


UNIVERSITY OF BERGEN



Department of Geoscience

MASTERS THESIS

---

**Quantifying Earth System interactions  
over the Mid-Pleistocene Transition**

---

*Author: Maria Salem*

*Supervisor: Bjarte Hannisdal*

*Co-supervisor: Kristian Agasøster Haaga, Jo Brendryen*

December 2, 2020

"The heart of the scientific enterprise is a rational effort to understand the causes behind the phenomena we observe. In large-scale complex dynamical systems such as the Earth system, real experiments are rarely feasible. However, a rapidly increasing amount of observational and simulated data opens up the use of novel data-driven causal methods beyond the commonly adopted correlation techniques."

*(Runge et al. 2019)*

# Summary

Many different mechanisms have been proposed to explain the change in glacial intensity and duration that occurred during the Mid-Pleistocene Transition (MPT), 1250-700 thousand years ago. Despite a growing volume of paleoclimate archives and advances in climate modeling, however, there is no consensus on what caused the MPT. In this thesis I take state-of-the-art paleoclimate records as a starting point, and address the following questions: (1) According to the observed data, what were the main causal interactions among key climate system components in the Pleistocene? (2) Did the strength or directionality of these interactions change across the MPT? I use available time series of atmospheric CO<sub>2</sub> concentration and Northern Hemisphere summer insolation, as well as proxy records for global ice volume and Southern Hemisphere aeolian dust. To analyze causal connections, I use the Predictive Asymmetry method, a novel technique that combines concepts from dynamical systems theory and information theory to quantify causal coupling directly from observed time series without modeling the unknown mechanisms. My results show that there is empirical evidence in the paleoclimate records for causal connectivity among some of the key components of the Pleistocene climate system. Both atmospheric *p*CO<sub>2</sub> and Southern Ocean dust flux were important factors influencing ice volume changes in the Late Pleistocene, with limited evidence of external forcing by insolation. Furthermore, dynamical information in the observed records indicate the some of the causal interactions in the climate system changed across the MPT. Specifically, forcing of ice volume changes by Southern Ocean dust appeared during the MPT and became prominent in the Late Pleistocene. These findings may help evaluate the competing hypotheses for explaining the MPT. This thesis is the first to make extensive use of Predictive Asymmetry on paleoclimate records, thus contributing valuable insights into the merits and limitations of this recently developed technique when studying the complex interconnectivity of the Earth System.

# Acknowledgment

I would like to extend my most heartfelt thanks to my excellent supervisors Bjarte Hannisdal, Jo Brendryen and Kristian Agasøster Haaga. It has been a true pleasure to be a part of the Earth System Interactions team. You have allowed me to explore and meander, all the while generously offering your time, patience and good humour in both teaching and feedback. Your support during the work of the thesis has been inestimable. Thank you for the stimulating discussions, and for your contagious passion and vision for science. Not least, thank you for the bone-dry jokes that keep the spirit alive in the wee hours of flicking.

M.S.



# Contents

<b>Summary</b>	<b>ii</b>
<b>Acknowledgment</b>	<b>iii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Background</b>	<b>5</b>
2.1 Key aspects of the Pleistocene climate system . . . . .	5
2.2 Causal hypotheses for the Mid-Pleistocene Transition . . . . .	6
2.3 Causality from time series - a dynamical systems approach . . . . .	9
<b>3 Data</b>	<b>11</b>
3.1 Proxies for global ice volume . . . . .	11
3.1.1 Ice volume proxy records used in this project . . . . .	12
3.2 Insolation . . . . .	17
3.3 Atmospheric CO <sub>2</sub> concentration . . . . .	18
3.4 Southern hemisphere dust records . . . . .	19
<b>4 Causal Inference Framework</b>	<b>22</b>
4.1 Dynamical system reconstruction . . . . .	22
4.2 Transfer Entropy . . . . .	24
4.3 Predictive asymmetry . . . . .	26
4.4 Analysis design and work flow . . . . .	27
4.4.1 Defining the pre-, syn-, and post-MPT time windows . . . . .	27

---

4.4.2	Data wrangling . . . . .	29
4.4.3	Predictive asymmetry analysis . . . . .	31
4.4.4	Conventions for interpreting the results . . . . .	34
<b>5</b>	<b>Results</b>	<b>35</b>
5.1	Empirical evidence for causal connections in Pleistocene climate records	36
5.1.1	The pre-MPT time window . . . . .	36
5.1.2	The syn-MPT time window . . . . .	36
5.1.3	The post-MPT time window . . . . .	40
5.2	Sampling-standardized comparison of the pre- and post-MPT dynamics	42
<b>6</b>	<b>Discussion</b>	<b>48</b>
6.1	Implications for our understanding of Pleistocene climate system interactions . . . . .	48
6.2	Implications for the use of Predictive Asymmetry on paleoclimate records	52
6.3	Way forward . . . . .	55
<b>7</b>	<b>Conclusions</b>	<b>57</b>
	<b>References</b>	<b>58</b>
<b>A</b>	<b>Appendix 1: Predictive asymmetry results</b>	<b>66</b>
A.1	Predictive asymmetry results for the pre-MPT windows . . . . .	66
A.2	Predictive asymmetry results for the syn-MPT windows . . . . .	69
A.3	Predictive asymmetry results for the post-MPT windows . . . . .	72
A.4	Comparative analysis for pre- and post-MPT windows . . . . .	75
A.5	Sensitivity analyses . . . . .	77
A.5.1	Temporal binning of time series . . . . .	77
A.5.2	Binning of the delay reconstruction ( $\epsilon$ ) . . . . .	80
A.5.3	Time series length ( $N$ ) . . . . .	81

# List of Figures

3.1	Ice volume proxy records used in this study . . . . .	13
3.2	Northern hemisphere summer insolation . . . . .	17
3.3	Atmospheric CO <sub>2</sub> records. . . . .	18
3.4	Atmospheric dust and Fe supply to the Southern Ocean over the last 4 Ma	20
4.1	Time series and attractor of a deterministic system with chaotic behaviour.	23
4.2	Delay reconstruction and the invariant probability distribution . . . . .	24
4.3	Calculating transfer entropy from the invariant probability distribution.	25
4.4	Transfer entropy and predictive asymmetry on short and noisy time series	28
4.5	Definition of time windows and overview of time series . . . . .	30
5.1	Overview of the inferred causal networks . . . . .	37
5.2	Predictive asymmetry results for the pre-MPT time window . . . . .	38
5.3	Predictive asymmetry results for the syn-MPT time window . . . . .	41
5.4	Predictive asymmetry results for the post-MPT time window at 1 kyr resolution . . . . .	43
5.5	Predictive asymmetry results for the post-MPT time window at 500 yr resolution . . . . .	44
5.6	Sampling-standardized comparison of the pre- and post-MPT dynamics	47
6.1	Qualitative summary of causal inferences . . . . .	52
A.1	Predictive asymmetry results for the pre-MPT time window (15740-1250 ka BP) . . . . .	67



---

A.2	Predictive asymmetry results for the longer pre-MPT control window (4000-1250 ka BP) . . . . .	68
A.3	Predictive asymmetry results for the syn-MPT time window (1240-1092 ka BP) . . . . .	70
A.4	Predictive asymmetry results for the longer MPT control window (1250-700 ka BP) . . . . .	71
A.5	Predictive asymmetry results for the post-MPT time window (492-13 ka BP) . . . . .	73
A.6	Predictive asymmetry results for the longer post-MPT control interval (797-13 ka BP) . . . . .	74
A.7	Sampling-standardized comparison of the pre- and post-MPT dynamics	76
A.8	Predictive asymmetry results for the post-MPT time window (492-13 ka BP) at 500 yr resolution . . . . .	77
A.9	Sensitivity analysis on temporal resolution of over-sampled ChaCO <sub>2</sub> record. . . . .	78
A.10	Sensitivity analysis on temporal resolution of GraSL record. . . . .	79
A.11	Sensitivity analysis on $\epsilon$ over the synMPT and control MPT time intervals ( $N=149$ and $N=551$ , respectively). . . . .	80



# Chapter 1

## Introduction

The Mid-Pleistocene Transition (1250-700 ka BP, hereafter MPT) is an unsolved mystery in paleoclimate research. During this transition, the duration and intensity of glacial intervals increased without any corresponding change in orbital forcing. Ongoing efforts to understand what caused the change in climate dynamics across the MPT have resulted in a plethora of proposed hypotheses but no conclusive answers so far, which highlights a gap in our understanding of the natural variability of Earth's climate system. The aim of this thesis is to use dynamical information in available paleoclimate records to help constrain the number of plausible hypotheses for the change in climate dynamics across the MPT.

The two main strategies for solving paleoclimate problems are to *gather more data*, and/or to *build better models*. Here, I take a step back and try a different approach to the problem. Using state-of-the-art empirical time series data on key climatic variables as a point of departure, I ask the following research questions:

1. According to the observed data, what were the main causal interactions among key climate system components in the Pleistocene?
2. Did the strength or directionality of these interactions change across the MPT?

Note that these questions involve *ontological* components (i.e. what were the causal interactions in the past?) that are inextricably linked to *epistemological* components (i.e. how do we extract causal evidence from observations?).

The data sets used are state-of-the-art time series records of the hypothesized key climatic variables in the Pleistocene climate system, namely: proxy data for ice volume ( $\delta^{18}\text{O}$  and global sea level), atmospheric  $\text{CO}_2$  concentration ( $p\text{CO}_2$ ), Northern

Hemisphere summer insolation, and proxies for marine export productivity (Southern Ocean Fe flux and Antarctic dust deposition). The choice of variables is based on the current understanding of the climate system, reflecting the processes that are considered key for explaining the transition according to the hypotheses with the most traction (more on this in chapter 2).

In order to address my research questions, I will use a novel, model-free method for detecting and quantifying causal connectivity in dynamical systems, directly from time series data. A model-free approach in this context means that I do not make any model assumptions about the underlying mechanisms (governing equations are not specified). Instead, I try to leverage dynamical information in the time series data themselves to detect causal directionality and characterize the strength of coupling.

The main contributions of this thesis have the potential to advance the status quo by

- Testing whether or not available records actually contain sufficient information to determine causal connectivity, which in itself may be of value for future studies of Pleistocene climate dynamics.
  - If state-of-the-art data *do not* contain sufficient information, then more and better data may indeed be required in order to constrain and validate climate models.
  - Conversely, if the data contain evidence for dynamical coupling between key climatic variables, then this evidence may be useful for guiding mechanistic modeling and hypothesis testing.
- Testing whether we can detect any changes in the underlying dynamics (i.e. changes in strength and directionality of interactions between the variables) across the MPT. If so, then the results may help constrain the range of plausible hypotheses for explaining the MPT.
- Quantitatively characterizing uncertainties in both measurement values and chronology (age estimates) and propagating these uncertainties into the causal analyses

My main findings suggest that there is information in the data to determine causal coupling between the variables: Both Southern Ocean Fe flux and  $p\text{CO}_2$  were significant drivers of ice volume changes in the late Pleistocene, whereas forcing from insolation was weak or insignificant on the time scales studied here. Furthermore, I find that the (strong) influence of Southern Ocean Fe flux on ice volume is a coupling that emerged during the MPT transition. This finding is robust regardless of differences in data quality for the time intervals before and after the MPT.

The remainder of this thesis is structured as follows:

Chapter 2 presents a brief review of the key aspects of the Pleistocene climate system, and the principal causal hypotheses proposed to explain the MPT climatic transition. I also introduce the notion of model-free causal analysis of time series.

Chapter 3 presents the data used in the analyses, including a discussion of the assumptions and uncertainties underlying each time series.

In chapter 4 I present the causal analysis framework and the method of predictive asymmetry. Because the method is new and unfamiliar, I find it useful to introduce the underlying concepts of dynamical systems reconstruction and transfer entropy. I also describe the analysis design and work flow.

In chapter 5 I present the predictive asymmetry results of the pre-, syn-, and post-MPT time windows as well as an analysis allowing comparison of the relative magnitudes of coupling between before and after the MPT.

Chapter 6 discusses the findings in the context of the previously proposed causal hypotheses. Additionally, I will discuss some of the merits and challenges of the predictive asymmetry method, and point to some important avenues for further study.

Finally, I summarize my main conclusions in Chapter 7. All code needed to reproduce the results in this thesis is attached in the appendices.



# Chapter 2

## Background

### 2.1 Key aspects of the Pleistocene climate system

Ever since the first continental glaciation hypothesis was put forward by Jens Esmark (1826), the recurrent glaciations of the Northern Hemisphere in the last  $\sim 3$  Myrs have remained a paleoclimatic puzzle (Raymo and Huybers, 2008). Historically, two main theories have been proposed to explain the Ice Ages: one astronomical, positing that orbital control of insolation intensity is the pacemaker of ice ages, and one geochemical, positing that minor imbalances in the carbon cycle causes the climatic oscillations (Paillard, 2015).

Although the orbital hypothesis was proposed much earlier, Milankovitch (1941) was the first to make precise calculations of how the orbital parameters eccentricity, obliquity, and precession modulated the regional distribution of insolation quantities received on Earth, thus formalizing a theory of how this could affect Earth's long-term climate. Milankovitch theory, which is the favoured hypothesis to this day (Raymo and Huybers, 2008), posits that glaciations occur when insolation intensities at high northern latitudes is weak, allowing snow and ice to survive through the melting season and gradually accumulate into an ice sheet. These conditions are favoured when Earth's spin axis is less tilted (governed by obliquity) and when the Earth is far from the sun (governed by precession), which follows periodicities of  $\sim 41$  kyr and  $\sim 21$  kyr, respectively.

The impact of greenhouse gases on Earth's climate had been suspected since the work of Fourier (Fourier, 1824). Aiming to solve the problem of ice ages, Arrhenius (Arrhenius, 1896) was the first to make calculations of the role of  $p\text{CO}_2$  on climate, finding that the estimated  $3^\circ\text{C}$  global cooling estimated for the ice ages could be reproduced

by reducing  $p\text{CO}_2$  by a third. Arrhenius also pointed out that, because the atmosphere was a relatively small carbon reservoir compared to the ocean or rocks,  $p\text{CO}_2$  could easily be changed over time by accumulating small imbalances between the reservoirs.

In the 20th century, the accumulation of paleoceanographic data, as well as major progress in geochemistry and dating methods, allowed for a more rigorous study of the Northern Hemisphere glaciation cycles. Evidence backing the orbital hypothesis was unearthed in the 1970s, when oxygen isotope ratios in marine carbonate ( $\delta^{18}\text{O}$ ) from a deep-sea sediment core revealed the same  $\sim 41$  kyr periodicities as predicted by Milankovitch theory (Hays et al., 1976). However, Hays et al. (1976) also observed a switch to a dominant  $\sim 100$  kyr cyclicity in the last 800 kyrs of the record, with slow build-up and rapid collapse of ice sheets. This  $\sim 100$ -kyr periodicity, although matching Earth's eccentricity periodicity, was perplexing, because the effect of eccentricity variability on radiative forcing is negligible and therefore cannot alone explain glacial-interglacial cycles as a linear response to insolation forcing. This led Hays et al. (1976) to conclude that some non-linear mechanism is needed to account for the phase-locking of glacials to the 100-kyr eccentricity cycles. Since then, several hypotheses have been proposed to explain the deep glacial maxima and their abrupt terminations.

When ice core data revealed that atmospheric  $\text{CO}_2$  concentrations had fluctuated in close step with the glacial-interglacial cycles, the role of  $p\text{CO}_2$  as a greenhouse gas was implicated in the glacial-interglacial cycles (Petit et al., 1999; Shackleton, 2000).

**The Mid-Pleistocene transition.** The switch from roughly symmetric glacial cycles with a dominant  $\sim 41$  kyr periodicity to higher-amplitude, asymmetric cycles with a dominant  $\sim 100$  kyr periodicity occurred in the interval 1.25-0.7 Ma, and is referred to as the Mid-Pleistocene Transition (MPT). A number of mechanisms have been put forward to explain the MPT, and here I briefly outline the main hypotheses that can be directly tied to the available paleoclimate data used in this thesis.

## 2.2 Causal hypotheses for the Mid-Pleistocene Transition

**Hypotheses involving only ice sheet dynamics.** There are several hypotheses that invoke only internal ice sheet mechanisms with insolation as the only external forcing mechanism, to explain the transition to 100-kyr periodicity of the late Pleistocene. One of these hypotheses is the 'regolith hypothesis' by Clark and Pollard (1998a), which posits that a change in the basal boundary condition of the ice sheet could explain the transition. The hypothesis builds on the seemingly contradicting observations that,



despite  $\delta^{18}\text{O}$  indicating smaller ice volumes in the Early Pleistocene, glaciological evidence suggest ice sheets expanded as far south as  $40^\circ\text{N}$ , which leads the authors to conclude that Early Pleistocene ice sheets were thin and expansive. Clark and Pollard (1998b) therefore posit that the basal flow of ice was elevated in the Early Pleistocene, and developed a hypothesis that an easily deformable layer of regolith had been substrata of the Early Pleistocene ice sheets, allowing little friction and a relatively fast flow of the ice, resulting in thinner, more extensive ice sheets which responded to the insolation cycles in a linear way. The subsequent erosion of the regolith would expose the crystalline basement rock and give more friction to the ice sheets, slowing their flow and allowing them to grow thicker, which would change the mass-balance both by reducing ablation and increasing accumulation, respectively. Because thicker ice takes longer to melt, the ice sheets would 'survive' the Milankovitch-predicted melting periods, which could explain the rise of the 100-kyr cycles of the Late-Pleistocene.

Coupled climate — ice-sheet models have been able to reproduce both the 41-kyr cyclicity observed prior to the MPT Huybers and Tziperman (2008) and the  $\sim 100$ -kyr cyclicity after the MPT Abe-Ouchi et al. (2013a) without the implication of  $p\text{CO}_2$ . Huybers and Tziperman (2008) have reproduced the 41-kyr cycles with ice sheet models using high basal sliding, reminiscent of the regolith hypothesis. Abe-Ouchi et al. (2013a) reproduced the 100-kyr cycles with a coupled climate — ice-sheet model, where the rapid glacial terminations can be explained by how growth of ice sheets across some threshold increases the sensitivity to insolation, triggering internal feedbacks between climate, ice sheets and isostatic rebound. Carbon dioxide is involved in this model, but is not determinative.

If ice sheet dynamics is in fact the driver of the 100-kyr cycles, and not just a response to some other factor responsible for the MPT, then we expect to find information in the ice volume proxy time series that helps predict the evolution of the other climate system parameters. More specifically, we would suspect to detect an influence from GSL to  $p\text{CO}_2$  and dust time series, rather than vice versa. We would also expect an influence from insolation to GSL, at least for the time intervals preceding the MPT. In the post-MPT, the effect of insolation on ice sheets may be insignificant, because the proposed linear response of ice sheets to insolation (Milankovitch, 1941) becomes more complicated (e.g. Hays et al., 1976).

**Hypotheses involving ice sheets, ocean circulation, and  $\text{CO}_2$ .** Atmospheric  $\text{CO}_2$  and Antarctic temperatures covary with glacial-interglacial cycles, with  $\sim 100$  ppm of the  $\text{CO}_2$  sequestered in the ocean during glacial periods (Delmas et al., 1980; Neftel et al., 1982), implicating  $p\text{CO}_2$  in the ice age dynamics. However, the precise mechanisms for

CO<sub>2</sub> flux between the ocean and atmospheric reservoir are still debated (e.g. Paillard, 2015).

One class of hypotheses ascribes the decrease in  $p\text{CO}_2$  during glacial periods to changes in ocean circulation, either by Southern Ocean sea ice stratification or reduction in meridional overturning circulation (MOC) during glacial intervals, both which would reduce upwelling and thereby allow a deep-ocean build-up of carbon (e.g. Lear et al., 2016; Farmer et al., 2019; Pena and Goldstein, 2014).

If these hypotheses are true, then I expect an influence between  $p\text{CO}_2$  and GSL time series, where the influence could be in either direction. However, to make a robust test of this class of hypotheses, one would need to include time series proxies of ocean circulation.

**Hypotheses involving dust — climate couplings.** Atmospheric dust is acknowledged to have significant impact on global climate in several ways: Firstly, through its influence the planet's radiative balance, both directly, as an aerosol (Tegen and Lacis, 1996; Tegen, 2003), and indirectly, through its effect on cloud formation (Fuks et al., 2017). Secondly, through supplying iron and other limiting micronutrients to high-nutrient, low-chlorophyll zones of the ocean to increase marine productivity (Martin, 1990; Martin et al., 1990; Winckler et al., 2016). Marine productivity, in turn, controls the flux of carbon from the atmosphere to the deep ocean.

Available records show that dust supply to the Southern Ocean (where CO<sub>2</sub> leaks from upwelling deep-water) increases during ice ages Martínez-García et al. (2011). Variations in iron fertilization of the Subantarctic zone is estimated to account for ~20 ppm of the glacial-interglacial  $p\text{CO}_2$  variability Martínez-García (2015). The study of Martínez-García et al. (2011) found that, in contrast to the gradual increase in atmospheric dust content observed over the past 3 Ma, Southern Ocean dust and iron flux rose sharply at the onset of the MPT, which is used to argue for an association between high dust input to the Southern Ocean and the emergence of deep glaciations during the MPT. Martínez-García et al. (2011) hypothesize that the climate-dust-CO<sub>2</sub> feedback may be the ultimate cause for the transition to the ~100-kyr late Pleistocene climate variability. Evidence backing these hypotheses is that dust deposition increases drastically during cold stages of climate cycles (Shaffer and Lambert, 2018). Furthermore, modelling indicates that increased dust may be the deciding factor in reaching the coldest glacial stages, and maximizing CO<sub>2</sub> drawdown.

If the hypothesis that increased atmospheric dust concentration was the ultimate cause behind the MPT is true, then we expect to find an increased influence from dust to GSL over the post-MPT time window, as compared to the pre-MPT time window. Further-

more, if the responsible dust-climate cooling mechanism is increased marine export productivity through Fe fertilization, we would also expect to find an influence from dust to  $p\text{CO}_2$  records.

## 2.3 Causality from time series - a dynamical systems approach

The study of causal connections in large-scale, complex dynamical systems such as the Earth System poses a formidable challenge, because controlled experiments are either not feasible or ethically problematic. As a substitute, fully coupled Earth System models can be used to run numerical experiments (Stocker et al., 2013; Zelinka et al., 2020). However, the only perfect model of the Earth system is the real world itself, and the errors associated with model approximations can be difficult to assess (Babtie et al., 2014). Hence, a data-driven approach to detecting causal connections may be a valuable complement to explicit process modeling (Hannisdal and Liow, 2018; Runge et al., 2019).

Geologists have a longstanding tradition of inferring causal mechanisms by visual inspection of observed time series (a.k.a. 'eyeballing'), or – if more quantitatively minded – by statistical analyses of (cross-)correlation or periodicities. In addition to the well-known danger of equating correlation with causation, these traditional approaches are based on linear concepts, and are thus essentially limited to linear systems. Linear systems can be broken down into parts that can be analyzed separately and then put back together, but this luxury cannot be afforded to nonlinear dynamical systems (Kantz and Schreiber, 2004). Moreover, geologists have typically focused on explaining individual 'events' (e.g. mass extinctions, isotope excursions), and if we seriously believe that the Earth is a complex, nonlinear dynamical system, then we need to supplement the event-based causal focus with a more extensive, dynamical notion of causality (Hannisdal and Liow, 2018).

In this thesis, I use a data-driven (model-free) method that can infer causal relationships directly from time series data. Although they have their limitations, model-free approaches to causal analysis of time series can be a useful tool, both to test existing hypotheses, and to help formulate new hypotheses and guide models of how a system is coupled (Hannisdal and Liow, 2018; Runge et al., 2019). The particular method I will use here is the novel predictive asymmetry statistic developed by Haaga et al. (2020), which is able to quantify causal relationships directly from time series data, without

modelling the unknown mechanisms. The predictive asymmetry is mathematically related to the underlying dynamics of the system generating the data, and has proven to be robust even for short and noisy time series, making the method applicable to empirical time series from a wide range of natural systems (Haaga et al., 2020). The predictive asymmetry builds on two fundamental concepts from the mathematical fields of dynamical systems theory and information theory, namely *delay reconstruction* (Takens, 1981; Deyle and Sugihara, 2011) and *transfer entropy* (Schreiber, 2000). These concepts are presented in Chapter 4.

# Chapter 3

## Data

### 3.1 Proxies for global ice volume

Eustatic sea-level records are of paramount importance in the study of Quaternary climate change and its dynamics, because global sea level reflects global ice volume. In the following sections I address how GSL records can be reconstructed from  $\delta^{18}\text{O}$  and some of the main challenges for GSL reconstruction. This will serve as context for the following sections, where I present and discuss the time series that are used as proxies for ice volume in this study.

