INSTITÜT FÜR PHYSIK DER UNIVERSITÄT POTSDAM
UND
POTSDAM-INSTITUT FÜR KLIMAFOLGENFORSCHUNG (PIK)

UNCERTAINTIES IN CLIMATE DATA ANALYSIS

PERSPECTIVES ON WORKING WITH MEASUREMENT ERRORS AND OTHER UNKNOWNS

DISSERTATION
ZUR ERLANGUNG DER AKADEMISCHEN GRADES
"DOCTOR RERUM NATURALIUM" (DR. RER. NAT.)
IN DER WISSENSCHAFTSDISZIPLIN "KLIMAPHYSIK"

EINGEREICHT AN DER
MATHEMATISCH-NATURWISSENSCHAFTLICHEN FAKULTÄT
DER UNIVERSITÄT POTSDAM

VON

BEDARTHA GOSWAMI

POTSDAM, NOVEMBER 2014

Published online at the
Institutional Repository of the University of Potsdam:
URN urn:nbn:de:kobv:517-opus4-78312
http://nbn-resolving.de/urn:nbn:de:kobv:517-opus4-78312

For my parents

Abstract

Scientific inquiry requires that we formulate not only what we know, but also what we do not know and by how much. In climate data analysis, this involves an accurate specification of measured quantities and a consequent analysis that consciously propagates the measurement errors at each step. The dissertation presents a thorough analytical method to quantify errors of measurement inherent in paleoclimate data. An additional focus are the uncertainties in assessing the coupling between different factors that influence the global mean temperature (GMT).

Paleoclimate studies critically rely on 'proxy variables' that record climatic signals in natural archives. However, such proxy records inherently involve uncertainties in determining the age of the signal. We present a generic Bayesian approach to analytically determine the proxy record along with its associated uncertainty, resulting in a time-ordered sequence of correlated probability distributions rather than a precise time series. We further develop a recurrence based method to detect dynamical events from the proxy probability distributions. The methods are validated with synthetic examples and demonstrated with real-world proxy records. The proxy estimation step reveals the interrelations between proxy variability and uncertainty. The recurrence analysis of the East Asian Summer Monsoon during the last 9000 years confirms the well-known 'dry' events at 8200 and 4400 BP, plus an additional significantly dry event at 6900 BP.

We also analyze the network of dependencies surrounding GMT. We find an intricate, directed network with multiple links between the different factors at multiple time delays. We further uncover a significant feedback from the GMT to the El Niño Southern Oscillation at quasi-biennial timescales. The analysis highlights the need of a more nuanced formulation of influences between different climatic factors, as well as the limitations in trying to estimate such dependencies.

Zusammenfassung

Wissenschaftliche Untersuchungen setzen nicht nur eine Formulierung des gegenwärtigen Kenntnisstandes mitsamt ihrer Unsicherheiten voraus, sondern ebenso eine Eingrenzung des Unbekannten. Bezogen auf die Analyse von Klimadaten beinhaltet dies eine präzise Spezifikation gemessener Größen sowie eine durchgängige Berücksichtigung ihrer Messunsicherheiten in allen Schritten der Analyse. Diese Dissertation präsentiert eine analytische Methode zur Quantifizierung der in Paläoklimadaten inhärenten Messunsicherheiten. Ein weiterer Schwerpunkt liegt auf der Untersuchung von Unsicherheiten in der Kopplungsstruktur zwischen Klimafaktoren die bekanntermaßen die globale Durchschnittstemperatur (GMT *global mean temperature*) beeinflussen.

Die Paläoklimaforschung beruht in kritischem Maße auf der Analyse von Proxydaten welche die Klimaentwicklung dokumentieren. Allerdings sind Proxydaten mit inhärenten Datierungsunsicherheiten behaftet. Basierend auf einem generischen Bayes'schen Ansatz wird in dieser Dissertation eine analytische Methode vorgestellt um aus den zeitlich unsicheren Proxydaten zeitlich präzise Folgen korrelierter Wahrscheinlichkeitsverteilungen zu erhalten. Von diesen Verteilungen werden Proxyzeitreihen zusammen mit ihren Unsicherheiten berechnet. Weiterhin wird eine rekurrenzbasierte Methode zur Analyse von Proxydaten entwickelt welche anhand dieser Wahrscheinlichkeitsverteilungen plötzliche Änderungen in der Dynamik des Systems ermittelt. Beide Methoden werden mit Hilfe synthetischer Beispieldaten validiert und mit realen Proxydaten demonstriert. Diese statistische Analyse von Proxydaten deckt unteranderem die Beziehungen zwischen der Variabilität der Daten und zugehöriger Unsicherheiten der Proxyzeitreihen auf. Die Rekurrenzanalyse des Ostasiatischen Sommermonsuns bestätigt die bekannten Trockenzeiten der letzten 9.000 Jahre um 8.200 bzw. 4.400 Jahre vor unserer Zeit und deckt eine zusätzliche Trockenzeit um etwa 6.900 Jahre vor unserer Zeit auf.

Die Kopplungsstruktur zwischen Klimafaktoren die bekanntermaßen die GMT beeinflussen lässt sich als ein verworrenes, gerichtetes Netzwerk mit multiplen Links, welche zu verschiedenen Zeitskalen gehören, darstellen. Speziell ergibt sich eine signifikante Wechselwirkung zwischen der GMT und dem ENSO-Phänomen (*El Niño-Southern Oscillation*) auf einer quasi- zweijährigen Zeitskala. Einerseits beleuchtet diese Analyse die Notwendigkeit Einflüsse verschiedener Klimafaktoren auf die GMT nuancierter zu formulieren, andererseits werden die Grenzen der Quantifizierung dieser Einflüsse aufgezeigt.

Acknowledgments

I am extremely grateful to my supervisor Prof. Jürgen Kurths for first of all helping me to start my doctoral studies under his guidance. His invaluable support during the last three years gave me the courage to constantly form new ideas and pursue them to their conclusions to the best of my capabilities.

I am also deeply indebted to Norbert Marwan, who was always there as a friend, guide, and inspiration, helping me out through confused times, showing which threads are best left alone and which are the ones best picked up. I have to also thank Jobst Heitzig—without whom nothing would have ever been possible, for I would have messed up all the math. Jobst has inadvertently taught me lessons in critical thinking, maths, probabilities, integrals, correlations, courage and honesty that I will carry for the rest of my life. I also thank Kira Rehfeld, for being a friend and colleague from the very start to the very end. I also thank Sebastian Breitenbach for being the lively person that he is, and the discussions I shared with him helped me a great deal in understanding paleoclimate. I cannot thank Niklas Boers and Aljoscha Rheinwalt enough ever, for they became close friends who just happened to work together. Without these people, I cannot even begin to imagine what my PhD and dissertation would have been like.

Furthermore, I have Paul Schultz to thank especially for the German version of the abstract of this dissertation. I am also thankful to my colleagues both past and present—in particular, Nishant, Jonathan, Reik, Peng, Veronika, Lyuba, Deniz, Carsten, Nora, Marc, Kevin, Arghya, Sushma, Anoop, and Praveen. I also wish to thank the Potsdam-Institut für Klimafolgenforschung (PIK) for its vibrant research environment, for its beautiful campus, and for the numerous opportunities that came my way during my studies here. And I also thank Heike, Till, Gabi and Anja for always helping me out at PIK regarding anything even remotely procedural.

I also wish to thank DFG research group HIMPAC (FOR 1380), and the graduate school IRTG 1740/TRP 2011/50151-0, funded by the DFG/FAPESP, for providing financial support.

Finally, I want to thank my parents who set me on this path.

In the end, I want to thank Debarchana Baruah, my friend and partner, for being with me through it all.

List of Publications

This dissertation is partially based on the following publications.

- **B. Goswami**, J. Heitzig, K. Rehfeld, N. Marwan, A. Anoop, S. Prasad, J. Kurths, Estimation of sedimentary proxy records together with associated uncertainty, *Nonlinear Processes in Geophysics* **21** (2014) 1091–1111
  doi:10.5194/npg-21-1093-2014

- **B. Goswami**, N. Marwan, G. Feulner, J. Kurths, How do global temperature drivers influence each other? A network perspective using recurrences, *Eur. Phys. J. Special Topics* **222** (2013) 861–873
  doi:10.1140/epjst/e2013-01889-8

- S. F. M. Breitenbach, K. Rehfeld, **B. Goswami**, J. U. L. Baldini, H. E. Ridley, D. J. Kennett, K. M. Prufer, V. V. Aquino, Y. Asmerom, V. J. Polyak, H. Cheng, J. Kurths, N. Marwan, COnstructing Proxy Records from Age models (COPRA), *Clim. Past* **8** (2012) 1765–1779
  doi:10.5194/cp-8-1765-2012

- **B. Goswami**, G. Ambika, N. Marwan, J. Kurths, On interrelations of recurrences and connectivity trends between stock indices, *Physica A* **391** (2012) 4364–4376
  doi:10.1016/j.physa.2012.04.018

List of Figures

List of Tables

Frequently used notations

$P(a)$	Probability that a random variable A takes the value a, i.e, it is equivalent to $P(A = a)$
$P(a\|b)$	Probability that the random variable $A = a$, given that we know that the random variable B takes the value b.
$P(a, b)$	Joint probability that $A = a$ and $B = b$.
$E[X]$	Expectation of a random variable X with a probability distribution given by $P(x)$, i.e., it is equal to $\int dx\, x P(x)$
$\langle x(t) \rangle$	Time average of $x(t)$
$\varrho^X(\tau)$	Autocorrelation of X at lag τ
$\varrho^{X,Y}(\tau)$	Crosscorrelation between X and $Y(\tau)$ where Y is shifted τ steps ahead of X.
$\mathbf{A}$	Adjacency matrix of a given network.
$\mathbf{R}$	Recurrence matrix of a given dynamical system.

Contents

Chapter 1

Introduction

Uncertainty is fundamentally linked to scientific inquiry. What we know about a system is inevitably connected to *what we do not know* about it, and by how much. At the most basic level, uncertainty is involved by means of measurement, where every measurement by its very nature implicates not only the measured quantity but also the precision with which the measurement was made. Uncertainties also come up in the estimation of derived quantities obtained from the measurements themselves. Every estimate derived from measurements in principle has to reflect the precision to which the estimate can be derived— which is based on one hand on the derivation itself, and on the other, the precision of the measurements. Furthermore, in an epistemological sense, there are far deeper uncertainties that involve an inaccurate description of observations without the knowledges of having done so, or those that involve qualitative variables such as belief, perception, and choice that are not easy to describe quantitatively. In the study of climate change and its impacts on human society, all of these levels of uncertainty play their respective crucial roles. Thus, while it is true for scientific inquiry in general, it is even more vital in the context of climate analyses that uncertainties should be given due attention, and quantified, visualized and communicated accurately (Spiegelhalter et al., 2011; Spiegelhalter, 2014).

In recent years, the quantification (Stainforth et al., 2005; Stott and Kettleborough, 2002) and communication (Katz et al., 2013; Giles, 2002; Maslin and Austin, 2012) of uncertainty related to various aspects of global climate change have received much attention because of its implication in mediating and responding to the impacts of a changing climate. It is important to note that this focus on uncertainties involved in climate studies is fueled not only by the need to highlight uncertainties,

but also by the fact that it is particularly challenging to quantify the uncertainties involved in such research. While we know that in order to grasp the ongoing climate change of our planet and to anticipate the future, we need to have a clear idea of the various mechanisms operating at the global level, along with the mechanisms involved in periods of climate change in our planet's history, it is also true that the Earth's climate and its interactions are inherently complex, nonlinear, and entail a multitude of factors that are difficult to describe precisely. In such studies, quantification and representation of the final results and their uncertainties therefore have to be cautious and nuanced to avoid misleading interpretations.

Moreover, we should note that the term *uncertainty* itself has different meanings in different disciplines of science and even within disciplines, its implications vary from context to context. We can have the uncertainties involved in inter-model comparison scenarios and we can have uncertainties in sampling the parameter spaces of the models themselves. There are also crucial uncertainties that arise from the spread of the global temperature values from different models and under different forcings. A deeper uncertainty—one that is not easily quantifiable—is the uncertainty involved in the selection of the model themselves. However, irrespective of the particular implication of the idea of uncertainty, it is evident that a clear notion of the uncertainties involved in the analyses is critical to how we interpret the results and how much of the results can help in evidence-based policy making. This highlights the need to effectively quantify the uncertainties incumbent in climate studies. As a simple example, consider the fact that, if the uncertainty between two datasets is not quantified properly, an analysis of the correlation between the two datasets can lead to erroneous conclusions.

In contrast to the uncertainties of modeling studies that involve the spread of future predictions of globally relevant climatic variables based on model output, the focus of this dissertation is primarily on the uncertainties of derived quantities based on measured climate data. Uncertainties in climate data analyses such as those presented in this dissertation typically involve a knowledge of the uncertainties of measurement whenever they are non-negligible, and a consequent estimation procedure of various quantities of interest such as events, transitions, power spectra, correlations, and couplings. The goal of such climate data analyses is to finally provide a characterization of the climate system under study along with a

quantification of the uncertainty with which the characterization was made.

Specifically relevant to this dissertation are the uncertainties inherent in our knowledge of the past climates of the earth. Empirical knowledge about the climate of the past is primarily obtained using proxy records. Proxies are measurable quantities that record climatic signals in natural archives such as lake sediments, stalagmites, and peat. Some examples of paleoclimate proxies are isotopic fractions such as δ^{18}O and δ^{13}C, pollen, and dust grain sizes which could represent various climatic variables like precipitation, plant primary productivity, wind strengths, monsoon intensity and so on. The fundamental paradigm of empirical paleoclimate studies involves the measurement of proxy signals along the depth of a sediment core obtained from an archive and then dating the core at various depth levels so that we can ascertain the age of the proxy signal. The final proxy-vs.-time estimate obtained is commonly referred to as a *proxy record*. An analysis of the proxy record can then help us in gaining insight into past climate conditions and variability.

There are two aspects of uncertainty linked to such paleoclimate proxy records. First is the level of uncertainty in determining the extent to which the specific proxy record the climate signal under study. Second, there is an inherent uncertainty in assigning the measured proxy value (along depth) to a specific time point in the past. The second kind of uncertainty is a primary focus of this dissertation. This uncertainty arises due to the fact that the proxy values measured along the depth of an archive must themselves be dated with the use of different dating techniques, and these dating methods typically have non-negligible errors of measurement. Thus, the proxy record itself is a *derived estimate* and not a direct measurement. Moreover, due to the dating errors, the time axis of the proxy record is not error-free, and this makes it particularly challenging to estimate the final proxy record uncertainty as well as to make unambiguous comparisons between different proxy records.

We present in this dissertation a new approach based fundamentally on Bayes' theorem and conditional probabilities using which we estimate the final proxy record along with its associated uncertainty from the two sets of age–depth and proxy–depth measurements. The final proxy record uncertainties incorporate var-

ious sources of error involved in the measurement process. This analysis sheds light on how the uncertainty of dating, the uncertainty of proxy measurements (which are negligible at times), and finally the uncertainty of radiocarbon calibration impact the final proxy uncertainty. We thereafter represent the proxies on an error-free timescale which effectively denotes the precise time at which the climate signal was recorded/deposited in the sediment. The result of our analysis is a time-ordered sequence of probability distributions which tell us the probability of various proxy values at any chosen time point of interest in the past.

It is crucial, however, to interpret the probability distributions obtained from our analysis in a correct manner to avoid erroneous conclusions. The probability distributions are often correlated between any given pair of time points of the past because of correlated uncertainties inherent in estimating a statistical relation between the ages and the depths of the archive. Because of this, interpretations regarding trends and periodicities in the proxy estimates finally obtained have to be made with the utmost caution as, in principle, the probability distributions on their own have no way of providing us the answers to these questions without the precise knowledge of the correlations between them. We thus present a detailed analysis highlighting the various quantities that can be estimated without the knowledge of the correlations and also illustrate which factors in the estimation process contribute to these correlations.

A resulting challenge is to analyze the obtained sequence of time-ordered correlated probability distributions to infer interesting characteristics of the dynamics represented by the proxy record. We adopt the framework of recurrence analysis which exploits the patterns of recurrences of a dynamical system to previous states to infer different dynamical characteristics. We provide a new way in which bounds of the recurrence matrix are constructed from the probability distribution sequence, and these bounds are then used to approximate the unknown recurrence patterns of the system underlying the proxy data. We use the framework of recurrence networks in particular to identify sudden changes in dynamical behavior in the proxy data with the help of recurrence network modularity. Besides being a novel approach to study paleoclimate dynamics, this analysis also highlights the possibility of analyzing paleoclimate data while keeping in sight the uncertainties of the data along every step of the way.

Finally, we present an analysis of historic climate datasets involving the global mean temperature (GMT) and the climatic factors that impact its evolution. The goal of this analysis is to challenge the implicit assumption of independence between the various factors as made in earlier studies. We present an extension of a recurrence based measure of dependence which, when combined with a lagged analysis and a suitable significance test, allows us to quantify the statistically significant influences that various climatic factors exert on GMT and on each other. This study brings to fore the need for a more wholesome understanding of global climatic interactions involving the possibility of multiple interactions at different time delays along with critical feedbacks between climatic systems.

The remainder of the dissertation is organized as follows: Chapter 2 presents an introduction to sedimentary paleoclimate proxy records and the methods of age modeling typically used for their reconstruction. It also presents as a comparison a heuristic, numerical approach to quantify the uncertainty of proxy records which overcomes the limitations of the standard paradigm. Chapter 3 details the Bayesian approach put forward to estimate proxy records along with their associated uncertainty. Chapter 4 presents applications of the method to two synthetic examples as a proof of concept and thereafter uses it to estimate the uncertainty of two Holocene proxies from Lonar lake in central India. It also presents a detailed discussion on various aspects related to proxy uncertainty estimation. Chapter 5 deals with the limits and uses of the marginal probability distributions obtained using the Bayesian approach. It discusses the various cases where these distributions might be used in further analyses, and highlights the sources of correlation between the proxy distributions at two different time points. Next, Chapter 6 presents a new method to estimate the recurrence characteristics of a sequence of marginal probability densities. It uses recurrence network modularity to determine three highly significant dry events in the East Asian Summer Monsoon (EASM) that occurred within the last 9000 years. Finally, in Chapter 7, we estimate the network of global temperature drivers. Our approach is based on recurrence based similarity measures and we construct a weighted, directed network with time delays based on empirical data from the last 120 years. This illustrates the uncertainties in estimating connections among various climatic subsystems.

Chapter 2

Sedimentary proxy records and COPRA

The bulk of our empirical knowledge about past environmental and climatic conditions comes from the study and analysis of "paleoclimate proxies". Proxy variables are those measurable quantities that correlate to unobservable climatic variables such as temperature and precipitation (Wefer et al., 1999). Examples of proxies include grain sizes, pollen, elemental isotopic fractions (δ^{18}O and δ^{13}C), and volcanic ash layers. These proxies can give us clues to various climatic and environmental parameters of interest like wind speed, nutrient content, temperature, precipitation, and carbon dioxide concentration.

Proxies are obtained from "natural paleoclimate archives" such as sediments from lake and ocean floors, stalagmites in caves, peat and bogs, and trees. Amongst these, archives formed by sediment deposition are known as *sedimentary archives* and the proxies obtained from them are called *sedimentary proxies*.

The interpretation of proxy datasets is often not straightforward due to unknown factors (e.g., anthropogenic impacts, internal dynamics, level of correlation with the climatic variable – cf. Blaauw et al., 2007) that might have influenced the proxy. The situation is further complicated by the fact that every proxy record has, within it by construction, errors related to the dating of the archive from which it was obtained. Clearly, proxies cannot be measured directly along time. They are measured along the depth of an archive, and the archive depths must then themselves be dated in a separate set of observations. In this sense, the proxy record as a function of time is a derived estimate, and it is one in which the time axis is not error free, because the radiometric ages of the archive depths have non-negligible errors of measurement. As a consequence of the latter, it is challenging to estimate the

uncertainty of proxies as well as to carry out proxy intercomparisons.

This chapter outlines the typical existing paradigm that is used to construct a proxy-vs.-time curve from the age–depth and proxy–depth measurements. It then highlights the critical drawbacks that are inherent in this paradigm. Lastly, it presents the COPRA (Constructing Proxy Records from Age models) methodology we have put forward in Breitenbach et al. (2012) to overcome the limitations of existing age modeling and proxy construction techniques.

Note that from here on forward, we shall use the term *proxy record* to denote exclusively the set of proxy values (or their probabilities) over time, as opposed to the set of proxy measurements made along the depth of a sediment core.

2.1. The paradigm of constructing proxy records

The construction of a proxy record from a natural archive begins with two sets of measurements made on the paleoarchive. On one hand, the proxy signal itself is measured at a large number of points along the depth of the archive. However, this dataset can only tell us the climatic conditions that were prevalent at various depths of the archive. Since the primary goal is to associate the climatic indicators at different depths to different times in the past, we have to date the archive depths as well. This leads to the second set of measurements which involves a set of ages measured at several archive depths. However, due to various practical limitations of time, money and effort, it is not possible to date every single depth at which the proxy has been measured. Typically, only a handful of age measurements are made from which the ages of the proxy measurements are to be inferred. The most interesting part here is that the dates have non-negligible errors of measurement. This is the reason behind the fundamental problem of paleoclimate proxy record construction:

How can we reliably infer the dates of the proxy measurements with the smaller number of age measurements from the same natural archive?

To put it more mathematically, consider a proxy signal x_j, $j = 1, 2, \ldots, N$ measured in large number along the depths z_j of a sedimentary archive. Next, consider a set

of radiometric age measurements $r_i, i = 1,2,\ldots,M$ (which could be U/Th or ^{14}C ages) measured with errors of measurement δr_i at much fewer number of depths z_i such that $M \ll N$. Thus, the initial set of measurements are the sets of ordered pairs (x_j, z_j) and $(r_i \pm \delta r_i, z_i)$. If we denote the 'calendar age' as $t_k, k = 1,2,\ldots,L$, then the fundamental problem of proxy record construction can be thought of as the search for the best possible transformation/procedure $\mathscr{F}$ such that

$$\mathscr{F} : \left\{(x_j, z_j),\ (r_i \pm \delta r_i, z_i)\right\} \longrightarrow (x_k, t_k). \tag{2.1}$$

Note that the radiometric age r_i is not necessarily equal to the 'calendar' age t_i because the age given by ^{14}C for instance needs to be calibrated with an established 'calibration curve' that accounts for changing levels of ^{14}C in the atmosphere which could offset the measured radiocarbon age from the true calendar age (Reimer et al., 2009, 2013). However, the U/Th dating method does not require such a calibration and are taken to represent the calendar age by themselves.

The existing paradigm for a solution to this problem involves various state-of-the-art techniques that provide a distribution of age values for each depth of the archive. This relation, statistical or otherwise, that relates each depth layer of an archive to a set of ages based on the age–depth measurements is called an *age model* of the archive. In other words, an age model takes as an input the set of radiometric age measurements $(r_i \pm \delta r_i, z_i)$ and returns at best a probability distribution $P_z(r)$ that tells us how likely it is to get a radiometric age of r when the depth under consideration is z. That is, the age modeling step involves a mapping $\mathscr{A}$ from the set of 2-tuples $(r_i \pm \delta r_i, z_i)$ to a probability value between 0 and 1

$$\mathscr{A} : \left(r_i \pm \delta r_i, z_i\right) \longrightarrow P_z(r) \subseteq [0,1] \tag{2.2}$$

From this distribution, the mean or median age estimate for the proxy depths is inferred and associated with the proxy value at that depth. Thus, a proxy-vs.-time dataset is created, and we refer to this typically as the *proxy record*. For the sake of

simplicity, let us consider the case of U/Th dating, where we can identify the radiometric age r with the calendar age t such that the age modeling transformation $\mathscr{A}$ actually gives us the probability $P_z(t)$ of which calendar ages are most likely given a particular depth value z. Then, the final proxy record obtained by associating the mean calendar age with the proxy measurement values can be written as

$$\mathscr{P} : \big(x_j, z_j, P_z(r)\big) \longrightarrow \big(x_j, E_{P_z}[t]\big) \tag{2.3}$$

where $\mathscr{P}$ is the final proxy construction transformation assigning the calendar ages to the measured proxy values along depth on the basis of the age model $\mathscr{A}$, and $E_{P_z}[\cdot]$ is the expectation value estimated with respect to the probability distribution P_z. Taking in to account all three equations 2.1, 2.2, and 2.3, we find that the overall transformation $\mathscr{F}$ is equivalent to a composition of the age modeling step $\mathscr{A}$ and the proxy construction step $\mathscr{P}$

$$\mathscr{F} \equiv \mathscr{P}\big(\,\mathscr{A}\big(\,\{(x_j, z_j),\ (r_i \pm \delta r_i, z_i)\}\,\big)\big). \longrightarrow \big(x_j, E_{P_z}[t]\big) \tag{2.4}$$

The proxy construction paradigm is illustrated schematically in Fig. 2.1, which shows a proxy signal measured along the depth of the sedimentary archive, which is indicative of dry/wet paleoclimatic conditions in the region. Initially, we have a large number of proxy measurements made along the depth of the archive but we wish to transfer these proxy–depth measurements to a proxy–time record. For this, six age measurements are made at various places along the archive's depth and an age model is estimated based on these dating points. The age model is now used to infer the mean or median age of any depth layer of the archive, and this information is used to create the proxy record over time as shown in the bottom of Fig.2.1. However, to be conscious of the fact that the time axis of the final proxy record is not error-free, we mark the dating points on the time axis in the final record along with the error bars to give an indication of the levels of chronological uncertainty involved in various parts of the record.

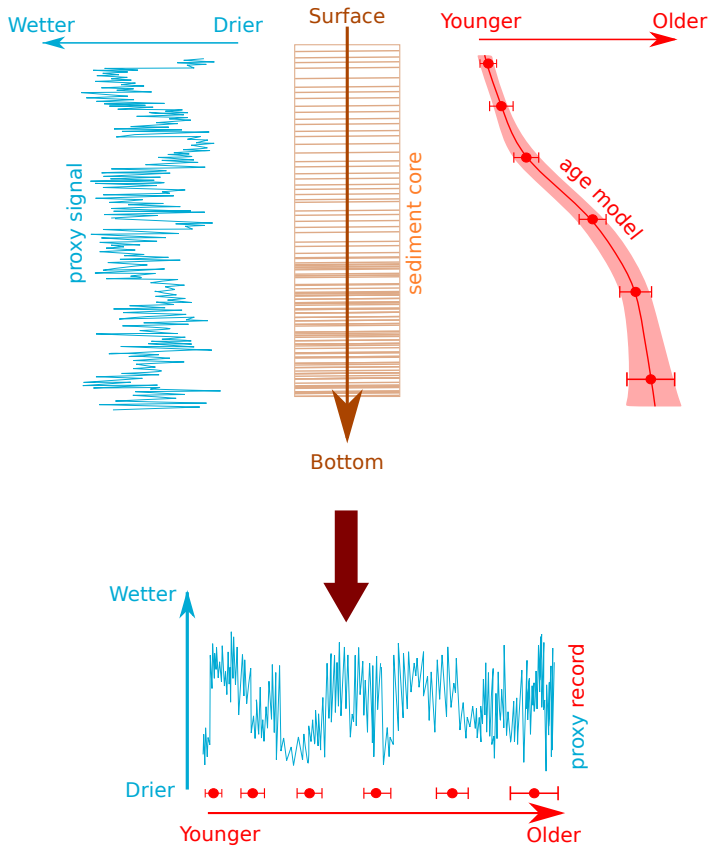

Figure 2.1. Proxy construction paradigm. A large number of proxy measurements (with tiny errors) are made along the depth of the sediment core (blue curve, top–left). A small number of age measurements (with non–negligible error) are made at several points on the archive's depth (red circles with error bars, top–right). From the dating points, an age model is constructed which gives a distribution of ages for each depth layer of the sediment core (red curve with shaded area for uncertainty, top–right). This is then combined with the proxy-depth data to obtain a final *proxy record* over time (bottom row).

The schema outlined in Figure 2.1 is typical to most studies related to sedimentary proxies. It would be slightly different in the case of proxies which take definite binary values, such as the existence/absence of flood layers in the archive. However, the basic procedure in constructing the final proxy record does remain the same. In the next section, we highlight the salient limitations of such proxy record constructions.

2.2. Limitations of the existing paradigm

There are two fundamental limitations in the method of proxy reconstruction as described in the previous section. These are outlined below.

2.2.1. Inadequate representation of uncertainty

The proxy record shown in the bottom of Fig. 2.1 has one critical aspect missing from it—*the proxy uncertainty*. In the current paradigm of proxy record construction, even though we know that the age measurements have non-negligible errors of measurement, the uncertainty carries forward only up to the level of the age model ($P_z(t)$ in equation 2.2) as shown in Fig. 2.1. The propagation of uncertainty stops at this point. In associating the proxy measurement depths with the mean age ($E_{P_z}[t]$ in equation 2.3) estimated from the age model, and thereby representing the proxy record as a single curve, we effectively present an inaccurate picture which seeks to imply that the proxy uncertainty is negligible. Such a representation has the potential to lead to unreliable conclusions in further analyses carried out on the proxy records.

Note that quantifying the uncertainty of the proxy at any particular time in the past is of paramount importance. This is because even though the proxy record is a derived measurement it is still, in its essence, a measurement; and the representation of any measured quantity without an uncertainty (error) of measurement can give rise to to misleading conclusions. For instance, assessment of correlations between datasets can be dramatically influenced by how the uncertainty of the data is represented (see Rehfeld and Kurths, 2014; Heitzig, 2013, Intro.). An additional point is that henceforth in this dissertation we will consider only the uncertainty associated with the estimation of the true value of the *proxy*. The estimation of the climate state from this proxy, or an ensemble of proxies, is a separate problem that follows after the proxy estimation step, and is beyond the scope of the analysis undertaken here.

2.2.2. Uncertain time axis

A second crucial limitation that comes out of the proxy representation in Fig 2.1 is that the time axis is not error-free. Here, the proxy record is not equivalent to a typical time series in which time is a precisely known control variable without errors. Since the age estimates from the age model themselves have distribution $P_z(t)$ of possible values, each point on the time axis is primarily a central estimate $E_{P_z}[t]$ of the age model. Typically, in order to deal with the time uncertainty, the age measurements along with error bars are shown in the same plot as the proxy record and all further inferences based on the proxy record thereafter are made keeping in mind a qualitative understanding of the extent of uncertainty in various portions of the record as denoted by the error bars of the age measurements (Fig. 2.1, bottom).

The major fallout of this is that the comparison of proxies from various places and from different archives becomes extremely non-trivial and cumbersome. Events showing up in a proxy record from an Archive A might appear related to events spotted in a record from another Archive B but their interrelation is not trivially interpretable as we are unable to precisely fix the timing of the events in the first place.

We present here an alternative representation in which the time axis is interpreted as being free of uncertainty and in which the entire burden of uncertainty is transferred on the shoulders of the proxy. We contend that such a representation has the advantage that it allows for clearer inter-record comparison.

2.3. The COPRA representation

Recent studies have raised the issues associated with conventional proxy record representation along with the fact that—although there are various approaches to estimate age–depth relationships for a given archive which help constrain their uncertainties—relatively less focus has been directed towards the propagation of the age uncertainties to the proxy record from the age model (Blaauw et al., 2007). Notably, in Fig. 2 of Blaauw et al. (2007) (reproduced here as Fig. 2.2), we find one

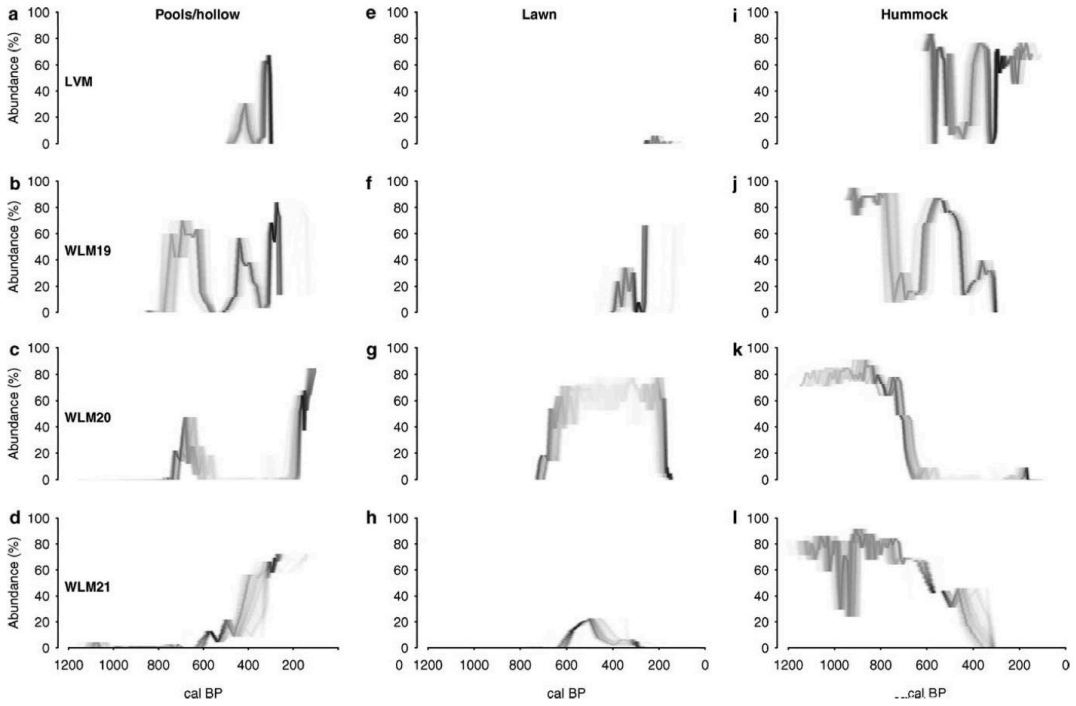

Figure 2.2. Visualization of proxy uncertainties (From Blaauw et al. (2007)). This is one of the first representations of the spread of uncertainty of the final proxy estimates by visualizing the posterior density (as a gray-scale) of proxy records as obtained from a large number of possible age models. COPRA extends this idea further to actually quantify the proxy uncertainty using the variance of the proxy records over time.

of the first representations of the uncertainty of a proxy record (shown as a gray-scale of possible proxy values) that quickly draws attention to the portions of the record that are more reliable and those parts that are not. The authors state that the

> Bayesian methods developed . . . form an important step towards more systematic assessments of links, leads and lags . . . Inherent chronological uncertainties need no longer be neglected in proxy-graphs nor in interpretations.

This idea was further extended and applied in later studies (Parnell et al., 2008; Charman et al., 2009; Blaauw et al., 2010; Swindles et al., 2012). More recently, Mudelsee et al. (2009, 2012) presents ideas on how to estimate the uncertainty of

the proxy record with regard to trend and periodicity estimation of the proxy.

However, even while these studies chose to highlight the non-negligible proxy uncertainties when duly represented, the focus was either on the visualization of the proxy uncertainty or on estimating the proxy uncertainties with respect to specific features such as simultaneity of events, or trend and periodicity estimation. In our approach, reported in Breitenbach et al. (2012), we take this idea forward with COPRA—an acronym for COnstructing Proxy Records from Age models—where we present a heuristic methodology that seeks to quantify and estimate the proxy uncertainty in a more general manner. COPRA is a numerical method that treats age modeling as an intermediate step in the construction of the proxy record, and thereafter represents the final proxy records along with their uncertainties of estimation.

The first step in COPRA is similar to that of conventional age modeling where Monte Carlo based numerics are used to estimate an ensemble of age models given the set of U/Th age measurements $(t_i \pm \delta t_i, z_i)$ against the depths z_i of the sediment core. However, it differs drastically from the existing paradigm in the second step that involves the construction of the proxy record from the age model ensemble. If we consider the age model ensemble as representing a probability distribution $P_z(t)$ similar to what was described in Section 2.1, then the first step in proxy record construction in COPRA involves a transformation $\mathscr{P}_{ens}$ that generates an ensemble of proxy records from each age model in the ensemble underlying $P_z(t)$ such that

$$\mathscr{P}_{ens} : \big(x_j, z_j, P_z(t)\big) \longrightarrow P_t(x) \tag{2.5}$$

where $P_t(x)$ is the probability of the proxy x at every time point t as estimated from the proxy record ensemble.

