems are complex, meaning that they are non-linear and do not show periodic behaviour. As expected, these systems are essentially unpredictable beyond a certain time-frame, because the initial conditions of any real-world system cannot be determined with perfect accuracy. The length of this time-frame, i.e., the time it takes for small-scale variations to have large-scale effects, depends on the details of the interactions that make the system non-linear – feedback being the most common mechanism. The knowledge and understanding of these details is therefore paramount not only for the attempting any kind of prediction of a complex system's behaviour, but also for determining the reliability of the prediction.

It was through simulations of a simplified mathematical model of atmospheric flow that Lorenz made his discovery. Although it was a highly simplified model, a small round-off error led to extremely different results of two 'identical' simulation runs. This suggests that the Earth's atmosphere *is* a complex system and that accurate long-term predictions are most likely impossible (*Lorenz*, 1963). Today, the maximum timescale for relatively accurate weather forecasts is believed to be approximately two weeks (*Shepherd*, 2003).

However, it is obvious that atmospheric dynamics are not completely random. The occurrence of cyclic trends such as the seasons which are generated by external solar forcing, is one example. Such general long-term patterns are more easily predicted than the short-term details. Thus, it is possible to identify and predict certain long-term patterns of atmospheric properties, as long as the timescale for which relative accuracy is desired is sufficiently large. This permits large-scale dynamic prediction, despite the complex nature of the underlying small-scale processes. The study of such long-term patterns (climatologies), is a common approach for understanding atmospheric dynamics. However, climatological analysis does not in general allow any accurate predictions. The regular climatological patterns are usually strongly affected by noise, i.e., they never occur in exactly the same way. Or as Mark Twain put it: "*Climate is what we expect, weather is what we get.*"

This illustrates that, while the knowledge of climatological averages is of course very useful information, it does not give any information about the causes of the variations from the average. Therefore, knowledge about the variability of climatological processes is very important for analysing and understanding atmospheric dynamics in order to improve climate models in general and their predictive capabilities in particular. Furthermore, a quantitative measure of the variability of atmospheric processes can also help to identify changes in this variability. The latter could be an indicator for and caused by large-scale climatological change.

Due to the large number of dynamical degrees-of-freedom and non-linear feedback processes in the atmosphere, its variability is very complex. A way of analysing the complexity of such a chaotic system (in a mathematical sense) is the use of statistical measures (*Beck and Schlögl*, 1993). Some of the main sources of complexity within the atmosphere are turbulent mixing and associated transport. Long-lived trace gases, such as methane, have been shown to be a means of detecting mixing (*Shepherd*, 2003). Methane is produced by biological processes in the troposphere (the lowest part of the atmosphere that extends from the surface to approximately 10 km altitude), and transported into the stratosphere (between 11 and 47 km altitude) in the tropics (*Randel et al.*, 1998). Therefore, methane distributions reflect the action of transport and mixing. The distributions of long-lived trace gases as well as aerosols can display large fluctuations at the boundary between distinct regions in the stratosphere (*Trepte and Hitchman*, 1992; *Plumb*, 1996). Dynamical variations in the boundary regions have the potential to cause sudden mixing of previously separated air masses, leading to noticeable changes in tracer concentrations.

The aim of this thesis is to study the utility of statistical complexity measures for understanding and analysing the mixing of chemical constituents in the stratosphere. The study considers the sample entropy (SE) and Rényi entropy (RE), which are both measures of complexity that have been successfully used in various fields of research. These measures examine the underlying distribution of a data set. The focus of this study is on the stratosphere as there are several well-known large-scale mixing patterns present, and because recent research indicates that this region of the atmosphere might have a significantly greater effect on near-surface dynamics than previously assumed (*Baldwin et al.*, 2003).

The stratosphere mainly affects the troposphere through multiple feedback processes. Changes in the troposphere cause changes in the stratosphere, which further affect the troposphere. One example is the Antarctic ozone hole. The ozone hole is caused by anthropogenic trace gases that are transported into the stratosphere. The severe reduction of stratospheric ozone in Antarctic spring leads to an increase in ultra-violet radiation at the surface and a consequent decrease in stratospheric heating. The reduced temperatures in the stratosphere result in stronger stratospheric polar winds, which in turn affect the troposphere potentially via modulation of planetary wave propagation. Although less severe, ozone depletion also occurs in the Northern Hemisphere, with similar impacts on the Northern Hemisphere polar vortex. A strong connection between the Northern Hemisphere vortex and the Northern-Annular-Mode (NAM) has been observed. The NAM is an observed sea-level pressure phenomenon that has large effects on the weather in the Northern Hemisphere (*Baldwin et al.*, 2003). A similar connection exists between the Antarctic ozone hole and the Southern Annular Mode (SAM) that affects Southern Hemisphere surface climate (*Thompson and Solomon*, 2002). This shows that the stratosphere has a strong effect on the troposphere and that understanding stratospheric dynamics can also improve the understanding of the tropospheric weather.

Significant debate exists regarding the mechanisms that drive many of these stratosphere-troposphere interactions. A more detailed understanding of the stratospheric dynamics is required as well as improved methods for analysing mixing. For this initial study, data from the SOCOL chemistry-climate model is utilized.

1.2 Structure

This thesis is divided into four main chapters and a concluding chapter that summarises the results. Chapter 2 presents background information which provides the foundation for the following chapters. These can be read individually, but at times draw on some of the results detailed in the preceding chapters.

Chapter 2 begins with a brief introduction to the structure and dynamics of the atmosphere, focusing particularly on the stratosphere. This is followed by an explanation of the use of stratospheric trace gases (tracers) for analysing dynamic processes, with a particular focus on methane, the tracer used for this study. With the help of a widely used example of a simple (mathematical) system that exhibits complex behaviour (the iteration of the logistic map), the principles of chaos and complexity are then introduced. The next part of chapter 2 provides details on how statistical entropy measures can be used to quantify the complexity of a dynamical system, and introduces the sample entropy (SE) and Rényi entropy (RE) measures. These are applied to stratospheric methane data from the SOCOL chemistry-climate model later in the thesis. The SOCOL model is explained in the final part of the chapter.

Chapter 3 starts with a detailed investigation of the optimal values for the parameters of the SE metric, which is a measure of the variability of time series. The quantification of complexity achieved by the SE is then tested with the logistic map process. Following this, the SE is applied to one year of methane data which results in a consistent high-entropy pattern in the equatorial region, possibly associated with the "tropical pipe" identified by other studies. To analyse this in more detail a slight alteration to the method of calculating the SE is performed. During the analysis, a problem with the application of the SE measure on a global scale is discovered, which is then examined in greater detail in the last part of the chapter and conclusions are drawn as to the utility of the SE measure for the analysis of stratospheric dynamics from tracer data.

In order to avoid the problems encountered with the *SE* measure, the *RE* measure is used in chapter 4. The *RE* statistical measure does not require long time series as input, but analyses the complexity of probability distribution functions (PDFs). After detailing the parameter values for the analysis, an explanation of the Bayesian optimal binning algorithm used for creating the PDFs is presented. The quantification of complexity achieved by the *RE* is then tested with the logistic map. The second half of the chapter details the results of the application of the *RE* to methane data from the SOCOL model. These are discussed in a general manner and found to correspond well to previous results of stratospheric studies, particularly in the Southern Hemisphere. This is then confirmed by comparisons with other studies of stratospheric mixing and transport, which leads to positive conclusions about the utility of the *RE* for the analysis of mixing and mixing barriers in the stratosphere.

Building on the previous chapter, chapter 5 verifies the results through a more detailed comparison of the location of the identified entropy structures and the location of the associated real-world structures, which have been identified in previous studies using different methodologies. In particular, the Southern Hemisphere winter polar vortex can be associated with a distinctive *RE* pattern from autumn to late spring. The analysis is then extended to a total of 11 years of SOCOL data and the results are compared and discussed. The identified patterns in the Southern Hemisphere are very self-consistent and show realistic variability. Therefore, the full data set is used to create a climatological mean of the *RE*. From this, additional features in the Northern Hemisphere are identified. The last part of this chapter gives a summary of the results and the utility of the *RE* for mixing analysis.

Chapter 6 summarises the conclusions of the previous chapters and indicates possible areas where future work could be pursued.

Chapter 2

Background

2.1 The Atmosphere

The Earth's atmosphere is a very large and complicated system. Understanding (and predicting) its dynamics has been of interest to human culture, ever since the cultivation of crops as the main source of food began more than 10,000 years ago. Agriculture, the origin of the word 'culture', requires a certain understanding of the cycle of the seasons in order to be successful. The ability to recognise and predict weather patterns can significantly improve the harvest. Since weather only occurs in the lowest few kilometres of the atmosphere (called the troposphere), until quite recently the majority of atmospheric research has been concerned with the dynamics in this region. The atmosphere actually has a thickness of several hundred kilometres, and processes that occur at higher altitudes can have a significant effect on the dynamics close to the Earth's surface (e.g., *Baldwin et al.*, 2003, 2007; *Matthes et al.*, 2006; *Perlwitz and Harnik*, 2003). Therefore, understanding the dynamics occurring at higher altitudes is of vital importance for understanding the atmosphere as a whole.

This section (section 2.1) gives a brief overview of the atmosphere and its large-scale structure. Following this, details of the dynamics of the stratosphere and their relationship with chemical trace gas distributions are discussed. Section 2.2 explains what dynamic complexity means by illustrating the evolution in the more theoretical context of the logistic map process, and then addresses how the concept relates to the analysis of atmospheric dynamics. Section 2.3 presents details about statistical complexity measures and how they may be applied to tracer data in order to analyse large-scale dynamics, which is the focus of this thesis. The SOCOL chemistry-climate model (CCM) is introduced in section 2.4 and is used as a source of atmospheric trace gas data in the following chapters.

2.1.1 General Atmospheric Structure

The atmosphere is vertically separated into five distinct regions. Beginning at the Earth's surface, these are the troposphere, stratosphere, mesosphere, thermosphere and the exosphere. The latter marks the transition between the Earth's atmosphere and space, and it extends from about 600 to 1000 km above the surface. The extent of the other layers is shown in figure 2.1. The boundary layers in between are called the tropopause, stratopause, mesopause, and thermopause respectively. This separation is based on the mean temperature

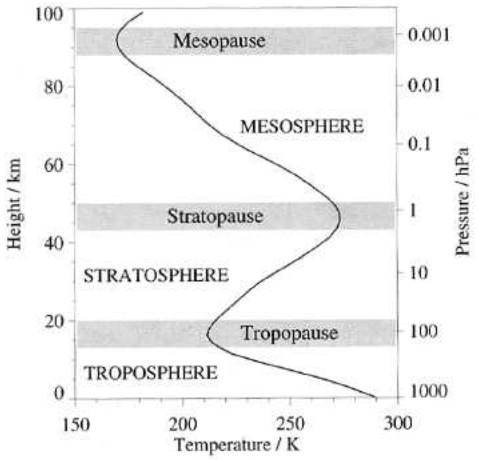

Figure 2.1: The lowest three layers of the atmosphere and the evolution of the (mean) temperature with altitude. The boundary layers are characterised by a reversal of the temperature gradient (from *Andrews*, 2000).

gradient (or lapse rate), which is also depicted in figure 2.1.

The boundary layers (the -pauses) are characterised by the reversal of the temperature gradient. In particular, the (thermal-)tropopause is defined as the region where the absolute value of the vertical temperature gradient is $2°\frac{K}{km}$ or less for a depth of at least 2 km (*Reiter et al.*, 1969). The sign of the temperature gradient has significant implications for the (thermo-)dynamics in the atmospheric layers. While regions with a positive temperature gradient are generally stable, a region with a negative lapse rate is relatively unstable with respect to vertical mixing and often associated with turbulent dynamics (*Andrews*, 2000). A stable region on the other hand, often shows a strong degree of stratification[1]. Due to the different dynamics, the four layers of the atmosphere are relatively isolated from one another and exchange between them through the boundary layers is limited.

The troposphere is dominated by large-scale convective mixing processes due to the strong radiative heating of the Earth's surface, with thorough global scale mixing. Intra-hemispheric mixing can occur on timescales as fast a few weeks (*Shepherd*, 2003), though inter-hemispheric mixing takes about one year (*Bowman and Cohen*, 1997). In contrast, the stratosphere can be considered to be in radiative equilibrium and stable with respect to convective mixing, and is therefore stably stratified in potential temperature surfaces. Potential temperature, θ, is an alternative vertical coordinate to pressure p or altitude. For an ideal gas it can be calculated by:

$$\theta = T\left(\frac{p_0}{p}\right)^{\frac{R}{c_p}} \tag{2.1}$$

T is the temperature of the air parcel, p_0 is the ground level pressure (usually 1013 mBar), R is the gas constant ($R \approx 287.05\frac{J}{kg\,K}$ for dry air) and $c_p \approx 1.03\frac{kJ}{kg\,K}$ is the specific heat capacity of air at constant pressure. θ can be interpreted as the temperature an air parcel would have if it were lowered to the Earth's surface adiabatically[2]. It is a useful vertical

---

[1] This is where the stratosphere gets its name from.

[2] *Adiabatic*: without adding or removing heat.

coordinate, as the motion of a stratospheric air parcel is most likely to occur on surfaces of constant potential temperature. This is because no heat has to be added or removed from the air parcel when moving along a potential temperature surface (or isentrope). As the stratosphere can be considered to be in radiative equilibrium, the additional heat required for non-adiabatic motion is not readily available. Hence, adiabatic motion is predominant for motion on relatively short timescales. The typical timescales of this adiabatic motion are between five and ten days (*Pierce et al.*, 1994). In general, the mixing timescales are rather longer than in the troposphere and global scale mixing can take several years (*Waugh and Hall*, 2002).

The relatively slow (mixing-) dynamics of the stratosphere are the reason for numerous large-scale climate phenomena, such as the Antarctic ozone hole or the long-term impacts of major volcanic eruptions (*Brasseur and Granier*, 1992). While dynamics in the stratosphere are generally simpler than in the troposphere and most of the basic processes behind the mentioned phenomena are quite well understood, there are still significant uncertainties as to the precise causes and interactions of many of the relevant processes. Therefore, this study focuses on the stratosphere and the usefulness of statistical complexity measures for understanding the known phenomena in detail. The following section gives a brief introduction into the known dynamics of the stratosphere and illustrates the importance of mixing processes.

2.1.2 Stratospheric Dynamics

The stratosphere, being defined by a positive temperature lapse rate (the temperature rises with altitude), is very stable with respect to convective mixing. Vertical motion is slow and the air is stably stratified along potential temperature surfaces. The air in the stratosphere is of relatively low density compared to the troposphere, as it ranges from 11 km to 47 km altitude.[3]

Horizontally, the stratosphere is commonly divided into four regions: the tropics, the (winter) mid-latitudes – often referred to as the “surf zone” – the polar vortex and the summer hemisphere (*Plumb*, 2002). It is apparent from the nomenclature that this separation is not permanent but seasonal. In winter, no sunlight reaches the polar region of the winter hemisphere, and it cools down significantly. This causes a very strong and persistent westerly (from west to east, eastward) wind to build up around the polar regions, which is often referred to as the polar night jet. These circumpolar stratospheric winds are strongest at around 32 km altitude ($\approx$ 10 mBar) and are a significant barrier to meridional (along a meridian, i.e. from north to south)[4] mixing, isolating the cold air inside the vortex from the surroundings. This isolation is a major prerequisite for the rapid springtime ozone depletion observed in the Southern Hemisphere (*Solomon*, 1999).

The surf zone denotes the middle latitudes (approximately from 20° to 60° latitude) in the winter hemisphere, separating the polar regions from the tropics. Here, significant planetary wave breaking occurs similar to water waves breaking on a beach, which is why it

---

[3] These are average values. Particularly the location of the lower boundary is very variable. It can be as low as 8 km (350 mBar) in the polar regions in winter and as high as 18km (100 mBar) in tropical latitudes.

[4] Similarly, *“zonal”* means along a circle of longitude, i.e., from east to west.

was termed the "surf zone" by *McIntyre and Palmer* (1984). Planetary waves are large-scale disturbances that can be caused by major topography or differences in radiative surface heating of landmasses compared to oceans. While planetary waves are created in the troposphere, they can, under the right conditions, propagate vertically and thereby reach the stratosphere. Rossby waves – a specific type of planetary wave, the restoring force of which is due to the conservation of the angular momentum associated with the Earth's rotational motion (*Shepherd*, 2003) – are the major driver of the large-scale stratospheric circulation known as the Brewer-Dobson circulation (*Brewer*, 1949 and *Dobson*, 1956). As planetary waves can only propagate vertically if the stratospheric zonal winds are westerly (*Holton*, 2004), the dominant summer easterlies do not allow such disturbances, making the summer hemisphere stratosphere rather calm. No large circulation patterns similar to the winter hemisphere can be detected.

The tropical region is relatively isolated from the rest of the stratosphere (*Trepte and Hitchman*, 1992), due to a net outward flow of air from this region. It is the major source of 'fresh' tropospheric air. In equatorial regions, the extremely strong convection caused by the intense heating of the sun, allows tropospheric air to rise high enough to reach the stratosphere. Therefore, one can imagine that tropospheric air is continuously pushed into the stratospheric tropics from below and is then further dispersed into the rest of the stratosphere. This image is not quite correct though, as the actual cause of this process is the Brewer-Dobson circulation and its driver lies in the surf zone.

Qualitatively, the Brewer-Dobson circulation can be described as tropical air rising from the troposphere to the stratosphere, from where it is then transported poleward at high altitudes and eventually descends back down. While air that descends at middle latitudes is recycled into the troposphere, air that descends at polar latitudes ($> 60°$) is stored in the polar lower stratosphere. In reality, Rossby wave breaking in the surf zone drives the whole process. Rossby waves are generally either stationary or propagate in an easterly direction. As they propagate vertically, they deposit their easterly momentum and *"exert a negative wave drag (...) which drives a poleward circulation within the stratosphere"* (*Shepherd*, 2003). Mass conservation requires the air that has moved poleward to be replaced, effectively sucking air up from tropics, this effect being called "wave drag". The air pumped to the polar regions cools (wintertime radiational cooling) and sinks. Again, the sinking air has to be replaced by more air from the middle stratosphere, which in turn has to be replaced from air at lower altitudes in the tropics. This feedback effect keeps a strong Brewer-Dobson circulation going continuously throughout the whole winter period. Figure 2.2 summarises the large-scale residual circulation of the middle atmosphere. The shaded region marked P denotes the region of planetary wave breaking.

The Brewer-Dobson circulation is limited to the winter hemisphere, as the persistent summer easterlies do not allow planetary wave propagation into the stratosphere (*Holton*, 2004). This flow is stronger in the Northern Hemisphere because of enhanced planetary wave activity in this hemisphere (there is more topography that can create planetary waves), leading to more wave breaking and hence stronger wave drag. Despite this constant large-scale pumping, the mixing timescales in the stratosphere are very long compared to the troposphere and it takes approximately four to five years for air to be transported from the tropical troposphere to the high latitude lower stratosphere (*Waugh and Hall*, 2002).

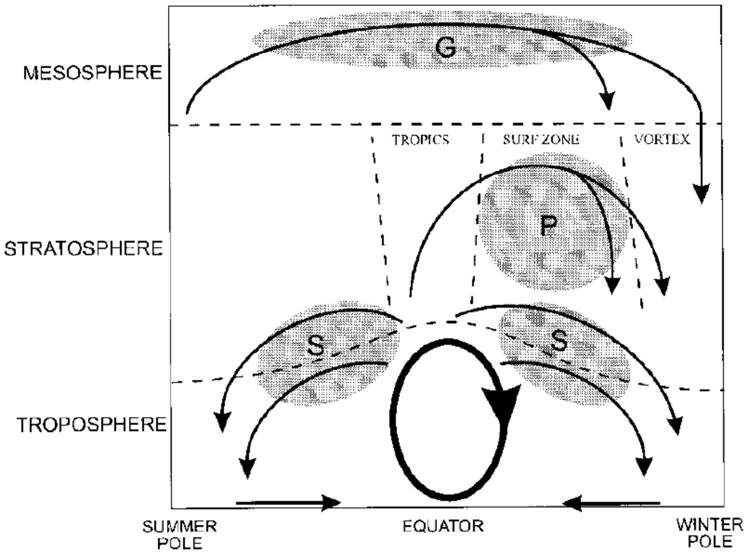

Figure 2.2: Large-scale (residual) circulation. The heavy ellipse denotes the thermally-driven Hadley circulation of the troposphere. The shaded regions (labelled “S”, “P”, and “G”) denote the regions of breaking waves (synoptic- and planetary-scale waves, and gravity waves, respectively), responsible for driving branches of the stratospheric and mesospheric circulation (from *Plumb*, 2002).

The constant planetary wave breaking at the edge of the polar vortex also acts to slow down the polar night jet and eventually causes the springtime break-up of the vortex. Particularly in the Northern Hemisphere the breaking of these wave motions considerably disturbs the polar vortex, distorts its shape and regularly displaces it off the pole. It significantly slows down the polar night jet, which results in the vortex being a more permeable transport barrier to horizontal mixing than its Southern Hemisphere equivalent. The higher planetary wave activity also accounts for the weaker and shorter boreal polar vortex period observed (*Waugh and Randel*, 1999).

The shaded regions, denoted S in figure 2.2, are regions where synoptic-scale waves break. These are important for the lower stratospheric and upper tropospheric circulation. The equatorial elliptical flow pattern symbolises the tropical Hadley circulation that transports air up in the tropics and down in the subtropics. Similarly, the mesospheric mean meridional circulation shown, transports air from the summer to the winter hemisphere at high altitudes (50-90 km). Rising (warm) air from the summer hemisphere replaces the sinking (cold) air in the winter hemisphere. This circulation is driven by small-scale gravity waves (shaded region region denoted G). These mechanisms are not discussed as this study focuses on large-scale stratospheric dynamical processes.

The following section describes how chemical trace gases, and particularly methane, can be used to analyse stratospheric dynamics.

2.1.3 Methane: a Chemical Tracer

Chemical trace gases, or tracers, are gaseous components of the atmosphere that only occur in small (trace) amounts in air. As the relative amounts of nitrogen (N_2) and oxygen (O_2)

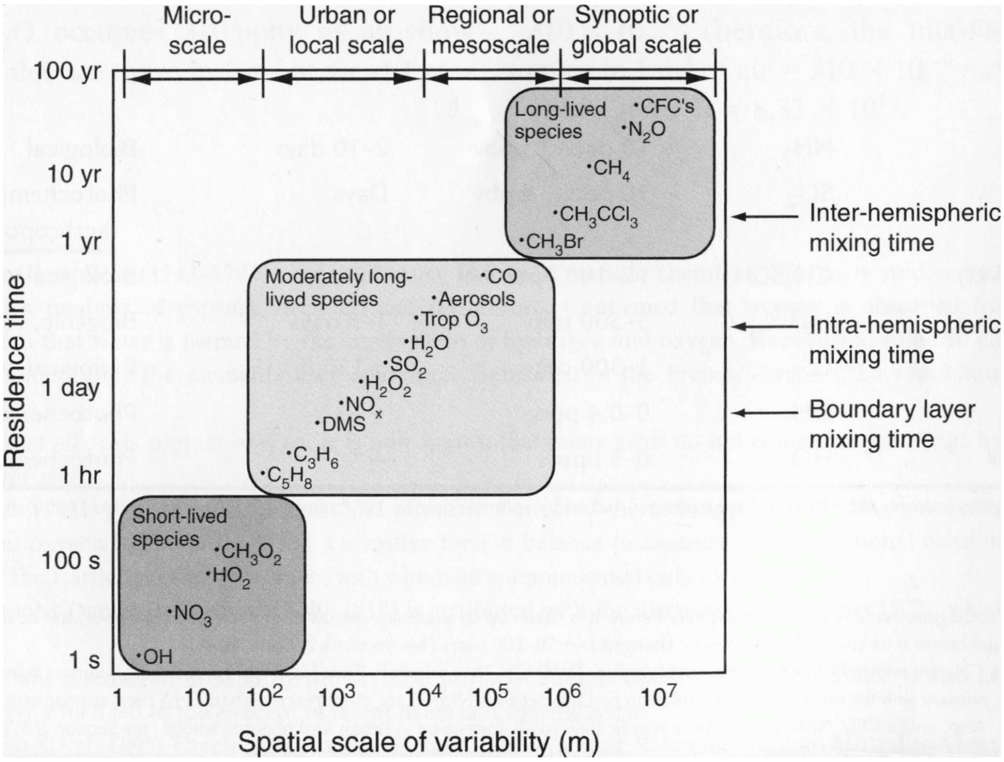

Figure 2.3: Atmospheric trace gas residence times and associated mixing scales (from *Wallace and Hobbs*, 2006, p. 156).

(78% and 21% respectively) are constant up to 120 km altitude and therefore make up 99% of the atmosphere, every gas other than nitrogen, oxygen and (arguably) argon (Ar, which contributes another 0.93%) is a trace gas. Due to the low relative amounts of trace gases in the air, a common unit for trace gas concentration is the mixing-ratio which can be given either in parts-per-million (ppm), i.e., the number of molecules of the tracer in one million molecules of air, or in parts-per-million-by-volume (ppmv), i.e., how many millionths of the total volume of the air are occupied by tracer molecules.

Trace gases can have very different lifetimes (or residence times)[5] in the atmosphere. These can reach from seconds (e.g., the hydroxyl radical (OH)) to more than a hundred years (e.g., some chloro-fluoro-carbons ($CFCs$)). Accordingly, the distributions of these gases are more or less affected by different scale processes. For example, a short-lived trace gas such as the hydroperoxyl radical (HO_2) will not be affected by advection over long distances. The time it would take for the tracer molecules to be transported (days) is much larger than its chemical lifetime (minutes) in the atmosphere, i.e., most molecules will have chemically reacted by the time the air has travelled an appreciable distance. On the other hand, the distributions of very long-lived tracers, such as methane, are not affected by local, small-scale mixing processes, as the tracer molecules have time to undergo several mixing cycles and are therefore relatively evenly distributed on small- or even synoptic-scales. Figure 2.3 gives an overview of the lifetimes of some important tracers and the associated spatial (mixing) scales that affect their distributions.

---

[5] The average time it takes for a molecule to be removed from the atmosphere. Removal is usually through some kind of chemical reaction, which is why it can also be referred to as the "chemical lifetime".

Methane in particular, which has an average atmospheric residence time of nine years and whose major source is biological decay (80%) (*Wallace and Hobbs*, 2006), is therefore distributed almost homogeneously throughout the whole troposphere (average mixing ratio in the troposphere ≈ 1.7 ppmv). In the stratosphere on the other hand, the methane distribution is much less uniform.

As mentioned in section 2.1, mixing between the troposphere and the stratosphere is rather limited. Methane is produced by biological processes in the troposphere and is transported into the stratosphere in the tropics (*Randel et al.*, 1998). It reaches high latitudes via a circulation that, whilst being primarily eddy-driven, is qualitatively similar to the Brewer-Dobson model. The major sink for stratospheric methane is photochemical destruction, which mainly occurs above 35 km via reactions with hydroxyl (OH) and oxygen radicals (O). The reaction can be summarised by equation 2.2 (*Randel et al.*, 1998).

$$CH_4 + 3O_2 + h\nu \longrightarrow 2H_2O + CO + O_3 \tag{2.2}$$

The light ($h\nu$) required for the reaction, needs to be of relatively short wavelength (high energy ultra-violet (UV) light), which is why the destruction of methane mainly occurs at high altitudes. At lower altitudes the amount of UV light is significantly reduced, as most of it has been absorbed by ozone.[6]

As stratospheric methane comes from the troposphere and is destroyed in the stratosphere, the amount of methane in a volume of stratospheric air is an indicator for how long that air has been in the stratosphere. The chemical destruction of methane via equation 2.2 reduces the methane content over time. This is referred to as chemical ageing. The "age of air" is a metric used for mixing analysis *(Waugh and Hall, 2002)* and can be linked to the methane content in stratospheric air *(Shepherd, 2003).* Combined with the relatively slow global scale dynamics in the stratosphere mentioned in the previous section, this leads to a typical pattern of methane concentrations in the stratosphere. This is illustrated in figure 2.4, which shows the zonal average methane distribution in the stratosphere and lower mesosphere in October. In general, the further away a point is from the lower tropical stratosphere entry point (bottom centre of the image), the lower (on the average) the methane mixing ratio.

While figure 2.4 is a snapshot, there is generally more methane in the tropical stratosphere than in the polar stratosphere, as methane is depleted during its progress to this region. The asymmetry of the contours can be explained by seasonal dynamics and hemispheric differences. In the Southern Hemisphere, the mixing ratio is generally lower than at the same latitude in the Northern Hemisphere. This is because the Brewer-Dobson circulation is stronger in the Northern Hemisphere winter and early spring due to higher planetary wave activity in that hemisphere. This allows faster transport of the tropical methane-rich air to higher latitudes which usually results in higher mixing ratios.

The strongest gradients in methane concentration occur at approximately 20°S and 60°S. This is due to the Brewer-Dobson circulation being active in the Southern Hemisphere and the gradients are associated with the transition regions between the tropical, surf zone, and polar vortex regions, as described in the previous section. Similar gradients can be found in

---

[6] The reactions leading to equation 2.2 also make methane the main source of stratospheric water vapour.

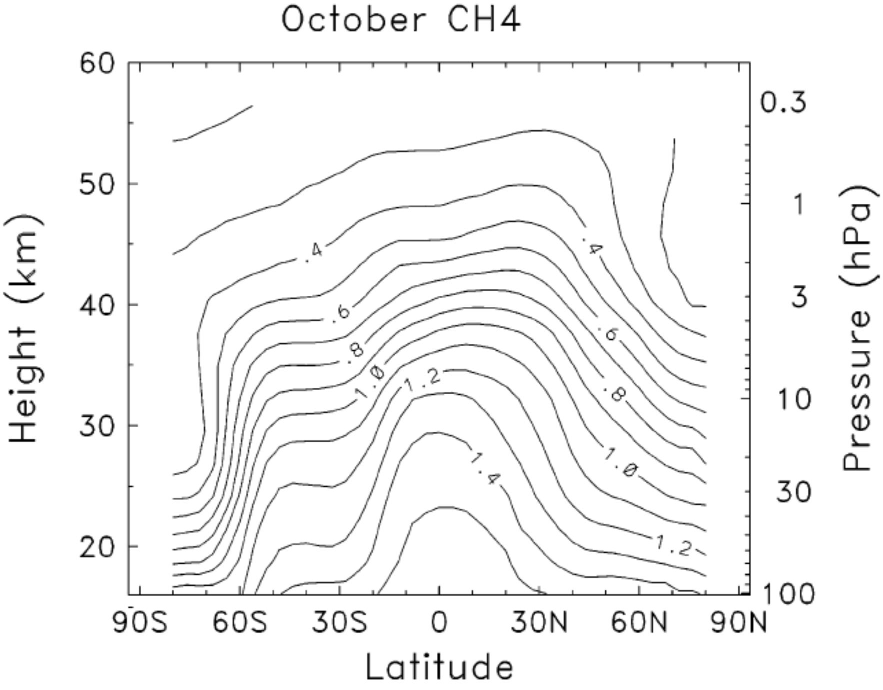

Figure 2.4: A typical latitude-altitude cross-section of mean methane mixing-ratio (contours in ppmv) for October from combined satellite measurements. The strong gradients in the Southern Hemisphere are associated with the transition regions between the tropical, surf zone, and polar vortex regions. The relatively high values in the Northern Hemisphere are due to generally stronger mixing dynamics (from *Shepherd*, 2003).

the Northern Hemisphere during its winter and spring period (not shown). Still, the methane gradients found in the Northern Hemisphere winter polar stratosphere are weaker than in the Southern Hemisphere winter polar stratosphere and the gradients also persist longer into spring in the Southern Hemisphere, because the austral polar vortex is significantly stronger than its boreal counterpart.

The connections between changes in methane concentrations and seasonal variations show that tracer distributions can give information about underlying dynamical processes. Generally, the distributions of both long-lived trace gases and also aerosols (solid and liquid particles) can display large fluctuations at the boundary between distinct regions in the stratosphere (*Trepte and Hitchman*, 1992; *Plumb*, 1996) and their analysis is a common approach to understanding atmospheric dynamics. The numerous interactions and feedbacks between transport, mixing, temperature, and chemistry in the atmosphere, lead to very complex, non-linear dynamics (*Shepherd*, 2003). Therefore appropriate non-linear analysis tools are required for studying atmospheric processes. The next section gives a brief introduction to the properties of non-linear systems, the concept of mathematical complexity, and statistical measures of complexity, which are applied to the analysis of atmospheric dynamics in the following chapters.

2.2 Complex Systems

In order to establish the notions of complexity and chaos, one of the most well known examples is introduced: the iteration of the logistic map. It is a simple non-linear system that is commonly used to illustrate these concepts and to test complexity measures (e.g., *Wackerbauer et al.*, 1994).

2.2.1 The Logistic Map

The logistic map was first used by P. F. Verhulst to describe the evolution of a self-regulating population x with time t (*Verhulst*, 1838). The two parts on the right hand side of the equation

$$x_{t+1} = 4\lambda x_t \cdot (1 - x_t) \tag{2.3}$$

can be interpreted as a population dependent birth-rate factor ($4\lambda x_t$), and a death-rate ($1-x_t$) that accounts for the limited resources and the consequent negative effects of too large a population. The latter factor can be interpreted as the feedback mechanism that makes the system non-linear. If $\lambda \epsilon [0; 1]$ and $x \epsilon [0; 1]$, this function's range is $[0; \lambda]$ and therefore can be iterated ad infinitum without leaving these boundaries. The "control parameter" λ determines the temporal evolution of the system as follows.

For $\lambda < 0.25$, values constantly decrease and tend towards zero. The population dies out.

For $0.25 < \lambda < 0.75$ the system tends towards a single stable state at which $x_{t+1} = x_t = x^*$. This is what one might expect in the case of an idealised population. At some point, the births would exactly equal the deaths and a constant population would be reached (see left panel of figure 2.5)[7]. This point is called an attractor, as the system is drawn to this point in phase space, independent of the seed value x_0.

For values larger than 0.75, the system tends towards periodic cycles. The period of these cycles increases with λ (e.g., see 4-cycle in right panel of figure 2.5), starting with 2, then 4, 6, 8, and so on, until a critical value $\lambda_c \approx 0.8924864...$ is reached. Beyond this point the dynamics become chaotic – the system does not settle down into any kind of recurring cycle, but continues on, in a non-periodic fashion, indefinitely. More precisely, after brief transient behaviour the system is confined to a strange attractor. Such an attractor, while including an infinite number of values, initially (i.e., for $\lambda_c > \lambda > 1$) covers only a limited range on the interval $[0, 1]$ (see figure 2.6). This range of the attractor increases with λ and spans the whole phase space[8] when $\lambda = 1$, which is illustrated with a bifurcation- or Feigenbaum-diagram (*Feigenbaum*, 1978) in figure 2.7. It plots the range of the attractor against the value of λ. The bifurcations and the onset of chaos can be clearly seen.