**Oxygen isotopes as a signal of global sea-level.** The concept of the oxygen isotope fractionation signal ( $\delta^{18}\text{O}$ ) as a global sea-level, and inversely, ice volume indicator, is based on the fractionation of oxygen that happens when water evaporates (Emiliani, 1955). Water molecules containing the lighter  $^{16}\text{O}$  isotope evaporate more readily than those containing the heavier  $^{18}\text{O}$  isotopes. During glacials, large amounts of freshwater (enriched in  $^{16}\text{O}$ ) are sequestered in ice sheets, leaving the isotopic composition of seawater ( $\delta^{18}\text{O}_w$ ) enriched in the heavier isotope. This signal will be reflected in the isotopic composition of marine calcifying organisms, e.g. foraminifera, because the oxygen entering the  $\text{CaCO}_3$  shells of the calcifying plankton captures the oxygen composition of the sea water at the time of formation. As these organisms die, their skeletons accumulate on the sea floor where they become part of the sediment record. With age constraints, these sediment records can therefore be used to construct time series showing how  $\delta^{18}\text{O}$ , and thus global sea-level and ice volume, has changed over time.


However,  $\delta^{18}\text{O}$  recorded in foraminiferal calcite is also affected by other factors than the  $\delta^{18}\text{O}$  of sea water. One important biological process contaminating the GSL signal

is that foraminifera preferentially incorporate the heavier isotope at colder water temperatures, known as temperature fractionation (Chappell and Shackleton, 1986). Another confounding factor is the local hydrography ( $\delta^{18}\text{O}$  of sea water varies spatially between regions that are affected by net evaporation or net precipitation) (Elderfield et al., 2012), deviating from the mean ocean  $\delta^{18}\text{O}$  governed by global ice volume. In order to get a "pure" signal representative of global sea-level from foraminiferal  $\delta^{18}\text{O}$ , it is therefore necessary to correct for the effects of both temperature fractionation and local hydrography.

The advantage of using GSL reconstructions as a record of ice volume dynamics, rather than the 'raw'  $\delta^{18}\text{O}$  signal recorded in sediment records, is that the temperature component and hydrographical component of the  $\delta^{18}\text{O}$  signal are corrected for, so that the reconstruction explicitly represent changes in ice volume. I have nevertheless chosen to include the LR04  $\delta^{18}\text{O}$  stack in the analyses, because this is the only proxy for ice volume allowing a comparison of results across all time windows. Additionally, because the LR04 is an important reference record in paleoclimatic studies, several of the other time series used in this project are tuned to its age model, which, as will be discussed below, introduces a caveat for the causal inference approach used in this project.

### 3.1.1 Ice volume proxy records used in this project

There are many sea-level indicators for more recent times ( $<0.5$  Ma), GSL records of the post-MPT interval are therefore numerous and robust. Further back in time ( $>0.5$  Myr), GSL reconstructions rely heavily on the oxygen isotope signal of deep-sea benthic foraminifera ( $\delta^{18}\text{O}_b$ ), which typically have much larger methodological uncertainties (Rohling et al., 2014). In particular, poor constraints of deep-water temperature influences on the fractionation process, as well as of the spatial variability in  $\delta^{18}\text{O}_b$ , increase the uncertainty in using  $\delta^{18}\text{O}$  stacks as proxies for GSL (Rohling et al., 2014). Therefore, the level of confidence that can be ascribed to GSL-reconstructions has long been constrained by the lack of  $\delta^{18}\text{O}_b$ -independent reconstruction methods, as well as limited temporal overlap between reconstructions (Rohling et al., 2014). This has motivated the development of a  $\delta^{18}\text{O}_b$ -independent sea-level reconstruction method (Rohling et al., 2014; Grant et al., 2014), which is used in the GraSL and RohSL sea-level reconstructions described below. The age models of the GraSL and RohSL time series are a particular advantage for the approach of this project, which will be further discussed below. However, these records have other disadvantages for this project (also described below), which is why I will also include several  $\delta^{18}\text{O}_b$  based records.



**Figure 3.1: Ice volume proxy records used in this study.** The records are plotted with the 95% confidence interval ( $\pm 2\sigma$ ) in both age and value.

To address the challenge of limited confidence in individual GSL reconstructions I have chosen to run the causal analyses of an ensemble of time series, arguing that this will allow me to draw more robust conclusions on the causal connections between ice volume and the other system components included in the analysis I have chosen to include 5 different records to proxy for ice volume. Each of the records have their own strengths and shortcomings with regards to the analysis, which will be discussed below in the presentation of each record.

The advantage of using GSL reconstructions rather than the more direct  $\delta^{18}\text{O}$  record is that the temperature component and hydrographical component of the  $\delta^{18}\text{O}$  signal are removed, so that the reconstruction explicitly represents changes in ice volume. I have nevertheless chosen to include the LR04  $\delta^{18}\text{O}$  stack in the analyses, because this is the only proxy for ice volume allowing a comparison of results across all time windows. Additionally, since the LR04 is an important reference record in paleoclimatic studies, several of the other time series used in this project are tuned to its age model, which, as will be discussed below, introduces a caveat for the causal inference approach used in this project.

**LR04: global  $\delta^{18}\text{O}_b$  stack.** The LR04 is a principal component analysis (PCA) of a stack of 57 globally distributed  $\delta^{18}\text{O}_b$  records (Lisiecki and Raymo, 2005) (Fig. 3.1). In princi-

ple, GSL reconstructions are better proxies of changes in ice volume than the pure  $\delta^{18}\text{O}$  signal recorded in sediment cores, as the temperature component and hydrographical component of the  $\delta^{18}\text{O}$  signal are corrected for in GSL reconstructions. I have nevertheless chosen to include the LR04  $\delta^{18}\text{O}$  stack in the analyses, because this is the only proxy for ice volume allowing a comparison of results across all time windows.

*Age model.* The LR04 age model is constructed based on average sedimentation rates between magnetostratigraphic and orbital tie-points with absolute age constraints. More specifically, the  $\delta^{18}\text{O}$  signal is tuned to a simple ice model based on June 21st insolation at  $65^\circ\text{N}$  (i.e., the La2004 time series used herein for northern hemisphere insolation). For our analyses, the LR04 age model comes with a caveat: Because the LR04 record is orbitally tuned, it inherently assumes a causal lag where  $\delta^{18}\text{O}$  follows insolation. This built-in assumption naturally poses a caveat to the validity of our causal analysis, which uses a lag-based definition of causality to infer causal connectivity between this very insolation signal and ice volume. Because several of the other records used in this study are tuned to the LR04 age model, entailing the same built-in assumptions of orbital forcing, this caveat also applies for those records.

Lisiecki and Raymo (2005) report the LR04 age model uncertainty as "Including all sources of error, we therefore estimate the uncertainty in the LR04 age model to be 40 kyr from 5.3-5Ma, 30 kyr from 5-4 Ma, 15 kyr from 4-3 Ma, 6 kyr from 3-1 Ma, and 4 kyr from 1-0 Ma", which I have interpreted as the maximum confidence envelope ( $\pm 2\sigma$ ) of potential systematic deviations in the age model. To propagate these uncertainties in the analyses, I have included the LR04 age model uncertainty in the preprocessing of the LR04 record and the other records that are tuned to the LR04 age model.

**SpraSL: global sea-level stack.** A "reference" record of global sea-level equivalent to the LR04 is the Spratt & Lisiecki sea-level stack (Fig. 3.1), hereafter denoted *SpraGSL* (Spratt and Lisiecki, 2016). This is a record built on PCA of 4-7 sea-level records spanning the last 800 kyrs.

The strength of using a stack of several GSL reconstructions is that it gives a robust record with conservative confidence, but a drawback is the length, the SpraSL GSL stack only covers the Late Pleistocene ice sheet dynamics. Another drawback with this GSL record is that it is tuned to the LR04 reference stack, and thereby raises the same caveat on orbital tuning of the age model.

**EldSL: GSL record from temperature-deconvolution of  $\delta^{18}\text{O}_b$ .** The global sea-level reconstruction by Elderfield et al. (2012), hereafter denoted *EldGSL*, is a high resolution, continuous sea-level record spanning the past 1.574 Myrs (Fig. 3.1). The long time span covered by the EldSL record enables comparison of all three time spans of



interest: the pre-MPT, syn-MPT and post-MPT.

The EldSL record is based on the  $\delta^{18}\text{O}_b$ -signal from a marine sediment core at the Chatnam Rise in the Southern Pacific Ocean. The temperature component of the  $\delta^{18}\text{O}$  signal has been deconvoluted through a separate Mg/Ca paleothermometry reconstruction of deep-sea temperatures at the site. Because the hydrographical component of the  $\delta^{18}\text{O}$  signal is estimated to be insignificant at the core site (Elderfield et al., 2012), the temperature deconvolution of  $\delta^{18}\text{O}$  should leave a pure signal for global sea-level.

*Uncertainties.* For the EldSL record, no uncertainty is included in the public dataset, but the article reports an "error in  $\delta^{18}\text{O}_W$  of  $\pm 0.2\%$  from propagation of estimated temperature and  $\delta^{18}\text{O}$  uncertainties" (Elderfield et al., 2012). However, according to Rohling et al. (2014), the uncertainty in the Elderfield record is poorly constrained. Rohling et al. (2014) performed a probabilistic assessment of the EldSL record, where they find a total uncertainty is of about  $\pm 35$  meters ( $1\sigma$ ), the bulk of which is random calibration uncertainty on temperature sensitivity ( $T_S$ ) due to many unknowns. In addition there is the  $\delta^{18}\text{O}_W$  sea-level conversion uncertainty of  $\pm 10\%$  (0.1 meters uncertainty per meter GSL change is the standard ratio). Rohling et al. (2014) further state that this uncertainty "may appear large, but there is strong autocorrelation in the record, which leads to considerably tighter uncertainty limits to underlying 'mean' trends." For the purpose of this project, the point of interest of the GSL record is not the absolute sea-level, but rather the dynamics, or relative change, of sea-level. Including the random calibration uncertainty would therefore be an exaggeration of uncertainty. A more appropriate uncertainty estimate would be to include the sea-level conversion uncertainty given by Rohling et al. (2014), which is set to the standard ratio of  $1\sigma = \pm 10\%$  (e.g. 0.1 meters uncertainty per meter GSL change).


*Age model.* The Elderfield  $\delta^{18}\text{O}$  signal is tuned to the LR04 age model. I have therefore added potential systematic deviations reported for the LR04 age model when defining the EldGSL age uncertainties. The same caveats about orbital tuning thus apply to the EldSL record.

**GraSL: Red Sea sea-level record.** Rohling et al. (2014) make a point that the uncertainty in the deep-water temperature influences on the fractionation process are poorly constrained (Rohling et al., 2014), increasing the uncertainty in using  $\delta^{18}\text{O}_b$  stacks as proxies for GSL. This has motivated them, to develop new  $\delta^{18}\text{O}_b$  independent methods to reconstruct global sea-level. An  $\delta^{18}\text{O}_b$ -independent method of reconstructing GSL was first proposed with the Red Sea relative sea-level record (Grant et al., 2014). The Grant sea-level record (GraSL) (Fig. 3.1) shows relative sea-level at the Strait of Bab el Mandab  $RSL_{BeM}$  which links the Red Sea to the Indian Ocean.  $RSL_{BeM}$  is

reconstructed  $\delta^{18}\text{O}$  of pelagic foraminifera, coupled with an hydraulic basin model of the Red Sea (Grant et al., 2014). Here, at lower sea-level, the sill at the straight of Bab el Mandab becomes shallower, which reduces the exchange of water between the Red Sea and the Indian Ocean. Hence, there is less water exchange over the straight at sea-level low stands. Since the rate of evaporation is assumed to remain roughly constant, the Red Sea  $\delta^{18}\text{O}$  composition becomes relatively more influenced by evaporation fractionation, a signal that is recorded in the foraminifera of the Red Sea. By implementing the measured  $\delta^{18}\text{O}$  values in a hydrological model of the Red Sea basin, the water exchange is reconstructed, and thereby also the sea-level stand at the straight of Bab el Mandab.

*Age model.* The GraSL age model is not tuned to insolation cycles, which lessens the age model bias for our causal time series analysis. Instead, the GraSL age model makes use of the monsoon-signal of the Inter-Tropical Convergence Zone (ITCZ), which is recorded both in the Red Sea sedimentary basin and by the Sanbao speleothem in China, providing tie-points in absolute age. The age model is then constructed calculating the average sedimentation rates between the radiometrically dated tie points. Drawbacks with using the GraSL is record is that it only covers the last 500 kyrs, and this record is therefore only suitable for analysing post-MPT dynamics.

**RohSL: Mediterranean sea-level record.** To the best of my knowledge, the RohSL record is the only long (>0.5 Ma) sea-level reconstruction that does not inherently assume an insolation forcing on ice volume (Fig. 3.1). The RohSL record reconstructs relative sea level at the Strait of Gibraltar following the same philosophy as the Grant record Rohling et al. (2014), with some adaptations to the the Mediterranean Basin. Although the Mediterranean Basin has a more complex hydrological model, it also has a much longer and very well constrained  $\delta^{18}\text{O}$  stack, spanning back 5.3 Myrs Rohling et al. (2014). However, parts of the record is riddled with intervals of major surface freshwater dilution, referred to as 'sapropel' intervals, which are associated with anoxic sea-floor conditions and identified visually as a dark olive/black in sediment cores Rohling et al. (2014). Rohling et al. (2014) have used a signal processing approach to correct for the sapropel intervals, but comparison with other available GSL reconstructions allows for the identification of three intervals where the signal-processing approach may have failed to filter out the  $\delta^{18}\text{O}$  anomalies relating to sapropel intervals. The temporal distribution of the identified remaining  $\delta^{18}\text{O}$  anomalies are between 400-700 ka BP, i.e. compromising RohSL analyses over the post-MPT window. To avoid most of the issues arising from sapropel intervals in the record, I have chosen not to include the the RohSL record in analyses over the post-MPT interval.



**Figure 3.2: Northern hemisphere summer insolation.** Numerical solution of the La2004 equations (Laskar et al., 2004), showing daily mean top of atmosphere (TOA) solar flux at 65°N on June 21 (summer solstice).

*Assumptions underlying the RohSL record.* (Rohling et al., 2014) present two assumptions that must hold true for the RohSL record to be representative of global sea level. Firstly, tectonic stability at the straight of Gibraltar is assumed, however, this is less certain early (>3.8 Ma) part of the record (Rohling et al., 2014). The second assumption concerns the ' $\delta^{18}\text{O}_b$ -to-RSL-converter', where it is assumed that the past parameter relationships of the foraminiferal species used remained within the same uncertainty ranges as found for the Late Pleistocene).

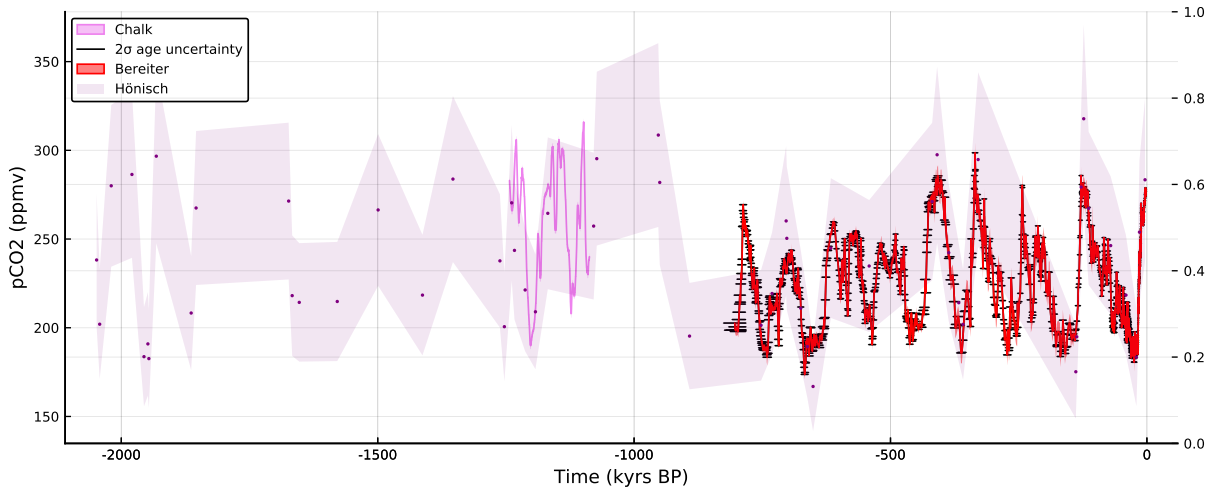
I take the liberty to add a third caveat to the RohSL record: there may be other sapropel intervals in the early parts of the record which, if missed by the signal processing approach, would go undetected because of limited temporal overlap between RohSL and other GSL records. I will therefore interpret the results from analyses with this record with caution.

In summary, none of the sea level records are perfect for studying the causal relationships in Pleistocene ice age dynamics. In this thesis, I therefore employ ensemble time series analysis over the selection of GSL records here presented, which presumably allows for more robust conclusions to be drawn.

## 3.2 Insolation

According to Milankovitch (1941), ice volume responds to summer energy at high northern latitudes. As a representative of northern latitude summer energy I here use the midsummer daily mean top of atmosphere (TOA) solar flux at 65°N, as given by the 'La2004' equations of Laskar et al. (2004) (Fig. 3.2). The numerical solution, hereafter denoted Ins, was computed using the AnalySeries software (Paillard et al., 1996).

*Uncertainties.* The chaotic behaviour of planetary orbits introduce uncertainties in the La2004 for projections further than 60 Ma into the past and future, but uncertainties are negligible over the time span considered here.



**Figure 3.3: Atmospheric CO<sub>2</sub> records.** a) Direct measurements from air bubbles in Antarctic ice cores (red with dots) (Bereiter et al., 2015), b)  $p\text{CO}_2$  reconstruction from the  $\delta^{11}\text{B}$  proxy (Chalk et al., 2017), c)  $p\text{CO}_2$  estimation from the  $\delta^{11}\text{B}$  proxy (Hönisch et al., 2009).

### 3.3 Atmospheric CO<sub>2</sub> concentration

For the Late Pleistocene climate system, available records of  $p\text{CO}_2$  are of high resolution and well constrained, owing to the great climate archives contained in ice sheets (Fig. 3.3). I use the Antarctic ice core record of  $p\text{CO}_2$  spanning the last 800 kyrs compiled by Bereiter et al. (2015, and refs. therein), here labeled BerCO<sub>2</sub>. Further back in time, the  $p\text{CO}_2$  data records are unfortunately either of low temporal resolution or short duration, owing both to the available climate archives and methodologies. A  $\delta^{11}\text{B}$ -based reconstruction of  $p\text{CO}_2$  by Hönisch et al. (2009) provide constraints on  $p\text{CO}_2$  over the past 2.1 Ma (Fig. 3.3). However, this record is not of sufficient resolution to contain dynamical information on the timescales studied here. To better resolve the role of  $p\text{CO}_2$  in the MPT, Chalk et al. (2017) reconstructed a high resolution  $p\text{CO}_2$  record (here labeled ChaCO<sub>2</sub>) spanning the early part of the MPT (ca 1.240-1.090 Ma).

**BerCO<sub>2</sub>: EDC ice core record of  $p\text{CO}_2$  (post-MPT).** The best record of  $p\text{CO}_2$  per today is the compilation of Antarctic ice core  $p\text{CO}_2$  records by Bereiter et al. (2015) and references therein, here labeled BerCO<sub>2</sub> (Fig. 3.3). BerCO<sub>2</sub> is a high resolution record with an uncertainty range of only a few ppm (Bereiter et al., 2015). The atmospheric composition of  $p\text{CO}_2$  through time is constructed by direct measurements of  $p\text{CO}_2$  from air bubbles trapped in the ice of Antarctic ice cores.

*Age model.* Absolute age constraints of the BerCO<sub>2</sub> record is derived from the AICC2012 age model (Bazin et al., 2013; Veres et al., 2013). The AICC2012 combines many different dating tools for different parts of the record, the details of which are outside the scope of this study. The AICC2012 age model has conservative estimates on the

age uncertainty, stemming from temporal divergences between Antarctic ice cores in recording of the same climatic events.

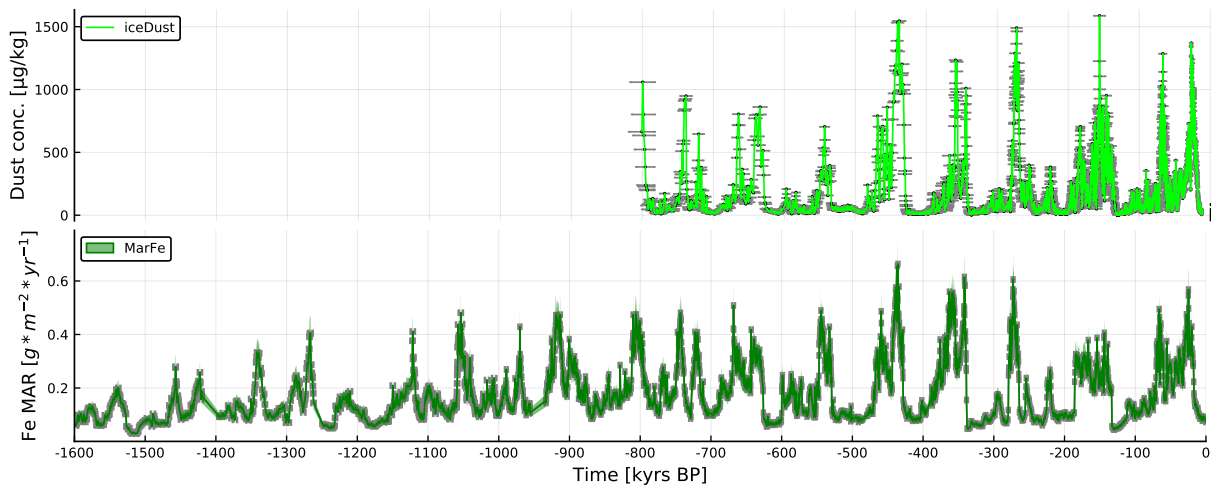
**ChaCO<sub>2</sub>:  $\delta^{11}\text{B}$  proxy record of  $p\text{CO}_2$ .** The ChaCO<sub>2</sub>  $p\text{CO}_2$  record is a  $\delta^{11}\text{B}$ -based record of atmospheric  $p\text{CO}_2$  spanning the early part of the MPT (ca 1.240-1.090 Ma)(Fig. 3.3). The boron isotope ratio ( $\delta^{11}\text{B}$ ) recorded in shells of foraminifera reflect sea-water pH at the time of formation (Hemming and Hanson, 1992). Because  $p\text{CO}_2$  in the surface ocean is in equilibrium with the atmosphere, atmospheric  $p\text{CO}_2$  will affect surface ocean pH, and in turn, the  $\delta^{11}\text{B}$  ratio in pelagic foraminifera. The ChaCO<sub>2</sub> record is communicated with associated uncertainty constraints quantified by Monte-Carlo resampling (Chalk et al., 2017). The record is resampled to a temporal resolution of 125 years Chalk et al. (2017), while the original sampling is one sample per 3.5-4.5 kyrs.

*Age model.* The ChaCO<sub>2</sub> age model was constructed by alignment of a local  $\delta^{18}\text{O}_b$  stack to the LR04 reference stack. I have therefore added the LR04 age model uncertainty to ChaCO<sub>2</sub> to propagate the uncertainties in the analyses.

### 3.4 Southern hemisphere dust records

The history of atmospheric dust is recorded in both ice cores and marine sediment cores. One advantage of using marine sediment cores to reconstruct dust flux, as opposed to ice core records, is that marine sediment cores cover much longer time spans. For example, the Antarctic EDC ice core record by Lambert et al. (2008) (IceDust) spans the last 800 kyrs, while the marine record (MarFe) yields dust flux data back to 4 Ma BP (Martínez-García et al., 2011). MarFe and IceDust show excellent agreement Martínez-García et al. (2011), indicating that they both record large scale depositional changes (Fig. 3.4).