Based on the proxy record probability distribution $P_t(x)$ we now estimate the mean proxy value $\mu(t) = E_{P_t}[x]$ and the corresponding variance $\sigma^2(t) = E_{P_t}[(x-\mu(t))^2]$ at every time instant t. We thus have an *estimate* $\mu(t)$ of the proxy value at every time point along with a quantifiable uncertainty of estimation $\sigma(t)$. This can then be

thought of as a final transformation $\mathscr{P}_{est}$ which gives us the final proxy record estimates and its associated uncertainties of estimation from the proxy record probability distributions $P_t(x)$:

$$\mathscr{P}_{est}: P_t(x) \longrightarrow (\mu(t), \sigma(t)). \tag{2.6}$$

Thus, in the COPRA approach, the final transformation $\mathscr{F}_{COPRA}$ as introduced in equation 2.1 is equivalent to a composition of three different transformations: the age modeling step, the proxy ensemble step, and finally the proxy estimate step,

$$\mathscr{F}_{COPRA} \equiv \mathscr{P}_{est}\left(\mathscr{P}_{ens}\left(\mathscr{A}\left(\left\{(x_j, z_j),\ (r_i \pm \delta r_i, z_i)\right\}\right)\right)\right) \longrightarrow (\mu(t), \sigma(t)). \tag{2.7}$$

[Note that the final output of the transformations given in equations 2.1 and 2.7 are also quite different, but they are similar in the sense that both attempt to construct a proxy- vs.-time relationship starting from the two sets of given measurements.]

In the COPRA approach, the uncertainty is effectively transferred from the time to the proxy axis, leaving the time axis error-free. We refer to this time axis as an "absolute timescale" in Breitenbach et al. (2012), corresponding to the time of deposition of the proxy material in the archive. However, in this dissertation, we will refer to this as the "error-free timescale" henceforth.

Another point regarding the COPRA representation is that representing the uncertainty on the proxy axis comes at price. The higher frequency variations of the proxy are seemingly washed out in the COPRA representation. Therefore, it effectively shows how much can be inferred from a given proxy record without going astray. The essential differences between a typical proxy representation and the COPRA representation is shown in Figure 2.3 as a schematic.

There are three practical points in which the two proxy record representations differ greatly. First is the uncertainty: in a typical proxy record the age model uncertainty is not propagated to the proxy variable and thus it is the age axis which

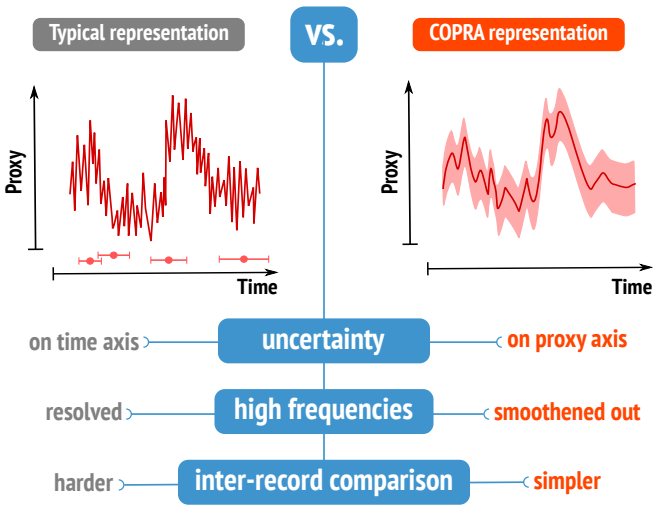

Figure 2.3. The COPRA representation. The essential differences between a typical proxy representation and the COPRA representation is highlighted in this schematic. On the left, we see that a typical representation entails uncertainty on the time axis, shows high frequencies (although time uncertain), but makes inter-record comparison harder. On the right, we find that the COPRA representation puts the uncertainty on the proxy axis, with the trade-off that inter-record comparisons become clearer at the cost that high frequencies are smoothened out.

contains the uncertainty, whereas in COPRA all the uncertainty is transferred to the proxy axis. The second aspect is a consequence of the transfer of uncertainty to the proxy axis: in a typical proxy record, the high frequency proxy variations are retained—in fact, the proxy record variance remains equal to the variance of the proxy over depth—but in the COPRA representation, the high frequency proxy variability is washed out, thereby indicating the extent of our precision in interpreting high frequency paleoclimate change. The third point in which the two representations differ is the clarity of inter-record comparisons. In the COPRA representation, because of the error-free timescale, it becomes relatively simpler to compare features observable in two different records as there is no ambiguity as to the timing of events.

Note: In Chapters 3 and 4, we present an analytical approach that achieves the same kind of proxy representation as COPRA, and we discuss in further detail the

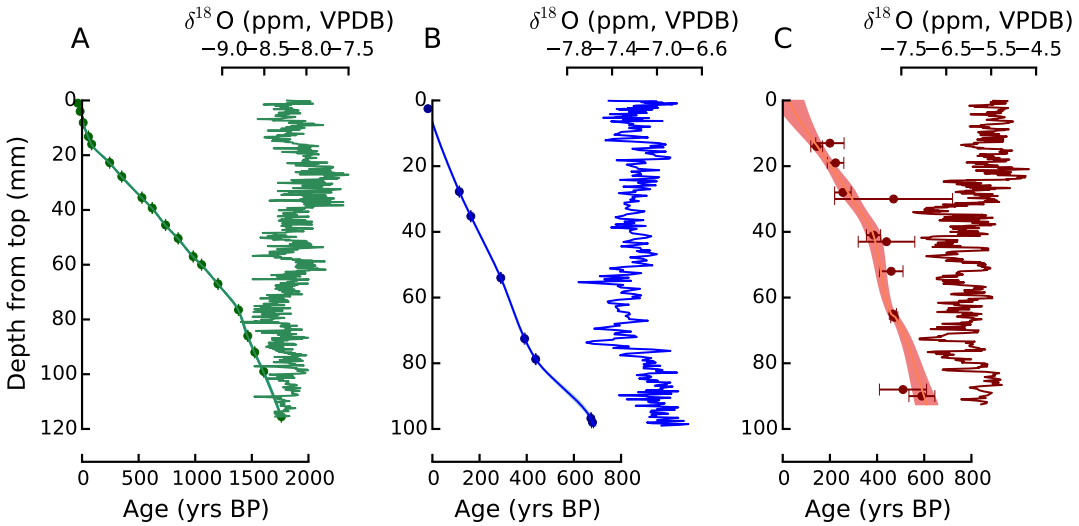

Figure 2.4. Proxy data and age models for three records from Asia. The three datasets from Wanxiang (Zhang et al., 2008) in **A** (green), Dayu (Tan et al., 2009) in **B** (blue), and Wah Shikar (Sinha et al., 2011) in **C** (maroon) are shown. The quality of the age measurements (circles with $\pm 2\sigma$ error bars) vary in all three cases, with Wanxiang being the most precise followed by Dayu, and then Wah Shikar. As a result the age models are also better constrained in that order. The age models are shows as shaded areas but these are barely visible in the case of Wanxiang and Dayu because of the precision of the U/Th ages.

implications of having such a proxy record representation in Section 4.2.

2.4. An illustrative example

To further demonstrate the extent to which typical proxy record representations can potentially lead to erroneous conclusions, we consider in this section three δ^{18}O isotopic proxy datasets from eastern Asia. These isotopic measurements are made on stalagmite cores obtained from caves. The datasets are respectively from Wanxiang (Zhang et al., 2008) and Dayu (Tan et al., 2009) caves in southeast China, and a third record from the Wah Shikar (Sinha et al., 2011) cave in northeastern India. The age–depth and proxy–depth measurements for each of these datasets are

shown in Fig. 2.4. The figure also shows the COPRA age model as a median curve ± 1 standard deviation error as a shaded area overlaying the age–depth measurements. At a first glance, it is clear that the magnitude of the age measurement errors increase on average as we proceed from the Wanxiang data to Dayu to Wah Shikar. Consequently, this order of precision shows up in the COPRA age model as well, with Wah Shikar presenting the largest spread of ages (hence larger age uncertainty) for any given depth layer in the sediment core. Note that the error bars (and thus the age model spread) of the Wanxiang and Dayu ages are not even visible in the plots in Fig. 2.4 due to their extremely small magnitude relative to the timespan of the dataset.

As discussed previously, in a typical proxy record representation this difference the levels of chronological uncertainty between these three records would not be apparent. On the other hand, the COPRA representation would effectively transfer the age uncertainties on the proxy δ^{18}O axis and reveal the reliability of the final proxy record. This difference in the two representations is shown in Fig. 2.5. In the left column of Fig. 2.5, we see the typical representation of the proxies and in the right column we see the corresponding COPRA proxy records along with the uncertainty estimates of the proxy at each time point.

As expected, the uncertainty in the final proxy estimates increase from Wanxiang to Dayu to Wah Shikar (panels B, D, and F of Fig. 2.5). In such a COPRA representation the spread of the age models directly translates to a proportional spread of the proxy records. This leads to a larger uncertainty estimate of the final proxy record, as measured by a dispersion estimate such as the variance. A larger uncertainty also concomitantly leads to a proxy estimate that has less variability on the time axis. This is because, we need higher precision (i.e. lower uncertainty) to be able to reliably extract the fast-varying components of the proxy record.

The final proxy uncertainty in the Wah Shikar record is much higher than that in the other two. Most of the fast varying components are washed out in the Wah Shikar record because of the high level of uncertainty, such that only the slowest trends in precipitation (linked to the δ^{18}O proxy) remain. However,the Wanxiang record, being the best constrained of all the three, reveals high frequency variability as well. In contrast, the typical proxy records in Fig. 2.5 (panels A, C, and E)

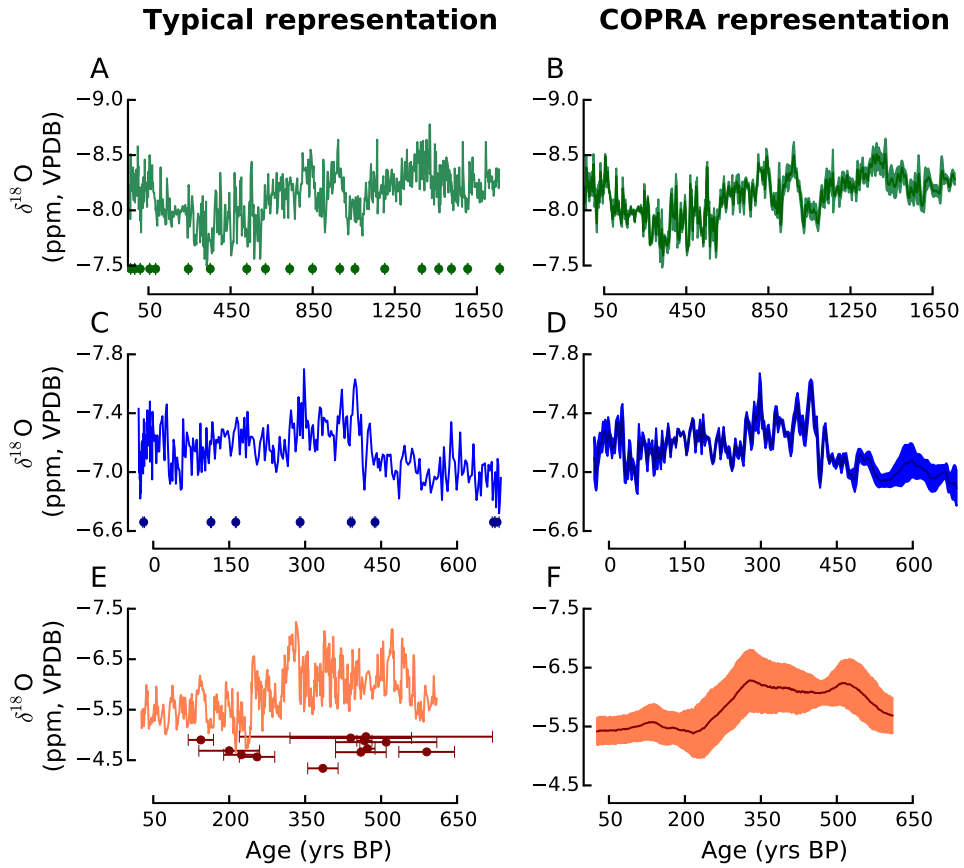

Figure 2.5. Comparison of proxy representations for east Asian δ^{18}O records. Typical proxy record representations (left column, **A**, **C**, **E**) as compared to the COPRA representation (right column, **B**, **D**, **F**) for the east Asian proxy records from Wanxiang (**A**, **B**), Dayu (**C**, **D**) and Wah Shikar (**E**, **F**) caves. It is clear that the proxy records on the left do not reflect the true impact of the age measurement uncertainty nor the final proxy estimate. On the other hand, the COPRA records on the right make it apparent which parts of the record are better constrained and which parts are not. The comparison is especially stark for the case of the Wah Shikar record.

have no way of communicating the way in which the age uncertainties impact the variability of the final proxy estimate. Nor do they communicate the error in the final proxy values at any given time point.

Moreover, a second crucial point to note is that, in Fig. 2.5, the time axis is error-

free so that the interpretations of visible past climatic events do not have to be made in the light of age uncertainties. Provided that the age measurements are well constrained and thus the resulting proxy estimate has sufficient high frequency variability, as is the case with the Wanxiang and Dayu caves, we hope that such an error-free timescale would allow for easier and more transparent inter-record comparisons.

A logical question that follows such an analysis is that: Is it possible to have a quantitative estimate of the relationship between the frequency of the final proxy estimate in terms of the age uncertainties made on the archive? A rough rule of thumb is highlighted in Section 4.2.3. However, a precise quantitave estimate of this relationship is till yet unknown and is non-trivial. This remains a focus of future study.

Chapter 3

A Bayesian approach to proxy uncertainty

Even though the COPRA approach provided one of first comprehensive methodologies to quantify proxy uncertainties that propagate age measurement and proxy measurement uncertainties into the final proxy record, it is not without its limitations. COPRA is based on heuristic numerical methods that rely on Monte Carlo procedures and thus the resulting ensemble of proxy records that it generates is always an approximation of the complete proxy record. Furthermore, COPRA in its current form is applicable solely to U/Th-dated archives as it cannot account for radiocarbon calibration and the uncertainties that are inherent in transforming a measured radiocarbon age to a calendar age (Blackwell and Buck, 2008). This is because radiocarbon age measurements need to be further "calibrated" in order to get the *calendar age* from the *radiocarbon age* (the two are not the same because of changing radiocarbon levels in the atmosphere). Blackwell and Buck (2008) contains a detailed overview of how this is achieved. The current version of COPRA does not have a routine that can estimate the calendar age distributions from the measured radiocabon ages and thus, it is not applicable to proxy datasets which have been dated with ^{14}C-dating.

A general, non-numerical framework for estimating a proxy record and its associated uncertainties that is based on thorough estimation-theoretic principles and that can account for radiocarbon calibration as well is still lacking. In this chapter, we present such an approach that provides analytical expressions for the final proxy record starting with the age–depth and proxy–depth measurements and that considers radiocarbon calibration as well if the archive is dated using radiocarbon dating methods. We have reported this approach in Goswami et al. (2014).

Our approach is founded on Bayesian estimation principles that first derives a probabilistic relation between the depth layers of the archive and its age which tells us which depth layers are most likely given a particular calendar age and not, as is typical in most age modeling studies till date, the other way round. In a spirit similar to COPRA, we use conditional probabilities to finally represent the proxy record on an error-free timescale which hopefully enables a much easier inter-comparison of proxy records from different archives.

The method discussed in the following sections of this chapter is generic in the sense that it can be applied to any archive provided that the age measurement errors are properly described by a suitable probability distribution. It is not in the least limited to U/Th-dated and ^{14}C-dated archives. Also, in our treatment of the growth of the sediment archive, we do not use an explicit growth model but rather consider the age measurements to be the sole source of growth information of the archive. Of course, when a more reasonable growth model is available, it can be incorporated in our approach in the form of prior belief about the growth of the sediment, as is already done in most age modeling approaches (Blaauw et al., 2007; Bronk Ramsey, 2008; Parnell et al., 2011). Thus, our method can also be potentially assimilated within existing age modeling frameworks such as COPRA (Breitenbach et al., 2012), StalAge (Scholz and Hoffmann, 2011), clam (Blaauw, 2010), or OxCal (Bronk Ramsey, 2008). Moreover, it could also be modified and extended to be applicable to other dating methods such dendrochronlogy, luminescence dating and tephrochronology.

Before we move ahead to a full description of the method, we wish to stress that the primary result of our analysis is the proxy record as represented by a series of probability distributions on an error-free timescale (a point that will be treated in more detail later in Section 4.2.4, and Chapter 5). Still, while visualizing the proxy distributions in a plot over time (and/or their mean/median), an observer might 'see' certain trends and variabilities in the record, and may be then tempted to interpret these patterns as trends and variabilities of the actual paleoclimatic conditions. We advise a word of caution in interpreting such trends from the plots because in principle, the posterior proxy distribution obtained at two different time points tell us almost nothing about the trend of the proxy. To actually comment on a particular relationship between the proxy at two different times, we need to,

in principle, rigorously derive the posterior distribution for precisely that relationship and for this, we would actually need to have the posterior joint distribution of the proxy record at two time points—and we do not focus on this aspect in our analysis present here. Note that a similar issue arises in the estimation of the radiocarbon calibration curve and is discussed in detail in Blackwell and Buck (2008).

Any interpretation regarding the proxy record based on our current analysis that involves the combined information of the final proxy distributions at two (or more) different time points can at best be a qualitative one, and will be strictly restricted by the uncertainty level of the final record. The greater the uncertainty, the less we can infer.

3.1. Theoretical background

3.1.1. Preliminary considerations

We differentiate between two radiometric dating method types: (i) the ones that do not require radiometric calibration, e.g., U/Th dating, and (ii) the ones that need calibration, e.g., ^{14}C dating. In the following, we restrict ourselves to the case of ^{14}C dating, since U/Th dating can be considered formally as a special case of the latter, i.e., with a trivial diagonal calibration curve without error.

The information given us is: (i) a set of radiometric ages for a few depth layers of an archive, (ii) a much larger number of proxy measurements at numerous depths of the same archive, and (iii) for the case of radiocarbon dated archives, a ^{14}C calibration curve, such as IntCal13 (Reimer et al., 2013). All three datasets have errors of measurement, although the errors in proxy measurements are negligible in several cases.

On the basis of the above information, we wish to answer the key question: *What were the most likely climatic conditions at a given time point in the past?* We then rephrase this question to make it more relevant to the given datasets as: "At a chosen time point, which proxy values were more likely than others?" We arrive at the answer using a probabilistic framework, i.e., we wish to now derive probability dis-

tributions for the proxy given a past time point. In particular, we adopt a Bayesian framework, as it quite naturally lends itself to transforming one kind of (given) information (such as the age–depth and proxy–depth measurements) into partial knowledge about other types of (desired) information (such as the age–depth and proxy–depth relationships). This is further exemplified by the historical prevalence of Bayesian methods in several studies that deal with radiocarbon dating, radiocarbon calibration, and age modeling (Niklaus et al., 1994; Blackwell and Buck, 2008; Bronk Ramsey, 2008; Niu et al., 2013).

Two fundamental concepts of probability are essential to our derivation.

- *Conditional probability.* Consider two random variables A and B that randomly takes values from a given set of possible values. If we know that there exists a relation between A and B, we can ask the question: knowing that B has taken a particular value b from its possible range of values, what is the probability that A takes a particular value a from its own range of values? This is idea of conditional probability which allows us to express our beliefs about how likely is it that a random variable A actually has the value *a conditioned on* our knowledge of that a related random variable B is equal to b. This is mathematically denoted as $P(A = a|b = b)$ or simply as $P(a|b)$ in shorthand, which reads "probability that $A = a$, *given* that $B = b$".

  [Note that we use the convention of denoting the random *variables* with uppercase letters A and B, whereas the *actual values* (which are real numbers) that they can take up are denoted with lowercase letters a and b.]

- *Bayes' theorem.* This allows us to combine two kinds of probabilistic information in order to obtain a third. The first kind is the *prior*—which encodes *a priori* knowledge or beliefs about a random variable A, denoted as $P(a)$. The second type is the conditional probability of an observation $B = b$ given a hypothesis $A = a$, $P(b|a)$, also known as the *likelihood* of the hypothesis $A = a$ given the observation $B = b$. Paraphrasing this in words, Bayes' theorem essentially states that *the posterior of a hypothesis is proportional to the product of its prior and its likelihood:*

OUTLINE OF THE DERIVATION FOR POSTERIOR PROXY PROBABILITIES

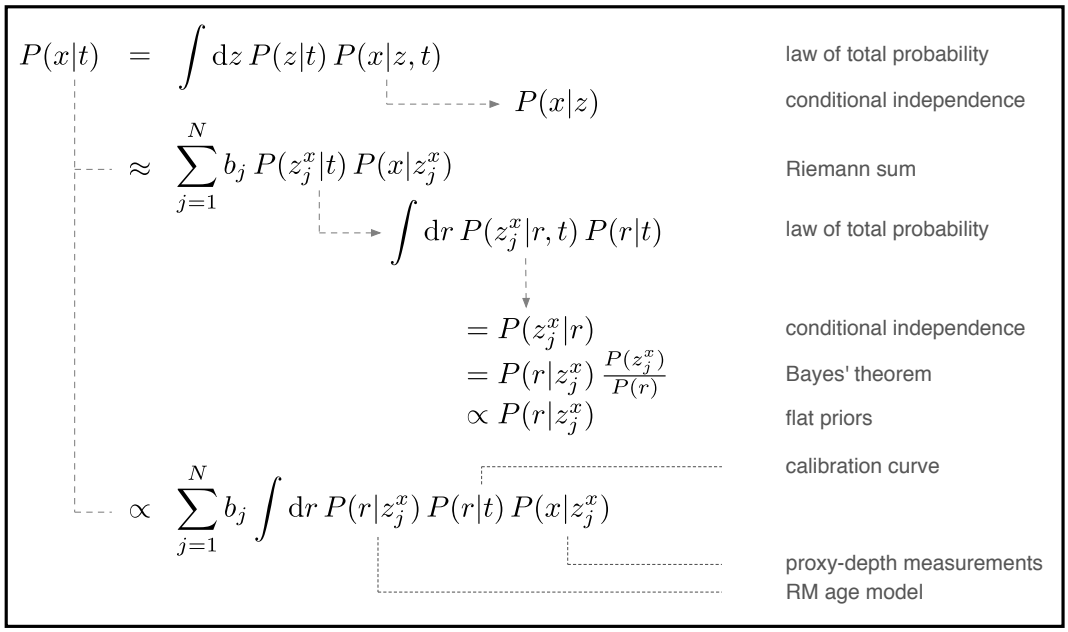

Figure 3.1. Outline of posterior proxy probability derivation. The box presents a rough sketch of the derivation of the posterior proxy probability densities. We denote the proxy, depth, time and radiometric age with the random variables X, Z, T, and R, respectively. The essential quantity of interest to be estimated is the posterior conditional density $P(x|t)$, which is finally expressed in terms of the measured/given quantities using the above steps. For a detailed explanation, please refer to Section 3.2 and Appendix A.

$$P(a|b) \propto P(a)P(b|a) \tag{3.1}$$

This is interpreted in the Bayesian framework as: given our prior assumptions/beliefs/knowledge about A, we use the observation $B = b$ to obtain the posterior $P(a|b)$ and thus *update* our prior beliefs about the value of A.

In the language of the terminology outlined in this section, the goal of this analysis is to arrive at the posterior distributions of the given proxy at all points in time. A

gist of the derivation of the posterior proxy probability distributions is presented in Fig. 3.1. A detailed explanation follows in the subsequent sections and also in the Appendix. Readers may also refer to the introductory sections of Bronk Ramsey (2008) and Parnell et al. (2011) for further explanatory notes on the use of Bayesian statistics in age modeling approaches.

3.1.2. Necessary assumptions

At the outset, we wish to state the assumptions incumbent in our approach and hold them up for review. We assume that:

- Each of the given ^{14}C ages, as well as each of the proxy measurements, are sufficiently well-described by a normal distribution with the mean at the measured value and a standard deviation equal to measurement error. This is motivated by the fact that, in most cases, observations (and associated errors) are adequately described by the mean (and standard deviation) of a Gaussian process (cf. Blaauw (2010); Hercman and Pawlak (2012); Scholz and Hoffmann (2011); Bronk Ramsey (2008)). Still, in principle, our method can be used for any general probability distribution.

- The errors in depth measurements are negligible. The precision of depth measurements motivate this second assumption and it is made for the sake of analytical simplicity. We note that most of the existing age modeling techniques involve this assumption, and also that our approach would have to be modified in cases where depth uncertainties cannot be neglected.

- All radiocarbon ages and depths are assumed to be equally likely *a priori*. This prior belief is then updated within the analysis using the set of radiocarbon age–depth measurements. Note that such a *prior* assumption can easily be modified in other cases and integrated into our approach. We note that, although a different prior can be assumed to incorporate more specific *beliefs* about the sediment growth, it does not change the main inferences that we wish to extract at the end of our analysis.

- The radiocarbon dates provided do not contain outlying values caused, e.g., by contaminated material. Dealing with outliers is beyond the scope of this analysis. If it is suspected that the dataset contains outliers, we suggest to identify and remove them with specific outlier-detection techniques and then use the remaining data as an input for our analysis.

- All obtainable growth related information about the archive is contained in the given set of radiocarbon dates. We note that although there are cases where additional sources of growth related information are available, we do not consider such situations in this preliminary exposition.

- The radiocarbon age of a given position in the archive is an unknown, smooth, continuous function of its depth. We denote this function as the *Radiometric (RM) age model*. The presence of a hiatus is a crucial issue regarding this. We advise that if a hiatus is known to have occurred at a particular depth (from sources other than the radiometric dates), the RM age model can be split at the hiatus depth into two smaller, independent RM age models—an approach used in Breitenbach et al. (2012).

3.2. Bayesian proxy estimation

3.2.1. Depth-spanning weight functions

In his discussion on deposition models for chronological records of paleoarchives, Bronk Ramsey (2008) has aptly articulated the fundamental idea behind the construction of age models as:

> What we are aiming to do mathematically is [...] to find a representative set of possible ages for each depth point in a sedimentary sequence.

Mathematically speaking, the fundamental idea behind our approach is similar as we too seek to establish a probabilistic relation between the depth and age as a first step. However, our approach differs from the idea above in that we obtain this relation *in the opposite direction, i.e., we aim to find an ensemble of depths for each calendar age*.

Let Z, R, and T be three variables that denote depth, radiocarbon age, and calendar age, respectively. Also, let the unknown proxy variable be denoted by X. We begin the task of finding such an ensemble of depths for each calendar age by asking: what is the probability that a depth $Z = z$ corresponds to a given calendar age $T = t$? By the law of total probability, this is proportional to the probability of finding a radiocarbon age $R = r$ at $T = t$ multiplied by the probability that the depth z corresponds to r, integrated over all possible values r of R. Formally,

$$P(z|t) \propto \int \mathrm{d}r \, P(r|t) \, P(z|r). \tag{3.2}$$

The expression on the right-hand side of equation 3.2 is similar to the calibration of radiocarbon dates. The first (calibration) term $P(r|t)$ involves the knowledge that not all radiocarbon ages are equally likely for a given calendar age and it thus filters out the most probable radiocarbon ages for that specific t. The second term $P(z|r)$ then filters out the most likely depths that could correspond to a given radiocarbon age r based on the estimated RM age model.

Now, using Bayes' Theorem, we see that the second term on the right-hand side of equation 3.2 is a form of a posterior distribution and is thus proportional to the product of some prior probability $P(z)$ and the likelihood $P(r|z)$ (for a more detailed discussion, see A). Assuming *a priori* that all depths are equally likely, we use the "flat" prior $P(z) \propto$ const. Thus, if the proxy is measured at depths $Z = z_j^x$ for $j = 1, 2, \ldots, N$, using the example of equation 3.2, we can define a *depth-spanning weight function (DWF)* on all proxy measurement depths z_j^x as:

$$w_t(z_j^x) := \int \mathrm{d}r \, P(r|t) \, P(r|z_j^x). \tag{3.3}$$

The term $P(r|t)$ is simply the radiocarbon calibration information which gives the probabilities of possible radiocarbon ages r given a particular calendar age t. The second term in equation 3.3, $P(r|z_j^x)$, is the RM age model which gives the probability of the radiocarbon age r given a particular depth z_j^x.

In combining these two terms under the integral, the DWF $w_t(z_j^x)$ constructed thus is proportional to the probability that a depth z_j^x corresponds to a given calendar age t in the archive. For each different value t of the calendar age T, a new DWF has to be constructed over the depths z_j^x. A schematic of two representative DWFs for a simulated archive is shown in Fig. 3.2 along with the information used in their construction. A couple of points to note from Fig. 3.2 are: (i) the shape of the DWF can be quite different for different calendar ages, and (ii) the peak of the DWF denotes the most probable value of Z for a given RM age model and a given value of T. This approximately corresponds to the value of Z at which the expected value of R equals the expected value of R for the given value of T. We illustrate this point with the help of the dashed arrows in Fig. 3.2. Conventionally, the calibration of a radiocarbon date involves estimating the probabilities of all possible calendar ages for a given radiocarbon age along with its uncertainty of estimation. This is analogous to proceeding in a counterclockwise direction in Fig. 3.2. In our analysis, we avoid 'calibration' in this sense by choosing to estimate a relationship between calendar ages and corresponding depths in the other direction—as shown by the clockwise sense of the arrows in Fig. 3.2.

Next we elaborate on the construction of the RM age model and its relevance in our method.

The RM age model

The construction of the DWFs involves the term $P(r|z_j^x)$ which has to be well-defined at every depth where the proxy is measured. However, in practice, the radiocarbon age measurements are limited to significantly fewer number of depth points. Let us say that the number of radiocarbon age measurements is M, where $M \ll N$, and the corresponding depths are denoted by z_k^r, $k = 1, 2, \ldots, M$. Then, from this set of measurements, we get M conditional probability distributions $P(r|z_k^r)$ for $k = 1, 2, \ldots, M$. Since M is much less than N, we need to be able to use the set of radiocarbon age–depth observations to construct a dataset that gives us the N probability distributions $P(r|z_j^x)$, $j = 1, 2, \ldots, N$. This is achieved by estimating a radiocarbon (or more generally, a radiometric) age–depth relation, i.e., the *radiometric (RM) age model*. In principle, the role of the RM age model can

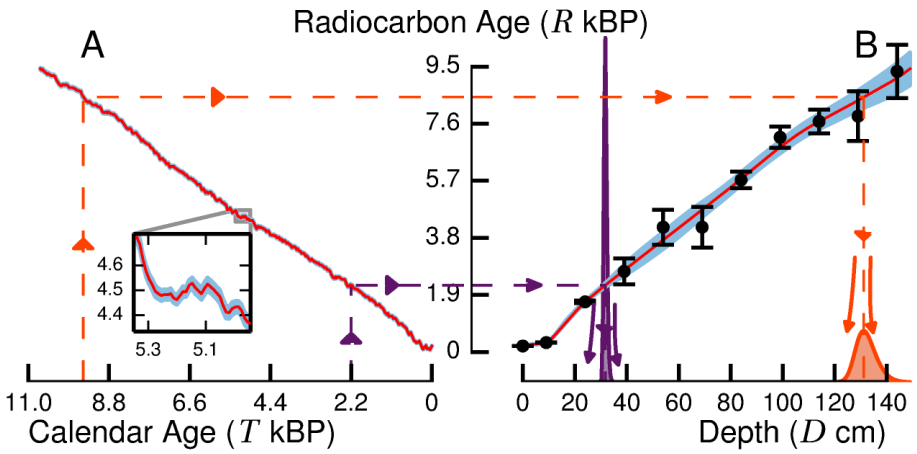

Figure 3.2. Depth-spanning weight functions (DWFs). A. Radiocarbon calibration curve representing the mean (red) and ± 2 standard deviations (blue) of the distribution $P(r|t)$. *Inset*: Calibration curve around 5.1 kBP detailing non-monotonicities. **B.** (i) ^{14}C measurements from a simulated archive (circles with error bars representing the distribution $P(r|z_k^r)$) and the RM age model obtained from them: mean (red) and ± 2 standard deviations (blue) of the distribution $P(r|z_j^x)$ that results from the used regression method. (ii) DWFs for the calendar ages 2.2 kBP (purple) and 9.5 kBP (orange). The corresponding colored dashed lines with arrows indicate how a given calendar age is related to a distribution over probable depths via the DWF. All uncertainty bounds correspond to $\approx$ 95% confidence.

be seen as a transformation of sediment growth information from the coarse age measurement depth-scale to the much finer proxy measurement depth-scale:

$$P(r|z_k^r), k = 1, \ldots, M \xrightarrow{\text{RM age model}} P(r|z_j^x), j = 1, \ldots, N, \tag{3.4}$$

which resonates with the conventional framework of age modeling. However, at this point, we use radiocarbon ages instead of calendar ages.

In our current analysis, we use a non-parametric Taylor-polynomial based regression method given by Heitzig (2013)—a data-driven approach that uses a combination of Bayesian updating and Taylor expansion about a point of interest to

provide an estimate of the smooth curve from which the observations have been sampled. However, any regression method that estimates a posterior probability distribution could be used equally, as well as any Monte Carlo based method that generates ensembles of interpolated RM age–depth relations from the set of observations to arrive at a mean/median estimate along with a standard deviation. Also note that in the case of dating methods that do not require radiometric calibration (e.g., U/Th dating), the first term of the right-hand side in equation 3.3 is equal to 1 if $r = t$ and 0 otherwise. In such cases, the RM age model involved is, in fact, the *calendar* age model—or what is simply known as the *age model* of the archive.

3.2.2. Incorporating monotonic growth

The DWFs constructed in Section 3.2.1 relate any given calendar age probabilistically to different depths based on how likely they are to correspond to that age. This relation does not, however, include one specific feature of sedimentary records: *the constraint of stratigraphically ordered growth*. In other words, a stratigraphically deeper layer of the sediment core cannot be younger than any layer above it. We thus need to incorporate this constraint into the *initial* DWFs derived in Eq. (3.3), henceforth denoted as w^i, and obtain a *final* set of DWFs (henceforth w^f) that takes this factor into account.

The DWFs are essentially a set of probability density functions and to impose an unambiguous monotonicity constraint on such a set is non-trivial. To overcome this we consider the cumulative probabilities of each DWF. A *Cumulative Depth-spanning Weight Function* (CDWF) is the probability that a given calendar age t corresponds to any depth less than or equal to a depth z_j^x. Formally, after ordering depths such that $z_1^x \le z_2^x \le \cdots \le z_N^x$, we put

$$W_t^i(z_j^x) = \sum_{l=1}^{j} w_t^i(z_l^x). \tag{3.5}$$

The CDWF $W_t^i(z_j^x)$ by definition increases along the depth axis z_j^x from 0 to 1.

[Note that we denote the cumulative probability distribution with an uppercase letter "W" to distinguish it from the corresponding probability distribution "w" denoted by the lowercase letter.]

Our task now is to ensure that the CDWF $W_t^i(z_j^x)$ is monotonically decreasing along the age axis T. This means that if we take a depth z_j^x and two ages t_1 and t_2 such that t_2 is greater (i.e., older) than t_1, the total probability that t_2 corresponds to a depth $\leq z_j^x$ cannot be more than the total probability that t_1 corresponds to a depth $\leq z_j^x$—which is the condition of monotonic growth. Formally: we would like that $W_{t_2}(z_j^x) \leq W_{t_1}(z_j^x)$.

The above condition for monotonic growth along the age axis is violated slightly but noticeably on many occasions. This is shown in Fig. 3.3A, where the wiggles in the white grid lines parallel to the calendar age axis illustrate the non-monotonicity. Since it would be quite difficult to enforce the desired monotonicity already in the step where $P(r|z_j^x)$ is estimated from $P(r|z_k^r)$, we instead fix the slight non-monotonicities after having derived W^i. This results in final CDWFs W^f that can then be transformed into the final DWFs w^f via

$$w_t^f(z_j^x) = W_t^f(z_j^x) - W_t^f(z_{j-1}^x). \tag{3.6}$$

To estimate the final W^f which adhere to monotonicity, we use the principles of relaxation dynamics, and the details of this estimation process are discussed in appendix A.3. In short, we start with a suitably chosen set of CDWFs which are already monotonic in T and then iteratively drag (pull-and-push) this function in order to minimize its distance from the initial CDWF, W^i, as far as it is possible to do so without violating its monotonic nature. The final equilibrium set of CDWFs is that which cannot be moved any closer to W^i by *any* form of dragging without compromising its monotonicity. We denote this set as W^f (shown in Fig. 3.3B).