---

[7] The iteration process is illustrated by the blue line. From the seed point x_0 on the abscissa, the value of the first iteration (x_1) can be found by simply drawing a vertical line until the graph of the function is reached. By drawing a horizontal line from x_1 to the line that bisects the coordinate axes, the value $x_1$is set as the starting point for the next iteration. From there, drawing a vertical line until the graph of the function is reached gives x_2, and so forth.

[8] Actually, it only occupies the interval $]0, 1[$, as the points 0 and 1 are unstable fix-points, so-called "repellors" (*Sweet and Ott*, 2000).

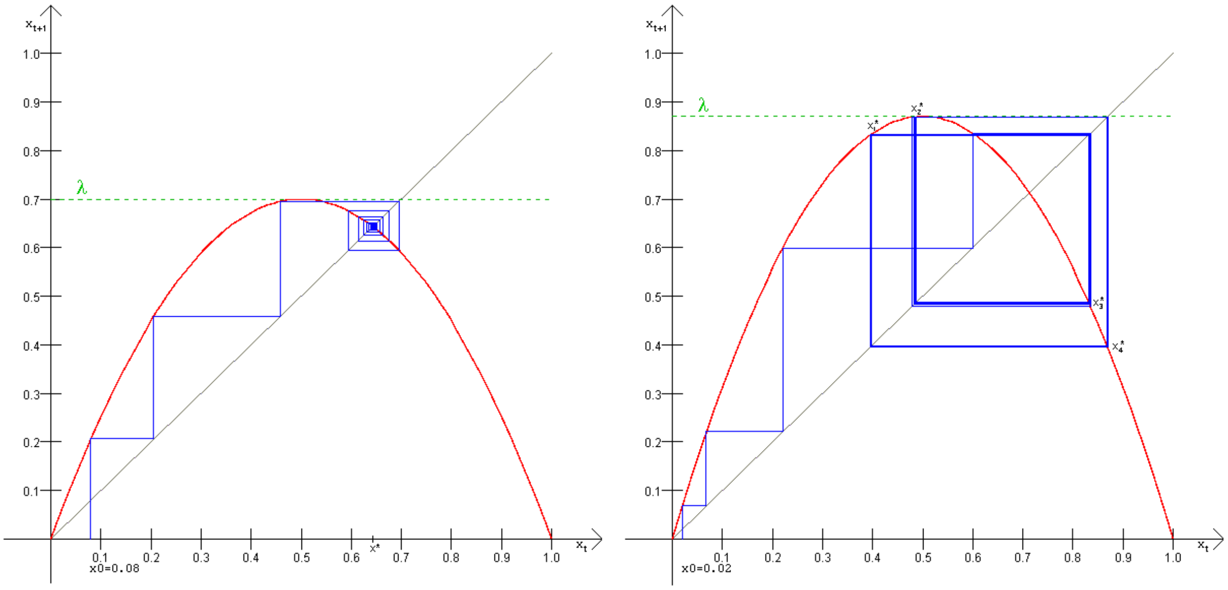

Figure 2.5: Left: iteration of the logistic map (see equation 2.3) with $\lambda = 0.7$. Starting from the seed point $x_0 = 0.08$ the system quickly approaches the attractive fixed point x^*. Right: iteration with $\lambda = 0.87$. The trajectory starting from the seed point $x_0 = 0.02$ this time ends up in an attractive cycle of 4 points x^*_{1-4}.

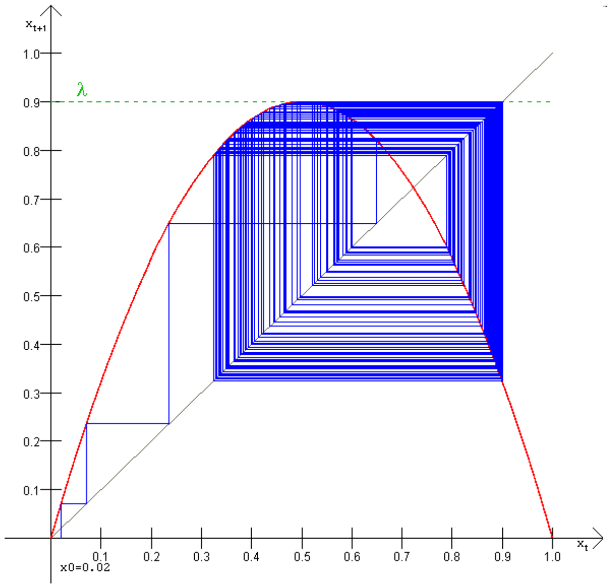

Figure 2.6: Iteration of the logistic map with $\lambda = 0.9$ until x_{150}. After the first few iterations the trajectory is 'caught' in a strange attractor which is limited to a distinct subset of the interval $[0, 1]$, but non-repetitive.

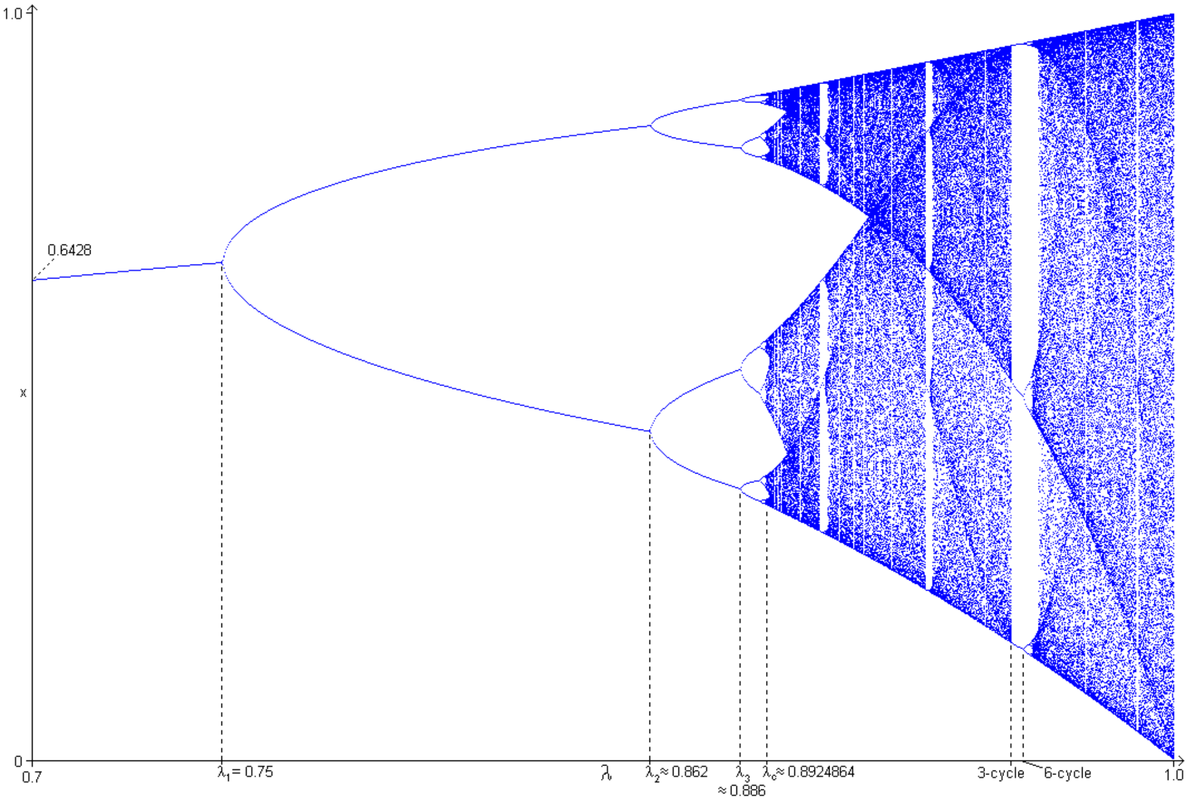

Figure 2.7: The Feigenbaum-diagram of the logistic map for $0.7 \leqq \lambda \leqq 1$ illustrates the transition from simple cyclic to complex chaotic dynamics.

Figure 2.7 also shows that the chaotic regime is interrupted by "windows of order" (the white stripes), in which cycles of higher periods appear (e.g., of period 3, 6, 12, etc. in the largest window), until another critical value is reached and (non-periodic) chaos takes over again.

2.2.2 Chaos and Complexity

The adjustable transition[9] of this system from simple fixed point dynamics, via cycles of increasing period to chaotic strange attractors, illustrates the concept of increasing dynamic complexity. The use of the word *chaotic* in this context is linked with predictability. While the system is fully predictable at first – it always tends towards the fixed point or cyclic attractors, independent of the seed value – the predictability is significantly reduced once λ_c is reached. In the chaotic regime, two trajectories with very similar seed values can diverge within only a few time-steps (see figure 2.8). This sensitive dependence on the initial conditions is a characteristic of chaotic systems. It implies that, given uncertainty about the exact seed value of the system, the only possible prediction that can be made is that the trajectory will end up in that part of the interval $[0, 1]$ that is occupied by the strange attractor. Such behaviour is considered to be *complex*. Therefore, if $\lambda = 1$, the dynamics of the system are essentially unpredictable.

Accordingly, the complexity of a system is linked to its predictability and is often inter-

---

[9] Via the control parameter λ.

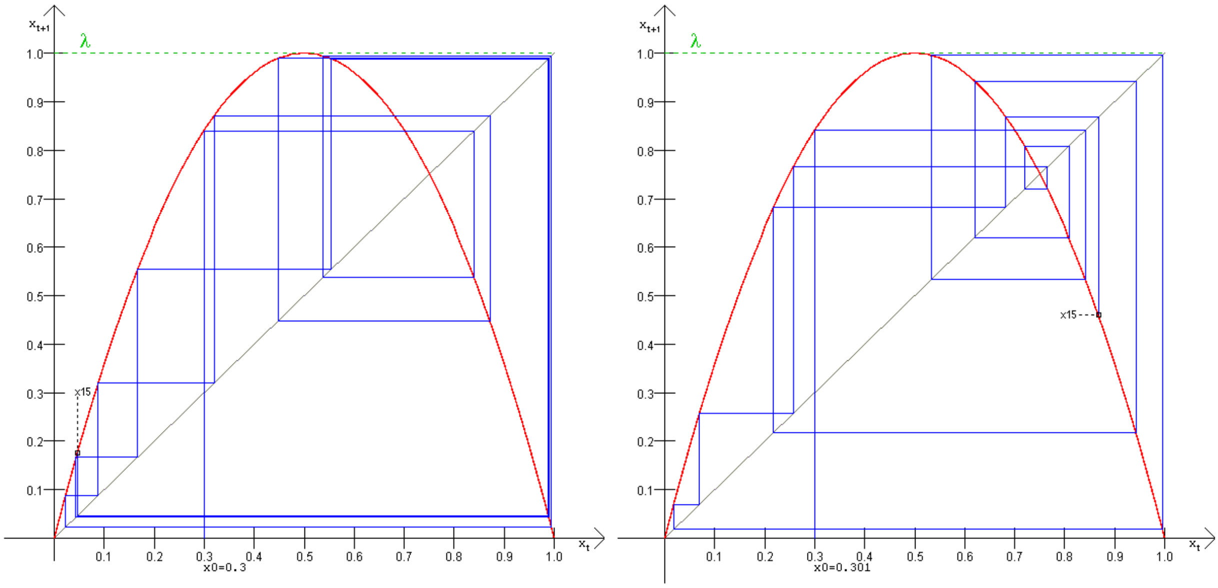

Figure 2.8: Iteration of the logistic map with $\lambda = 1$ until x_{15}, using two slightly different seed values x_0. Left: $x_0 = 0.3$, right: $x_0 = 0.301$. After 15 iterations, the trajectories show no more resemblance. The dynamics are highly unpredictable, i.e., complex.

preted as the degree of randomness of its dynamics. As the dynamics of the logistic map can be easily adjusted by the control parameter λ, it is an optimal system for testing tools that are intended for analysing complex dynamics, i.e., complexity measures such as the ones used in this thesis. There are a large number of complexity measures (different approaches to this quantification), such as Lyapunov exponents, correlation dimension and algorithmic complexity (*Wackerbauer et al.*, 1994), which use different approaches to quantify the behaviour of a system and are usually optimised for use on certain kinds of data. Several of these have been previously utilised for the analysis of atmospheric dynamics (e.g., *Garny et al.*, 2007). This thesis focuses particularly on statistical entropy measures.

2.3 Complexity Measures

2.3.1 Statistical Entropy

While the complexity arising in the dynamics of the logistic map can be quite directly linked to the predictability of its evolution, defining complexity in the general context of complexity measures is not a simple task. The notion of complexity has emerged from work in various fields (*Gallagher and Appenzeller*, 1999) and is therefore probably best defined in a functional way. Similar to the definition of intelligence in psychology: “intelligence is what intelligence tests measure”, complexity is probably best defined as “what complexity measures measure”. There is of course a kind of general consent as to what properties a complexity measure should have. Most complexity measures are in some way linked to predictability, although the precise details of this link can be rather diverse. A large number of complexity measures are limited to values between zero and one. This range suggests a very simple association with probabilities of some sort, which is usually not present. Another common property of complexity measures is that they are relative measures, i.e., they are dimensionless and their range does not depend on the absolute values of the underlying data set. Furthermore, many measures are specifically designed for the analysis of the temporal evolution of a system. None of these properties are mandatory though. Therefore, the two complexity measures discussed in this study will be introduced with a description of how they fit into this rather vague concept of complexity, and their ability to quantify complexity will be illustrated by application to the logistic map process.

Sample entropy (SE) and Rényi entropy (RE), the two measures addressed in this study, are statistical entropy measures. The term statistical entropy originally refers to the use of probability theory in thermodynamics, but is used in a more general sense today. One of the well known measures of statistical entropy is the Boltzmann entropy S:

$$S = -k_B \sum_i p_i \cdot log(p_i) \tag{2.4}$$

When all probabilities are equal $p_i = \frac{1}{\Omega} \, \forall \, i$, this turns into the famous formula on Ludwig Boltzmann's grave[10]:

$$S = k_B \log \Omega \tag{2.5}$$

where k_B is the Boltzmann constant, and Ω is the number of micro-states corresponding to the current macro state (*Beck and Schlögl*, 1993). S from equation 2.5 is also known as thermodynamic entropy (often the natural logarithm is used), as it can be used to describe an isolated system in thermal equilibrium. Changes in the entropy can be directly related to measurable quantities, i.e., if a gas in a cylinder is heated, the change of heat ΔQ is proportional to the change in thermodynamic entropy ΔS:

$$\Delta S = \frac{\Delta Q}{T} \tag{2.6}$$

Hence, the entropy S as defined by equation 2.5 has a clear physical meaning. This is not generally the case for statistical entropy measures.

A formula very similar to equation 2.4 was introduced by C. E. Shannon in the context of information theory (*Shannon*, 1948):

$$H(X) = -\sum_i \left(p\left(x_i\right) \cdot \log_2 p\left(x_i\right)\right) \tag{2.7}$$

where $p\left(x_i\right)$ are the probabilities of all i possible values of the random variable X. Shannon introduced it as a measure of the uncertainty associated with a random variable X or conversely the amount of information conveyed if the value x_i of a random variable becomes known. Due to the similarities between equations 2.7 and 2.4, H is called Shannon entropy. The use of base two for the logarithm is due to the common practise to use bits as units of information. In general, any logarithm can be used. When multiplied by k_B, H is identical to S.

---

[10] Actually it says $S = k. \log W$, but the difference is cosmetic.

H can be further generalised (*Rényi*, 1961) into a family of functions known as Rényi entropy (RE):

$$RE\left(\alpha, X\right) = \frac{1}{1-\alpha} \cdot \log \sum_i p^{\alpha}\left(x_i\right) \tag{2.8}$$

This is one of the metrics utilised in this thesis. If the exponent $\alpha = 1$, equation 2.8 is equal to H. While these information-theoretical entropy measures are all linked, and give similar results if applied to a system in the right manner, there are other definitions of statistical entropy such as approximate entropy or sample entropy, which give rather different results. These alternative definitions are usually associated with different 'questions' with respect to the data set. While RE can be considered a measure of *inhomogeneity* of the probability distribution function (PDF) of a system at a given point in time, other statistical measures, such as the sample entropy discussed in the next section, are designed as measures of the 'inhomogeneity' of a system's evolution with time, i.e., its variability and correlation properties.

A feature of most statistical entropy measures is that they are relative measures. This means that they do not depend on the exact values in a data set. Therefore, statistical measures are often unaffected by problems such as biases within the data. The latter are a particularly frequent problem in climatological data sets. Biases can be introduced to measurements through miscalibrations in the instruments or uncertainties in the retrieval algorithms (e.g., *Thomas*, 2003). Model data on the other hand, is more prone to errors through oversimplifications of the dynamics due to computational limitations, e.g., a cold bias in the stratosphere is observed in many climate models (*Eyring et al.*, 2006). This study applies statistical entropy measures to both model simulation data and measurements of chemical tracers in the atmosphere.

2.3.2 Sample Entropy

Sample entropy (SE) is a statistical measure, proposed by *Richman and Moorman* (2000), which quantifies the variability of a time series by comparing sequences of consecutive data points. It provides a measure of the regularity and hence the predictability of the time series (high sample entropy is related to low predictability / high complexity). Sample entropy is derived from the conditional probability that a sequence of data points stays within a certain tolerance range r for $(m+1)$ steps, given that it has already been for m steps. The tolerance r is commonly measured in units of the standard deviation of the time series. For use on finite data sets, sample entropy is approximated by the statistic $SE(m, r, N)$, which depends on the length of the data series N, the length m of sequences to be compared and the tolerance range r specified.

The SE is a development of the approximate entropy (AE) introduced by *Pincus et al.* (1991). AE was created for the analysis of noisy, and relatively short, physiological time series. It was designed to improve on other more sophisticated measures of a time series' complexity, such as Lyapunov exponents, correlation dimension or the Kolmogorov-Sinai entropy.[11] These other measures require very long or even infinite data lengths to be cal-

---

[11] Kolmogorov-Sinai entropy, or KS entropy, is closely related to Shannon entropy. AE can also be related

culated accurately. Also, they are strongly affected by noise, so low-noise or even noise-free data sets are needed (*Pincus and Kalman*, 2004; *Richman et al.*, 2004). Both requirements are not common traits of real-world time series. Therefore, AE was explicitly designed as a tool to estimate (approximate) statistical entropy from short, noisy data sets. It has since been applied successfully in a variety of fields, such as heart-rate analysis (e.g., *Ryan et al.*, 1994; *Goldberger et al.*, 1994) and quantification of financial market volatility (*Pincus and Kalman*, 2004).

SE improves on several shortcomings of AE, which has a low bias for very short record lengths. This is mainly related to "self-matches" (see below and *Richman and Moorman*, 2000). A more detailed discussion of SE is given in *Richman et al.* (2004). The algorithm for SE also increases calculation speed compared to AE. Both entropy measures quantify the recurrence of similar vectors of m consecutive data points in a time series. Here, "similar" is defined in terms of a tolerance interval r. A high entropy value is associated with a relatively low number of recurrences of similar vectors of length $(m + 1)$ compared to similar vectors of length m. SE is therefore a measure of the variability of a time series, similar to the standard deviation, but with respect to a certain timescale (related to the vector length m). SE has been successfully used in multiple disciplines, e.g., as a tool for early detection of neonatal sepsis from foetal heart-rate data (*Lake et al.*, 2002), as a hacking tool (*Fu et al.*, 2004), and for quantification of monsoon-related climate complexity over Southwest China (*Li et al.*, 2006).

The notion of complexity that the SE measure addresses is that of the 'inhomogeneity' of a process in time, i.e., its variability or inconsistency from one time step to another. It is defined as the negative logarithm of the probability that the process described by the time series stays within a tolerance r for $(m + 1)$ steps if it has already done so for the previous m steps. Therefore, SE is associated with predictability on a discrete timescale of m steps.[12] However, there is an approach called multi-scale entropy (MSE), developed by *Costa et al.* (2002), that uses multiple values of m to analyse signals with variability on several timescales. Usually the time steps are assumed to be identical in length, e.g., a regular sampling interval of measurements, but this is not strictly necessary. A high value of SE means that a process is very irregular and not likely to stay within certain boundaries - given by r - in the next time step, even if it has been for the previous m steps. This also means that the system's evolution is difficult to predict, as the following steps are not in general related to the previous ones.

SE is calculated by comparing vectors of length m, $\mathbf{x}_m$, from the time series and counting matches. A "*match*" is found if the elements of two vectors are all close to each other, "close" meaning that the difference between the elements is less than or equal to the tolerance r. For a time series of N elements, $u(i) : 1 \leq i \leq N$, $(N - m + 1)$ vectors of m consecutive data points, $\mathbf{x}_m(i) = \{u(i + k) : 0 \leq k \leq m - 1\}$, can be formed. A match between two vectors $\mathbf{x}_m(i)$ and $\mathbf{x}_m(j)$ is found, if the maximum absolute element-wise difference between them is smaller than or equal to r, i.e., if

---

to Shannon entropy as has been shown by *Rukhin* (2000).

[12] This also distinguishes it from simpler variability measures, such as the standard deviation, σ, which does not take into account any kind of timescale.

$$max\left[|u(i+k)-u(j+k)|:\ 0\leq k\leq m-1\right]\leq r \tag{2.9}$$

The number of times a match for a template vector $\mathbf{x}_m(i)$ is found in the full time series is counted and the relative frequency denoted $B_i^m(r)$ as it depends on the value of the tolerance r.

$$B_i^m(r)=\frac{1}{N-m-1}\#\left\{j:1\leq j\leq N-m, j\neq i,\ \mathbf{x}_m(i)\ matches\ \mathbf{x}_m(j)\right\} \tag{2.10}$$

Calculating this for all templates and taking the average leads to $B^m(r)$, "*the probability that two sequences will match for m points*" (*Richman and Moorman*, 2000).

$$B^m(r)=\frac{1}{N-m}\sum_{i=1}^{N-m}B_i^m(r) \tag{2.11}$$

The same value is derived for vectors of length $(m+1)$ and the resulting probabilities are $A_i^m(r)$ and $A^m(r)$, i.e.,

$$A_i^m(r)=\frac{1}{N-m-1}\#\left\{j:1\leq j\leq N-m, j\neq i,\ \mathbf{x}_{m+1}(i)\ matches\ \mathbf{x}_{m+1}(j)\right\} \tag{2.12}$$

and

$$A^m(r)=\frac{1}{N-m}\sum_{i=1}^{N-m}A_i^m(r) \tag{2.13}$$

Accordingly, $A^m(r)$ is the probability of a match of sequences of $(m+1)$ points.

Sample entropy is then formally defined as:

$$SampleEntropy(m,r)=\lim_{N\to\infty}\left[-\ln\left(\frac{A^m(r)}{B^m(r)}\right)\right] \tag{2.14}$$

For application to real-world, finite length data sets, this is approximated by the statistic SE, the negative logarithm of $A^m(r)/B^m(r)$.

$$SE(m,r,N)=-\ln\left(\frac{A^m(r)}{B^m(r)}\right) \tag{2.15}$$

Note that the last possible vector of length m, $\mathbf{x}_m(N-m+1)$, is not included in the calculation of $B^m(r)$. This is because there are only $(N-m)$ vectors of length $(m+1)$. The additional vector can therefore only contribute to $B^m(r)$ and could cause a bias, although this effect should be negligible in most cases (*Richman et al.*, 2004). Note also that matches of a vector with itself are not counted (equation 2.10). This is important, as it is one of the advantages of the SE algorithm over AE. Self-matching causes AE to be biased towards higher order (lower entropy), particularly in short time series. The two omissions are the reason for the standardisation factor in equation 2.10 being $(N-m-1)$ instead of $(N-m+1)$.

$A^m(r)$ is the probability to get a match of length $(m+1)$ and $B^m(r)$ is the probability to get a match of length m. As any two vectors that match for $(m+1)$ points must also match for m points, $B^m(r)\geq A^m(r)$, and therefore $SE\geq 0$. It follows that

$$CP=\frac{A^m(r)}{B^m(r)} \tag{2.16}$$

is the conditional probability that a sequence that has been within r for m steps, stays within the tolerance for another step. This illustrates that SE quantifies the variability of a time series and therefore is a measure of the complexity of the underlying dynamics. In order to allow comparison of the values of SE between very different data sets, the tolerance interval r is usually given as a fraction of the standard deviation, σ, of the time series. Thereby, the variability quantified by SE is relative to the overall variability σ of the time series itself. While any value is possible, the maximal useful range of r is between 0.01σ and 1.0σ according to *Richman and Moorman* (2000). The choice of r relative to the standard deviation makes SE a relative measure of complexity. Values of SE calculated for very different time series but with the same parameters m and r, can be used to directly compare the degree of complexity of the data sets.

Obviously, the choice of m and r has a significant impact on the value and meaning of SE calculations. Finding appropriate values of these parameters is crucial to a successful analysis of time series with SE. Therefore Monte-Carlo testing with artificial time series is used in chapter 3, in order to determine the optimal parameter values for analysing atmospheric dynamics with SE.

2.3.3 Rênyi Entropy

While the most common ways of analysing mixing are examination of trends in measurement time series (e.g., *Li et al.*, 2006) and analysis of trajectories of chemical tracers from measurements, or of air parcels from simulations (e.g., *Garny et al.*, 2007; *Haynes and Shuckburgh*, 2000), probability distribution functions (PDFs) of tracer concentrations offer an alternative point of view (*Sparling*, 2000). For this approach, zonal concatenation of data is common practice, i.e., using all data points on a circle of latitude to create a PDF representative of that latitude. Thereby longitudinal resolution is sacrificed in order to obtain sufficient data points for statistical analysis. Rényi entropy is a statistical measure of the homogeneity of a PDF, which can give information about the dynamics of the system that created the PDF.

Rényi entropy was introduced by A. Rényi (1961) as a generalised form of the Shannon-information-entropy H (see also equation 2.7). It can be defined as:

$$RE(\alpha, b, N) = \frac{1}{1-\alpha} \cdot \ln\left(\sum_{i=1}^{b} p_i^{\alpha}(b, N)\right) \tag{2.17}$$

where p_i are the probabilities of certain states i of an evolving system. These probabilities $p_i(b, N)$ are derived from a histogram of N data points using b bins of equal width. α is the exponent of RE. For the values $\alpha = 0$ and $\alpha = 1$ the RE is mathematically equivalent to Hartley entropy and Shannon entropy, respectively. The Hartley entropy is simply the logarithm of the number of bins, $\ln(b)$. It is also the maximum of all REs, independent of α, as RE is a monotonically decreasing function of α (*Beck and Schlögl*, 1993). As the probabilities in equation 2.17 are taken to the power of α, peaks in the distribution are highlighted for larger values of α.

The following examples illustrate how RE is calculated, when it is maximal, and how it can be applied in different simple situations. RE can be viewed as a measure of the

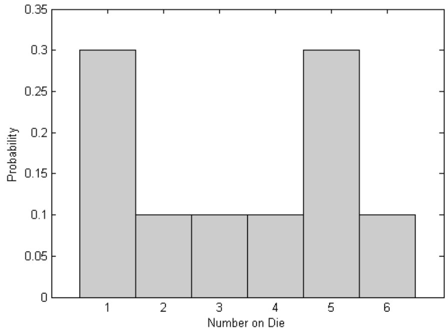

Figure 2.9: A possible PDF of a loaded die.

homogeneity of a PDF. The more homogeneous the PDF of a physical process (homogeneous meaning that the probabilities are evenly distributed), the higher its RE. If all states of a system are equally likely then RE is maximal. Take for example a purely random process such as the outcome of a coin toss. In this case the PDF would have two states ($b = 2$): heads or tails, each with the probability $p_{h/t} = 0.5$. The RE in this case would be:

$$\begin{aligned} RE_{coin} &= \tfrac{1}{1-\alpha} \cdot \ln\left(2 \cdot (0.5)^{\alpha}\right) = \tfrac{1}{1-\alpha} \cdot \ln\left((0.5)^{-1} \cdot (0.5)^{\alpha}\right) \\ &= \tfrac{1}{1-\alpha} \cdot \ln\left(0.5\right)^{\alpha-1} = \tfrac{\alpha-1}{1-\alpha} \cdot \ln\left(0.5\right) = \ln 2 = \ln b \end{aligned} \tag{2.18}$$

Similarly, the RE of an ideal die throw would be $RE_{die} = \ln(b) = \ln 6$. This shows that the absolute value of RE is different for different processes. The fact that $(RE_{coin} = \ln 2) < (RE_{die} = \ln 6)$ reflects that the result of throwing a die is less predictable than that of throwing a coin as there are six possible outcomes instead of two.

Accordingly, a loaded die with an adjusted PDF similar to figure 2.9 has a RE associated with it that is less than $\ln 6$, e.g., for $\alpha = 2$:

$$RE_{loadedDie} = -\ln \sum_{i=1}^{6} p_i^2 = -\ln\left(2 \cdot 0.3^2 + 4 \cdot 0.1^2\right) \approx \ln 4.5 \tag{2.19}$$

The $RE_{loadedDie}$ reflects the change in the dynamics compared to RE_{die}. The outcome of a throw with the loaded die is not pure chance. There is a ‘system’ behind it that is shown in the change of the PDF and can therefore be quantified by RE. On the other hand, $RE_{loadedDie}$ is still significantly larger than RE_{coin}, although the PDF of the coin is certainly more homogeneous, in the sense that the probabilities are more evenly distributed. Furthermore, while an ideal coin is more predictable than an ideal die – there are only two possible outcomes as opposed to six – the underlying dynamics are similar, i.e., completely random. The difference in RE arises due to the different number of bins used for the associated PDFs. Hence, it is useful to standardise the RE by dividing it by its maximum value, $\ln(b)$, thereby obtaining a relative measure. By using equation 2.20, RE can now be utilised for comparing the complexity of very different systems, independent of the way the PDF is created (e.g., the number of bins used).

$$RE(\alpha, b, N) = \frac{1}{\ln(b)} \cdot \frac{1}{1-\alpha} \cdot \ln\left(\sum_{i=1}^{b} p_i^{\alpha}(b, N)\right) \tag{2.20}$$

The previous example also illustrates that, while there is a relationship between the RE and the predictability of a system, RE in equation 2.20 is a complexity measure, not a direct measure of predictability. It is applied to atmospheric tracer data from climate model simulations in chapter 4. The next section introduces the SOCOL model used as a data source.

2.4 The SOCOL Chemistry-Climate Model

The aim of this study is to use the complexity measures introduced in the previous section, to calculate the complexity of the dynamics of tracer distributions in the stratosphere. In turn, the complexity measures will be used to analyse the underlying large-scale mixing processes. The required tracer data, methane is used in this case, are obtained from a control integration of the SOCOL v1.0 chemistry-climate model (CCM), run on the University of Canterbury Super Computer. Models are commonly used in studies of atmospheric dynamics to verify current theories by comparing measurement time series with simulation outcomes and for forecasting. Model data has the advantage that it is more consistent than most measurement data sets. There are no data gaps due to malfunctioning equipment or repositioning of the satellite, which are common in measurements. Also, the number of continuous multi-year global measurement campaigns is very low. As this study uses statistical complexity measures, which require relatively large amounts of data, the more complete and readily available model data is used.

The SOCOL model (SOlar Climate Ozone Links) is a model designed specifically to simulate the dynamics in the middle atmosphere (stratosphere and mesosphere), with particular focus on ozone dynamics. It is a modification of the MA-ECHAM4 (Middle Atmosphere version of the European Centre Hamburg Atmosphere Model 4) general circulation model combined with the MEZON (Model for the Evaluation of oZONe trends) chemistry-transport model (see *Egorova et al.* (2005) and references therein for details).

Derived from the general circulation component MA-ECHAM4, SOCOL has a spectral resolution of T30, with 39 vertical levels spanning the model atmosphere from the surface to 0.01 mBar. The T30 horizontal truncation approximately equates to a spatial resolution of 3.75° in latitude and longitude. The corresponding physical altitude of the 39 vertical levels varies dynamically. The output data used for the analysis is re-gridded to a fixed latitude-longitude-pressure grid, with 48 points in latitude, 96 in longitude, and 30 pressure levels from 1000 mBar to 0.01 mBar.

The MEZON chemistry component of the SOCOL v1.0 simulates 41 chemical species, with 118 gas phase, 33 photolysis, and 16 heterogeneous reactions. The methane data used is given as a mixing ratio in moles of methane per mole of all chemical constituents combined $(\frac{mol}{mol})$.[13] While the calculations for the physical and chemical processes are calculated at 15 minute intervals, the output of the model used in this study is daily scale data. Note that the model year is simplified to 12 months of 30 days each, resulting in a total of only 360 days per year. Therefore, the complete data sets used for all calculations in the following

---

[13] Multiplying this by $1 \cdot 10^6$ gives the mixing ratio in the more commonly used units of parts per million (ppm).

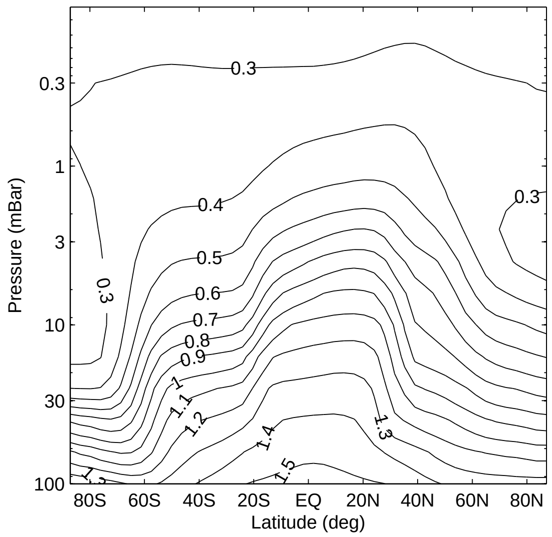

Figure 2.10: Latitude-altitude cross-section of mean methane mixing-ratio (contours in ppm) for SOCOL 1980 October. The vertical range is equal to that of figure 2.4 (100 mBar to 0.1 mBar ≈ 16 km to 60 km).

chapters are on a 360 × 48 × 96 × 30 data point grid.