The marine sediment core used in this project (MarFe, described below) is a (quasi-) continuous high resolution record spanning the last 4 Ma, and thereby covers all the time windows of analysis in this project. MarFe represents dust flux spatially delimited to the Southern ocean, which is identified as a key high-nutrient, low-chlorophyll zone, meaning that increased Fe flux has the potential to increase marine productivity here (Martin, 1990; Winckler et al., 2016). Therefore, I consider MarFe to be a well-suited time series to test the role of Fe fertilization-mediated drawdown of  $p\text{CO}_2$  across the MPT. The MarFe record also has high temporal resolution, and is continuous over the studied time windows, which allows the testing of climate-dust interactions across the studied time intervals. Additionally, I include the EDC ice core record of dust deposi-



**Figure 3.4: Atmospheric dust and Fe supply to the Southern Ocean over the last 4 Ma.**

tion by Lambert et al. (2008) (IceDust), because IceDust and BerCO<sub>2</sub> are both records from the same EDC ice core. Inclusion of the IceDust record therefore significantly reduces age uncertainties between  $p\text{CO}_2$  and dust in this project. Tighter relative age constraints is a great advantage for the sensitivity of the method used for causal inference, which operates on a lag-based definition of causality (more on this in chapter 4).

**MarFe: Marine sediment core of Southern Ocean Fe flux record.** The principal time series chosen to represent atmospheric dust is the Fe-component of a high-resolution dust flux record from a sediment core in the sub-Antarctic Atlantic ocean (Martínez-García et al., 2011), dubbed MarFe (Fig. 3.4). The MarFe record spans the last 4 Ma, and is thereby the first record to give insights to Southern Ocean dust deposition throughout the entire time span of Quaternary climatic transitions (Martínez-García et al., 2011).

*Age model.* The MarFe record is divided into three intervals that each are tuned to different age models. The first 800 kyrs of the record are tuned to the EDC3 age model (Parrenin et al., 2007) by graphic correlation of XRF Fe measurements to the IceDust record. The remainder of the record is on a core site-specific  $\delta^{18}\text{O}$ -based age model, with a slight modification in the latter part of the record (>2.9 Ma) based on biostratigraphy

**IceDust: EDC Ice core record of Antarctic dust deposition (post-MPT).** The IceDust record is a high-resolution record of Antarctic aeolian dust record as measured by dust concentration in the EDC ice core by Lambert et al. (2008) (Fig. 3.4). In their paper, Lambert et al. (2008) refer to the dust *flux* [ $\mu/\text{kg}/\text{year}$ ], which is the quantity that is implicated in the dust-climate hypotheses. Unfortunately only dust concentration in

snow [ $\mu/kg$ ] is reported in the published dataset. Flux and concentration are not necessarily equivalent in terms of recording dynamics (years with little snowfall will give a higher dust concentration even though they may not have a higher dust flux). Unfortunately, the dynamical parameters used by Lambert et al. (2008) to reconstruct variations snowfall could not be found, therefore the IceDust time series as used in this project is of dust concentration. Furthermore, it has been cautioned that the larger-scale climatic signal contained in Antarctic dust records may be contaminated by local glacial transport mechanisms (Shaffer and Lambert, 2018). I have nevertheless chosen to include this time series in the analyses to see if it still contains some dynamical information that can help resolve the relationship between southern hemisphere aeolian dust and  $pCO_2$ . As mentioned above, the IceDust record has the advantage that it allows for substantial reduction of age model uncertainties in analyses with  $pCO_2$  in this project, because IceDust and BerCO<sub>2</sub> are both from EDC ice cores, which reduces the relative age uncertainty to the lock-in depth of gas bubbles in the ice. Even if I will not hypothesize on what mechanistic process or regional extent the IceDust record may represent, I wish to explore the possibility that the IceDust record nevertheless may contain some dynamical information that will leave us better equipped to resolve the causal relationship between  $pCO_2$  and aeolian dust.

*Data.* To construct the Antarctic dust flux record (not found published, ref. section above), Lambert et al. (2008) collected both laser scatter (lpc) and coulter counter (cc) data of the dust concentration as well as fine particle percentage (fpp). However, I have chosen only to use the data collected with the laser scatter method (lpc) data, as this dataset has the highest resolution for the older parts of the record.

*Uncertainties and age model.* No uncertainties were reported for the original Lambert dataset, neither in age nor value. In order to align the IceDust record on the same age model as BerCO<sub>2</sub>, I have transferred the IceDust data from the original EDC3 age model (Parrenin et al., 2007; Lambert et al., 2008) to the AICC2012 age model (Bazin et al., 2013), yielding conservative estimates of age uncertainty.

## Chapter 4

# Causal Inference Framework

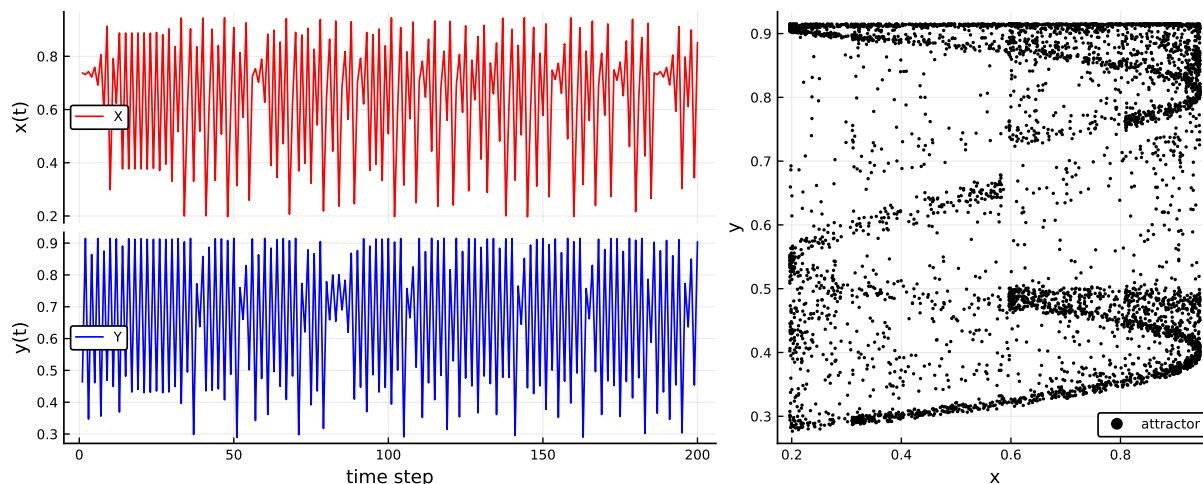
In the context of time series analysis, a wide range of methods have been designed to detect causal connectivity (see Hlaváčková-Schindler et al., 2007; Hannisdal and Liow, 2018; Runge et al., 2019, for reviews). Among these methods, the data-driven, model-free (a.k.a. non-parametric, or equation-free) approaches are of particular interest when studying natural systems, because they do not require mechanistic model assumptions, that is, we don't have to specify a model of how the system works. All of these methods have both theoretical and practical limitations, however, and to overcome some of these limitations, Haaga et al. (2020) recently proposed a novel approach called the Predictive Asymmetry, which I have used in this thesis.

In their paper, Haaga et al. (2020) show how the strength and directionality of causal interactions in dynamical systems can be studied by taking a difference between two information-theoretic quantities that are computed directly from observed time series. Crucially, they mathematically prove that the value of this difference is fundamentally related to the dynamics of the unknown system that generated the time series. The two main theoretical concepts underpinning the predictive asymmetry are dynamical systems reconstruction and transfer entropy. Because these concepts lie outside the traditional disciplinary boundaries of geology (and my own comfort zone), I will briefly introduce them in the following sections before moving onto the predictive asymmetry method itself.

### 4.1 Dynamical system reconstruction

Dynamical systems theory is the study of time-dependent processes (Strogatz, 2018). In nature, dynamical systems often show complex, nonlinear, and chaotic behaviour,

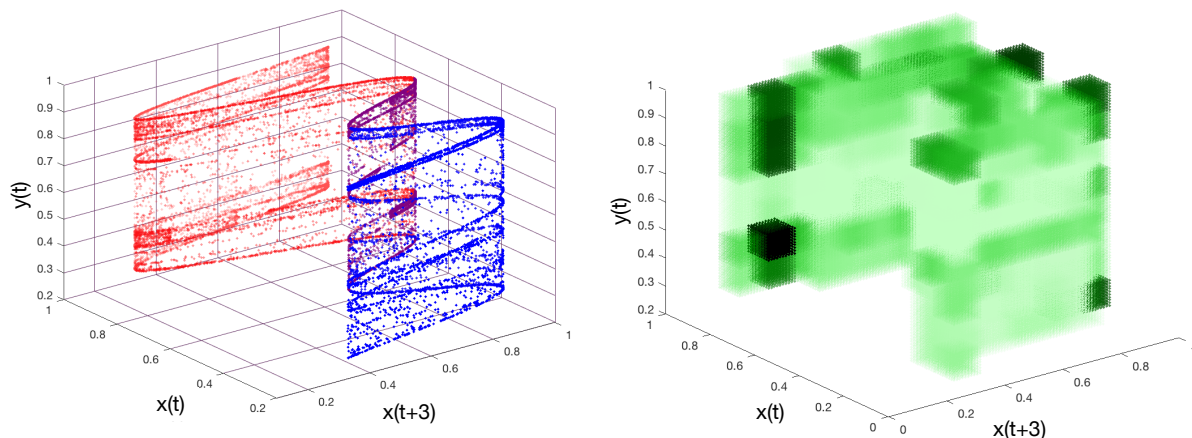




**Figure 4.1: Time series and attractor of a deterministic system with chaotic behaviour.** Left: Time series X and Y of a fully deterministic system with chaotic behaviour (logistic map). Right: State space configuration (attractor) of the system.

which can be difficult to model. Instead, we can take a geometric approach and describe the evolution of the system as trajectories in a state space defined by the components of the system (Fig. 4.1). These trajectories in the state space allow us to study the *invariant* properties of the dynamics (the attractor) describing the long-term behaviour of the system (well-known examples are the "limit cycles" of predator-prey systems, or the famous "butterfly" attractor of the Lorenz system). However, we typically don't have (or even want) access to all the components of a system, so how can we then obtain a faithful state-space representation of the system?

Four decades ago, mathematicians were able to show that a faithful representation of the state space (attractor) of a dynamical system can be obtained from a single time series realization of the system by *delay reconstruction* (Takens, 1981): by using lagged (time-delayed) coordinates of a time series to define reconstructed states, we obtain a state-space configuration that is topologically equivalent (in 1:1 correspondence) to the original system's attractor (Fig. 4.2). This result implies that we can study the invariant properties of the dynamics without having access to all the components of the system (Fig. 4.2). Although proven in the context of deterministic systems, equivalent concepts hold true also for stochastic systems, as long as they have some deterministic component. Delay reconstruction is the starting point for model-free approaches to detecting causality from time series. However, to make the leap to causality, we need the concept of transfer entropy.



**Figure 4.2:** Delay reconstruction and the invariant probability distribution. Left panel: Delay reconstruction of the coupled logistic map, now including an element of time, by adding a time delay to one of the time series. Right panel: An estimate of the invariant probability distribution of the system obtained by gridding the delay reconstruction and counting the number of times the system visits each cell in the grid (the visitation frequency). Figures courtesy of David Diego.

## 4.2 Transfer Entropy

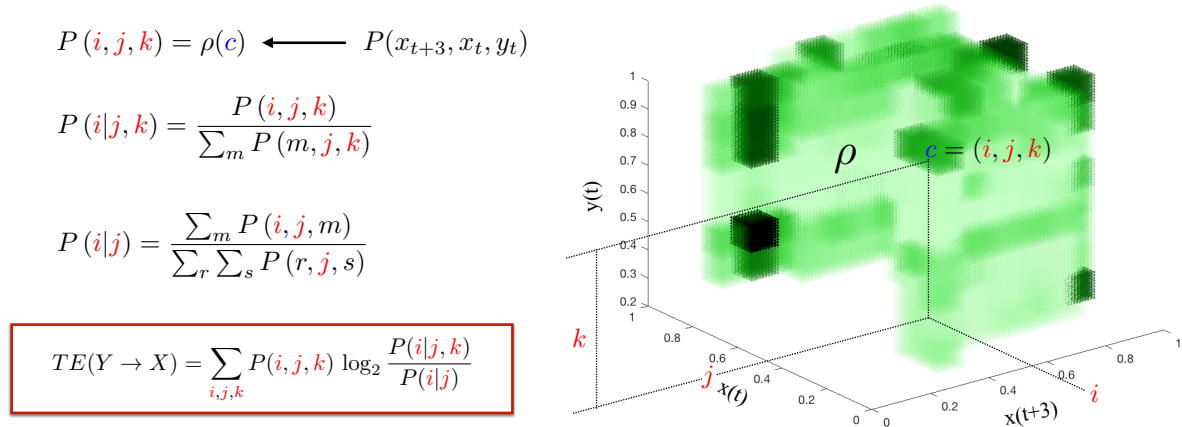
Causality is a concept that has eluded scientists and philosophers for centuries, but there is general agreement at least on two criteria that should be met in order to infer a causal relationship: (1) the cause should precede the effect, and (2) observing the cause should increase the probability of observing the effect (Hannisdal and Liow, 2018). In the context of time series, these two criteria were formalized by Schreiber (2000) in a quantity he called *transfer entropy*.

Transfer entropy from a source time series  $X$  to a target time series  $Y$ , both synchronously measured and consisting of  $N$  data points, is defined as

$$TE_{X \rightarrow Y}(\eta, \epsilon) = \sum p(x_t, y_t, y_{t+\eta}) \log \frac{p(y_{t+\eta} | y_t, x_t)}{p(y_{t+\eta} | y_t)}. \quad (4.1)$$

Equation 4.1 says that if knowing the state  $x_t$  increases the probability of knowing the state  $y_t$   $\eta$  steps in the future ( $y_{t+\eta}$ ), i.e.  $p(y_{t+\eta} | y_t, x_t) > p(y_{t+\eta} | y_t)$ , then the transfer entropy will be a positive number. Conversely, if (for all prediction lags  $\eta$ ), knowing  $x$  does not increase the probability of knowing  $y$  in the next prediction lag  $\eta$ , i.e. ( $p(y_{t+\eta} | y_t, x_t) = p(y_{t+\eta} | y_t)$ ), then the transfer entropy will (as  $N \rightarrow \infty$ ) equal 0.

Transfer entropy is thus a property of multivariate probability distributions constructed from multivariate time series data. In the context of dynamical systems, the relevant probability distribution is the invariant distribution that we estimate from the delay reconstruction (Fig. 4.2). By superimposing a grid on the delay reconstruction, we can



**Figure 4.3:** Calculating transfer entropy from the invariant probability distribution.

count the number of times the system visits each cell in the grid (i.e. the visitation frequency), which yields a multivariate histogram that estimates the invariant probability distribution (the probability of the system occupying different regions of the state space; Fig. 4.2). When we use a generalized delay reconstruction that includes both  $x$  and  $y$  components (Deyle and Sugihara, 2011), then the approximation of the invariant probability distribution contains the information we need to estimate the transfer entropy (Fig. 4.3).

To compute the transfer entropy in practice, we restate eq. 4.1 as a sum of information entropies:

$$TE_{X \rightarrow Y}(\eta, \epsilon) = -H(x_t, y_t, y_{t+\eta}) + H(y_{t+\eta}, y_t) + H(x_t, y_t) - H(y_t), \quad (4.2)$$

where  $H = -\sum p(x) \log p(x)$  is the information entropy as defined by Shannon (1948).

The reformulated transfer entropy (eq. 4.2) does not involve conditional probabilities and can be readily calculated from the joint probability distributions estimated by visitation frequency. With sufficient time series length, the estimated probability distributions will approach the true distributions, and thus transfer entropy (being a property of probability distributions) will also converge on the true transfer entropy.

Despite its intuitive appeal, however, the transfer entropy can suffer from statistical biases and yield ambiguous or misleading results in practical applications, as illustrated in the supplement of Haaga et al. (2020). It turns out, however, that by doing the seemingly non-sensical exercise of comparing the values of transfer entropy for positive prediction lags (predicting into the future) with the values for negative prediction lags (predicting into the past), we obtain a much more robust causality statistic: the predictive asymmetry (Haaga et al., 2020).

### 4.3 Predictive asymmetry

The predictive asymmetry ( $\mathbb{A}$ ) is defined as the difference between TE forwards in time and TE backwards in time:

$$\mathbb{A}_{X \rightarrow Y}(\eta, \epsilon) = \int_0^{\eta_{\max}} TE_{X \rightarrow Y}(\eta) - \int_0^{\eta_{\max}} TE_{X \rightarrow Y}(-\eta) \quad (4.3)$$

where  $\eta$  is the prediction lag, and  $\epsilon$  is the size of the cells used for gridding the delay reconstruction to estimate the invariant probability distribution. Because  $\mathbb{A}$  is defined using the transfer entropy, it is tempting to interpret the values as some kind of measure of "information flow", as Schreiber (2000) interpreted transfer entropy. Unlike transfer entropy, however, the predictive asymmetry is theoretically supported by a deep mathematical result relating the value of  $\mathbb{A}$  to the underlying dynamics of the (unknown) system that generated the time series (Haaga et al., 2020). Both analytical (exact) and numerical results in Haaga et al. (2020) show that:

- If there is no coupling between  $X$  and  $Y$ , then  $\mathbb{A}_{X \rightarrow Y}$  and  $\mathbb{A}_{Y \rightarrow X}$  converge on zero.
- If there is a unidirectional coupling, say  $X \rightarrow Y$ , then  $\mathbb{A}_{X \rightarrow Y}$  is positive and  $\mathbb{A}_{Y \rightarrow X} \leq 0$ .
- If there is a (strong) bidirectional coupling,  $X \leftrightarrow Y$ , then both  $\mathbb{A}_{X \rightarrow Y}$  and  $\mathbb{A}_{Y \rightarrow X}$  are positive.

Although  $\mathbb{A}$  is much more robust than transfer entropy itself, numerical estimation from short and noisy time series inevitably induces statistical fluctuations (Fig. 4.4). In addition, the intrinsic magnitude of the transfer entropy used in calculating  $\mathbb{A}$  is system-dependent, making it difficult to interpret the relative magnitude of  $\mathbb{A}$  in terms of coupling strength. To rectify this, Haaga et al. (2020) propose a normalization of  $\mathbb{A}$  to the system-characteristic magnitude of transfer entropy:

$$\mathcal{A}_{X \rightarrow Y}^f(\eta) := \frac{\mathbb{A}_{X \rightarrow Y}(\eta)}{\frac{f}{\eta} \int_{-\eta_{\max}}^{\eta_{\max}} TE_{X \rightarrow Y}(\eta) d\eta}. \quad (4.4)$$

Now, the relative magnitude of the normalized predictive asymmetry  $\mathcal{A}$  yields a measure of relative coupling strength that can be compared across different systems. Moreover, the normalization factor  $f$  equips the  $\mathcal{A}$  statistic with a built-in significance threshold with which we can statistically accept or reject the hypothesis of a causal coupling.

Increasing the normalization factor  $f$  will increase the specificity of our analyses (i.e. fewer false positives), but does so at the cost of the method's sensitivity (i.e. increase in false negatives). Appropriate significance thresholds have been determined heuristically with synthetic systems by analyzing false positive rates as a function of time series length and system types, indicating that  $f=1$  is sufficient in most cases (Haaga et al., 2020). That said, Haaga et al. (2020) stress that for time series with  $N=150$  data points or less, it is recommended to raise  $f$  to 1.5 for most systems. The conventions used in this thesis for interpreting the significance of the  $\mathcal{A}$  results is described at the end of this chapter.

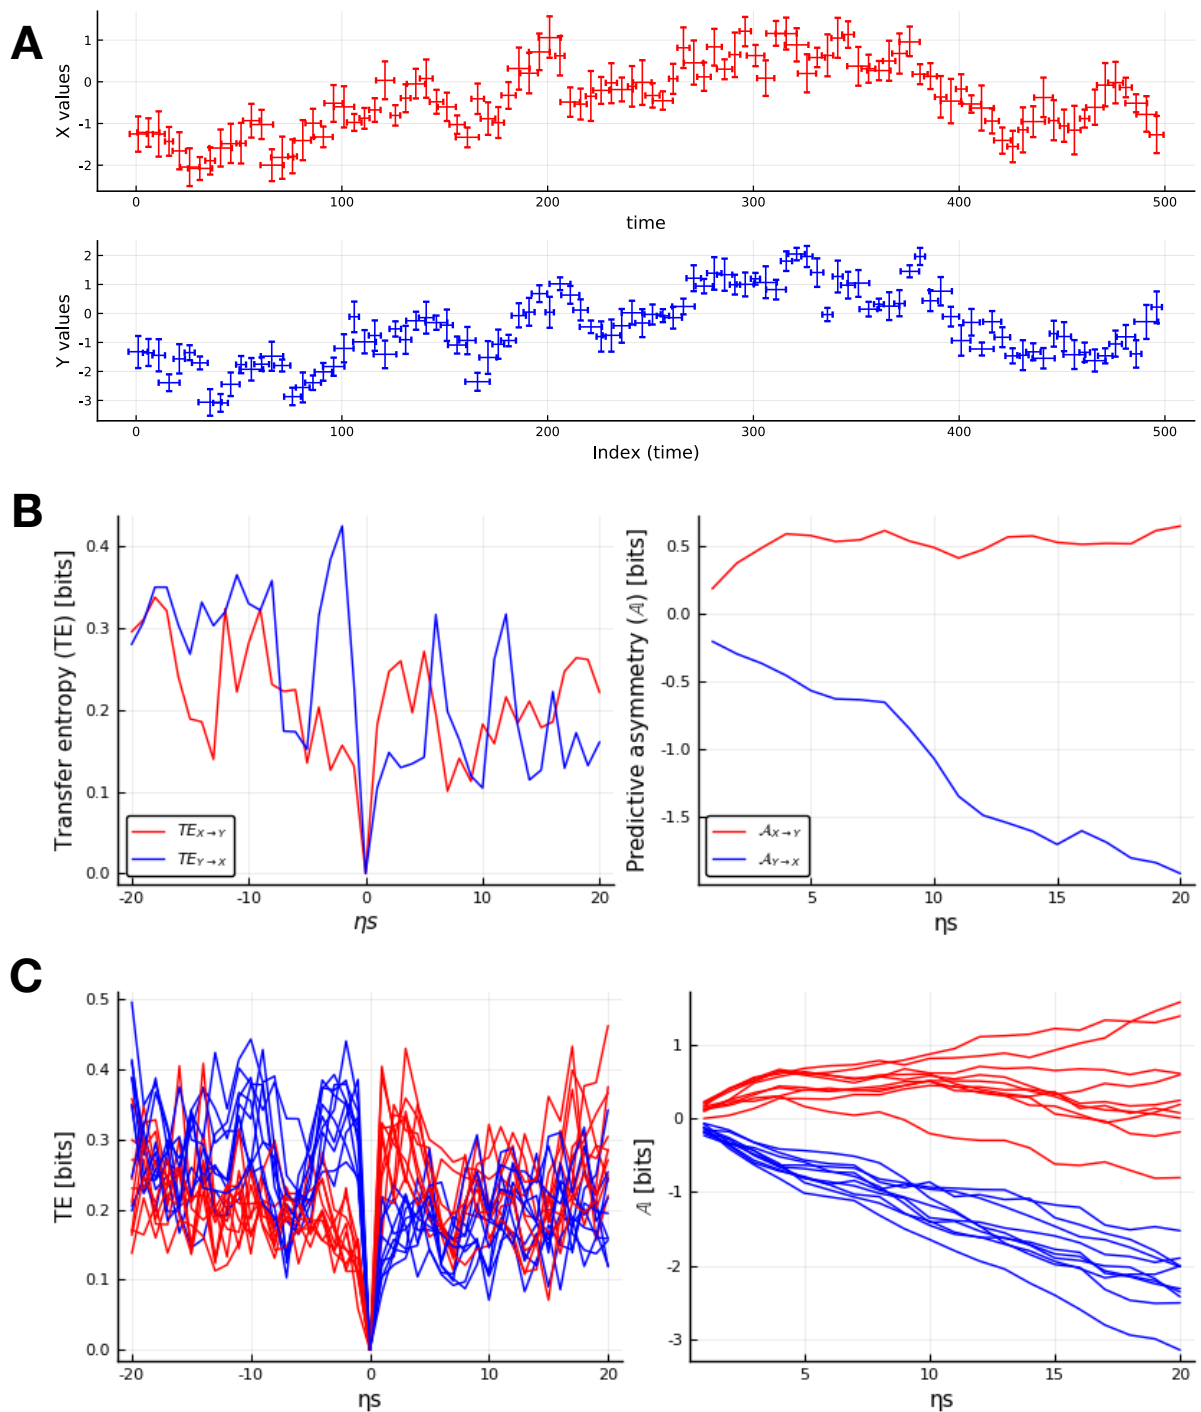
## 4.4 Analysis design and work flow

The main objectives of this project are to i) resolve whether or not there exists information in the data to determine causal connectivity between the Earth system components, and ii) assess whether or not there is evidence in the data of a change in the dynamics (change in causal coupling) of the climate system before and after the MPT. I have approached both objectives by predictive asymmetry analyses of the chosen time series in separate with time windows before, during and after the MPT (Fig. 4.5). In this section, I present the way I have chosen to design and implement a work flow for my analyses.