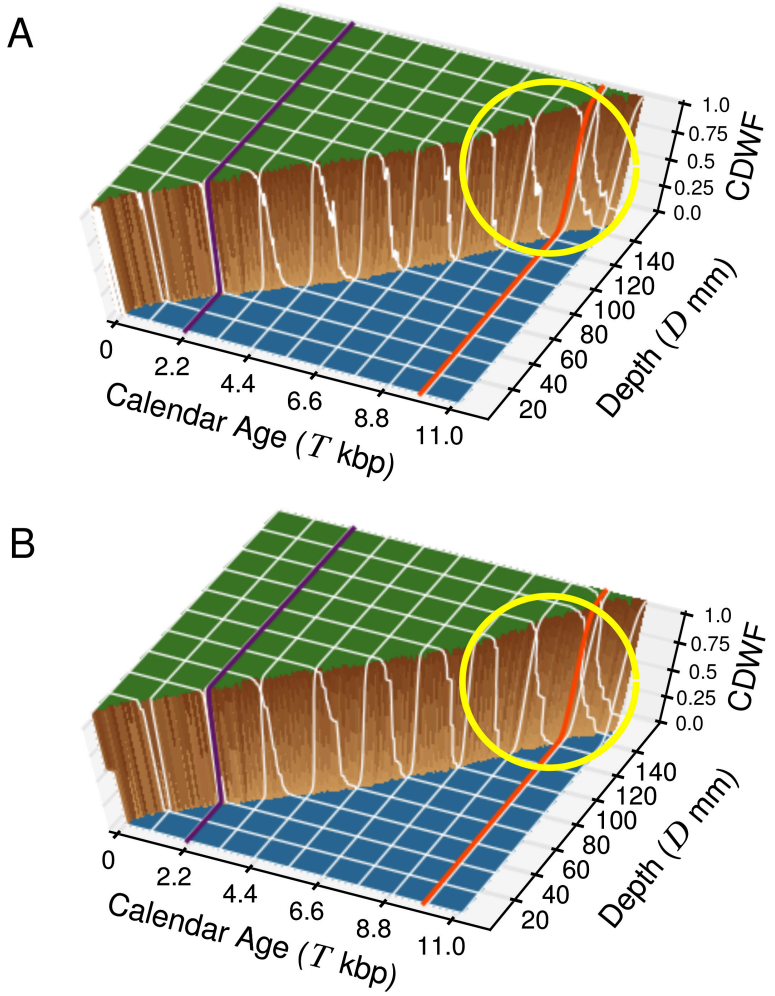

Figure 3.3. The age–depth sea cliff. A. The initial CDWF W^i obtained using equation 3.5, and **B.** the final CDWF W^f obtained after imposing monotonic growth, for the same example as used in Fig. 3.2. $W = 0$ is shown in blue and $W = 1$ are shown in green. Intermediate values of W are shown with increasing darkness from light brown to dark brown, as the value increases from 0 to 1. The overlay of white lines over the CDWF surface highlight the differences between the pre-monotonicity and post-monotonicity scenarios. The purple and orange lines are the CDWFs for the ages 2.2 kBP and 9.5 kBP as used in Fig. 3.2 respectively. The effects of imposing the condition of monotonic growth on the CDWFs is seen more prominently in the white grid lines within the marked yellow circles.

The age–depth sea cliff

The visualization of the CDWFs in Fig. 3.3 is likened to a sea cliff where $W = 0$ is shown as the blue sea, and $W = 1$ as the green highland. All intermediate values of W are contained in the sudden rise of the brown cliffs.

As stated earlier, the fundamental idea behind age modeling is to arrive at a relationship between the calendar ages and the depths of a paleo-archive. A set of functions that perform this function can be thought of, in a broader sense, as an 'age model'. The CDWFs visualized in Fig. 3.3 are thus analogous to an age model in our analysis. However, we wish to emphasize that the construction of the CDWF relations did not involve assumptions about the growth and/or sediment accumulation of the archive and was entirely a data-driven estimation with the sole input of the principle of monotonic growth of the core.

In this sense, the age–depth sea cliff is a formal age model which, in future studies, could be developed further to incorporate specific growth conditions, leading to a better estimation of the (calendar) age–depth relation of the archive.

3.2.3. Estimating the proxy record and its associated uncertainty

Once we have estimated the set of monotonic w^f functions, it is relatively straightforward to estimate the proxy and its uncertainty. For this, we need to consider the probability encoded in $w_t^f(z_j^x)$, for each depth z_j^x, as a weight for the corresponding proxy measured at that depth. Since X denotes the unknown proxy at a given calendar age t, we thus estimate the probability $P(x|t)$ as the Riemann sum fraction,

$$P(x|t) \approx \frac{\sum_{j=1}^{N} b_j w_t^f(z_j^x)\, P(x|z_j^x)}{\sum_{j=1}^{N} b_j w_t^f(z_j^x)},, \tag{3.7}$$

where b_j is the width of the depth interval represented by z_j^x:

$$b_j = \frac{1}{2}\begin{cases} z_2^x - z_1^x & j = 1 \\ z_{j+1}^x - z_{j-1}^x & 1 < j < N \\ z_N^x - z_{N-1}^x & j = N\,. \end{cases} \tag{3.8}$$

[A detailed derivation of the estimation is outlined in appendix A.4.]

We now have a probability distribution $P(x|t)$ for the proxy values $X = x$ at any given time t and using this, we can estimate the mean/median, as well as uncertainty bounds constructed using percentiles or variance. In this study we restrict ourselves to median proxy values and represent the associated uncertainty of estimation with (i) a 95% confidence band constructed from the region lying between the 2.5^{th} and 97.5^{th} percentiles, and (ii) a 50% confidence band constructed from the region lying between the 25^{th} and 75^{th} percentiles of the distribution $P(x|t)$. Furthermore, in the following chapter, the mean/median curves are represented as dotted lines rather than continuous curves to emphasize that they are derived from the probability distribution at individual time points and do not have any relation to values at other time points.

3.3. Review of the steps involved

Before we present the applications and results in the next chapter, it is helpful to briefly summarize the salient steps involved in estimating the proxy records and their associated uncertainty using our approach.

1. We construct the RM age model using an appropriate regression method that provides the posterior distributions of radiometric ages at the proxy measurement depths (Equation 3.4).

2. We estimate the weight function DWF that relates any given calendar age to the proxy measurement depths in terms of the likelihood that they correspond to the chosen calendar age (Equation 3.3).

3. Next, we construct cumulative weight functions CDWFs from the initial set of DWFs obtained in the previous step. The CDWFs are used to impose the constraint of stratigraphically ordered growth of the archive (Equation 3.5).

4. We obtain a final set of CDWFs that are consistent with such monotonic growth over time, and we derive a final set of DWFs from them (Equation 3.6).

5. For each chosen calendar age, we use the corresponding stratigraphically ordered DWF to weight the proxy measurements over depth and thus obtain a posterior proxy probability distribution (Equation 3.7).

6. We use the posterior proxy probability distribution at each calendar age to estimate quantities of interest such as the mean or median proxy values for that age. Furthermore, we also estimate uncertainty measures such as quantile ranges or variance.

Chapter 4

Aspects of proxy record uncertainty

In this chapter, we present the results of applying the Bayesian proxy estimation approach described in Chapter 3. In order to illustrate the applicability of the method, we first consider two synthetic examples where we simulate numerically the growth of a stalagmite and a lake sediment which are then dated with U/Th and ^{14}C dating methods respectively. Both of these synthetic paleoarchives have a sinusoidal proxy deposited in every depth layer, and using a set of given age–depth and proxy–depth measurements from these archives, we estimate the posterior proxy record distributions using our method. Besides telling us how the various sources of error impact the uncertainty of the final records, these examples also help clearly demonstrate the validity of our analysis.

In a next step, we consider groundwater inflow and surface erosion proxies from Lonar lake in central India that were reported earlier in Anoop et al. (2013) and Prasad et al. (2014). We then compare our results with the typical age modeling approach based on OxCal (Bronk Ramsey, 2008) that was used in these studies. The results help us understand the interplay between proxy variations and the final proxy record uncertainty. Furthermore, based on a thought experiment where we imagine the age–depth relation of the Lonar lake sediment to be known with absolute precision, we are able to show that in the specific case of the Lonar lake proxies, it is the proxy variations rather than the age measurements which impact the final proxy uncertainty to a greater degree.

Section 4.1 presents the results of applying our method to the various scenarios while Section 4.2 presents a comprehensive discussion of the various aspects incumbent in estimating proxy records and the nuances therein.

4.1. Applications

4.1.1. Synthetic examples

To illustrate our method as well as to test its efficacy, we consider two types of paleoarchives: (i) a stalagmite extending over 0–28 kBP dated with U/Th, and (ii) a lake sediment core extending over 0–11 kBP dated with ^{14}C. From our perspective, the crucial difference between the two is that for the lake sediment, the radiocarbon ages have to be calibrated using IntCal13, whereas this is not needed for the U/Th ages. To simulate sediment growth, we follow Blaauw (2010) such that the sediment accumulates with an initial growth rate of 20 yr/cm. At subsequent depths, a non-negative growth rate is chosen from a normal distribution that has the growth rate of the previous year as its mean and a fixed standard deviation of 7 yr/cm. In both cases, the proxy values are simulated as a sinusoidal signal consisting of two components with different time periods. Also, the proxy datasets were generated annually, i.e., with a proxy value for every year. We simulate a few noisy radiometric age measurements and a much higher number of almost perfect proxy measurements (error ~ 0.001). These 'observations' are then used to estimate the proxy record with our method. [Note: Detailed equations of the growth of the synthetic archives are provided in appendix D.]

U/Th dated archives

The results for the synthetic stalagmite are shown in Fig. 4.1. In this case, the calibration curve (as shown earlier in Fig. 3.2) is replaced by a straight line of slope one without any error (Fig. 4.1A). This is possible because the U/Th radiometric ages can be identified with the calendar ages. The observational noise for the U/Th age measurements increases with the depth of the stalagmite to a maximum of 5% (Fig. 4.1B). The proxy signal has two components with time periods 2000 years and 400 years (Fig. 4.1D). Note, from panels C and D of Fig. 4.1, that the proxy signal can be distorted in the depth domain—depending on the nature of the actual age–depth relation.

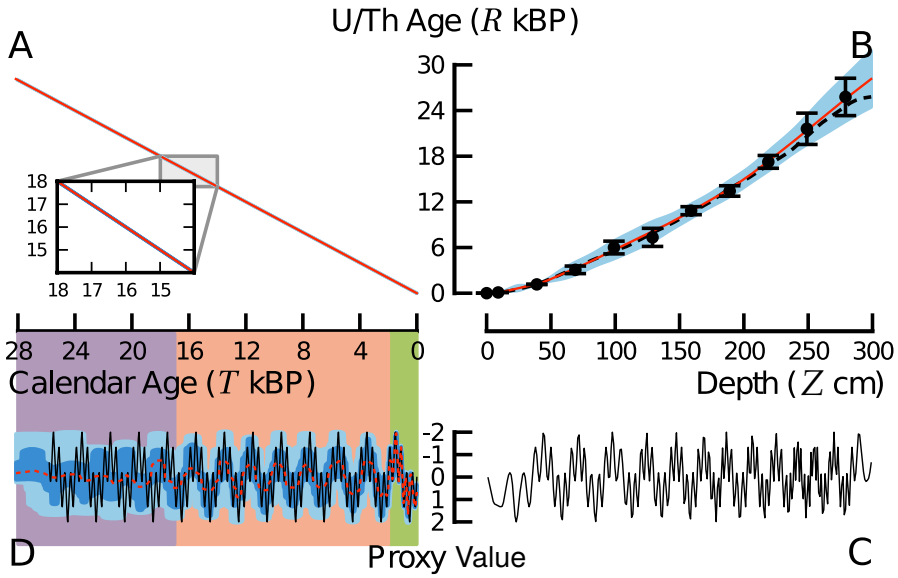

Figure 4.1. U/Th dated synthetic stalagmite. Legend as in Fig. 3.2. **A.** The straight line of slope one used in place of a calibration curve. *Inset*: Unlike a real calibration curve, this line has no error. **B.** U/Th measurements from a synthetic stalagmite (circles with error bars) and the estimated RM age model obtained using regression (red line with blue ≈ 95% confidence band), along with true age model (black dashes). **C.** The proxy curve along stalagmite depth obtained from noise-free measurements. **D.** The actual proxy record (black), shown alongside the estimated median proxy record (red, dotted) along with associated uncertainty of estimation (sky blue denotes the interquartile range, i.e., 50% confidence, whereas light blue denotes the region between the 97.5$^{\text{th}}$ and the 2.5$^{\text{th}}$ percentiles, i.e., 95% confidence). The estimated record in D is demarcated into three distinct regions in terms of the frequencies it resolves: (a) both frequencies of the true sinusoidal proxy (green), (b) only the lower frequency (orange), and (c) unable to resolve either of the two frequencies (purple).

^{14}C dated archives

Figure 4.2 shows the results of our method as applied to the synthetic lake sediment core. In panel A, we see the irregularities of the radiocarbon calibration curve and its estimation uncertainty. The error in radiocarbon age measurements

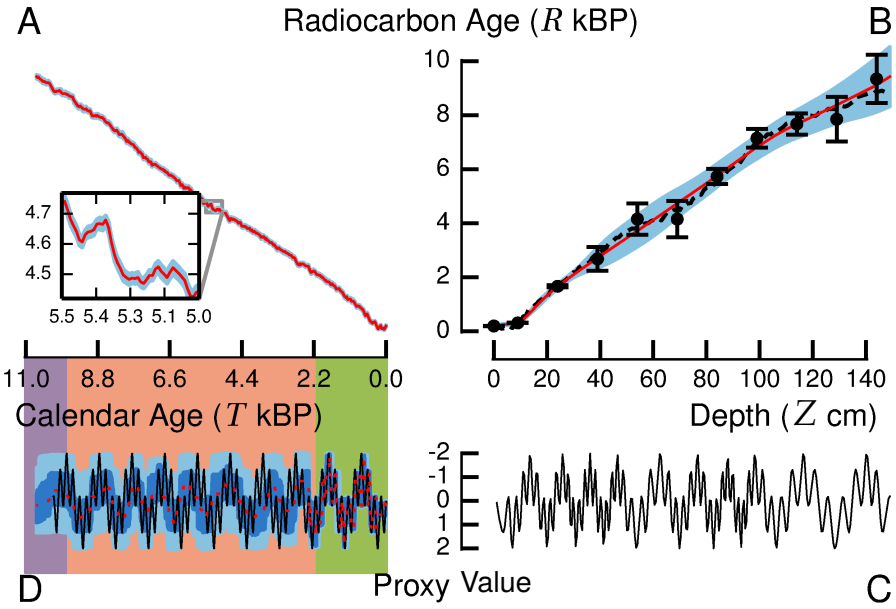

Figure 4.2. ^{14}C dated synthetic lake sediment. Legend as in Fig. 4.1.

in Fig. 4.2B increases with depth as in the previous case. The proxy signal used in this case has two time periods of 1000 years and 200 years. Here too, one can see that the proxy signal is distorted in panel C, when compared to the one in panel D; however, the distortion in this case is mediated not only by the irregular RM age–depth relation but also by the calibration curve.

4.1.2. Holocene proxies from central India

As an application to a real-world scenario, we consider the set of age–depth ^{14}C measurements from the Lonar lake in central India (Anoop et al., 2013; Prasad et al., 2014). The radiocarbon ages involved are tabulated in appendix C.2 and labeled in Fig. 4.3B. This included two ^{14}C ages after 1950, L21 and L20a, for which we use the Northern Hemisphere 3 (NH3) "post-bomb" calibration curve (Hua et al., 2013), and 17 pre-1950 ages, L19–L1, for which we use IntCal13 (Reimer et al., 2013).

For the proxy records, we take the Ca-area proxy for groundwater inflow and the Al-

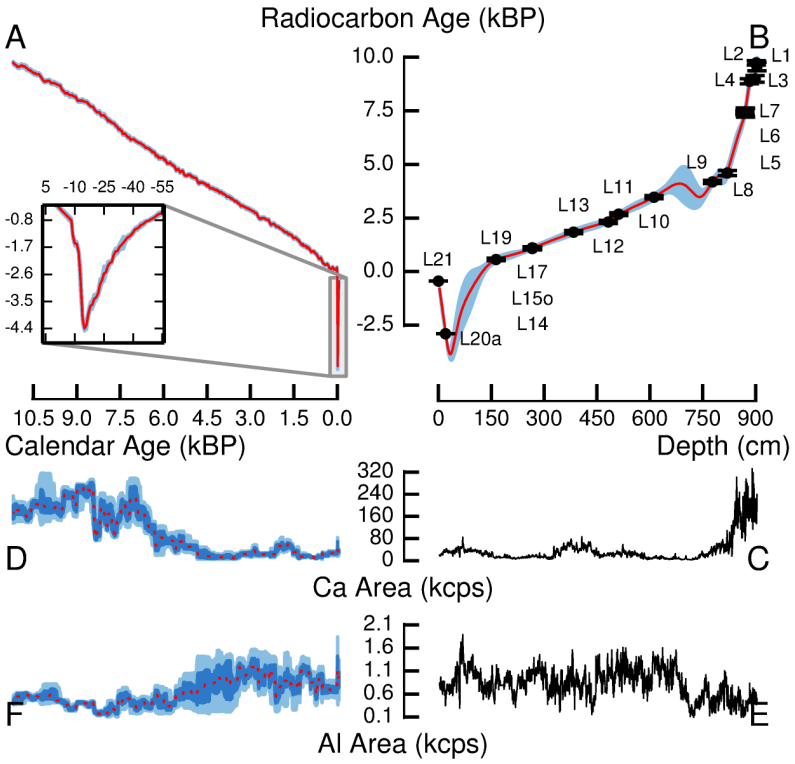

Figure 4.3. Holocene proxy records from central India. **A.** The IntCal09 calibration curve. *Inset*: The post-bomb calibration curve that was used. **B.** The RM age model for the Lonar lake sediment core. For details of the radiocarbon age–depth measurements, c.f., appendix C.2. **C, E.** The groundwater inflow (Ca area) and surface erosion (Al area) proxies measured along the depth of the core. **D, F.** Corresponding proxy record estimates as obtained using the Bayesian approach detailed in the text. Legend in these panels is the same as that of Fig. 4.1 D; 'kcps' denotes 'kilo counts per seconds'.

area proxy for surface erosion from the same archive at Lonar. The links of both the Ca to groundwater inflow (evaporitic carbonate ($CaCO_3$) formed during periods of low lake levels) and that of the Al to surface erosion (lithogenics brought on by rain events) have been validated in Basavaiah et al. (2014).

The proxies were obtained from a continuous down-core X-ray fluorescence (XRF) (Avaatech XRF Core Scanner III) scanning of the Lonar lake sediment core surface. The relative abundances of the elements (Ca, Al, Ti, Si, and K) were recorded ev-

ery 5 mm with the X-PIPS SXP5C-200-1500 detector from Canberra, while the tube voltage was kept at 10 kV (Prasad et al., 2014). The Al counts were found to be strongly correlated with the Ti, Si and K counts obtained from the XRF scanning (see appendix C.2). Due to this, and combined with the findings of Basavaiah et al. (2014), where they show the relation of the Al abundance to catchment erosion as well as the lithogenic contents, we choose this as a representative proxy for the Lonar lake surface erosion. We note that due to the difficulties in representing errors in XRF measurements, we consider the proxy observations along depth to be error free. This, however, does not change the fundamental objective of our analysis, which is to estimate the final proxy uncertainties in an analytical fashion and to investigate how they are impacted by proxy–depth variability. If the proxy measurements were to have errors, these would simply be added to the final errors, as is indicated by equation A.14.

The final proxy records estimated are shown in Fig. 4.3. We compare our results with proxy records obtained from a typical mean age model of the archive. The age model involved OxCal P-sequence modeling with three sedimentological boundaries imposed *a priori*. Figure 4.4A and B compares the final proxy estimates obtained using the OxCal P-sequence age model with those that were obtained using our approach—for both Ca area and Al area. These figures indicate that the results of our approach and the concomitant uncertainties contain within them the estimates obtained by using median age estiamtes from OxCal. This is, in a sense, a corroboration of our approach in the sense that it manages to provide resonable proxy and uncertainty estimates while at the same time being true to the inherent age unceratinties involved in the analysis.

4.2. Features of proxy uncertainty estimation

4.2.1. Proof of concept

The synthetic examples shown in Sec. 4.1.1 illustrate the validity of our approach. In panel D of Figs. 4.1 and 4.2, the first finding to note is that the 95% confidence band consistently contains well over 95% of the black curve, and the 50% band consistently contains about half of the black curve. In addition to this general fit

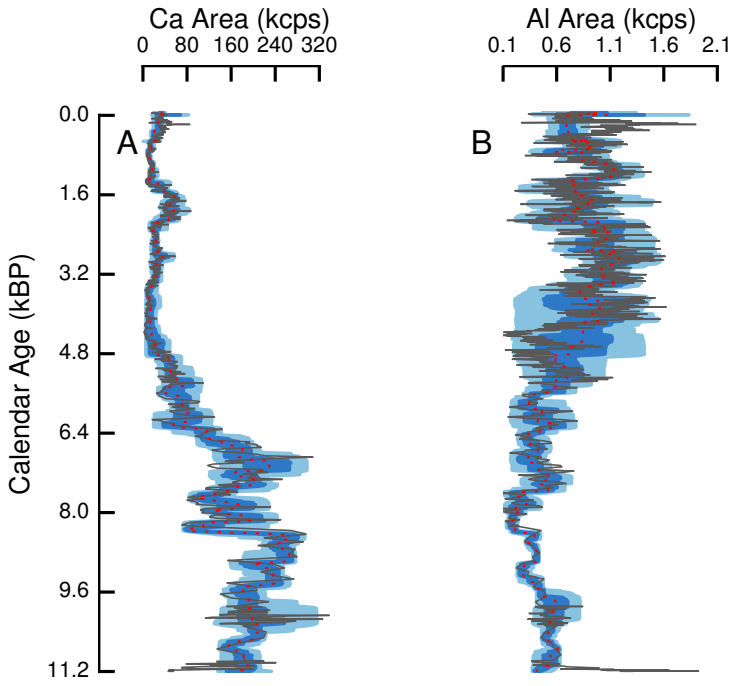

Figure 4.4. Comparison of results with OxCal. The proxy records estimated by the present approach (red, dotted) compared to proxy records obtained by using an OxCal P-sequence age model (dark gray) for the groundwater inflow (Ca area, **A**) and surface erosion (Al area, **B**) proxies from Lonar lake. Legend for the confidence bounds to the Bayesian proxy estimate is same as in Fig. 4.3, panels D and F.

between the true record and the estimated confidence bands, one can compare the median estimate (red dotted curve) with the true record (black curve), and distinguish three broad regions: (a) the youngest portion of the proxy records (green region), where the median estimates follow the true proxy series closely, and reproduces even faster oscillations of $1/400$ yr^{-1} (for the stalagmite) and $1/200$ yr^{-1} (for the lake sediment core) accurately; (b) the intermediate portion of the proxy records (orange region) where, for the most part, the median estimates show only the slower sinusoidal component due to larger dating uncertainties but follow the lower frequencies of the true proxy curves ($1/2000$ yr^{-1} for the stalagmite, and $1/1000$ yr^{-1} for the lake sediment core) closely, and (c) the oldest portion of the records (purple region), where the median estimates are almost flat curves due

to the high uncertainties. The differences among these three regions are due to the associated uncertainty of estimation of the proxy record which increases (as seen from the confidence bands) progressively from the youngest to the oldest portions of the record, becoming as large as the range of values in the end. The proxy uncertainty depends strongly on the errors of the corresponding RM age models (Figs. 4.1 and 4.2, panel B), which increases towards the oldest portion of the cores as well, and the errors of the RM age models are themselves influenced by the errors of the radiometric age measurements. *We would like to emphasize that the objective of our method is not to estimate frequencies or variability but to represent the available knowledge about the proxy value itself at each given point in time in the best possible way.* The seeming inability of the proxy estimates to reproduce oscillations is a necessary consequence of the posed research question when dating uncertainties are large. Still, in regions where the errors of measurement are small, the estimates reproduce the oscillations at both frequencies.

This discussion highlights two crucial factors: (i) proxy estimation errors depend on contingent errors of age measurements (and, to a large extent, the errors of calibration and proxy measurements); and (ii) the interrelation between the estimated proxy uncertainty and the variations that are resolved in the record. Both these issues are discussed in the subsequent sections.

4.2.2. Factors impacting the uncertainty

Figs. 4.1–4.3, it is apparent that even though the proxies are measured to high precision along the depth of the respective archives, *the uncertainty in the proxy value for any given time point is not negligible.* This is in agreement with the results of earlier studies, e.g., Blaauw et al. (2007) and Breitenbach et al. (2012) and the illustrative example in Sec. 2.4. Also, the final uncertainty is not the same as the error of the proxy-depth measurements. The confidence bands span the whole range of values of the proxy for error levels of ≈ 5–10% of the radiometric age measurements.

At a first glance at Figs. 4.1–4.3, it is obvious that the final proxy uncertainty is influenced by the calibration uncertainty, RM age model uncertainty, and the proxy

measurement error (if any). However, a closer inspection of Fig. 4.3 D and F around 3–4.5 kBP reveals an additional factor. At around 3–4.5 kBP, we find that Al-area has much higher uncertainty in comparison to Ca-area even though both have the same calibration curve and age model. Moreover, even the proxy measurement error for both were considered negligible in the analysis. *Then why is the uncertainty much higher for Al-area?*

To answer this, we proceed clockwise in Fig. 4.3 from the interval around 3–4.5 kBP to find that this range of calendar ages would roughly correspond to the depth range of around 500–800 cm using the given calibration curve and RM age model. The critical difference between the two proxies in this depth range is that the Al-area has a much larger variance in comparison to the Ca-area. Thus, given the same DWF, the Al-area proxy estimates would have a much larger spread relative to the Ca-area.

To understand this in more detail, consider the schematic in Fig. 4.5. The proxy curve shown in the figure has distinctively high fluctuations in the purple portion of the curve and is then confined within a narrow band of values in the orange portion. On the left side of the figure, we consider two points t_1 and t_2 on the calendar age axis, such that the DWF $w^f_{t_1}(d^x_j)$ of t_1 covers mainly the high variability region A in Fig. 4.5 and similarly the DWF of t_2 covers the low variability region B in Fig. 4.5. Thus, we get a smaller uncertainty for the proxy at t_2 than at t_1 when we weight the proxy values with the height of the respective DWFs. Hence, the variability of the proxy measured along the depth also contributes to the final uncertainty of its estimation.

The final uncertainty thus depends on four contingent factors:

- calibration uncertainty (if any),
- uncertainty of the RM age model,
- errors of measuring the proxy along depth, and
- variability of the proxy signal along the depth domain.

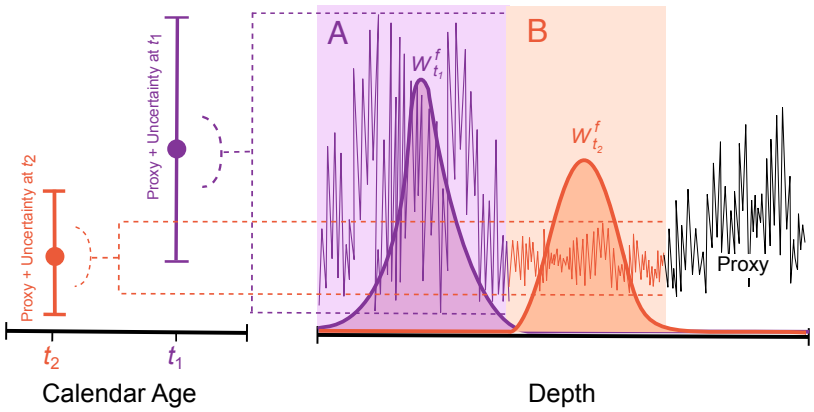

Figure 4.5. Uncertainty of proxy estimations (schematic). The proxy curve has high variations in the region **A** (purple) and low variations in region **B** (orange). Two time points t_1 and t_2 are chosen such that their corresponding DWFs cover most parts of regions **A** and **B** respectively. The estimation uncertainty is far greater for the proxy estimate at t_1 than for t_2 because of the corresponding high variability of the proxy in region **A**. *Note*: The final uncertainty of the proxy estimate is *not equal* to the corresponding variability in the depth domain but is only proportional to it. Hence, we denote the impact of proxy variability over depth (marked with the dashed lines) on the final proxy uncertainty with an open-arc.

4.2.3. Variability of the proxy record

From Section 4.2.2, it is evident that the uncertainty of proxy estimation is influenced by its own variability in the depth domain. However, the proxy uncertainty is also closely related to the variability of the median proxy estimate as well. This is seen in Figs. 4.1–4.3, where regions of high estimation uncertainty are associated with low variability *of the median estimates along time*, and less resolution of higher frequencies, and vice versa. This does not mean that the proxy itself *does not have* faster variations. It simply implies that, in regions with high proxy uncertainty, we cannot *reliably comment* on the fast variations of the proxy. In order to analytically quantify the various fast/slow varying components of the proxy and their uncertainties in a thorough fashion we should in principle, as stated before, proceed with a separate analysis. This is because the knowledge of the marginal

proxy probability distributions are not sufficient to comment on the variations, especially in the presence of non-negligible uncertainty of estimation.

The best one could say about the variability of the proxy-vs.-time in regions of high dating uncertainty would be to estimate some aggregate measures of variability such as the slope or curvature of the proxy curve or the momentary amplitude of a certain sinusoidal component (at each calendar age). For example, a simple way to obtain a "central" estimate of the slope $\mathrm{d}x/\mathrm{d}t$ would be to use MoTaBaR to find the mean estimates $R = r(t)$, $D = d(r)$, and $X = x(d)$ that correspond to given values $T = t$, $R = r$, or $D = d$, and then use the chain rule to calculate $\mathrm{d}x/\mathrm{d}t \approx x'(d(r(t)))\, d'(r(t))\, r'(t)$. Equivalently, in graphical terms, follow the dashed lines from t via r and d to x and multiply the corresponding slopes of the calibration curve $r(t)$, the RM age model $d(r)$, and the proxy-vs.-depth curve $x(d)$ that you encounter on the way (the calibration curve might have to be smoothened for this). In a similar fashion, the second derivative can be estimated by applying the product rule: $\mathrm{d}^2x/\mathrm{d}t^2 \approx x''(d(r(t)))\, d'(r(t))^2\, r'(t)^2 + x'(d(r(t)))\, d''(r(t))\, r'(t)^2 + x'(d(r(t)))\, d'(r(t))\, r''(t)$.

Finally, if the proxy-vs.-depth curve shows a sinusoidal component of amplitude ξ and period length Δd around depth $D = d(r(t))$ (as could be seen, e.g., from a wavelet analysis), one can conclude that the true climate-vs.-time curve contains a sinusoidal component around time $T = t$ of the same amplitude ξ and a period length that can be estimated as $\Delta t \approx \Delta d / d'(r(t))\, r'(t)$.

4.2.4. Interpreting the posterior probabilities

A critical point arising out of the discussions in the previous subsections is that the final proxy estimate—such as the mean/median—when visualized over a contiguous period of time, may not reveal short-time variations. For paleoclimatic studies focussed on transitions taking place over short timescales this can be a major hurdle. Even in studies that wish to address climatic patterns operating in the higher frequency region, a proxy estimate which does not resolve such frequencies is of little practical utility.

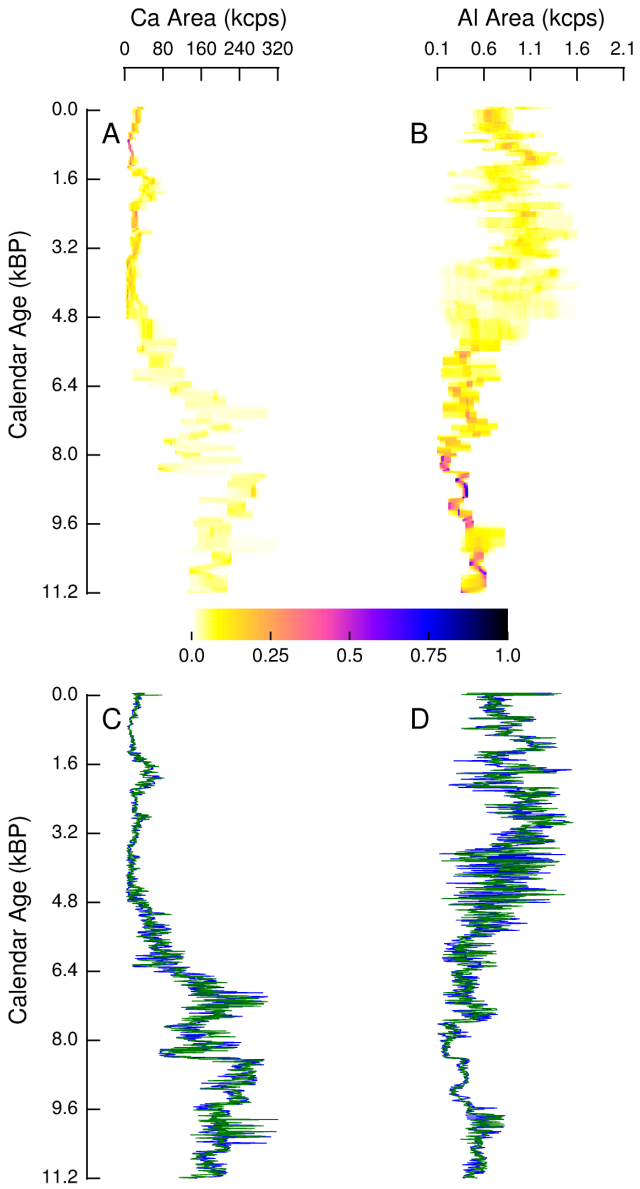

Figure 4.6. Interpreting the posterior proxy distributions. A, B. The posterior distributions for Ca-area (in A) and Al-area (in B). For each calendar age, we provide a probability distribution (indicated by the color-bar) for the proxy along the horizontal proxy axis. **C, D.** Two randomly chosen ensemble members (blue and green curves) for Ca-area (in C) and Al-area (in D) out of all possible proxy records given the probability distributions in A and B respectively. They retain the high frequencies, indicating that the high frequency information is still contained in the posterior probabilities.

We stress that it is misleading to conclude that the proxy record does not contain high frequency components based on figures such as Figs. 4.1–4.3. As already stated in the previous subsection, the fast varying components of the proxy are not ruled out by the probability distributions. Rather, only in estimating the mean or median, we might be unable to say anything about them with confidence.

To understand how this is possible, note that the primary and foremost result of our approach is a *marginal probability distribution* of the proxy values at each value of calendar age (shown in Fig. 4.6 A, B as a colormap). Such a visualization is in principle similar to Figure 2 of Blaauw et al. (2007)—only that we obtain the visualization from mathematical expressions, and not as a histogram of ensemble members.

We now interpret the distributions as representing the probability densities of an ensemble of possible proxy records. This means that each member of this ensemble is a record that is a proxy for one of many possible past climatic histories that fit the available set of measurements and data. This is shown in Fig. 4.6 C and D, in which two such members of the ensemble are shown for each of the groundwater inflow and surface erosion proxies from Lonar lake. These records were constructed by drawing random numbers from each proxy probability distribution at every calendar age under consideration. It is immediately clear that the individual members of the ensemble retain the high frequency components as well. However, since we have no way of knowing which of the infinite possible ensemble members actually constituted the actual climatic history, we have to estimate the mean/median climatic history and estimate our confidence on this mean value. The uncertainty bounds shown in this study represent the impossibility (given a set of measurements) of determining precisely the mean proxy value and hence, by extension, the mean paleoclimatic condition that it would represent.

In order to be able to have a narrow uncertainty range, efforts must be taken to reduce the various sources of error that contribute to the final proxy error. We discuss the possibilities and limitations of this in the next section.

4.2.5. Reduction of uncertainty

Let us take the example of the Lonar lake observations and ask: *how can we reduce the final proxy uncertainty?* For this, we have to look at the four factors that determine it. Among these, the calibration uncertainty cannot be reduced until a more tightly constrained calibration curve is released, the proxy-depth variance is beyond our control, and the proxy measurement error is already set to zero. Thus, we are left with the sole option of reducing the RM age model error (cf. Section 3.2.1). This can be achieved in several ways, such as additional radiometric dating of the archive, or by incorporating layer counted segments of the record that can provide age–depth information with relatively less magnitudes of error. However, as we do not consider layer counted data in our approach, we will consider below only the effect of adding more radiocarbon dating points.

We might plan to make a few more measurements especially around those depths where the RM age model is not very precise, e.g., at around 700 cm (Fig. 4.3B). Still, the fact remains that a significant portion of the final uncertainty might also be due to the intrinsic variance of the proxy along depth and we thus need to fully understand exactly how much of the final uncertainty is contributed by the age measurement errors. The highly non-trivial way in which the final uncertainty is related to the RM age model uncertainty (via the DWFs) makes it almost impossible to find a precise analytical answer to questions of the type: if we make two ^{14}C age measurements at depths d_1 and d_2 with a maximum error of ε, what is the fraction z by which the uncertainty at calendar age t is going to be reduced?