The model is forced by values of AMIP II sea surface temperatures and sea ice concentration observations for the full simulation period (1980 – 1990).[14] Changes in chemical sources and sinks and the variability in solar forcing are according to E. Rozanov and T. Egorova, PMODWRC (Physikalisch-Meteorologisches Observatorium Davos / World Radiation Center), Davos, March 2003. An 11 year control integration for conditions representative of the 1980 - 1990 period, was performed on the University of Canterbury supercomputer. Initially, only the data for the first two years was available for this study, which is why much of the analyses in chapters 3 and 4 refer to the results from 1980 and 1981. Also, most plots shown in these chapters are for the 10 mBar level, as it lies in the middle of the stratosphere and its dynamics should be representative.

Figure 2.10 shows a contour plot of SOCOL mean methane mixing-ratio (in ppm[15]) in the stratosphere and lower mesosphere (100 to 0.1 mBar ≈ 16 to 60 km) for October 1980. This latitude-altitude cross-section can be compared with the methane measurement data shown in figure 2.4. The methane distribution is qualitatively similar to the measurement data, particularly at lower altitudes, although the SOCOL model seems to have a global low bias in that all contours appear at slightly lower altitudes. Still, the structures are equivalent, and mixing-ratio contours of the same value follow similar paths. Both figures show the tropical maximum leaning towards the Northern Hemisphere with a bump around $30°N$. In the Southern Hemisphere, the mixing ratio is lower on average than in the Northern Hemisphere, and the strongest gradients occur at approximately $20°S$ and $60°S$.The similarities between

---

[14] Atmospheric Model Intercomparison Project data sourced from http://www-pcmdi.llnl.gov/projects/amip/AMIP2EXPDSN/BCS_OBS/amip2_bcs.htm

[15] Note that the contours in figure 2.4 are given in parts per million by volume (ppmv), which is identical to ppm for an ideal gas and can also be considered identical for stratospheric gases, due to the low density of air at high altitudes.

the model data and the measurements suggest that SOCOL methane simulations are realistic and thus potentially useful for further examination.

In the following two chapters, the complexity measures introduced in section 2.3 are discussed in more detail and the results of the application of these measures to the methane data from the SOCOL model runs are shown and analysed.

Chapter 3

Sample Entropy (SE)

Sample entropy (*Richman and Moorman*, 2000), a statistical entropy measure, was introduced in the previous chapter (section 2.3.2) and is based on the approximate entropy (AE) algorithm suggested by *Pincus* (1991). Sample entropy compares sections of time series, or templates, with each other to quantify the variability of the time series with respect to a certain tolerance interval, r, and a timescale determined by the parameter m. For use on finite data sets, sample entropy is approximated by the statistic $SE(m, r, N)$, which depends on the length of the data series N, the length of the templates to be compared, m, and the tolerance range, r, commonly specified in terms of the standard deviation of the time series, σ. It has been previously applied successfully for analysis of heart-rate variability (*Lake et al.*, 2002) and as a measure of monsoon-related climate complexity (*Li et al.*, 2006). In this chapter, SE is applied to atmospheric tracer data from the SOCOL chemistry-climate model (CCM) (see section 2.4) after appropriate values for the parameters N, m, and r are determined.

Atmospheric time series can contain relatively regular signals that are superimposed on highly variable noise. The noise can often be attributed to mixing processes in the atmosphere. A quantitative measure of the variability, or complexity, of atmospheric time series potentially can be used to analyse the underlying regular dynamics and also the mixing that causes the complexity.

Using artificial time series, the sensitivity of the SE to noise and noisy data is tested in section 3.1. The parameters m and r are optimised in order to maximise the distinction between plain noise and noisy signals with relatively few data points N. SE is then tested on the logistic map process. In the second part of the chapter (section 3.2), the determined parameter values are used to apply SE to SOCOL methane data. The results are discussed and problems identified.

3.1 Monte-Carlo Testing with Artificial Time Series

3.1.1 Choice of $m = 2$

The value of m determines the length of the templates to be compared for the calculation of SE (see equations 2.9 to 2.15). It has been suggested by *Richman et al.* (2004) to *"choose m based upon knowledge of the timescale of the underlying process, and r based upon knowledge of the scale of signal noise".* For atmospheric tracer data the variations can be very erratic. While a certain degree of regular variability can be expected, such as daily, seasonal, or annual cycles, the large-scale (i.e., long-term) dynamics are likely to be a highly complex summation of quasi-periodic signals with superimposed noise. Therefore, there is no scale that would be particularly useful to take into account when choosing m. For such cases, where the underlying process is essentially unknown, choosing small values of m and relatively coarse r is advocated by *Pincus and Kalman* (2004) *"to ensure good replicability"*.[1]

Most published studies that utilise AE or SE, use either just $m = 2$ or both $m = 1$ and $m = 2$ for the analysis, and r-values that range between 0.1 and 0.25 times the standard deviation of the time series, (e.g., *Alcaraz and Rieta*, 2007; *Costa et al.*, 2002; *Li et al.*, 2006; *Pincus and Kalman*, 2004.). Unusually, *Lake et al.* (2002) use $m = 3$. As the variability can be on any timescale, a small value of m is more likely to pick up regular variability on several scales. For this study the widely used value $m = 2$ is selected. This effectively assumes that the variability of interest is of the order of a few days. While an approach for using SE on signals with variability on several timescales was developed by *Costa et al.* (2002), this is not an option in this case. It involves severely truncating the data and therefore requires very long data sets, making it impractical for the analysis of one year of daily data from SOCOL (see section 2.4). Considering that the complexity measures that are used here on climate model data should also be applicable to measurement data, choosing large values of m also does not seem practical. High m would mean that relatively long vectors are compared and checked for similarity. Due to noise in any realistic data set, a match becomes less likely the longer the vectors that are compared (as all components have to match individually). Smaller values of m on the other hand, do not give much information about the dynamics of the system. For $m = 0$, the conditional probability in equation 2.16 simply becomes the probability that two data points match by chance, as there is *no* previous condition in this case: $B^0(r) = 1 \; \forall \; r$ as $B_i^0(r) = 1 \; \forall \; i$ (vectors of length zero always match; see also equations 2.11 and 2.10). This could be interpreted as the (conditional) probability of someone making a correct forecast of sun or rain from within a sealed underground bunker, that is, without any information. Similarly, for $m = 1$ the conditional probability is that of predicting tomorrow's weather given today's. While some information about the system can be gained from such a measure, "accidental matches" or "non-matches"[2] can easily lead to wrong conclusions about the complexity (or regularity) of the underlying dynamics. Particularly when using relatively short time series (small N) the number of accidental matches and non-matches will not average out. Therefore, $m = 2$ is taken as the smallest value for which SE calculations are likely to be robust measures of complexity for most data sets. As this choice can be considered rather arbitrary, results of using different values of m are also shown and discussed in the following section.

---

[1] While the paper actually refers to approximate entropy (AE), this rule can be applied to SE as well, as the definition of a match is the same for both measures. They only differ in the way the matches are added up. The parameters only affect detection of matches as such, not the way they are counted.

[2] That is, the matches that are not warranted by the underlying dynamics (accidental matches) or the cases where two vectors really are similar, but do not match due to noise in the data (accidental non-matches).

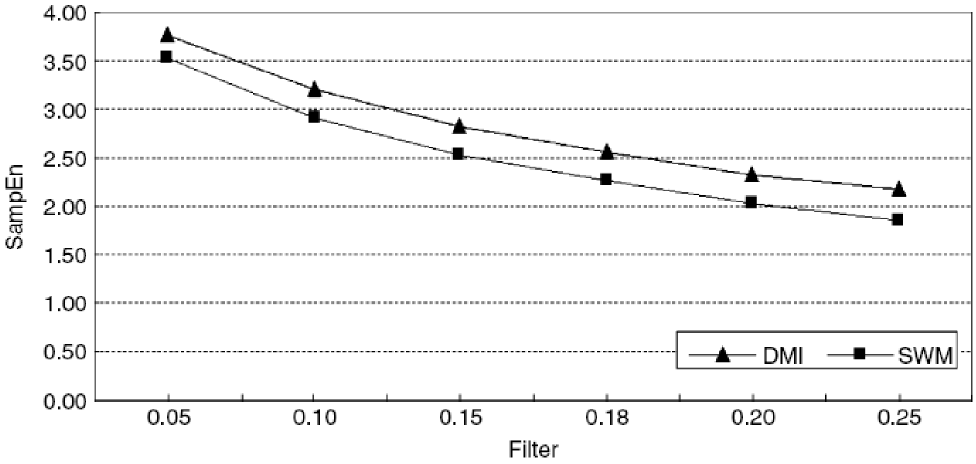

Figure 3.1: Figure 6 from *Li et al.*, 2006. Comparison of results of calculating SE with different values of the parameter r. The data sets used are temperature measurements from a region where different air masses interact (DMI) and a region that is mainly under Southwest Monsoon control (SWM).

3.1.2 Determination of an Optimal Value for r and Minimal Data Length N

For fixed m, the other two parameters r and N must be optimised. The minimum number of data points N required to distinguish between time series of different complexity is potentially dependent on the tolerance interval r. The previously mentioned studies have typically used values between 0.1 and 0.25 times the standard deviation (σ). The reasons given for this choice are diverse and usually not universally applicable. The study of temperature data by *Li et al.* (2006) uses $SE(m = 2)$ to quantify climate complexity on a local scale. From figure 3.1 and by "*analysing finer patterns*" in their results, they claim that $r = 0.18\sigma$ is optimal, although they seem to use $r = 0.15\sigma$ throughout the rest of the paper. However, their reasons for this claim are rather inconclusive, particularly considering the dubious scale of the x-axis (apparently linear with an increment of 0.025 per tick except between 0.15 and 0.20, where there are three ticks with 0.18 at the center) in figure 3.1, which throws doubt on their choice. Aiming to improve this reasoning, Monte-Carlo tests are devised for determining appropriate values of r for this study.

The aim of this study is to determine the usefulness of complexity measures such as SE for analysing stratospheric dynamics from tracer data. Therefore the parameter r is optimised with respect to SE's ability to distinguish between complex patterns and different types of noise. A simple example of a very complex but non-random time series is a sine superimposed with uniformly or Gaussian distributed random numbers (noise), sampled at a regular interval that is not proportional to the frequencies. For testing purposes eight different artificial time series are created. White noise, autocorrelated noise, three sines of different frequency with superimposed white noise, and the same sines with autocorrelated noise superimposed. The white noise simply consists of uniformly distributed random numbers between zero and one. The autocorrelated noise is created by averaging over three data points of the white noise. For the noisy signals a normal sine wave (values between -1 and

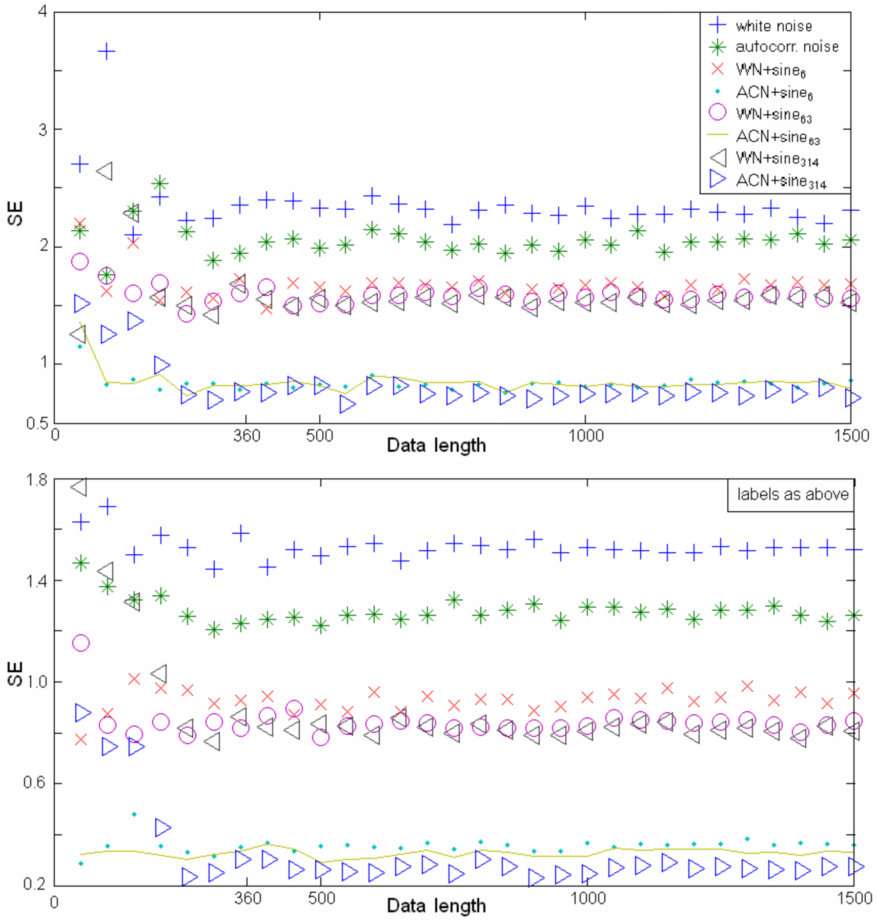

Figure 3.2: Dependence of $SE(m=2)$ of plain noise and noisy sines on data length N for $r = 0.18\sigma$ (top panel) and $r = 0.4\sigma$ (bottom panel). The signal-to-noise ratio for the noisy sines is 2. Given a sufficient number of data points, SE clearly distinguishes between white noise, autocorrelated noise, sines with white noise ($WN + sine$), and sines with autocorrelated noise ($ACN + sine$) superimposed. The different wavelengths of the sines used are not resolved.

1) is created and either white noise ($WN + sine$) or autocorrelated noise ($ACN + sine$) is added to it. The sampling is kept constant but non-periodic by using the integers from 1 to N as values in radians and calculating the sine of these.[3] Lower frequencies are achieved by dividing the integers by 10 and 50.[4] Thus, the three different signals have a period close (but not equal) to 6, 63 and 314 time steps (i.e., data points) and are labelled $sine_6$, $sine_{63}$, and $sine_{314}$, respectively. The signal-to-noise ratio for the artificial signals created in this way is two.

Figure 3.2 shows the dependence of SE on data length N for eight different artificial time series, using $r = 0.18\sigma$ (top) and $r = 0.4\sigma$ (bottom). Comparison of the two diagrams shows that the absolute value of SE, the relative separation of the signals, and therefore the sensitivity to changes in the dynamics of the time series, strongly depend on r. In both cases SE successfully distinguishes between noise and noisy signals, given sufficient data

---

[3] This means, the first value of the sine signal is $\sin(1rad) \approx 0.666$, the second value is $\sin(2rad) \approx 0.909$, the third value is $\sin(3rad) \approx 0.141$, and so forth.

[4] For example, in the latter case the first point of the sine signal is $\sin(\frac{1}{50}rad) \approx 0.020$, the second value is $\sin(\frac{2}{50}rad) \approx 0.040$, etc.

points, but convergence to a relatively stable value occurs faster (smaller N) for $r = 0.4\sigma$ (see also figures 3.3 and 3.4). While $N \geq 500$ data points is found to be optimal for clearly separating signals of different complexity, reasonable differentiation can be achieved even with 365 data points (e.g., daily data for one year). It is interesting to note that SE converges to almost equal values for signals with the same type of added noise but different frequencies. It follows, that SE cannot be used to distinguish between signals of similar form but different frequency. This is not unexpected, a signal which contains noise and has a certain superimposed periodicity is no more complex than the same quantity of data but with a different periodicity – assuming both periods are sampled adequately. The latter condition is likely the reason for the slightly higher value of SE for the $sine_6$ series compared to the ones with lower frequency. In general six data points per period cannot be considered an adequate sampling for a sine, though the sampling does meet the Nyquist criterion.

In order to give some additional justification for the choice of $m = 2$ and to further justify the value of $r = 0.4$ (the "σ" is omitted from this point on for simplicity), which is larger than the value used in most other studies, the dependence of SE on data length, N, shown in figure 3.2 is calculated for several combinations of m and r, and also for a higher signal-to-noise ratio.

Figure 3.3 shows the same plots as figure 3.2, but without the lower frequency signals, for $m = 2$, $m = 3$, and $m = 5$ (rows, top to bottom), combined with $r = 0.18$ (left column) and $r = 0.4$ (right column). By comparing rows (the results that use the same r but different m), it can be seen that the absolute value of SE is not significantly affected by the choice of m. The values SE converges to for large N are (almost) independent of m. On the other hand, for higher values of m, SE seems to converge slightly faster to its large-N-value for the $ACN + sine$ series, while the opposite is true for most other cases.

In contrast, comparing columns (same m, different r) shows again that the value of r has a significant impact on the absolute value of SE. Larger r leads to lower values of SE, because a large value of r (larger tolerance interval) will allow more matches to be found, indicating simpler dynamics. Figure 3.3 also confirms the findings from figure 3.2, that a higher r-value leads to faster convergence and reduced variability of the value of SE, for higher m.

As the signal-to-noise ratio (SNR) of the artificial signals can be considered relatively low, the same tests are repeated for time series with a higher SNR. Therefore, the absolute values of the sines used are multiplied by two, while the noise is left unchanged. The new time series has a SNR of eight. Figure 3.4 shows the application of SE to the higher SNR data.

For the higher SNR, the general behaviour of SE is essentially the same, only that the absolute value of SE is decreased and that the large-N-value is approached faster (for smaller N). For example, $SE(m = 2, r = 0.4, N = 1500) \approx 0.45$ for $WN + sine$, while it is approximately 0.95 when the $SNR = 2$. Also the fluctuations around this value are larger for $N = 500$ when the $SNR = 2$, than when the $SNR = 8$. The faster convergence is a straight forward effect of the (relative) reduction of noise (less noise in the data => less noise in SE). Similarly, the lower absolute values can be explained by the higher SNR. The relative reduction of the high frequency variability of the time series means that the simple periodic sine wave dominates and therefore the signal is found to be less complex. This also

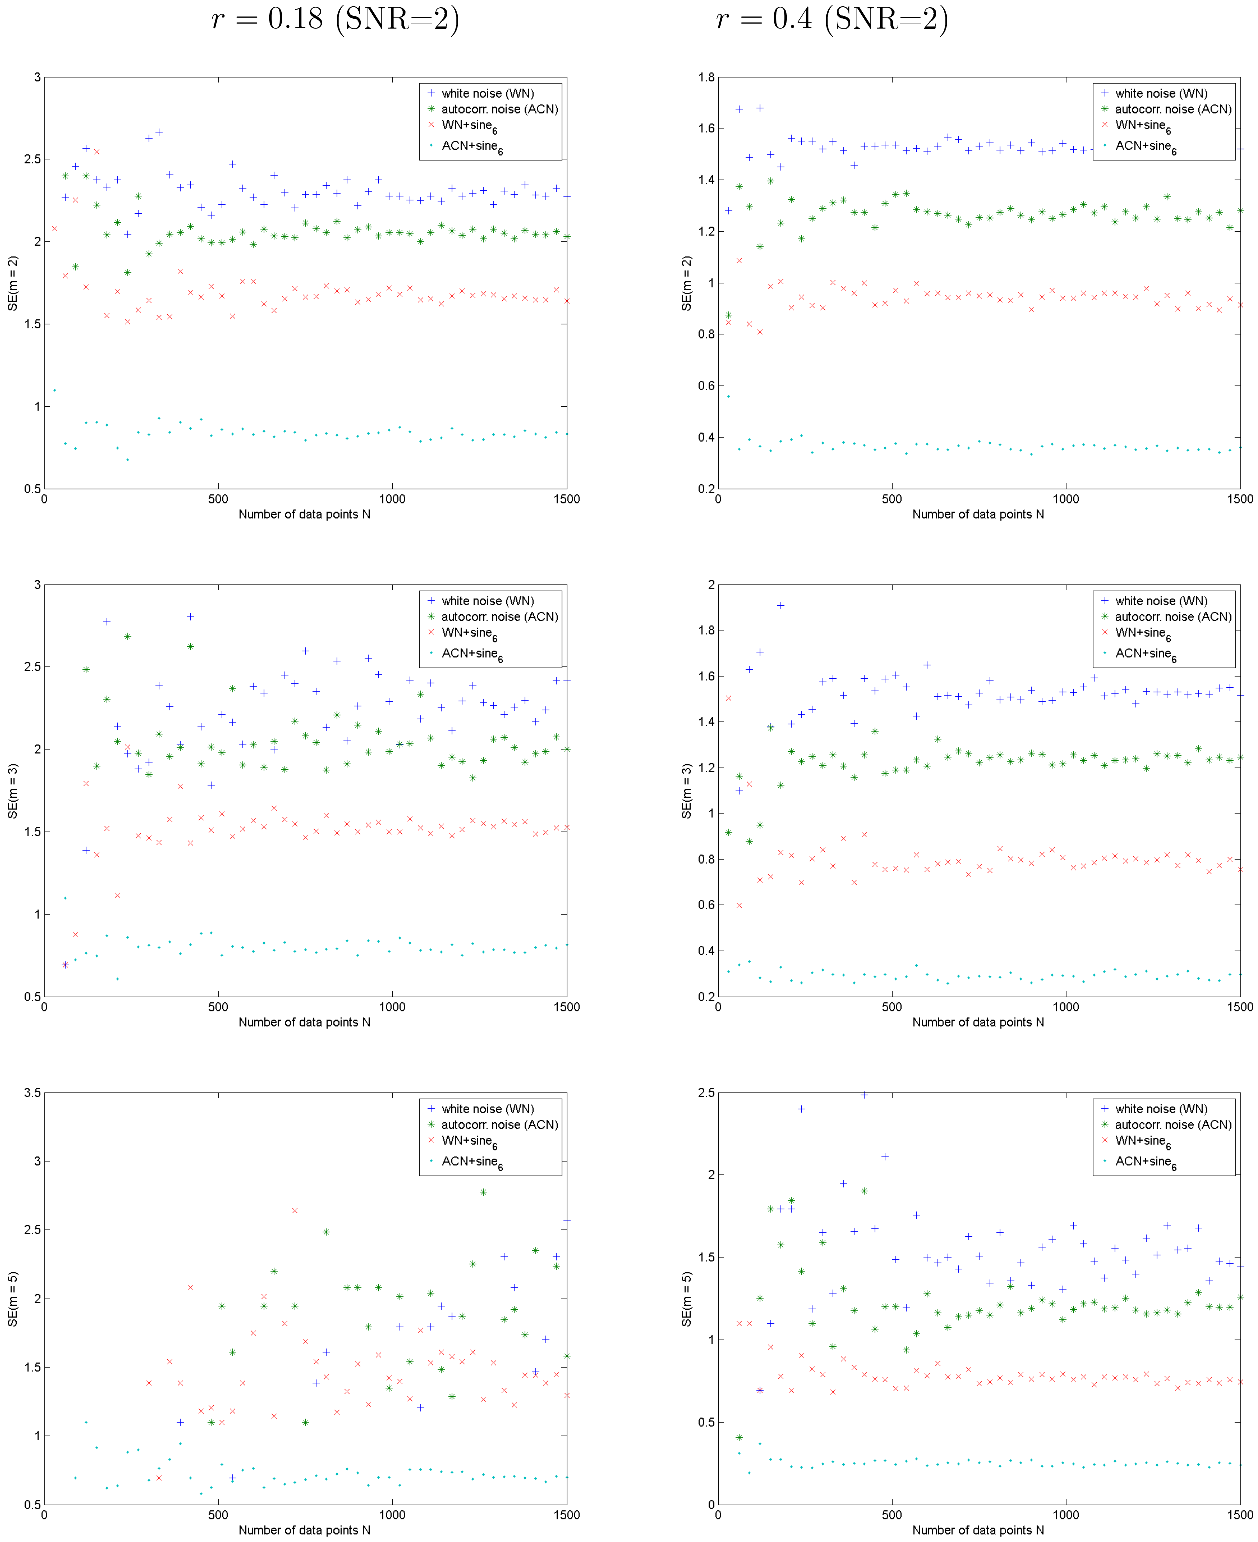

Figure 3.3: SE parameter tests for $r = 0.18$ (left column) and $r = 0.4$ (right column). From top to bottom $m = 2$, $m = 3$, and $m = 5$.

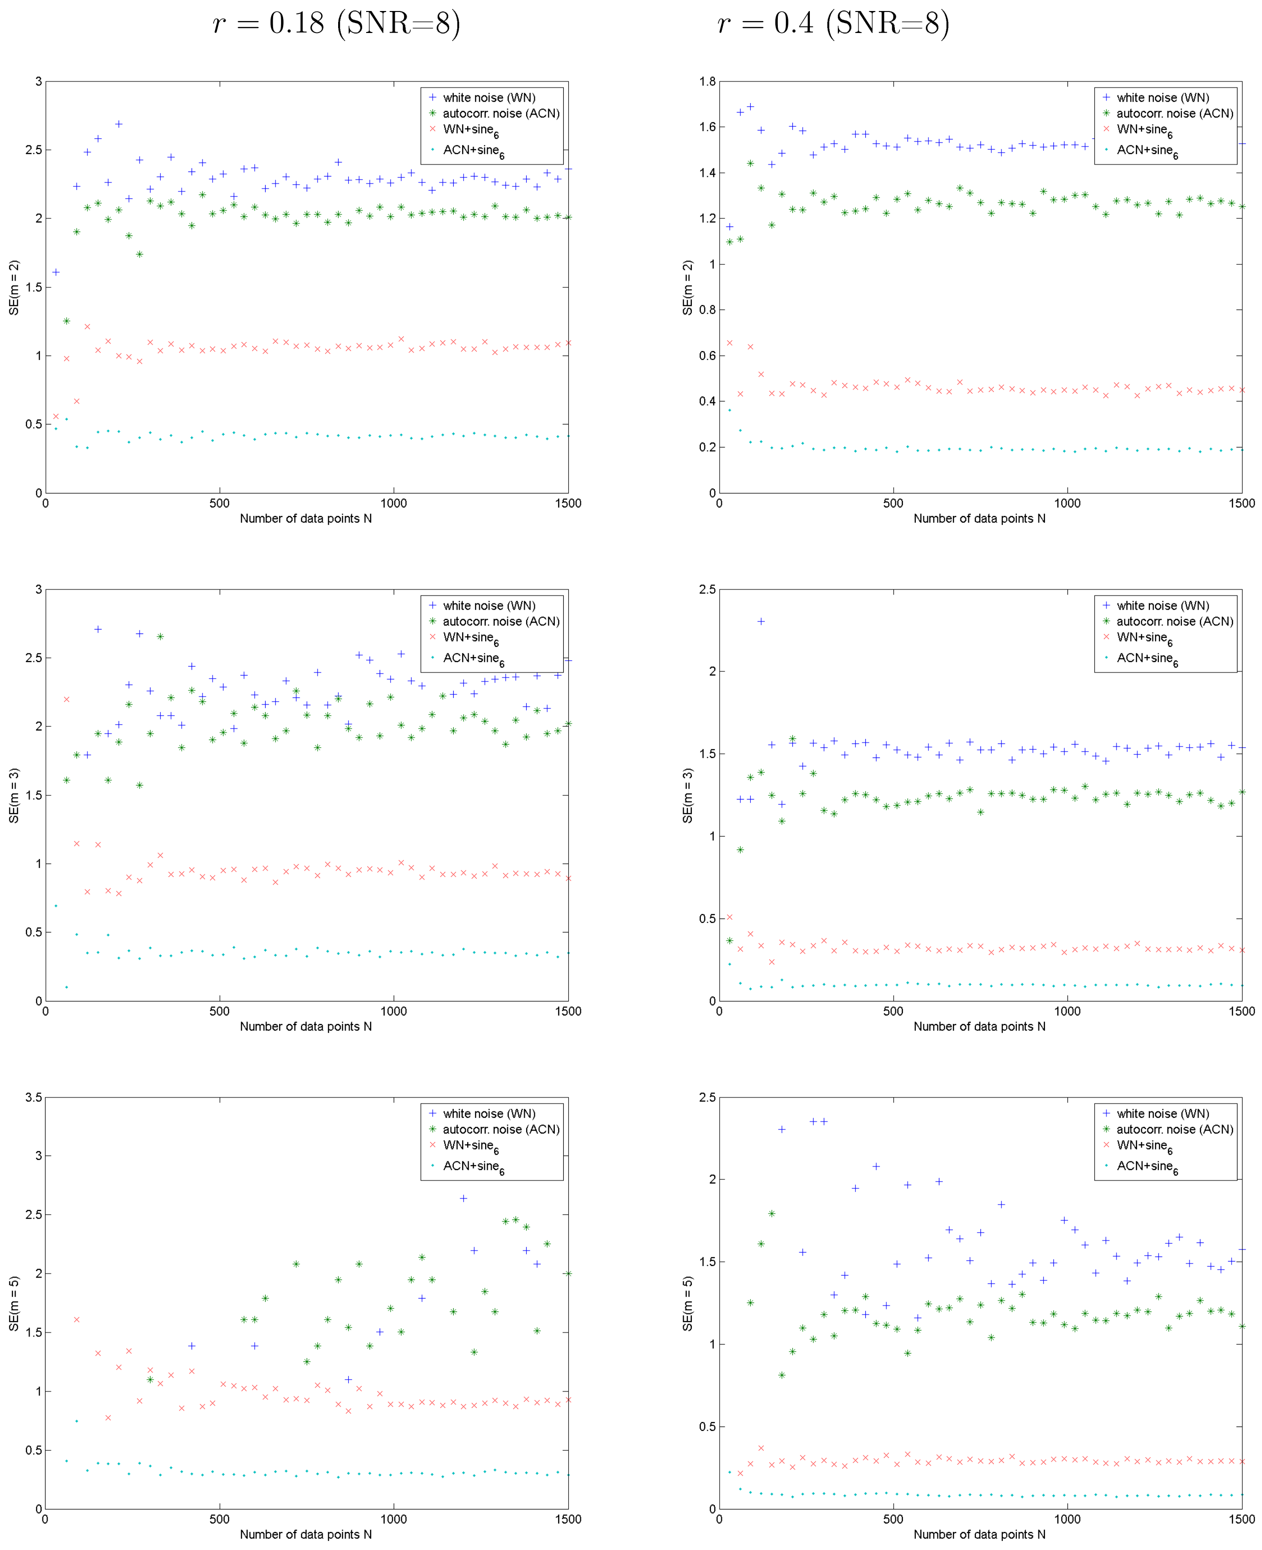

$r = 0.18$ (SNR=8)

$r = 0.4$ (SNR=8)

Figure 3.4: Same as figure 3.3, but for a $SNR = 8$.

implies that the signals with higher noise levels are not distinguished as well as before. On the other hand, the distinction between noisy signals and plain noise is improved. With a higher SNR, the SE of the signals goes down, while the SE of the noise remains the same. Accordingly, SE can also be interpreted as a signal-to-noise measure.

Unfortunately, the previous tests do not allow any decisions about the optimal value of either m or r. The differences between the behaviour of SE with different values of m are too subtle to be conclusive. One problem is that the range of values of SE continuously changes with changing tolerance r, as does its variability. Because of this, it is rather difficult to judge for which r-value the separation of the different time series is optimal relative to the variability of the SE value itself. To find the best value of r, a measure is needed that quantifies the separation of the different signals identified by SE, relative to the uncertainty (or variability) of the SE values. For this purpose, two relative differentiation metrics, $\triangle_1$ and $\triangle_2$, are introduced:

$$\triangle_1(S1, S2) = \frac{\overline{SE(S1)} - \overline{SE(S2)}}{\sqrt{\sigma^2_{SE(S1)} + \sigma^2_{SE(S2)}}} \tag{3.1}$$

and

$$\triangle_2(S1, S2) = \left(\overline{SE(S1)} - \sigma_{SE(S1)}\right) - \left(\overline{SE(S2)} + \sigma_{SE(S2)}\right). \tag{3.2}$$

In both cases it is assumed that the time series $S1$ is more complex than $S2$, i.e., $\overline{SE(S1)} > \overline{SE(S2)}$. $\triangle_1(S1, S2)$ is the difference of the mean value of SE of two time series (denoted $S1$ and $S2$), relative to the combined standard deviation (or standard error) of the values of SE.[5] $\triangle_2(S1, S2)$ is the distance between the "lower standard deviation" of the mean SE value of the more complex time series and the "upper standard deviation" of the mean SE of the less complex one. For both metrics a higher value indicates a better discrimination of the two time series.

The mean value of SE of a time series S, $\overline{SE(S)}$, and the standard deviation of this value, $\sigma_{SE(s)}$, are calculated by adding 200 different uniformly distributed noise series to the sine signals, calculating SE for each case and then taking the mean and standard deviation of the derived SE. As the value of SE is dependent on m, r, and N, the value of the $\triangle$s is implicitly dependent on these as well. The results for different values of r and m are shown in figure 3.5. All calculations were completed using time series of length $N = 360$, as this is the data length that SE will be applied to (one year of daily SOCOL output, see section 2.4). In all plots in figure 3.5, the green continuous line is the value of $\triangle_i$ for white noise and autocorrelated noise, the red dashed line is $\triangle_i$ for autocorrelated noise and a sine of period six with white noise superimposed ($WN + sine_6$), and the cyan line with dots is $\triangle_i$ for $WN + sine_6$ and a sine with autocorrelated noise superimposed ($ACN + sine_6$). The blue crosses are the sums of the three values of $\triangle_i$ (denoted $Total\triangle_1$ and $Total\triangle_2$), which are used as an indicator of the overall ability of SE to distinguish between time series of different complexity.

---

[5] This is similar to a Student's t-test.