The notebooks (NBs) attached to the thesis contain all the work done in this project and allow for all of the results to be reproduced. Preparation of the data for analysis is documented in NB1. NB2 focuses on explaining the method through examples of synthetic systems, and serves as supplement to this chapter. The code used for the predictive asymmetry analyses is summarized in NB3. Finally, all the analyses and results for the different time intervals are found in the NBR notebooks.

### 4.4.1 Defining the pre-, syn-, and post-MPT time windows

**Pre-MPT.** To represent the climate system prior to the MPT, I have chosen the time window 1574-1250 ka BP, labeled pre-MPT (NBR-A results in appendix A.1). This time interval is chosen to allow the inclusion of three ice volume proxies in the analyses (EldSL, in addition to RohSL and LR04) allowing for more robust conclusions to be drawn regarding ice volume dynamics. Because the robustness of  $\mathcal{A}$  depends to some extent on time series length, additional analyses covering a longer time window (4000-1250 ka BP) are also performed for comparison, but this window only includes LR04



**Figure 4.4:** Transfer entropy and predictive asymmetry on short and noisy time series. (A) Two time series with uncertainties in both age and value. The time series belong to a first-order autoregressive system of the first order (AR1) and are dynamically coupled in the direction  $X \rightarrow Y$ . (B) Left panel: Transfer entropy between the two time series in (A), computed for positive and negative lags. Right panel: Predictive asymmetry ( $\mathcal{A}$ ) between the time series computed from the transfer entropy values in the left panel. (C) Transfer entropy (left) and predictive asymmetry (right) for ten different time series realizations generated by sampling within the uncertainties in (A).

and RohSL as ice volume proxies (NBR-D results in appendix A.2). As mentioned in chapter 3, there exists no record of  $p\text{CO}_2$  prior to the MPT with sufficient resolution for the time scale of dynamics studied here. Therefore, my analyses of this time interval are limited to exploring the causal relationships between ice volume proxies, insolation, and Southern Ocean Fe flux.

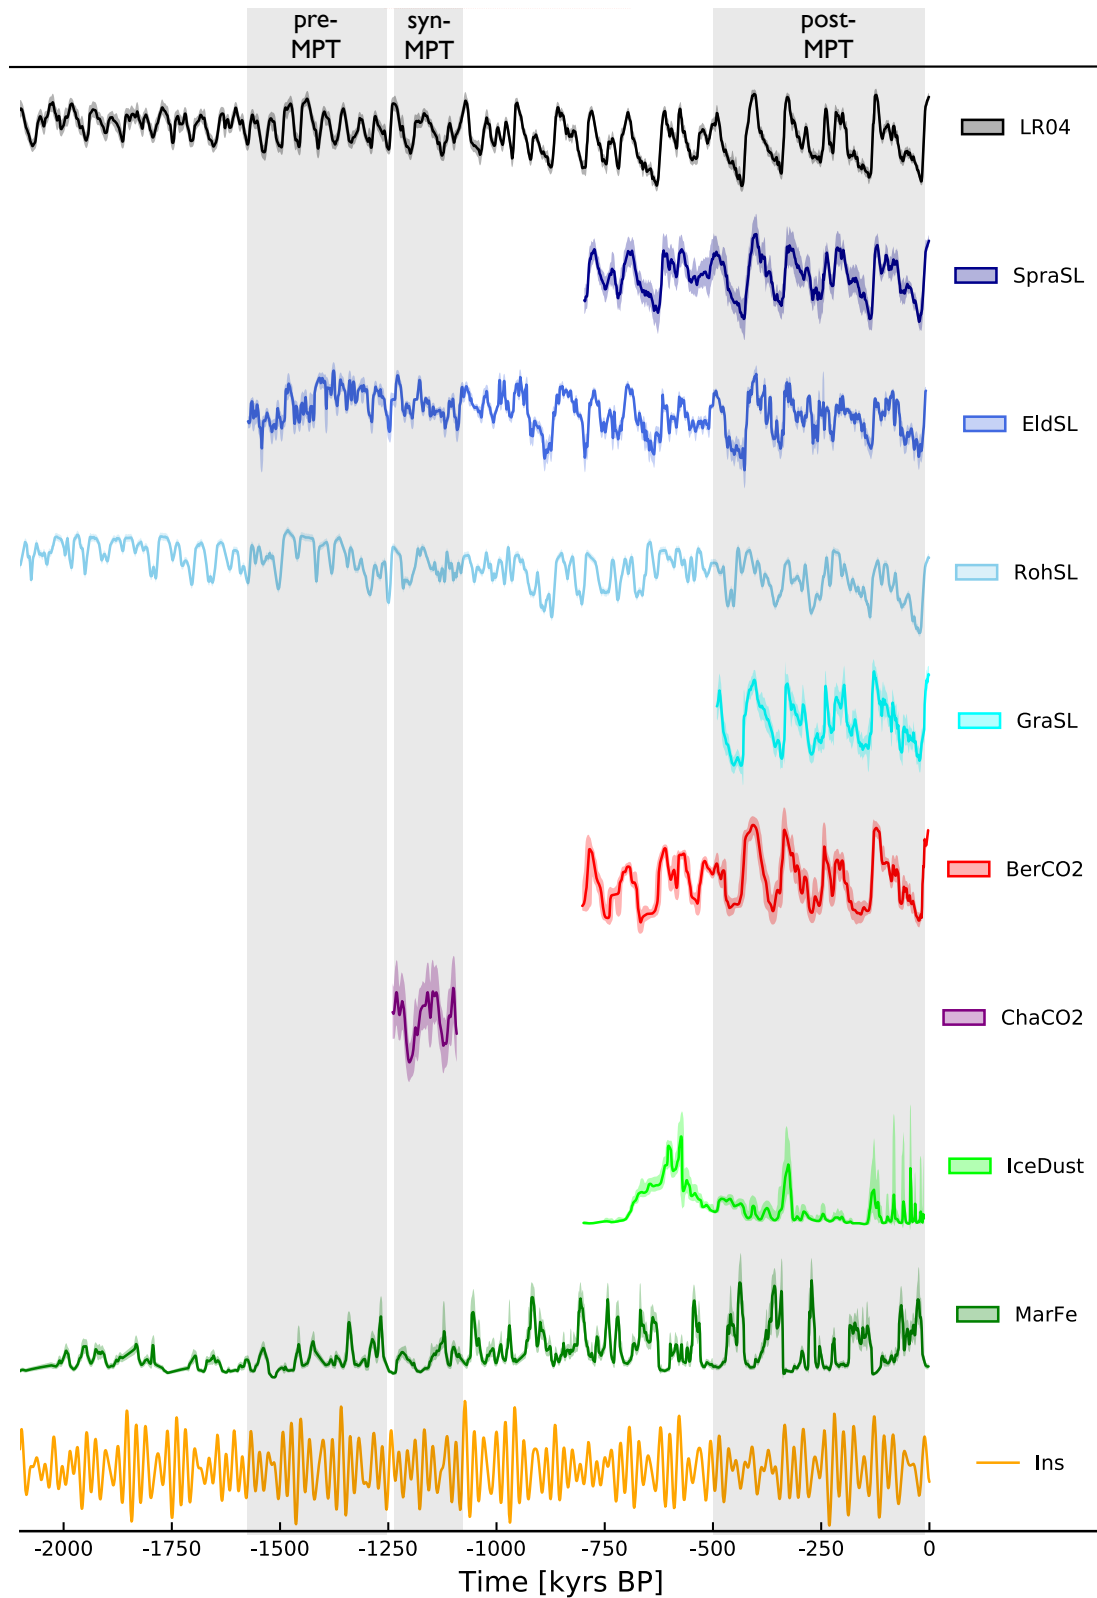
**Syn-MPT.** To analyse the climate system dynamics during the MPT, I have chosen the time interval from 1240-1092 ka BP, labeled syn-MPT, which is defined according to available high resolution data for  $p\text{CO}_2$  (ChaCO<sub>2</sub>). Note that this time interval is relatively short (149 data points at 1 kyr resolution), and spans only the early part of the MPT. Again, additional  $\mathbb{A}$  analyses over the full MPT time interval (1250-700 ka BP) are also computed for comparison, albeit without a suitable  $p\text{CO}_2$  record over this longer interval (NBR-E results in appendix A.4).

**Syn-MPT.** To analyse climate dynamics after the MPT, I have chosen the time window 492-13 ka BP, labeled post-MPT. This time interval is chosen to allow inclusion of the GraSL sea level record in the analyses, which has two main advantages. First, the GraSL record age model does not inherently assume any lagged response of ice volume to insolation, which is the case for the  $\delta^{18}\text{O}_b$ -based proxies. Second, the GraSL record has a far better temporal resolution than the other ice volume proxy records (125 yrs vs 1000 yrs). Additional  $\mathbb{A}$  analyses over a longer time interval (792-800 ka BP) are also computed for comparison, albeit without the GraSL record (NBR-F results in appendix A.6).

#### 4.4.2 Data wrangling

In notebook NB1, I collect the relevant data from global data repositories, and construct data sets combining average (mean or median) values with uncertainties. For handling the uncertainties in a streamlined way, I use the software package `UncertainData.jl` Haaga (2019), which offers seamless integration with the `CausalityTools.jl` software used for the causal analyses.

**Interpolation.** Unfortunately, hiatuses and irregular temporal sampling are typical characteristics of geological time series, which presents a challenge for methods based on delay reconstruction. For records with relatively high temporal resolution relative to the age model uncertainty, irregular sampling can be ameliorated through resampling. For records in which hiatuses are larger than the age model uncertainty, interpolation of missing data is a possible solution. Here I have interpolated values for all records that have gaps larger than 1 kyr in the original resolution (EldSL, BerCO<sub>2</sub>,



**Figure 4.5:** Definition of time windows and overview of time series. The empirical records are interpolated and resampled to a regular time grid with one value per 1 kyr. Uncertainties in age and value are incorporated through Monte Carlo analysis. Time series used in analyses are LR04 (Lisiecki and Raymo, 2005), SpraSL (Spratt and Lisiecki, 2016), EldSL (Elderfield et al., 2012), RohSL Rohling et al. (2014) and GraSL (Grant et al., 2014) as proxies of global ice volume; BerCO<sub>2</sub> (Bereiter et al., 2015) and ChaCO<sub>2</sub> (Chalk et al., 2017) representing  $p\text{CO}_2$ ; IceDust (Lambert et al., 2008) and MarFe (Martínez-García et al., 2011) representing Southern Hemisphere aeolian dust.



IceDust and MarFe). I have then used linear interpolation of mean and standard deviations ( $\sigma$ ) in both age and value and created new data arrays of mean and ( $\sigma$ ) containing values interpolated for every 100 yr. These interpolated arrays are in turn resampled on a coarser grid according to the same procedure as for the non-interpolated data, as described below. To guard against potential adverse effects of interpolation on the  $\mathcal{A}$  results, I have set the mean temporal resolution of each record's original sampling as the highest bound for the temporal bin resolution of the resampling grid. Thus, only the records that had a high temporal resolution (mean > 500 yrs) in the original sampling are included in the high resolution analyses.

**Redefining data as probability distributions.** `UncertainData.jl` provides a framework for handling of uncertainties in both age and value, through the data type *UncertainIndexValueDataset* (hereafter referred to as uivDs), which redefines the data as probability distributions by kernel density estimation (KDE), allowing for a convenient handling of data with uncertainties of different types. Each data point in the uivD consists of two probability distributions: one describing the value of the data point (e.g. sea-level in meters) and one describing the timing of the data point (age in kyr). Next, the uivD is resampled by Monte Carlo analysis, following the procedure detailed below.

**Binned resampling of time series data on a regular time grid.** To ensure that each data point in one time series  $X$  will represent the exact same time point as each data point in another time series  $Y$ , the probability distribution (KDE) of the observations are resampled on an equidistant time grid (hereafter referred to as binned resampling) using the *BinnedResampling* function from `UncertainData.jl`.

There is a trade-off to be made when deciding on a temporal bin size for resampling on the equidistant time grid. On the one hand, longer time bins (lower temporal resolution) on the final time series grid may obviate the need to interpolate data. On the other hand, the time grid has to be fine enough to preserve the dynamical information encoded in the time series. The choice of temporal bin size is also made as a trade-off between datasets: There is overall higher temporal resolution in the data spanning the post-MPT interval compared to the records spanning further back in time. As an overall compromise I have chosen a common bin size of 1 kyr that will work for all the records across all the time windows.

### 4.4.3 Predictive asymmetry analysis

To run the predictive asymmetry analyses, I have written a 'Toolbox' notebook (NB3 in attachments), which includes all the ingredients necessary to perform the predictive

asymmetry analyses in this thesis: I import all the libraries that are needed downstream, and read the wrangled time series data into the notebook. I have also written shorthand functions used in the analyses and results notebooks (NBRs). Specifically, the *computePredictiveAsymmetry* function implements the *VisitationFrequencyTest* and *PredictiveAsymmetryTest* from the *CausalityTools.jl* package, which are described in more detail below.

Each results notebook (NBR) is defined by its time window of analysis. In total, there are 7 NBRs; one notebook for each of the 6 time windows presented above, plus one NBR for the comparative analysis between the pre-MPT and post-MPT intervals, which is described in the last section of this chapter. In the NBRs, the first step is to select and cut the time series to the time window that will be analysed. Again, because transfer entropy operates with a lag-based definition of causality, it is crucial that the corresponding indexes in the binned resampled time series represent simultaneous time points, if not, the values computed for  $\mathcal{A}$  will not make any sense.

**Estimating transfer entropy between time series pairs.** The transfer entropy from source to target is estimated using the visitation frequency test, which takes four input arguments: a source time series  $X$ , a target time series  $Y$ , a bin size argument ( $\epsilon$ ) which will determine the partitioning of the delay reconstruction, and an array of prediction lags ( $\eta_s = -\eta_{max}:\eta_{max}$ ) for which we will compute transfer entropies. In the *VisitationFrequencyTest* function, the time series  $X$  and  $Y$  are used to make the delay reconstruction of the dynamics. Next, a rectangular grid is projected onto the delay reconstruction, with  $\epsilon$  defining the number of bins (i.e. the level of ‘coarse-graining’ of the state space). The *VisitationFrequencyTest* then counts how many times the delay reconstruction visits each partition in the grid, yielding an estimate of the invariant probability distribution over the grid. The estimated invariant distributions, together with the prediction lags  $\eta_s$  are then used to compute transfer entropy, according to eq. 4.2.

**Choice of bin size argument  $\epsilon$ .** In the same way that the number of bins in a histogram affects the shape of probability distribution that the histogram approximates, the estimation of the invariant probability distribution will be sensitive to the binning resolution of the delay reconstruction, which is given by  $\epsilon$ . I have chosen  $\epsilon$  according to the heuristic rule proposed by Paluš (1995), hereafter referred to as the ‘Palus horizon’:

$$\epsilon = \left[ N^{1/(D+1)} \right], \quad (4.5)$$

where  $N$  is the time series length,  $D$  is the embedding dimension, and  $[x]$  is the round-

ing operator which rounds the number  $x$  to the nearest integer. I have used the default embedding dimension of  $D = 3$  in all my analyses, following the recommendations of Haaga et al. (2020). Note that for very short time series, the Palus horizon is too coarse-grained and does not give reliable results (A.11). I have therefore decided to operate with a lower limit of  $\epsilon=4$  ( $4^3=64$  bins). In practice, this choice affects only the short syn-MPT interval. My preliminary sensitivity analysis suggests that the choice of  $\epsilon$  mainly affects the confidence interval of the  $\mathcal{A}$  results, but seems to have little effect on the inferred directionality of coupling (A.11).

**Choice of prediction lag range  $\eta_{max}$ .** The predictive asymmetry is computed by integrating over a spectrum of prediction lags  $\eta$  (eq. 4.3). In synthetic systems that are not periodic (e.g. AR1) the  $\mathcal{A}$  nicely converges with increasing  $\eta_{max}$ , but in systems with strong periodicity, the  $\mathcal{A}$  can also appear quasi-periodic (Haaga et al., 2020). To minimize the potential effects of large-scale (glacial-interglacial) periodicity on the predictive asymmetry, I decided to limit  $\eta_{max}$  to the equivalent of 20 kyrs (the scale of a glacial cycle). For time series that are resampled with a temporal bin resolution of 1 kyr, I thus use  $\eta_{max} = 20$ , and for time series resampled to a 500-yr and 125-yr temporal resolution I use  $\eta_{max}$  values of 40 and 160, respectively.

**Accounting for uncertainty.** Because of the uncertainty in the time series, the data are passed as `uivD` to the *PredictiveAsymmetryTest*, which then generates a distribution of  $\mathcal{A}$  values for each value of  $\eta$  (see Fig. 4.4c for an example with only ten realizations). In this way, we ensure that the uncertainty in the values and ages of the observed time series are propagated through the analyses. However, I add another layer of uncertainty to the analyses by randomly drawing sub-segments of the time series that are slightly shorter than the original (down to 70% of the total number of observations), using the *RandomSequencesTest*. I draw a family of 150 random segments, from which I extract the 0.025, 0.5 and 0.975 quantiles of the  $\mathcal{A}$  values, shown as 95% confidence ribbons on the result plots.

**Choice of significance level.** I normalize the predictive asymmetry according to eq. 4.4 using the *normalizePredictiveAsymmetry* function defined in NB3. As mentioned above, Haaga et al. (2020) propose  $f = 1$  as an acceptable threshold for most types of systems, as long as time series are above ca. 100-150 data points long. To be conservative, I have chosen to raise the significance threshold to  $f = 1.5$  for time series consisting of less than 200 observations, which affects only the syn-MPT window, where time series have 149 observations on the 1-kyr grid.

#### 4.4.4 Conventions for interpreting the results

- Positive  $\mathcal{A}$  above the significance threshold indicates a causal coupling in the direction  $\mathcal{A}_{source \rightarrow target}$ . To infer a causal coupling, I require that more than half of the random sequence iterations (i.e. the median  $\mathcal{A}$ ) must be above the significance threshold by  $\eta = \eta_{max}$ .
- Because the predictive asymmetry is normalized,  $\mathcal{A}$  is comparable across the analyses in terms of coupling strength, as long as the time series length and the parameters chosen for the analyses are the same.
- A complete divergence of the 95% confidence ribbons of the  $\mathcal{A}$  in either direction (i.e. no overlap) is taken to be unequivocal evidence of unidirectional coupling.
- If  $\mathcal{A}$  is not above the significance threshold by  $\eta = \eta_{max}$ , then there is not enough information in the data to infer a causal coupling between the time series on the given time scales. As always, the absence of evidence is not evidence of absence.

# Chapter 5

## Results

In this chapter I will present the results of the predictive asymmetry analyses between each of the observed climate system components in three different time intervals of the Pleistocene. At the end of the chapter, I also present results from a comparative analysis between pre-MPT and post-MPT intervals, where I standardize sampling to ensure that the records from the two time intervals have the same number of observations (notebook NBR-G attached to the thesis).

Before presenting the results, I will briefly explain my choice of wording in the description of the results. Strictly speaking, the results consist of values of the normalized predictive asymmetry  $\mathcal{A}$  relative to the built-in significance level, according to the criteria described in Section 4.4.4. However, to avoid repetition and in an attempt to increase the readability of this chapter, I have chosen to describe the results in terms of the presence and strength of evidence for causal coupling. This phrasing is analogous to the convention used in statistical correlation analyses, e.g. instead of repeatedly referring to the results as "a Spearman  $\rho$  value that exceeds the 95th percentile of a uniformly random null distribution", we simply refer to "a significant correlation". The predictive asymmetry is a statistical measure of causal coupling, not correlation, hence I will typically describe the results in terms of evidence for a significant causal coupling. In cases where the causal inference from  $\mathcal{A}$  is ambiguous I will briefly argue for the inference made, to avoid unnecessary repetition of results descriptions in the Discussion chapter.

## 5.1 Empirical evidence for causal connections in Pleistocene climate records

The results in this section address the first objective of the thesis, namely to quantify the directionality and relative strength of coupling in the Pleistocene climate system based on empirical evidence in the best available data. As described in Section 4.4, I have run separate analyses on three time windows in the Pleistocene, and the empirical evidence for causal connections varies somewhat for the different time windows. Before I describe the results for each time window in detail, I provide a graphical summary of all the analyses and causal inferences (Fig. 5.1).

### 5.1.1 The pre-MPT time window

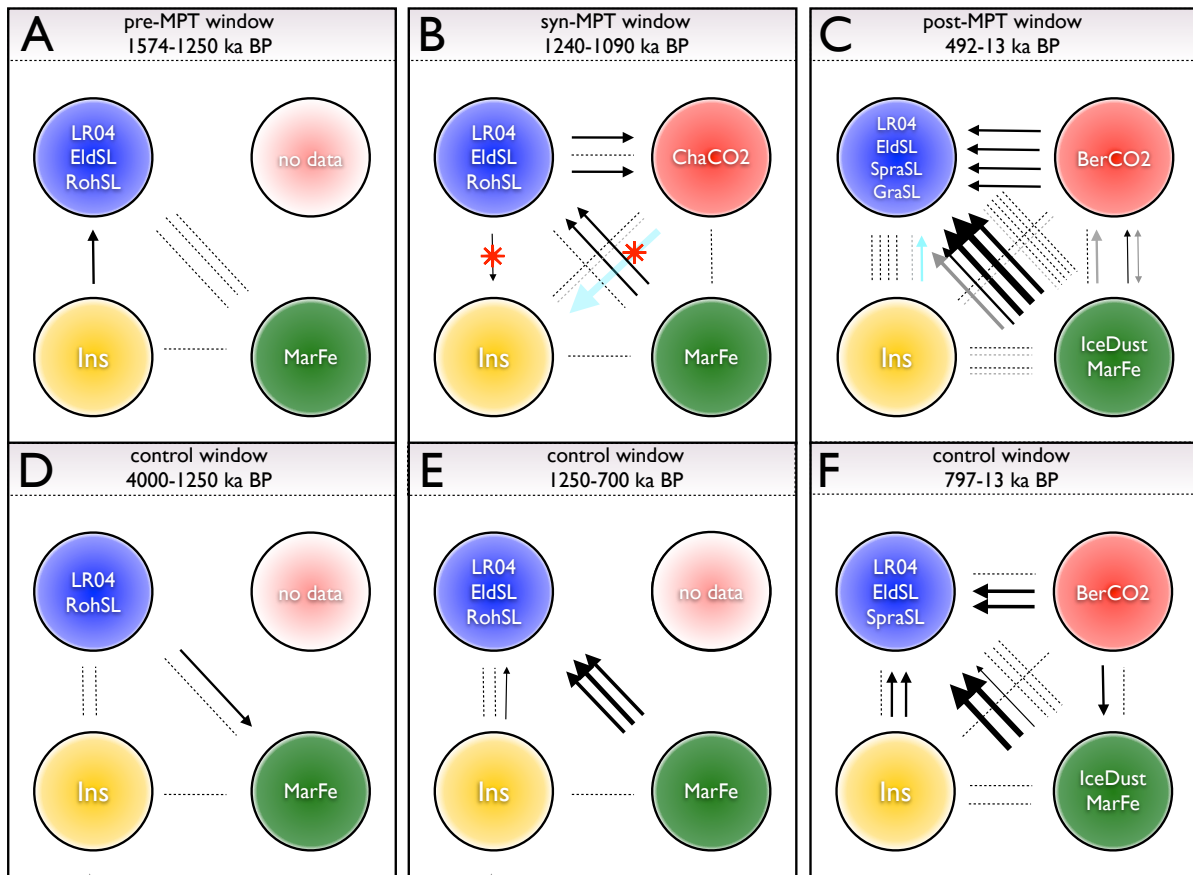
**Ice volume — insolation.** Analyses of the pre-MPT time window (1574-1250 ka BP) show a significant positive predictive asymmetry from Ins to LR04 for well over half of the random sequences (Fig. 5.2a). There is no significant  $\mathcal{A}$  between insolation and the pure GSL time series however (Fig. 5.2b and c). Over the longer pre-MPT control window (4-1.25 Ma BP), there is no empirical evidence in these records for a causal coupling between the insolation and ice volume time series (Appendix A.2a,b).