We can nevertheless get some insight into how much error is contributed by the age uncertainty by considering a simple thought experiment. Let us assume that we are able to reduce the RM age model uncertainty to zero by taking N error-free radiocarbon age measurements at the precise depths of proxy measurements. In this scenario, the variance of the DWF will solely depend on the calibration uncertainty, and this, in conjunction with the proxy's intrinsic variance, will determine the final proxy uncertainty. We can then compare the uncertainty levels of the proxy before and after setting the age model error to zero. This is shown in Fig. 4.7 for the Ca-area (panels A and B) and Al-area proxies (panels C and D). We

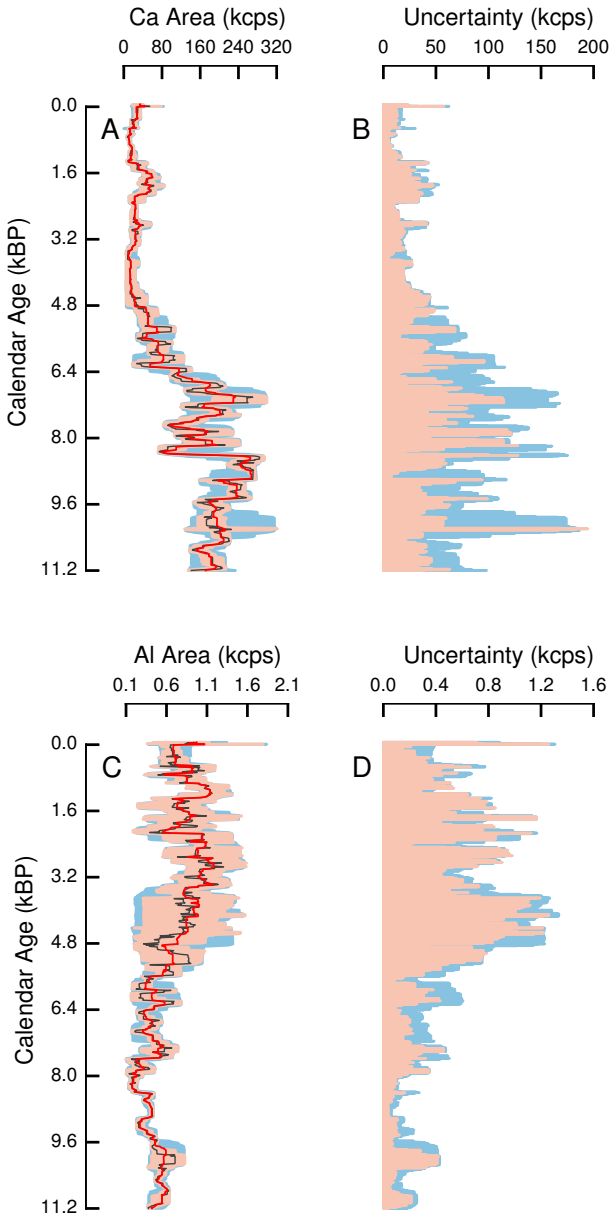

Figure 4.7. Contribution of age uncertainty to proxy estimation uncertainty. A, C. Proxy records for Ca-area (in A) and Al-area (in C) for the original set of observations (red curves for the median, light blue area for 95% confidence bands) along with the proxy records after setting the RM age model uncertainty to zero (dark gray curves for the median, orange area for 95% confidence bands). **B, D.** Uncertainties for each time point (95% confidence bands shown in A and C) and for the two cases: original observations (light blue); and after removing RM age model error (orange).

can see from the figure that the final uncertainty of the proxy is not reduced by a great amount (panels B and D)—even when the uncertainty of the RM age model had been set to zero. Among the two proxies, the reduction of uncertainty in the Ca-area record is more than that of the Al-area record. This is because the Al-area signal has relatively higher variability than the Ca-area signal (c.f., Fig. 4.3 C and E), and so the relative contribution of the age uncertainties to the final proxy uncertainty is less for the Al-area record than the Ca-area record. In both records, we find that a reduction of age uncertainty helps to resolve some higher frequency variations than before. However, this improvement is not a drastic one. Further, even if age uncertainty is reduced to zero, the proxy records still differ a great deal from the records constructed by using the OxCal P-sequence model as shown in Fig. 4.4.

Coming back to the issue of improving the Lonar proxy records with the help of additional measurements at around 700 cm, we first note that these depths would roughly correspond to calendar ages of around 4–5 kBP (see Fig. 4.3B: starting at around 700 cm, and going counter-clockwise from the depth axis to the RM age model curve to the calibration curve to the calendar age axis). From Fig. 4.3 we find that the region around 4–5 kBP show almost no improvement by setting the age uncertainty to zero! This indicates that the Lonar lake radiocarbon age measurements are not the primary source of the final proxy uncertainties. Rather, it is more likely that the major part of the proxy uncertainties in the Lonar records are due to the proxy fluctuations in the depth domain and the calibration uncertainties.

The thought experiment illustrates several points:

- Even though the final uncertainty of proxy estimation is linked to the age uncertainty in a complicated manner, it is possible to understand the relative contributions of the age errors by setting them to zero.

- The final proxy uncertainty can, in some cases, be determined more by its own variability in the depth domain, rather than the age uncertainty.

- Given a set of observations and a RM age model, it is possible to obtain a

limit of the precision with which the proxy can be estimated.

- The variability of a proxy signal is inherently linked to the kind of paleoclimate variations that it will allow to be investigated, and also to the level of precision with which such studies can be carried out.

A thought experiment such as this one can also be potentially used to devise an optimal strategy for dating a newly discovered paleoarchive such that redundant dates are not measured. This is the scope of further future work.

4.2.6. Precise, error-free timescale

One important consequence of having an age uncertain timescale for representing proxy records is that the inter-comparison of records from different archives become difficult and ambiguous. In the approaches outlined in this dissertation, we overcome this difficulty because of the use of conditional probabilities. The use of conditional probability implies that every time we consider a particular calendar age and then estimate the DWF for it, we 'know' the calendar age precisely and then obtain the likelihoods of the proxy depths for that age. This means that the final proxy estimate we obtain is represented on a timescale that is error-free, i.e., it is without any uncertainty. Such a notion of a *precise timescale* has already been introduced for speleothems in Breitenbach et al. (2012), where it is termed as an "absolute" timescale which corresponds effectively to the time of deposition of the proxy material on the the archive. They illustrate the utility of an error-free timescale with the help of Monte Carlo age modeling approaches. In the present paper, we generalize this idea for all radiometrically dated archives and also provide an analytical framework for it.

Typically, if we represent paleoclimate proxies on an age uncertain time scale, the interpretations of paleoclimate events are constrained by not being able to know when exactly an event took place in the past. This is overcome if the uncertainty is transferred from the age axis to the proxy axis ensuring that the timescale of representation is always error-free. An analogy to visualize this process is to think of the uncertainty as a bag of error that can be carried either by the time axis or the

proxy axis. In paleo-investigations till now, this bag of error had been left on the shoulders of the time axis, but we choose to transfer it to the proxy instead. This comes at a price—because in doing so, we are not certain any longer about the high frequency variability of the proxy. Thus, it is not a question of whether this particular representation of proxies is more correct than the conventional age-uncertain one—the choice of representation is context dependent and is determined by the goals of the paleoclimatic investigation. One can conceive a study in which the interest in the high frequency variability is outweighed by the problems of vagueness induced by an age-uncertain time axis. In such a scenario, it is reasonable to use an existing framework of age modeling that establishes a representative set of ages for each depth level of the core, provided the approximations used for dealing with the irregular calibrated age distributions are reasonable.

A methodological advantage of having an error-free time axis is that time series analysis methods are more readily applicable to them because most time series analysis techniques are conceived with precise time ordering of data sets in mind. Although there has been recent work that are able to extract the climate spectrum (Mudelsee et al., 2009) as well as timescale dependent trends with errors (Mudelsee et al., 2012), interpretations of time series analyses such as the construction of paleoclimate networks (Rehfeld et al., 2012) have to be cautious of the time uncertainties that were inherent in original data and were left unresolved. Our approach provides a clear platform to carry out such analyses and interpret the results.

Chapter 5

The limits of given information

Up until now, we have dealt with various issues regarding the estimation of proxy records and their associated uncertainties. In the process, we presented a sequence of time-ordered probability distributions from which the mean/median proxy estimates and related measures of dispersion such as the variance or interquartile ranges can be obtained. However, in order to learn more about the paleoclimatic conditions and variability underlying the proxy records, we need to carry out further analyses such as the estimation of autocovariance function, power spectra, correlations between different proxy records, and so on. In order to do this, we first need to understand the utility as well as the limitations of what these proxy distributions represent and how they can be used. To some extent, we tried to stress this point in Section 4.2.4. In this chapter, we reiterate this fact and try to highlight the cases in which the proxy distributions are potentially useful. We also discuss the additional information required to extend the horizon of utility of the posterior distributions.

5.1. Non-independence between proxy estimates at two time points

Figure 5.1 illustrates the structure of information that we have in our hands at the end of applying our methods presented earlier to proxy observations obtained from natural archives. Before we actually derive the mean/median estimates and associated parameters of uncertainty, we fundamentally have a sequence of posterior probability distributions. The first fact to note is that these are marginal distributions with respect to time in that they tell us only the probability of getting a

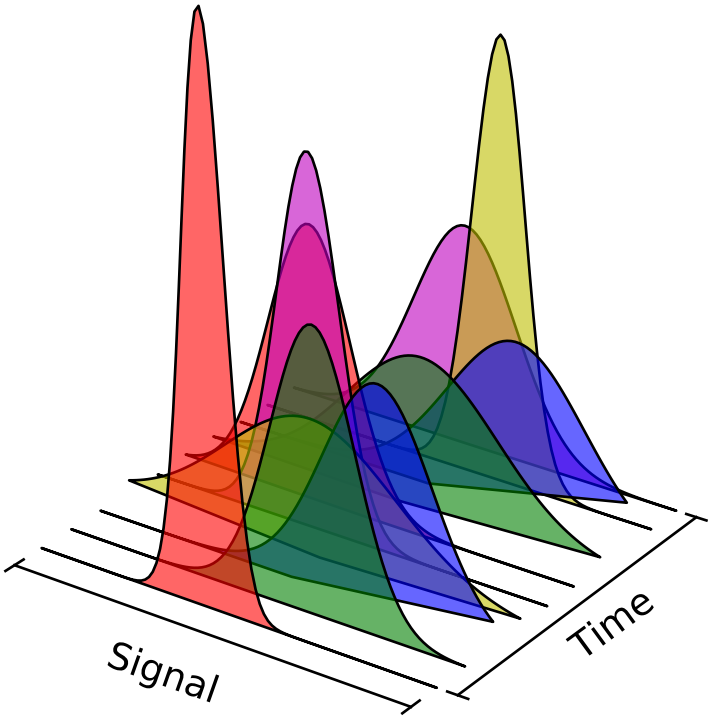

Figure 5.1. Time-ordered sequence of probability densities. This schematic illustrates the information at hand after we apply the methods described earlier. On the bottom plane, we have the time axis and the axis denoting the values of the signal (i.e., the proxy). The colored curves represent the posterior marginal probability densities derived in Chapter 3. The height of the curve at a given time denotes the probability of the value of the signal/proxy at that time.

particular proxy value at a chosen point of time. They do not provide any information about the joint probability of the proxy values at two different points of time. This is because we have no knowledge of the (non-zero) covariance between the marginal distributions at two time points. In special cases, it might turn out that the derived marginals are, in fact, independent of each other (i.e. the covariance between the proxy probabilities at two different time points is zero), however, it is unjustified to assume that this is true in general. Second, it is necessary to understand that we have no way of estimating the covariances between the given set of time-ordered marginal densities from the densities themselves. This is because:

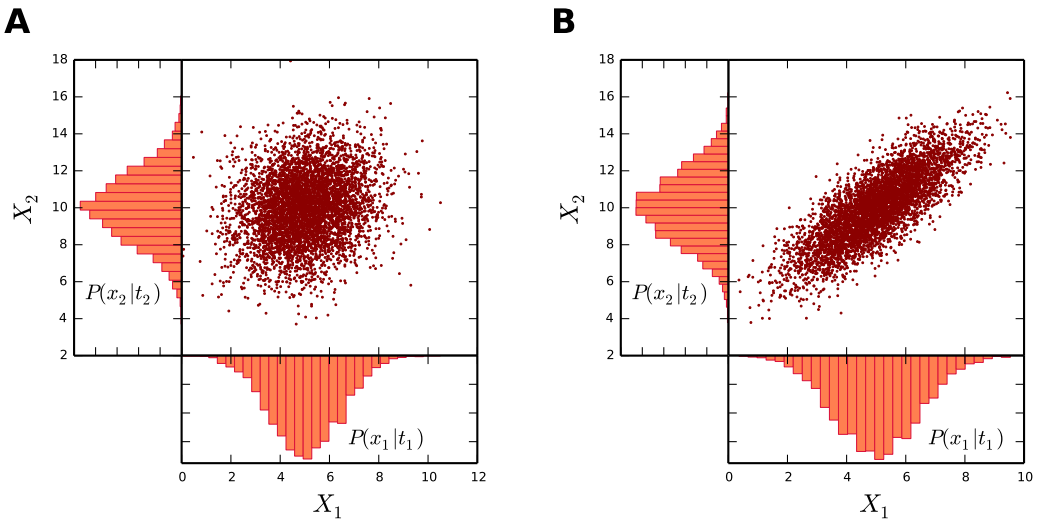

Figure 5.2. The importance of knowing the covariance. Two marginal distributions $P(x_1|t_1)$ and $P(x_2|t_2)$ are chosen for two given time points $T = t_1$ and $T = t_2$ such that they are described by the normal distributions $\mathcal{N}(5,2)$ and $\mathcal{N}(10,3.5)$ respectively. The marginal distributions are shown as orange bars whereas the bivariate normal joint distribution is shown as a scatter plot for correlation values 0.2 in **A** and 0.8 in **B**.

(a) we do not have observations, i.e. samples, recorded from these probability densities that might reveal to us the underlying covariance structure among them, and (b) it is not possible to uniquely determine the covariance structure (and by extension, the joint distribution) for a given set of marginals from the marginals themselves.

To further highlight why the knowledge of the covariance is crucial in carrying out further analyses, let us consider the following example (as shown in Figure 5.2) of just a pair of marginal densities corresponding to two different time points t_1 and t_2, where the indices 1 and 2 denote time instances. The corresponding proxy values at t_1 and t_2 are then represented by the random variables X_1 and X_2 respectively. According to our earlier analysis on proxy estimation we are given the posterior marginal densities $P(x_1|t_1)$ and $P(x_2|t_2)$. Let us assume, for the sake of illustration, that these two marginals are well-described by two normal distributions $\mathcal{N}(5,2)$ and $\mathcal{N}(10,3.5)$. Here $\mathcal{N}(\mu,\sigma^2)$ denotes a normal distribution with

mean μ and variance σ^2). Given the correlation (which is proportional to the covariance) between the two marginals, we can sample the joint distribution numerically using the bivariate normal form. Figure 5.2 shows the two marginals on the two axis and a sample of 5000 points on the (X_1, X_2) planes for two different values of the correlation—0.2 and 0.8. We find that the resulting joint distribution, as observed by the density of points on the (X_1, X_2) plane is very different in the two cases. From this, we can infer that (a) the joint distribution is not unique for the given pair of marginals, and (b) further analyses carried with an incorrect estimate of the covariance might lead to drastically different conclusions.

At this point, a question that arises is: Can we generate ensemble members from the given marginal distributions and numerically estimate all further time series characteristics from the ensemble? The answer to this is that, without the covariances, this too is not possible. This relates to the second point mentioned above—because the joint distribution is not uniquely determined by just the marginals, it is not possible to generate an ensemble based on the marginals alone as each ensemble member is effectively the joint probability of the proxy value at all the time points. If we were to generate an ensemble by choosing random numbers, at each time point, that conform to the marginal distribution at that time point, this would be equivalent to assuming that marginal densities are independent!

In the following section, we describe a few time series characteristics and the ways in which the joint distribution (of the proxy at pairs of time points) plays a role in determining these characteristics.

5.2. The role of the joint distribution

To illustrate the importance of the joint distribution of the proxy values at two time points, we consider several time series characteristics and derive the mathematical expressions for them based on the given sequence of posterior proxy densities. One of the first characteristics that might be of interest is the expectation value of the proxy with respect to time. Note that this expectation is different from the expectation of the proxy at a given point of time (as shown in equation A.13). The expectation value of a random variable depends on the probability density used to estimate it and since we are now interested in obtaining the expectation value over

time, we have to first derive the total probability of a proxy value over a period of time. This is given as

$$P(x) = \frac{\int \mathrm{d}t\, P(x|t)}{\int \mathrm{d}t}. \tag{5.1}$$

Using the above probability density, it is possible to show that the expectation of the proxy over a period of time is equal to the time average of the expectation values of the proxy at each time point (i.e. $\bar{x}(t)$ from equation A.13).

$$\mu_X = E[X] = \int \mathrm{d}x\, x\, P(x) \tag{5.2}$$

$$= < \bar{x}(t) > . \tag{5.3}$$

Similarly, we can derive the variance σ_X^2 of the proxy over a period of time and with respect to the expectation estimated above.

$$\sigma_X^2 = \int \mathrm{d}x\, (x - \mu_X)^2\, P(x) \tag{5.4}$$

$$= < \sigma^2(t) + \bar{x}(t)^2 > - \mu_X^2. \tag{5.5}$$

We thus find that the expectation value and the variance of the proxy over an extended period of time is uniquely determined from the expected values and variance of the proxy at individual time points.

Next we estimate the autocovariance function of the proxy over a period of time. For this let us assume that X_t and $X_{t+\tau}$ are random variables denoting the proxy value at a given time t and and at times $t+\tau$ respectively. In such a notation, the autocovariance function $\varrho^X(\tau)$ is simply

$$\varrho^X(\tau) = E[X_t X_{t+\tau}] - E[X_t] E[X_{t+\tau}] \tag{5.6}$$

The first term of equation 5.6 involves the joint distribution $P_{X_t X_{t+\tau}}(x_t, x_{t+\tau})$ of the proxy value at times t and $t+\tau$ which is unknown to us and thus, the autocovariance function of the proxy cannot be estimated exactly, in principle. If, in special cases, we can motivate that the proxy values at two distinct points are sufficiently independent we can hope to factorize the joint distribution as $P_{X_t}(x_t)P_{X_{t+\tau}}(x_{t+\tau})$ and thus estimate the above integral. A more detailed derivation of the autocovariance relation is given in appendix B.2.

However, the case of the crosscovariance $\varrho^{XY}(\tau)$ between two proxy records X and Y is slightly different,

$$\varrho^{XY}(\tau) = E[X_t Y_{t+\tau}] - E[X_t]E[Y_{t+\tau}] \tag{5.7}$$

where the joint probability $P_{X_t Y_{t+\tau}}(x_t, y_{t+\tau})$ can in fact be broken down to *conditionally independent* factors $P_{X_t}(x_t)P_{Y_{t+\tau}}(y_{t+\tau})$ such that the final expression for the crosscovariance reduces to

$$\varrho^{XY}(\tau) = \langle \bar{x}(t)\,\bar{y}(t+\tau)\rangle - \langle \bar{x}(t)\rangle\,\langle \bar{y}(t+\tau)\rangle. \tag{5.8}$$

Next, to derive the power spectrum of the proxy, we can use the Wiener-Khinchin Theorem to explicitly write the power spectral density $S^X(\omega)$ in terms of the autocovariance function as,

$$S^X(\omega) = \mathrm{mod}\Big[\int \mathrm{d}\tau\, \varrho^X(\tau)\exp(-i\omega\tau)\Big]^2. \tag{5.9}$$

Equation 5.9 shows that, since the autocovariance function cannot be estimated without the joint distribution or the assumption of independence between the proxy probabilities at two time points, the power $S^X(\omega)$ faces the same limitation as well.

Another relevant characteristic that we wish to estimate is the recurrence plot of the proxy dynamics. In principle, a recurrence plot constructed from an observed time series tries to first encode the recurrence information of the dynamics in a binary recurrence matrix $\mathbf{R}_{ij}$ with entries 1 or 0. $\mathbf{R}_{ij} = 1$ implies a recurrence between the phase space trajectory of the system at times i and j respectively, and $\mathbf{R}_{ij} = 0$ implies the converse. To understand the recurrence properties of the proxy given the set of posterior probabilities, we first need to estimate the recurrence probabilities based on them. The probability that the proxy value X_i at time i recurs within an ε-neighborhood of the proxy value X_j at time j is given by,

$$P_{ij}(\varepsilon) = \int \mathrm{d}x_i \int_{\mathcal{N}_\varepsilon(x_i)} \mathrm{d}x_j \, P_{X_i X_j}(X_i, X_j), \tag{5.10}$$

where $\mathcal{N}_\varepsilon(x_i)$ denotes the ε-neighborhood of x_i. Here too, we see that the joint distribution is the central quantity of interest required to estimate the recurrence probabilities and, in the absence of which, we are either unable to estimate it precisely or are forced to assume independence.

The dynamical characteristics considered in this section and the role of the joint distribution in their estimation are listed in Table 5.1. The table summarizes the fact that, barring the time expectation, variance, and crosscovariance, the other characteristics cannot be estimated without the knowledge of the joint probability density of the proxy at a pair of time points.

5.3. Sources of covariance: Sediment growth and calibration

At this point, it is essential that we understand as to why is it that we cannot justify the assumption of independence between the posterior proxy densities. This is illustrated by a derivation of the posterior joint probability density of the proxy for a pair of time points, in which we try to express the joint probability in terms of the measured information such as the age-depth measurements, the proxy measurements, and the calibration curve.

Table 5.1. The role of the joint distribution in system characteristics

Name	Symbol	Mathematical expression	Joint density
Time expectation	μ_X	$< \bar{x}(t) >$	No
Time variance	σ_X^2	$< \sigma^2(t) + \bar{x}(t)^2 > - \mu_X^2$	No
Autocovariance	$\varrho^X(\tau)$	$E[X_t X_{t+\tau}] - E[X_t]E[X_{t+\tau}]$	Yes
Crosscovariance	$\varrho^{XY}(\tau)$	$\langle \bar{x}(t)\bar{y}(t+\tau)\rangle - \langle \bar{x}(t)\rangle \langle \bar{y}(t+\tau)\rangle$	No
Power spectrum	$S^X(\omega)$	$\mathrm{mod}\left[\int \mathrm{d}\tau\, \varrho^X(\tau) \exp(-i\omega\tau)\right]^2$	Yes
Recurrence probability	$P_{ij}(\varepsilon)$	$\int \mathrm{d}x_i \int_{\mathcal{N}_\varepsilon(x_i)} \mathrm{d}x_j\, P_{X_i X_j}(X_i, X_j)$	Yes

For this, let X_i, $i = 1,2$, be the two random variables that describe the proxy values at two time instants $T = t_1$ and $T = t_2$ respectively. Also, let Z_i and R_i, $i = 1,2$, be the random variables related to the depth and radiocarbon axes for the X_1 and X_2 respectively. Proceeding analogously to the single proxy case, we can determine the joint posterior density of the proxy at the two time points as,

$$P(x_1, x_2|t_1, t_2) = \text{Weighted mean of the product } P(x_1|z_1^{x_1})P(x_2|z_2^{x_2}), \tag{5.11}$$

where the weights are given by a DWF $w^f_{t_1 t_2}(z_i^{x_1}, z_j^{x_2})$ defined on both the depth variables Z_1 and Z_2,

$$w^f_{t_1 t_2}(z_{j_1}^{x_1}, z_{j_2}^{x_2}) := \iint \mathrm{d}r_1 \mathrm{d}r_2\, P(r_1, r_2|z_{j_1}^{x_1}, z_{j_2}^{x_2})\, P(r_1, r_2|t_1, t_2). \tag{5.12}$$

[*Note*: A more detailed derivation of the above is provided in appendix B.]

We now have to consider the two terms on the right-hand side of Equation 5.12 as the DWF defined here is crucial in estimating the posterior joint densities. In order to understand how we could go about estimating these joint densities we have to

first attempt to understand the underlying random processes that give rise to such joint probabilities.

Sediment growth

The first term on the right-hand side of Equation 5.12, i.e., $P(r_1, r_2 | z_{j_1}^{x_1}, z_{j_2}^{x_2})$, denotes the joint probability of the radiocarbon ages given two distinct depth points along the core. This is inherently related to the question: if we know that the radiocarbon age at depth $z_{j_1}^{x_1}$ is r_1, what can we say about the radiocarbon age r_2 at a second depth $z_{j_2}^{x_2}$? Does this belief about the value of r_2 change if we had not known anything about r_1? The answer to the second question is a 'yes' in most situations, as the growth of a sediment core is not completely random and thus, knowing the age at a particular depth influences our beliefs about the age at any other given depth. This implies that the joint probability cannot be broken down to the product of the marginal probabilities $P(r_1 | z_{j_1}^{x_1})$ and $P(r_2 | z_{j_2}^{x_2})$ because we cannot assume independence in this context. Thus we have to either estimate the joint probabilities from the measured age–depth radiocarbon measurements themselves using regression (along with some a prior assumptions) or assume a parametric growth model, such as a Markov Chain, and then use Bayesian parametric estimation methods to sample the joint probabilities. Out of these two options, to the best of our knowledge, almost all existing regression techniques are inadequate to provide the joint probabilities, while on the other hand, parametric growth models with Bayesian estimation methods are already quite popular in the age-depth research community (such as, e.g., Heegaard et al. (2005)).

Radiometric calibration

The second term on the right-hand side of Equation 5.12, i.e., $P(r_1, r_2 | t_1, t_2)$, denotes the joint probability of the radiocarbon ages given two distinct time points, i.e., calendar ages. This information is essentially related to the radiocarbon calibration method. Based on the model outlined in Heaton et al. (2009) it is clear that the calibration curve is based on a Markov Chain model where the the beliefs about the radiocarbon age at one calendar age is critically dependent on our

knowledge of the radiocarbon age at any other calendar age. As in the previous case, this would mean that the joint density in question cannot be factorized into two independent components. However, the critical difference is that we are dependent on the IntCal researchers to provide us with the necessary covariance information which can help to estimate/derive the joint probabilities. Heaton et al. (2009) state in their introduction that "by generating plausible samples, given the calibration data [...] one is able to obtain posterior covariances between all points on the curve together with any other summary statistic desired". However, even until the recent release of the IntCal calibration curve, *viz.* IntCal13, the covariance information of the radiocarbon ages with respect to calendar ages has not been made available.

From the above discussion, we can conclude that the covariance of the proxy probabilities at two given calendar ages arise from two sources: (a) the sediment growth and (b) the radiocarbon calibration. Out of the two, it is possible, with the help of an appropriate growth model, to estimate the covariances that arise from the growth processes of the core. However, the covariance information related to the radiocarbon calibration has to be provided by the IntCal research group and it is, in principle, possible to estimate this as well. Therefore, for archives that are dated with U/Th dating and that require no calibration of the measured radiometric dates, the covariance of the proxy probabilities at two different times arise entirely from the sediment growth and it is thus possible, in theory, to carry out all the necessary data analysis methods mentioned in Section 5.2. For the case of radiocarbon dating, we propose that the radiocarbon, and especially the IntCal, communities acknowledge the importance and the necessity of the covariance information related to calibration in the analysis of proxy dynamics and uncertainty, and thereby make this information available in a form that can be used by paleoclimate researchers.

Notwithstanding these limitations, in the following chapter, we present some ideas on how one can proceed to analyze the recurrence characteristics of the proxy record without the knowledge of the joint distribution. We shall attempt to use the recurrence analysis framework in order to infer dynamical characteristics from paleoclimate proxy records as obtained by our approach.

Chapter 6

Recurrence networks of paleoclimate data

The primary concern of this chapter is to analyze dynamical characteristics of paleoclimate proxies keeping in mind the uncertainties that are inherent in it. We present here a new approach to reconstruct the dynamical recurrences of a dataset which is provided as a time-ordered sequence of correlated probability densities rather than a precise time series. This is particularly challenging as the uncertainties of the probability densities at different time points are correlated (cf. Chapter 5). Using bounds on the difference distributions of each pair of probability densities, we are able to approximate the qualitative features of the recurrence network adjacency matrix **A** of the dynamical system assumed to be underlying the given paleoclimatic dataset. In the following sections, we first provide a brief overview of the bounds of the recurrence probabilities that are critical to the rest of the analysis, and a consequent validation using the bivariate normal distribution. We then demonstrate the ability of the method to reproduce qualitative features of the original recurrences of the system (for low levels of uncertainty) using representative exemplary systems.

In the context of sedimentary paleoclimate proxy records, our analysis is particularly relevant and significant because now more than ever, increasing number of paleoclimate records have enabled researchers to gain newer insights about past climates. In particular, speleothem-based isotopic climate proxies from various parts of the world have provided well-dated and progressively precise climatic records at timescales ranging from a few decades to hundreds of thousands of years (McDermott, 2004). It is necessary to develop theoretically sound and practically pertinent methods that are able to reveal depper insights from paleo-datasets regarding paleoclimatic conditions and variability.

Note that we use the framework of recurrence networks in the analysis presented in this chapter. Recurrence networks allow the analysis of individual time series with complex network tools and techniques (Marwan et al., 2009). By representing the recurrences of a given system as a complex network embedded in the phase space, we can characterize dynamical features and transitions from a given dataset (Donner et al., 2010). Recent development of its theoretical underpinnings have furthered the understanding of the complex network measures derivable from a given time series (Donges et al., 2012). Recurrence network analysis of paleoclimate data has also recently shed light on the impact of climatic-variability transitions on human evolution over the last 5 million years (Donges et al., 2011).

We posit that communities in recurrence networks correspond to parts of the phase space trajectory where the dynamics are locally similar. In the context of real-world complex networks, communities typically denote parts of the network that are more strongly connected within themselves than to the rest of the network (Newman, 2004c, 2006). Although there are various ways of identifying communities in a given complex network, we use a modularity-based approach as proposed by Newman (2004b). In particular, we identify the probability of recurrence matrix as a weighted adjacency matrix and obtain the modularity with respect to a chosen partition (Newman, 2004a).

6.1. Recurrence networks

Recurrence patterns of dynamical systems provide a useful tool for the characterization of dynamical behavior by observing time series generated from them. In the past few decades, recurrence plots in particular have been used to analyze and study a wide variety of complex systems. Essentially, a *recurrence* is said to occur when the phase trajectory of the dynamical system gets close to a previously visited part of the phase space (shown in Fig. 6.1). By considering all sampled points of the phase trajectory, and with a predefined recurrence threshold ε we can define a recurrence matrix, each entry of which is either 1 whenever a given pair of time points i and j fall within each other's ε-neighborhoods, or 0 otherwise (cf. Marwan et al. (2007) and Chapter 7 for a more detailed treatment of recurrence plots).

Analogously, recent studies have put forth the idea of equating the recurrence ma-

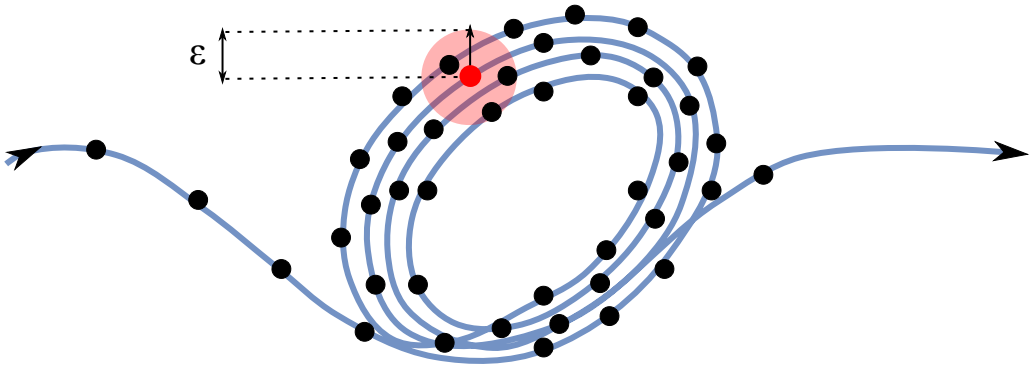

Figure 6.1. Recurrences of dynamical systems. An idealized schematic of a phase space trajectory of a dynamical system is shown here in blue. The black circles represent points sampled on the trajectory that give rise to the observed trajectory. the light red shaded circle around the point marked in red denotes its *recurrence* neighborhood of size ε. We can see that it has three sampled points as its neighbors according to this definition.

trix with an adjacency matrix of a network whose nodes are the sampled phase trajectory time points, and whose edges represent links between points in the phase space that are within each other's ε-neighborhoods (Marwan et al., 2009; Donner et al., 2010). Thus the three black points that fall within the light red circle around the red point in Fig. 6.1 would be joined to the red point with links in the network corresponding to the depicted neighborhood size ε. Such a complex network is referred to as a *recurrence network*, whose adjacency matrix $\mathbf{A}$ is given by,

$$\mathbf{A} = \begin{cases} \Theta(\varepsilon - \|\vec{x}_i - \vec{x}_j\|), & i \neq j \\ 0, & \text{otherwise} \end{cases} \tag{6.1}$$

where $i, j = 1, 2, \ldots, N$ represent time indices, $\Theta(\cdot)$ the Heaviside function (i.e., $\Theta(\cdot) = 1$ for non-negative values and 0 otherwise) and $\|\cdot\|$ denotes an appropriate distance metric defined for a point $\vec{x}_i$ in the phase space trajectory of the dynamical system X. Note that $A_{ii} = 0$ by definition in order to avoid self-loops.

A crucial advantage of representing a given time series (and by extension the corresponding phase space dynamics) as a network in phase space is that it allows us to study the geometrical and topological properties of the dynamics with complex networks tools and techniques (Donner et al., 2011). The analytical theory behind recurrence networks were further elaborated in Donges et al. (2012), illustrating the link between the discrete complex network based measures such as degree, clustering and betweenness to their continuous equivalents in terms of invariant measures of the phase space trajectory. The studies have collectively shown that recurrence networks are useful in classifying different dynamical behavioral regimes from observed time series (cf. e.g., Zou et al. (2010)).

Recurrence networks have also been applied to paleoclimate datasets and have helped identify nonlinear shifts in paleoclimate variability from an analysis of dust flux proxies going back around 5 Ma BP from marine sediment cores around north and east Africa (Donges et al., 2011). This study reports three major epochs of climatic regimes in the Middle Pliocene, Early Pleistocene and Late Pleistocene, the onsets and ends of which were possible related to significant events in the history of hominid evolution. Another recent study has used recurrence networks to similarly identify nonlinear regime shifts in the Asian monsoon during the Holocene from an analysis of 10 spatially separated records spread over the Asian monsoon domain (Donges et al., 2014).

Similarly, in the following sections, we use recurrence networks to identify dynamical transition events in paleoclimate datasets. However, we use the concept of communities in networks in contrast to the earlier studies, which classified the dynamics based on complex network measures such as clustering coefficient and average path length. Note that the *recurrences* in our analysis are defined with respect to an ε-ball whereas there are, in practice, many different approaches in which we can potentially define the recurrence of a system (Marwan et al., 2007). However, since we have the problem of dealing with a sequence of probability distributions instead of time series, we choose, as a first step, the simplest approach of defining recurrences with an ε-ball.

6.2. Bounds on pairwise recurrence probabilities

To restate our problem of analyzing the probability distributions arising from the method in Chapter 3: in lieu of a time series $x_t,\ t = 1,2,\ldots,N$, we are provided a sequence of (marginal) probability distributions $P_t(x_t),\ t = 1,2,\ldots,N$, with the corresponding probability densities $\varrho_t(x_t)$. The random variables X_t at a given time $T = t$ are neither identical nor independently distributed. Neither do we have any information regarding the joint distributions of the variables X_i and X_j at times $t = i$ and $t = j$. The challenge is to analyze the recurrence properties of the given data, particularly in the context of the framework of recurrence networks, with proper estimations of the recurrence network measures along with uncertainties of estimation. These uncertainties could arise both due to imprecise measurement as well as due to incomplete specification of covariance relations.

The main idea behind our approach is: since the recurrence properties depend primarily on the pairwise differences $x_i - x_j$ we use established theorems from the domain of copula theory to obtain the upper and lower bounds for the probability distributions of these pairwise differences. These bounds depend solely on the marginal distributions and they allow us to estimate bounds on the probability that the difference $x_i - x_j$ is less than a given threshold ε.