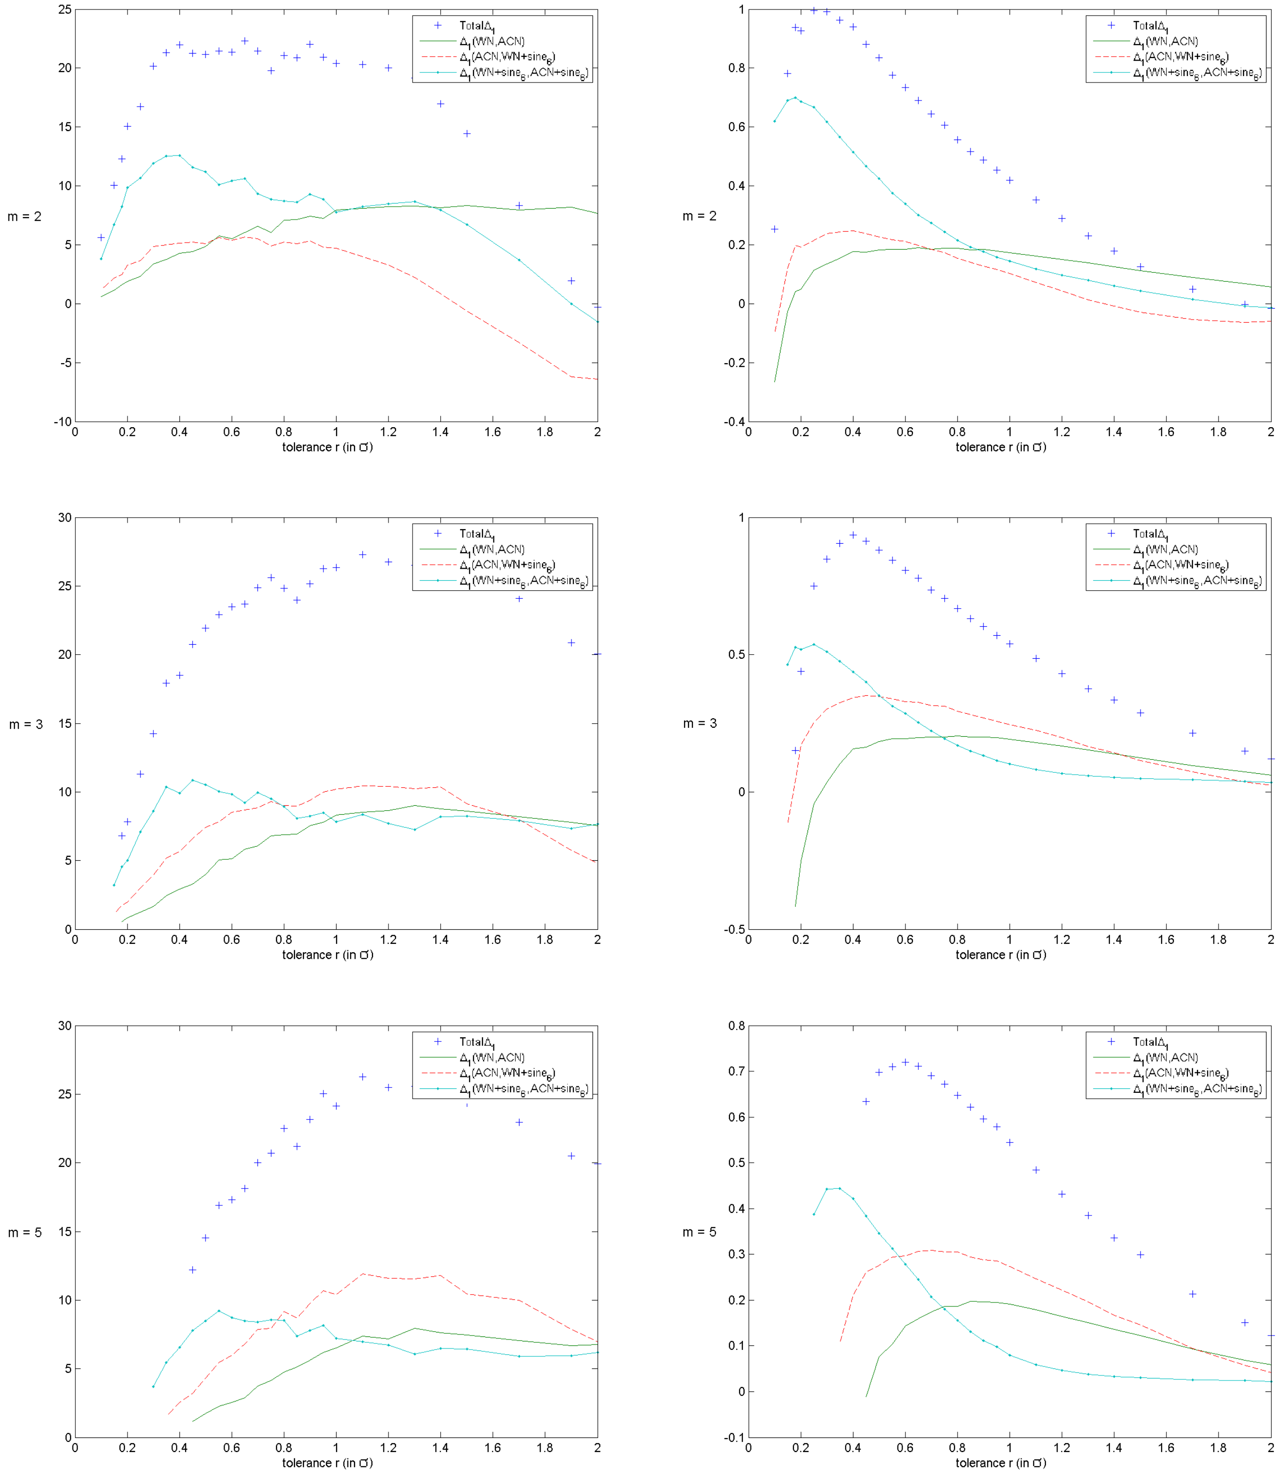

Figure 3.5: SE-$\triangle_i$ signal differentiation tests. $N = 360$, Left column: $\triangle_1$. Right column: $\triangle_2$. Top row: $m = 2$. Middle row: $m = 3$. Bottom row: $m = 5$. While $m = 2$ and $m = 3$ have similar differentiation capabilities, $m = 5$ is not a useful parameter as rather high r-values are required (some $\triangle_i$s are not defined for $r < 0.45$).

The top left panel in figure 3.5 shows that the total differentiation capability, $Total\triangle_1$, of SE for $m = 2$ reaches its maximum for about $r = 0.4$ and stays close to this maximum until about $r = 1$. For $m = 3$ the maximum of $Total\triangle_1$ is reached at $r = 1.1$, and relatively high values begin to occur for $r \geq 0.6$ (see middle left panel in figure 3.5). For $m = 5$ (bottom row) some $\triangle_i$s are not defined for values smaller than $r = 0.45$ and the $Total\triangle_1$ indicates that a value of $r = 1.1$ or larger is optimal. This seems to suggest that a wider tolerance is better for use with higher m-values. Very large tolerance intervals are unlikely to be useful since they can cause ubiquitous matches in the entropy calculations that are not warranted by any kind of order in the series. This indicates that $m = 5$ or larger is not a useful choice for calculating SE when using relatively short data sets. It supports the choice of $m = 2$ for this study.

In both cases, the contributions to $Total\triangle_1$ change with r. For low values of r the ability of SE to differentiate between $WN + sine_6$ and $ACN + sine_6$ contributes most strongly to $Total\triangle_1$. $\triangle_1(WN+sine_6, ACN+sine_6)$ starts decreasing at around $r = 0.4$ for both values of m, while the relative differentiation between the other pairs of time series still improves for increasing r. For $m = 2$, values of r between 0.4 and 0.6 offer the best compromise between the different $\triangle_1$s, while values between 0.6 and 1.0 seem best for $m = 3$. As explained before, high values of r can be problematic and therefore the analysis with $\triangle_1$ suggests that $m = 2$ combined with $r = 0.4$ is the best choice to optimise the signal-separation capabilities of SE based on these representative time series.

However, this selection is not confirmed from the plots of $\triangle_2$ (right column of figure 3.5). $Total\triangle_2$ peaks at $r = 0.4$ for $m = 3$, while it peaks at $r = 0.25$ for $m = 2$. In the latter case the peak is relatively broad and $Total\triangle_2$ is only slightly reduced for $r = 0.4$. Similar to $Total\triangle_1$, $\triangle_2(WN+sine_6, ACN+sine_6)$ contributes most strongly to $Total\triangle_2$ for low values of r. For $m = 2$, most of the curve's peak at $r = 0.2$ comes from the separation of the two signals, but SE's ability to distinguish between a signal and noise, $\triangle_2(ACN, WN + sine_6)$, is somewhat improved for $r = 0.4$. In general though, it can be argued from the plots of $\triangle_2$ that the combination of $m = 3$ and $r = 0.4$ is optimal. Compared to $m = 2$, the contributions to $Total\triangle_2$ are more evenly spread and therefore allow for better discrimination of a wider range of different signals.

While high values of m can be ruled out, $m = 2$ is not special and it is shown that $m = 3$ could also be a valid choice. Although other parameter combinations might be equally useful, $m = 2$ and $r = 0.4$ are selected as the best compromise for distinguishing between all the different time series tested. This choice takes into account the absolute values of SE for different time series, their variability, and the length of the data series that SE will be used on (daily data points for one year).

Unfortunately, the reasoning for the choice of parameters is rather unsatisfactory and it has become clear that the usefulness of some of the parameters chosen depends on the signal the SE is calculated for. Therefore, SE is applied to a (mathematically ideal) complex system in the next section, in order to get a more practical understanding of how SE quantifies complexity.

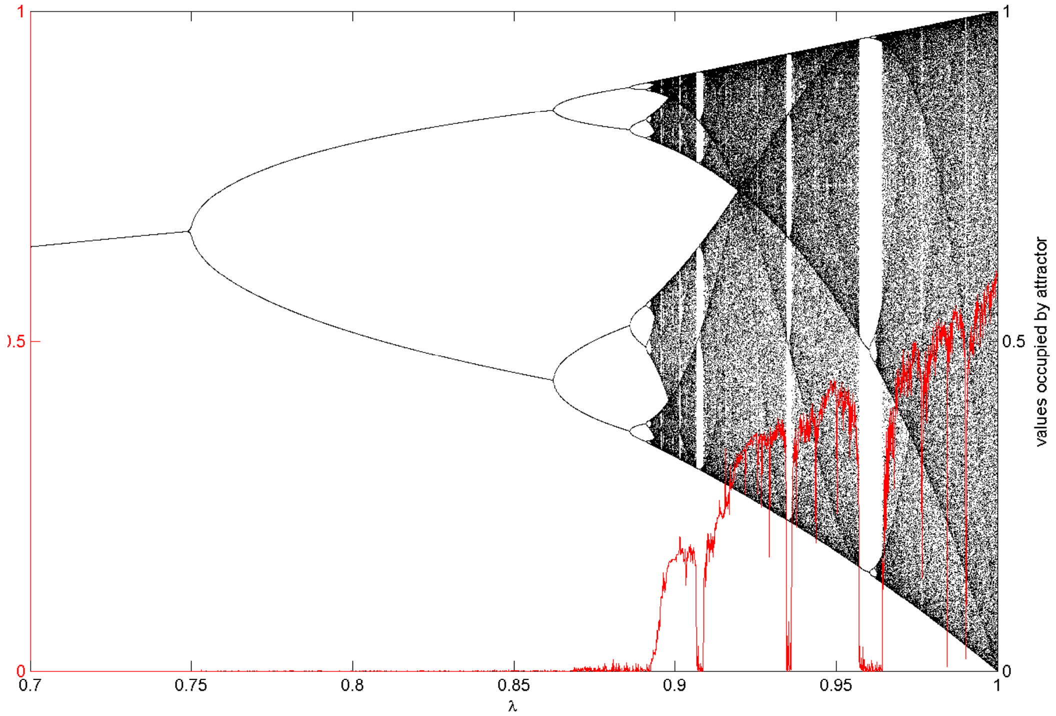

Figure 3.6: $SE(2, 0.4, 1000)$ of the logistic map for increasing values of λ superimposed on the corresponding Feigenbaum diagram.

3.1.3 Application of SE to the Logistic Map

Using the parameters determined in the previous section, SE is applied to the logistic map (see equation 2.3) in order to test its quality as a measure of complexity on an ideal complex system. Figure 3.6 shows the SE of the logistic map for $\lambda \,\epsilon\, [0.7; 1]$ superimposed on the corresponding Feigenbaum diagram. $SE(m = 2, r = 0.4)$ is calculated using the parameter values determined previously for time series consisting of 1000 iterations of equation 2.3 from random seed points.

For $\lambda < \lambda_c$, values of SE are very close to zero. This means that, in an ideal (noise-free) deterministic system, SE does not distinguish between fixed point attractors and regular cycles of higher period. This is in accordance with the findings of *Pincus* (1991) who tested AE on the logistic map and reports $AE = 0$ for a period-four cycle. It also confirms the results of section 3.1.2 that SE has almost identical values for noisy sine signals of different frequency.

The values of SE become slightly larger than zero for values of λ close to (but smaller than) λ_c. This is probably due to the fact that in this region convergence to the attractor cycles is relatively slow and therefore the time series are found to be slightly more complex. This can be seen by neglecting the transient phase from the calculations of SE. When omitting the first 500 data points and calculating SE for the following 1000 time steps, values appreciably larger than zero, e.g., greater than 0.01, do not occur until $\lambda > \lambda_c$. This is shown in figure 3.7. Thus, SE can be used to detect the onset of chaos in the logistic map

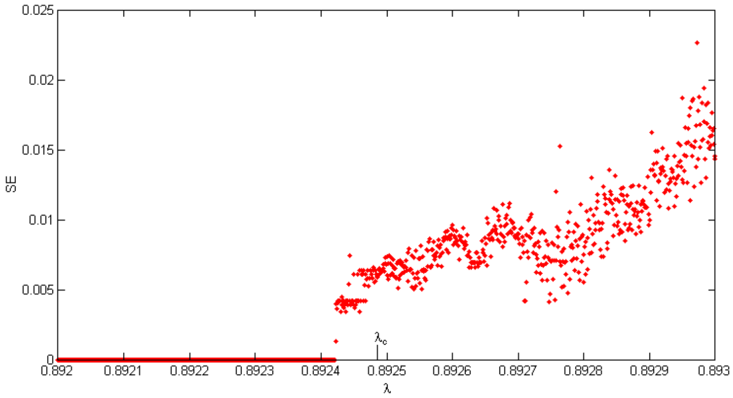

Figure 3.7: Zoom into the region close to $\lambda_c = 0.8924864...$ of figure 3.6. $SE(2, 0.4, 500)$ is calculated neglecting the first 500 iteration steps in order to get rid of most of the initial transient behaviour of the system in this region. SE detects the onset of chaos at λ_c by a sudden increase in its value.

system. When $\lambda = \lambda_c$ the value of SE becomes non-zero, mirroring the sudden increase in system complexity.[6] As explained in section 2.2.2, the essential change occurring at λ_c, is a loss of predictability for similar initial conditions. In these terms, SE qualifies as a measure of the (un-)predictability for noise-free data sets of deterministic systems.

For $\lambda > \lambda_c$, SE increases with λ, and therefore with increasing complexity, while dropping back to very small values in the sections where the logistic map system exhibits periodic behaviour again, the “windows of order”. The maximum value of SE, reached for $\lambda = 1$, is approximately 0.625. It is interesting to note that this value is significantly smaller than the value of SE for white noise, which is approximately 1.5 (see figure 3.2). This shows that, while the evolution of the logistic map is essentially chaotic for $\lambda = 1$ (as explained in section 2.2.2), it is not completely random and SE reflects this. This also suggests that, when using SE as a complexity measure of time series, noisy signals may be difficult to separate from chaotic dynamical systems. This is a rather undesirable property of SE, particularly attempting to interpret SE in terms of predictability, as fairly accurate long-term predictions might be possible for a noisy, but regular signal, while they are impossible in a chaotic system.

However, the results show that SE generally detects complex dynamics, increases with increasing complexity, and is sensitive enough to distinguish between deterministic chaos and statistical randomness. Therefore it is a valid and potentially useful complexity measure, and can be applied to stratospheric tracer data for analysing atmospheric dynamics. This is tested in the next section.

---

[6] Note that this is only the case when omitting the first 500 data points from the SE calculation. This removes the transient behaviour of the logistic map iteration which otherwise causes small, but non-zero values of SE much earlier (see figure 3.6). Close investigation of figure 3.7 also shows that, even when removing the first 500 data points, the increase in SE occurs slightly before λ reaches λ_c. As the λ-value, at which the SE has its first non-zero value, moves closer to λ_c the more ‘early’ data points are omitted (not shown), the premature increase in SE is probably related to even longer transient behaviour when λ is very close to λ_c, or to round-off errors in the calculations.

3.2 Application of the Sample Entropy Metric to SOCOL Tracer Data

3.2.1 Tropical Pipe and Surf-Zone

SE has previously been used in a geophysical context by *Li et al.* (2006) to analyse tropospheric daily temperature measurements in Southwest China and used to identify regions associated with climate conditions of different complexity. They successfully identified two distinct regions of high and lower SE in Southwest China. The higher values are attributed to the interactions of different air masses. While the region associated with SE is under the influence of at least two air masses at any time of the year (Northwest and Northeast Monsoon in winter, Southwest and Southeast Monsoon in summer; *Li et al.*, 2006), the area of lower SE is mainly affected by the Southwest Monsoon throughout the whole year. Interactions between air masses seem to lead to strong variations in temperature. Such interactions of the air masses are likely to be a highly non-linear combination of vertical and horizontal eddy-mixing, and therefore the increase in complexity reflected by SE is a reasonable result.

This study attempts to apply SE to stratospheric tracer data on a global scale in order to analyse large-scale complexity patterns and to test the ability of this measure to identify mixed regions and transport barriers (complex tracer time series are likely to be caused by mixing processes). Methane is used because it is a representative long-lived tracer (*Randel et al.*, 1998) and its stratospheric distribution is therefore mainly affected by large-scale mixing and transport. This application of complexity measures to tracer data is potentially a novel way to identify and quantify mixing in the middle atmosphere. Using the parameters determined previously, SE is calculated from methane (CH_4) data from a control integration of the SOCOL chemistry-climate model (CCM) (see section 2.4).

Figure 3.8 shows the results of calculating $SE(m = 2, r = 0.4, N = 360)$ of the annual global methane field at a pressure of 10 mBar ($\approx$ 31 km altitude) from a SOCOL simulation representative of 1980 conditions. SE is calculated at each model grid point (48 points in latitude and 96 in longitude) for each pressure level, using the daily data points of one year (360 data points, see section 2.4 for details). Several consistent patterns can be observed. A zonal region of high SE values ($0.6 - 1.0$) from the equator to approximately $\pm 10^\circ$ latitude followed by a similar region of low SE ($0 - 0.3$) can be seen in both hemispheres. In the Northern Hemisphere the low values extend to approximately $25^\circ N$ and somewhat further in the Southern Hemisphere ($\approx 30^\circ S$). Methane is produced by biological processes in the troposphere and atmospheric upwelling in the tropics is its main transport pathway into the stratosphere. Therefore, the tropics can be viewed as the source of stratospheric methane (see section 2.1.3). The high values of SE in this region could be associated with high variability due to the constant mixing that is occurring. The upwelling branch of the Brewer-Dobson circulation combined with diffusive mixing and ongoing chemical processes, are likely to cause strong fluctuations in tracer concentrations in this area and therefore complex dynamics. These fluctuations are successfully identified by SE. The regions of lower complexity are in accordance with previous studies which have highlighted the existence of a one-way barrier, which mainly allows poleward flow, thereby separating the tropics from the rest of the stratosphere. Due to the isolation, this region of the stratosphere is sometimes

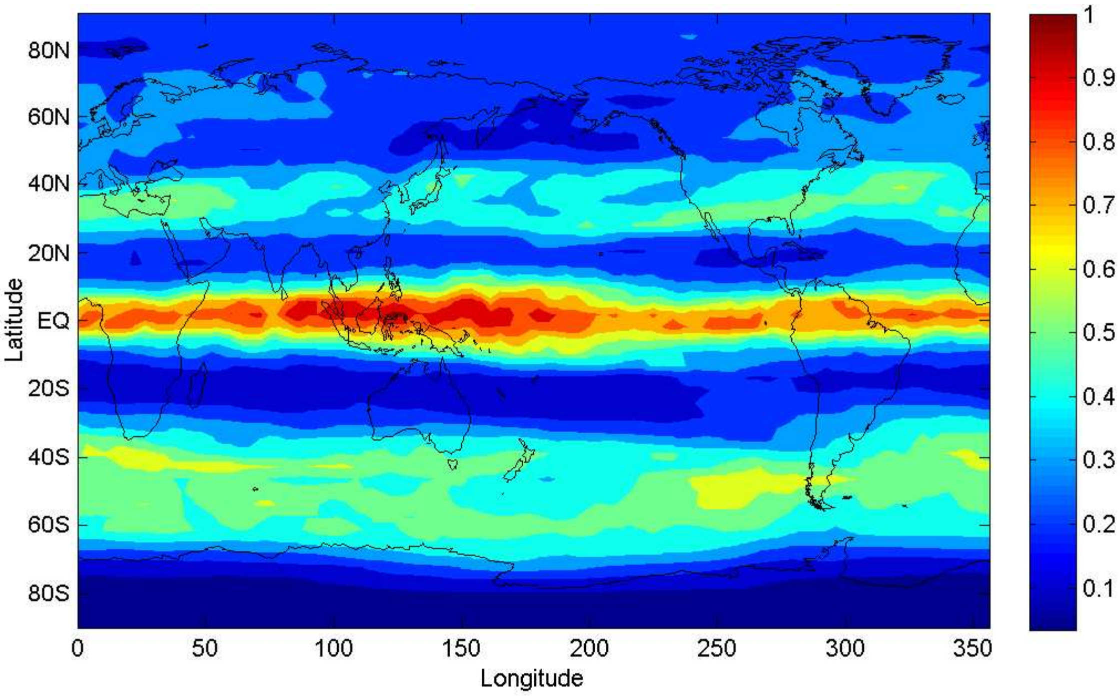

Figure 3.8: Contour plot of global $SE(m = 2, r = 0.4, N = 360)$, calculated using methane mixing-ratio data from SOCOL at approximately 31 km altitude (10 mBar) for 1980 conditions. A distinctive high-low-high pattern from the equator to the extra-tropics can be observed in both hemispheres. The equatorial high SE could be associated with mixing due to the upwelling arm of the Brewer-Dobson circulation.

referred to as the "tropical pipe" (e.g., *Plumb*, 1996; *Neu et al.*, 2003). The fact that the barriers (subtropical edges) appear wider than the isolated region itself can be explained by the variability of its location. Throughout the year the "pipe" and its edges move back and forth between the hemispheres, with a tendency of resting in the summer hemisphere until late autumn. For example, *Neu et al.* (2003) report that in July 1992 the tropical pipe was centred in the Northern Hemisphere and little change was observed until November. As a full year of data is used to calculate SE, the change of latitudinal location is likely to have the effect of increasing the standard deviation of the time series[7], thereby increasing the likelihood of finding matches, and thus leading to the broad regions of lower SE.

The intermediate values of SE, between $0.3 - 0.6$, at latitudes of $25°N - 45°N$ and $30°S - 65°S$ in the respective hemispheres, could be associated with the dynamics caused by planetary wave breaking, as these are the regions known as "surf zones" in winter (*McIntyre and Palmer*, 1984). Additionally, the air brought upwards in the tropics and moving poleward at high altitudes due to the Brewer-Dobson circulation in winter, starts descending at mid-latitudes. This continues during summer (while the Brewer-Dobson circulation is active in the other hemisphere), and while the mixing processes involved occur on relatively longer timescales, they can be a further source of dynamic complexity.

The equatorial low-high-low pattern of SE is part of a coherent pattern that spans multiple altitude levels of the SOCOL model. The altitude-latitude plots of SE shown in figure 3.9 illustrates this point. The left panel of the figure is an altitude-latitude contour

---

[7] This effect is explored in more detail below.

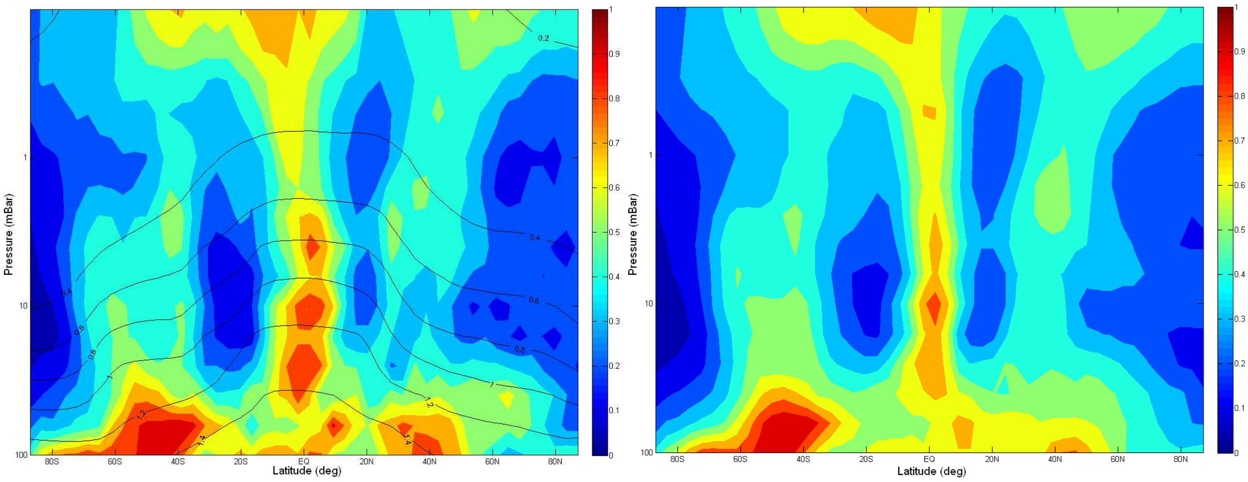

Figure 3.9: Altitude-latitude filled contour plot of $SE(2, 0.4, 360)$ derived from 1980 SOCOL methane data. The left figure shows the results at 176° longitude (including mean methane contours) and the right figure shows the zonal average of SE. The pressure range, 100 mBar ($\approx$ 16 km altitude) to 0.1 mBar ($\approx$ 45 km altitude), covers most of the stratosphere and the lower mesosphere.

plot of SE, with mean methane contours superimposed, at 176.25° longitude (approximately New Zealand's longitude). The right panel shows the zonal-mean results.

Figure 3.9 shows that the distinct patterns seen at 10 mBar in figure 3.8 are not limited to a single pressure level and can be found throughout a large part of the stratosphere. The regions of very low SE, which might be associated with the effects of the tropical edges, reach down to approximately 40 mBar. This is in accordance with *Randel et al.* (1998) who state that the lower boundary of the edge region is located at around 22 km, which approximately corresponds to the 40 mBar pressure level. Similarly, the study of the "edge" of the tropical pipe by *Neu et al.* (2003) also focuses on the levels between 31.2 mBar and 2.2 mBar, although these limits might be due to the data available for that study. In figure 3.9 the low-latitude regions of small SE values extend to very high levels (about 0.3 mBar). This already lies in the mesosphere and is potentially associated with poor simulation of the tropical flow and the lack of a QBO (Quasi-Biennial-Oscillation) in the SOCOL model version used. Thus, the SE might be highlighting a model artifact. The model mesosphere is free at the top and the dynamics of the highest levels are generally less reliable in SOCOL.

In order to test the validity of the association of the equatorial high entropy region with the tropical pipe at stratospheric altitudes, a slight variation to the way SE is calculated is used in the next section to analyse the variability of the tropical pipe.

3.2.2 Examination of the Variability of the Tropical Pipe with SE

In contrast to the results of section 3.1, *Richman and Moorman* (2000) claim that 10^m data points can be sufficient for calculating SE. Considering the degree of signal-separation when using the parameters $m = 2$ and $r = 0.4$ (bottom panel in figure 3.2), calculations of SE from only 100 data points, could give valid results, though a large uncertainty must be expected. This allows the calculation of daily values of SE by using the data points of a

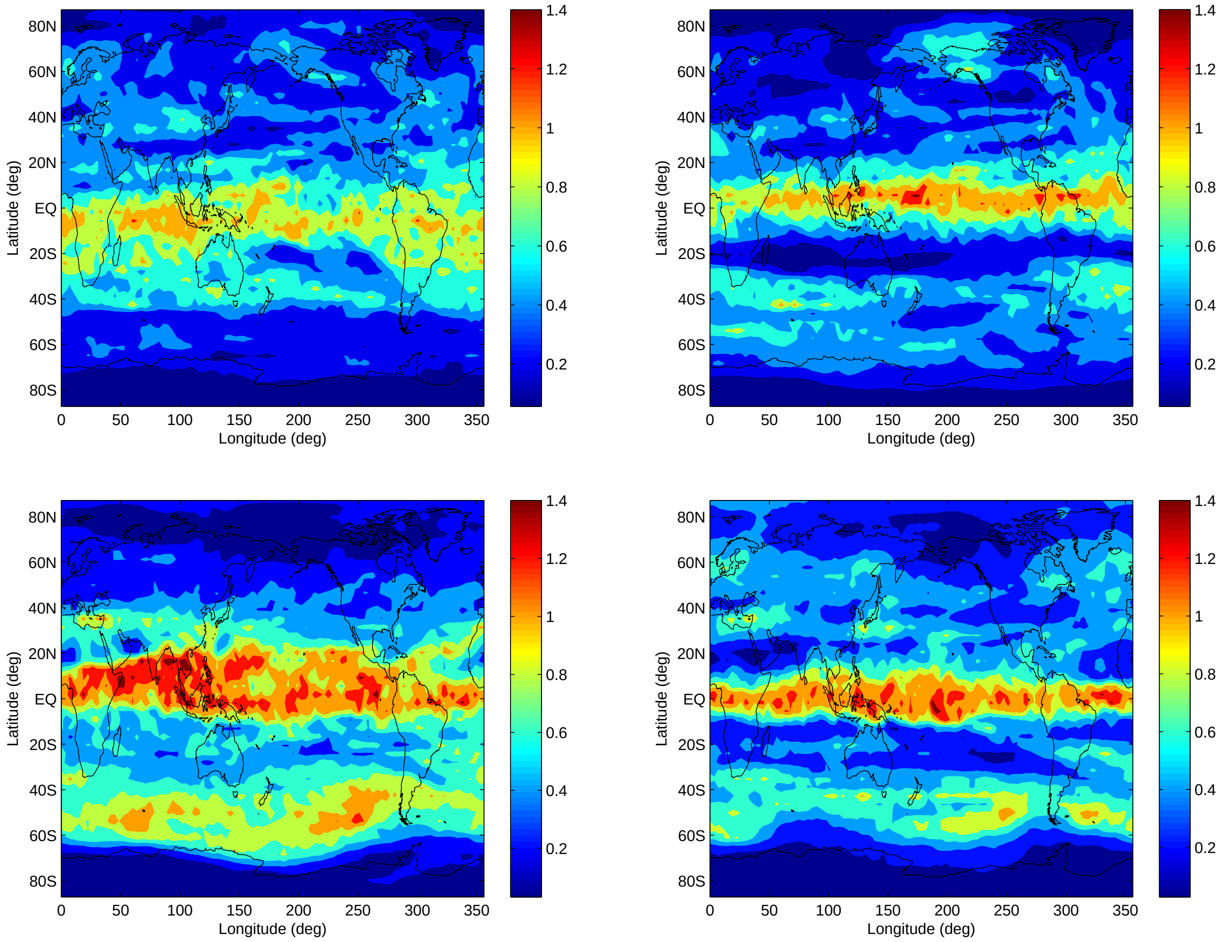

Figure 3.10: World-map-view of the temporal evolution of $SE(2, 0.4, 100)$ for 1980 SOCOL methane data at 10 mBar. Snapshots of $SE(2, 0.4, 100)$ are shown for early May (day 126, top left), beginning of June (day 150, top right), mid-November (day 315, bottom left), and late December (day 350, bottom right). The seasonal shift of the region of high entropy (possibly the centre of the tropical pipe) occurring in late Autumn can be seen in both hemispheres and is qualitatively similar to the evolution of the tropical pipe region described by *Neu et al.* (2003).

certain day and the previous 99 days. Using such a “100-day moving window”, the temporal evolution of SE can be calculated. It is important to note that $SE(m, r, 100)$ can only be considered an estimate of $SE(m, r, 360)$. It is not a fully equivalent metric and will give different information about the system from which the time series are derived.

Calculating $SE(2, 0.4, 100)$ using SOCOL methane mixing ratio data for a 100-day moving window at the 10 mBar level for 1980, results in a time-evolution of $SE(2, 0.4, 100)$ starting 100 days after the first data point, that is, in mid-April. “Snapshots” of this evolution at four different days (day 126, 150, 315 and 350) are presented in figure 3.10.

As was to be expected, the results are significantly noisier than when using a full year of data. All four images in figure 3.10 show a longitudinal band of high SE values in the equatorial region similar to those in figure 3.8, although the location of the centre of the band changes significantly with time. The bands of low SE surrounding it, can also be observed at most times, but with varying width and "strength". *Neu et al.* (2003) give a detailed description of the latitudinal migration of the tropical pipe from July 1992 to April 1993. Unfortunately, no SOCOL data for this time period was available for analysis, but the evolution of the central region can be assumed to be qualitatively similar every year.[8] *Neu et al.* (2003) find that

> *"in July the tropical region (as defined by chemical composition) is centred in the Northern Hemisphere (...). There is little change in the position of the edges until November, when, above about 10 mBar, the edge moves rapidly equatorward. (...) By April, the tropics are centred in the Southern Hemisphere. (...) at the onset of winter, the tropical region becomes very narrow."*

The equatorial region of high SE in figure 3.10 displays a similar evolution. In early May (top left panel of figure 3.10), the high entropy band is centred in the Southern Hemisphere. Less than a month later, by the beginning of June (top right panel), the band has drifted north and is now centred around the equator. It has also become significantly narrower. In mid-November (bottom left panel) the high entropy band is much wider again and clearly centred in the Northern Hemisphere. By late December (bottom right panel) it has again moved toward the equator. Both the latitudinal evolution and the change in width of the high SE region are very similar to the dynamics of the tropical pipe region identified in *Neu et al.* (2003). This supports the interpretations of SE in the previous section.

3.2.3 Polar Regions

The previous results suggest that SE can be used to identify mixing barriers. The equatorial and mid-latitude structures found in figures 3.8 through 3.10 are relatively consistent with the current understanding of stratospheric dynamics in these regions. On the other hand, the values of SE at polar latitudes, and particularly over the Antarctic, are surprisingly low. Considering the large-scale changes occurring throughout a year – vortex build-up, planetary wave breaking and sudden stratospheric warmings (particularly in the Northern Hemisphere), followed by final vortex break-up – the associated variability might be expected to lead to significant system complexity and therefore to relatively high SE in these regions. Instead, the polar regions appear as rather simple and homogeneous areas of low SE. The fact that a full year of data is used could mean that the mentioned dynamic processes are 'averaged out', as they occur on shorter timescales. This could mean that short-term dynamical information might be lost in the calculation of SE. Given that this is a particular focus of the research and an area of active investigation by other research groups, this is a significant flaw.

---

[8] *Neu et al.* (2003) show that the location of the "edge" of the tropical pipe is slightly affected by the quasi-biennial oscillation (QBO), but otherwise has a low inter-annual variability. The version of the SOCOL model used for this study does not exhibit a QBO. Therefore, assuming relatively similar dynamics of the tropical pipe for every year is justified.