**Ice volume — Fe flux.** No causal coupling is inferred between ice volume and Fe flux in the pre-MPT, as the median  $\mathcal{A}$  between the time series does not cross the set significance threshold in either direction (Fig. A.1d-f). Over the longer pre-MPT control window (4-1.25 Ma BP), however, a positive  $\mathcal{A}$  from LR04 to MarFe suggests that information contained in the  $\delta^{18}\text{O}$  time series significantly helps predict the evolution of Fe flux, thus implying an underlying causal connection.

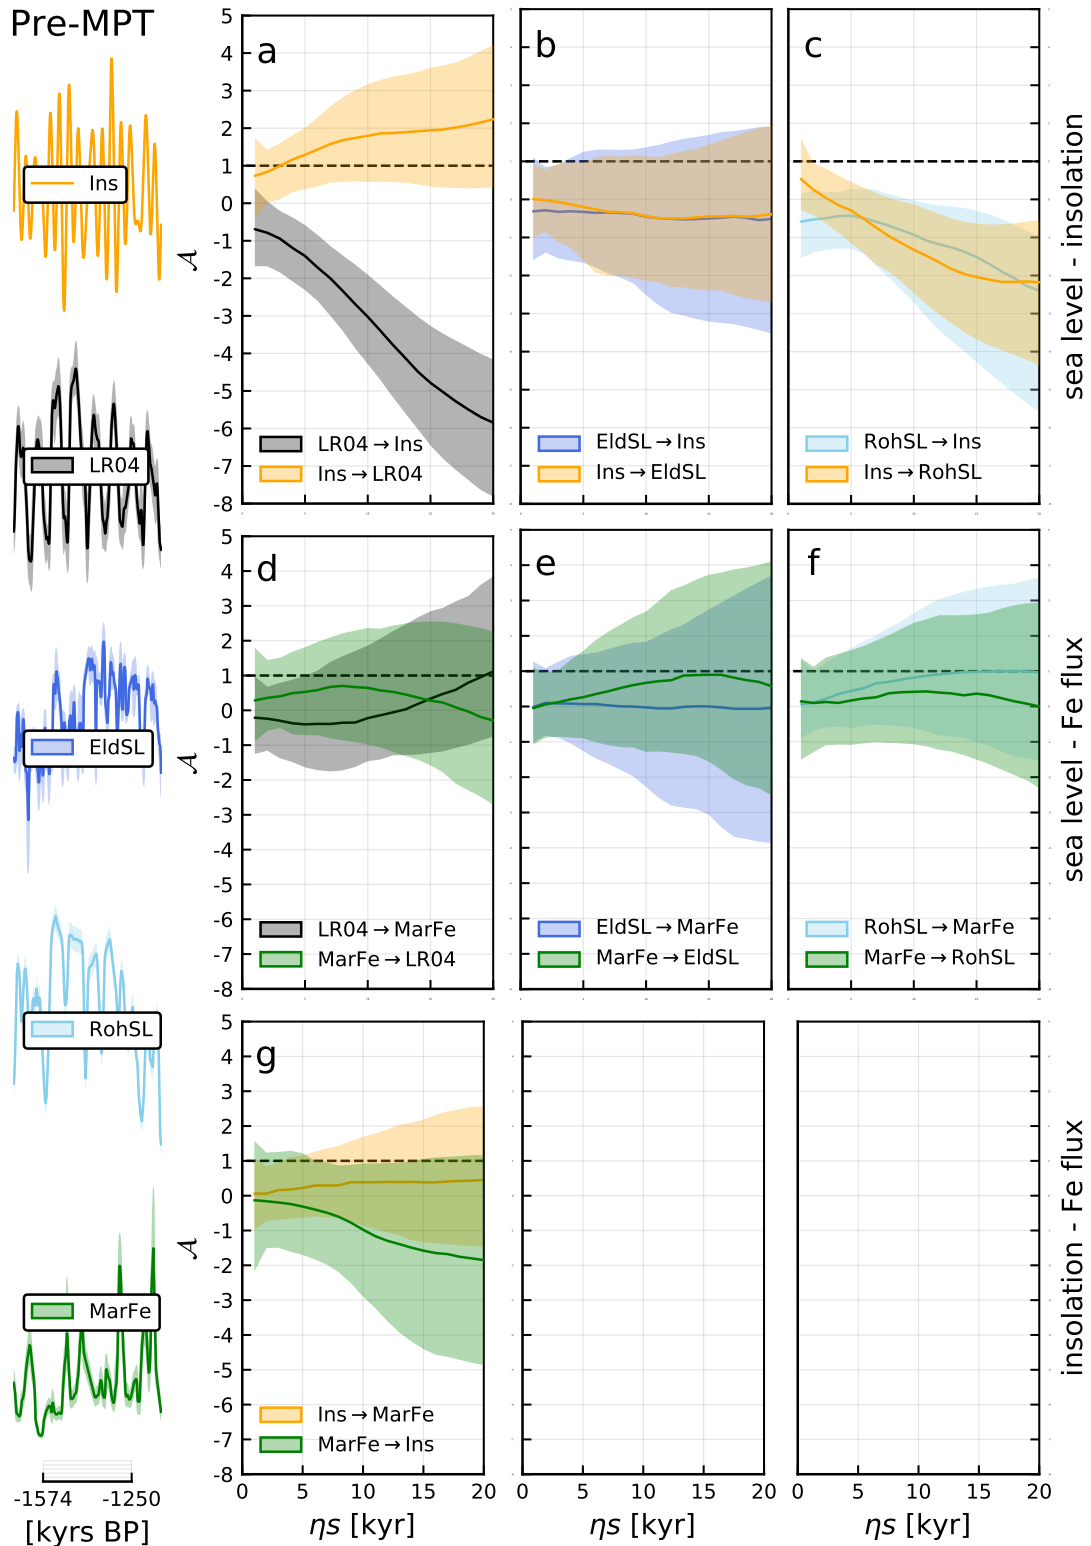
**Insolation — Fe flux.** Similarly, no causal coupling is inferred between the insolation and Fe flux time series (Fig. 5.2g) over the pre-MPT interval, as median  $\mathcal{A}$  between the time series does not cross the set significance threshold in either direction. Over the longer pre-MPT control window, the empirical evidence of an influence of insolation on Fe flux is somewhat more suggestive, but the median  $\mathcal{A}$  does not cross the significance threshold (appendix A.2a,b).

### 5.1.2 The syn-MPT time window

**Ice volume — insolation.** Over the syn-MPT window (1240-1090 ka BP), the predictive asymmetry is below the significance threshold for both LR04 and RohSL time series



**Figure 5.1: Overview of the inferred causal networks.** Time windows of analysis are chosen based on a hypothesized change in dynamical coupling in the climate system before and after the Mid-Pleistocene Transition, as well as according to time series overlap. The time windows analysed are 1574-1250 ka BP (A), 1240-1092 ka BP (B) and 492-13 ka BP (C) (labeled pre-MPT, syn-MPT and post-MPT, respectively), as well as three longer control windows (D, E, F). Time series are grouped according to which processes they are thought to represent: global ice volume (blue), northern hemisphere summer insolation (yellow), atmospheric CO<sub>2</sub> concentration (red), and atmospheric dust (green). See Fig. 4.5 for time series acronyms. Arrows represent the inferred causal connections, with the line thickness reflecting the relative magnitude of the predictive asymmetry. Stippled lines indicate negative results, i.e. where no significant coupling is detected between the time series. Line colours are used to distinguish the three different temporal resolutions used in binning the time series for analysis: 1 kyr (black), 500 yrs (gray) and 125 yrs (light blue). Red stars mark inferred causal couplings that are known to be false positives. The predictive asymmetry analyses used to make this overview can be studied in detail in the following figures: (A) 5.2; (B) 5.3 and A.9; (C) 5.4, 5.5 and A.10; (D) A.2; (E) A.4; (F) A.6.



**Figure 5.2: Predictive asymmetry results for the pre-MPT time window.** Left column: time series for the pre-MPT interval on a temporal bin resolution of 1 kyr. See Fig. 4.5 for time series acronyms. (a - g) Predictive asymmetry results. Values are medians and 95% ranges of  $\mathcal{A}$  values computed according to the analysis design described in chapter 4. Dashed line represents the significance threshold using a normalization factor of  $f = 1$  (eq. 4.4).



5.3a and c, respectively), hence there is no empirical support for a causal coupling between insolation and ice volume in this time window. Interestingly, median predictive asymmetry from EldSL to Ins is above the significance threshold (Fig. 5.3b), which we know to be a false positive, because a GSL forcing on insolation is non-sensical. This false positive persists despite taking the precautions of setting a lower bound on  $\epsilon = 4$ , and using the conservative significance threshold (raising the normalization factor  $f = 1.5$ ). Possible reasons for this false positive will be discussed in the next chapter. Meanwhile, I note that over the longer control time interval (1250-700 ka BP), predictive asymmetry is below the significance interval for all ice volume proxies, implying that any dynamical information in these time series do not support a causal coupling between ice volume and insolation in this time window (Fig. A.4a-c).

**Ice volume —  $p\text{CO}_2$ .** Over the syn-MPT window, we find significant predictive asymmetry from ice volume proxies to  $p\text{CO}_2$  with the LR04 and RohSL time series (Fig. ??d and f, respectively). In the case of the EldSL time series, however, there is no evidence for a causal coupling (Fig. 5.3e). There is no suitable  $p\text{CO}_2$  record to confirm the results over the longer control interval.

**Ice volume — Fe flux.** There is also significant positive predictive asymmetry from Fe flux to ice volume, where both  $\mathcal{A}_{\text{MarFe} \rightarrow \text{LR04}}$  and  $\mathcal{A}_{\text{MarFe} \rightarrow \text{EldSL}}$  cross the significance threshold in a large majority of the random sequences (Fig. 5.3g and h, respectively).  $\mathcal{A}_{\text{MarFe} \rightarrow \text{RohSL}}$  crosses the significance threshold at  $\eta=20$  in just under half of the iterations (Fig. 5.3i), and is therefore just short of the significance threshold set for causal inference, although the overall pattern of predictive asymmetry matches that of the other ice volume proxies. The same analyses performed over the longer control interval give unambiguous results of a causal coupling from Fe flux to ice volume, with most or all of the 95% confidence envelope rising above the significance threshold for all three ice volume proxies (Fig. A.4d-f, appendix A).

**Fe flux —  $p\text{CO}_2$ .** The predictive asymmetry between ChaCO2 and MarFe is below the significance threshold for the syn-MPT-interval (Fig. 5.3j), implying that no causal coupling between Fe flux and  $p\text{CO}_2$  can be inferred from the data in this time window. No equivalent analysis is possible over the longer control interval, due to the lack of a suitable  $p\text{CO}_2$  record.

**Fe flux — insolation.** The predictive asymmetry between MarFe and insolation is below the significance threshold for the syn-MPT-interval (Fig. 5.3l), meaning there is not sufficient dynamical evidence in the time series to infer a causal coupling between these components. The same result holds for the longer control interval (Fig. A.4g, appendix A).

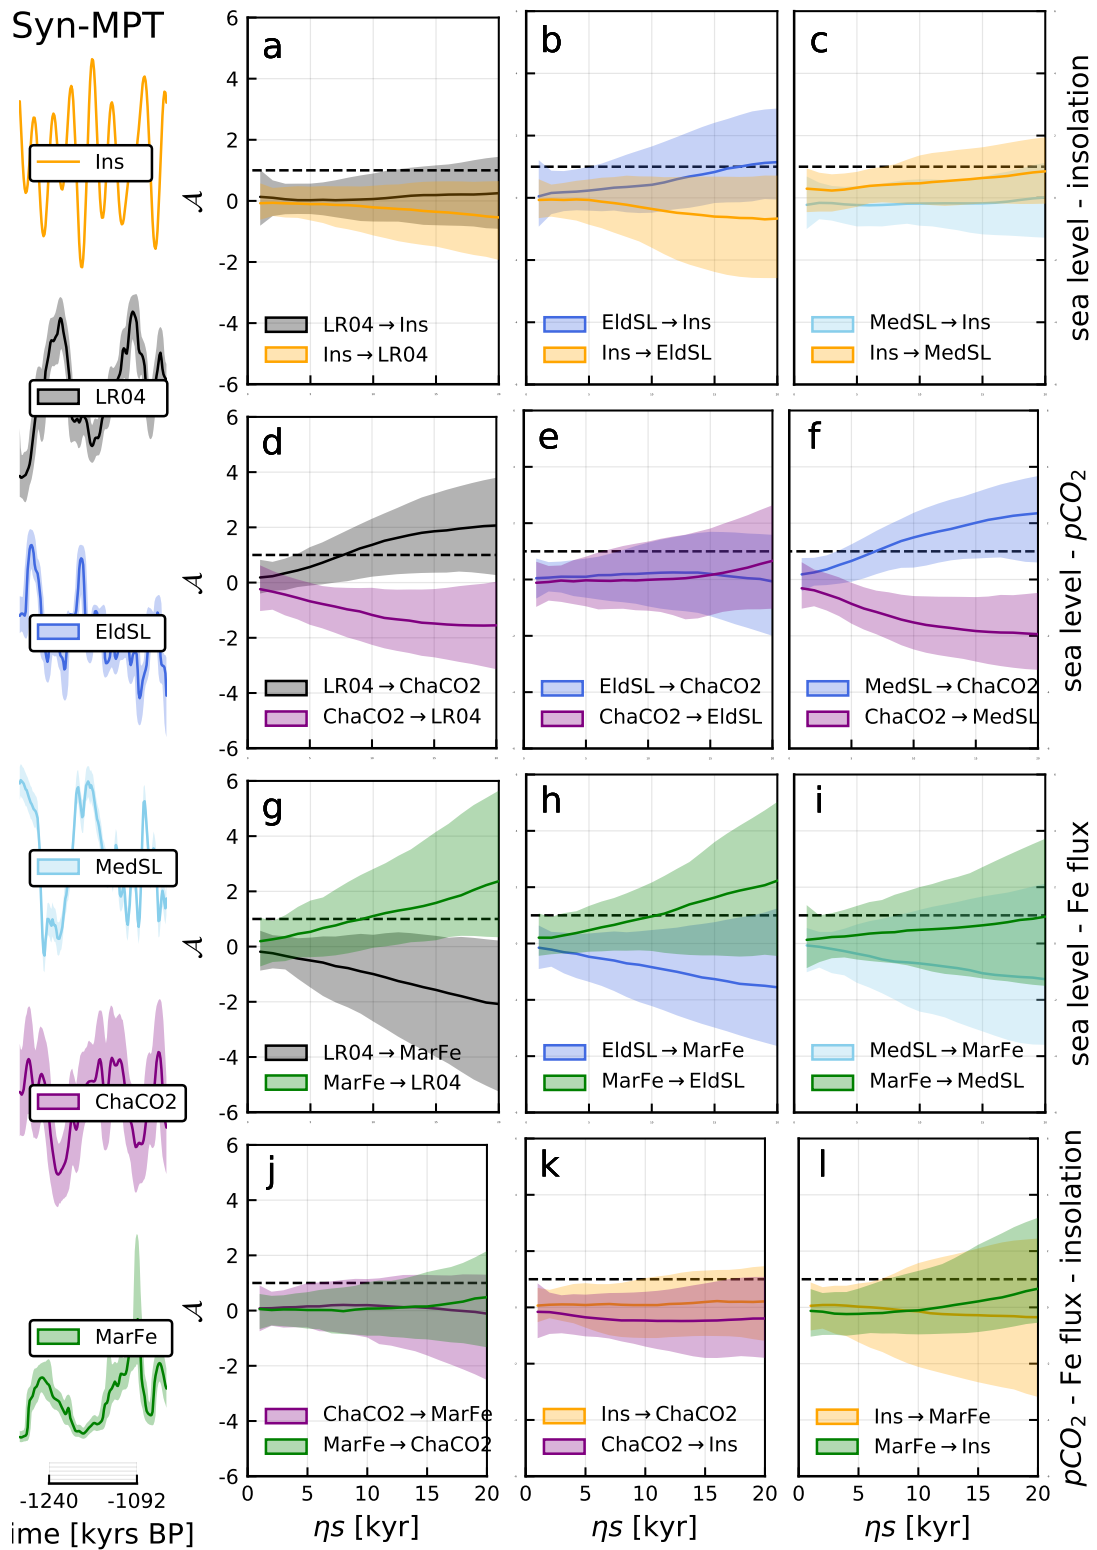
**Insolation —  $p\text{CO}_2$ .** Predictive asymmetry between ChaCO<sub>2</sub>-Ins time series indicates no significant coupling when computed on a 1 kyr time bin resolution (Fig. 5.3k). However, the same analysis computed on a higher-resolution binning of 125 yrs gives a median predictive asymmetry above the significance threshold, (A.9c) implying a causal influence from  $p\text{CO}_2$  to insolation. Again, we know this is a false positive, and the possible reasons behind this result will be discussed in the next chapter.

### 5.1.3 The post-MPT time window

**Ice volume — insolation.** Using a time bin resolution of 1 kyr, the median predictive asymmetry does not cross above the significance level for any of the ice volume records (Fig. 5.4a-d). The same result holds for the GraSL record on a binning resolution of 500 yr (Fig. 5.5h). These findings suggest that there is not sufficient dynamical evidence for a causal coupling between insolation and ice volume. However, for a higher resolution analysis with a temporal bin duration of 125 years, there is a positive predictive asymmetry from Ins to GraSL with a median  $\mathcal{A}$  that crosses above the significance threshold (Fig. 5.5c, appendix A). This result implies that if the temporal bin resolution is sufficiently high (125 yr bins), then there is empirical evidence for insolation as a forcing of ice volume dynamics. Note, however, that the predictive asymmetry decreases after about  $\eta=40$  and dips below the significance threshold again at  $\eta=60$ .

**Ice volume —  $p\text{CO}_2$ .** Median predictive asymmetry is significant from BerCO<sub>2</sub> to all the ice volume proxy time series in analyses with a time bin resolution of 1 kyr (Fig. 5.4e-h). The higher resolution analysis between GraSL and BerCO<sub>2</sub> with a temporal bin duration of 500 years also indicates a coupling from  $p\text{CO}_2$  to ice volume (Fig. 5.5e). This agreement between independent ice volume proxies, with independent age models, provide robust evidence that  $p\text{CO}_2$  is a driver of ice volume dynamics over the post-MPT time interval. Common to all the  $p\text{CO}_2$ — ice volume analyses over the post-MPT window is  $\mathcal{A}$  values of around 1-2 by 20 prediction lags. Among the records analyzed here, only Fe flux has a stronger relative forcing on ice volume.

**Fe/dust flux —  $p\text{CO}_2$ .** Predictive asymmetry between  $p\text{CO}_2$  and Fe/dust records analysed on 1 kyr time bin resolution approaches the significance threshold with both the MarFe and the IceDust time series (Fig. 5.4s and t, respectively), implying that the information in these records at this resolution is not sufficient to determine any clear causal connection. A higher resolution analysis (500-yr bins) finds predictive asymmetry that reaches just above the significance threshold in both directions, BerCO<sub>2</sub> → MarFe at  $\sim \eta = 20$ , and MarFe → BerCO<sub>2</sub> at  $\sim \eta = 40$ , with confidence intervals largely



**Figure 5.3: Predictive asymmetry results for the syn-MPT time window.** Left column: time series for the syn-MPT interval on a temporal bin resolution of 1 kyr. See Fig. 4.5 for time series acronyms. (a - l) Predictive asymmetry results. Values are medians and 95% ranges of  $\mathcal{A}$  values computed according to the analysis design described in chapter 4. Dashed line represents the significance threshold using a normalization factor of  $f = 1.5$  (eq. 4.4).

overlapping, (Fig. A.8e). This result could imply a bidirectional coupling between Fe flux and  $p\text{CO}_2$ , but it could also stem from chronological uncertainties (i.e. age uncertainties may obscure a lead-lag relationship). The IceDust record can help resolve this issue, because the IceDust and BerCO2 records are derived from the same ice core, which greatly reduces the relative age uncertainties between dust and  $p\text{CO}_2$  dynamics. Indeed, the predictive asymmetries between BerCO2 and IceDust clearly diverge, suggesting a clear directionality of coupling from IceDust to BerCO2 (Fig. 5.5b), implying that Fe/dust dynamics is the driver and  $p\text{CO}_2$  changes is the response.

**Insolation —  $p\text{CO}_2$ .** Predictive asymmetry is insignificant between Ins and BerCO2, both in analysis with 1 kyr time bins (Fig. 5.4r) and 500 yr bins (Fig. 5.5i), implying that there is no evidence in the data of a coupling between  $p\text{CO}_2$  and insolation.

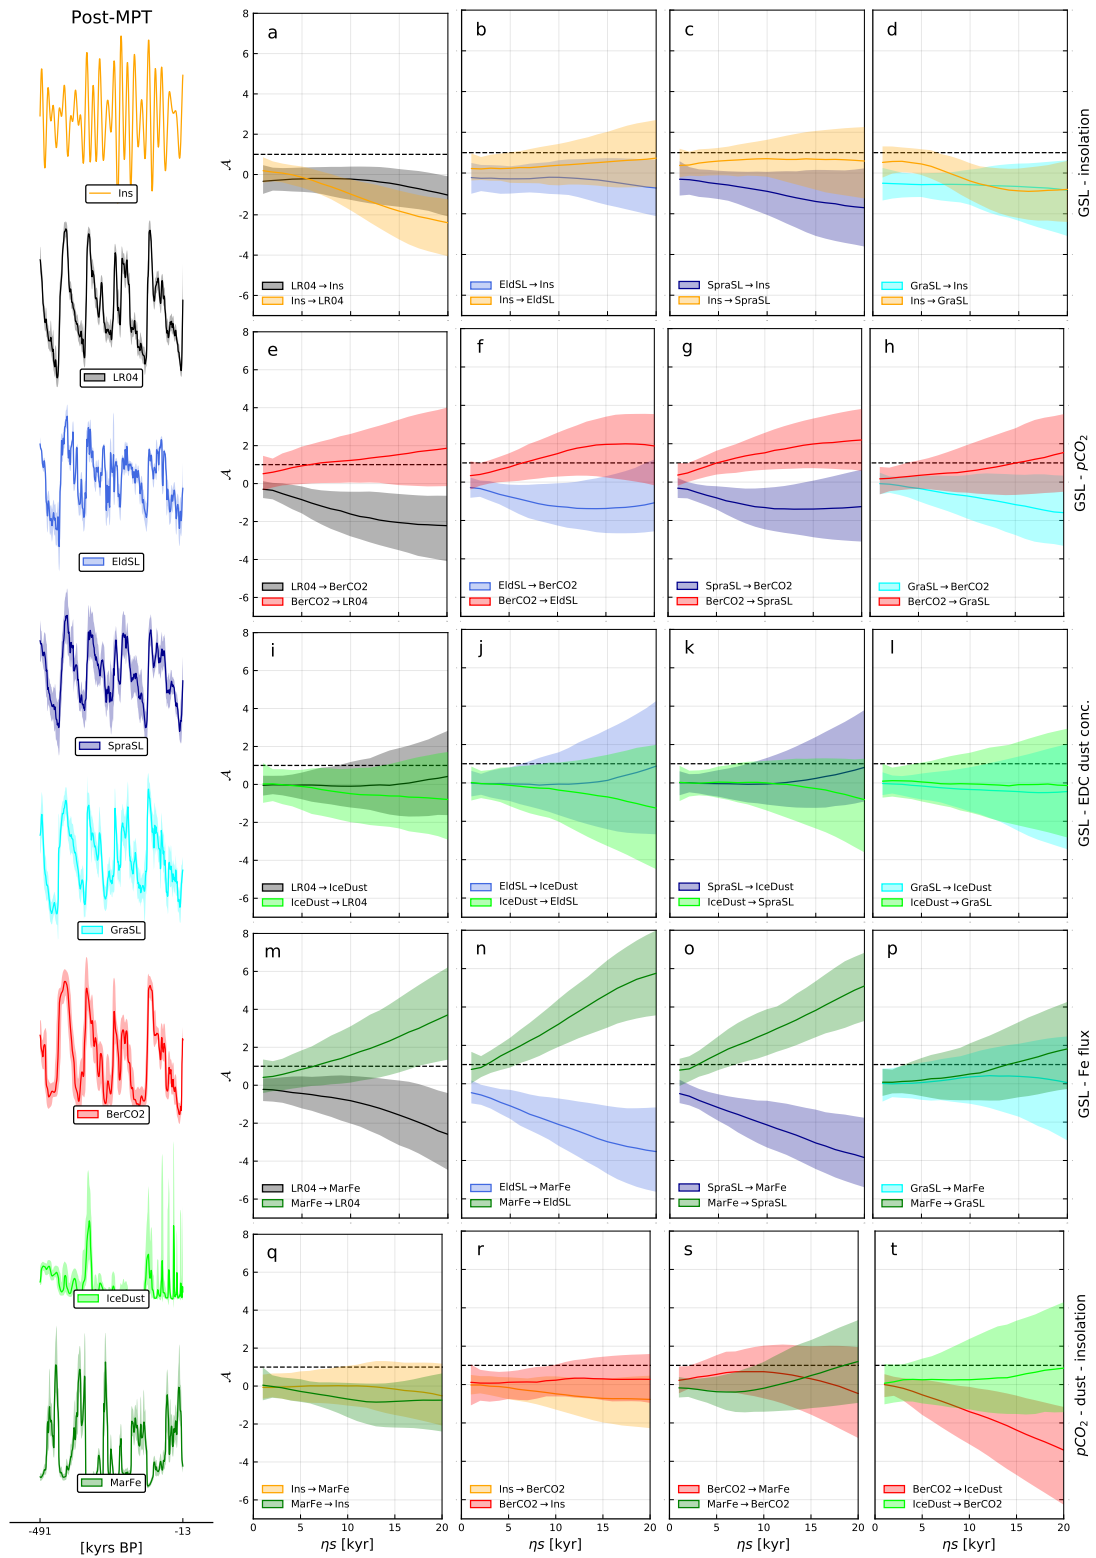
**Insolation — dust.** Similarly, predictive asymmetry is insignificant between Ins with MarFe and IceDust, both in analysis with 1 kyr prediction lag (Fig. 5.4q) and 500 yr prediction lag (Fig. 5.5c,f), implying that there is no evidence in these records of a coupling between Fe/dust flux and insolation.