Consider the random variable $Z_{ij} = X_i - X_j$ whose cumulative distribution $P_{Z_{ij}}(z_{ij})$ denotes the total probability that $x_i - x_j \leq z_{ij}$. In this terminology, the probability of recurrence within a neighborhood of size ε is given by:

$$P_{ij}(\varepsilon) = P_{Z_{ij}}(\varepsilon) - P_{Z_{ij}}(-\varepsilon). \tag{6.2}$$

The best-possible upper and lower bounds m_{ij} and M_{ij} for the distribution Z_{ij} are given in Williamson and Downs (1990):

$$m_{ij}(z_{ij}) \leq P_{Z_{ij}}(z_{ij}) \leq M_{ij}(z_{ij}), \tag{6.3}$$

where

$$m_{ij}(z) = \max\{\sup_u f_{ij}(u, z_{ij}),\, 0\} \tag{6.4}$$

$$M_{ij}(z) = \min\{\inf_u f_{ij}(u, z_{ij}),\, 0\} + 1 \tag{6.5}$$

$$f_{ij}(u, z_{ij}) = P_i(u) - P_j(u - z_{ij}). \tag{6.6}$$

Using Eqs.6.2-6.6, we can construct the bounds on $P_{ij}(\varepsilon)$

$$P^l_{ij}(\varepsilon) \leq P_{ij}(\varepsilon) \leq P^u_{ij}(\varepsilon), \tag{6.7}$$

where

$$P^l_{ij}(\varepsilon) = \max\{m_{ij}(\varepsilon) - M_{ij}(-\varepsilon), 0\} \tag{6.8}$$

$$P^u_{ij}(\varepsilon) = \min\{M_{ij}(\varepsilon) - m_{ij}(-\varepsilon), 1\}. \tag{6.9}$$

To illustrate the validity of the above bounds on the recurrence probabilities, we consider two normal distributions with different means and standard deviations equal to 1. We then vary the correlation between these two distributions from -1 to 1 in steps of 0.1 and sample their joint distribution using a standard routine for the sampling a bivariate normal distribution available in Numeric Python (NumPy). Following this, we estimate the mean probability of recurrence from 500 samples of size 1000. Also, for a given set of marginal distributions (specified by the marginal means and standard deviations) we estimate the upper and lower bounds from the marginals for different values of the neighborhood size ε using Eqs. 6.7-6.9.

We observe from the results (Fig. 6.2) that the numerically obtained recurrence rates (colored curves) are always within the analytically derived upper and lower

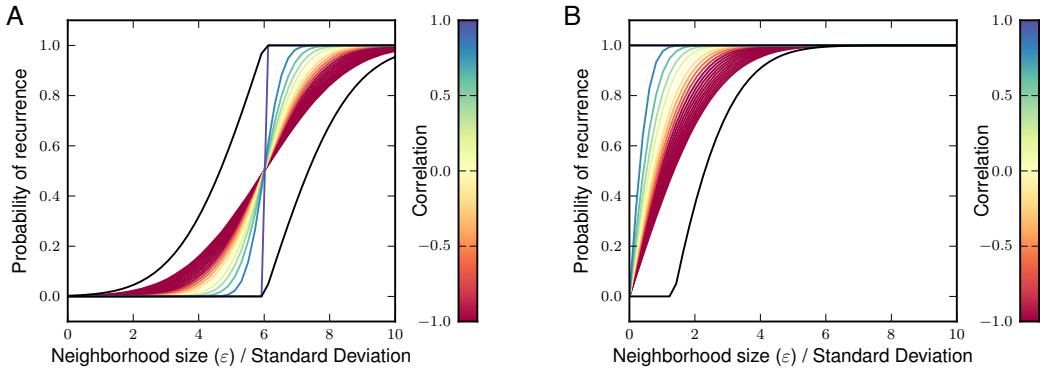

Figure 6.2. Pairwise bounds on recurrence probabilities. The probabilities of recurrences for different values of ε are shown for different values of the correlation of the marginals (shown here as colored curves). The upper and lower bounds for the probabilities of recurrence according to Eqs. 6.7-6.9 are shown as black curves. The means of the marginal distributions are $\mu_1 = -4$, $\mu_2 = 2$ for **A** and $\mu_1 = \mu_2 = 0$ for **B**.

bounds (black curves). However, the estimation of the bound for low values of ε is extremely bad for the case when the means of the marginals are the same (Fig. 6.2B). Another point to note from Fig. 6.2 is that the true recurrence rates do not show any particular pattern of distribution between the two analytical bounds.

6.3. Recurrence networks from pairwise bounds

The pairwise bounds on the recurrence probabilities only give us an upper and lower limit on the *probability* of recurrence of the system between a given pair of time points. In other words, they give us *limits to our confidence* as to *how likely is it* that a given pair of two points could have recurred, provided that we know, with some uncertainty, how they are distributed at those times. By themselves, these bounds do not tell us exactly how likely it it is that the two points could have recurred given the uncertainties of the data.

What we essentially learn from Eq. 6.7 is that given the sequence of probability distributions $P_t(x_t)$, we can confine the probability of recurrence $P_{ij}(\varepsilon)$ only up to a

finite interval $[P^l_{ij}(\varepsilon), P^u_{ij}(\varepsilon)] \subseteq [0,1]$. Note also that we have no way of knowing how the *actual* probability of recurrence is distributed within the bounds $P^l_{ij}(\varepsilon)$ and $P^u_{ij}(\varepsilon)$. Still, if we consider $\mathbf{A}$ to be the true adjacency matrix of the system underlying the given data, we can compute the total marginal probability that $\mathbf{A}_{ij} = 1$. For this we need to take the marginal probability $P_{ij}(p_{ij})$ which is obtained by integrating out all the other remaining $N^2 - 1$ combinations of i, j from the full joint probability $P(p_{11}, p_{12}, \ldots, p_{ij}, \ldots, p_{NN})$.

$$P(\mathbf{A}_{ij} = 1) = \int_{P^l_{ij}}^{P^u_{ij}} dp_{ij}\, P(\mathbf{A}_{ij} = 1 | p_{ij})\, P_{ij}(p_{ij}) \tag{6.10}$$

$$= \int_{P^l_{ij}}^{P^u_{ij}} dp_{ij}\, p_{ij}\, P_{ij}(p_{ij}) \tag{6.11}$$

$$= \mathbb{E}_{P_{ij}}[p_{ij}] \tag{6.12}$$

Thus, the total probability that $A_{ij}(\varepsilon)$ equals one is simply the expectation value of the random variable $P_{ij}(\varepsilon)$ over the interval $[P^l_{ij}(\varepsilon), P^u_{ij}(\varepsilon)]$. Assuming that $P_{ij}(\varepsilon)$ is symmetrically distributed in this interval we find that

$$\varrho(\mathbf{A}_{ij} = 1) = \frac{P^l_{ij} + P^u_{ij}}{2}. \tag{6.13}$$

We define a *probability of recurrence matrix* $\mathbf{P}$ whose elements $\mathbf{P}_{ij}$ equals the above probability given in equation. 6.13. In this framework, we postulate that recurrence characteristics of the 'true' adjacency matrix of the recurrence network $\mathbf{A}$ is qualitatively and accurately describable by $\mathbf{P}$. For low to mid levels of noise, the true adjacency matrix is described $\mathbf{A}$ as

$$\mathbf{A}_{ij} = \begin{cases} 1, & \text{with probability } \mathbf{P}_{ij} \\ 0, & \text{with probability } 1 - \mathbf{P}_{ij}. \end{cases} \tag{6.14}$$

Note that we keep $\mathbf{P}_{ii} = 0$ by definition in order to avoid self-loops in the recurrence network.

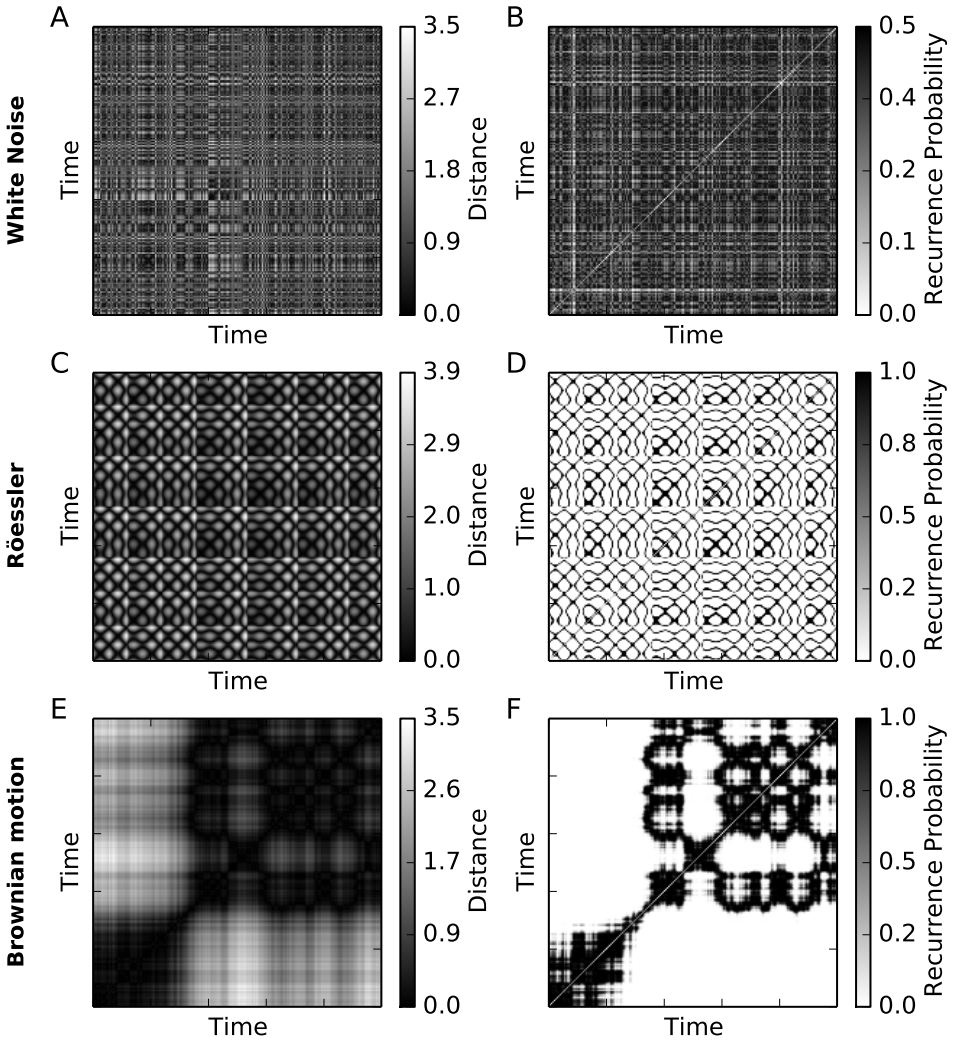

Figure 6.3. The probability of recurrence matrix $\mathbf{P}$. We choose three representative systems at low sampling noise levels—white noise (**A**, **B**), x-component of the chaotic Rössler with $a = b = 0.2$, $c = 5.7$ (**C**, **D**), and Brownian motion (**E**, **F**). The left column shows the 1-dimensional Euclidean distance matrices (**A**, **C**, **E**), whereas the right column shows the derived probability of recurrence matrices as given in Eq . 6.12 (**B**, **D**, **F**). Note the reversed color bars between the two columns, to allow better visual comparison. Low distances correspond to high probabilities of recurrence.

The entries of the probability of recurrence matrix $\mathbf{P}_{ij}$ represent our confidence as to whether the time points $t = i$ and $t = j$ could have recurred or not. This 'confidence' is always with respect to a given choice of the recurrence threshold ε and the marginal distributions $P_t(x_t)$ which encode the uncertainties of measurement. $\mathbf{P}_{ij} = 1$ implies that we are absolutely sure that i and j are recurring time points, whereas $\mathbf{P}_{ij} = 0$ implies that we are absolutely sure that i and j are non-recurring time points.

P incorporates the qualitative features of the recurrences of the dynamical system from which the initial marginal distributions are obtained (Fig . 6.3). We interpret it as a weighted adjacency matrix representative of the recurrence network fo the dynamical system, and whose entries represent the strength of the links between two nodes in the recurrence network. In Fig. 6.3, we take three representative datasets— white noise, the chaotic Rössler system (equation D.5), and Brownian motion (equation D.6), and show the distance matrices along with the derived probability of recurrence matrices **P** for low levels of noise. To simulate correlated time-ordered probability distributions, we first generate a 'true' time series of the chosen system, and then derive a 'sampled' time series from a Gaussian kernel filter of a fixed bandwidth over the 'true' time series. Using the mean and standard deviations of the 'sampled' time series as the parameters for the initial given marginals $P_t(x_t)$, we obtain the probabilities of recurrence as discussed in this chapter. Note that we do not embed the datasets to obtain the recurrences as it is non- trivial to define embedding for a series of non-independent probability densities, and is beyond the scope of this analysis. Apart from there being a general visual correspondence between the left and right columns of Fig. 6.3, we find that low distances in the true times series (left column) correspond to high probabilities of recurrence as obtained using our method (right column). This shows that our method has the ability to reconstruct the qualitative features of dynamical recurrences even in the presence of correlated noisy measurements.

We further look at the local ε-degree density (as defined in Donges et al. (2012)) of the true recurrence network and the estimated probabilities of recurrence for low, medium and high noise levels of sampling (Fig. 6.4) for these representative systems. The local degree density is defined as the degree that can be attributed to a point in the phase space—which is also a node in the recurrence network. Since

we are interested in a qualitative comparison, we can get an idea of the local degree density function by plotting the degree sequence k_t against the corresponding values x_t. For the case of the sampled datasets, we use the mean $\bar{x}_t$ since the real x_t is then not a given quantity.

The results are shown in Fig. 6.4. For white noise, the estimated degrees do not capture the qualitatively flat nature of the true degree density curve. This is most likely due to suprious correlations induced by the Gaussian kernal sampling scheme we use. On the other hand, for the chaotic Rössler and Brownian motion, and for low levels of noise (blue dots), the qualitative features are well captured by the estimated weighted degree from **P**. Note the difference in scale between the true degree and estimated degree. Another point to keep in mind is that the recurrence threshold ε are different for the true recurrence network **A** and the probability of recurrence matrix **P**. This is because the degrees of the true recurrence network are the maximum attainable degrees from the estimated probability of recurrence matrix, in the case where there is no uncertainty. Typically, however, the estimated degrees from **P** are lower than those obtained from **A**. For the comparison in Fig. 6.4, we choose ε such that we get discernible recurrent features in both cases for low sampling noise levels. For medium to high noise levels, the recurrence threshold for **P** are kept at the same level as that for low noise level.

6.4. Detecting paleoclimate transitions using modularity

In this section, we develop a new idea to detect paleoclimate transitions using recurrence network modularity. In the analysis of complex networks, network modularity (Newman, 2004a,b) is widely used to detect communities in a given network. A *community* is a part of the network in which the nodes have a stronger density of edges between them as compared to the density of edges they share with the rest of the network. Communities are an inherent feature of a wide range of real-world complex networks. The edges of most real-world networks are inhomogenously distributed, thereby making it possible to consider the network as being composed of several sub-networks which are strongly connected within themselves and have relatively fewer links among them. For instance, the global

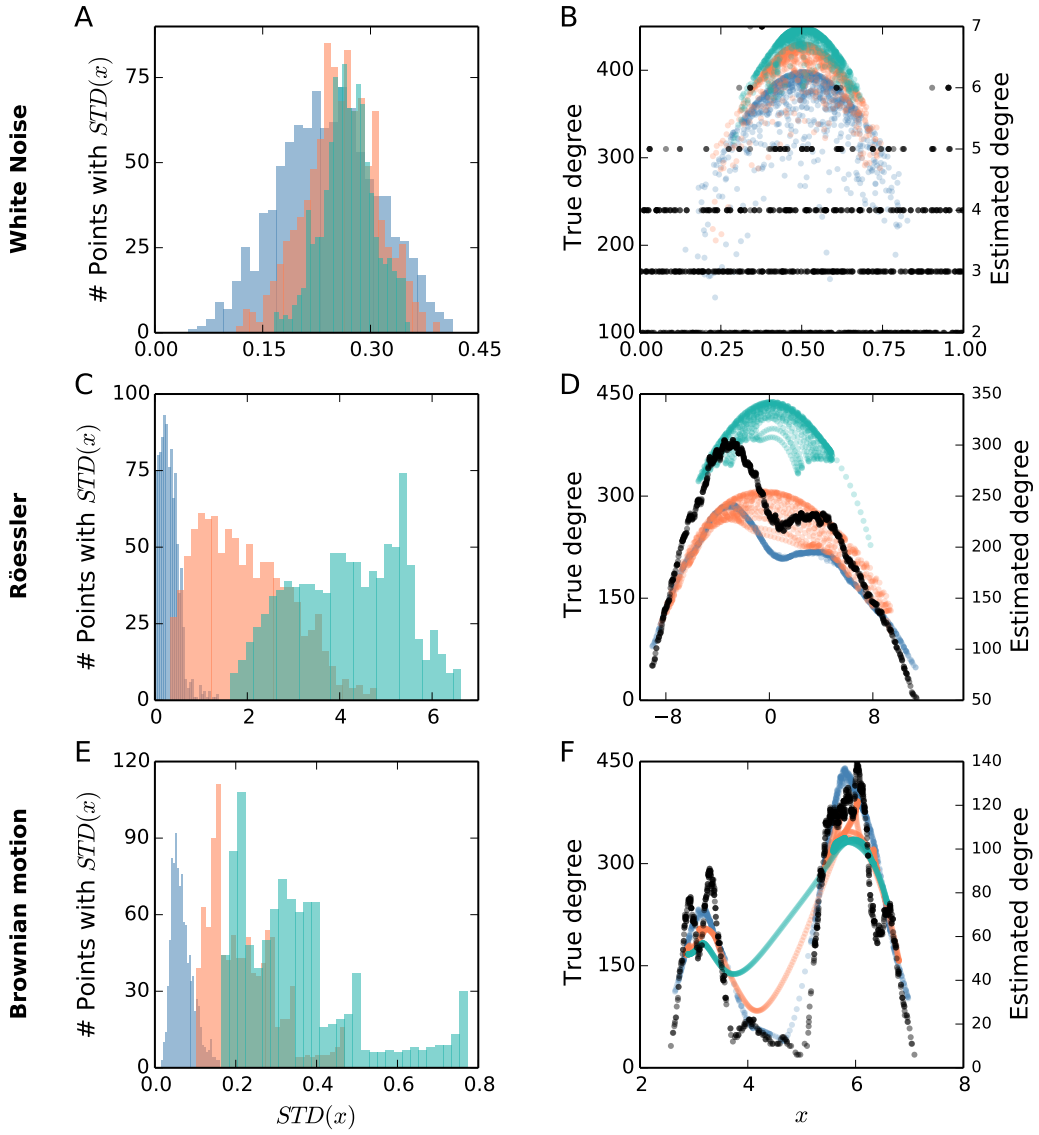

Figure 6.4. Local ε-degree density. Histogram of standard deviation of sampling in the left column (**A**, **C**, **E**) indicates sampling noise —low (blue), medium(orange), and high(green). The right column (**B**, **D**, **F**) shows local ε-degree density against system variable x. For low noise levels (blue dots), the 2-peak nature of the true degree density (black dots) is preserved for the Rössler (**D**) and Brownian (**F**) datasets, whereas for medium(orange dots) to high (green dots) noise levels, this feature is not recovered. The true degree density for white noise (**B**) is not recovered for all noise levels.

airport network can be considered as being composed of communities of smaller domestic airport networks with a much higher number of flights (network edges) operating within each country as compared to the number of international flights between countries.

According to Newman (2004a), the modularity Q_{Π_r} of a given (weighted) undirected network $\mathbf{A}$ with respect to a partition Π_r that divides $\mathbf{A}$ into r different communities, can be formally defined as

$$Q_{\Pi_r} = \frac{1}{2m} \sum_{ij} \left(\mathbf{A}_{ij} - \frac{k_i k_j}{2m} \right) \delta(c_i, c_j), \tag{6.15}$$

where k_i is the (weighted) degree of the i$^{\text{th}}$ vertex, and $m = \frac{1}{2} \sum_{ij} \mathbf{A}_{ij}$ is the number of edges of the network. The community to which a vertex i is assigned to by the partition Π_r is denoted by the *identification vector* $\{c_i\}$. The first term in Eq. 6.15 denotes the fraction of edges that fall within communities whereas the second term denotes the same fraction as would be expected if we were to connect the edges of $\mathbf{A}$ at random.

In the context of recurrence networks, communities have a different implication. The edges of recurrence networks essentially link two points in the phase space which are closer than a chosen distance threshold. Thus, recurrence network communities refer to a sub-spaces of the phase space with higher densities of states than the neighbouring regions of the phase space. This means that if we identify communities in the recurrence network using a community detection algorithm, and without caring for the order of the nodes, each identified community would effectively represent the various basins visited by the phase space trajectory. However, we note that the nodes of our recurrence network are ordered in time. Thus, if we use a two-partition Π_2 that divides the recurrence network into two communities while, at the same time, preserving the time ordering of the nodes, we are able to partition the phase space trajectory in two halves in each of which the system trajectory was confined for some time and then jumped to the other one. Practically, this is achieved by placing a Π_2 partition at a given time point and estimating Q_{Π_2} and repeating this for all N time points. The maximum value of Q_{Π_2} corre-

sponds to the best possible division of the recurrence network in two halves.

$$Q_{\Pi_2}^{\max} = \max\left\{Q_{\Pi_2}(t)\right\}. \tag{6.16}$$

In this way, given a sequence of time-ordered probability distributions, we can construct the probability of recurrence matrix **P** from the densities and then use recurrence network modularity to estimate maximum possible modularity $Q_{\Pi_2}^{\max}$ where a dynamical transition might have occurred. By comparing the maximum modularities of different parts of the trajectory we can thus find out those parts of the trajectory where the maximum modularity achieves extreme values as compared to the rest of the trajectory. These regions are then most likely the times where a sudden change in dynamical behavior could have occurred.

To illustrate this approach, we create a synthetic paleoclimate East Asian Summer Monsoon (EASM) dataset which has an internal dynamics composed of an AR1 component and a sinusoidal component (cf. appendix D for detailed set of equations). This is then forced by a much slower varying sine forcing which can be thought as being similar to the solar forcing component of the EASM in the Holocene (Fig. 6.5A). Finally, the artificial EASM dataset is interrupted by three events—two short 100 year duration events of lower intensity at 7500 and 5000 BP respectively and one relatively stronger event of 500 years starting at 2750 BP and ending at 2250 BP (Fig. 6.5A). The complete *true* monsoon signal is obtained by a superposition of these component signals (Fig. 6.5 B). As with the earlier examples of this chapter, we use a Gaussian kernel filter to simulate the measurement process of the paleo-signal and to ensure that we get a sequence of correlated probability densities at each time point (Fig. 6.5C). The probability of recurrence matrix **P** obtained from these probability densities is shown in Fig. 6.5D. A simple visual inspection of **P** reveals certain 'ruptures' in the general recurrence patterns at the times when the events took place. We show that the events are accurately quantifiable as dynamical events through our recurrence network modularity approach.

We choose a sliding window of size 250 years and move this window from the oldest time point at around 10000 BP to the present with a step size of 10 years. In each window, we estimate $Q_{\Pi_2}^{\max}$, thus obtaining a time-dependent curve $Q_{\Pi_2}^{\max}(t)$

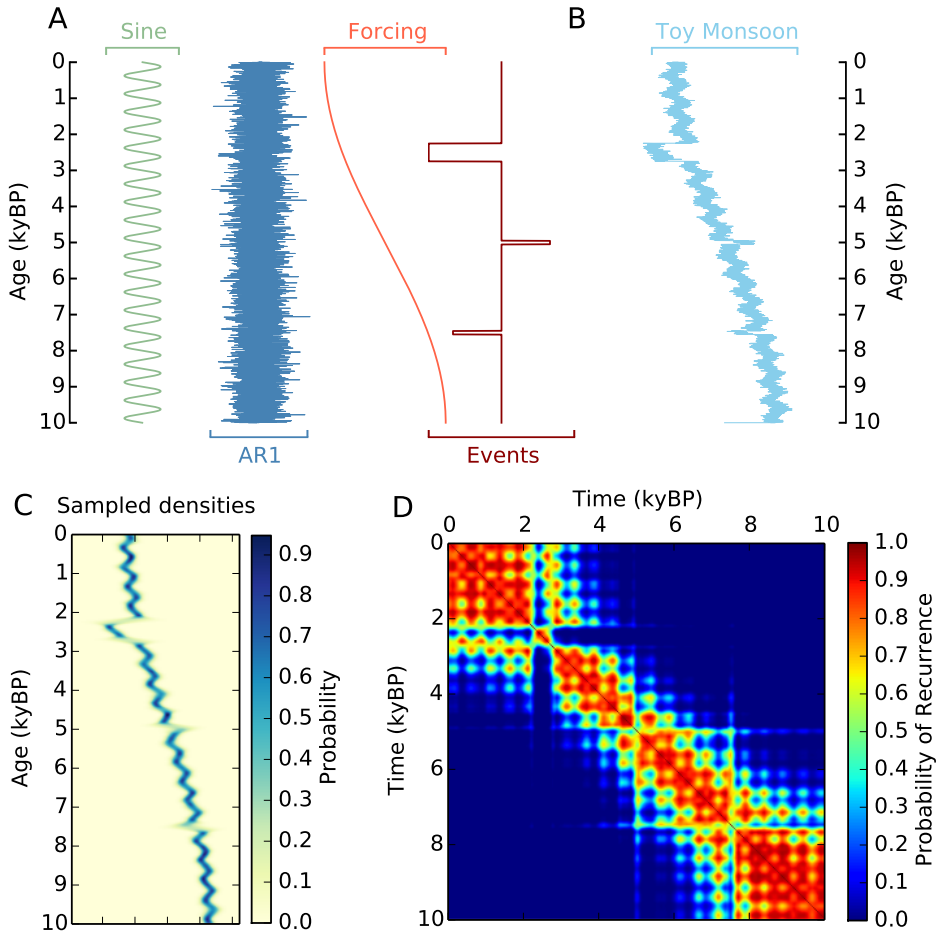

Figure 6.5. Synthetic EASM example. A. The components of the synthetic East Asian Summer Monsoon (EASM)—the internal fast-varying sine component (green), the AR1 component (blue), the slow varying solar forcing component (red) and the event component (maroon). The three events are centered at 7500, 5000, and 2500 BP, and last for 100, 100 and 500 years respectively. The event at 5000 BP has opposite sign than the other two. **B.** The *true* monsoon signal as a superposition of all four components shown in **A**. **C.** Probability densities obtained by sampling the signal in **B** with a Gaussian kernel of 7 years width. **D.** The probability of recurrence matrix $\mathbf{P}$ obtained by using the approach discussed in this chapter from the sequence of densities shown in **C**.

which shows the maximum modularity possible for each time window centered at time t. Furthermore, to ensure that the modularity values obtained in the analysis could not have arose by chance, we shuffle the time order of the probability densities in each window and estimate the maximum modularity for each shuffling. This is repeated for a 100 times in each time window. This leads to a null distribution of modularity values obtainable for a random ordering of the probability densities given in the dataset. The null distribution is obtained by considering the randomly shuffled densities in each window as being from a single underlying distribution corresponding to the null hypothesis that there was no transition in any chosen window of length 250 years. Then, at a 10% significance level, we check for those modularity values in the curve $Q^{\max}_{\Pi_2}(t)$ which are greater than the significance threshold value of modularity. This means that these significance values have only a 10% probability to have come up by random chance.

The curve $Q^{\max}_{\Pi_2}(t)$ is shown in Fig. 6.6A. We clearly observe four significant peaks at 7562, 4922, 2732, and 2242 BP respectively. These correspond to the events in the original monsoon signal at 7500, 5000, and around 2500 BP. Note that we get two peaks for the event occurring from 2750 BP to 2250 BP because the window size of 250 years is small enough to record *the start and the end* of this event as two separate transitions. For the other two older events centered at 7500 BP and 5000 BP, each lasting for a 100 years, the window size of 250 years fails to resolve the start and the end of these short events and thus they show up as a single peak. Fig. 6.6B shows the relative position of the significant modularity values in the null distribution of modularity values, highlighting the fact that the chances of obtaining these values from the null distribution are very low.

Another point that we learn from this analysis is that our approach cannot distinguish between events of different direction. The event at 5000 BP was of a different sign than the ones at 7500 and 2500 BP but in the final modularity curve, they all show up simply as events, without any information as to the direction in which the system was forced by the event. In the context of analyzing paleo-monsoon dynamics, this means that we are unable to differentiate between “dry” and “wet” events.

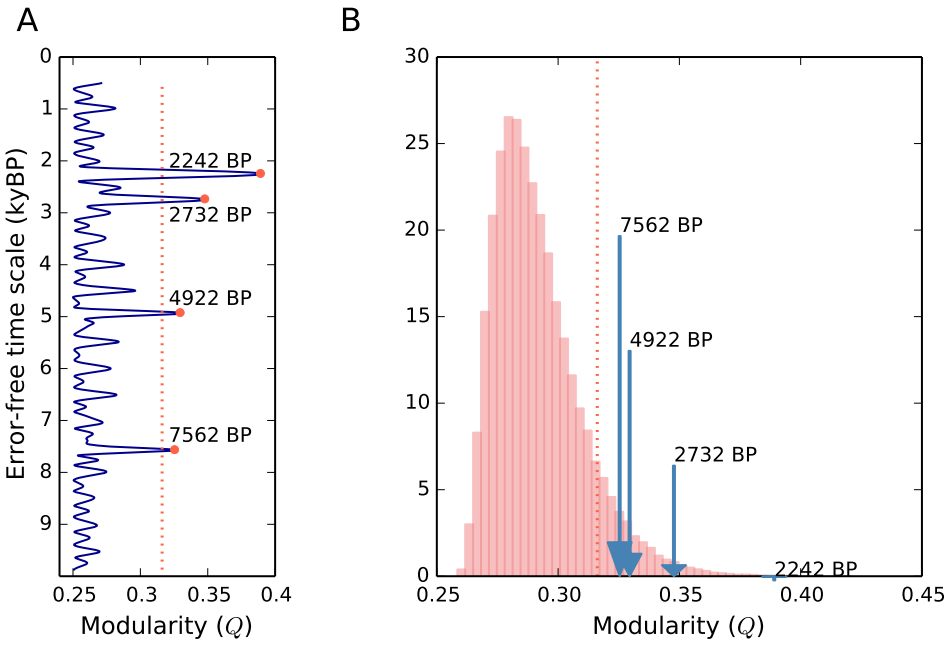

Figure 6.6. Detection of events in the synthetic EASM dataset. A. The curve for maximum modularity $Q_{\Pi_2}^{\max}(t)$ for each window of 250 years centered at time t. The four significant peaks (at 10% significance) at 7562, 4922, 2732, and 2242 BP are marked as red dots. **B.** The location of the significant modularity values in the null distribution. The vertical red dotted line in both **A** and **B** denotes the significance threshold modularity value at 10% significance level.

6.5. Holocene events in the East Asian Monsoon

The EASM is the eastern branch of the Asian Monsoon domain which roughly approaches towards the mainland via the South China Sea. It is a critical climatic subsystem which is impacted by various other global climatic factors. While it is highly essential to understand its dynamics and how they relate to climate forcing such as solar variability, greenhouse gases, and stratospheric aerosols, there is still much that is left to be understood. In the past decade, a surge of new paleoclimate proxy records from the EASM domain have shed light on its behavior in the past, and especially in the Holocene. However, there is much debate on the interpretation of these records and how the behavior shown by these isotopic proxies are to be interpreted in the context of past climatic conditions (see, e.g., Maher (2008)).

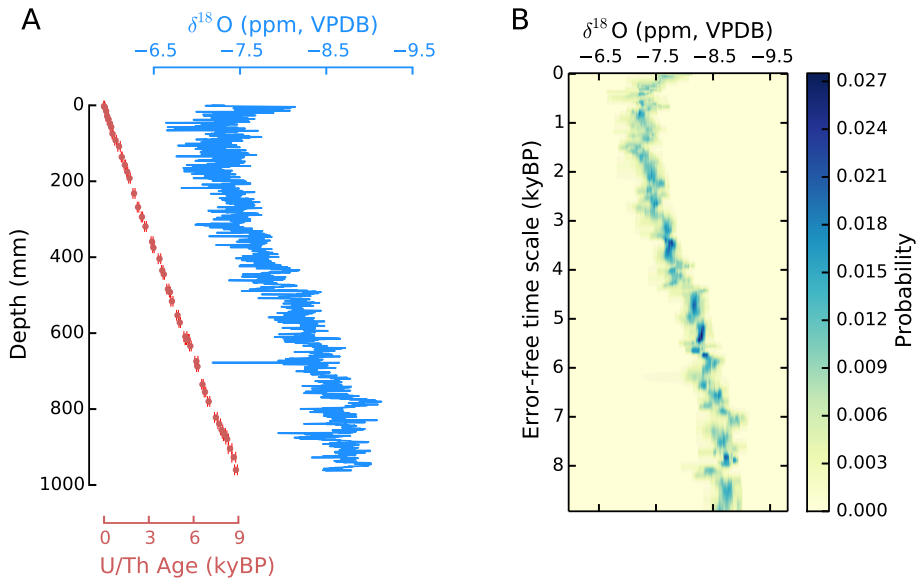

Figure 6.7. The Dongge cave dataset. A. The U/Th age vs. depth measurements (red circles with error bars) and the δ^{18}O proxy measurements along depth (blue curve) from Dongge cave in south-eastern China. **B.** The proxy record probabilities estimated according to the approach described in Chapter 3. The resulting probability densities are represented on an error-free time scale as all the uncertainties are shifted to the proxy axis.

In general, the EASM shows a long term trend towards a weaker monsoon intensity over the past 10000 years after reaching what is called the 'Holocene climate optimum' at around 8000 years before present (Wang et al., 2005a; Hu et al., 2008). However, this trend is interrupted by several 'dry' monsoon events where, for durations that range from one to five centuries, the monsoon drops down to almost drought-like conditions. These 'events' have been further linked by Wang et al. (2005a) to societal change, and even collapse in the case of the Chinese Neolithic culture at around 4500 BP. Such a collapse has been also reported by other studies such as that of Liu and Feng (2012).

Other aspects of the dynamics of the EASM, such as its teleconnections to other global climatic systems (Johnson, 2011; Liu et al., 2013), its connection to solar variability (Duan et al., 2014; Bard and Frank, 2006), and the connections between rapid climate change (RCC) events in the EASM in the Holocene to other global

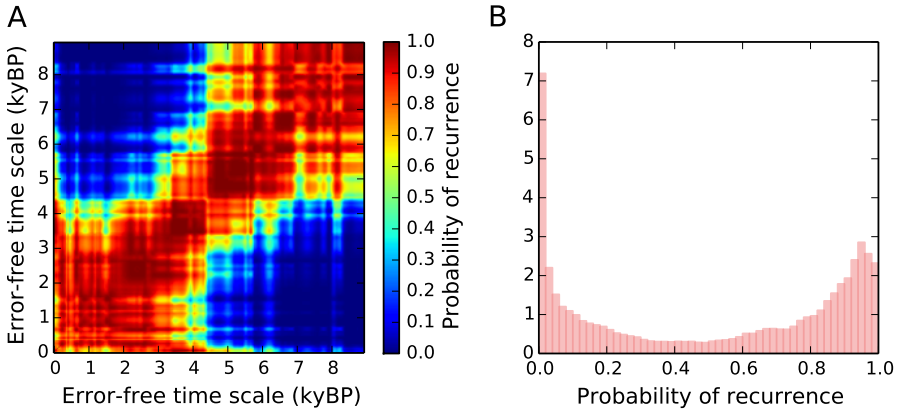

Figure 6.8. Recurrence probabilities for Dongge Cave. A. The probability of recurrence matrix $\mathbf{P}$ estimated from the proxy densities shown in Figure 6.7B according to Eq. 6.13. **B.** The histogram of the entries shown in **A**. Most values are either close to zero or one.

RCCs (Wang et al., 2013; Overpeck and Webb, 2000), are still not well-understood. Particularly relevant for our analysis are the RCCs reported by Bond et al. (2001), referred to commonly as *Bond events*, the signatures of which have since then been reported in a wide range of records from across the globe. In their paper, Wang et al. (2005a) too have reported the 'synchronous' timing of several dry monsoon events in their dataset with the Bond events based on a visual inspection of their EASM proxy record. In this study, we consider the Dongge cave record from the same study by Wang et al. (2005a) from southeastern China which covers approximately a little over the last 9000 years. This is a δ^{18}O isotopic proxy record, obtained from a stalagmite in the Dongge Cave, that is anti-correlated to the summer precipitation in that region. It is an extremely well- dated EASM record with 45 U/Th dates used to establish the chronology of the archive, and 2124 δ^{18}O measurements along the depth of the stalagmite (Fig. 6.7A).

We apply the proxy estimation method discussed in Chapter 3 to reconstruct the proxy record for the δ^{18}O EASM proxy from Dongge cave and thus represent it on an error-free time scale (Fig. 6.7B). We then use the method discussed in the pre-

vious sections of this chapter to estimate the dynamical recurrences of the dataset with the probability of recurrence matrix **P** (Fig. 6.8A). A histogram of the probability of recurrence values shows that most of the probabilities are either very close to zero or one (Fig. 6.8B).