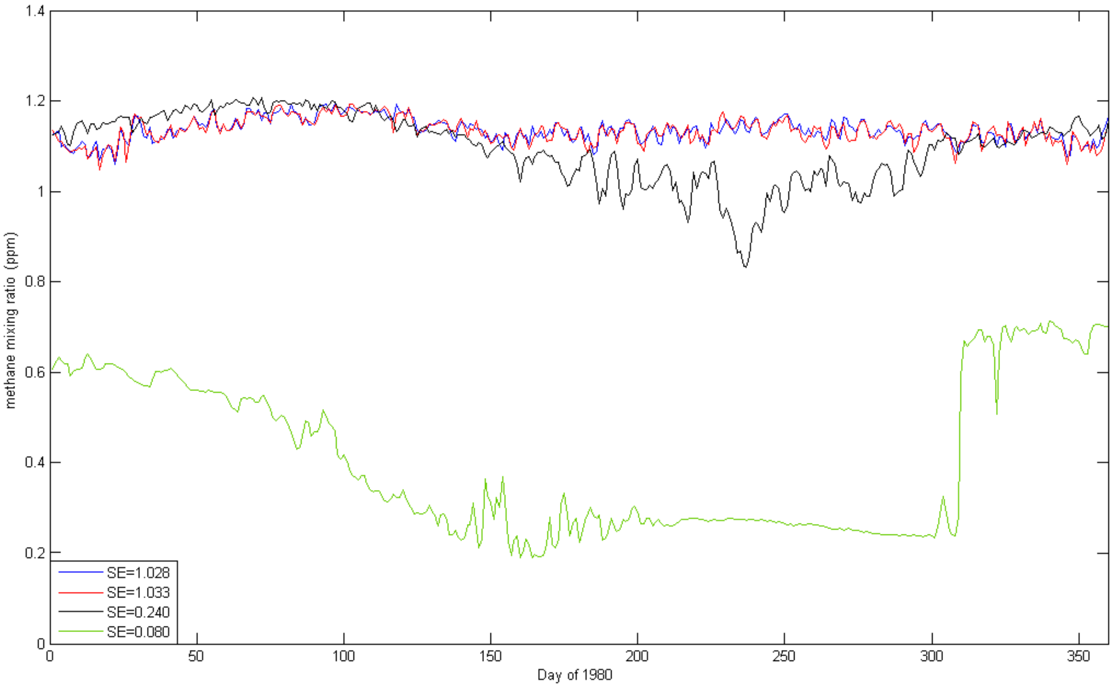

Figure 3.11: The time series of methane mixing ratio (in ppm) at four different locations at the 10 mBar level. These are examples of the data underlying the SE calculations shown in figure 3.8. The red and the blue line are from 1.9°N, 123.75° and 127.5° longitude, respectively. The black line corresponds to 13.0°S , 123.75° longitude, and the green line is from 79.8°S and 123.75° longitude.

3.2.4 The Standard Deviation Problem - Deficiencies of the SE

In order to identify the reasons for the very low values of SE in polar regions (and the potentially associated loss of information on short-term dynamics), the actual data underlying the calculations are analysed. Figure 3.11 displays examples of the methane data at different positions used for calculating the SE in figure 3.8. The red and the blue line in figure 3.11 are from neighbouring points close to the equator at 1.9°N, 123.75° and 127.5° longitude, respectively. The black line represents the methane mixing ratio at 13.0°S and 123.75° longitude and the green line shows the evolution at 79.8°S and 123.75° longitude. The calculated values of SE at these locations are 1.033, 1.028, 0.24, and 0.08, respectively.

The two equatorial time series (red and blue line in figure 3.11) show rapid fluctuations around a mean value of approximately 1.1 ppm. These are probably associated with fast, turbulent processes as the air entering from the troposphere continues to rise, and warrant the high value of SE observed. It should be noted that the values of SE for these two time series are significantly higher than that of the most chaotic dynamics of the logistic map ($\approx$ 0.65, see figure 3.6) and almost as high as the SE of autocorrelated noise ($\approx$ 1.1, see figure 3.2). This suggests that the mixing occurring in the tropical stratosphere is highly unpredictable and creates very erratic day-to-day changes in the methane distributions. This underlines the fact mentioned at the end of section 3.1, namely that SE does not distinguish between chaotic dynamics and noisy signals, is potentially a problem. It severely limits the explanatory power of the SE when it is applied to atmospheric tracer data.

During the first half of the year, the data from slightly higher latitude (black line) shows similar variations, but with smaller day-to-day changes. This is followed by a dip to lower mixing ratio (0.8 ppm) in the second half of the year. This is likely to be due to the tropical pipe moving into the Northern Hemisphere at this time. As the location is fixed in latitude, it probably first lies inside the central part of the tropical pipe and then becomes part of the tropical "edge" region. During this period the variations significantly increase in magnitude. However, the overall dynamics are less erratic than in the equatorial time series, justifying a lower value of SE, although the magnitude of the decrease seems disproportionate. A possible reason for this discrepancy is discussed in the following paragraphs.

The evolution of the methane mixing ratio at polar latitudes (green line in figure 3.11) is heavily influenced by the polar vortex. Until about day 100 (mid-April), the vortex builds up, causing a decrease in mixing ratio from 0.6 ppm to 0.4 ppm with minor fluctuations. With the establishment of the vortex in April (*Huck et al.*, 2007), the decrease speeds up, due to the increased isolation from mid-latitude air masses, down to 0.2 ppm. Several dynamical disturbances (temporary increases in methane) are observed in mid-winter. In late winter and early spring (day 220-300), the vortex is an extremely strong barrier and the methane mixing ratio is extremely stable with a slight downward trend. This changes very suddenly in November which is highlighted by a discontinuity in the methane data. This discontinuity is associated with the period when the vortex breaks up and mid-latitude air masses quickly replenish the polar air. The mixing ratio recovers to even higher values than at the beginning of the year (0.7 ppm). Methane levels remain high until the end of the year and fluctuations are relatively rapid.

The very low value of SE calculated for this time series (≈ 0.08) is in great contrast to these very complex dynamics. The reason for this discrepancy is the method of determining the tolerance interval r for the calculations. As this is taken to be proportional to the standard deviation of the time series, the large changes in the methane values observed during the vortex period lead to a large standard deviation. Therefore the tolerance interval is significantly larger than in the other cases mentioned above (e.g., the red, equatorial time series in figure 3.11 has a standard deviation of $2.59 \cdot 10^{-2}$ ppm as opposed to $17.1 \cdot 10^{-2}$ ppm for the green, polar one). Accordingly, the tolerance interval used for calculating SE of the polar time series is significantly larger than the average day-to-day variations in methane. This drastically increases the likelihood that the mixing ratio stays within the tolerance ($r = 0.4\sigma$) for another day, given that it already has been for the last two. In more technical terms, while the number of matches of length two, B^2 (see equation 2.15), is probably significantly smaller for the polar than for the equatorial time series (due to the large absolute differences in mixing ratio at different times), given a match of length two, a match of length three is much more likely for larger r. Consequently, the number of matches of length three, A^2, is going to be very similar to B^2, which leads to the very small value of SE ($= -ln\left(\frac{A^2}{B^2}\right)$) for the polar time series.

The same effect probably causes the value of SE to be so much smaller for the black tropical time series than for the red and blue ones. The temporal dip in mixing-ratio during the second half of the year inflates the standard deviation, σ, of the time series, which leads to the drop in SE. This effect of temporary spikes and non-stationary data on the values of SE is a known problem (e.g. *Lake et al.*, 2002). Possibilities for correcting this 'standard deviation problem' have been discussed in *Richman et al.* (2004) and include averaging multiple SE values and re-sampling methods. These are not valid options for the present study, as they all effectively truncate the data set, and therefore require much larger initial data sets than those available. When working with non-stationary data, another suggested approach uses surrogate time series (*Schreiber and Schmitz*, 2000). The idea is to adjust the used parameters of SE in such a way as to maximise the difference between the SE of the actual data and that of a surrogate time series that includes the linear trend of the data. Combined with a piecewise analysis this could improve matters somewhat, but at the cost of significantly increased computing times. Additionally, the merit of this approach is rather uncertain, as the use of surrogates that include the trend of the real data, requires a priori knowledge of this trend. This raises the question of how much of the information that might be derived from the calculations of SE afterwards, might be due to the surrogate trend assumed. Therefore, SE can only be used in a limited (tropical) region to determine the complexity of atmospheric dynamics from tracer data. When including non-stationary time series – such as the ones created by relatively long-term but transient phenomena, like the polar vortex – the values of SE do not adequately represent the relative dynamic complexity of different parts of the data set.

3.3 Conclusion

It follows from the above analysis that the SE statistical measure cannot be used to analyse climate complexity on a global scale and is particularly unsuitable for use in polar regions, where the variation in tracer concentrations is particularly non-stationary. As the mixing processes in these regions are very important, i.e., they have a significant impact on the global climate, a different approach to quantifying dynamic complexity in the stratosphere is required. Sudden changes such as the rapid increase in methane at vortex break-up (see figure 3.11), are likely to cause problems with most complexity measures that are based on the analysis of time series. Thus, a metric that is non-temporal, i.e., does not require chronological data, is desirable. *Sparling* (2000) suggests the use of probability distribution functions (PDFs) of tracer concentrations as an alternative point of view on stratospheric processes. PDFs can be created from almost any kind of data set, temporal or not, and offer a way to avoid the mentioned problems. Therefore, the focus is turned to Rényi entropy, which can be used as a complexity measure of PDFs.

Chapter 4

Rényi Entropy (*RE*)

Large parts of this chapter are from the paper "Identifying mixing in chemistry-climate model simulations using Rényi entropy" published in Geophysical Research Letters (Krützmann et al., 2008).

Understanding the dynamics of the stratosphere has become increasingly important since the 1980s. The discovery of the mechanisms that lead to the Antarctic ozone hole showed that stratospheric dynamics could have a direct impact on the surface. It was found that the chloro-fluoro-carbons (CFCs) emitted by anthropogenic sources are transported from the surface into the stratosphere and accumulate in the polar regions, where they cause major ozone depletion in Springtime (*Huck*, 2007), leading to a significant increase in UV-irradiance on the surface. In the stratosphere, the dynamics that cause the transport of chemical species are large-scale advection and small-scale eddy-mixing. As mentioned in section 2.1.2, the stratosphere can be divided into four regions with different mixing properties, which are separated by more or less distinct mixing barriers. The most common ways of analysing mixing in the atmosphere are examination of trends and variability of measurement time series (e.g., *Li et al.*, 2006), and analysis of the trajectories of chemical tracers from measurements, or of air parcels from simulations (e.g., *Garny et al.*, 2007; *Haynes and Shuckburgh*, 2000). However, the examination of the probability distribution functions (PDFs) of tracer concentrations offers an alternative point of view (*Sparling*, 2000).

Air parcel trajectory analysis, where the paths of air parcels are calculated from wind speeds derived from measurements or model simulations, involves time series similar to the ones analysed in the previous chapter, but with the coordinate system following along the path of a volume of air (a Lagrangian framework) instead of being fixed to a longitude, latitude, altitude grid (an Eulerian framework). The changes in the trajectories of air parcels from similar starting conditions and/or changes in the chemical composition of the air parcels can be used to derive mixing metrics such as finite time Lyapunov exponents (*Garny et al.*, 2007), equivalent length (*Lee et al.*, 2001), and effective diffusivity (*Haynes and Shuckburgh*, 2000). These usually involve computationally demanding calculations and the application of such measures to large amounts of data (e.g., a decade of atmospheric tracer data from measurements or simulations) can require very large computer facilities and a significant amount of time. Other, more statistical, approaches also require very large amounts of data, such as multi-decadal time series (e.g., *Li et al.*, 2006). However, creating PDFs in a fixed coordinate system is much easier and has the advantage that zonal concatenation can be used in order to obtain the number of data points required for statistical analysis. In particular, for long-term analysis, where predictions on a daily scale are likely to be inaccurate anyway, statistical mean and trend analysis via PDFs is much less computationally demanding than large-scale trajectory simulations. Therefore, the utility of the Rényi entropy statistical measure (*RE*), a measure of the complexity of a PDF (see section 2.3.3), as a tool for analysing climate dynamics with tracer data is tested in this chapter. The Shannon entropy, one special case of the *RE*, has been successfully used for analysis of climate data by *Knuth et al.* (2005). It has also been used in theoretical studies of mixing in fluids by *Everson et al.* (1998) and *Camesasca et al.* (2006).

The following section discusses the methodology used to calculate the *RE*, which is then tested on the logistic map process. Section 4.2.2 presents the results of applying *RE* to tracer data from SOCOL simulations (see section 2.4) and contrasts these with the results of using the better known Shannon entropy. In section 4.2.3 these results are compared with the findings of previous studies that use different methodologies for analysing atmospheric mixing. Due to the significant impact of the Antarctic polar vortex on stratospheric dynamics, the initial focus is on the Southern Hemisphere stratosphere. In this region, characteristic patterns can be expected, similar to the ones found by *Haynes and Shuckburgh* (2000), *Allen and Nakamura* (2001) and *Lee et al.* (2001). Section 4.3 draws conclusions about the utility of the *RE* for this analysis.

4.1 Methodology

Analysing tracer distributions using PDFs can be useful for identifying large-scale transport in the atmosphere (*Sparling*, 2000) as the concentrations of long-lived tracers, such as methane, are heavily dependent on the transport occurring. In order to calculate *RE* for PDFs that are derived from histograms of tracer mixing-ratios, the parameter α (see equation 2.20) and a method for binning the tracer data into PDFs must be determined. For this analysis, α is set to two. In section 4.2.2 it is shown that using $\alpha = 1$ in equation 2.20 results in a significantly reduced dynamic range (the range of values of *RE*) and is therefore sub-optimal for the intended application. As *RE* is a strictly decreasing function of α (*Beck and Schlögl*, 1993), using values higher than two does not intrinsically provide any additional information except that strong peaks in a distribution become emphasised. Structures similar to those identified in the analysis in section 4.2.2 can still be recognised when using higher values of α (not shown) but are less concise and appear noisier. Hence, this study suggests that the value $\alpha = 2$ provides the most useful bin-weighting. Equation 4.1 shows *RE* in the form used for the calculations described in section 4.2.2. Unless stated otherwise, the abbreviation *"RE"* also refers to this form.

$$RE(b,N) = -\frac{1}{\ln(b)} \ln\left(\sum_{i=1}^{b} p_i^2(b,N)\right) \tag{4.1}$$

The method used for deriving the PDF, $p_i(b,N)$, required for the calculation of the *RE*, is detailed in the following section.

4.1.1 The Optimal Binning Algorithm

When working with data from geophysical measurements (e.g., tracer concentrations from a satellite with global coverage) one usually faces a large number of data points. Using this data for analyses of circulation properties often requires some kind of binning of the data. In some cases the system might provide "natural bins", e.g., heads and tails for a coin toss. Generally though, the binning depends heavily on the data itself and choosing an appropriate number of bins for an unknown PDF is not trivial. This study uses an optimal binning algorithm to estimate the PDFs required for calculating RE from equation 4.1. More accurately, the data is used to construct a histogram (with an optimal number of bins), which is very similar to a piecewise constant PDF. Any histogram can be converted to a piece-wise constant PDF by dividing the bin counts by the total number of data points. In this sense, the words "PDF" and "histogram" can be considered synonymous and are used interchangeably.

The optimal binning algorithm (optBINS) proposed by *Knuth* (2006) determines the optimal number of histogram bins, b, under the assumption that all bins are of equal width. If the number of bins, b, is too small the detailed structure of the histogram (or PDF), such as minor peaks, can be lost due to coarse binning. If b is chosen too large, noisy variations in the data can be overemphasised or the distribution can become disjoint and large-scale structure is lost. Therefore, it is essential to choose a value for b small enough to correctly represent the overall structure but still large enough to resolve small-scale characteristics.

When estimating the PDF of a data set (i.e., tracer concentrations on a circle of latitude) the optimal bin width (and thereby the optimal number of bins)[1] depends on the range of values covered by the data. This can vary by several orders of magnitude depending on the latitude and altitude considered (see also section 3.2.4). The optBINS algorithm applies Bayesian probability theory to calculate the "*posterior probability of the number of bins given the data*" (*Knuth et al.*, 2005). The "posterior probability" is the probability that the data was drawn as a sample from the PDF that uses b bins. optBINS calculates this for a large number of values of b and selects the value that maximises the posterior probability. This value of b is then the "*most probable solution given all the information about the state*" (*Thomas*, 2003).

Figure 4.1 illustrates the usefulness of the optBINS algorithm. A set of points is created that are centred around the value 1.5 and have added random noise with a magnitude between ±0.5. Additional sets with different numbers of points centred around the values 2.5, 3.5, ..., 7.5 are created in the same manner (see top left panel of figure 4.1). From the combination of all of these points a histogram can be created. When using 90 bins for the histogram (top right panel of figure 4.1), spikes can be observed. This binning clearly identifies structures that are not warranted by the data. Similarly, distortion of the 'real' structure can also be seen when using too few bins. With only four bins (bottom left panel of figure 4.1), the distribution has its maximum in the bin that ranges from 2.75 to 4.5, although the largest number of points was originally created in the interval between 4 and 5. Using the optBINS algorithm an optimal number of seven bins is determined, giving the

---

[1] Throughout this study, PDFs are always created using bins of equal width. This is a simplifying, but common assumption, e.g., *Sparling* (2000).

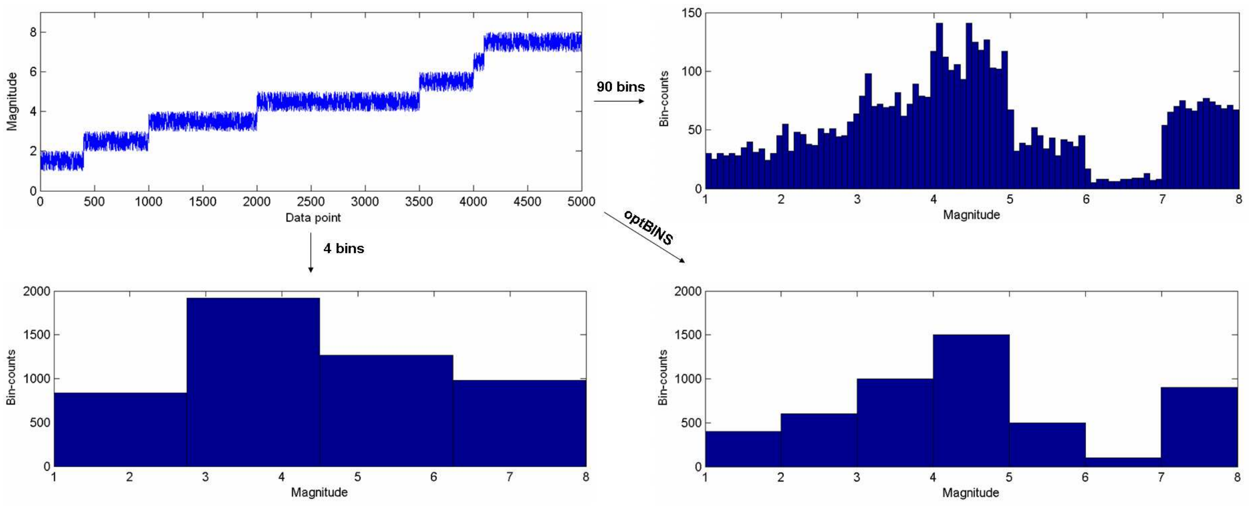

Figure 4.1: Test of the optimal binning algorithm, optBINS, with artificial data (top left). The algorithm determines an optimal number of seven bins (bottom right), which agrees with intuition as the data was created in seven distinct ranges. The top right and bottom left panels show histograms with too many and too few bins, respectively.

histogram in the bottom right panel of figure 4.1. This is obviously the most sensible value and would be expected intuitively from examination of the input, as the data was created in seven distinct ranges. As mentioned before, these histograms could be easily converted to the PDFs required for calculation of the RE of the distribution by simply dividing the bin-counts by the total number of points, 5000 in this case. Therefore, PDFs created from the histogram determined with the optBINS algorithm should give the most reliable estimates of the RE, given the data.

Figure 4.2 shows an example of the application of the optBINS algorithm to a distribution of an atmospheric trace gas. The left panel shows a sample Southern Hemisphere summer profile (monthly zonal mean) of a stratospheric tracer at 10 mBar, derived from SOCOL data. The right panel in figure 4.2 shows the histogram derived from the data by using the optBINS algorithm. In this case, 14 bins are found to be optimal. Using this histogram and using equation 4.1 ($\alpha = 2$) the RE of the distribution is calculated to be ≈ 0.786.

Monte-Carlo simulations indicate that the optimal binning technique requires a minimum of 150 to 200 data points to provide stable values for Gaussian distributed data (*Knuth et al.*, 2005, and *Knuth*, 2006) making the technique applicable to relatively small data sets. Whether this limit is also sensible when using optBINS for calculating RE is tested in the next section.

4.1.2 RE of the Logistic Map

In order to test how the suggested form of RE, given in equation 4.1, characterises system complexity it is applied to the logistic map (see equation 2.3). For this test, the evolution of the logistic map is calculated over 1000 steps for values of the parameter λ between 0.7 and 1.0, starting from a random seed point. The resulting data points were binned into a histogram using the optBINS algorithm (*Knuth*, 2006). The bin counts are converted to

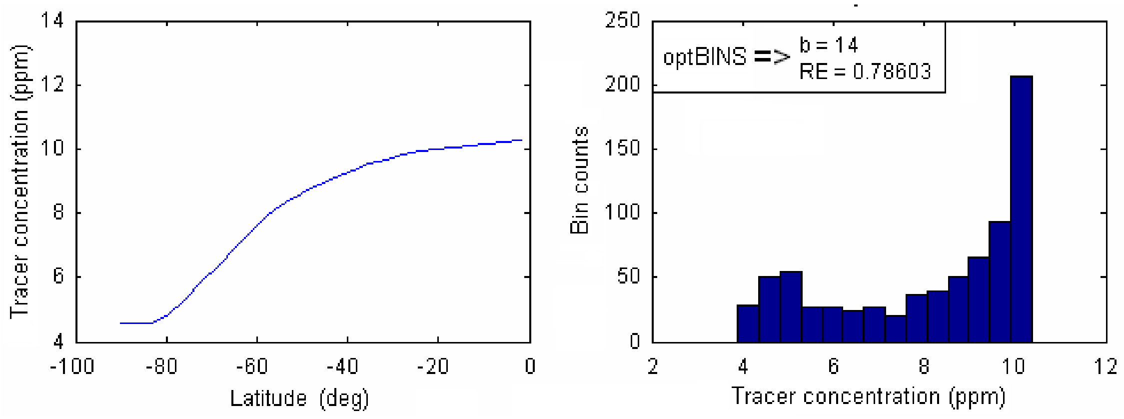

Figure 4.2: Example of the use of the optBINS algorithm. The left panel shows a sample Southern Hemisphere summer profile (monthly zonal mean) of a stratospheric tracer at 10 mBar derived from SOCOL data. The right panel shows the histogram derived from the data by using the optBINS algorithm. $b = 14$ is found to be optimal. With $\alpha = 2$, the RE of the distribution is calculated to be ≈ 0.786.

probabilities by dividing the number of counts by 1000 (the total number of data points). The resulting piecewise constant PDF is used to calculate the RE. Figure 4.3 shows the value of RE plotted over λ, superimposed on the Feigenbaum diagram of logistic map.

A relation between changes in the value of RE and the changes in the dynamics of the logistic map (discussed in section 2.2.2) can be clearly seen. Below $\lambda = \lambda_c$, RE increases stepwise with every bifurcation of the attractor cycle. Except very close to the bifurcation points, the value of RE is almost constant and only minor fluctuations are observed. For example, RE never drops below 0.125 between the first and the second bifurcation. After the second bifurcation a similar "plateau" can be seen with a value of approximately 0.25, and also at 0.375 after the third bifurcation (cycle of period eight). The end of the plateau at 0.375 is relatively close to λ_c, but by increasing the resolution in λ another plateau can be observed (see figure 4.4). No further plateaus can be discerned with any certainty, even when further increasing the resolution (not shown). This is probably caused by the optBINS algorithm. It cannot resolve the additional points of the attractor after the fifth bifurcation (from a 16- to a 32-cycle). At this point, the peaks are very close together and most additional peaks are likely to fall into the bins of other peaks. Further investigation of this behaviour of the optBINS algorithm is omitted here, as it seems rather theoretical. Distributions with 32 or more peaks are not likely to be a regular feature of the tracer distribution considered.

Further increasing λ beyond λ_c rapidly increases the RE. It permanently stays above values of approximately 0.5 for $\lambda > \lambda_c$, except in the "windows of order", where it drops to lower values again. In the largest window a similar stepping process towards higher values can be seen again. This suggests that, for the logistic map, $RE = 0.5$ can be used as an approximate threshold value, distinguishing relatively simple from complex behaviour.[2] The strong peak of RE at $\lambda = 0.75$ (first bifurcation) shows that RE can be used to differentiate

---

[2] It is important to note that the absolute values of RE mentioned in this section, e.g., 0.125, 0.25, 0.375, and 0.5, are not in any general way meaningful. Interpretations based on absolute values of RE cannot be generalised, not even to cases where only the number of data points is changed, much less when applying RE to a different system.

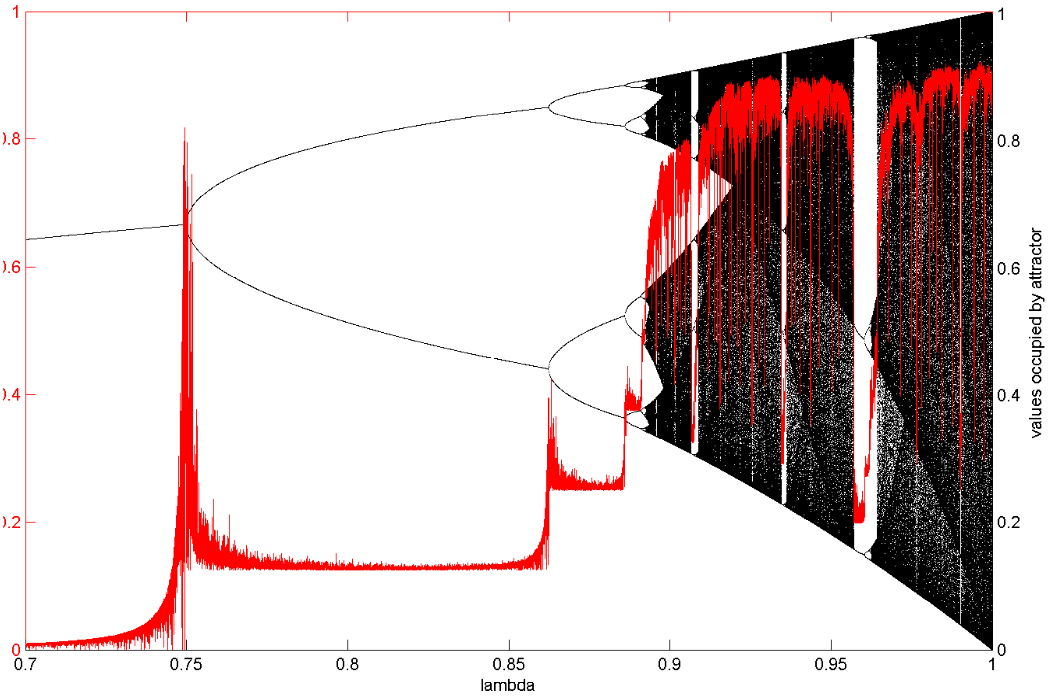

Figure 4.3: RE of the logistic map superimposed on the Feigenbaum diagram. Values above $RE = 0.5$ seem to indicate complex behaviour due either to transitional- (bifurcation) or chaotic dynamics.

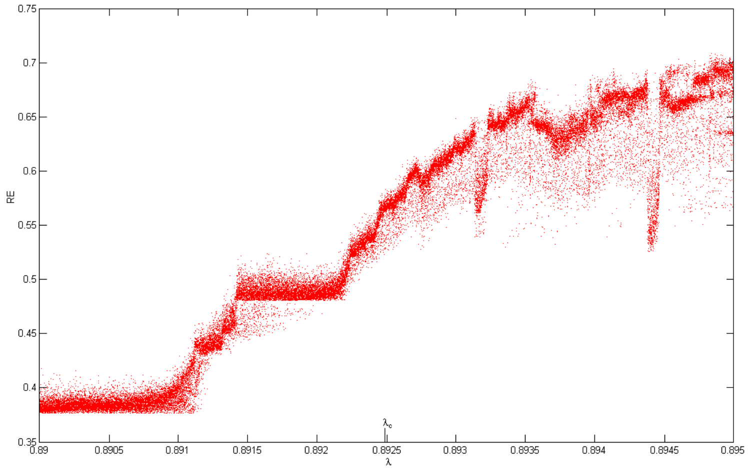

Figure 4.4: Zoom into the region close to $\lambda_c = 0.8924864...$ of figure 4.3.

when the system is undergoing a transition from one state to another. Similar, weaker, peaks can be observed at the other visible bifurcations. In the transitional phase, trajectories take a significantly longer time to converge on the attractor cycle, which leads to a higher *RE*, signalling higher complexity.

At values $\lambda > \lambda_c$, there are similarities between the *RE*-patterns and sample entropy (*SE*) behaviour. Like the *SE* (discussed in section 3.1.3), the *RE* gradually increases with lambda towards a maximum at $\lambda = 1$ (see figure 3.6). In the windows of order, both metrics display sudden decreases in their values. While the *SE* drops almost down to zero in the larger windows, the *RE* is only reduced to values that are related to the periodicity of the cycles in that window. For example, in the largest window of order in figure 4.3, when a period three cycle occurs the *RE* values drop down to values slightly greater than 0.2. This is consistent with the plateau structures found for $\lambda < \lambda_c$. For the 2-cycle, the plateau is at 0.125 and at 0.25 for the 4-cycle (see above). The value $RE \approx 0.2$ for the 3-cycle lies in between the values of the higher and the lower order cycles. Hence, the value of the *RE* can be approximately associated with a certain period cycle of the logistic map. This is very different from the behaviour of the *SE*, which is close to zero for any kind of periodic behaviour (see figure 3.6).

Overall the *RE* metric, derived using the optBINS algorithm, shows a monotonically increasing relationship with system complexity for the logistic map.

To determine the impact and importance of the optimal binning algorithm for the analytic properties of *RE*, the calculations is repeated (with slightly reduced resolution in λ) and the *RE* is calculated by using the optimal number of bins as well as by using a fixed number of 15 bins (see figure 4.5). While the graphs show very similar qualitative features, the differentiation between different types of behaviour is significantly better when using the optBINS algorithm. This demonstrates that the use of the optimal binning algorithm improves the sensitivity of *RE* as a complexity measure.

In order to test the minimum number of data points required to achieve reasonable estimates of the PDF of the data, the calculations for the logistic map are repeated with significantly shortened data sets. This is necessary as the limit of 150-200 data points, given by *Knuth et al.* (2005), is for optimal identification of Gaussian distributed data and the limit of 150 points for data of unknown distribution, mentioned in *Knuth* (2006), is only approximate. The PDFs underlying tracer distributions cannot generally be assumed as Gaussian and may have multi modal distributions. Therefore, separate testing is prudent. Most PDFs created from the logistic map will have several peaks, particularly in the cyclic attractor region, and therefore are useful for the intended test. Figure 4.6 shows the *RE* calculated from data sets with 50, 70, 90, 110, 150, and 200 data points.

The results using 50 to 70 data points (left panel in figure 4.6) are very noisy. Values of *RE* fluctuate significantly when using so few data points. With 90 and 110 data points (middle panel), a distinct variation of *RE* with changing system dynamics, similar to figure 4.3, becomes apparent. When using 150 or 200 data points (right panel in figure 4.6), the *RE* values are more consistent and a reasonably reliable distinction between simple and complex behaviour is achieved. While having more data points is desirable, as it improves the quality of the analysis of complexity with *RE*, the 150 data points limit determined for the optBINS algorithm by *Knuth* (2006), also seems a sensible minimum for general application of *RE*.

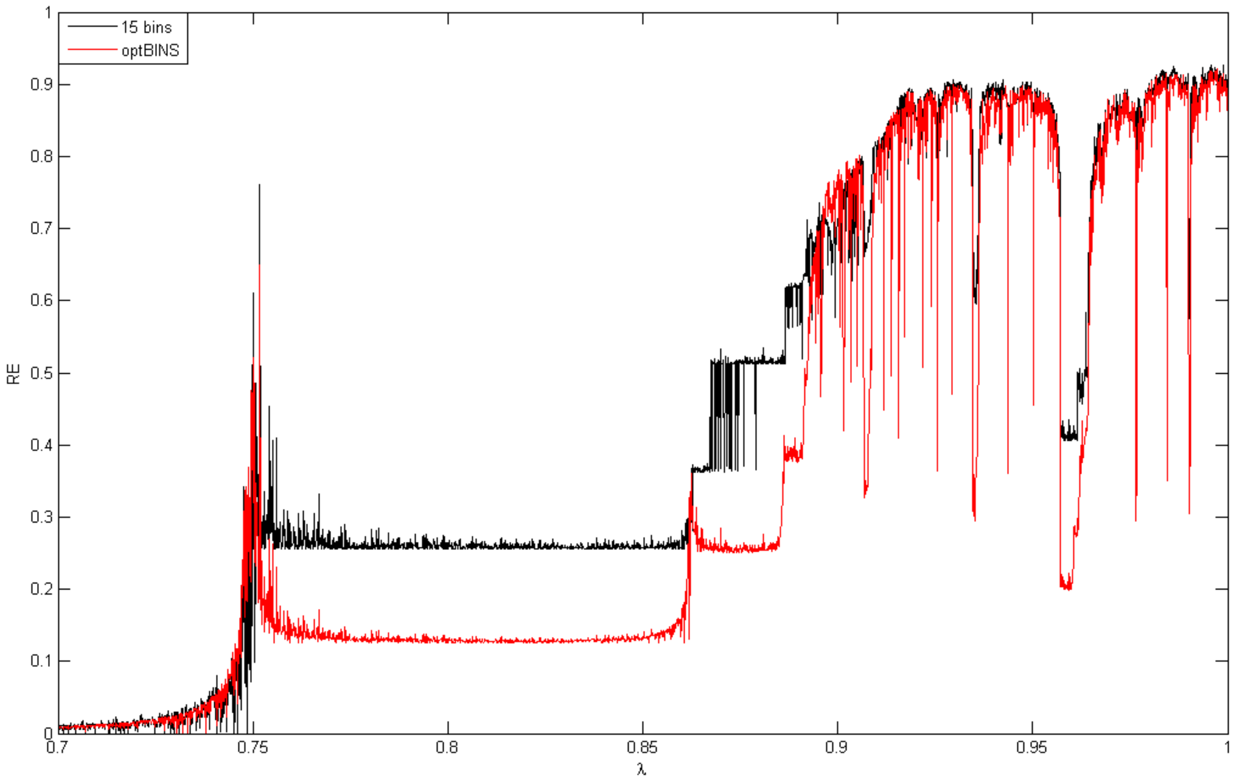

Figure 4.5: Comparison of RE of the logistic map when using the optBINS algorithm (red, same as in figure 4.3) or a fixed number of 15 bins (black). Clearly, the differentiation between simple and complex dynamics is significantly improved by the use of the optimal binning technique.