## 5.2 Sampling-standardized comparison of the pre- and post-MPT dynamics

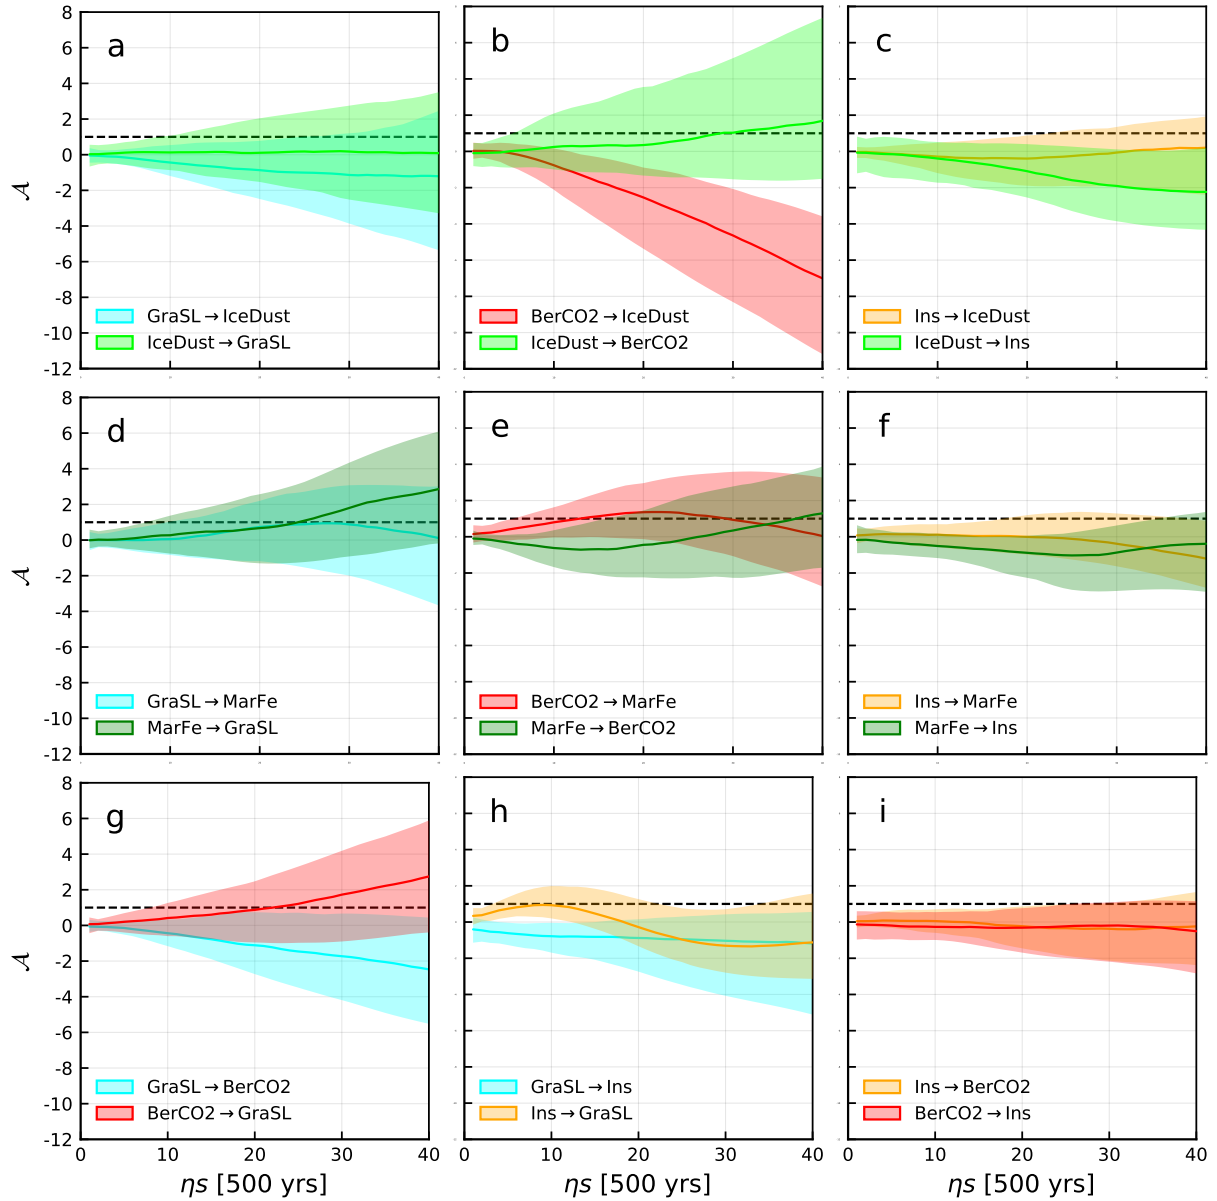
In the previous section I presented causal couplings inferred between different components of the climate system analyzed in different time windows. However, any inferred differences between the time windows are possibly confounded by the fact that the duration of the time windows (and hence the number of observations) and the system components available for analysis are not the same.

I have therefore set up a second analysis to better distinguish between differences that might arise from unequal sampling and differences that might arise from an underlying change in the system dynamics. This section presents the results of this sampling-standardized comparison between the pre-MPT and post-MPT, which addresses the second objective of this thesis: Is there evidence in the data that the causal coupling of the climate system has *changed* across the MPT?

In this analysis I have only used time series that span both the post-MPT and pre-MPT windows, including LR04, EldSL, MarFe, and Ins. These records have the same underlying assumptions and roughly the same resolution in both post- and pre-MPT time intervals. Hence, differences in data quality are primarily due to differences in time series length, i.e the number of available data points ( $N$ ). The pre-MPT window



**Figure 5.4: Predictive asymmetry results for the post-MPT time window at 1 kyr resolution.** Left column: time series for the post-MPT interval on a temporal bin resolution of 1 kyr. See Fig. 4.5 for time series acronyms. (a - t) Predictive asymmetry results. Values are medians and 95% ranges of  $\mathcal{A}$  values computed according to the analysis design described in chapter 4. Dashed line represents the significance threshold using a normalization factor of  $f = 1.0$  (eq. 4.4).



**Figure 5.5: Predictive asymmetry results for the post-MPT time window at 500 yr resolution.** Left column: time series for the post-MPT interval on a temporal bin resolution of 500 yrs. See Fig. 4.5 for time series acronyms. (a - i) Predictive asymmetry results. Values are medians and 95% ranges of  $\mathcal{A}$  values computed according to the analysis design described in chapter 4. Dashed line represents the significance threshold using a normalization factor of  $f = 1.0$  (eq. 4.4).

has a shorter time series length as compared to the post-MPT window. Time series length is a key factor in the statistical robustness of  $\mathcal{A}$ , as shown by sensitivity analyses in Haaga et al. (2020).

To make the analyses of post- and pre-MPT windows more equitable, and thus allow for a comparison of the results, I ran  $\mathcal{A}$  analyses with the common time series (LR04, EldSL, MarFe, Ins) and the same bin resolution (1 kyr), but now with random sequences of equal length. I chose to set the length of the random sequences to the same length as the pre-MPT interval ( $N = 325$ ). I used a set of 150 random sequences of length  $N=325$ , yielding an ensemble of 150 delay reconstructions for transfer entropy estimation and computation of predictive asymmetries. For the post-MPT window, the 150 sequences were chosen at random within the larger post-MPT window ( $N = 479$ ). For the pre-MPT window, each random sequence spanned the same window (1574-1250 kaBP), but different realizations were obtained by drawing randomly within the uncertainties associated with the time series, thus yielding different transfer entropy estimates in each iteration.

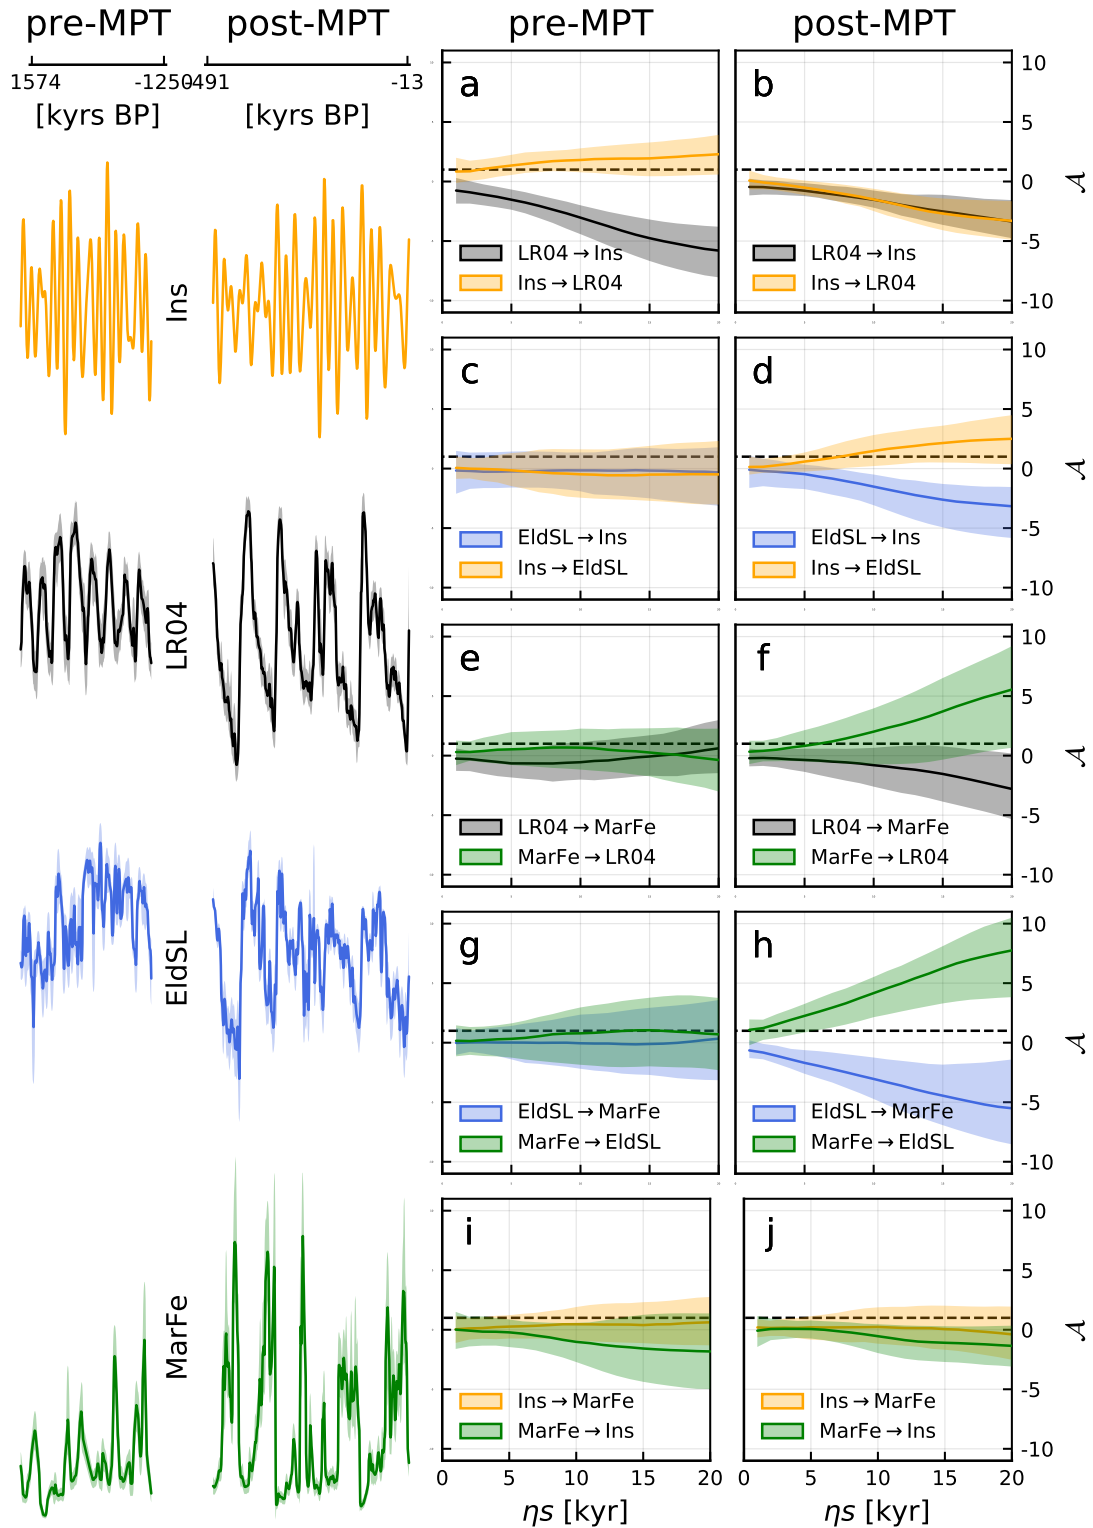
Where predictive asymmetries of pre-MPT and post-MPT overlap, the information of causal coupling is equal between the two periods. Where the predictive asymmetries are different, this means that information flow between the parameters have changed between the two time intervals.

**Ice volume — insolation.** Analysis of LR04 and Ins shows that empirical evidence for a causal influence from the insolation time series to ice volume proxy is present in the pre-MPT interval, but not in the post-MPT interval (Fig. 5.6a and b, respectively). Analysis of EldSL and Ins, on the other hand, detects a coupling from the insolation time series to the ice volume proxy only for the post-MPT interval, while no causal coupling can be detected in the pre-MPT interval (Fig. 5.6d and c, respectively). There are similarities between the results where a causal coupling is inferred - both diverge after just a few prediction lags, and reaches a prediction skill of around 2 by 20  $\eta$ s (Fig. 5.6a and d), indicating a comparable coupling strength.

**Ice volume — Fe flux.** In the case of LR04 and MarFe, no coupling is detected in the pre-MPT window, but a causal coupling is inferred from MarFe to LR04 is present in the post-MPT window (Fig. 5.6e and f, respectively). The same change in inferred causal interactions is found in analyses of MarFe and EldSL, with no coupling detected by  $\mathcal{A}$  in the pre-MPT window, but there is strong support for a causal coupling from MarFe to EldSL in the post-MPT window (Fig. 5.6g and h respectively). Given that the analyses are computed in the same way, and with sampling standardization, these results provide compelling evidence for a significant change in the causal role of

Southern Ocean Fe flux in the climate system across the MPT.





**Figure 5.6: Sampling-standardized comparison of the pre- and post-MPT dynamics.** Left column: time series for the pre- and post-MPT intervals on a temporal bin resolution of 1 kyr. The post-MPT time series are constrained to match the number of observations in the pre-MPT window ( $N=325$ ). See Fig. 4.5 for time series acronyms. (a - l) Predictive asymmetry results. Values are medians and 95% ranges of  $\mathcal{A}$  values computed according to the analysis design described in chapter 4. Dashed line represents the significance threshold using a normalization factor of  $f = 1.0$  (eq. 4.4).

# Chapter 6

## Discussion

### 6.1 Implications for our understanding of Pleistocene climate system interactions

**The role of Southern Ocean Fe flux.** Dynamical evidence in the records I have analyzed suggest that some of the climate system interactions changed across the MPT. Southern ocean dust flux emerged as a key player affecting changes in global sea level after the MPT (Fig. 5.6e-h). Prior to the MPT, the causal relationship between dust and glacial-interglacial fluctuations was either not significant (Fig. 5.6e,g), or possibly of opposite directionality (Fig.A.2c), i.e. changes in ice volume influenced changes in dust flux. The roles seem to have shifted already during the MPT (Figs. A.3g-i and A.4d-f). In the post-MPT time window,  $\mathcal{A}$  results imply that Southern Ocean Fe flux was the most important climatic forcing among the system components considered here.

In light of the mechanistic hypotheses outlined in Chapter 2, these results may support the notion that the relative importance of feedback mechanisms between climate and dust shifted in the MPT. More specifically, the feedbacks between dust and climate prior to the MPT may have been more dominated by glacial climates impact on dust flux, where a generally cold and dry glacial climate, less vegetation cover and strong winds would increase the atmospheric dust flux (Shaffer and Lambert, 2018). In contrast, climate-dust feedbacks in the opposite direction could become more prominent in the post-MPT, perhaps triggered by the nonlinear amplification of atmospheric dust load induced by the increasingly deeper glacial climates. This amplification could tip the balance in favour of dust-driven mechanisms, for example through increased global cooling due to aerosol scattering of incoming solar radiation (e.g. Tegen and Lacis, 1996), or through indirect interactions with biogeochemical cycles (Martin, 1990).

In principle, the analyses performed in my thesis could help quantify the empirical support for this latter hypothesis of increased glacial dust flux increasing marine productivity. However, according to the relative coupling strength inferred from  $\mathcal{A}$  results, the dust forcing on ice volume was substantially greater than the  $p\text{CO}_2$  forcing in the post-MPT interval. Dynamical information in the available records thus indicate that dust-climate feedbacks other than increased marine productivity (e.g. aerosol scattering or cloud formation) were important factors determining the nature of the Late Pleistocene glacial cycles.

**The role of  $p\text{CO}_2$ .** My results also strongly suggest that atmospheric  $p\text{CO}_2$  was an important forcing of ice volume changes after the MPT (5.4e-h), although the causal relationship might have been of opposite directionality during the MPT (5.3d,f). The lack of high resolution, extensive records of  $p\text{CO}_2$  limits what can be inferred of causal couplings in the climate system prior to the MPT. Atmospheric  $p\text{CO}_2$  has been implicated in hypotheses of glacial-interglacial sea level change ever since the strong covariation of the two variables became known two decades ago (Petit et al., 1999; Shackleton, 2000). Nevertheless, the causal relationship between the two has remained controversial (Abe-Ouchi et al., 2013b). The findings presented in this thesis show that there is significant information in the data supporting the causal hypothesis that  $p\text{CO}_2$  was an important driver of glacial-interglacial changes in the Late Pleistocene climate system.

With the time series available in this study, I have only tested one of the proposed mechanisms by which  $p\text{CO}_2$  itself may plausibly have been controlled, namely sequestration of carbon from the atmospheric reservoir by increased marine productivity, if enabled by Fe fertilization of the Southern Ocean. The relationship between  $p\text{CO}_2$  and Southern Ocean Fe flux is difficult to resolve when taking into account the uncertainties associated with the data. I do, however, find that when minimizing age uncertainties by analysing Antarctic dust and  $p\text{CO}_2$  from the same core, and choosing higher-resolution time bins, there is dynamical information in the Antarctic aeolian dust record that helps predict changes in  $p\text{CO}_2$ . Here, the inferred coupling is unidirectional (A.8e).

Given the good agreement between the Southern Ocean marine sediment Fe record and the Antarctic ice core dust record (MarFe and IceDust), my expectation would be that this direction of coupling would also be found between Southern ocean dust flux and  $p\text{CO}_2$  if the records had tighter age constraints. If so, and if the IceDust record actually represents a regionally coherent atmospheric dust signal as proposed by Martínez-García et al. (2011), then such a finding would provide evidence in support of the hypothesis that Fe fertilization increased marine productivity during glacial periods,

leading to a drawdown of  $p\text{CO}_2$  that was a forcing factor of Late Pleistocene changes in ice volume. I stress, however, that this argument assumes the equivalence of the Antarctic dust *concentration* record with the Southern ocean dust *flux* record, which is not straightforward, as pointed out by Shaffer and Lambert (2018).

**The role of insolation.** Whether or not the predictive asymmetry detects a causal relationship from the insolation time series to the ice volume proxy records appears to depend the time window selected and on the temporal binning resolution. Causal coupling is detected over the last 800 kyrs (Fig. A.6b-c), but not over the last 500 kyrs (Fig. 5.4a-d). Interestingly, this time-window sensitivity is not obviously related to the number of observations, because  $\mathcal{A}$  values indicate causal coupling on the shorter (N=325 data points) pre-MPT time interval (Fig. 5.2), but not on the longer (N=2750 data points) pre-MPT control interval (Fig. A.2). The possible role of age model differences is discussed in the next section.

Higher-resolution analysis detects a coupling from Northern Hemisphere insolation to global ice volume (Fig. A.10c). This preliminary sensitivity analysis of the effect of temporal binning suggests that a time bin resolution of 1 kyr is not the ideal resolution to resolve the dynamics between insolation and ice volume. It might also suggest that the coupling between insolation and ice volume operates with shorter ( $\sim$ centennial) response times, and I will thus argue that if higher resolution ice volume proxy records were available they might help resolve a coupling between insolation and ice volume further back in time (e.g. in syn-MPT and pre-MPT intervals).

**Limitations of the pre-MPT time window** There is information in the data suggesting ice volume forcing of Southern ocean dust flux in the longer control time window analysis (Fig. A.2c), but not in the shorter pre-MPT window analysis (Fig. 5.2d-f). As discussed below, this result is consistent with the expectation that a sufficiently long window of observation is needed for the empirical delay reconstruction to capture the dynamics. However, this expectation does not explain the finding that Northern Hemisphere insolation helps predict changes in  $\delta^{18}\text{O}$  values in the shorter time window (Fig. 5.2a), but not in the longer control time window (Fig. A.2a). It is also notable that I do not detect this causal coupling when using the RohSL sea level reconstruction, only the LR04  $\delta^{18}\text{O}$  record. As mentioned in Chapter 3, the RohSL record has its issues both with possible tectonic instability in early parts of the record, and with the presence of sapropel intervals, both of which might render the RohSL record less representative of ice volume dynamics. A third option, discussed below, is that the temporal binning resolution is too coarse to capture the relevant time scales. Overall, however, with the time series used in these analyses, there is limited dynamical information to support

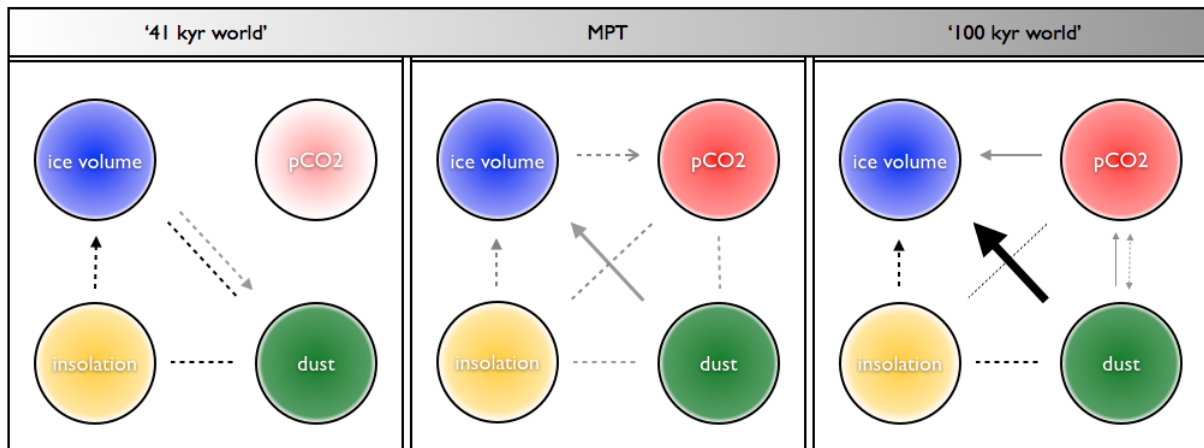
the inference of causal couplings in the pre-MPT time interval.