From the recurrence network represented by **P**, we then estimate the maximum modularity for a 2-partition $Q^{\max}_{\Pi_2}(t)$ for a sliding window of 250 years centered at time t BP. The window is moved from the oldest time point in the record ($\approx$ 8500 BP) to the present with a step size of 10 years. As with the illustrative example in Section 6.4, we estimate the statistical significance of the obtained $Q^{\max}_{\Pi_2}(t)$ values using a null hypothesis that there was no event in the randomly chosen window of 250 years width. Our approach presents a novel way of accurately determining events in the paleoclimate record represented by the Dongge cave δ^{18}O data.

We identify three highly significant events in the windows centered at 8224, 6874, and 4364 BP (Fig. 6.9A). These events are significant according to our null hypothesis at a significance level of 1% (Fig. 6.9B). The uncertainty of the timing of these events as obtained from our method is equal to the width of the window size considered for the analysis i.e., 250 years. In other words, we are able to narrow down the timing of these events roughly to periods of 8349–8099 BP, 6999–6750 BP, and 4489–4239 BP. The 4.4 kBP and 8.2 kBP events have also been noted by Wang et al. (2005a) in their study and have been connected to the collapse of the Chinese Neolithic Culture and an ice rafting event of the Laurentide sheet respectively. These are connected in turn to the Bond events numbered 3 and 5 respectively as given in Bond et al. (2001). Both these events correspond to extremely weak monsoons leading to drought-like conditions which lasted for a few centuries. Our method is thus able to determine these widely known RCCs of the Holocene period from a consideration of the recurrence characteristics of the dynamics of the dataset.

The event centered around 6.9 kBP has not yet been reported to the best of our knowledge. However, this event is quite similar to the events at 8.2 kBP and 4.4 kBP in the sense that it led to a sudden drop in the monsoon intensity over a short period of time (Fig. 6.10). The δ^{18}O value before the event was approximately -8.84 and this value increased to around -8.50 at the end of the event. This corresponds to roughly a 3.8% change in the magnitude of the mean δ^{18}O value. Note that

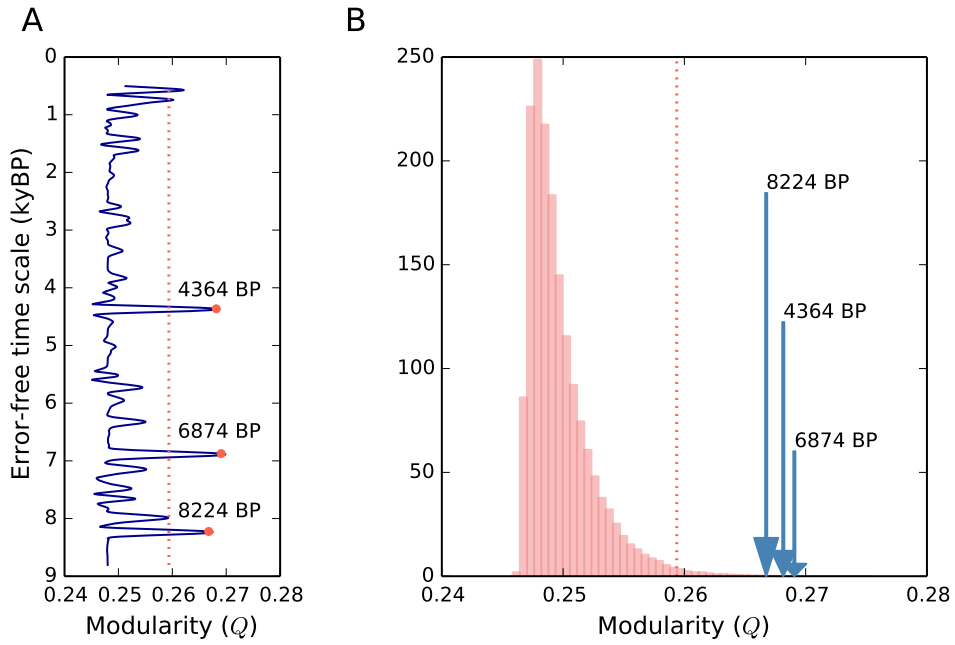

Figure 6.9. Events in the EASM in the Holocene. A. Maximum modularity $Q_{\Pi_2}^{\max}(t)$ for each window of 250 years centered at time t for the Dongge cave δ^{18}O EASM proxy. The three significant peaks (at 1% significance) at 8224, 6874, and 4364 BP are marked as red circles. **B.** The location of the significant modularity values in the null distribution. The vertical red dotted line in both **A** and **B** denotes the significance threshold modularity value at 1% significance level.

the similar magnitudes of change in the mean δ^{18}O values for the 4.4 kBP and 8.2 kBP events are 3.2% and 3.4% respectively. The fact that the 6.9 kBP event has not been reported yet indicates the need for developing and applying quantitative methods for analyzing occurrences of climate shifts to a much larger paleoclimatic database. We intend to carry out a similar analysis of detecting events in other monsoon records spread over the Asian monsoon domain in the hope that the results thereafter will help shed light on the 6.9 kBP dry event that we find in the Dongge cave record.

In general, the analysis presented in this chapter helps illustrate several key points:

- It is possible to analyze and study paleoclimatic proxy datasets keeping in

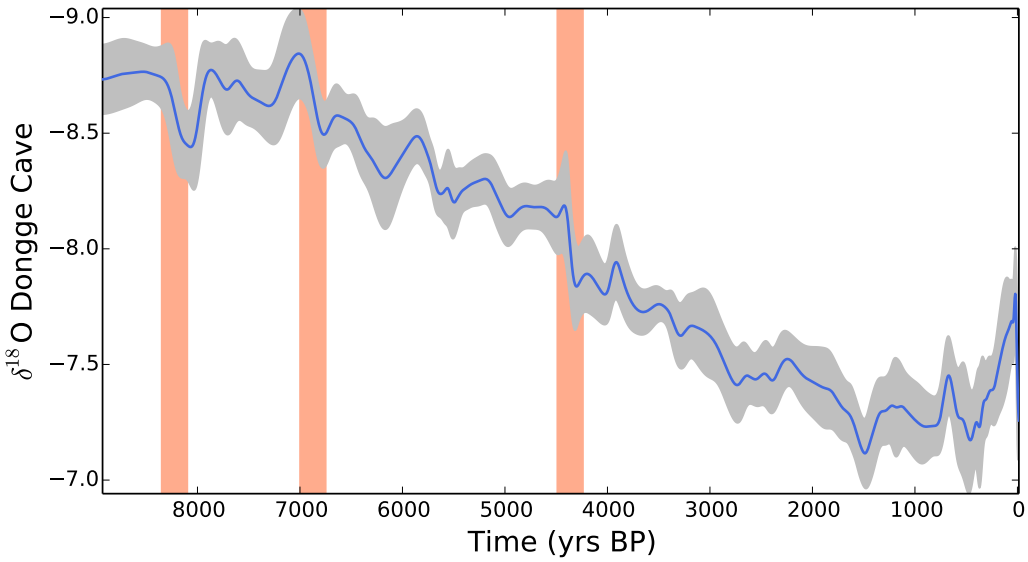

Figure 6.10. The three most strongest events in the Dongge dataset. The mean expected value of the Dongge cave δ^{18}O values at each time point (blue) along with ± 1 standard deviation (gray) of the posterior proxy probabilities at that time point (from Fig. 6.7B). The three most strongest and statistically significant (at 1% significance level) events (from Fig. 6.9) are shown as vertical orange bars of 250 years width centered at the time points of the window.

mind their uncertainties of estimation and in spite of unknown correlations between them.

- Such paleoclimate studies can reveal new features of the dynamics underlying the considered dataset, along with a suitable confidence in (or conversely, the uncertainty of) the deduced results.

- There is pressing need to apply mathematically rigorous methods to the ever increasing corpus of paleoclimate proxy data so that we may understand and infer paleoclimatic conditions more reliably.

Chapter 7

Directed network of global temperature drivers

Till this point, we were primarily concerned with (a) paleoclimate datasets, and (b) uncertainties that arise due to inherent measurement errors in such datasets. This is but only one of the many possible sources of uncertainties that we encounter in our efforts to understand the climatic systems and mechanisms. A crucial problem in the study of climate is the correct estimation of the interrelations that exist between various climatic systems. While in principle it is possible to understand climatic interrelations with a thorough treatment of the underlying physical mechanisms, for a complex system such as climate, it is highly non-trivial to do so in practice. That is why we supplement physical modeling of the climate with empirical analyses that uncover its interrelation patterns. However, as with any empirical analyses, we end up with the following uncertainties: the uncertainty of not knowing precisely the intensity with which various climatic subsystems impact each other, and also the uncertainty if not knowing precisely how these interactions are mediated (involving, e.g., the time delay of interactions).

In this chapter, we attempt to answer this question by focusing on the global mean temperature (GMT). We consider a GMT dataset that goes back to around a 120 years from the present and try to understand the various other relevant climatic subsystems that could influence the GMT. In the light of earlier studies like Lean and Rind (2008), we consider four major factors that could potentially impact GMT, viz., greenhouse gases, the El Niño Souther Oscillation (ENSO), volcanic activity and solar irradiance. However, in addition to earlier work on this issue, we present in Goswami et al. (2013) a new nonlinear measure of influence that, when combined with a time-delayed analyses, allows us to infer directed influences between the various datasets considered. Also, besides trying to estimate those influences

which could potentially impact GMT, we also estimate influences that the potential GMT drivers could have among themselves, with the understanding that the results would present a more holistic view of the network of interrelations that governed the dynamics of GMT over the past century.

In agreement with earlier studies, we find the major drivers for GMT to be ENSO, greenhouse gases, solar irradiance and volcanic activity, on timescales that range from a few months to decades. However, we are also able to uncover a critical feedback between the ENSO and GMT. Our results reiterate the pressing need to revise our imagination of climatic interrelations to one that involves multiple, delayed interactions within the drivers of global mean temperature.

7.1. Introduction and Background

Global mean temperature is perhaps the most critical measurable quantity that serves as an indicator of our planet's climate and modern climate change. It remains till date the most crucial variable of interest in the research that is focused on global climate change and variability. Particularly of interest— especially to the media and the larger society outside the scientific community —has been the role of anthropogenic forcing in the warming trend of GMT. The Fourth IPCC Assessment report clarifies this matter to a large extent, especially in the context of the time period of the last century (Hegerl et al., 2007). However, even while there is a scientific consensus on how anthropogenic greenhouse gases dominate the observed rise in GMT, there are persistent uncertainties in the answers to questions such as: *how much* of the temperature change is induced *by which factors*, and *how* do they change the temperature. These two questions combine two interwoven threads of GMT studies till date: (i) the effort to estimate and understand GMT variability of the past, and (ii) the the effort to predict how the climate (as mostly represented by the GMT) is going to change in the future under different scenarios of climate change.

Increasing number of observational records and 'reanalysis' studies for the last century have enabled us to get more meaningful estimates of the GMT that better represent the state of the Earth's climate (cf. Section 3.2 of Trenberth et al. (2007), and Hansen et al. (2010)). Such records have allowed us to narrow down the list of

potential drivers of GMT to primarily four important climatic factors: the ENSO, greenhouse gases, volcanoes, and solar activity. Still, there is significant amount of debate over the precise nature of the temperature response, exemplified by the diverse explanations put forward to explain the 'deceleration' of global warming in the first decade of the 21st century (Benestad, 2012). In this context, the studies by Lean and Rind (2008, 2009) and Foster and Rahmstorf (2011) play a critical role in shedding light on how the various climatic factors contribute to GMT variability. Lean and Rind use multivariate regression models to estimate the fraction of variability of the GMt that arises due to the variability of anthropogenic greenhouse gases, ENSO, solar irradiance and volcanic aerosols. They further project the GMT to the future based on their empirically derived model under various climate change scenarios. Complementary to this, Foster and Rahmstorf present a detailed analyses that estimates the residual anthropogenic impact on the warming trend of the GMT in the period from 1979 to 2010. They achieve this by first removing the variability that arises due to all other known natural forcing of GMT.

However, there are two fundamental assumptions underlying these studies: (i) independence *between* the factors that impact temperature, and (ii) a linear superposition of their effects on temperature variations. Note that it might be the case that both of the above assumptions hold true (to a reasonable degree of approximation). Still, we feel that a proper assessment of their validity is necessary and more so, such a verification is lacking in the existing body of literature to the best of our knowledge.

The main focus of the following sections of this chapter is to assess the validity of the first of the above two assumptions. In other words, we wish to study the way in which the various drivers of GMT interact with each other and whether it is reasonable to assume that these interactions could be negligible while trying to study GMT variability. In order to do so, we use an empirical approach based on the recurrences of dynamical systems. In particular, we use a recurrence-plot-based similarity measure to analyze the level of *non-independence* among the considered climatic datasets. Recurrence plots are an empirical time series analysis tool that find usage in studying a wide variety of studies, ranging from cardiogram data (Zbilut et al., 2002), stock markets (Bastos and Caiado, 2011; Goswami et al., 2012), speech comprehension (Richardson and Dale, 2005), and even for

cover song identification (Serrà et al., 2009). Marwan et al. (2007) provides a comprehensive overview of the theory, practice, and applications of recurrence plots, while a history of the field of recurrence-plot-based time series analyses can be found in Marwan (2008). In climate analyses, cross recurrence plots have helped to uncover interrelations between the ENSO and northwest Argentinian precipitation in both modern and paleoclimatic timescales (Marwan et al., 2003). More recently, joint recurrence plots have been used to study how the links between climatic variables (such as temperature and precipitation) and vegetation are geographically distributed over China (Li et al., 2011).

We propose in Goswami et al. (2013) an extension of the measure for lag and generalized synchronization as developed in Romano et al. (2005). Since our measure is exclusively based on joint recurrence plots, it is suitable for studying the interactions between structurally different systems as in our case. Moreover, we provide a more general, lagged version of our similarity measure which helps reveal delayed influences from one time series to the other, indicating the driver and the driven systems as a result.

7.2. Lagged dependencies from recurrences

Recurrence plots were first introduced by Eckmann et al. (1987) with a one-line abstract that said:

> A new graphical tool for measuring the time constancy of dynamical systems is presented and illustrated with typical examples.

Theoretically, a dynamical system *recurs* to an earlier state if it comes arbitrarily close to it (up to a tiny deviation). However, in practice, recurrences are estimated from a given time series when the system underlying the time series returns to an earlier *neighborhood* in its phase space.

Consider a system X where $\mathscr{X}$ is the set of all possible trajectories. We can then define a recurrence matrix $\mathbf{R}^X$ for a given trajectory $\vec{x} = (\vec{x}_1, \vec{x}_2, \ldots, \vec{x}_N) \in \mathscr{X} \subseteq \mathbb{R}^m$ as:

$$\mathbf{R}^X_{ij}(\varepsilon) = \Theta(\varepsilon - \| \vec{x}_i - \vec{x}_j \|), \quad i,j = 1,\ldots,N, \tag{7.1}$$

where N is the trajectory length, ε the neighborhood size, $\Theta(\cdot)$ the Heaviside function (i.e., $\Theta(\cdot) = 1$ for non-negative values and 0 otherwise) and $\| \cdot \|$ denotes an appropriate distance metric defined for $\vec{x} \in \mathscr{X}$. A *recurrence plot* (RP) is a visual representation of $\mathbf{R}$ typically constructed by putting a black dot for every 1 in $\mathbf{R}$ (cf. Marwan et al. (2007) for an overview of recurrence plots).

Now consider an arbitrary trajectory $\vec{x} \in \mathscr{X}$. The probability $P(\vec{x} \sim x_i)$ that the trajectory $\vec{x}$ visits the neighborhood of $x_i \in \vec{x}$ is equal to the column-sum of the recurrence matrix $\mathbf{R}^X$ (from Eq.7.1),

$$P(\vec{x}_i) = \frac{1}{N} \sum_{j=1}^{N} \mathbf{R}^X_{ij}. \tag{7.2}$$

where $P(\vec{x}_i)$ is used as shorthand for $P(\vec{x} \sim \vec{x}_i)$. In this formulation, the mean probability that *any* trajectory of the system recurs to *any* given state (given the observed time series) is the mean $\langle P(\vec{x}_i) \rangle = \sum_{i=1}^{N} P(\vec{x}_i)/N$, and is known as the (global) *recurrence rate* (RR^X) of X.

In our analysis, we wish to compare the recurrence structures of two structurally different systems such as the ENSO and solar irradiance which, in principle, exist in completely different phases spaces. For situations such as this, we can use joint recurrences which simply represents the co-occurrences of recurrences in both systems. A joint recurrence matrix $\mathbf{JR}^{XY}_{ij}$ for a given pair of systems X and Y is defined as

$$\mathbf{JR}^{XY}_{ij}(\varepsilon_x, \varepsilon_y) = \Theta(\varepsilon_x - \| \vec{x}_i - \vec{x}_j \|)\, \Theta(\varepsilon_y - \| \vec{y}_i - \vec{y}_j \|) \quad i,j = 1,\ldots,N, \tag{7.3}$$

where $\vec{x}$ and $\vec{y}$ are two given trajectories of systems X and Y respectively. A non-zero value $\mathbf{JR}^{XY}_{ij}$ thus captures co- occurring recurrences of trajectories $\vec{x} \in \mathscr{X}$ and

$\vec{y} \in \mathscr{Y}$ in corresponding neighborhoods of systems X and Y respectively, i.e., $\vec{x}_i \sim \vec{x}_j$ and $\vec{y}_i \sim \vec{y}_j$ at the same time instant i (or j). Note that, due to the different dynamical properties of systems X and Y, the neighborhood sizes ε_x and ε_y used to define recurrences are not identical in general.

Based on the joint recurrence matrix $\mathbf{JR}^{XY}$, we can now define the joint probability that two arbitrary trajectories $\vec{x} \in \mathscr{X}$ and $\vec{y} \in \mathscr{Y}$ recur in the neighborhood of $\vec{x}_i$ and $\vec{y}_i$ simultaneously,

$$P(\vec{x}_i, \vec{y}_i) = \frac{1}{N} \sum_{j=1}^{N} \mathbf{JR}_{i,j}^{XY}, \tag{7.4}$$

where, similar to Equation 7.3, $P(\vec{x}_i, \vec{y}_i)$ is used as shorthand to denote the joint probability $P(\vec{x} \sim \vec{x}_i, \vec{y} \sim \vec{y}_i)$.

Romano et al. (2005) put forth a way in which joint recurrences could be used to estimate whether two systems where in generalized or lag synchronization from their observed time series. Their idea was founded on two crucial results: (i) it is possible to reconstruct the topology of the phase space from the recurrence plot (Thiel et al., 2004), and (ii) if two systems are related by a functional relationship (which is the case for general or lagged synchronization), then they will have similar recurrence plots. Moreover, they argue that if two systems are specifically in generalized synchrony, their respective global recurrence rates $\langle P(\vec{x}_i) \rangle$ and $\langle P(\vec{y}_i) \rangle$ are approximately equal to each other and are in turn also approximately equal to the joint recurrence rate $\langle P(\vec{x}_i, \vec{y}_i) \rangle$. Using the constant recurrence rate method for constructing recurrence plots (cf. (Marwan et al., 2007)), they define a measure for synchronization

$$JPR = \max_{\tau} \frac{S(\tau) - R_0}{1 - R_0}, \tag{7.5}$$

where $R_0 N = n_0 = \langle P(\vec{x}_i) \rangle = \langle P(\vec{y}_i) \rangle$ by construction, and

$$S(\tau) = \frac{\langle P(\vec{x}_i, \vec{y}_i(\tau))\rangle}{R_0}. \tag{7.6}$$

Here, $\langle P(\vec{x}_i, \vec{y}_i(\tau))\rangle$ represents the joint recurrence rate obtained when the time series form system Y is shifted by τ time steps ahead of X.

We propose an extension of the above idea to the more general case where the overall recurrence rates $\langle P(\vec{x}_i)\rangle$ of X and $\langle P(\vec{y}_i)\rangle$ of Y are not equal to each other. Moreover, we also consider the situation that $P(\vec{x}_i) \neq P(\vec{y}_i)$. In particular, we define a *recurrence-based measure of dependence* RMD as

$$RMD_i = \frac{P(\vec{x}_i, \vec{y}_i)}{P(\vec{x}_i)P(\vec{y}_i)}, \tag{7.7}$$

which quantifies the level of 'non-independence' between the recurrence characteristics of systems X and Y as quantifiable from the given observed pair of time series. Note the probabilistic ratio term on the right-hand side of Equation 7.7 is similar to the *odd-to-expected ratio* (Miettinen and Wang, 1981), a quantity used in medical analyses to estimate the dependence between multiple measured variables. This point also highlights a subtle difference between RMD_i in Equation 7.7 and the earlier synchronization estimator JPR in Equation 7.5. Although, RMD_i is motivated from a consideration of the dynamical recurrence characteristics, it is more of a statistical measure of the dependence between the given pair of systems under study.

An additional requirement of estimating global climatic interrelations as in our case is to be able to account for delayed interactions, as climatic subsystems are often coupled to each other with considerable lag. We argue that, similar to $S(\tau)$ in Equation 7.6, it is natural to incorporate lagged probabilities in RMD_i simply by introducing a relevant lag in one of the systems. We thus define the log-mean $RMD(\tau)$ at lag τ as:

$$RMD(\tau) = \log_2\Big(\frac{1}{N'}\sum_{i=1}^{N'} RMD_i(\tau)\Big), \tag{7.8}$$

where $N' = N - \tau$, and $RMD_i(\tau) = P(\vec{x}_i, \vec{y}_i(\tau))/(P(\vec{x}_i)P(\vec{y}_i(\tau)))$, if we consider Y to be shifted by τ units. For two independent systems X and $Y(\tau)$, $P(\vec{x}_i, \vec{y}_i(\tau)) = P(\vec{x}_i)P(\vec{y}_i(\tau))$, which implies that $RMD(\tau) = 0$. For $\tau > 0$, non-zero RMD implies that Y is dependent on X and the converse is true for $\tau < 0$. Note that RMD can quantify both uni- and bi-directional dependencies as well as multiple lags at which the systems might influence each other. Based on the lagged analysis of RMD as given in Equation 7.8, we also get an indication of the direction of influence between the two systems. Note that in earlier studies, recurrence-based measures such as the *Mean Conditional Probability of Recurrence* (MCR) (Romano et al., 2007; Zou et al., 2011) were proposed to infer the coupling direction between two systems, however the MCR measure was not defined for an explicitly lagged analysis. In our approach, which is somewhat analogous to a lagged correlation analysis, it is easy to infer at which time delays are the values of RMD the most significant.

7.3. Testing for significance of observed values

In analyses of connectivity between pairs of experimental datasets, the measures used to quantify dynamical similarity often yield values that are intermediate, and we then cannot easily conclude whether the pair of time series are strongly or weakly connected. Even if we choose to define our measure so that it always lies within a finite interval – such as between 0 and 1 – experimental datasets can typically give ambiguous results of intermediate connectivity values such as 0.6. The connectivity measures obtained from such *passive experiments* cannot thus provide an unambiguous interpretation due to the lack of a comparative 'test case'. This is in contrast to *active experiments* (made, e.g., from laboratory experiments or numerical simulations of models) where the obtained values can be used for a consistent interpretation by its comparison to the uncoupled case. In passive experiments, the observed values of connectivity have to be statistically tested to ensure that they could not have been obtained by random chance. The statistical

test is carried out using surrogate datasets generated from the observed time series in conjunction with an appropriate null hypothesis.

Surrogate time series are different from the original, observed ones, and yet preserve essential dynamical properties. There are several ways of generating surrogates and each method has its respective null hypothesis. The test statistic (which is RMD in our case) is calculated for a sufficiently high number of surrogate pairs and the observed value is tested for statistical significance using the obtained distribution and the appropriate null hypothesis.

For our analysis, we use a recurrence-based approach to generate *twin surrogates* (TS) from the observed time series (Thiel et al., 2006, 2008). *Twins* are two points $\vec{x}_i$ and $\vec{x}_j$ of X such that they share the same neighborhood up to the limit ε, i.e., for $k = 1,2,\ldots,N$, $\mathbf{R}^X_{k,i} = \mathbf{R}^X_{k,j}$. The TS method requires that we first identify all possible twins given an observed trajectory $\vec{x}$. To generate the surrogate series $\vec{s}$, we then choose an arbitrary random point $\vec{x}_k \in \vec{x}$ and set it as $\vec{s}_1$. Now, given $\vec{s}_i = \vec{x}_l$, we append subsequent points to $\vec{s}$ iteratively according to the following rule: when $\vec{x}_l$ has no twins, $\vec{s}_{i+1} = \vec{x}_{l+1}$; on the other hand when $\vec{x}_l \in T$ such that $T = \{\vec{x}_l\} \cup \{\vec{x}_m : \vec{x}_m \text{ and } \vec{x}_l \text{ are twins}\}$ and the number of elements in T is n, then $\vec{s}_{i+1} = \vec{x}_{k+1}$ where $\vec{x}_k \in T$ with probability $1/n$.

The null hypothesis for TS is that each surrogate trajectory is an independent realization of the system corresponding to a different initial condition. To test whether the observed value of $RMD(\tau)$, between X and $Y(\tau)$, is a statistically significant measure of X driving Y we do the following: (i) generate TS of Y, (ii) obtain a test distribution of $RMD(\tau)$ using the observed time series of X and the surrogates of Y, (iii) construct a 95% confidence band from the area between the 2.5^{th} and 97.5^{th} percentiles. This interval represents the region where we fail to reject the null hypothesis.

Observed values of RMD outside the confidence band imply a statistically significant dependence between X and Y at delay τ.

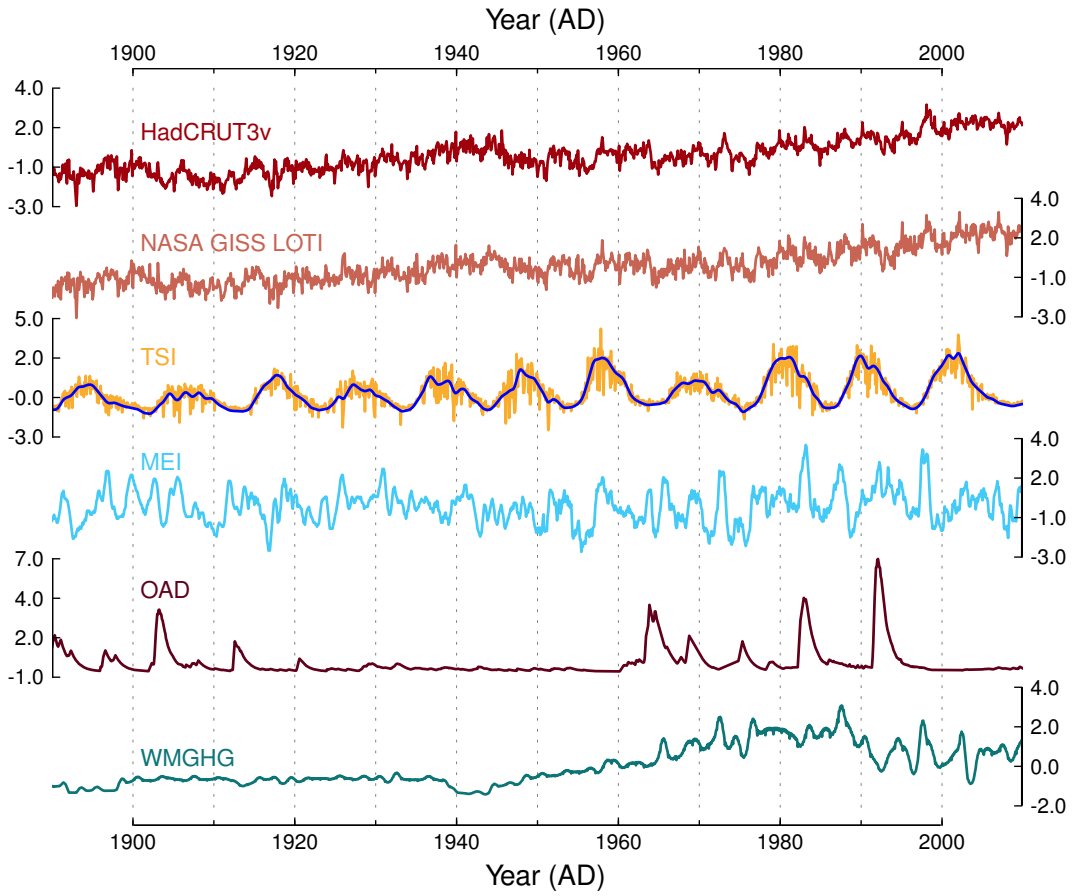

Figure 7.1. Global time series data for temperature and its climatic drivers. From top to bottom: global surface temperature from the CRU at East Anglia (HadCRUT3v, deep red); NASA GISS Land-Ocean Temperature Index (NASA GISS LOTI, light red); Wang et al.'s reconstruction of total solar irradiance (TSI, deep yellow) with the Gaussian-kernel (bandwidth of $\approx$ 11 years) filtered curve (in blue); the Multivariate ENSO Index (MEI, light blue); Sato et al.'s volcanic activity index based on Optical Aerosol Depth (OAD, brown); monthly changes in the NASA GISS's modelE global mean forcing values for Well-Mixed GreenHouse Gases (WMGHG, teal).

7.4. Datasets: global temperature drivers

Following Lean & Rind Lean and Rind (2008), we consider four critical contenders that have the potential to influence global mean temperature: (i) the ENSO, (ii) vol-

canic activity, (iii) solar irradiance, and (iv) the concentration of well- mixed greenhouse gases in the atmosphere. To characterize these phenomena, we choose: a combination of the multivariate El Niño index (Wolter and Timlin, 1998) (which extends from 1950 till present) and the index of Meyers et al. (Meyers et al., 1999) (for data prior to 1950) for the ENSO; stratospheric aerosol optical thickness measurements compiled by Sato et al. Sato and Hansen (1993) for volcanic activity; Wang et al.'s (Wang et al., 2005b) reconstruction of total solar irradiance that extends to around 300 years back; annual global mean forcing values used for well-mixed greenhouse gases as used in the NASA GISS modelE (Hansen et al., 2007) which were then interpolated to a monthly resolution. For the global temperature itself, we use two datasets: the NASA GISS Land-Ocean Temperature Index (LOTI) Hansen et al. (2010), and the surface temperature record HadCRUT3v assimilated by the Climate Research Unit (CRU) at East Anglia University (Brohan et al., 2006). (The source of the datasets are listed in Appendix C.3). We consider the 120 year period from January 1890 to December 2009 in which all the datasets are resolved on a monthly basis for the purpose of this study. Moreover, for the TSI dataset, we use a low pass filter using a Gaussian kernel of bandwidth $\approx$ 11 year, and we consider the first derivative of WMGHG instead of the almost monotonically increasing raw dataset (figure not shown). Note that, in contrast to Lean & Rind Lean and Rind (2008), we do not use the full anthropogenic forcing but just the well-mixed greenhouse gases. The final time series used in the analysis (normalized to mean zero and variance one) are shown in Fig. 7.1. For brevity, we henceforth refer to the datasets using the following labels: *CRUTA* for the HadCRUT3v data; *GISS* for NASA GISS LOTI; *TSI* for total solar irradiance; *VOLC* for the optical aerosol depth data; and *WMGHG* for the well-mixed greenhouse gas series.

7.5. Details of the numerical analysis

The analysis is divided into two parallel components corresponding to the CRUTA and GISS temperature datasets. For each of these we consider delays of up to 150 months and estimate RMD, and test the observed values for significance. Similarly, we analyze all possible combinations for the forcing datasets as well as between the forcing and the temperature datasets. All results are grouped according to the 'driven' dataset.

We do not embed the time series while constructing the recurrence matrices. Several recurrence properties are invariant under embedding and it is not essential to embed the time series. In particular, time delay embedding was not feasible for all of the datasets considered in this study and was hence avoided. The recurrence threshold was based on fixed amount of nearest neighbors, which was kept at 5% for all datasets. The qualitative nature of our results are robust to the choice of this threshold value because small changes in the threshold do not alter the qualitative features of the recurrence plot (result not shown).

The significance tests were carried out using 500 TS of the 'driven' dataset. Significant values were then considered to construct an approximate network visualization of the dependencies involving the temperature dataset as well as the various radiative forcings.

7.6. Network of dependencies around global temperature

The results of the analysis for the temperature datasets are shown in Fig. 7.2, from which several points are evident.

- TSI appears to impact the GISS dataset more than it does the CRUTA data. In fact, even on short time scales of around a few months, the impact of TSI on CRUTA is barely significant. In previous studies by Lean & Rind Lean and Rind (2008); Lean (2009) (done with the CRU data), TSI is found to influence variations in GMT on the scale of a month from their analysis based on multivariate regression.

- ENSO clearly has a sharp impact on the CRUTA series at a short delay of around 5 months (similar to Lean & Rind.). However, we also find further significant influences at delays of (approximately) 130 months for the CRU data, and at 90 months and 125–130 months for the GISS LOTI data. ENSO, however, does not have any impact on GISS temperatures at shorter time scales.

- Even though volcanic aerosols impact both temperature series on short time

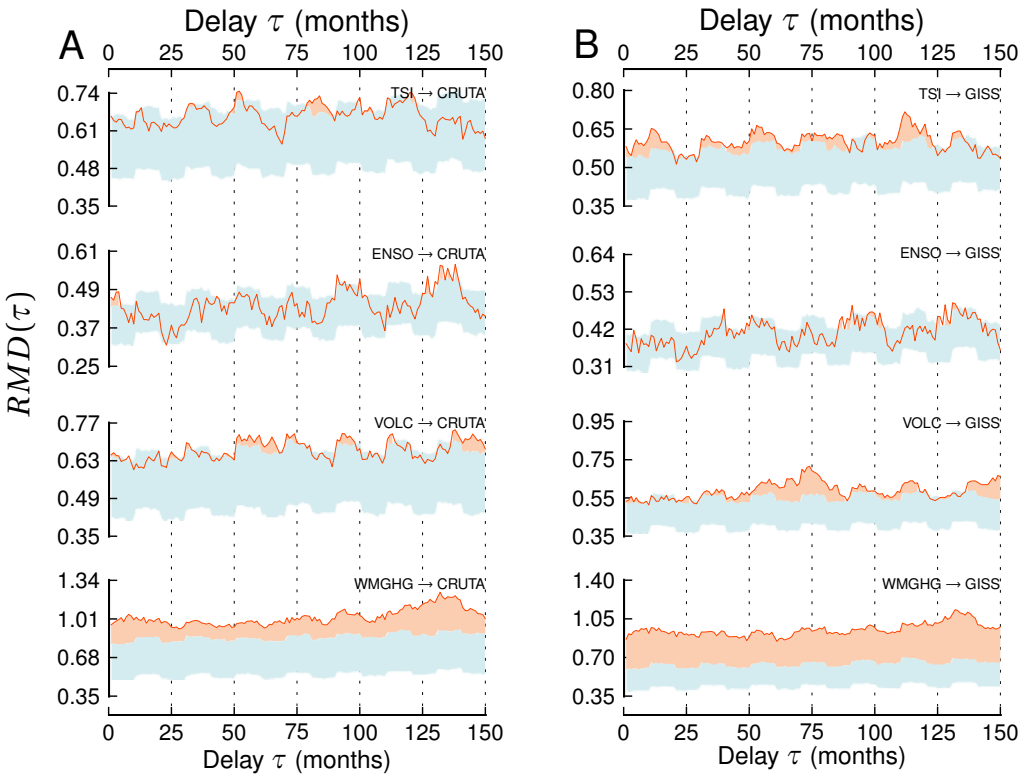

Figure 7.2. Lagged influences on GMT. A. The results for the HadCRUT3v dataset. **B.** The results for the NASA GISS LOTI dataset. The 95 % confidence band obtained from the significance test using TS is in light blue. The observed values of $RMD(\tau)$ is in dark orange. Regions where the value of $RMD(\tau)$ falls outside the confidence band are shaded in light orange. The datasets are labeled as mentioned in Sec. 7.4.

scales of around 1–10 months (Lean & Rind find 6 months), here too, there are further delayed interactions significant in the 50-75 month period and even as late as 140–150 months. This is discussed in more detail later.

- The influence of greenhouse gases on GMT is statistically significant for all values of delay considered and this influence peaks at around a delay of 130 months (approximately 11 years). This value is close to the previously considered value of 120 months in the Lean & Rind study.