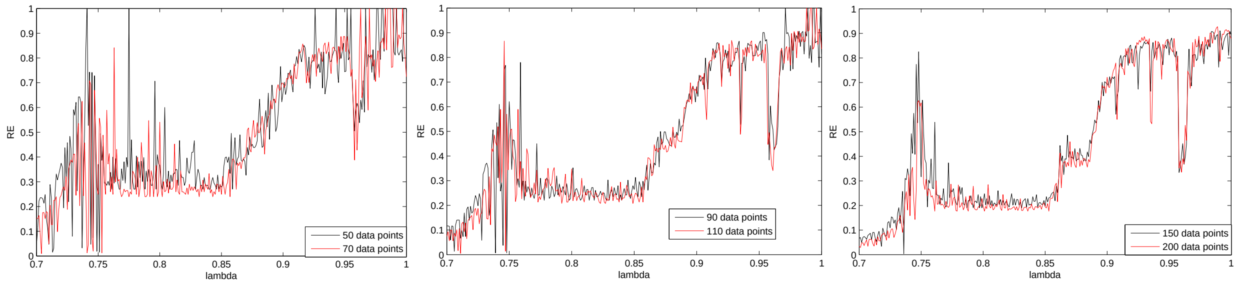

Figure 4.6: RE of the logistic map using 50, 70, 90, 110, 150 and 200 data points. Too small number of data points (left and middle panel) lead to highly variable and therefore unreliable results. A minimum of 150 to 200 data points (right panel) is a useful limit for reasonable estimates of RE.

This is underlined by the results shown in section *4.2.3.1*. There, the discovered patterns disappear when the number of data points drops below 150 due to latitudinal truncation of the data sets. However, in order to avoid any problems associated with a lack of data, the PDFs constructed for the stratospheric analysis rely on significantly more data points than the minimum requirements.

4.2 Application of the Rényi Entropy Metric to SOCOL Tracer Data

Using the optBINS algorithm and equation 4.1 ($\alpha = 2$), the RE measure is now applied to SOCOL tracer data.

4.2.1 Data Selection

Thus far, the RE statistical measure has been applied to well-understood non-linear dynamical systems. The utility of the Bayesian optimal binning algorithm for this analysis was demonstrated in section 4.1. The remainder of this chapter describes the application of the RE for identifying barriers to transport in the stratosphere. This methodology is motivated by previous work detailed in *Camesasca et al.* (2006), *Everson et al.* (1998), *Patmore and Toumi* (2006), and *Sparling* (2000), which has demonstrated that complexity measures can be utilised to identify and quantify mixing in the atmosphere and in fluids. In the present analysis, the RE is applied to tracer concentration (or mixing-ratio) data produced by the SOCOL model (see section 2.4). In particular, RE is applied to methane data in the stratosphere.

The daily mixing ratio of methane (in ppm) is calculated with the SOCOL model at every grid point (96 points in longitude, 48 in latitude, and 30 in altitude, see section 2.4). Daily PDFs for each latitude are estimated for a moving "10-day-window", i.e., by using all data points around the respective circle of latitude and also the values of the previous nine days. The PDFs are created using the optimal binning algorithm, optBINS, by *Knuth* (2006) that was introduced in section 4.1.1. Thus, RE is calculated from a zonal "10-day-PDF" (RE_{z10}) consisting of 960 data points. While this causes a certain degree of smoothing (the data of each day is used ten times), it is a useful approach as large-scale mixing processes in the stratosphere are likely to occur on a timescale of several days. Also, this timescale is much shorter than the chemical lifetime of methane, which is, on average, nine years for the whole atmosphere (*Wallace and Hobbs*, 2006) and more than 100 days below 40 km (*Randel et al.*, 1998). Therefore, it can be assumed that the distribution is defined by transport and mixing. Thus, an understanding of the complexity of the tracer PDF provides a measure of transport and mixing properties.

The tracer PDFs used in this study were not weighted according to latitude. While this is necessary when creating and analysing global PDFs, as one model grid-point corresponds to a larger area at lower latitudes (*Sparling*, 2000), this study does not directly combine data across latitudes, but only uses data points of identical latitude to create the PDFs from which RE is calculated. This might lead to increased smoothness of the PDFs at high latitudes, due to a denser sampling in longitude. However, the better sampled distribution should still have the same relative proportions, which is the most relevant factor for the calculation of the RE. Hence, the value of the RE should not be affected significantly by the denser sampling of the high-latitude regions, making latitudinal weighting unnecessary. This is confirmed by findings in section 4.2.3.1.

The following section presents the results of the calculations of RE_{z10} of one year of SOCOL methane data. These are then discussed in section 4.2.3. The focus is on the Southern Hemisphere stratosphere because of the significant impact of the winter Antarctic polar vortex on the dynamics in this region.

4.2.2 Results

The top panel in figure 4.7 shows a contour plot of RE_{z10} derived from SOCOL methane distributions in the Southern Hemisphere, representative of conditions associated with 1980. The superimposed contours represent ten-day zonal-mean zonal wind values. The calculations are performed on the 10 mBar pressure level. During the Antarctic autumn, winter and early spring, a band of high RE_{z10} ($60°S - 70°S$) separates intermediate values at mid-latitudes from lower values at southern high latitudes ($70°S - 80°S$). The maxima of zonal wind speed are observed to be co-located with regions of high entropy. The northern polar regions (not shown) display a much weaker seasonal evolution of entropy partitioning. This is probably explained by the weaker and more distorted boreal polar vortex, which allows significantly more mixing with mid-latitude air.

The bottom panel of figure 4.7 shows the results when using $\alpha = 1$ in equation 4.1 (Shannon entropy) for comparison. While the same large-scale entropy patterns can be identified, they appear rather broad and detailed structure is lost. This is due to a significant reduction in dynamic range, e.g., at the 10 mBar level considered the values of RE_{z10} lie between 1 and 0.15, while the Shannon entropy covers values between 1 and 0.3. The average reduction of the dynamic range is 20%. When $\alpha = 2$ is used (see equation 4.1) a higher dynamic range is observed and the measure is more sensitive to small changes in the PDFs and hence to smaller scale dynamics. While the range can be further increased by using even higher values of α, the gain in range is less substantial and the results (not shown) appear strongly fragmented, thereby losing some of the large-scale structure.

The "white band" in figure 4.7, the centre of which is interpreted as a mixing barrier (see section 4.2.3), can be identified more clearly in the seasonal-average cross sections given in figure 4.8. RE is calculated by using the full zonal data for each month (RE_{zm}) and subsequent averaging over a season (3 months). The left panel in figure 4.8 shows strong entropy gradients in austral winter (June, July, August; 1980), associated with the structure of the polar night jet. In contrast, the right panel in figure 4.8 shows weaker entropy gradients in austral summer (December, January, February; 1980-81), these being associated with the weaker summer easterly winds.

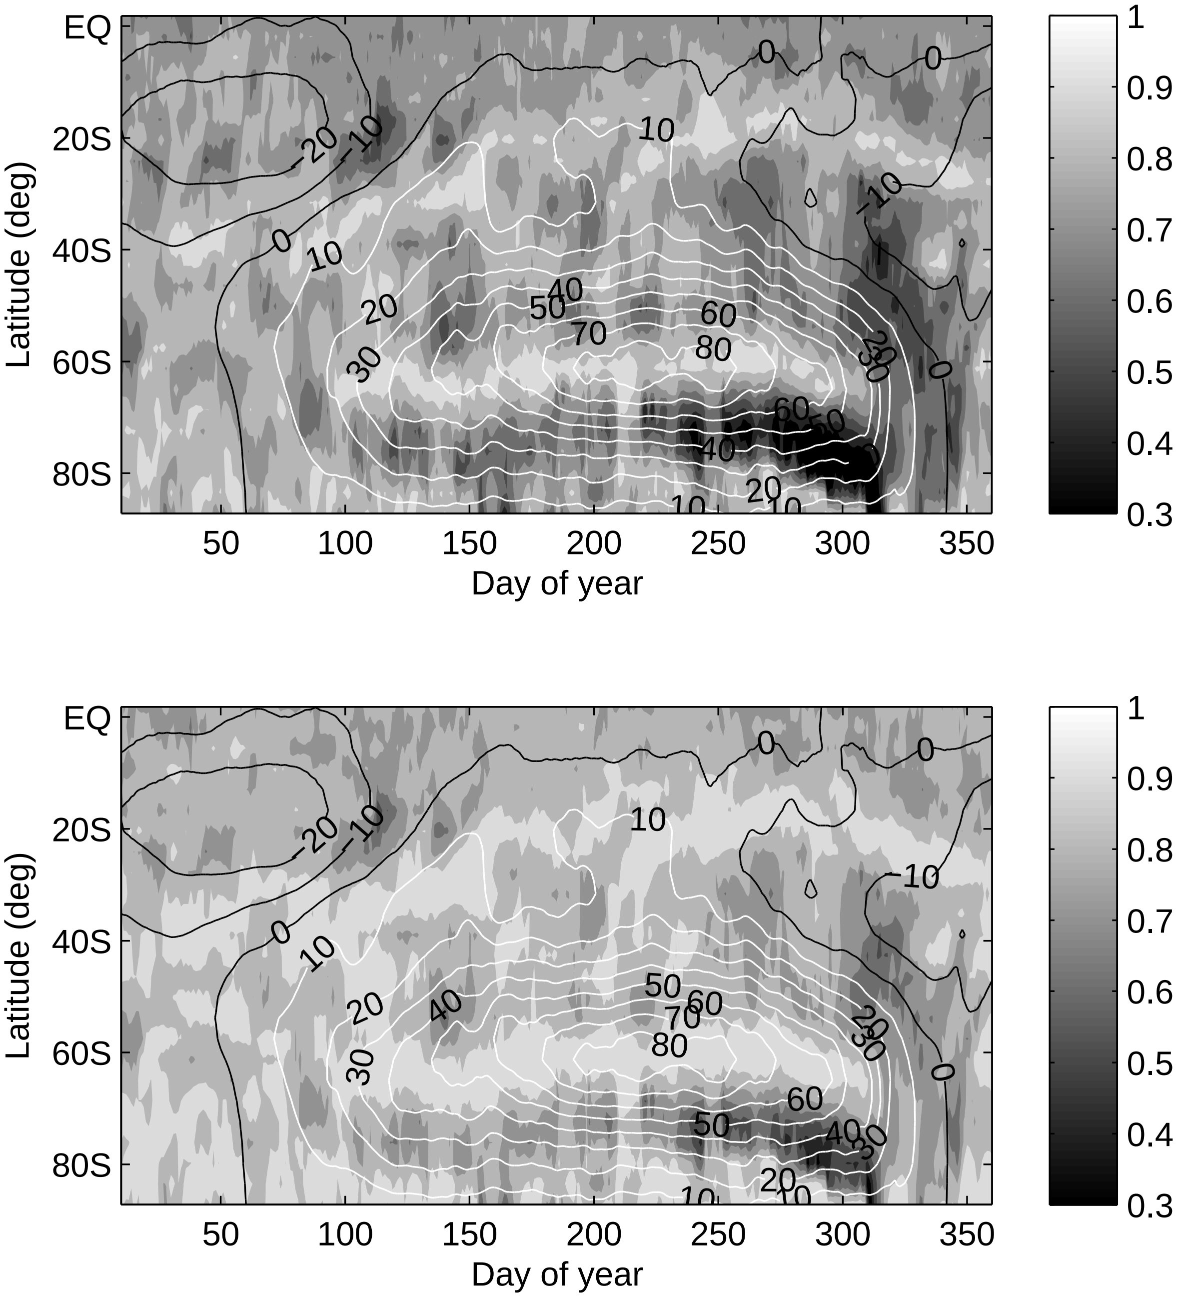

Figure 4.7: Time-latitude filled contour plots of RE_{z10} with $\alpha = 2$ (top panel), and RE_{z10} with $\alpha = 1$ (Shannon entropy, bottom panel), derived from SOCOL methane data (representative of 1980 conditions) in the Southern Hemisphere with zonal wind contours in m/s (white contours are westerly) at 10 mBar ($\approx$ 31 km altitude). A band of high RE_{z10} ($60°S - 70°S$) and a band of low RE_{z10} ($70°S - 80°S$) can be identified in both images during austral winter and early spring. The colour range is set to be equal in both images, in order to improve the grey-scale contrast. The actual dynamic ranges are as mentioned in the text.

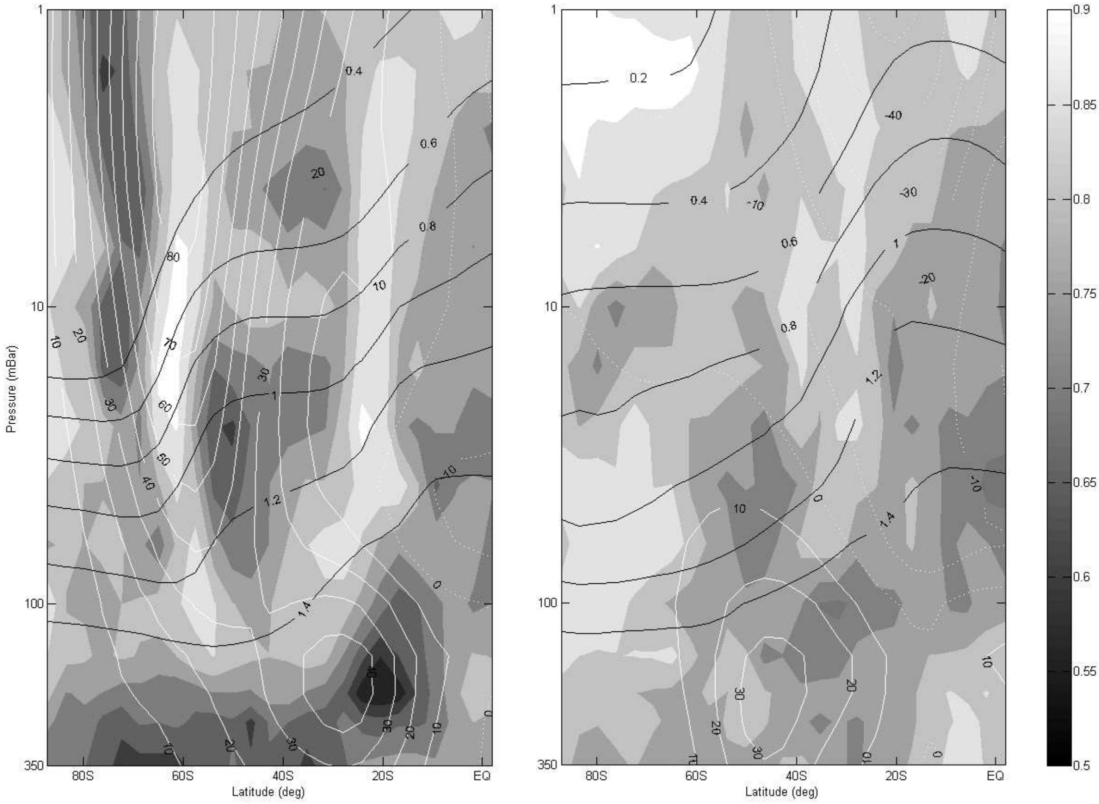

Figure 4.8: Latitude-altitude (350 mBar ≈ 7.5 km; 10 mBar ≈ 31 km) filled contour plots of seasonally averaged RE_{zm} derived from SOCOL methane data, with superimposed mean methane contours in ppm (black) and zonal wind contours in m/s (white; negative values are easterly). The bands of high and low RE_{zm} at high latitudes from figure 4.7 can be clearly identified in the left panel (June, July, August; 1980), while the right panel (December, January, February; 1980-81) does not show any similar patterns.

4.2.3 General Discussion

4.2.3.1 Latitudinal weighting

As mentioned previously, the PDFs used in this analysis are not weighted according to latitude. The applied method of estimating RE does not directly combine data across latitudes, but only uses data points of identical latitude to create the PDFs from which RE_{z10} is calculated. Hence, latitudinal weighting should be unnecessary. This can be tested by applying a latitudinal weighting and comparing the results with those from the full data set.

To achieve latitudinal weighting, random data points are removed at every latitude so that the remaining number of points is proportional to the cosine of the latitude. RE_{z10} is then calculated from the PDF of the remaining data points (see figure 4.9). As expected, the resulting structures are almost identical to the ones from calculations using all data points (see top panel in figure 4.7). However, at very high latitudes ($75°S - 85°S$), the patterns are noticeably less coherent. This is simply associated with the fact that the number of remaining data points drops below 200 for latitudes higher than $78°S$ and below 150 for latitudes higher than $81°S$. This confirms the result in section 4.1.2, that no less than 150

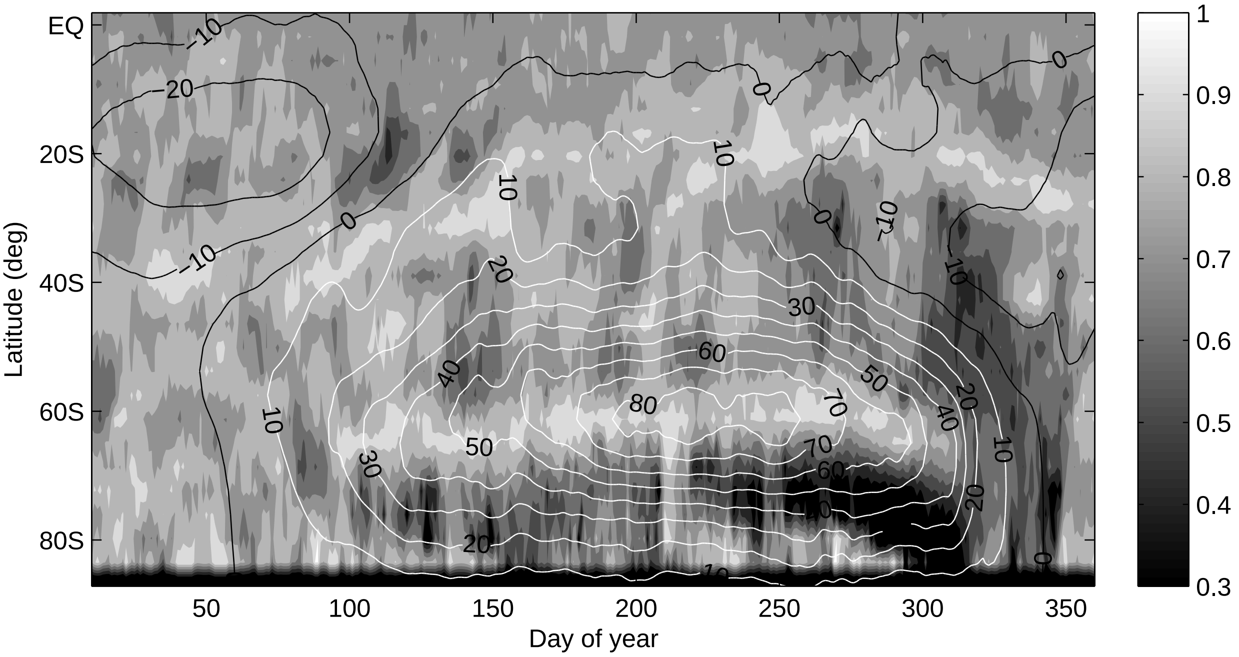

Figure 4.9: RE_{z10} calculated from latitudinally weighted PDFs at the 10 mBar level. The number of data points used for estimation of the PDFs is decreased as the latitude increases, to take into account the different sizes of the model grid boxes. The effect this has on the entropy patterns can be observed at high latitudes. There, the RE_{z10} patterns are significantly less coherent due to the number of data points falling below 150.

data points should be used for calculating RE. The strong similarity of the results from truncated PDFs with those using the full data set also shows that the calculations are quite robust. This analysis validates the assumption in section 4.2.1 that latitudinal weighting is unnecessary.

In this context, it should be emphasised that there is an important difference between the use of PDFs in this study and the applications proposed by *Sparling* (2000). While Sparling advocates the use of global scale PDFs because they have the advantage "*that they are less sensitive to outliers and hence to the effect of intrusions*" (*Shepherd*, 2007), this study uses PDFs that are localised in latitude, and whose shape is particularly strongly affected by intrusions. It is this distortion of the PDFs by dynamic changes that is exploited to infer transport properties. In the presence of strong mixing, the PDFs will show very strong peaks with few distortions, as most of the air will have very similar values of methane mixing ratio (i.e., it is well mixed). If mixing is weak the presence of tracer gradients should lead to a wider range of mixing ratios in the PDF. This means the air is not well mixed and the tracer amounts are very different at different locations and times. The tracer PDF will be rather 'widely spread' and/or show multiple weak peaks. Therefore, the PDF distortions caused by the presence or absence of large-scale mixing processes can be identified with RE.

4.2.3.2 RE and Mixing Barriers

RE can be considered to be a measure of the (non-)uniformity of a dynamical system (*Bialas et al.*, 2006). The strong polar vortex in the Southern Hemisphere leads to significant inhomogeneity in the large-scale distribution of stratospheric tracers (*Haynes and Shuckburgh*, 2000). The vortex-edge separates the cold polar winter air from warmer mid-latitude air. Its intensification in austral autumn severely restricts meridional mixing, thereby creating strong gradients in tracer concentrations (*Sparling*, 2000). This isolation ultimately also contributes to the magnitude of the Antarctic Ozone Hole (*Lee et al.*, 2001).

The altitude-latitude cross-sections of seasonally averaged RE_{zm} in figure 4.8 also have mean zonal wind (white) and mean methane contours superimposed (black). In summer (right panel), these display a simple layering of methane with a tendency to lower values at high latitude, while the winter distribution (left panel) shows a distinctive cascaded pattern with three main levels of mixing ratio. This is in accordance with the common separation of the stratosphere into summer hemisphere, tropical region, "surf zone", and polar vortex (*Plumb*, 2002). It also agrees with the work mentioned in section 2.1.3, which indicates that methane reaches the stratosphere in the tropics and is transported poleward via a circulation that is qualitatively similar to the Brewer-Dobson circulation. This suggests that the methane dynamics underlying the RE calculations are reasonably well simulated by SOCOL (see also *Egorova et al.*, 2005).

In winter (left panel of figure 4.8), the strongest gradients in methane concentration coincide with areas of high RE_{zm}, corresponding to relatively uniform PDFs, which are associated with unmixed air masses or a dynamic barrier. This is underlined by the fact that the band of high RE_{zm} extends over a large range of altitudes during austral winter, as do the highest wind contours (see left panel of figure 4.8). This wind structure is the polar night jet, which is a strong barrier to horizontal mixing in that region, and the contours are centred around the band of high RE_{zm}. A second high RE_{zm} band can be observed around $20°S$, which can be associated with the mixing boundary of the "tropical pipe" (*Plumb*, 1996). This is discussed in more detail in section 4.2.4. While these are not the only region of high entropy, they are the most salient and persistent features, indicating that the RE of tracer fields can be associated with atmospheric transport. During summer (right panel in figure 4.8), there is no polar vortex acting as a mixing barrier. Accordingly, the RE-patterns are, on the average, more diverse and no structure similar to that in winter can be found.

Other potential sources of changes in RE are transport of air masses to higher latitudes by planetary waves, or turbulent processes on a smaller chemical scale, e.g., at very high altitudes where the mixing ratio is low. The latter could be the reason for the area of high entropy at high latitude and altitude in summer (right panel in figure 4.8). Here, the concentration of methane is particularly low (note the 0.2 ppm contour, which is not visible in the winter distribution). Associated variations are on such a small-scale that the high entropy could be caused by problems with either the SOCOL model or the optimal binning algorithm. This issue potentially highlights a problem with the SOCOL runs that requires further investigation, but is beyond the scope of the current study.

In section 4.2.4 the observed entropy patterns are compared to previous work on stratospheric tracer dynamics and initial results of calculating RE_{z10} for data from satellite measurements are presented.

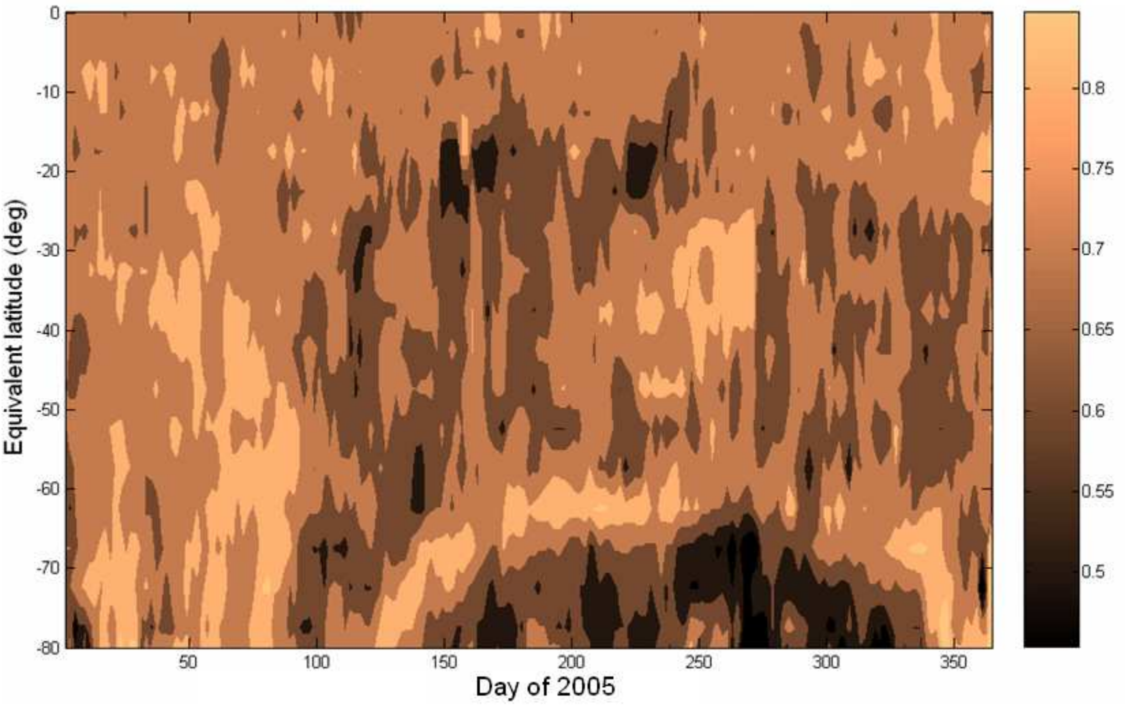

Figure 4.10: RE_{z10} of MLS N_2O measurements on the 850 K potential temperature surface ($\approx$ 13 mBar) in the Southern Hemisphere in 2005. Note that the vertical axis is equivalent latitude, an alternative meridional coordinate. Qualitatively similar, but weaker entropy patterns than those in figure 4.7 are observed at high latitudes during the polar vortex period.

4.2.4 Further Results and Comparison with Other Studies

4.2.4.1 Preliminary Results with MLS Measurements

In order to determine whether the other discovered patterns are realistic and not artifacts originating from the workings of the SOCOL model, RE_{z10} is calculated for EOS-MLS (Earth Observing System - Microwave Limb Sounder) *(Waters et al., 2006)* satellite measurements of nitrous oxide (N_2O), a tracer with properties very similar to methane. Figure 4.10 shows the results for the 850 K potential temperature surface ($\approx$ 13 mBar) in the Southern Hemisphere in 2005. Note that the vertical axis is equivalent latitude, an alternative meridional coordinate that is related to potential vorticity and potential temperature. The results display similar structures to the ones mentioned above, the main differences are the lack of a tropical structure and a significantly weaker high entropy band at high southern latitudes (see figure 4.10). The latter may be attributed to the effects of the model cold-pole bias, a common deficiency in CCMs (*Eyring et al.*, 2006). Stronger meridional temperature gradients than those observed lead to an intensification of the polar vortex and an associated reduction in meridional mixing. Nevertheless, a qualitative comparison of the patterns shows that the mid- to high-latitude RE_{z10} structures are reasonably similar and not simulation artifacts.

Further comparison with this data set is desirable but lies beyond the scope of this thesis. At the time, the data set used for figure 4.10 was the only full annual data set available. It also had several data gaps and, therefore, more and better data would be appropriate for a

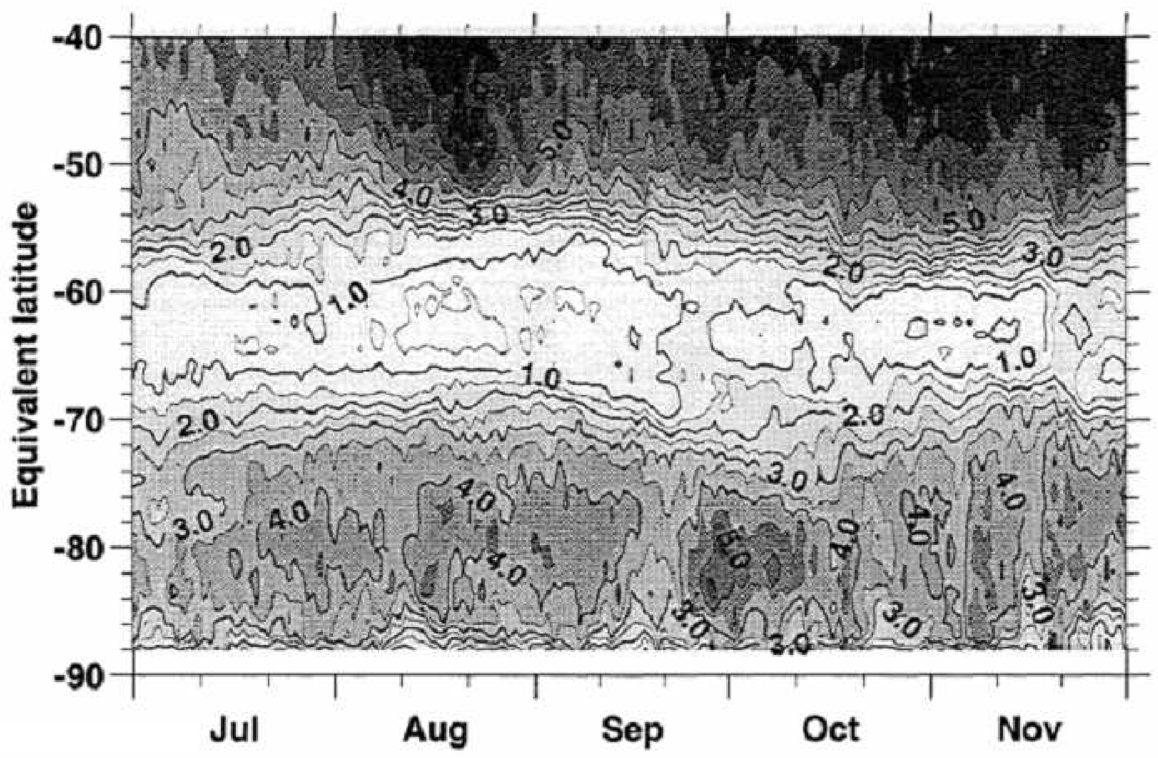

Figure 4.11: Evolution of the lognormalized equivalent length on the 480 K potential temperature surface ($\approx$ 70 mBar) in the Southern Hemisphere for a tracer advected by the winds observed in the 1996 winter and spring period (from *Lee et al.*, 2001).

potential future investigation.

4.2.4.2 Comparison with Other Studies of Mixing

Comparison of the results with previous studies that have attempted to quantify mixing in the Antarctic stratosphere allows examination of the utility of RE. With respect to mixing, *Lee et al.* (2001) distinguish two parts of the Antarctic polar vortex – the vortex-core and the vortex-edge. Air in the vortex-core is well mixed but isolated from air at lower latitudes by the vortex-edge region. This edge region restricts mixing between high- and mid-latitude air, is particularly unmixed itself, and is associated with very strong zonal winds. The white band (high RE_{z10}) at the centre of the strongest winds in figure 4.7, resembles the region of low lognormalised equivalent length[3] which characterises the vortex-edge in figure 1 of *Lee et al.* (2001) (reproduced in figure 4.11). The lognormalised equivalent length was in this case calculated for the 450 K potential temperature surface ($\approx$ 70 mBar). While this lies significantly lower than the 10 mBar level discussed, it was shown in figure 4.8 that the band of high RE extends over a large range of altitudes during austral winter, as does the polar vortex. The similarity of the structures suggests that the high-entropy region is associated with the vortex-edge.

Furthermore, examination of figures 4.7 and 4.8 shows another consistent band of high

---

[3] Equivalent length is based on the length of tracer contours on isentropic surfaces (*Nakamura*, 1996). It *"is a measure of the geometric structure of the tracer and hence of mixing"* (*Lee et al.*, 2001).