**Limitations of the syn-MPT time window** The false positives found in analyses of the syn-MPT time window indicate that one should be careful in accepting the results for this time interval. In the case of the EldSL — Ins result (Fig. 5.3b), it is plausible that the time series length is too short for the  $\mathcal{A}$  test to be statistically robust, as the syn-MPT window only contains 149 observations when using 1 kyr time bins. Furthermore, the random sequences used to compute the predictive asymmetry further sub-sample down to 70% of this time series length, leaving the  $\mathcal{A}$  computation at the very lower limit of what is required according to comprehensive sensitivity analyses of synthetic systems (Haaga et al., 2020). Because the statistical robustness of  $\mathcal{A}$  is system dependent (e.g. dimensionality, nonlinearity, chaos, noise properties) we have no exact estimate of the lower data limit for real systems of unknown complexity.

Should we then reject all couplings inferred for this short time window? An argument can be made that the  $\mathcal{A}$  results showing a causal coupling from ChaCO<sub>2</sub> and MarFe to the ice volume records are qualitatively more trustworthy, because the predictive asymmetries fully diverge, with non-overlapping confidence ribbons (Fig. 5.3d, f and g). This is not the case for the borderline significant false positive between EldSL-Ins, which just barely crosses the significance threshold. Furthermore, the Fe flux forcing of ice volume is confirmed by the same analyses performed over the longer control interval (Fig. A.4d-f). The tentative conclusion that  $p\text{CO}_2$  and Fe flux acted as a forcing on ice volume in the syn-MPT interval would arguably also be somewhat more robust, given that the couplings are detected with two out of the three ice volume proxy records.

The false positive between ChaCO<sub>2</sub> — Ins on the other hand, is unlikely to be the result of insufficient time series length, because with time bins of 125 years the time series consisted of 1192 observations, with random sequences down to a minimum of 834 data points, well above the lower data limit for all systems tested in Haaga et al. (2020). Instead, there are two possible culprits of this false positive. Firstly, the ChaCO<sub>2</sub> time series was resampled by Chalk et al. (2017) and published with a resolution of 125 years, although the original observations had one data point every 3.5-4.5 kyrs, i.e. an order of magnitude lower (Chalk et al., 2017). Although the effects cannot be assessed without knowing the resampling technique and running sensitivity tests on  $\mathcal{A}$ , it is not unlikely that this degree of oversampling may have introduced biases in the ChaCO<sub>2</sub> record.

The other possibility is that the syn-MPT time window simply is not wide enough to capture the relevant dynamics of the system (regardless of temporal resolution and



**Figure 6.1: Qualitative summary of causal inferences.** Arrows represent the inferred causal connections, with the line thickness reflecting the relative magnitude of the predictive asymmetry. Black arrows/lines show results that are comparable across time intervals. Grey arrows/lines can only be compared within the respective time windows. Stippled arrows show causal couplings that are only detected under specific conditions (only one time series, or sensitive to time window or temporal bin). Solid lines show that the signal is detected across many different records, suggesting a more robust conclusion.

number of observations). An insufficient window of observation would imply that the delay reconstruction from the time series would be unable to capture the dynamics of the system in such a way that causal couplings can be determined. If this were the case then it would compromise all causal inferences made by  $\mathcal{A}$  over the syn-MPT time interval. I will therefore interpret the results of the syn-MPT interval with caution, and focus instead on the results for the longer control interval (Fig. A.4). Over this longer syn-MPT control interval, there is no dynamical evidence of coupling between insolation and ice volume (Fig. A.4a-c), but there is solid dynamical evidence in the data that Southern Ocean Fe flux was a driver of global ice volume during the MPT (Fig. A.4d-f).

## 6.2 Implications for the use of Predictive Asymmetry on paleoclimate records

As this is the first study that makes extensive use of the novel predictive asymmetry method on empirical data and ‘real world’ questions where causal couplings are unknown, it is of particular relevance to reflect upon the insights gained concerning the method and its further development.

Using an model-free approach like the  $\mathcal{A}$  to infer causality makes two assumptions: Firstly, it must be assumed that the dynamics of the observed system has settled on

some typical behaviour, or global attractor. Secondly, it must be assumed that the system has been observed long enough to be able to describe this behaviour, such that a delay reconstruction is able to characterise relevant aspects of the invariant properties of the dynamics. Only then can the visitation frequency estimator approximate the true probability distribution of the system's states and provide accurate estimates of the transfer entropy. It follows from this second point that a sufficiently long window of observation and a sufficient number of observations within that window are crucial for  $\mathcal{A}$  to yield statistically robust causal inferences.

**Window of observation.** Although sensitivity analyses on synthetic systems have shown the  $\mathcal{A}$  test to be remarkably robust to sampling, with only a few hundred observations being sufficient for many complex systems, there is another aspect of sampling that may be as important in real-world applications, namely the duration of the window of observation. Even if we are safe in assuming that the system under study has settled on its attractor, we need to observe the system over enough time to adequately sample the state space for delay reconstruction. For example, the syn-MPT time window may not fully capture the relevant dynamics for some of the processes involved, as suggested by the visible trend in the syn-MPT EldSL record, which may only represent a fragment of longer-term oscillations. Hence, the false positive EldSL — Ins result disappears when the analysis is performed on the extended MPT control interval (1250-700 ka BP; Appendix A.4).

Originally, the pre-, syn-, and post-MPT time intervals were chosen for the analysis design of this study, based on the overlap of the available paleoclimate records. However, the additional analyses performed to explore the effect of sampling on the robustness of results has led to the insight that not only the number of observations, but also the choice of time window has an impact on what causal relationships can be inferred by the predictive asymmetry (I here refer to the seemingly ephemeral results of causal coupling between ice volume records and insolation both before and after the MPT).

**Age models.** Assuming causal mechanisms when constructing the age models (e.g. LR04) naturally poses a caveat when using a lag-based causal analysis framework such as the predictive asymmetry. A predictive asymmetry analysis previously done on some of the same time series used in this thesis (i.e. SpraSL, GraSL and Ins) has attributed differences in causal inference between different ice volume proxies to the effect of age model assumptions (Haaga et al., 2020). Haaga et al. (2020) found significant  $\mathcal{A}$  from Ins to the orbitally tuned SpraSL over the past 800 kyrs, while  $\mathcal{A}$  from Ins to the orbitally independent GraSL over the past 500 kyrs was found to be insignificant. They attributed this result to bias in the SpraSL age model, which inherently

assumes a lagged response of sea level to orbital forcing. What I find in this thesis work, however, is that analyzing SprasL and the other ice volume proxies used herein over the same time intervals yield consistent results of orbital forcing over the past 800 kyrs and no orbital forcing over the past 500 kyrs, which rejects the argument that the differing results are due to different assumptions in the age models. Instead, I will argue that further work is needed to assess the impact of choice of time window on the outcome of the predictive asymmetry. However this does not take away from the admonition that age model construct should be taken into special consideration when using lag-based causal inference methods.

**Temporal binning, interpolation and over-sampling.** Geological records are often riddled with hiatuses and irregular sampling resolution. In order to apply a method such as the  $\mathcal{A}$  that requires time series sampled on a regular time grid, geological records will often require some kind of regularization by resampling or interpolation. I have chosen to run analyses mainly using a temporal bin resolution of 1 kyr, because this was more compatible with the resolution of the time series used here. However, I have found that running the analyses with shorter prediction lags for the time series where higher resolution data is available can in some cases give further revelations of causal coupling (e.g. fig. A.10).

One may question what would be the most meaningful temporal resolution for a causal analysis. As is well known, the Earth system and the climate system operate on many different time scales. Consider, for example, atmospheric  $p\text{CO}_2$ , which on tectonic timescales is regulated by erosion and uplift, while on seasonal or even diurnal time scales  $p\text{CO}_2$  fluctuates in response to the respiration of plants. For dynamics that operate across very large time scales, the time series used here will not be long enough to contain the relevant dynamical information. At the other extreme, where the response time is shorter than the temporal bin duration used in the analysis, we intuitively expect that the relevant dynamics is not captured. However, when there is memory in the system (expressed as autocorrelation in the records), dynamics that operate faster than the time bin duration may still be picked up.

Arguably, although the predictive asymmetry is a data-driven method which in theory requires no initial model, the a priori choices made concerning temporal binning and the range of prediction lags  $\eta_{max}$  will have implications for what causal relationships can be uncovered. The analysis design will thus inevitably carry some 'model' assumptions about the system to be studied. Choices regarding temporal binning and  $\eta_{max}$  will inevitably be based on what process we wish to study and our understanding of what time scale it operates on. To assess the effect of making such choices, one



solution could be to compute the predictive asymmetry over an array of different time bin resolutions, as far as our data resolution allows.

Furthermore, if we assume that the false positive result between ChaCO<sub>2</sub> and Ins in the syn-MPT is not due to an inadequate window of observations, then a possible explanation is the fact that at a resolution of 125 yrs, the ChaCO<sub>2</sub> record provided by Chalk et al. (2017) is highly over-sampled. The time series resolution used to compute the predictive asymmetry here is an order of magnitude higher than the original data, which raises a warning flag for data-driven causal analysis.

Further work is needed to better understand how data interpolation and over-sampling affect the predictive asymmetry, which is outside the scope of this thesis, but I offer some suggestions in the next section.

### 6.3 Way forward

**Quantifying the robustness of our results with regards to the number of observations.** How many observations are required to go from saying the results are sensitive to sampling to claiming the results are robust? This would be an important goal for further work. One approach would be to do a sensitivity analysis on the most robust post-MPT results with the same analysis by gradually down-sampling the time series.

**Understanding the effect of uncertainties.** For time series with associated uncertainties, which is typically the case in geological records, a computational framework for dealing with uncertainties and propagating them into the analysis has already been made available (Haaga, 2019), which I have used in this thesis. However, further work is needed to understand how these uncertainties affect causal inferences, especially when trying to reconstruct causal networks among variables with different levels of uncertainty.

**Investigating the effect of interpolation.** Short time series and discontinuous sampling are inherent problems in geological records. Time series length is probably the most important determinant of the robustness of dynamical system reconstruction of the Earth system. Furthermore, the temporal resolution determines the time scales of the dynamics to be studied, but naïve interpolation may have unintended consequences for data-driven causal analysis. A systematic approach is needed to quantitatively assess the effect of time series interpolation on the predictive asymmetry.

**Lessons learned from false positives.** The false positives draw extra attention to where there is a need for further work in understanding the robustness of  $\mathcal{A}$ . In partic-

ular, understanding how incomplete delay reconstruction due to an insufficient window of observation would affect  $\mathcal{A}$  (would it sometimes detect a coupling that isn't there?) could be investigated by sensitivity analyses on synthetic systems. Likewise, further work is needed to push the boundaries of uncertainty handling by resampling (e.g. how much resampling is over-sampling?).

**Better understanding of how the choice of time intervals may affect  $\mathcal{A}$ .** We have seen that slight changes in the choice of time interval may have an impact in whether or not a causal coupling is detected (e.g. analyses between Ins and ice volume proxy records in different time windows). As this does not necessarily relate to the number of observations or data quality, further work is needed to understand how to make informed choices for analysis design, and to assess the impact of choice of time intervals on the outcome of the predictive asymmetry.

## Chapter 7

# Conclusions

In this thesis I have used a novel causal analysis framework on state-of-the-art paleoclimate records to uncover empirical evidence for causal connectivity among some of the key components of the Pleistocene climate system. Specifically, the dynamical information in the time series suggests that atmospheric  $p\text{CO}_2$  and Southern Ocean dust flux were important factors influencing ice volume changes in the Late Pleistocene. In contrast, there is limited support for external insolation forcing, at least on the time scales I have targeted here. Furthermore, a comparative analysis of pre-MPT and post-MPT time windows suggests that the strength and directionality of some of the causal interactions in the climate system changed across the MPT. Specifically, Southern Ocean dust seems to have emerged as a forcing factor during the MPT and became a significant forcing of ice volume in the Late Pleistocene. These findings can help constrain the number of hypotheses currently competing to explain the rise of the  $\sim 100$ -year periodicity of Late Pleistocene glacial cycles. Moreover, this data-driven approach may be valuable for assessing the level of dynamical information contained in available data, and for guiding mechanistic modeling efforts. In order to further quantify the relative empirical support for different proposed causal mechanisms, I suggest including proxy records for ocean circulation, which has been implicated in deep-ocean carbon storage.

This thesis is the first to use the predictive asymmetry method on a comprehensive set of paleoclimate records since it was developed by Haaga et al. (2020). The ability to determine the strength and directionality of causal interactions directly from observed time series, without specifying a model, is remarkable, but it also places tough demands on the observational records and on the analysis design. In my experience, the two most prominent challenges to making robust causal inferences are associated with data pre-processing (temporal binning, resampling, and interpolation), and with

the assumption that the time window of observation is sufficiently long to capture the dynamics of the system. If future developments of the analytical framework can help overcome some of these challenges, then the method can expand our toolkit for both testing and building hypotheses on causal interactions in the complex Earth system.

## References

- Abe-Ouchi, A., F. Saito, K. Kawamura, M. E. Raymo, J. Okuno, K. Takahashi, and H. Blatter (2013a). Insolation-driven 100,000-year glacial cycles and hysteresis of ice-sheet volume. *nature* 500(7461), 190–193.
- Abe-Ouchi, A., F. Saito, K. Kawamura, M. E. Raymo, J. Okuno, K. Takahashi, and H. Blatter (2013b, jul). Insolation-driven 100,000-year glacial cycles and hysteresis of ice-sheet volume. *Nature* 500(7461), 190–193.
- Arrhenius, S. (1896). Xxxi. on the influence of carbonic acid in the air upon the temperature of the ground. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 41(251), 237–276.
- Babtie, A. C., P. Kirk, and M. P. Stumpf (2014). Topological sensitivity analysis for systems biology. *Proceedings of the National Academy of Sciences* 111(52), 18507–18512.
- Bazin, L., A. Landais, B. Lemieux-Dudon, H. Toyé Mahamadou Kele, D. Veres, F. Parrenin, P. Martinerie, C. Ritz, E. Capron, V. Y. Lipenkov, M.-F. Loutre, D. Raynaud, B. M. Vinther, A. M. Svensson, S. O. Rasmussen, M. Severi, T. Blunier, M. C. Leuenberger, H. Fischer, V. Masson-Delmotte, J. A. Chappellaz, and E. W. Wolff (2013). The Antarctic ice core chronology (AICC2012).
- Bereiter, B., S. Eggleston, J. Schmitt, C. Nehrbass-Ahles, T. F. Stocker, H. Fischer, S. Kipfstuhl, and J. Chappellaz (2015). Revision of the epica dome c co2 record from 800 to 600 kyr before present. *Geophysical Research Letters* 42(2), 542–549.
- Chalk, T. B., M. P. Hain, G. L. Foster, E. J. Rohling, P. F. Sexton, M. P. Badger, S. G. Cherry, A. P. Hasenfratz, G. H. Haug, S. L. Jaccard, et al. (2017). Causes of ice age intensification across the mid-pleistocene transition. *Proceedings of the National Academy of Sciences* 114(50), 13114–13119.
- Chalk, T. B., M. P. Hain, G. L. Foster, E. J. Rohling, P. F. Sexton, M. P. S. Badger, S. G. Cherry, A. P. Hasenfratz, G. H. Haug, S. H. Jaccard, A. Martínez-García,

- J. Crespin, R. D. Pancost, and P. A. Wilson (2017). Early Mid-Pleistocene Transition (MPT) carbon dioxide from ODP Site 165-999. PANGAEA. In supplement to: Chalk, Thomas B; Hain, Mathis P; Foster, Gavin L; Rohling, Eelco J; Sexton, Philip F; Badger, Marcus P S; Cherry, Soraya G; Hasenfratz, Adam P; Haug, Gerald H; Jaccard, Samuel L; Martínez-García, Alfredo; Pälike, Heiko; Pancost, Richard D; Wilson, Paul A (2017): Causes of ice age intensification across the Mid-Pleistocene Transition. *Proceedings of the National Academy of Sciences*, 114(50), 13114-13119, <https://doi.org/10.1073/pnas.1702143114>.
- Chappell, J. and N. Shackleton (1986). Oxygen isotopes and sea level. *Nature* 324(6093), 137–140.
- Clark, P. U. and D. Pollard (1998a, feb). Origin of the Middle Pleistocene Transition by ice sheet erosion of regolith. *Paleoceanography* 13(1), 1–9.
- Clark, P. U. and D. Pollard (1998b). Origin of the middle pleistocene transition by ice sheet erosion of regolith. *Paleoceanography* 13(1), 1–9.
- Delmas, R. J., J.-M. Ascencio, and M. Legrand (1980). Polar ice evidence that atmospheric  $\text{CO}_2$  20,000 yr bp was 50% of present. *Nature* 284(5752), 155–157.
- Deyle, E. R. and G. Sugihara (2011). Generalized theorems for nonlinear state space reconstruction. *Plos one* 6(3), e18295.
- Elderfield, H., P. Ferretti, M. Greaves, S. Crowhurst, I. N. McCave, D. Hodell, and A. M. Piotrowski (2012). Evolution of Ocean Temperature and Ice Volume Through the Mid-Pleistocene Climate. *New Series* 337(6095), 704–709.
- Emiliani, C. (1955). Pleistocene temperatures. *The Journal of Geology* 63(6), 538–578.
- Esmark, J. (1826). Remarks tending to explain the geological history of the earth. *The Edinburgh New Philosophical Journal* 2, 107–121.
- Farmer, J., B. Hönisch, L. Haynes, D. Kroon, S. Jung, H. Ford, M. Raymo, M. Jaime-Seguí, D. Bell, S. Goldstein, et al. (2019). Deep atlantic ocean carbon storage and the rise of 100,000-year glacial cycles. *Nature Geoscience* 12(5), 355–360.
- Fourier, J. (1824). Remarques générales sur les températures du globe terrestre et des espaces planétaires. In *Annales de Chimie et de Physique*, Volume 27, pp. 136–167.
- Fuks, D., O. Ackermann, A. Ayalon, M. Bar-Matthews, G. Bar-Oz, Y. Levi, A. Maeir, E. Weiss, T. Zilberman, and Z. Safrai (2017). Dust clouds, climate change and coins:

- consilience of palaeoclimate and economy in the late antique southern levant. *Levant* 49(2), 205–223.
- Grant, K., E. Rohling, C. B. Ramsey, H. Cheng, R. Edwards, F. Florindo, D. Heslop, F. Marra, A. Roberts, M. E. Tamisiea, et al. (2014). Sea-level variability over five glacial cycles. *Nature communications* 5(1), 1–9.
- Haaga, K. A. (2019). Uncertaindata.jl: a julia package for working with measurements and datasets with uncertainties. *Journal of Open Source Software* 4(43), 1666.
- Haaga, K. A., D. Diego, J. Brendryen, and B. Hannisdal (2020). A simple test for causality in complex systems.
- Hannisdal, B. and L. H. Liow (2018). Causality from palaeontological time series. *Palaeontology* 61(4), 495–509.
- Hays, J. D., J. Imbrie, and N. J. Shackleton (1976). Variations in the earth’s orbit: Pacesetter of the ice ages. *Science* 194(4270), 1121–1132.
- Hemming, N. and G. Hanson (1992). Boron isotopic composition and concentration in modern marine carbonates. *Geochimica et Cosmochimica Acta* 56(1), 537–543.
- Hlaváčková-Schindler, K., M. Paluš, M. Vejmelka, and J. Bhattacharya (2007). Causality detection based on information-theoretic approaches in time series analysis. *Physics Reports* 441(1), 1–46.
- Hönisch, B., N. G. Hemming, D. Archer, M. Siddall, and J. F. McManus (2009). Atmospheric carbon dioxide concentration across the mid-pleistocene transition. *Science* 324(5934), 1551–1554.
- Huybers, P. and E. Tziperman (2008). Integrated summer insolation forcing and 40,000-year glacial cycles: The perspective from an ice-sheet/energy-balance model. *Paleoceanography* 23(1).
- Kantz, H. and T. Schreiber (2004). *Nonlinear time series analysis*, Volume 7. Cambridge university press.
- Lambert, F., B. Delmonte, J.-R. Petit, M. Bigler, P. R. Kaufmann, M. A. Hutterli, T. F. Stocker, U. Ruth, J. P. Steffensen, and V. Maggi (2008). Dust-climate couplings over the past 800,000 years from the epica dome c ice core. *Nature* 452(7187), 616–619.
- Laskar, J., P. Robutel, F. Joutel, M. Gastineau, A. Correia, and B. Lestrade (2004). A long-term numerical solution for the insolation quantities of the earth. *Astronomy & Astrophysics* 428(1), 261–285.

- Lear, C. H., K. Billups, R. E. Rickaby, L. Diester-Haass, E. M. Mawbey, and S. M. Soudan (2016). Breathing more deeply: Deep ocean carbon storage during the mid-pleistocene climate transition. *Geology* 44(12), 1035–1038.
- Lisiecki, L. E. and M. E. Raymo (2005). A pliocene-pleistocene stack of 57 globally distributed benthic  $\delta^{18}\text{O}$  records. *Paleoceanography* 20(1).
- Martin, J. H. (1990). Glacial-interglacial  $\text{CO}_2$  change: The iron hypothesis. *Paleoceanography* 5(1), 1–13.
- Martin, J. H., S. E. Fitzwater, and R. M. Gordon (1990). Iron deficiency limits phytoplankton growth in antarctic waters. *Global Biogeochemical Cycles* 4(1), 5–12.
- Martinez-Garcia, A. (2015). Iron fertilization of the subantarctic ocean during the last ice age. *AGUFM 2015*, A21L–02.
- Martínez-García, A., A. Rosell-Melé, S. L. Jaccard, W. Geibert, D. M. Sigman, and G. H. Haug (2011). Southern ocean dust–climate coupling over the past four million years. *Nature* 476(7360), 312–315.
- Milankovitch, M. (1941). *Kanon der Erdebestrahlung und seine Anwendung auf das Eiszeitenproblem*. Königlich Serbische Akademie.
- Neftel, A., H. Oeschger, J. Schwander, B. Stauffer, and R. Zimbrunn (1982). Ice core sample measurements give atmospheric  $\text{CO}_2$  content during the past 40,000 yr. *Nature* 295(5846), 220–223.
- Paillard, D. (2015). Quaternary glaciations: from observations to theories. *Quaternary Science Reviews* 107, 11–24.
- Paillard, D., L. Labeyrie, and P. Yiou (1996). Macintosh program performs time-series analysis. *Eos, Transactions American Geophysical Union* 77(39), 379–379.
- Paluš, M. (1995). Testing for nonlinearity using redundancies: Quantitative and qualitative aspects. *Physica D: Nonlinear Phenomena* 80(1-2), 186–205.
- Parrenin, F., J.-M. Barnola, J. Beer, T. Blunier, E. Castellano, J. Chappellaz, G. Dreyfus, H. Fischer, S. Fujita, J. Jouzel, et al. (2007). The edc3 chronology for the epica dome c ice core. *Climate of the Past* 3(3), 485–497.
- Pena, L. D. and S. L. Goldstein (2014). Thermohaline circulation crisis and impacts during the mid-pleistocene transition. *Science* 345(6194), 318–322.