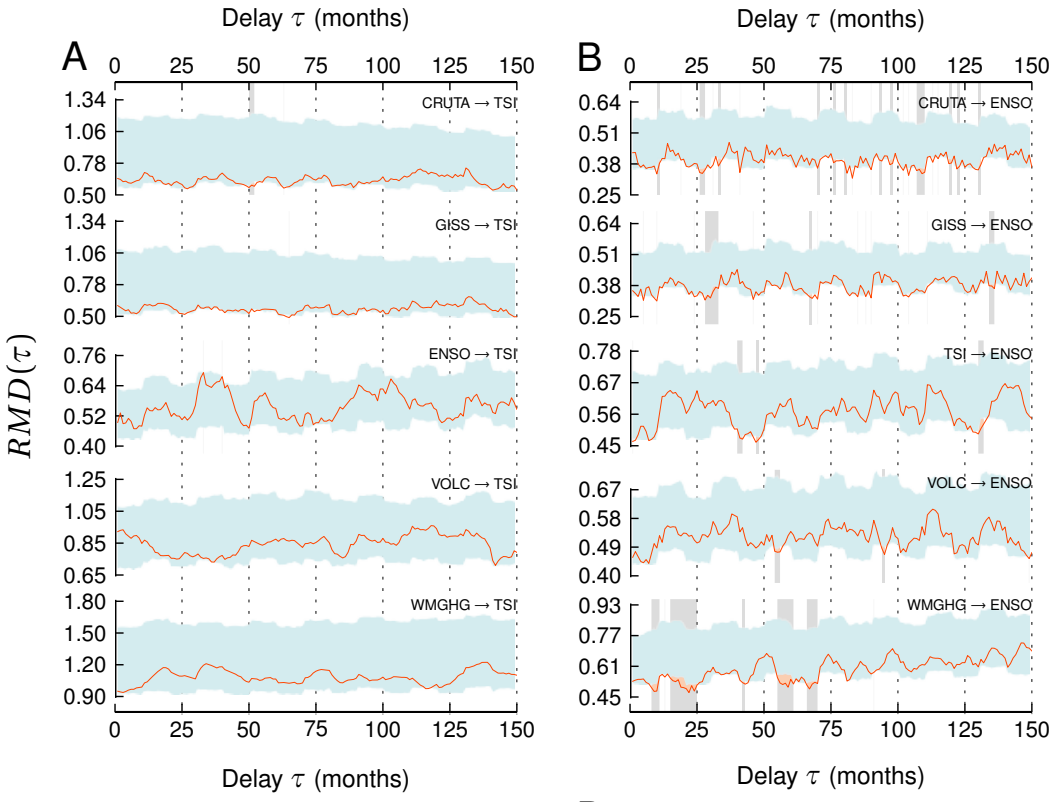

Figure 7.3. Lagged dependencies among the forcing datasets. Results for the other datasets **A.** Driving TSI. **B.** Driving ENSO. Legends and keys to the figure are same as in Fig. 7.2. The gray bars highlight the values of τ for which the observed values of RMD are statistically significant.

The results of our analysis for the rest of the data are given in Figs. 7.3 & 7.4, grouped according to the 'driven' dataset. The following points are discernible from Figs. 7.3 & 7.4.

- None of the datasets show a significant result when being tested against the hypothesis that they drive TSI (Fig. 7.3 A). This demonstrates that, provided the datasets are accurate, our approach is able to rule out physically unreasonable connections.

- Both the temperature datasets seem to drive the ENSO time series at several values of τ (Fig. 7.3B). However, a noticeable dark band around the pe-

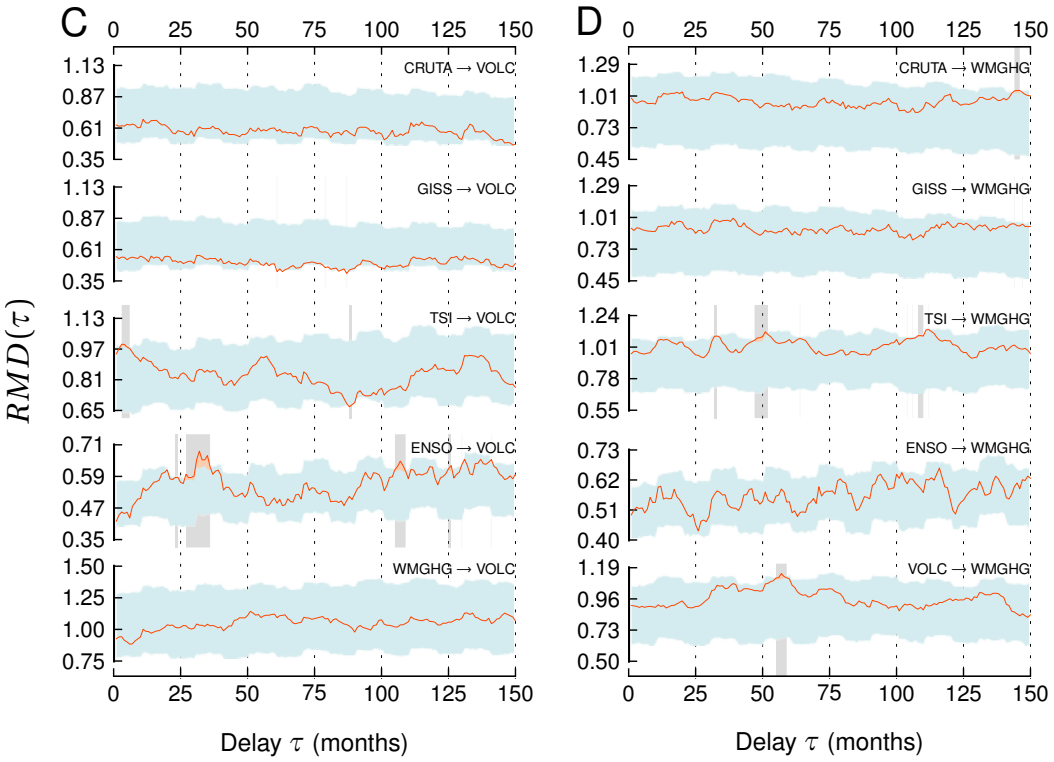

Figure 7.4. Lagged dependencies among the forcing datasets. Results for the other datasets **C.** Driving VOLC. **D.** Driving WMGHG. Legends and keys to the figure are same as in Fig. 7.2. The gray bars highlight the values of τ for which the observed values of RMD are statistically significant.

riod of 25–30 months for both CRUTA and GISS suggest a strong influence of GMT on ENSO around the quasi-biennial oscillatory period. The reverse connection from ENSO to temperature has been discussed in earlier studies (cf. (Ghil and Vautard, 1991) and (Moron et al., 1998)). However, this feedback from the temperature to ENSO around the quasi-biennial period has not been considered in much detail and thus needs more careful investigation. The greenhouse gases too seem to impact ENSO in around this period. The greenhouse gas series also impacts ENSO around 55-60 month period which might be linked to a quasi-quadrennial kind of phenomenon (Jiang et al., 1995). Similar influence of greenhouse gas emissions on the ENSO have been discussed elsewhere (Stevenson, 2012; Stevenson et al., 2012).

- Only TSI and ENSO seem to have an effect on the volcanic activity dataset (Fig. 7.4C) with the former influencing VOLC at a short time scale of around 1-5 months and the latter influencing it rather strongly at around 30 months. The idea that climatic phenomena could influence volcanic activity is controversial and debatable (e.g., Rampino et al. (1979) postulates one possible mechanism based on lithospheric stress). However, Robock (2000) challenges the more popular idea that ENSO is influenced by volcanic eruptions and suggests a converse dependency, with one possible physical mechanism being the influence of oceanic angular momentum on the earth's rotation rate at subdecadal scales (c.f. Marcus et al. Marcus et al. (1998)) which in turn might influence seismic activity on the mantle. Alternatively, the transport of volcanic aerosols in the atmosphere could be influenced as well which might lead to the results observed here (also discussed below). Our results indicate that it might be worthwhile to study this more closely in order to gain a better understanding of globally relevant climatic phenomena such as volcanoes and the ENSO.

- The greenhouse data is influenced by only the TSI and the VOLC datasets (Fig. 7.4D) at similar values of τ around 50–55 months. The impact of VOLC on the greenhouse gases most likely reflects the response of the Earth's carbon-cycle to volcanic eruptions (e.g., Jones and Cox (2001)). The observed influence of TSI on WMGHG could be an example of indirect influence of TSI on the WMGHG data via the VOLC dataset as the TSI impacts VOLC at much shorter time scales of a few months.

The above results are summarized as a network of dependencies in Fig. 7.5, where two different directed networks are visualized: one each for the two temperature datasets—CRUTA and GISS. Only statistically significant values are used to construct these networks. However, from Figs. 7.2, 7.3 & 7.4, we see that, in many cases, there are multiple values of τ at which estimated values of RMD are significant for a given pair of datasets. In such cases, for the sake of visual clarity, a single value of τ is chosen as representative of (nearly) continuous intervals of τ in which RMD values are significant. In case the interval is too wide for a single value to be representative, the interval itself is mentioned. Two notable exceptions to this are the links between WMGHG and CRUTA/GISS, and TSI and GISS. In these cases,

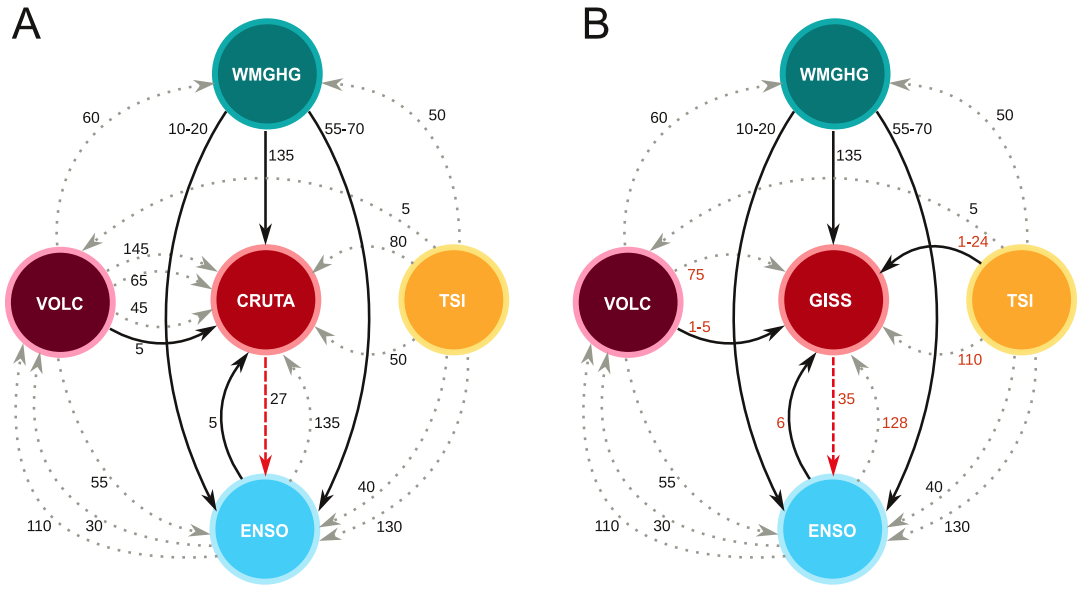

Figure 7.5. Network of dependencies surrounding GMT. A. For the HadCRUT3v temperature series. **B.** For the NASA GISS LOTI series. The arrows correspond to a statistically significant influence. The numbers beside the arrows denote the (interval of) delay(s) at which the statistically significant link is found. In **B**, the tau values that differ from **A** are given in red. The solid black arrows correspond to links that are known and/or are discussed to some extent in the literature. The gray dotted arrows correspond to links that are not easily explained. The red dashed arrow is the feedback from temperature to ENSO that is uncovered by our analysis. The datasets are labeled as given in Sec. 7.4: *CRUTA* for the HadCRUT3v data; *GISS* for NASA GISS LOTI; *TSI* for total solar irradiance; *VOLC* for the optical aerosol depth data; and *WMGHG* for the well-mixed greenhouse gas series. *Note*: For the links between WMGHG and temperature, and TSI and temperature, only a few of the statistically significant values are shown for visual clarity.

RMD values are significant for almost all values of τ considered—and hence only the presumably physically relevant values of τ are shown in Fig. 7.5 for clarity. We choose here $\tau = 135$ months for the links between WMGHG and CRUTA, and between WMGHG and GISS, as this is approximately the peak value of $RMD(\tau)$ for these pairs and is also close to the earlier estimated value of $\tau = 120$ months in

Lean and Rind (2008). We note that the networks shown in Fig. 7.5 are not constructed on the basis of a mathematical procedure. They are, in fact, intended to reveal the *qualitative* structure of the network of interactions among global temperature drivers. We can infer the following from Fig. 7.5:

- The results for the CRUTA and GISS temperature datasets are reasonably consistent for the links with ENSO, WMGHG and VOLC, but differ in the case of their respective links with TSI.

- TSI seems to drive all the remaining four datasets. However, the delays at which TSI impacts them are much larger than what is normally assumed in studies involving TSI and GMT. Our results indicate that it might be possible that the longer time scale periodicities of solar irradiance might be an important factor influencing global climatic phenomena.

- The GMT feeds back into *only* the ENSO series near the quasibiennial period (the results are slightly different between CRUTA and GISS here). This might be a crucial point to consider in both (semi-) empirical as well as theoretical modeling of GMT in terms of natural forcings.

- Apart from having a distinctively significant impact on the GMT, greenhouse gases influence the ENSO as well at multiple delays. Two distinctive periods in which WMGHG drives the ENSO are around 1–2 years and 4–6 years.

- Volcanic aerosols seem to impact GMT not only at the evident short time scales of a few months, but also at other time scales ranging from 45 months to 145 months, which is relatively hard to justify considering that volcanic aerosols do not last in the atmosphere up to 145 months. This example demonstrates the limitations of our approach and is discussed below.

One obvious limitation of our approach is the uncertainty of the statistical test itself. Thus, even when we find a link to be significant (with respect to the previously stated null hypothesis) at a 5% level of significance, we could be wrong 5 times out of 100. A second limitation, and one that might cause, e.g., the suspicious result of VOLC impacting GMT at 145 months, is that the 'volcanic activity' is represented here by the optical depth of stratospheric aerosols. This quantity, in

time periods located away from major volcanic eruptions, might be influenced by the other climatic factors. However, a third, and more critical limitation is that our approach cannot distinguish indirect/spurious links from direct ones. This means that if we have: (i) X drives A and A drives Y, or (ii) A drives both X and Y, in both cases our method would show a link between X and Y even though there is no direct connection between them. The direction of this link might then be influenced by random noise (if any) in the systems. We feel that the impact of volcanic aerosols on the GMT at delays as large as 145 months might most likely be due to spurious links that are not removed. There are other methods, e.g., multivariate transfer entropy (Runge et al., 2012), which can detect and remove spurious links. It is computationally intensive and difficult to incorporate such a principle into the current recurrence-based approach and is thus intended as a goal of future investigations.

Appendix A

Mathematical derivation of posterior proxy distributions

In the following sections, we outline the detailed derivations of the ideas described in Chapter 2. After first detailing the basic terminology that we intend to use, we state the given information in methematical terms, and then as a first step, we derive the depth-spaccing weight functions described in Section 3.2.1. Thereafter, we show how we impose the monotonicity condition of stratigraphically ordered sediment growth on the set of obtained DWFs, which result in a stratigraphically conforming post-DWFs. Using these we finally show how the posterior proxy distributions are obtained and can then be used to estimate posterior proxy means and variances at any given time point in the past.

A.1. Basic terminology

Consider Z, R and T to be the variables denoting depth, radiocarbon age, and calendar age respectively. The variable X is used to denote the unknown proxy. The set of observations correspond to the radiocarbon age-depth measurements, and the proxy-depth measurements. In other words, we have

- a time point of interest t which runs over a fine, regular grid. The latter is, in fact, the *precise* calendar age interval for which we wish to estimate the proxy probability distribution.

- a known *calibration curve* r_t defined for all these values of t, where t denotes radiocarbon age, and the associated Gaussian error estimates σ_t^{C} (the

superscript "C" denotes "calibration").

- a (typically small) number M of ^{14}C *dating points* (z_k^r, r_k), along with estimates σ_k^{R} of the standard deviation of the individual Gaussian ^{14}C dating errors.

- a (typically large) number $N \gg M$ of *proxy measurement points* (z_j^x, x_j) for another (typically different) set of depths z_j^x, where x_j denotes proxy values, along with estimates σ_j^{X} of the standard deviation of the individual Gaussian proxy measurements errors.

Thus, we have the following Gaussian conditional probability distributions:

- The calibration curve specifies the conditional density of r given t as $P(r|t) \propto \exp[-(r-r_t)^2/2(\sigma_t^{\mathrm{C}})^2]$.

- The ^{14}C age data specify the conditional density of r given $Z = z_k^r$ as $P(r|z_k^r) \propto \exp[-(r-r_k)^2/2(\sigma_k^{\mathrm{R}})^2]$.

- The proxy data specify the conditional density of x given $Z = z_j^x$ as $P(x|z_j^x) \propto \exp[-(x-x_j)^2/2(\sigma_j^{\mathrm{X}})^2]$.

We only state proportionalities ($\propto$) here, taking care of proper normalization only in the end.

A.2. Estimating Depth-spanning Weight Functions

To answer the question: *which among the proxy measurement depths are more likely than others to correspond to a given true age,* we apply the law of total probability and Bayes' Theorem to combine the calibration curve distribution $P(r|t)$ and the output of the RM age model (which gives $P(r|z_j^x)$ from the measured data $P(r|z_k^r)$):

$$P(z_j^x|t) = \int \mathrm{d}r\, P(z_j^x|r)\, P(r|t) = \int \mathrm{d}r\, \frac{P(z_j^x)\, P(r|z_j^x)}{P(r)}\, P(r|t) \tag{A.1}$$

Assuming all ages and depths are equally likely *a priori*, we use the (flat) prior distributions $P(r) \propto P(z_j^x) \propto 1$ and see that $P(z_j^x|t)$ is proportional to the weight

$$P(z_j^x|t) \propto \int \mathrm{d}r\, P(r|z_j^x)\, P(r|t) \qquad =: w_t(z_j^x). \tag{A.2}$$

where the weight $w_t(z_j^x)$ is the *Depth-spanning Weight Function*. Since the RM age model returns the means r_j and standard deviations σ_j^{R} of the radiocarbon age R for the depths z_j^x, we can substitute the relevant Gaussian functions in equation A.2 to get

$$w_t(z_j^x) \propto \int \mathrm{d}r\, \exp\left(\frac{-(r-r_j)^2}{2(\sigma_j^{\mathrm{R}})^2}\right) \exp\left(\frac{-(r-r_t)^2}{2(\sigma_t^{\mathrm{C}})^2}\right). \tag{A.3}$$

The first term in equation A.3 is the output of the RM age model while the second term represents the calibration curve. *Note*: For the RM age model, we use the MoTaBaR regression method from Heitzig (2013) with order parameter $p = 2$ and a data-driven choice of prior distributions.

A.3. Imposing monotonic growth using force-based relaxation dynamics

Since it is difficult to define unambiguously a monotonic condition relation in T for the set of DWFs defined in equation A.3, we choose to work with the cumulative distributions while imposing the constraint of monotonic growth of the paleoarchive. The first step is to convert the initially estimated DWFs, $w_t^i(z_j^x)$ into corresponding cumulative distributions—

$$W_t^i(z_j^x) = \sum_{l=1}^{j} w_t^i(z_l^x). \tag{A.4}$$

The initial cumulative distributions $W_t^i(z_j^x)$ are by construction weakly monotonic (non-decreasing) over the depths z_j^x. Our goal is to find the a final set of CDWFs, $W_t^f(z_j^x)$, that are weakly monotonic (non-increasing) over t as well.

The final CDWFs are those that satisfy the above condition of monotonicity and are at a minimal distance from the initially estimated $W_t^i(z_j^x)$. Ideally, we can estimate $W_t^f(z_j^x)$ by minimizing the functional $\sum_j \sum_t [W_t^f(z_j^x) - W_t^i(z_j^x)]^2$. However, this is not straightforward and computationally expensive and we thus use a force-based relaxation dynamics under reasonable assumptions to estimate $W_t^f(z_j^x)$.

We introduce an artificial 'time' variable τ (not related to the calendar age T) such that the cumulative distribution function that is given by $W_t^f(z_j^x, \tau \to \infty)$. We choose the initial condition at the starting point of the artificial time $\tau = 0$ as

$$W_t^f(z_j^x, 0) = \frac{1}{2}\left[\min_{t' \le t} W_{t'}^i(z_j^x) + \max_{t' \ge t} W_{t'}^i(z_j^x)\right]. \tag{A.5}$$

Such an initial condition ensures that $W_t^f(z_j^x, 0)$ is monotonic from the start. The next step is to 'drag' this function towards the non-monotonic $W_{t'}^i(z_j^x)$ as far as monotonicity permits. For this we first evaluate

$$\delta(\tau) = W_{t'}^i(z_j^x) - W_t^f(z_j^x, \tau) \tag{A.6}$$

which is the distance of the monotonic function from the original function after τ time steps. The evolution of $W_t^f(z_j^x, \tau)$ over τ is formulated as:

$$\frac{\mathrm{d}}{\mathrm{d}\tau}[W_t^f(z_j^x,\tau)] = \begin{cases} \delta(\tau) \begin{cases} \delta(\tau) > 0 \;\&\; W_t^f(z_j^x,\tau) < \min\left(W_t^f(z_{j+1}^x,\tau), W_{t-1}^f(z_j^x,\tau)\right) \\ \delta(\tau) < 0 \;\&\; W_t^f(z_j^x,\tau) > \max\left(W_t^f(z_{j-1}^x,\tau), W_{t+1}^f(z_j^x,\tau)\right) \end{cases} \\ 0, \quad \text{otherwise.} \end{cases} \tag{A.7}$$

Integrating over τ, we obtain the monotonic W^f in the limit $\tau \to \infty$. This particular approach of estimating the monotonic W^f functions has the advantage that even when the process is stopped before convergence due to a too slow rate of convergence, the result is still monotonic and at least as close to $W_{t'}^i(z_j^x)$ as the initial guess $W_t^f(z_j^x, 0)$. Preliminary tests (results not shown) with example data sets suggest that a step size of $\mathrm{d}\tau = 10/s$ with $s = 1000$ steps might lead to sufficient results.

A.4. Posterior proxy distributions, their means and variances

After estimating $W_t^f(z_j^x)$, we obtain the corresponding $w_t^f(z_j^x)$ by the first-order difference of $W_t^f(z_j^x)$ along the z_j^x axis, (the reverse of equation A.4), i.e.,

$$w_t^f(z_j^x) = W_t^f(z_j^x) - W_t^f(z_{j-1}^x),\ j = 2,\ldots,N \tag{A.8}$$

with the first value set as $w_t^f(z_1^x) = W_t^f(z_1^x)$.

We now use these as weights to estimate the posterior proxy distributions $P(x|t)$. For this, we again apply the law of total probability and approximate the integral over depth by a Riemann sum using the proxy measurement depths:

$$P(x|t) = \int \mathrm{d}z\, P(x|z)\, P(z|t) \approx \sum_{j=1}^{N} b_j P(x|z_j^x)\, P(z_j^x|t), \tag{A.9}$$

where b_j is the width of the depth interval represented by z_j^x:

$$b_j = \frac{1}{2}\begin{cases} z_2^x - z_1^x & b_j = 1 \\ z_{j+1}^x - z_{j-1}^x & 1 < b_j < N \\ z_N^x - z_{N-1}^x & b_j = N\,. \end{cases} \tag{A.10}$$

Finally, we plug in the final DWFs, $w_t^f(z_j^x)$, taking care of a correct normalization, and find that the proxy probability density at t is simply a weighted mean of the densities corresponding to the individual proxy measurements, i.e., a Gaussian mixture:

$$P(x|t) \approx \frac{\sum_{j=1}^{N} b_j\, w_t^f(z_j^x)\, P(x|z_j^x)}{\sum_{j=1}^{N} b_j\, w_t^f(z_j^x)}, \tag{A.11}$$

where $P(x|z_j^x)$ is the Gaussian

$$P(x|z_j^x) = \frac{\exp[-(x-x_j)^2/2(\sigma_j^{\mathrm{X}})^2]}{\sqrt{2\pi}\sigma_j^{\mathrm{X}}}. \tag{A.12}$$

Once the probability distribution of the proxy at a given t is known (as given by equations A.11 and A.12) estimates of the mean and median, and respective estimates of uncertainty such as the variance and the inter-quartile range can be calculated. E.g., the mean $\bar{x}(t)$ of $P(x|t)$ is estimated as

$$\bar{x}(t) = \int \mathrm{d}x\, x\, P(x|t) \approx \frac{\sum_{j=1}^{N} b_j x_j w_t^f(z_j^x)}{\sum_{j=1}^{M} b_j w_t^f(z_j^x)}, \tag{A.13}$$

using which the variance $\sigma^2(t)$ of $P(x|t)$ can be computed as:

$$\begin{aligned} \sigma^2(t) &= \int \mathrm{d}x\, (x - \bar{x}(t))^2\, P(x|t) \\ &\approx \frac{\sum_{j=1}^{N} b_j w_t^f(z_j^x) \left[(x_j - \bar{x}(t))^2 + (\sigma_j^{\mathrm{X}})^2\right]}{\sum_{j=1}^{N} b_j w_t^f(z_j^x)} \end{aligned} \tag{A.14}$$

Note however, that because $P(x|t)$ is not necessarily Gaussian, $\sigma^2(t)$ might not give reliable confidence bounds, which is why we use the exact quantiles of $P(x|t)$ instead to construct confidence bounds for X.

Appendix B

Mathematical derivation of the role of joint distributions

B.1. Expectation value and variance over time

In this part, we are concerned with estimating the expectation value of the proxy over a period of time given the (marginal) posterior proxy densities $P(x|t)$ at each time point t in the past. In order to estimate the expectation value $E[X]$ of the proxy X over time, we have to first arrive at a sensible probability density that expresses the total probability of having a proxy value $X = x$ over the time period under consideration. This turns out to be,

$$P(x) = \frac{\int \mathrm{d}t\, P(x|t)}{\int \mathrm{d}t}. \tag{B.1}$$

Once we have a well-defined probability distribution $P(x)$ at hand, we can now estimate the expectation value according to the definition,

$$\mu_X = E[X] = \int \mathrm{d}x\, x\, P(x). \tag{B.2}$$

Using the expression for the probability density from Equation B.1 and rearranging the terms we find that the expectation value is expressible as the time average of the expectation values of the proxy at each point of time $\bar{x}(t)$. Formally

$$E[X] = \frac{\int \mathrm{d}t \left[\int \mathrm{d}x \, x \, P(x) \right]}{\int \mathrm{d}t} \tag{B.3}$$

$$= \langle \bar{x}(t) \rangle, \tag{B.4}$$

where $\int \mathrm{d}x \, x \, P(x) = \bar{x}(t)$ from equation A.12 and $\langle \cdot \rangle$ denotes the time average.

After estimating the expected value, we are now in a position to derive the expression for the variance of the proxy values over an extended period of time with respect to the expectation μ_X derived above. To do this, we use the standard definition of the variance,

$$\sigma_X^2 = \int \mathrm{d}x \, (x - \mu_X)^2 \, P(x), \tag{B.5}$$

from which, by using the expression in Equation B.1 we get,

$$\sigma_X^2 = \int \mathrm{d}x \, x^2 \, P(x) - \mu_X^2 \tag{B.6}$$

$$= \frac{\int \mathrm{d}t \left[\int \mathrm{d}x \, x^2 P(x|t) \right]}{\int \mathrm{d}t} - \mu_X^2. \tag{B.7}$$

Now, from Appendix A.4, we find that $\int \mathrm{d}x \, x^2 P(x|t) = \sigma^2(t) + \bar{x}(t)^2$, which results in the expression for the variance over time as,

$$\sigma_X^2 = < \sigma^2(t) + \bar{x}(t)^2 > - \, \mu_X^2. \tag{B.8}$$

Equations B.4 & B.8 link the statistical quantities over a period of time to the corresponding statistical quantities at individual time points. Note that we had no reason to assume that this should have been the case, and these results are, in fact, representative of the fact that the instantaneous characteristics of the proxy distributions are linked to its long-term characteristics.

B.2. Autocovariance and crosscovariance

Other crucial quantities of interest in time series analyses are the autocorrelation, crosscorrelation and the power spectrum of a given dataset. In our case, since we have a series of probability densities rather than a time series on our hands, we have to take care to derive from first principles these quantities in order to avoid mathematical oversights.

Let us first consider the autocorrelation of the series of probability densities represented as $P(x|t)$. To avoid ambiguity and to adopt a shorthand notation for these densities, let us denote the random variable $X = x$ given $T = t$ as X_t and its corresponding probability density $P(x|t)$ as $P_t(x_t)$. Now, since it is not guaranteed that the process $X = (X_t)$ is not stationary, we define the two time-dependent means for X_t and $X_{t+\tau}$ as the expectation values analogous to equation B.3

$$E[X_t] = \frac{\int \mathrm{d}t \left[\int \mathrm{d}x_t \, x_t \, P_t(x_t) \right]}{\int \mathrm{d}t} \tag{B.9}$$

$$E[X_{t+\tau}] = \frac{\int \mathrm{d}t \left[\int \mathrm{d}x_{t+\tau} \, x_{t+\tau} \, P_{t+\tau}(x_{t+\tau}) \right]}{\int \mathrm{d}t}. \tag{B.10}$$

Then, the autocovariance $\varrho^X(\tau)$ at lag τ is given by

$$\varrho^X(\tau) = E[X_t, X_{t+\tau}] - E[X_t] \, E[X_{t+\tau}], \tag{B.11}$$

where the joint expectation value $E[X_t, X_{t+\tau}]$ involves the total joint probability $P(x_t, x_{t+\tau})$ which we can write down as

$$P(x_t, x_{t+\tau}) = \frac{\int \mathrm{d}t \, P(x_t, x_{t+\tau}|t)}{\int \mathrm{d}t}. \tag{B.12}$$

The term on the right-hand side of equation B.12 is unknown to us and thus the autocovariance cannot be estimated without approximations and assumptions about the data.

B. Mathematical derivation of the role of joint distributions

The lagged crosscovariance $\varrho^{XY}(\tau)$ between two different proxy records can also be written down in terms of the expectation values as

$$\varrho^{XY}(\tau) = E[X_t, Y_{t+\tau}] - E[X_t]\,E[Y_{t+\tau}] \tag{B.13}$$

$$= \iint \mathrm{d}x_t \mathrm{d}y_{t+\tau}\, x_t y_{t+\tau}\, P_{X_t,Y_{t+\tau}}(x_t, y_{t+\tau}) - E[X_t]\,E[Y_{t+\tau}], \tag{B.14}$$

which we find is critical dependent on the total joint probability term

$$P_{X_t,Y_{t+\tau}}(x_t, y_{t+\tau}) = \int \mathrm{d}t \frac{P_{X_t,Y_{t+\tau}}(x_t, y_{t+\tau}|t)}{\int \mathrm{d}t}. \tag{B.15}$$

However, in contrast to the case of autocivariance, in this case, we can assume *conditional independence* of the two proxies such that

$$P_{X_t,Y_{t+\tau}}(x_t, y_{t+\tau}|t) = P_{X_t}(x_t|t)\, P_{Y_{t+\tau}}(y_{t+\tau}|t). \tag{B.16}$$

Now, combining the equations B.14, B.15 and B.16, we get the crosscovariance as

$$\varrho^{XY}(\tau) = \iint \mathrm{d}x_t \mathrm{d}y_{t+\tau}\, x_t y_{t+\tau} \frac{\int \mathrm{d}t P_{X_t}(x_t|t)\, P_{Y_{t+\tau}}(y_{t+\tau}|t)}{\int \mathrm{d}t} - E[X_t]\,E[Y_{t+\tau}], \tag{B.17}$$

which can be simply rearranged as

$$\varrho^{XY}(\tau) = \langle \bar{x}(t)\bar{y}(t+\tau)\rangle - \langle \bar{x}(t)\rangle\langle \bar{y}(t+\tau)\rangle. \tag{B.18}$$

B.3. Posterior joint probability for proxy at two time points

This section focusses on the derivation of the posterior joint probability density for the proxy values at two different time points under consideration. As with the derivation of the posterior probability densities derived in Appendix A, we shall consider the more general case of an archive dated with radiocarbon dating, and which thus needs calibration, as this helps to illustrates all the sources of covariance that are incumbent in the joint probability estimation. For this, let us assume that X_i, $i = 1,2$, are the random variable representing the proxy values at the two time points T_i. Next, let Z_i and R_i represent the depth and radiocarbon axes related to the random variables X_i. Note that, for a single archive and a single proxy, the depth and radiocarbon measurements are the same, but we still denote them with two separate variables in the following derivation for the sake of clarity and in the anticipation of necessary double integrals over the depth and radiocarbon axes.

For the sake of clarity, the joint probability density function is for the proxy values at T_1 and T_2 is represented as $P(x_1, x_2|t_1, t_2)$. Using this nomenclature, we now want to answer the question: *what is the probability that the proxy X_1 at T_1 had the value x_1 and the proxy X_2 at T_2 had the value x_2?* Formally, this can be written down as,

$$P(x_1, x_2|t_1, t_2) = \iint \mathrm{d}z_1 \mathrm{d}z_2 \, P(x_1, x_2|z_1, z_2) \, P(z_1, z_2|t_1, t_2). \tag{B.19}$$

Next, consider the first term on the right-hand side of Equation B.19, which effectively requires us to provide the joint probability of making two proxy observations x_1 and x_2 at two corresponding depths z_1 and z_2. Now, we assume the proxy measurements along the depth of the archive are made with instruments with negligible systematic errors, and thus, we can assume that each proxy measurement at a given depth is independent of any other measurement at other depths, i.e., knowing the proxy measurement at particular depth does not change or effect our beliefs about a subsequent proxy measurement at any other depth. In probabilistic

terms, this implies that $P(x_1, x_2|z_1, z_2) = P(x_1|z_1)P(x_1|z_1)$, i.e., the proxy measurements along depth are conditionally independent. Now, considering $z_{j_i}^{X_i}$, $i = 1, 2$, as the depths at which the proxy was was measured to get the proxy value X_i at the times T_i, we can estimate the integral given by Equation B.19 using the Riemann sum:

$$P(x_1, x_2|t_1, t_2) \approx \sum_{j_1=1}^{N} \sum_{j_2=1}^{N} b_{j_1} b_{j_2} \, P(x_1|z_{j_1}^{x_1}) \, P(x_1|z_{j_2}^{x_2}) \, P(z_{j_1}^{x_1}, z_{j_2}^{x_2}|t_1, t_2), \tag{B.20}$$

where, analogous to Equation A.10, b_{j_i} are the widths of the depth intervals represented by $z_{j_i}^{X_i}$:

$$b_{j_i} = \frac{1}{2} \begin{cases} z_2^{x_i} - z_1^{x_i} & b_{j_i} = 1 \\ z_{j_i+1}^{x_i} - z_{j_i-1}^{x_i} & 1 < b_{j_i} < N \\ z_N^{x_i} - z_{N-1}^{x_i} & b_{j_i} = N \,. \end{cases} \tag{B.21}$$

Now, let us focus on the second term $P(z_{j_1}^{x_1}, z_{j_2}^{x_2}|t_1, t_2)$ on the right-hand side of Equation B.20, which requires us to estimate the joint probability of the most likely depths z_1 and z_2 given two calendar ages t_1 and t_2. We can express this probability as a double integral over the radiocarbon measurement axes and then use Bayes' Theorem to obtain a proportional relationship of this probability in terms of quantities that are potentially accessible for us:

$$P(z_{j_1}^{x_1}, z_{j_2}^{x_2}|t_1, t_2) = \iint \mathrm{d}r_1 \mathrm{d}r_2 \, P(z_{j_1}^{x_1}, z_{j_2}^{x_2}|r_1, r_2) \, P(r_1, r_2|t_1, t_2) \tag{B.22}$$

$$= \iint \mathrm{d}r_1 \mathrm{d}r_2 \, \frac{P(r_1, r_2|z_{j_1}^{x_1}, z_{j_2}^{x_2}) \, P(z_{j_1}^{x_1}, z_{j_2}^{x_2})}{P(r_1, r_2)} \, P(r_1, r_2|t_1, t_2). \tag{B.23}$$

Now, assuming that all combinations of ages, and all combinations of depths are equally likely *a priori,* we use (flat) prior distributions $P(z_{j_1}^{x_1}, z_{j_2}^{x_2}) \propto P(r_1, r_2) \propto 1$,

we find, from the expression in Equation B.23, that $P(z_{j_1}^{x_1}, z_{j_2}^{x_2} | t_1, t_2)$ is proportional to the weight

$$P(z_{j_1}^{x_1}, z_{j_2}^{x_2} | t_1, t_2) \propto \iint \mathrm{d}r_1 \mathrm{d}r_2 \, P(r_1, r_2 | z_{j_1}^{x_1}, z_{j_2}^{x_2}) \, P(r_1, r_2 | t_1, t_2) \tag{B.24}$$

$$=: w_{t_1 t_2}(z_{j_1}^{x_1}, z_{j_2}^{x_2}). \tag{B.25}$$

Finally, plugging in the weight function obtained in equation B.25 into the Riemann sum defined for the final proxy joint density in equation B.20, we get the joint proxy probability at times t_1 and t_2 as the weighted mean of the conditionally independent densitites corresponding to individual proxy measuements

$$P(x_1, x_2 | t_1, t_2) \approx \frac{\sum_{j_1=1}^{N} \sum_{j_2=1}^{N} b_{j_1} b_{j_2} \, w_{t_1 t_2}(z_{j_1}^{x_1}, z_{j_2}^{x_2}) \, P(x_1 | z_{j_1}^{x_1}) \, P(x_1 | z_{j_2}^{x_2})}{\sum_{j_1=1}^{N} \sum_{j_2=1}^{N} b_{j_1} b_{j_2} \, w_{t_1 t_2}(z_{j_1}^{x_1}, z_{j_2}^{x_2})}. \tag{B.26}$$

Note that the primary source of covariance in the final joint proxy probability in the above equation arises from the weigth function $w_{t_1 t_2}(z_{j_1}^{x_1}, z_{j_2}^{x_2})$ which in turn has two main components as can be seen from equation B.24. The first term $P(r_1, r_2 | z_{j_1}^{x_1}, z_{j_2}^{x_2})$ reflects the growth relation of the sediment core and thus implies that the growth of the sediment core strongly correlates the uncertainties between any given pair of depth layers. The second term $P(r_1, r_2 | t_1, t_2)$ is the calibration term which implies that uncertainties in determining the precise calendar ages for a given pair of radiocarbon ages are correlated.