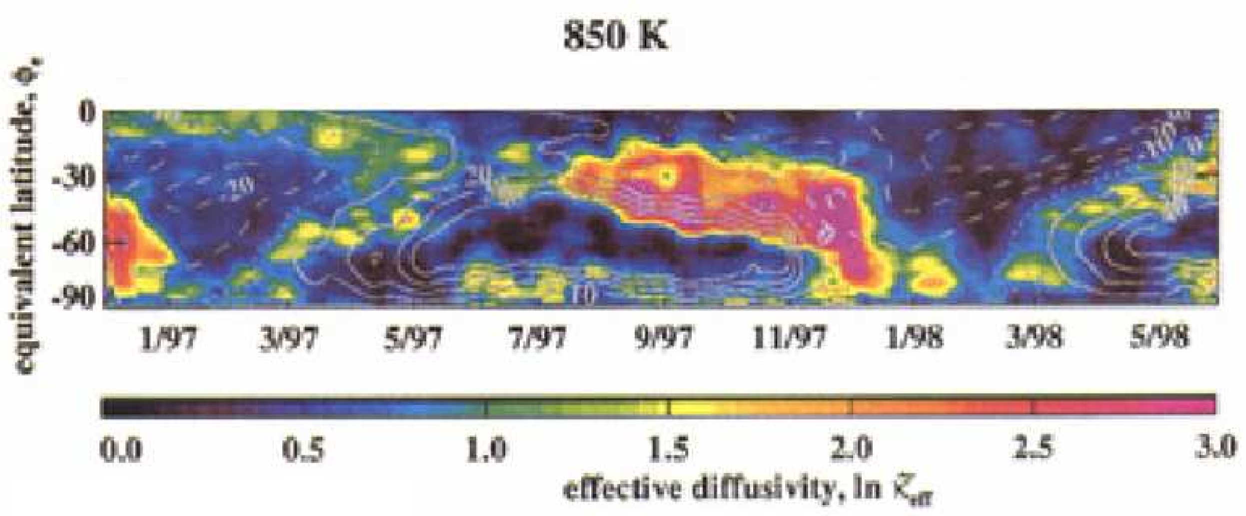

Figure 4.12: Evolution of effective diffusivity on the 850 K potential temperature surface (≈ 13 mBar) in the Southern Hemisphere calculated for a tracer advected by the winds observed in 1997/98 (adapted from *Haynes and Shuckburgh*, 2000).

values at lower latitudes, which separates tropical from mid-latitude air. This is in accordance with previous studies that have highlighted the existence of a "tropical pipe" region of relatively isolated air at the equator (*Trepte and Hitchman*, 1992; *Plumb*, 1996), and is underlined by the methane gradient that can be seen in this area in both panels of figure 4.8, indicating another mixing barrier. More recently, *Neu et al.* (2003) have discussed the tropical boundary in more detail. Considering that the present study uses data from model simulations, while the analysis of the previous work mentioned is largely based on measurements, it is noteworthy that the results show similar seasonal variations and altitudinal structure. This will be discussed in more detail in chapter 5.

The L-shaped region of intermediate RE_{z10} north of the white band in figure 4.7 compares well with a similarly shaped area of high effective diffusivity[4] in the bottom panel of plate 4 in *Haynes and Shuckburgh* (2000) (reproduced in figure 4.12), relating to the distinctness of surf zone and vortex until vortex break-up . The separation of air masses between high- and mid-latitudes, and the merging of them at the end of the year, can be clearly seen in both figures. The area of high effective diffusivity in figure 4.12 (and low relatively low RE_{z10} in the top panel of figure 4.7) at mid-latitudes, is a region of relatively low (westerly) wind speeds. Here, planetary wave breaking occurs throughout winter and spring, producing the well-known mid-latitude surf zone first described by *McIntyre and Palmer* (1984).

In addition, the general RE-patterns in figures 4.7 and 4.8 are qualitatively similar to the findings of *Allen and Nakamura* (2001) and *Garny et al.* (2007) (not shown), who use equivalent length and finite time Lyapunov exponents, respectively. Although such approaches for analysing mixing are very different and significantly more computationally demanding than the derivation of RE, the dynamic patterns are similar in all these studies. This strongly suggests that the RE is a promising tool for the analysis of mixing and transport in climate model data.

---

[4] Effective diffusivity is a measure of the diffusion of a tracer on isentropic surfaces, based on equivalent length. It is associated with non-advective mixing (*Haynes and Shuckburgh*, 2000).

4.3 Conclusion

The RE statistical measure was tested on artificial data, the logistic map and two years of SOCOL tracer data (methane was used in this case). An optimal binning algorithm (*Knuth*, 2006) was employed for creating zonal PDFs of ten consecutive days of methane data, from which the RE was calculated for every day (moving 10-day window). The results indicate that RE is a promising tool for identifying mixing barriers in the stratosphere, such as the Southern Hemisphere polar vortex. In Southern Hemisphere winter and early spring, RE produces patterns similar to those found in other studies of stratospheric mixing, e.g., *Allen and Nakamura* (2001), *Haynes and Shuckburgh* (2000), and *Lee et al.* (2001). Bands of high entropy are found centred around the strongest winds of the polar night jet. The strong wind current poses a significant barrier to stratospheric mixing in winter, isolating the polar regions from mid-latitudes. This suggests that the RE can be used as a novel way to analyse mixing in the stratosphere from tracer data. High values of RE are indicative of strong fluctuations in the tracer distributions associated with gradients in the vicinity of mixing barriers, while lower values suggest more thoroughly mixed air masses.

While comparable results are obtained when using Shannon entropy ($\alpha = 1$ in equation 4.1), it can be argued that results with $\alpha = 2$ offer the best compromise between large-scale consistency and sensitivity to small-scale variability. RE is well suited for analysing changes in the PDFs of all kinds of tracer data, such as zonal ten-day (RE_{z10}) and monthly (RE_{zm}) tracer fields. The results clearly show the separation of high- and mid-latitude air masses in the Antarctic stratosphere, associated with the strong polar-vortex in winter and early spring. Our analysis shows RE has an advantage over previous techniques (e.g., equivalent length, effective diffusivity, and Lyapunov exponents) because it is data driven and requires significantly less computational effort; given that the previous measures of mixing generally rely on the advection of air parcels or artificial tracers by modelled or reanalysed wind fields. Using the RE measure for comparing observations of tracer distributions with those from model simulations potentially presents a new approach for analysing mixing in the stratosphere that is independent of model wind biases.

It should be noted that at least one of the problems identified in *Haynes and Shuckburgh* (2000) associated with the calculation of effective diffusivity (*Nakamura and Ma*, 1997) from observations rather than simulations is also mitigated using RE. Namely, because the RE is based on PDFs of tracer data, gaps in the dataset do not require interpolation which can cause significant errors in the effective diffusivity measurements. RE only requires a suitably large number of points to give accurate entropy estimates. However, more subtle issues associated with sampling the tracer field may need to be considered if the tracer field sampled is geographically biased significantly.

In the following chapter, the analysis of stratospheric dynamics using the RE will be expanded to both hemispheres and to eleven years of SOCOL data.

Chapter 5

Extended Analysis of Rényi Entropy

In the previous chapter, the Rényi entropy (*RE*) statistical measure was introduced as a method for analysing mixing in the stratosphere from tracer probability distribution functions (PDFs). It was applied to methane fields from runs of the SOCOL chemistry-climate model (CCM) (*Egorova et al.*, 2005) for 1980 and 1981 conditions. In the Southern Hemisphere, the *RE* was able to identify mixing patterns similar to those reported by *Haynes and Shuckburgh* (2000), *Allen and Nakamura* (2001) and *Lee et al.* (2001). Thus far, most of the analysis described has been limited to a single pressure level (10 mBar), which is assumed to be representative. In order to test the interpretation of the *RE* in the previous chapter, the results from different pressure levels are inspected in detail in section 5.1. The results are further verified by comparing multiple years in section 5.2, which are then used to create a climatological mean of the *RE* that is discussed in section 5.3.

5.1 *RE* Comparison Throughout the Stratosphere

The seasonal average plots of *RE* for the Southern Hemisphere winter shown in the left panel of figure 4.8 indicate that the *RE* structures identified at the 10 mBar level (top panel of figure 4.7) are not confined to a single pressure level. The band of high *RE* values at the centre of the highest wind speeds stretches up to almost 1 mBar ($\approx$ 45 km altitude) and down to approximately 40 mBar ($\approx$ 22 km altitude). The latitude-time contour plots of *RE* for 1980 SOCOL methane data corresponding to several intermediate pressure levels are shown in figure 5.1.

At the 1 mBar level (top left panel) no clear continuous patterns of *RE* can be identified. Similarly, the wind patterns are noticeably more variable and, while a vortex-like structure can be identified in the Southern Hemisphere winter, it is very distorted. While this region has an overall relatively high *RE*, the pattern is far less consistent than the continuous mixing-barrier found at lower altitudes. Relatively strong westerlies are also found in the Northern Hemisphere winter and early springtime. These are associated with the Northern Hemisphere polar night jet. The wind contours are less coherent than in the Southern Hemisphere, which is due to the strong impact of planetary wave breaking on the dynamics in the Northern Hemisphere (see section 2.1.2). Only the tropical regions, on the average, display relatively lower *RE* values similar to the 10 mBar level (top panel of figure 4.7), but

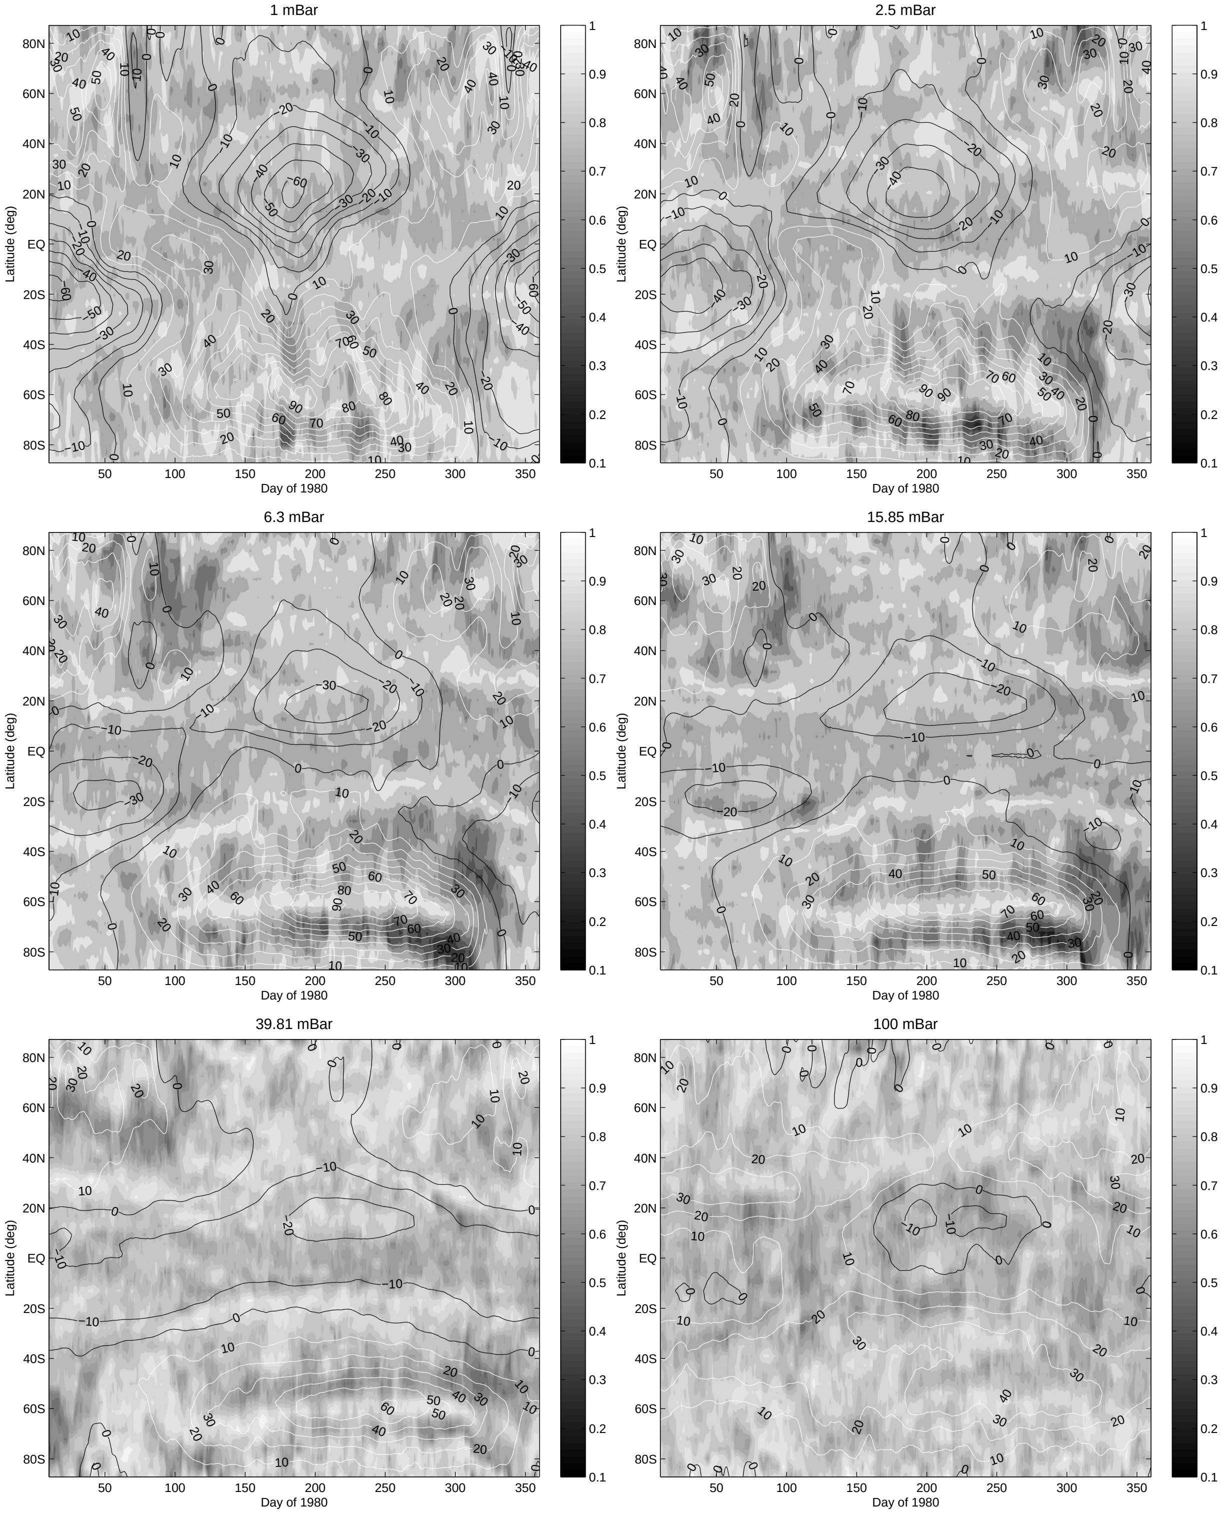

1 mBar

2.5 mBar

6.3 mBar

15.85 mBar

39.81 mBar

100 mBar

Figure 5.1: Comparison of RE patterns at pressure levels between 1 mBar and 100 mBar for 1980 SOCOL methane. Contours are mean zonal winds (white = westerly).

no clear boundary of high *RE* can be identified in either of the hemispheres at 1 mBar.

At 2.5 mBar (≈ 40 km; top right panel of figure 5.1), *RE* displays significantly more structure. The characteristic band of high *RE* at the centre of the, now much smoother, high wind contours in the Southern Hemisphere winter can be identified, along with areas of relatively low *RE* on either side of this band. Additionally, variations in the location of the centre of the tropical low *RE* region can also be observed. However, none of these structures are as clear as at the 10 mBar level. In comparison, the high *RE* bands appear broadened and are interrupted more regularly by lower values. Interestingly, at vortex break-up, the two bands of low *RE* do not merge as observed at lower levels (see below). Relatively high values of *RE* rapidly shift poleward in early November and are then replaced by lower values a few days later. Note that the most intense planetary wave breaking, which weakens the vortex and leads to its final break-up, occurs at high altitudes (*Smith*, 2003). The constant deterioration by planetary waves is likely to make the polar vortex a weaker mixing barrier at high altitudes. Thus, the more variable behaviour of the *RE* in the polar regions is to be expected. Furthermore, a temporary general decrease in the *RE* at mid- and high latitudes can be found in both hemispheres in springtime, when the polar winter low pressure system is replaced by the summer high.

At the 6.3 mBar (≈ 34.5 km) and 15.85 mBar (≈ 28 km) levels (left and right panels in the middle of figure 5.1, respectively), the structures strongly resemble those identified for the intermediate 10 mBar level. At these levels, consistently high *RE* bands can be found at the centre of the polar night jet during Southern Hemisphere winter and early spring. In November, when the vortex breaks up, the high *RE* band disappears and the two areas of relatively low *RE* on either side of the polar night jet merge. Furthermore, the seasonal oscillation of the tropical low *RE* area and its encompassing high *RE* regions is clearly visible at these pressure levels, though the Northern Hemisphere boundary is rather indistinct, particularly in the middle of the year. Generally, the Northern Hemisphere does not display very coherent patterns, but some structures similar to those seen in the Southern Hemisphere are suggestive of similar dynamics, e.g., the decrease of *RE* polewards of $30°N$ in March and April after vortex break-up resembles the dynamics in Southern Hemisphere November. The low values of *RE* indicate relatively well mixed air masses, between 30° and 60° in both hemispheres during spring. In the Southern Hemisphere, the low values of *RE* can even be observed throughout winter. As mentioned in section 2.1.2, this is the region where the continuous planetary wave breaking occurs that drives the Brewer-Dobson circulation. In this "surf zone" (*McIntyre and Palmer*, 1984), mixing is increased, which explains the lower *RE* values. The associated PDFs have a strong peak, as most of the air will have approximately the same methane mixing ratio (see also section 5.3).

As expected from the findings from figure 4.8, the polar and tropical patterns can be seen down to 39.81 mBar (bottom left panel of figure 5.1), with gradually decreasing wind speeds. At this level, the structures, while still being visible, appear diffuse, indicating a weakening of the processes that create the structures. The behaviour of *RE* at vortex break-up is also slightly different from that observed at higher levels. The band of high values continues almost to the end of the year with intermediate values "interrupting" it briefly in November. Additionally, the tropical boundary is still visible in the Southern Hemisphere and at this level is surrounded by the 0 and -10 m/s (easterly) wind contours. In the

Northern Hemisphere, little structure can be identified with any certainty, although again some similarity to Southern Hemisphere patterns can be found.

At 100 mBar (≈ 16 km; bottom right panel of figure 5.1), most of the patterns have disappeared, and the wind structure changes noticeably. This is probably related to the fact that the "lowermost stratosphere" (or the tropopause at tropical latitudes) is relatively close to this level (*Holton et al.*, 1995). Therefore the local dynamics are affected by stratosphere-troposphere exchange dynamics *(Stohl et al., 2003).* This perturbs the distinctively stratospheric dynamical patterns observed at higher altitudes, which explains the reduced structure in the *RE* plots.

It is interesting to note that the band of high *RE* associated with the Southern Hemisphere polar vortex appears and disappears earlier at higher altitudes. This is in accordance with the fact that the vortex builds up "from the top" in autumn and spring break-up also progresses from top to bottom (*Palmer*, 1959), which is reflected in the temporal evolution of the wind patterns at different altitudes. The downward propagation of the degradation of the vortex in spring is due to planetary waves depositing most of their momentum at high altitudes (*Rao et al.*, 2003).

These results are in accordance with the literature with respect to the observed polar vortex dynamics (e.g., *Palmer*, 1959; *Rao et al.*, 2003) and the vertical extent of the tropical-pipe region (e.g., *Neu et al.*, 2003) showing that the interpretation of the high *RE* bands as mixing barriers is justified. Furthermore, the structures of *RE* found at 10 mBar are shown to be representative of the *RE* patterns at most levels in the stratosphere, which supports the assumption that *RE* is associated with real-world dynamics, particularly large-scale mixing. It therefore seems that the *RE* metric can be used to analyse dynamics, by providing information about the statistical complexity of long-lived tracer distributions. An association between the complexity of tracer distributions and mixing has been suggested previously by *Haynes and Shuckburgh* (2000). Changes in complexity can be related to changes in the underlying processes and a measure of the complexity of the tracer field is therefore a potentially useful tool. More specifically, the fact that the Southern Hemisphere vortex break-up coincides with a reduction of *RE* at all but the lowest altitudes examined, suggests that there is a direct correspondence between (high) *RE* and vortex strength, which in turn is related to mixing barrier strength (e.g., *Allen and Nakamura*, 2001; *Garny et al.*, 2007). Hence, *RE* is potentially a measure of mixing barrier strength.

5.2 Application to 11 Years of SOCOL Data

In order to further test whether there is a consistent and reliable relationship between bands of high *RE* and stratospheric mixing barriers, the *RE* is applied to all eleven years of available SOCOL data. The simulations are representative of the conditions in the years from 1980 to 1990. Variations of *RE* and whether these correspond to physical variations in the atmosphere is of primary interest. In the previous section, the patterns associated with large-scale mixing dynamics were found to be most distinct at the 10 mBar level. Additionally, the Northern Hemisphere structures are not as clear as those found in the Southern Hemisphere. Therefore, this section focuses on the results from the Southern

Hemisphere 10 mBar level. Latitude-time contour plots of RE are shown in figure 5.2 for the years from 1981 to 1990.

Examination of figure 5.2 shows that the coincidence of high RE values with the highest wind speeds at the centre of the polar night jet is clearly a persistent feature as it occurs each year. Similarly, the band of high RE separating tropical- from mid-latitudes, and the seasonal oscillation of its location, can be seen in all of the plots. Year-to-year changes in the vortex dynamics are always accompanied by matching changes in the RE patterns. Particularly extreme years, e.g., the very early break-up in 1981 or the very late break-up in 1987, are found to be associated with equally early or late disappearance of the band of high RE values. The final break-up of the vortex and the associated final warming is defined as the time of change of the wind regime at very high latitudes (*Rao et al.*, 2003), i.e., when the transition from the winter to the summer circulation, or eastward to westward flow, is observed. In figure 5.2, this is the point in time when the (black) 'zero-line' (= $0\frac{m}{s}$ wind speed) reaches the pole. In all years, the low RE values have replaced the higher ones by this point in time, although in some years the zero-line never reaches the pole. For example, in 1990 (a very strong vortex year) the zero-line still has not reached the pole by the end of December.

It appears that odd entropy patterns can be linked to unusual dynamics in that year. For example, in 1987 the band of high RE is particularly persistent and very high RE values (brightest contour level at approximately $60°S$) can be seen almost until December, while in other years these already disappear in October or early November. The wind contours show that high wind speeds also last until very late in the year, indicating a very strong vortex. This should cause a high degree of isolation of polar regions in that year. This is in accordance with the unusually high ozone mass deficit observed in 1987 (*Huck et al.*, 2005). Ozone depletion is heavily dependent on low temperatures and therefore on the degree of isolation of the polar regions. Thus, very high ozone depletion is indicative of a strong vortex. Hence, the RE measure can at least be linked to vortex strength empirically.

The existence of a connection between RE and mixing barrier strength can also be observed in figure 5.3, which shows a time-latitude contour plot of the RE for 1987 methane, with mean methane contours superimposed. Around day 300 (late October / early November) there is a temporary reduction in the RE values of the vortex band. At the same time, the methane contours dip towards the pole. This means a sudden increase in methane mixing ratio inside the vortex area, which either indicates an influx of air with higher methane concentrations from lower latitudes or strong diabatic descent. Therefore, it seems likely that there was a temporary 'hole' in the mixing barrier, allowing mixing with mid-latitude air masses to occur. Hence, the drop in RE at this time can be connected to the reduction of the barrier strength of the vortex.

As the application of the RE to eleven years of (CCM) SOCOL methane data show multiple consistent features, the results are combined in a climatological mean in the following section.

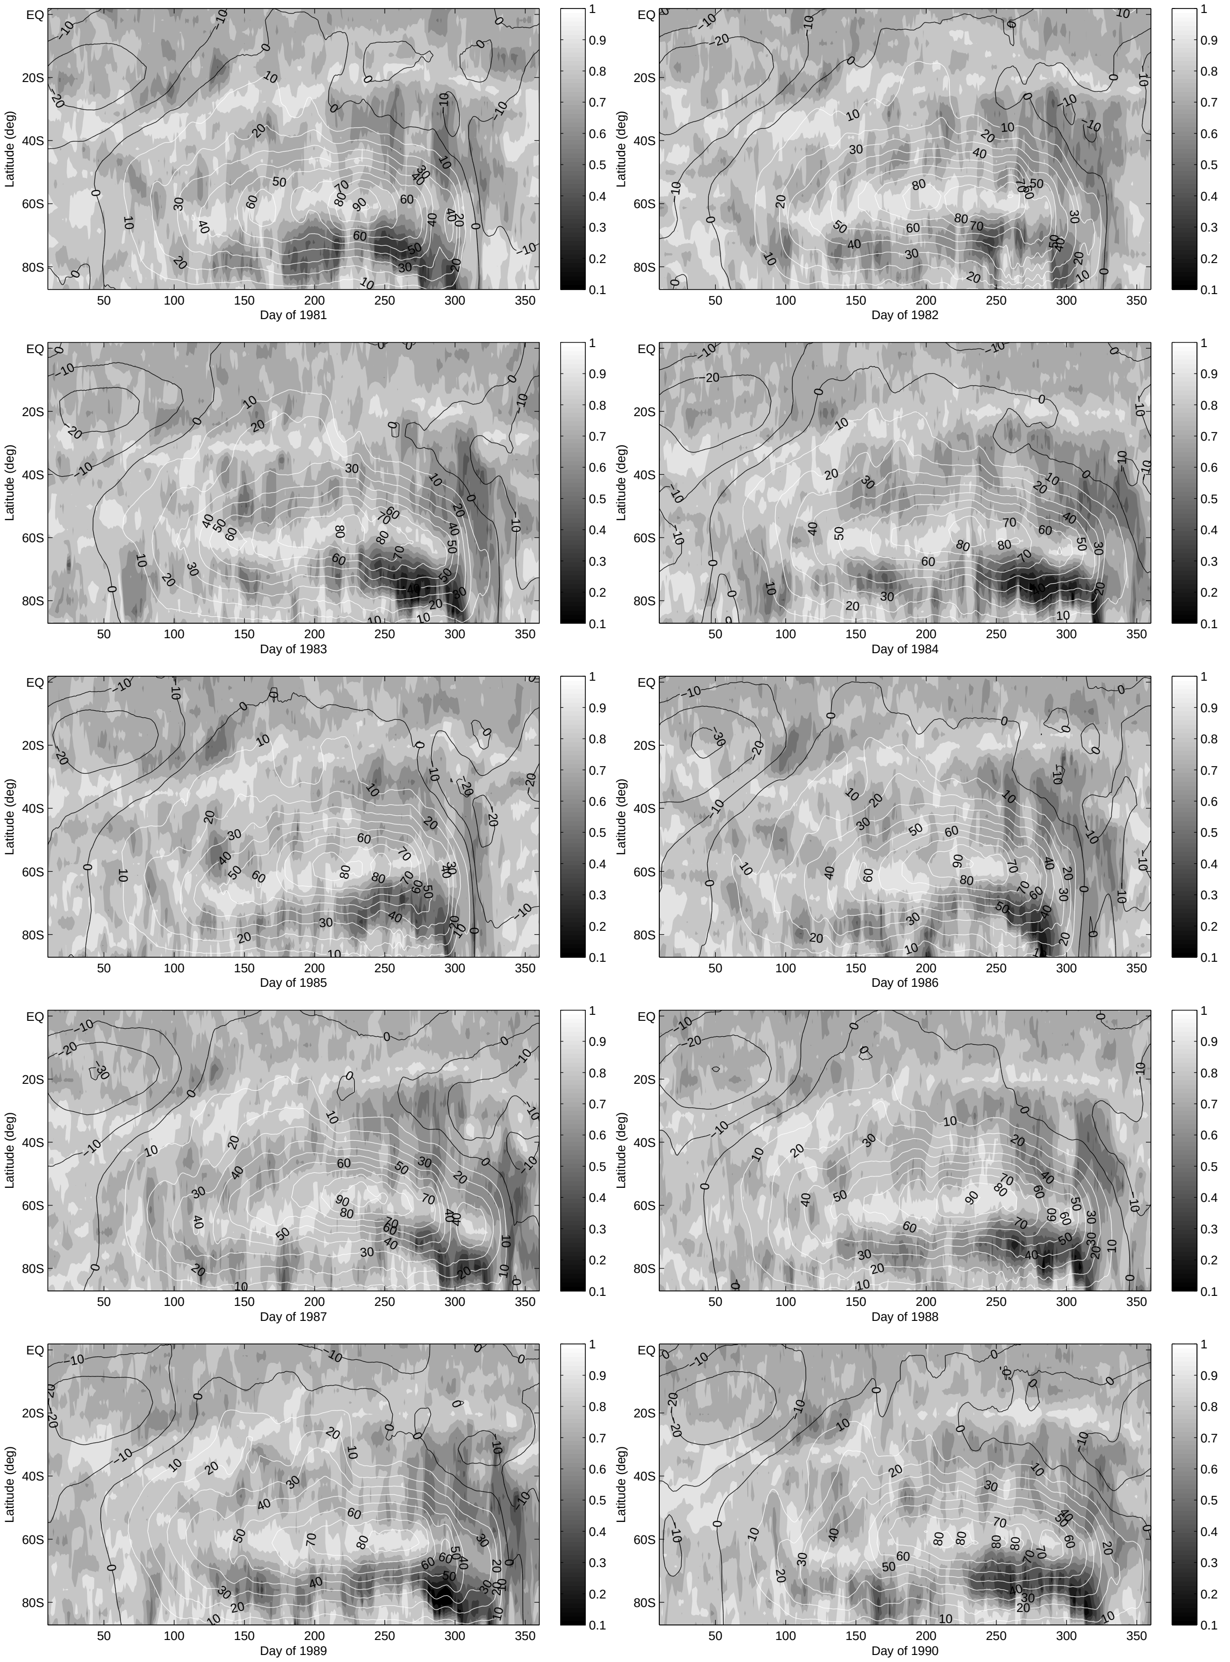

Figure 5.2: *RE* of methane of 1981 to 1990 SOCOL data at 10 mBar.

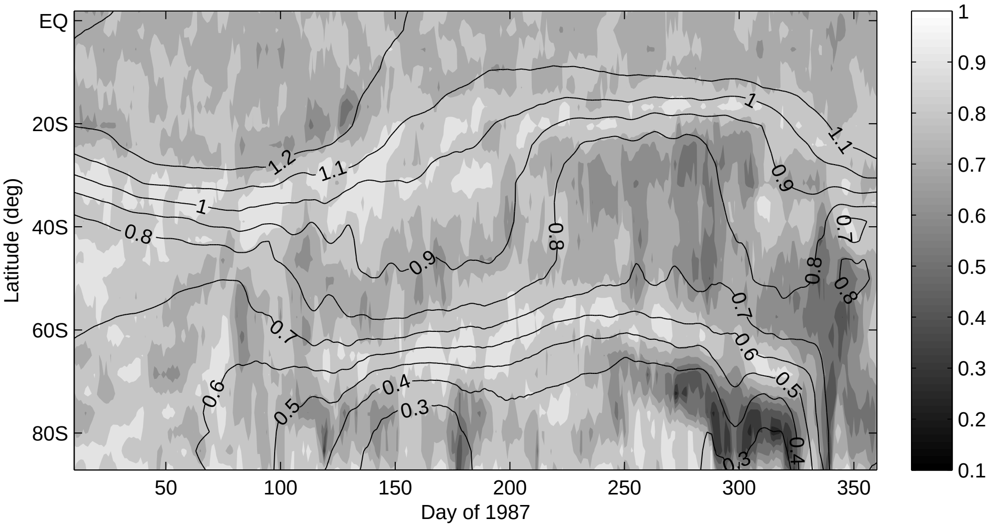

Figure 5.3: *RE* of SOCOL methane for 1987 with mean methane mixing-ratio contours (in ppm) at 10 mBar.

5.3 A Climatology of Rényi Entropy

Using all eleven years of data, a climatological average of *RE* can be calculated. The results for the 10 mBar pressure level are shown in figure 5.4 with mean zonal wind contours (top panel) and mean methane contours (bottom panel). A different colour-map is used to emphasise the difference between this average result and the single years. Additionally the colour range is adjusted to improve contrast as the range of *RE* values is reduced by the averaging process. The structures identified in previous sections are now very easy to observe, as the averaging process removes transient behaviour which may have blurred the pattern in individual years. This supports their association with the recurring circulation features identified previously.

The Southern Hemisphere displays the two bands of high *RE*, previously interpreted as mixing barriers associated with the polar vortex (polar night jet) and the tropical pipe. Additionally, there are two wide bands of relatively low *RE*, one between about $30°S$ and $60°S$ and one at very high latitudes, poleward of the centre of the polar night jet. Both bands begin to stand out at the same time as the high *RE* values associated with the jet. The low values further decrease towards the end of winter until the two bands merge at vortex break-up, replacing the high *RE* band. This development reflects the known dynamics of winter hemisphere mixing in the stratosphere. In autumn, the vortex begins spinning up, creating a mixing barrier between the mid-latitude surf zone and the polar region. It becomes stronger throughout the winter, leading to increased isolation of the polar regions. Significant planetary wave activity during winter gradually creates increasingly more homogeneous tracer distributions in the surf zone (see the gradual poleward migration of the 0.8 ppm contour in the bottom panel of figure 5.4, indicating that most of the surf zone has the same methane mixing ratio), which decreases the overall *RE* in the area. It follows that the vortex-core region is also well mixed by the end of winter, as the high latitudes show even lower *RE* values than the surf zone. In spring, planetary wave breaking slows down the polar night jet westerlies, eventually causing the break-up of the vortex (and the high

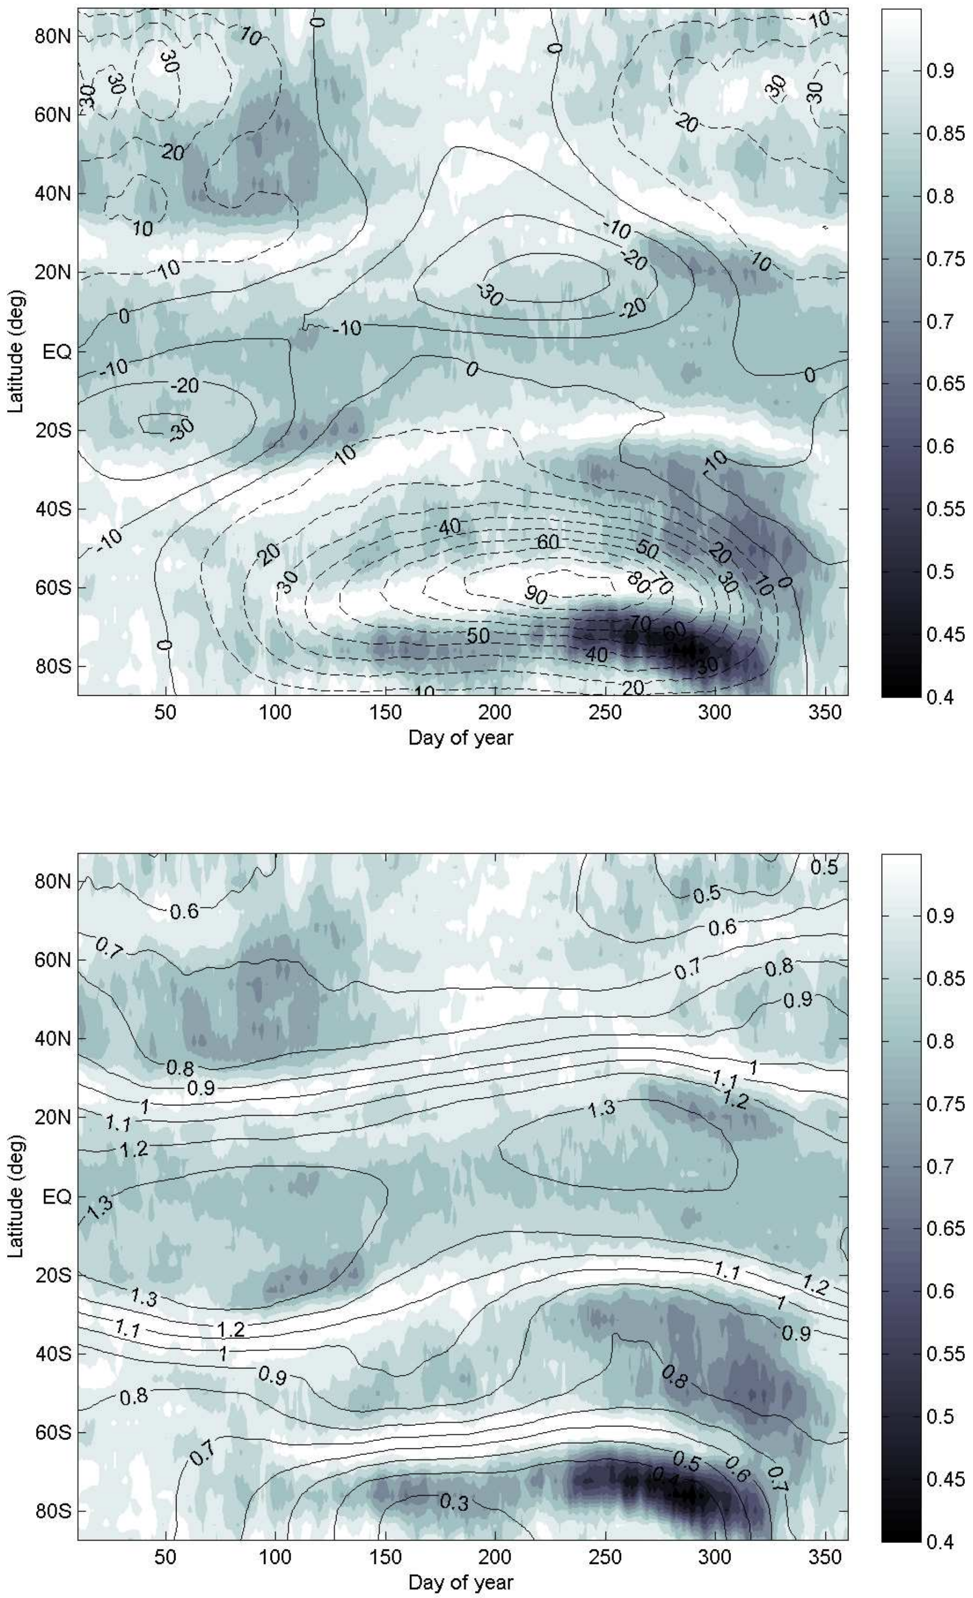

Figure 5.4: Climatological mean of *RE* of methane at 10 mBar with climatological mean contours of wind (top) and methane mixing ratio in ppm (bottom), derived from a SOCOL simulation run from 1980 to 1990.

RE band) in November and allowing mid-latitude air masses to replenish the polar methane concentrations (lowest methane contours disappear). Before the break-up, the vortex-core is known to be relatively well mixed (*Lee et al.*, 2001), which fits the low RE values in this region at the end of the vortex period. However, the known cold bias of the SOCOL model is likely to cause this region to be even more isolated from higher latitudes than is realistic. Therefore, the very low RE values might be rather excessive and could be a subject of future investigation.

In general, the Southern Hemisphere polar vortex is characterised by a band of very high RE values with a width of $5° - 10°$, bordered by regions of relatively low RE on both sides. A qualitatively similar pattern can be found in figure 5.4 in the Northern Hemisphere. With the winter westerlies starting to build up in October and November, a band of about 20° width appears at high northern latitudes (between $55°N$ and $75°N$), consisting of high and sometimes very high values. It is flanked by two regions of middle to low RE, although these are not as low as those witnessed inside the vortex in the Southern Hemisphere. This band moves slightly poleward with time and by January it is centred around $70°N$. While the low values in the surf zone are quite similar to those in the Southern Hemisphere, the very high latitudes do not show as small values in winter or spring. The high-value band of RE disappears around mid March in the Northern Hemisphere at the same time as the westerlies begin to slow down, signalling vortex break-up. While the winds never reach the high speeds occurring in Southern Hemisphere winter, a distorted but relatively similar pattern of wind contours indicates the polar night jet.

Both the variability in the jet winds and in the RE values fits well with the current theories about the dynamics of the Northern Hemisphere polar vortex. Namely that it is significantly more affected by the breaking of planetary waves, which can displace the vortex from the pole, slow down the jet or cause sudden warmings which also slow down the circumpolar winds. In particular, statistics show that major stratospheric warmings are observed in the Northern Hemisphere approximately every two years (e.g., *Schoeberl*, 1978; *Dunkerton and Baldwin*, 1991), while in the Southern Hemisphere the only observed major stratospheric warming occurred in 2002 (*Allen et al.*, 2003). Therefore, the Northern Hemisphere vortex cannot build up as much strength as its southern equivalent and does not pose such a strong barrier to mixing. Hence, it is suggested that the relatively lower values of the RE and the larger width of the high value band, can be associated with the weaker mixing barrier due to reduced wind speeds, and regular displacement of the centre of the vortex from the pole. Both of these effects are likely to be associated with planetary wave activity (*Waugh and Randel*, 1999).

The RE structures previously associated with the tropical pipe and its boundaries, also show up more clearly in the climatological average. Even in the Northern Hemisphere, a tropical band of very high RE, with seasonally varying latitudinal position, is clearly visible for most of the year. This further supports the interpretation of RE being a measure of the complexity created by large-scale mixing dynamics and regions of high RE being associated with mixing barriers. It also shows that the high RE band around $60°S$ is not simply an artifact created by strong winds, as there are no strong winds that could explain the tropical high RE band. The fact that the RE patterns in the Northern Hemisphere show up much more clearly in the climatological mean than in any single year, can probably be attributed to the significantly higher day-to-day variability, caused by stronger planetary wave activity compared to the Southern Hemisphere.

Another interesting feature of the climatological mean RE (see figure 5.4) is the behaviour of RE at mid- and high latitudes in the summer months. During this time, the latitudes above 40° in both hemispheres are characterised by high values of RE without any detailed patterns. Examination of the methane contours in the bottom panel of figure 5.4 shows that these latitudes are relatively well mixed throughout the summer (most of the area stays within the same methane contours). However, RE seems to detect a qualitative difference between the 'well-mixedness' of the surf zone in winter (low RE) and the 'well-mixedness' of the summer hemisphere (high RE). This could indicate different mixing mechanisms and/or timescales dominating during these periods, which is potentially linked to the absence of planetary wave activity in summer.

5.4 Conclusion

The analysis of stratospheric mixing and transport, and the positions of mixing barriers derived from methane tracer data using the RE statistical measure, introduced in chapter 4, has been extended to a larger data set. The RE, calculated from SOCOL methane data at multiple stratospheric altitude levels and for multiple years, shows that the structures identified in chapter 4 are large-scale recurring features. These can be associated with previously identified large-scale atmospheric regions of different mixing characteristics, such as the winter polar vortices, the surf zones, and the tropical pipe *(Plumb, 2002).* Areas of low RE values can be associated with well mixed air masses, while high RE values indicate unmixed air due to low dynamic variability or the presence of mixing barriers.

Analysis of the RE of 11 years of SOCOL data (representative of conditions between 1980 and 1990) at the 10 mBar pressure level in the Southern Hemisphere shows that the RE structures exhibit realistic inter-annual variability. For example, the band of high RE values associated with the mixing barrier created by the Antarctic polar vortex, varies in strength and duration in a similar fashion to the (high) wind speeds of the polar night jet. In years with relatively short vortex periods, e.g., in 1985 (middle left panel in figure 5.2), the band of very high RE values (centred around $60°S$) disappears relatively early, as do the very high wind speeds. Similarly, this band of high RE can be much more persistent in years with a strong, long-lasting polar vortex, e.g., in 1987 (lower middle left panel in figure 5.2).

While the structures associated with the Southern Hemisphere polar vortex are the most salient features identified by RE, the Southern Hemisphere surf zone and the tropical boundary are also observed each year. The variability of the tropical pipe region and its boundary is surprisingly low. This may be explained by the fact that the version of the SOCOL model used for this study exhibits no QBO (Quasi-Biennial-Oscillation). The QBO has a significant effect on the location of the tropical boundary (*Neu et al.*, 2003). Therefore, the observed low variability of the tropical high-RE region may not be representative of real conditions.

The Northern Hemisphere on the other hand, does not display such clear annual patterns, although qualitatively similar structures can be observed (see figure 5.1). The clarity of these structures is significantly improved by calculating the climatological average of the RE field. In plots of the 11-year average of the RE, the Southern Hemisphere tropical, mid-latitude, and polar structures stand out clearly and very similar Northern Hemisphere structures are also observed (see figure 5.4). Accordingly, these can be associated with the Northern Hemisphere tropical boundary, winter surf zone, and polar vortex. The observation that these structures are so much weaker in any single year than their Southern Hemisphere counter parts, can be explained by the significantly higher planetary wave activity in the Northern Hemisphere. This can be expected to cause higher dynamical variability and reduced zonal symmetry, particularly of the polar vortex region.

The fact that the main features of stratospheric dynamics can be associated with distinct structures of RE, strongly suggests that the RE is a useful metric for analysing stratospheric mixing and its variability.

Chapter 6

Final Conclusions and Potential Future Work

6.1 Summary and Final Conclusions

The aim of this project was to investigate the utility of mathematical complexity measures for the analysis of stratospheric dynamics. In particular, the statistical entropy measures sample entropy (SE) and Rényi entropy (RE) were examined. Through theoretical considerations and tests with artificial data sets, e.g., the iteration of the logistic map, suitable parameters were determined for the application of the two measures to stratospheric tracer data. Stratospheric tracers have been previously used to identify large-scale atmospheric regions with different mixing properties, such as the winter polar vortices, the "surf zones", and the "tropical pipe" *(Plumb, 2002)*. Therefore, the SE and the RE were applied to methane (a long-lived tracer) data from simulations of the SOCOL chemistry-climate model (CCM).

The SE is most useful for quantifying the variability of recurring patterns in a time series, by comparing sections of m consecutive data points of the series and quantifying the variability with respect to a tolerance interval r (usually given in terms of the standard deviation, σ, of the time series). Using the parameters $m = 2$ and $r = 0.4\sigma$, the SE was able to identify tropical patterns similar to those reported by studies of the tropical pipe region (*Neu et al.*, 2003). However, the SE statistical measure was found to be particularly unsuitable for use in polar regions and therefore cannot be used to analyse climate complexity on a global scale. This turned out to be related to a general problem of the SE measure when calculated for non-stationary data. Methane distributions at extra-tropical latitudes show deep troughs due to the impact of the polar vortices on mixing during wintertime. As the mixing processes in these regions are very important, i.e., they have a significant impact on the global climate, a different approach to quantifying dynamic complexity in the stratosphere was taken.

Sparling (2000) suggests the use of probability distribution functions (PDFs) of tracer concentrations as an alternative point of view on stratospheric processes. PDFs can be created from almost any kind of data set, and offer a way to avoid problems with discontinuities and non-stationary behaviour. Therefore, the focus was turned to the RE, which is a complexity measure of PDFs. Using the second order $RE(\alpha = 2)$ and a normalisation factor (see equation 4.1), zonal PDFs of ten consecutive days of methane data were created with an optimal binning technique (*Knuth*, 2006). From these, the RE was calculated for every day using a "moving 10-day window".

The results were published (see *Krützmann et al.*, 2008) as they indicate that the RE is a promising tool for identifying stratospheric mixing barriers, such as the Southern Hemisphere polar vortex. In Southern Hemisphere winter and early spring, the RE produces patterns similar to those reported in other studies of stratospheric mixing, e.g., *Allen and Nakamura* (2001), *Haynes and Shuckburgh* (2000), and *Lee et al.* (2001). Bands of high RE are found centred around the strongest winds of the polar night jet. The strong wind current poses a significant barrier to stratospheric mixing in winter, isolating the polar regions from mid-latitudes. High values of RE are indicative of the strong fluctuations in tracer distributions associated with relatively unmixed air in general, and with gradients in the vicinity of mixing barriers, in particular. Lower values suggest more thoroughly mixed air masses.

The validity of the RE-analysis was underlined by showing that qualitatively similar patterns can be seen when using observational satellite data (EOS-MLS data) of a different tracer. This suggests that RE can be used as a novel way to analyse mixing in the stratosphere from tracer data. Its major qualities are summarised in the following:

- The minimum data length required for a reasonably accurate estimate of the RE lies between 150 and 200 data points. This makes the RE applicable to relatively small data sets compared to other statistical measures.
- When applied to stratospheric tracer data, the RE shows distinct structures that exhibit realistic inter-annual variability. Patterns identified with the Antarctic polar vortex mixing barrier are the most prominent features and consistent with observations.
- RE seems to be a useful way to identify mixing barriers. Bands of high RE values can be observed coincident with two mixing barriers of very different origin (the polar vortex and the tropical pipe). In the Southern Hemisphere, the barriers and other structures observed, e.g., the well-mixed surf zone, match relatively well with the results of previous alternative methods of studying mixing, such as effective diffusivity.
- The use of PDFs makes RE particularly robust with respect to data gaps, as long as the total number of data points is sufficiently large. This solves a problem identified in *Haynes and Shuckburgh* (2000) associated with the calculation of effective diffusivity (*Nakamura and Ma*, 1997) from observations rather than simulations. For effective diffusivity calculations from measurements, gaps in the data set require interpolation which can cause significant errors. RE only requires a suitably large number of points to give accurate estimates of the entropy. However, more subtle issues associated with sampling the tracer field may need to be considered if the tracer field sampled is significantly geographically biased. Namely, because the RE is based on PDFs of tracer data.

- Compared to previous techniques (e.g., equivalent length, effective diffusivity, and Lyapunov exponents), the *RE* has the advantage that it requires significantly less computational effort, given that the previous measures of mixing generally rely on the advection of air parcels or artificial tracers by model wind fields. Hence, using the *RE* measure for comparing observations of tracer distributions with those from model simulations potentially presents a new approach for analysing mixing in the stratosphere that is independent of model wind biases.

6.2 Potential Future Work

The application of the *RE* statistical measure to SOCOL CCM methane data shows distinct structures that can be associated with every main feature of stratospheric dynamics. This strongly suggests that the *RE* is a useful new metric for analysing stratospheric mixing and its variability. It has the unique feature that it can be used to derive dynamical information from model or measurement data without relying on any other input such as wind fields. Furthermore, *RE* is very robust with respect to data gaps, which makes it ideal for application to observations.

Accordingly, future work should focus on improving the understanding of mixing in the stratosphere by applying *RE* to observational data of other tracers, e.g., more and improved MLS data, and comparing them with more SOCOL model simulations. In order to make these comparisons, it might be useful to transform the SOCOL tracer fields into potential temperature and equivalent latitude coordinates. While these coordinates are very common for studies of atmospheric dynamics, pressure and latitude were used for this initial analysis, as the derivation of the dynamic coordinates introduces further uncertainties to the data, e.g., through multiple interpolation (*Randel et al.*, 1998). Although some uncertainties, such as offsets, are unlikely to affect the *RE* – it depends solely on the probabilities of certain values, not the values themselves – others, such as fast, small-scale variations (e.g., at low stratospheric altitudes), seem to cause unusually high values of the *RE*. Similarly, very high *RE* values found at high polar altitudes in summer could be linked to diabatic descent. The reasons for this are also an area of potential future investigation.

Previous studies, such as *Garny et al.* (2007), have examined the variation of mixing barrier strength as a function of time and shown a trend in the strength of the barrier. The results of the present study suggest a possible link between the *RE* and vortex strength in the Southern Hemisphere. Given the shorter data set available for the present study, no conclusions can be made on this trend. Therefore, a long-term trend analysis from an extended model run may be a useful area of further research. By using an updated version of the SOCOL CCM that exhibits a QBO (Quasi-Biennial-Oscillation), this could also include a long-term trend analysis of the tropical pipe region and its variability.

Another area of potential future work is the investigation of the deviations from the climatological mean of the *RE*. This approach is more promising for the investigation of Northern Hemisphere dynamics and could potentially be useful for long-term trend analysis as well.

Shukla et al. (2006) have used entropy measures in a different context, namely to examine the fidelity of climate models. Hence the application of complexity measures to model data in order to examine whether they may be used to test the underlying reality of the dynamics in these models, is another possible application of the statistical entropy measures introduced in this study. By comparing model entropy fields with entropy fields from satellite observations, a test of the mixing processes parameterised in models is possible and unexplained differences can be indicative of previously unknown phenomena.

Acknowledgements

First of all, I want to thank my supervisors Dr. Adrian McDonald and Dr. Steve George for providing the idea for this research and for giving me the opportunity to pursue it. Their help and support throughout this year was great and made it a very fruitful experience.

I would also like to thank the University of Canterbury and particularly the Department of Physics & Astronomy for Master's funding, thereby enabling me to pursue this project.

Furthermore, I want to express my greatest thanks to Robyn Nightowl Woollands for making me run and for the full on sprint at the end of this thesis work, Cpt. Clare Worley for comments and comas, (K)Athleen Monahan for lots of brownies and new words, André LyX Geldenhuis for his middle name and associated support, and Matthias Mason Danninger for supplying physical knowledge and physical education. Without their 24/7 support throughout this year and particularly during the final stages of this thesis, it would not have been completed.

I also want to acknowledge Ernesto Illy, who passed away in February this year. Conducting this research and the creating this thesis was greatly supported by the constant assistance of his produce. But the connection even goes beyond the pure consumption of liquid inspiration, as his publication on "The complexity of coffee" (*Illy*, 2002) has also provided scientific stimulation.

Finally, I want to thank my "godfather" Claus C. Schroeder, for essentially getting me into this whole chaos. It's all your fault!

Without all these people, this whole project would not have been possible, much less successful. Once again, THANK YOU ALL!!!

Bibliography

Alcaraz, R., and J. J. Rieta (2007), Bidomain sample entropy to predict termination of atrial arrhythmias, *2007 IEEE International Symposium on Intelligent Signal Processing, Conference Proceedings Book*, pp. 259–264.

Allen, D. R., and N. Nakamura (2001), A seasonal climatology of effective diffusivity in the stratosphere, *Journal of Geophysical Research-Atmospheres*, *106*(D8), 7917–7935.

Allen, D. R., R. M. Bevilacqua, G. E. Nedoluha, C. E. Randall, and G. L. Manney (2003), Unusual stratospheric transport and mixing during the 2002 Antarctic winter, *Geophysical Research Letters*, *30*(12), 1599, doi:10.1029/2003GL017117.

Andrews, D. G. (2000), *An introduction to atmospheric physics*, Cambridge University Press, Cambridge, UK; New York, NY.

Baldwin, M. P., D. W. J. Thompson, E. F. Shuckburgh, W. A. Norton, and N. P. Gillett (2003), Weather from the stratosphere?, *Science*, *301*(5631), 317–318.

Baldwin, M. P., M. Dameris, and T. G. Shepherd (2007), Atmosphere - How will the stratosphere affect climate change?, *Science*, *316*(5831), 1576–1577.

Beck, C., and F. Schlögl (1993), *Thermodynamics of chaotic systems: an introduction*, Cambridge nonlinear science series; 4, Cambridge University Press, Cambridge, UK; New York, NY.

Bialas, A., W. Czyz, and K. Zalewski (2006), The Renyi entropy H-2 as a rigorous, measurable lower bound for the entropy of the interaction region in multi-particle production processes, *European Physical Journal C*, *48*(1), 215–221.

Bowman, K. P., and P. J. Cohen (1997), Interhemispheric exchange by seasonal modulation of the Hadley circulation, *Journal of the Atmospheric Sciences*, *54*(16), 2045–2059.

Brasseur, G., and C. Granier (1992), Mount Pinatubo Aerosols, Chlorofluorocarbons, and Ozone Depletion, *Science*, *257*(5074), 1239–1242.

Brewer, A. W. (1949), Evidence for a World Circulation Provided by the Measurements of Helium and Water Vapour Distribution in the Stratosphere, *Quarterly Journal of the Royal Meteorological Society*, *75*(326), 351–363.

Camesasca, M., M. Kaufman, and I. Manas-Zloczower (2006), Quantifying fluid mixing with the Shannon entropy, *Macromolecular Theory and Simulations*, *15*(8), 595–607.

Costa, M., A. L. Goldberger, and C. K. Peng (2002), Multiscale entropy analysis of complex physiologic time series, *Physical Review Letters*, *89*(6), 068102, doi:10.1103/PhysRevLett.89.068102.

Dobson, G. M. B. (1956), Origin and Distribution of the Polyatomic Molecules in the Atmosphere, *Proceedings of the Royal Society of London Series A - Mathematical and Physical Sciences*, *236*(1205), 187–193.

Dunkerton, T. J., and M. P. Baldwin (1991), Quasi-Biennial Modulation of Planetary-Wave Fluxes in the Northern-Hemisphere Winter, *Journal of the Atmospheric Sciences*, *48*(8), 1043–1061.

Egorova, T., E. Rozanov, V. Zubov, E. Manzini, W. Schmutz, and T. Peter (2005), Chemistry-climate model SOCOL: a validation of the present-day climatology, *Atmospheric Chemistry and Physics*, *5*, 1557–1576.

Everson, R., D. Manin, L. Sirovich, and M. Winter (1998), Quantification of mixing and mixing rate from experimental observations, *AIAA Journal*, *36*(2), 121–127.

Eyring, V., N. Butchart, D. W. Waugh, H. Akiyoshi, J. Austin, S. Bekki, et al. (2006), Assessment of temperature, trace species, and ozone in chemistry-climate model simulations of the recent past, *Journal of Geophysical Research-Atmospheres*, *111*(D22), doi:10.1029/2006JD007327.

Feigenbaum, M. J. (1978), Quantitative Universality for a Class of Non-Linear Transformations, *Journal of Statistical Physics*, *19*(1), 25–52.

Fu, X. W., B. Graham, D. Xuan, R. Bettati, and W. Zhao (2004), Empirical and theoretical evaluation of active probing attacks and their countermeasures, in *Information Hiding*, *Lecture Notes in Computer Science*, vol. 3200, pp. 266–281.

Gallagher, R., and T. Appenzeller (1999), Beyond reductionism - Introduction, *Science*, *284*(5411), 79.

Garny, H., G. E. Bodeker, and M. Dameris (2007), Trends and variability in stratospheric mixing: 1979-2005, *Atmos. Chem. Phys.*, *7*(21), 5611–5624.

Goldberger, A. L., J. E. Mietus, D. R. Rigney, M. L. Wood, and S. M. Fortney (1994), Effects of Head-Down Bed Rest on Complex Heart-Rate-Variability - Response to Lbnp Testing, *Journal of Applied Physiology*, *77*(6), 2863–2869.

Haynes, P., and E. Shuckburgh (2000), Effective diffusivity as a diagnostic of atmospheric transport 1. Stratosphere, *Journal of Geophysical Research-Atmospheres*, *105*(D18), 22,777–22,794.

Holton, J. R. (2004), *An introduction to dynamic meteorology*, 4th ed., Elsevier Academic Press, Burlington, Mass.

Holton, J. R., P. H. Haynes, M. E. McIntyre, A. R. Douglass, R. B. Rood, and L. Pfister (1995), Stratosphere-Troposphere Exchange, *Reviews of Geophysics*, *33*(4), 403–439.

Huck, P. E. (2007), The coupling of dynamics and chemistry in the Antarctic stratosphere: A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Physics in the University of Canterbury, Ph.D. thesis.

Huck, P. E., A. J. McDonald, G. E. Bodeker, and H. Struthers (2005), Interannual variability in Antarctic ozone depletion controlled by planetary waves and polar temperature, *Geophysical Research Letters*, *32*(13), L13819, doi:10.1029/2005GL022943.

Huck, P. E., S. Tilmes, G. E. Bodeker, W. J. Randel, A. J. McDonald, and H. Nakajima (2007), An improved measure of ozone depletion in the Antarctic stratosphere, *Journal of Geophysical Research-Atmospheres*, *112*(D11), doi:10.1029/2006JD007860.

Illy, E. (2002), The complexity of coffee, *Scientific American*, *286*(6), 86–91.

Knuth, K. H. (2006), Optimal Data-Based Binning for Histograms, arXiv:physics/0605197v1.

Knuth, K. H., A. Gotera, C. T. Curry, K. A. Huyser, K. R. Wheeler, and W. B. Rossow (2005), *Revealing Relationships among Relevant Climate Variables with Information Theory*, paper presented at Earth-Sun Systems Technology Conference, Earth Science Technology Office, NASA, Adelphi, Md.

Krützmann, N. C., A. J. McDonald, and S. E. George (2008), Identification of mixing barriers in chemistry-climate model simulations using Rényi entropy, *Geophysical Research Letters*, *35*(6), L06806, doi:10.1029/2007GL032829.

Lake, D. E., J. S. Richman, M. P. Griffin, and J. R. Moorman (2002), Sample entropy analysis of neonatal heart rate variability, *American Journal of Physiology-Regulatory Integrative and Comparative Physiology*, *283*(3), R789–R797.

Lee, A. M., H. K. Roscoe, A. E. Jones, P. H. Haynes, E. F. Shuckburgh, M. W. Morrey, and H. C. Pumphrey (2001), The impact of the mixing properties within the Antarctic stratospheric vortex on ozone loss in spring, *Journal of Geophysical Research-Atmospheres*, *106*(D3), 3203–3211.

Li, S. C., Q. F. Zhou, S. H. Wu, and E. F. Dai (2006), Measurement of climate complexity using sample entropy, *International Journal of Climatology*, *26*(15), 2131–2139.

Lorenz, E. N. (1963), Deterministic Nonperiodic Flow, *Journal of the Atmospheric Sciences*, *20*(2), 130–141.

Matthes, K., Y. Kuroda, K. Kodera, and U. Langematz (2006), Transfer of the solar signal from the stratosphere to the troposphere: Northern winter, *Journal of Geophysical Research-Atmospheres*, *111*(D6), D06108, doi:10.1029/2005JD006283.

McIntyre, M. E., and T. N. Palmer (1984), The Surf Zone in the Stratosphere, *Journal of Atmospheric and Terrestrial Physics*, *46*(9), 825–849.

Nakamura, N. (1996), Two-dimensional mixing, edge formation, and permeability diagnosed in an area coordinate, *Journal of the Atmospheric Sciences*, *53*(11), 1524–1537.

Nakamura, N., and J. Ma (1997), Modified Lagrangian-mean diagnostics of the stratospheric polar vortices. 2. Nitrous oxide and seasonal barrier migration in the cryogenic limb array etalon spectrometer and SKYHI general circulation model, *Journal of Geophysical Research-Atmospheres*, *102*(D22), 25,721–25,735.

Neu, J. L., L. C. Sparling, and R. A. Plumb (2003), Variability of the subtropical "edges" in the stratosphere, *Journal of Geophysical Research-Atmospheres*, *108*(D15), 4482, doi:10.1029/2002JD002706.

Palmer, C. E. (1959), The Stratospheric Polar Vortex in Winter, *Journal of Geophysical Research*, *64*(7), 749–764.

Patmore, N., and R. Toumi (2006), An entropy-based measure of mixing at the tropopause, *Quarterly Journal of the Royal Meteorological Society*, *132*(619), 1949–1967.

Perlwitz, J., and N. Harnik (2003), Observational evidence of a stratospheric influence on the troposphere by planetary wave reflection, *Journal of Climate*, *16*(18), 3011–3026.

Pierce, R. B., W. L. Grose, J. M. Russell, and A. F. Tuck (1994), Evolution of Southern-Hemisphere Spring Air Masses Observed by HALOE, *Geophysical Research Letters*, *21*(3), 213–216.

Pincus, S., and R. E. Kalman (2004), Irregularity, volatility, risk, and financial market time series, *Proceedings of the National Academy of Sciences of the United States of America*, *101*(38), 13,709–13,714.

Pincus, S. M. (1991), Approximate Entropy as a Measure of System-Complexity, *Proceedings of the National Academy of Sciences of the United States of America*, *88*(6), 2297–2301.

Pincus, S. M., I. M. Gladstone, and R. A. Ehrenkranz (1991), A Regularity Statistic for Medical Data-Analysis, *Journal of Clinical Monitoring*, *7*(4), 335–345.

Plumb, R. A. (1996), A "tropical pipe" model of stratospheric transport, *Journal of Geophysical Research-Atmospheres*, *101*(D2), 3957–3972.

Plumb, R. A. (2002), Stratospheric transport, *Journal of the Meteorological Society of Japan*, *80*(4B), 793–809.

Randel, W. J., F. Wu, J. M. Russell, A. Roche, and J. W. Waters (1998), Seasonal cycles and QBO variations in stratospheric CH_4 and H_2O observed in UARS HALOE data, *Journal of the Atmospheric Sciences*, *55*(2), 163–185.

Rao, V. B., M. B. Rosa, J. P. Bonatti, and S. H. Franchito (2003), Stratospheric final warmings in the Southern Hemisphere and their energetics, *Meteorology and Atmospheric Physics*, *83*(1-2), 35–49.

Reiter, E. R., M. E. Glasser, and J. D. Mahlman (1969), Role of Tropopause in Stratospheric-Tropospheric Exchange Processes, *Pure and Applied Geophysics*, *75*(4), 185–218.

Rényi, A. (1961), On measures of entropy and information, in *Proceedings of the 4th Berkeley Symposium on Mathematics, Statistics and Probability 1960*, pp. 547–561, University of California Press.

Richman, J. S., and J. R. Moorman (2000), Physiological time-series analysis using approximate entropy and sample entropy, *American Journal of Physiology-Heart and Circulatory Physiology*, *278*(6), H2039–H2049.

Richman, J. S., D. E. Lake, and J. R. Moorman (2004), Sample entropy, in *Numerical Computer Methods, Pt E, Methods in Enzymology*, vol. 384, pp. 172–184.

Rukhin, A. L. (2000), Approximate Entropy for Testing Randomness, *Journal of Applied Probability*, *37*(1), 88–100.

Ryan, S. M., A. L. Goldberger, S. M. Pincus, J. Mietus, and L. A. Lipsitz (1994), Gender-Related and Age-Related Differences in Heart-Rate Dynamics - Are Women More Complex Than Men, *Journal of the American College of Cardiology*, *24*(7), 1700–1707.

Schoeberl, M. R. (1978), Stratospheric Warmings - Observations and Theory, *Reviews of Geophysics*, *16*(4), 521–538.

Schreiber, T., and A. Schmitz (2000), Surrogate time series, *Physica D-Nonlinear Phenomena*, *142*(3-4), 346–382.

Shannon, C. E. (1948), A Mathematical Theory of Communication, *Bell System Technical Journal*, *27*(3), 379–423.

Shepherd, T. G. (2003), Large-scale atmospheric dynamics for atmospheric chemists, *Chemical Reviews*, *103*(12), 4509–4531.

Shepherd, T. G. (2007), Transport in the middle atmosphere, *Journal of the Meteorological Society of Japan*, *85B*, 165–191.

Shukla, J., T. DelSole, M. Fennessy, J. Kinter, and D. Paolino (2006), Climate model fidelity and projections of climate change, *Geophysical Research Letters*, *33*(7), L07702, doi:10.1029/2005GL025579.

Smith, A. K. (2003), The origin of stationary planetary waves in the upper mesosphere, *Journal of the Atmospheric Sciences*, *60*(24), 3033–3041.

Solomon, S. (1999), Stratospheric ozone depletion: A review of concepts and history, *Reviews of Geophysics*, *37*(3), 275–316.

Sparling, L. C. (2000), Statistical perspectives on stratospheric transport, *Reviews of Geophysics*, *38*(3), 417–436.

Stohl, A., P. Bonasoni, P. Cristofanelli, W. Collins, J. Feichter, A. Frank, et al. (2003), Stratosphere-troposphere exchange: A review, and what we have learned from STACCATO, *Journal of Geophysical Research-Atmospheres*, *108*(D12), 8516, doi:10.1029/2002JD002490.

Sweet, D., and E. Ott (2000), Fractal dimension of higher-dimensional chaotic repellors, *Physica D-Nonlinear Phenomena*, *139*(1-2), 1–27.

Thomas, G. (2003), A New Instrument for Atmospheric Aerosol Measurement: A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in Physics in the University of Canterbury, Ph.D., University of Canterbury.

Thompson, D. W. J., and S. Solomon (2002), Interpretation of recent Southern Hemisphere climate change, *Science*, *296*(5569), 895–899.

Trepte, C. R., and M. H. Hitchman (1992), Tropical Stratospheric Circulation Deduced from Satellite Aerosol Data, *Nature*, *355*(6361), 626–628.

Verhulst, P. F. (1838), Notice sur la loi que la population pursuit dans son accroissement, *Correspondance mathématique et physique*, *10*, 113–121.

Wackerbauer, R., A. Witt, H. Altmanspacher, J. Kurths, and H. Scheingraber (1994), A Comparative Classification of Complexity-Measures, *Chaos, Solitons & Fractals*, *4*(1), 133–173.

Wallace, J. M., and P. V. Hobbs (2006), *Atmospheric science : an introductory survey*, 2nd ed., Academic Press, Burlington, Mass.

Waters, J. W., L. Froidevaux, R. S. Harwood, R. F. Jarnot, H. M. Pickett, W. G. Read, et al. (2006), The Earth Observing System Microwave Limb Sounder (EOS MLS) on the Aura satellite, *IEEE Transactions on Geoscience and Remote Sensing*, *44*(5), 1075–1092.

Waugh, D. W., and T. M. Hall (2002), Age of stratospheric air: Theory, observations, and models, *Reviews of Geophysics*, *40*(4), doi:10.1029/2000RG000101.

Waugh, D. W., and W. J. Randel (1999), Climatology of Arctic and Antarctic polar vortices using elliptical diagnostics, *Journal of the Atmospheric Sciences*, *56*(11), 1594–1613.