Appendix C

Data sets used

We tabulate here the datasets that were used in the dissertation. Although we present here the age–depth measurements that were used, due to the long length of the proxy–depth measurements (typically of around a 1000 points long), we do not tabulate them here. However, for the Wanxiang and Dayu caves, they are available as supplementary information to their respective publications. They are also available for download in ASCII format from:
`http://www.ncdc.noaa.gov/data-access/paleoclimatology-data/datasets/speleothem`.
The data for the Wah Shikar cave were obtained from the authors by request.

C.1. Data from Wanxiang, Dayu and Wah Shikar

Wanxiang cave

The Wanxiang cave is located in in Wudu County, Gansu Province, China between the Qinghai-Tibetan Plateau and the Chinese Loess Plateau at 33°19′N, 105°00′E and at an elevation of 1200 m. We use the 118 mm long WX42B stalagmite from the Wanxiang cave reported in Zhang et al. (2008). The stalagmite grew roughly from 190 to 2003 AD without any hiatus. In total there were a total of 703 δ^{18}O measurements made along its depth along with 19 U/Th age measurements. The U/Th age measurements are reported in Table C.1.

Table C.1. U/Th age measurements from Wanxiang cave

Sample	Depth (mm)	U/Th age (yBP)	Error (± yBP)
WX42B-2-2	1.0	13	0.5
WX42B-0	4.0	33	0.5
WX42B-2-2	8.1	60	0.5
WX42B-1	13.2	106	0.5
WX42B-2-3	16.0	135	0.5
WX42B-2	22.7	294	0.5
WX42B2-4	27.8	401	1.0
WX42B-3	35.5	579	0.5
WX42B2-5	39.3	670	1.0
WX42B-4	45.5	788	1.0
WX42B2-6	50.4	899	1.0
WX42B-5	57.0	1032	1.5
WX42B2-7	60.0	1106	1.5
WX42B-7	67.0	1251	1.0
WX42B-8	76.5	1432	1.5
WX42B-9	86.0	1514	2.5
WX42B2-9	92.0	1576	1.5
WX42B-10	99.0	1655	2.0
WX42B-11	115.5	1811	2.0

Dayu cave

The Dayu cave is located on the southern side of the Qinling mountains in central China at $33^{o}08'$N, $106^{o}18'$E. The δ^{18}O signal record in the stalagmite DY-1 as reported in Tan et al. (2009) was used in our analysis. This stalgmite grew from roughly from 1271 AD to 1969 AD without interruptions. The dating of DY-1 involves a set of 8 U/Th ages with an approaximate mean error of 1.7 years. The set of ages are tabulated in Table C.2.

Table C.2. U/Th age measurements from Dayu cave

Depth (mm)	U/Th age (yBP)	Error (± yBP)
2.50	-19	1.5
27.75	114	1.0
35.25	163	1.0
54.00	290	1.5
72.50	391	2.0
78.75	438	1.5
96.75	672	2.5
98.00	679	4.0

Wah Shikar cave

The Wah Shikar cave is located in northeastern India around 30 km from the town of Shillong at 25°15′N, 91°52′E and at an elevation of 1290 m. We use the δ^{18}O signal recorded in the stalagmite WS-B as reported in Sinha et al. (2011). The stalagmite itself is around 93 mm long and grew without any hiatuses from 1360 AD to 1750 AD. The average error of the 12 U/Th ages measured along the depth of the stalagmite is around 65 years. The set of U/Th age measurements are tabulated in Table C.3.

C.2. Data from Lonar lake, central India

The Lonar lake is a circular crater located in central India at 19°58′N, 76°30′E. The lake itself is annually fed by precipitation from the Indian summer monsoon and the runoff from three major rivers in the region: Dhara, Sitanahani and Ramgaya. The diamter of the crater is approxiamtely 1.8 km and it has an average depth of around 135 m. In our analysis in Chapter 4, we use the final set of ^{14}C age measurements made along the composite core depth of 10.04 m combined from two nearby sediments cores separated by 50 cm (Anoop et al., 2013; Prasad et al.,

Table C.3. U/Th age measurements from Wah Shikar cave

Depth (mm)	U/Th age (yBP)	Error (± yBP)
13	200	60
14	144	25
19	224	35
28	255	35
30	470	250
41	385	30
43	440	120
52	460	50
65	467	15
66	473	15
88	510	100
90	590	55

2014).

This consists of 17 pre-1950 ages, L19–L1, and two ^{14}C ages after 1950, L21 and L20a. The entire set of age measurements finally used for our analysis are given in the following section.

^{14}C measurements from Lonar

The ^{14}C age depth measurements used in the analysis (as shown in Fig. 4.3) are given in Table C.4. The data reported in Table C.4 lists only those measurement samples from Anoop et al. (2013) which were finally used in their analysis.

Table C.4. ^{14}C age measurements from Lonar Lake, central India

Sample	Depth (cm)	^{14}C age (yBP)	Error (± yBP)
L21[1]	0	40	-
L20a[2]	20	-2902	0.01
L19	163.5	564	30
L17	266	1105	30
L15o	266.5	1075	30
L14	267.5	1100	30
L13	383.5	1840	35
L12	482	2315	35
L11	511.5	2680	35
L10	612	3470	35
L9	778	4185	35
L8	820	4600	60
L7	870	7420	40
L6	870.5	7460	90
L5	872	7410	100
L4	882.5	8880	60
L3	899	8990	80
L2	902	9740	50
L1	904	9570	100

[1] Point at surface fixed at 2007 AD (≈ 40 ^{14}C yBP)
[2] 143.51 ± 0.0043 pMC

Note on Al area as a surface erosion proxy

Figure 4.3E reveals that the Al total counts are extremely low as compared to those of Ca, for instance. It is advisable to treat such data with caution, as XRF measurements give only a qualitative overview of the elemental abundances. However, as shown in the earlier study by Prasad et al. (2014), the Al counts are strongly correlated with the Si, Ti and K counts obtained from the same core (shown in Fig. C.1). This determined our choice of taking Al as a representative proxy for the whole

core. Even though the actual magnitude of the Al XRF estimates might be low, its variations mimic the variations of the Si, Ti and K counts, which have relatively higher magnitudes of XRF estimates. We note that this fact, combined with the study of Basavaiah et al. (2014), validates the use of the Al counts as a proxy for the surface erosion in the Lonar catchment area.

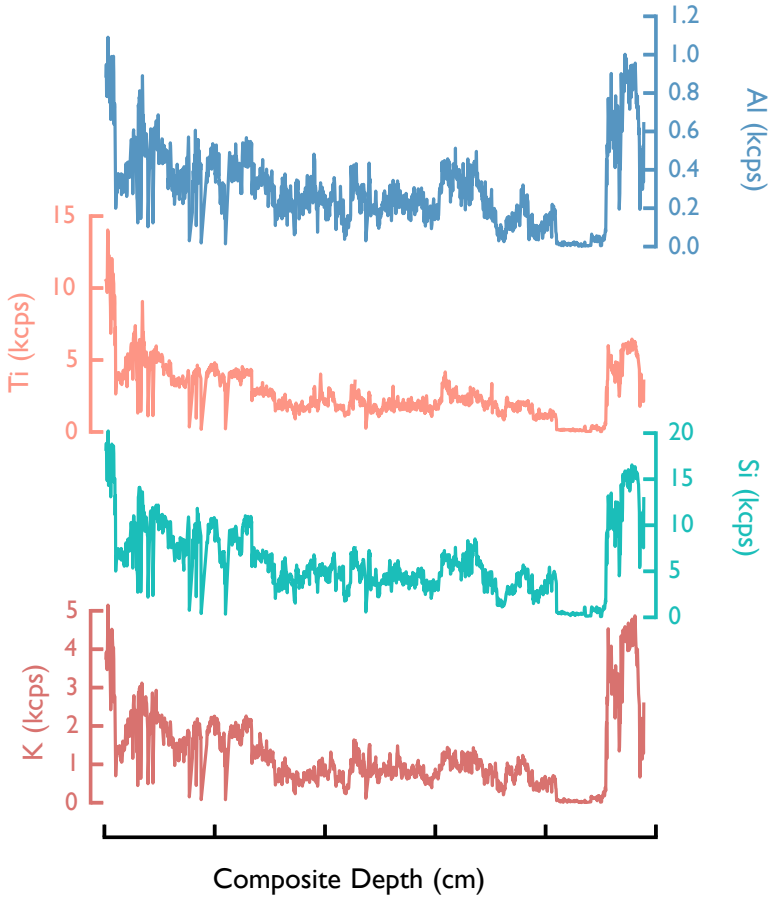

Figure C.1. Al counts as a representative proxy for the Lonar lake. Al counts (top panel, in blue) are highly correlated with the Ti (orange), Si (green) and K (red) counts obtained from XRF scanning of the Lonar lake sediment. This correlation with the other elements, combined with the fact that the Al counts arise due to the catchment erosion, forms the basis for choosing it as a representative proxy for this analysis. This choice also helps us to illustrate the impacts of proxy–depth variability on the final proxy estimate uncertainties.

C.3. Global time series datasets for the past 120 years

The datasets used in this analysis are available (as of December 1, 2012) for download at the web addresses listed below.

- Multivariate ENSO Index
  - From 1950 to present
    `http://www.esrl.noaa.gov/psd/enso/mei/table.html`
  - Prior to 1950
    `ftp://www.coaps.fsu.edu/pub/JMA_SST_Index/`
- Stratospheric optical aerosol depth
  `http://data.giss.nasa.gov/modelforce/strataer/tau_line.txt`
- Total Solar Irradiance
  `http://lasp.colorado.edu/sorce/tsi_data/TSI_TIM_Reconstruction.txt`
- WMGHG forcing
  `http://data.giss.nasa.gov/modelforce/RadF.txt`
- NASA GISS LOTI
  `http://data.giss.nasa.gov/gistemp/tabledata_v3/GLB.Ts+dSST.txt`
- HadCRUT3v global mean temperature
  `http://www.cru.uea.ac.uk/cru/data/temperature/hadcrut3v.zip`

Appendix D

Mathematical models

Here we give the equations for mathematical models and synthetic examples used in the dissertation.

- Growth of a synthetic archive with a sinusoidal proxy, as discussed in Sec. 4.1.1. The calendar age t_i of a the i^{th} depth layer is

$$t_i = t_{i-1} + r_{i-1} z_{i-1}, \tag{D.1}$$

  where the growth rate r_i of the i^{th} depth layer is chosen as $r_i \sim \mathcal{N}(r_{i-1}, \sigma_r)$. In case a negative or zero r_i is drawn, we repeat the draw until we get a non-zero r_i. We choose $r_0 = 20\,\text{yr/cm}$ and $\sigma_r = 7\,\text{yr/cm}$.The sinusoidal proxy x_t with two superimposing frequencies are given by

$$x_t = \sin(2\pi\omega_1 t) + \sin(2\pi\omega_2 t), \tag{D.2}$$

  with $(\omega_1, \omega_2) = (0.0005, 0.0025)$ for the U/Th-dated archive (Sec. 4.1.1) and $(\omega_1, \omega_2) = (0.001, 0.005)$ for the ^{14}C-dated archive (Sec. 4.1.1).

- The chaotic Rössler system,

$$\dot{x} = -y - z \tag{D.3}$$

$$\dot{y} = x + ay \tag{D.4}$$

$$\dot{z} = b + z(x - c). \tag{D.5}$$

with the parameters $(a, b, c) = (0.2, 0.2, 5.7)$ (Sec. 6.3).

- Brownian motion,

$$x(t + dt) = x(t) + \mathcal{N}(0, \delta^2 dt), \tag{D.6}$$

where δ is the standard deviation of the normal distribution and dt is the time step of estimation. We choose $\delta = 0.1$, $dt = 0.5$, and $x(0) = 0.5$ (Sec. 6.3).

- Synthetic paleomonsoon with AR1 an sinusoidal components, as discussed in Sec. 6.4. The paleomonsoon variable $x(t)$ is given by,

$$x(t-1) = x_P(t) + \alpha x_A(t) + s\mathcal{N}(0,1) + I(t) + F(t), \tag{D.7}$$

where the periodic sine component $x_P(t)$, the solar forcing component $F(t)$, and the event component $I(t)$ are,

$$x_P(t) = A_x \sin\left(\frac{2\pi}{T_x} t\right) \tag{D.8}$$

$$F(t) = -A_F \cos\left(\frac{2\pi}{T_F} t\right) \tag{D.9}$$

$$I(t) = \begin{cases} -3, \ 2250 \le t \le 2750 \\ 2, \ 4950 \le t \le 5050 \\ -2, \ 7450 \le t \le 7550. \end{cases} \tag{D.10}$$

Note in equation D.7, we proceed with time $t = N, N-1, \ldots, 3, 2, 1, 0$ as the time axis in this paleomonsoon example is years before present (BP), which is simply the difference of the AD year from 1950, e.g., –2000 AD equals 3950 BP. Also, the set of parameter values used in the example are $\alpha = 0.2$, $s = 0.5$ for the AR1 component $x_A(t) + s\mathcal{N}(0,1)$; $A_x = 0.5$, $T_x = 500$ for the sinusoid component $x_P(t)$; and $A_F = 5.$, $T_F = 20,000$ for the solar forcing $F(t)$.

Bibliography

Anoop, A., Prasad, S., Plessen, B., Basavaiah, N., Gaye, B., Naumann, R., Menzel, P., Weise, S., and Brauer, A. (2013). Palaeoenvironmental implications of evaporative Gaylussite crystals from Lonar lake, Central India. *Journal of Quaternary Science*, 28:349–359.

Bard, E. and Frank, M. (2006). Climate change and solar variability: What's new under the sun? *Earth and Planetary Science Letters*, 248:1–14.

Basavaiah, N., Wiesner, M., Anoop, A., Menzel, P., Nowaczyk, N., Deenadayalan, K., Brauer, A., Gaye, B., Naumann, R., Riedel, N., Stebich, M., and Prasad, S. (2014). Physicochemical analyses of surface sediments from the Lonar Lake, central India – implications for palaeoenvironmental reconstruction. *Fundamental and Applied Limnology*, 184:55–68.

Bastos, J. A. and Caiado, J. (2011). Recurrence quantification analysis of global stock markets. *Physica A*, 390:1315–1325.

Benestad, R. E. (2012). Reconciliation of global temperatures. *Environmental Research Letters*, 7:011002.

Blaauw, M. (2010). Methods and code for 'classical' age-modelling of radiocarbon sequences. *Quaternary Geochronology*, 5:512–518.

Blaauw, M., Christen, J. A., Mauquoy, D., van der Plicht, J., and Bennett, K. (2007). Testing the timing of radiocarbon-dated events between proxy archives. *The Holocene*, 17:283–288.

Blaauw, M., Wohlfarth, B., Christen, J., Ampel, L., Veres, D., Hughen, K., Preusser, F.,

and Svensson, A. (2010). Were last glacial climate events simultaneous between Greenland and France? A quantitative comparison using non-tuned chronologies. *Journal of Quaternary Science*, 25:387–394.

Blackwell, P. and Buck, C. (2008). Estimating Radiocarbon Calibration Curves. *Bayesian Analysis*, 3:225–248.

Bond, G., Kromer, B., Beer, J., Muscheler, R., Evans, M., Showers, W., Hoffmann, S., Lotti-Bond, R., Hajdas, I., and Bonani, G. (2001). Persistent solar influence on North Atlantic climate during the Holocene. *Science*, 294:2130–6.

Breitenbach, S. F. M., Rehfeld, K., Goswami, B., Baldini, J. U. L., Ridley, H. E., Kennett, D., Prufer, K., Aquino, V. V., Asmerom, Y., Polyak, V. J., Cheng, H., Kurths, J., and Marwan, N. (2012). COnstructing Proxy-Record Age models (COPRA). *Climate of the Past*, 8:1765–1779.

Brohan, P., Kennedy, J. J., Harris, I., Tett, S. F. B., and Jones, P. D. (2006). Uncertainty estimates in regional and global observed temperature changes: A new data set from 1850. *Journal of Geophysical Research*, 111:D12106.

Bronk Ramsey, C. (2008). Deposition models for chronological records. *Quaternary Science Reviews*, 27:42–60.

Charman, D., Barber, K., Blaauw, M., Langdon, P., Mauquoy, D., Daley, T., Hughes, P., and Karofeld, E. (2009). Climate drivers for peatland palaeoclimate records. *Quaternary Science Reviews*, 28:1811–1819.

Donges, J. F., Donner, R. V., Marwan, N., Breitenbach, S. F. M., Rehfeld, K., and Kurths, J. (2014). Nonlinear regime shifts in holocene asian monsoon variability: potential impacts on cultural change and migratory patterns. *Climate of the Past Discussions*, 10:895–975.

Donges, J. F., Donner, R. V., Trauth, M. H., Marwan, N., Schellnhuber, H.-J., and Kurths, J. (2011). Nonlinear detection of paleoclimate-variability transitions possibly related to human evolution. *Proceedings of the National Academy of Sciences*, 108:20422–20427.

Donges, J. F., Heitzig, J., and Donner, R. V. (2012). Analytical framework for recurrence network analysis of time series. *Physical Review E*, 85:1–16.

Donner, R. V., Heitzig, J., Donges, J. F., Zou, Y., Marwan, N., and Kurths, J. (2011). The geometry of chaotic dynamics – a complex network perspective. *European Physical Journal B*, 84:653–672.

Donner, R. V., Zou, Y., Donges, J. F., Marwan, N., and Kurths, J. (2010). Recurrence networks – a novel paradigm for nonlinear time series analysis. *New Journal of Physics*, 12:033025.

Duan, F., Wang, Y., Shen, C., Wang, Y., Cheng, H., Wu, C., Hu, H., Kong, X., Liu, D., and Zhao, K. (2014). Evidence for solar cycles in a late Holocene speleothem record from Dongge Cave, China. *Scientific Reports*, 4:5159.

Eckmann, J. P., Kamphorst, S. O., and Ruelle, D. (1987). Recurrence plots of dynamical systems. *Europhysics Letters*, 4:973–977.

Foster, G. and Rahmstorf, S. (2011). Global temperature evolution 1979–2010. *Environmental Research Letters*, 6:044022.

Ghil, M. and Vautard, R. (1991). Interdecadal oscillations and the warming trend in global temperature time series. *Nature*, 350:324–327.

Giles, J. (2002). When doubt is a sure thing. *Nature*, 418:476–478.

Goswami, B., Ambika, G., Marwan, N., and Kurths, J. (2012). On interrelations of recurrences and connectivity trends between stock indices. *Physica A*, 391:4364–4376.

Goswami, B., Heitzig, J., Rehfeld, K., Marwan, N., Anoop, A., Prasad, S., and Kurths, J. (2014). Estimation of sedimentary proxy records together with associated uncertainty. *Nonlinear Processes in Geophysics*, 21:1093–1111.

Goswami, B., Marwan, N., Feulner, G., and Kurths, J. (2013). How do global tem-

perature drivers influence each other? *European Physical Journal Special Topics*, 222:861–873.

Hansen, J., Ruedy, R., Sato, M., and Lo, K. (2010). Global Surface Temperature Change. *Reviews of Geophysics*, 48:RG4004.

Hansen, J., Sato, M., Ruedy, R., Kharecha, P., Lacis, A., Miller, R., Nazarenko, L., Lo, K., Schmidt, G. a., Russell, G., Aleinov, I., Bauer, S., Baum, E., Cairns, B., Canuto, V., Chandler, M., Cheng, Y., Cohen, a., Del Genio, a., Faluvegi, G., Fleming, E., Friend, a., Hall, T., Jackman, C., Jonas, J., Kelley, M., Kiang, N. Y., Koch, D., Labow, G., Lerner, J., Menon, S., Novakov, T., Oinas, V., Perlwitz, J., Perlwitz, J., Rind, D., Romanou, a., Schmunk, R., Shindell, D., Stone, P., Sun, S., Streets, D., Tausnev, N., Thresher, D., Unger, N., Yao, M., and Zhang, S. (2007). Climate simulations for 1880–2003 with GISS modelE. *Climate Dynamics*, 29:661–696.

Heaton, T., Blackwell, P., and Buck, C. (2009). A Bayesian approach to the estimation of radiocarbon calibration curves: The INTCAL09 methodology. *Radiocarbon*, 51:1151–1164.

Heegaard, E., Birks, H., and Telford, R. J. (2005). Relationships between calibrated ages and depth in stratigraphical sequences: an estimation procedure by mixed-effect regression. *The Holocene*, 15:612–618.

Hegerl, G., Zwiers, F., Braconot, P., Gillett, N., Luo, Y., Marengo Orsini, J., Nicholis, N., Penner, J., and Stott, P. (2007). Understanding and attributing climate change. In Solomon, S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, K., Tignor, M., and Miller, H., editors, *Climate Change 2007: The Physical Science Basis.*, pages 663–746. Cambridge University Press, Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change edition.

Heitzig, J. (2013). Moving Taylor Bayesian Regression for nonparametric multidimensional function estimation with possibly correlated errors. *SIAM Journal on Scientific Computing*, 35:A1928–A1950.

Hercman, H. and Pawlak, J. (2012). MOD–AGE: An age-depth model construction

algorithm. *Quaternary Geochronology*, In Press.

Hu, C., Henderson, G., Huang, J., Xie, S., Sun, Y., and Johnson, K. (2008). Quantification of Holocene Asian monsoon rainfall from spatially separated cave records. *Earth and Planetary Science Letters*, 266:221–232.

Hua, Q., Barbetti, M., and Rakowski, A. (2013). Atmospheric rediocarbon for the period 1950–2010. *Radiocarbon*, 55:2059–2072.

Jiang, N., Neelin, J., and Ghil, M. (1995). Quasi-quadrennial and quasi-biennial variability in the equatorial pacific. *Climate Dynamics*, 12:101–112.

Johnson, K. (2011). Palaeoclimate: Long-distance relationship. *Nature Geoscience*, 4:426–427.

Jones, C. and Cox, P. (2001). Modeling the volcanic signal in the atmospheric CO2 record. *Global Biogeochemical Cycles*, 15:453–465.

Katz, R. W., Craigmile, P. F., Guttorp, P., Haran, M., Sansó, B., and Stein, M. L. (2013). Uncertainty analysis in climate change assessments. *Nature Climate Change*, 3:769–771.

Lean, J. (2009). Cycles and trends in solar irradiance and climate. *Wiley Interdisciplinary Reviews: Climate Change*, 1:111–122.

Lean, J. L. and Rind, D. H. (2008). How natural and anthropogenic influences alter global and regional surface temperatures: 1889 to 2006. *Geophysical Research Letters*, 35:L18701.

Lean, J. L. and Rind, D. H. (2009). How will Earth's surface temperature change in future decades? *Geophysical Research Letters*, 36:1–5.

Li, S., Zhao, Z., Wang, Y., and Wang, Y. (2011). Identifying spatial patterns of synchronization between NDVI and climatic determinants using joint recurrence plots. *Environmental Earth Sciences*, 64:851–859.

Liu, F. and Feng, Z. (2012). A dramatic climatic transition at ~4000 cal. yr BP and its cultural responses in Chinese cultural domains. *The Holocene*, 22:1181–1197.

Liu, Y.-H., Henderson, G. M., Hu, C.-Y., Mason, a. J., Charnley, N., Johnson, K. R., and Xie, S.-C. (2013). Links between the East Asian monsoon and North Atlantic climate during the 8,200 year event. *Nature Geoscience*, 6:117–120.

Maher, B. (2008). Holocene variability of the East Asian summer monsoon from Chinese cave records: a re-assessment. *The Holocene*, 18:861–866.

Marcus, S., Chao, Y., Dickey, J., and Gegout, P. (1998). Detection and modeling of nontidal oceanic effects on Earth's rotation rate. *Science*, 281:1656–1659.

Marwan, N. (2008). A historical review of recurrence plots. *European Physical Journal Special Topics*, 164:3–12.

Marwan, N., Donges, J. F., Zou, Y., Donner, R. V., and Kurths, J. (2009). Complex network approach for recurrence analysis of time series. *Physics Letters A*, 373:4246 – 4254.

Marwan, N., Romano, M., Thiel, M., and Kurths, J. (2007). Recurrence plots for the analysis of complex systems. *Physics Reports*, 438:237–329.

Marwan, N., Trauth, M. H., Vuille, M., and Kurths, J. (2003). Comparing modern and Pleistocene ENSO-like influences in NW Argentina using nonlinear time series analysis methods. *Climate Dynamics*, 21:317–326.

Maslin, M. and Austin, P. (2012). Climate models at their limit? *Nature*, 486:183–184.

McDermott, F. (2004). Palaeo-climate reconstruction from stable isotope variations in speleothems: a review. *Quaternary Science Reviews*, 23:901–918.

Meyers, S., O'Brien, J., and Thelin, E. (1999). Reconstruction of monthly SST in the tropical pacific ocean during 1868–1993 using adaptive climate basis functions. *Monthly Weather Review*, 127:1599–1612.

Miettinen, O. and Wang, J. (1981). An alternative to the proportionate mortality ratio. *American Journal of Epidemiology*, 114:144–148.

Moron, V., Vautard, R., and Ghil, M. (1998). Trends, interdecadal and interannual oscillations in global sea-surface temperatures. *Climate Dynamics*, 14:545–569.

Mudelsee, M., Fohlmeister, J., and Scholz, D. (2012). Effects of dating errors on nonparametric trend analyses of speleothem time series. *Climate of the Past*, 8:1637–1648.

Mudelsee, M., Scholz, D., Röthlisberger, R., Fleitmann, D., Mangini, A., and Wolff, E. (2009). Climate spectrum estimation in the presence of timescale errors. *Nonlinear Processes in Geophysics*, 16:43–56.

Newman, M. (2004a). Analysis of weighted networks. *Physical Review E*, 70:056131.

Newman, M. (2004b). Detecting community structure in networks. *European Physical Journal B*, 38:321–330.

Newman, M. E. J. (2004c). Detecting community structure in networks. *European Physical Journal B*, 38:321–330.

Newman, M. E. J. (2006). Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103:8577–82.

Niklaus, T., Bonani, G., Suter, M., and Wölfli, W. (1994). Systematic investigation of uncertainties in radiocarbon dating due to fluctuations in the calibration curve. *Nuclear Instruments and Methods in Physics Research B*, 92:194–200.

Niu, M., Heaton, T., Blackwell, P., and Buck, C. (2013). The Bayesian approach to radiocarbon calibration curve estimation: The IntCal13, Marine13, SHCal13 methodologies. *Radiocarbon*, 55:1905–1922.

Overpeck, J. and Webb, R. (2000). Nonglacial rapid climate events: Past and future. *Proceedings of the National Academy of Sciences*, 97:1135–1338.

Parnell, A., Haslett, J., Allen, J., Buck, C., and Huntley, B. (2008). A flexible approach

to assessing synchroneity of past events using Bayesian reconstructions of sedimentation history. *Quaternary Science Reviews*, 27:1872–1885.

Parnell, A. C., Buck, C. E., and Doan, T. (2011). A review of statistical chronology models for high–resolution, proxy–based holocene palaeoenvironmental reconstruction. *Quaternary Science Reviews*, 30:2948–2960.

Prasad, S., Anoop, A., Riedel, N., Sarkar, S., Menzel, P., Basavaiah, N., Krishnan, R., Fuller, D., Plessen, B., Gaye, B., Röhl, U., Wilkes, H., Sachse, D., Sawant, R., Wiesner, M., and Stebich, M. (2014). Prolonged monsoon droughts and links to Indo-Pacific warm pool: A Holocene record from Lonar Lake, central India. *Earth and Planetary Science Letters*, 391:171–182.

Rampino, M., Self, S., and Fairbridge, R. (1979). Can rapid climate change cause volcanic eruptions? *Science*, 206:826–829.

Rehfeld, K. and Kurths, J. (2014). Similarity estimators for irregular and age-uncertain time series. *Climate of the Past*, 10:107–122.

Rehfeld, K., Marwan, N., Breitenbach, S. F. M., and Kurths, J. (2012). Late Holocene Asian summer monsoon dynamics from small but complex networks of paleoclimate data. *Climate Dynamics*, 41:3–19.

Reimer, P. J., Baillie, M. G. L., Bard, E., and Bayliss, A. (2009). Intcal09 and Marine09 radiocarbon age calibration curves, 0-50,000 years cal BP. *Radiocarbon*, 51:1111–1150.

Reimer, P. J., Bard, E., Bayliss, A., Beck, J. W., Blackwell, P. G., Bronk Ramsey, C., Buck, C. E., Cheng, H., Edwards, R. L., Friedrich, M., Grootes, P. M., Guilderson, T. P., Haflidason, H., Hajdas, I., Hatte, C., Heaton, T. J., Hoffmann, D. L., Hogg, A. G., Hughen, K. A., Kaiser, K. F., Kromer, B., Manning, S. W., Niu, M., Reimer, R., Richards, D. A., Scott, E. M., Southon, J. R., Staff, R. A., Turney, C. S. M., and van der Plicht, J. (2013). Intcal13 and Marine13 radiocarbon age calibration curves, 0-50,000 years cal BP. *Radiocarbon*, 55:1869–1887.

Richardson, D. C. and Dale, R. (2005). Looking to understand: the coupling be-

tween speakers' and listeners' eye movements and its relationship to discourse comprehension. *Cognitive science*, 29:1045–60.

Robock, A. (2000). Volcanic eruptions and climate. *Reviews of Geophysics*, 38:191–219.

Romano, M., Thiel, M., Kurths, J., and Grebogi, C. (2007). Estimation of the direction of the coupling by conditional probabilities of recurrence. *Physical Review E*, 76:036211.

Romano, M. C., Thiel, M., Kurths, J., Kiss, I. Z., and Hudson, J. L. (2005). Detection of synchronization for non-phase-coherent and non-stationary data. *Europhysics Letters*, 71:466–472.

Runge, J., Heitzig, J., and Petoukhov, V. (2012). Escaping the curse of dimensionality in estimating multivariate transfer entropy. *Physical Review Letters*, 108:1–5.

Sato, M. and Hansen, J. E. (1993). Stratospheric aerosol optical depths, 1850–1990. *Journal of Geophysical Research*, 98:987–994.

Scholz, D. and Hoffmann, D. L. (2011). StalAge — An algorithm designed for construction of speleothem age models. *Quaternary Geochronology*, 6:369–382.

Serrà, J., Serra, X., and Andrzejak, R. G. (2009). Cross recurrence quantification for cover song identification. *New Journal of Physics*, 11:093017.

Sinha, A., Berkelhammer, M., Stott, L., Mudelsee, M., Cheng, H., and Biswas, J. (2011). The leading mode of Indian Summer Monsoon precipitation variability during the last millennium. *Geophysical Research Letters*, 38:L15703.

Spiegelhalter, D., Pearson, M., and Short, I. (2011). Visualizing uncertainty about the future. *Science*, 333:1393–1400.

Spiegelhalter, D. J. (2014). The future lies in uncertainty. *Science*, 345:264–265.

Stainforth, D. A., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, D. J., Kettleborough, J. A., Knight, S., Martin, A., Murphy, J. M., Piani, C., Sexton, D., Smith,

L. A., Spicer, R. A., Thorpe, A. J., and Allen, M. R. (2005). Uncertainty in predictions of the climate response to rising levels of greenhouse gases. *Nature*, 433:403–6.

Stevenson, S. (2012). *The Past , Present and Future of the El Niño / Southern Oscillation*. PhD thesis, Faculty of the Graduate School, University of Colorado.

Stevenson, S., Fox-Kemper, B., Jochum, M., Neale, R., Deser, C., and Meehl, G. (2012). Will there be a significant change to El Niño in the twenty-first century? *Journal of Climate*, 25:2129–2145.

Stott, P. and Kettleborough, J. (2002). Origins and estimates of uncertainty in predictions of twenty-first century temperature rise. *Nature*, 1068:723–727.

Swindles, G., Blaauw, M., Blundell, A., and Turner, T. (2012). Examining the uncertainties in a 'tuned and stacked' peatland water table reconstruction. *Quaternary International*, 268:58–64.

Tan, L., Cai, Y., Cheng, H., An, Z., and Edwards, R. L. (2009). Summer monsoon precipitation variations in central China over the past 750 years derived from a high-resolution absolute-dated stalagmite. *Palaeogeography, Palaeoclimatology, Palaeoecology*, 280:432–439.

Thiel, M., Romano, M. C., and Kurths, J. (2004). How much information is contained in a recurrence plot? *Physics Letters A*, 330:343 – 349.

Thiel, M., Romano, M. C., Kurths, J., Rolfs, M., and Kliegl, R. (2006). Twin surrogates to test for complex synchronisation. *Europhysics Letters*, 75:535–541.

Thiel, M., Romano, M. C., Kurths, J., Rolfs, M., and Kliegl, R. (2008). Generating surrogates from recurrences. *Philosophical Transactions of the Royal Society A*, 366:545–57.

Trenberth, K. E., Jones, P. D., Ambenje, P., Bojariu, R., Easterling, D., Klein Tank, A., Parker, D., Rahimzadeh, F., Renwick, J. A., Rusticucci, M., Soden, B., and Zhai, P. (2007). Observations: Surface and atmospheric climate change. In Solomon,

S., Qin, D., Manning, M., Chen, Z., Marquis, M., Averyt, K. B., Tignor, M., and Miller, H. L., editors, *Climate Change 2007: The Physical Science Basis.*, pages 235–336. Cambridge University Press, Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change edition.

Wang, S., Ge, Q., Wang, F., Wen, X., and Huang, J. (2013). Abrupt climate changes of Holocene. *Chinese Geographical Science*, 23:1–12.

Wang, Y., Cheng, H., Edwards, R., He, Y., Kong, X., An, Z., Wu, J., Kelly, M., Dykoski, C., and Li, X. (2005a). The Holocene Asian Monsoon: Links to solar changes and North Atlantic climate. *Science*, 308:854–857.

Wang, Y.-M., Lean, J. L., and Sheeley, Jr., N. R. (2005b). Modeling the sun's magnetic field and irradiance since 1713. *The Astrophysical Journal*, 625:522–538.

Wefer, G., Berger, W. H., Bijma, J., and Fischer, G. (1999). Clues to ocean history: a brief overview of proxies. In Fischer, G. and Wefer, G., editors, *Use of Proxies in Paleoceanography: Examples from the South Atlantic*, pages 1–68. Springer Verlag, Berlin Heidelberg.

Williamson, R. and Downs, T. (1990). Probabilistic arithmetic I. Numerical methods for calculating convolutions and dependency bounds. *International Journal of Approximate Reasoning*, 4:89–158.

Wolter, K. and Timlin, M. (1998). Measuring the strength of ENSO events: How does the 1997/98 rank? *Weather*, 53:315–324.

Zbilut, J. P., Thomasson, N., and Webber, C. L. (2002). Recurrence quantification analysis as a tool for nonlinear exploration of nonstationary cardiac signals. *Medical Engineering & Physics*, 24:53–60.

Zhang, P., Cheng, H., Edwards, R. L., Chen, F., Wang, Y., Yang, X., Liu, J., Tan, M., Wang, X., Liu, J., An, C., Dai, Z., Zhou, J., Zhang, D., Jia, J., Jin, L., and Johnson, K. R. (2008). A test of climate, sun, and culture relationships from an 1810-year Chinese cave record. *Science*, 322:940–942.

Zou, Y., Donner, R. V., and Donges, J. F. (2010). Identifying complex periodic windows in continuous-time dynamical systems using recurrence-based methods. *Chaos*, 20:043130.

Zou, Y., Romano, M. C., Thiel, M., Marwan, N., and Kurths, J. (2011). Inferring Indirect Coupling By Means of Recurrences. *International Journal of Bifurcation and Chaos*, 21:1099–1111.