- Petit, J.-R., J. Jouzel, D. Raynaud, N. I. Barkov, J.-M. Barnola, I. Basile, M. Bender, J. Chappellaz, M. Davis, G. Delaygue, and Others (1999). Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature* 399(6735), 429.
- Raymo, M. E. and P. Huybers (2008). Unlocking the mysteries of the ice ages. *Nature* 451(7176), 284–285.
- Rohling, E., G. L. Foster, K. Grant, G. Marino, A. Roberts, M. E. Tamisiea, and F. Williams (2014). Sea-level and deep-sea-temperature variability over the past 5.3 million years. *Nature* 508(7497), 477–482.
- Runge, J., S. Bathiany, E. Bollt, G. Camps-Valls, D. Coumou, E. Deyle, C. Glymour, M. Kretschmer, M. D. Mahecha, J. Muñoz-Marí, et al. (2019). Inferring causation from time series in earth system sciences. *Nature communications* 10(1), 1–13.
- Schreiber, T. (2000). Measuring information transfer. *Physical review letters* 85(2), 461.
- Shackleton, N. J. (2000). The 100,000-year ice-age cycle identified and found to lag temperature, carbon dioxide, and orbital eccentricity. *Science* 289(5486), 1897–1902.
- Shaffer, G. and F. Lambert (2018). In and out of glacial extremes by way of dust- climate feedbacks. *Proceedings of the National Academy of Sciences* 115(9), 2026–2031.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal* 27(3), 379–423.
- Spratt, R. M. and L. E. Lisiecki (2016). A late pleistocene sea level stack. *Climate of the Past* 12(4), 1079–1092.
- Stocker, T. F., D. Qin, G. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley (2013). Climate change 2013: the physical science basis. intergovernmental panel on climate change, working group i contribution to the ipcc fifth assessment report (ar5). *New York*.
- Strogatz, S. H. (2018). *Nonlinear dynamics and chaos with student solutions manual: With applications to physics, biology, chemistry, and engineering*. CRC press.
- Takens, F. (1981). Detecting strange attractors in turbulence. In *Dynamical systems and turbulence, Warwick 1980*, pp. 366–381. Springer.
- Tegen, I. (2003). Modeling the mineral dust aerosol cycle in the climate system. *Quaternary Science Reviews* 22(18-19), 1821–1834.

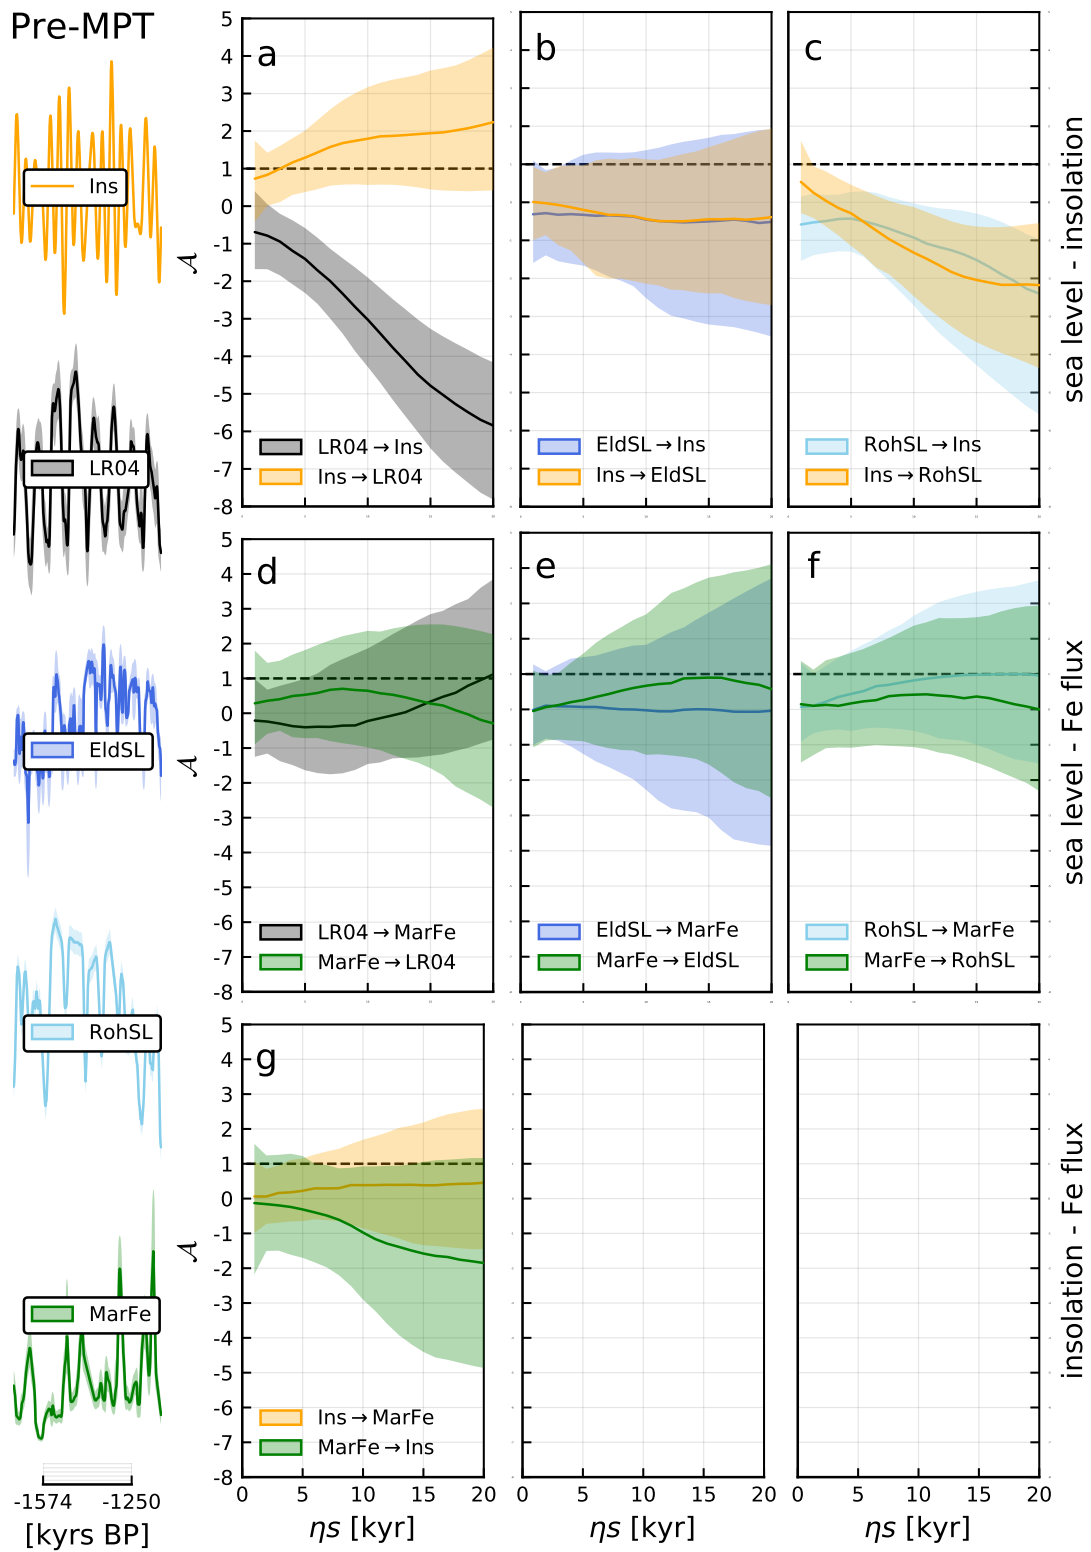
- Tegen, I. and A. A. Lacis (1996). Modeling of particle size distribution and its influence on the radiative properties of mineral dust aerosol. *Journal of Geophysical Research: Atmospheres* 101(D14), 19237–19244.
- Veres, D., L. Bazin, A. Landais, H. Toyé Mahamadou Kele, B. Lemieux-Dudon, F. Parrenin, P. Martinerie, E. Blayo, T. Blunier, E. Capron, et al. (2013). The antarctic ice core chronology (aicc2012): an optimized multi-parameter and multi-site dating approach for the last 120 thousand years. *Climate of the Past* 9(4), 1733–1748.
- Winckler, G., R. F. Anderson, S. L. Jaccard, and F. Marcantonio (2016). Ocean dynamics, not dust, have controlled equatorial Pacific productivity over the past 500,000 years. *Proceedings of the National Academy of Sciences of the United States of America* 113(22).
- Zelinka, M. D., T. A. Myers, D. T. McCoy, S. Po-Chedley, P. M. Caldwell, P. Ceppi, S. A. Klein, and K. E. Taylor (2020). Causes of higher climate sensitivity in cmip6 models. *Geophysical Research Letters* 47(1), e2019GL085782.



## **Appendix A**

### **Appendix 1: Predictive asymmetry results**

#### **A.1 Predictive asymmetry results for the pre-MPT windows**



**Figure A.1:** Predictive asymmetry results for the pre-MPT time window (15740-1250 ka BP). This figure is included here to facilitate comparison with the corresponding analysis over the longer control window.

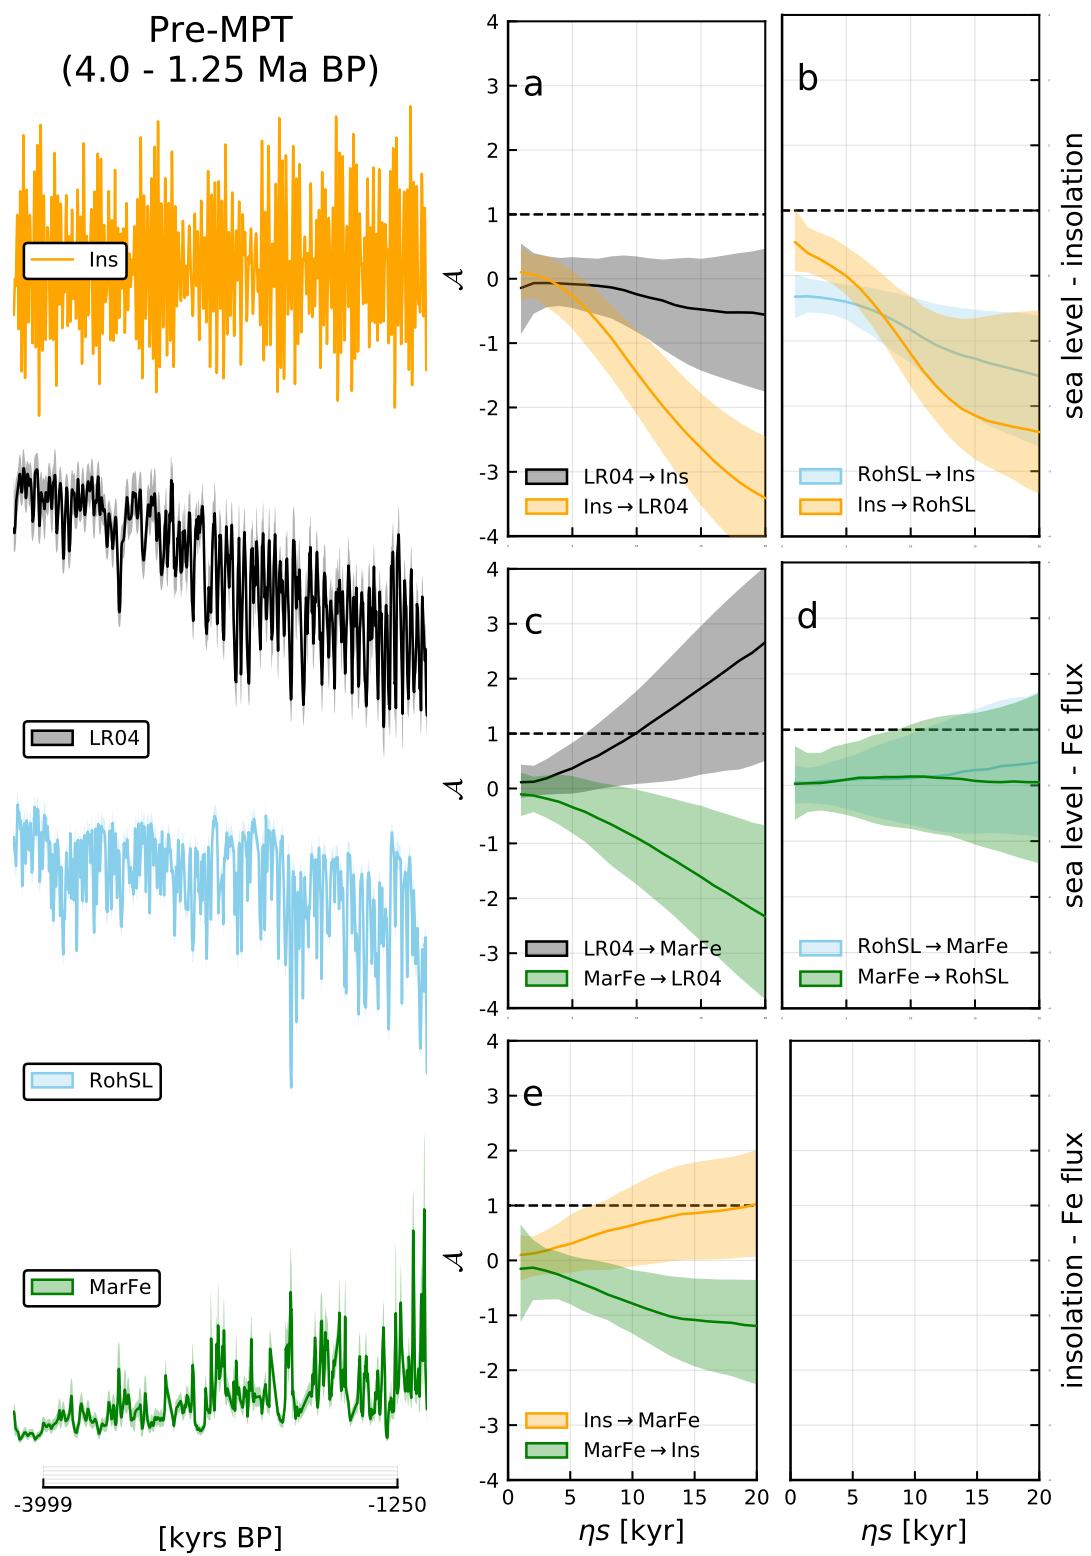
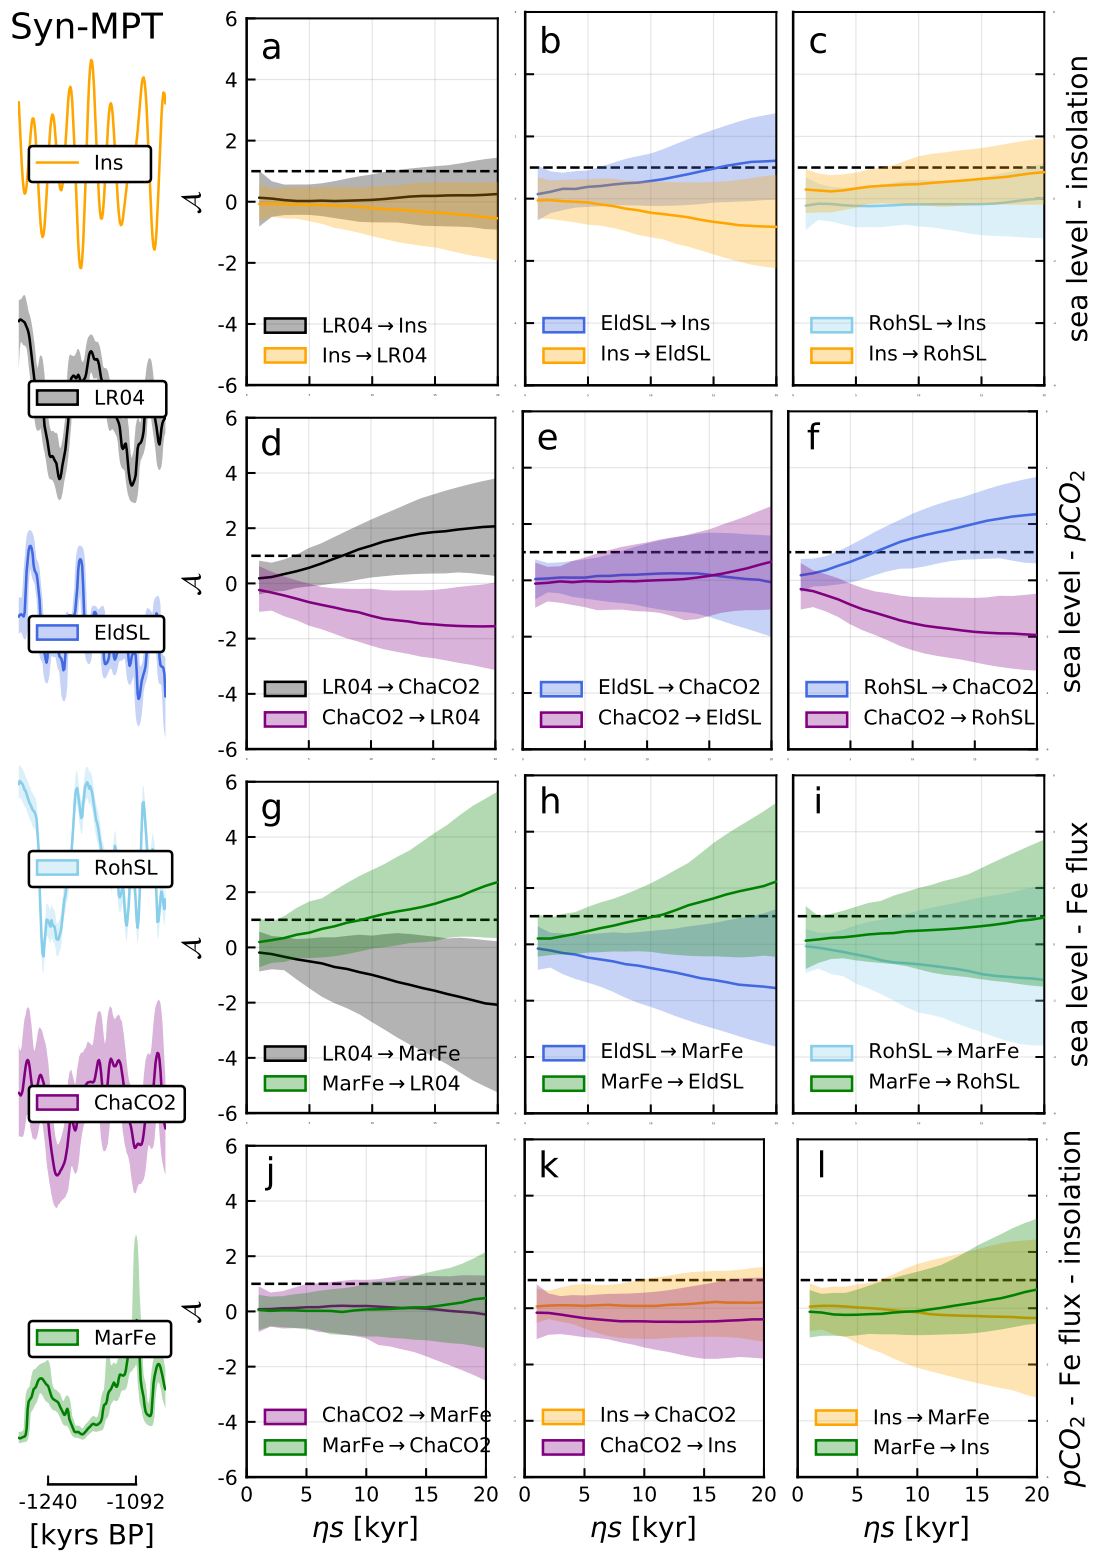


Figure A.2: Predictive asymmetry results for the longer pre-MPT control window (4000-1250 ka BP).

## **A.2 Predictive asymmetry results for the syn-MPT windows**



**Figure A.3:** Predictive asymmetry results for the syn-MPT time window (1240-1092 ka BP). This figure is included here to facilitate comparison with the corresponding analysis over the longer control window.



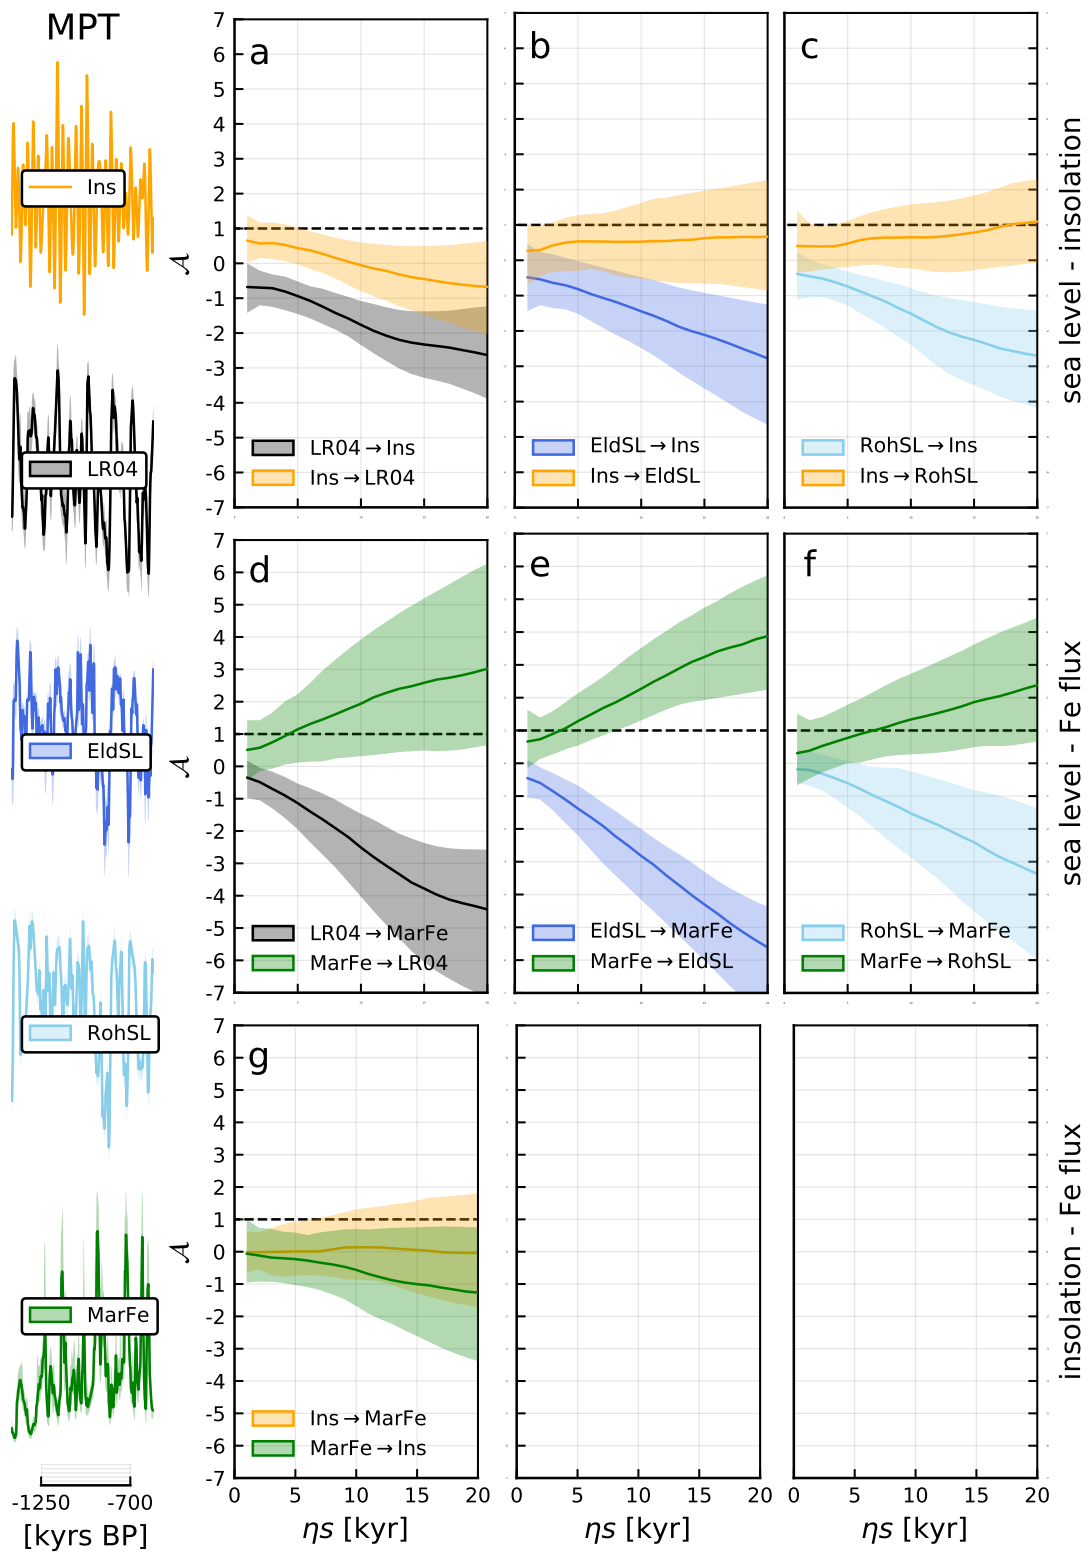
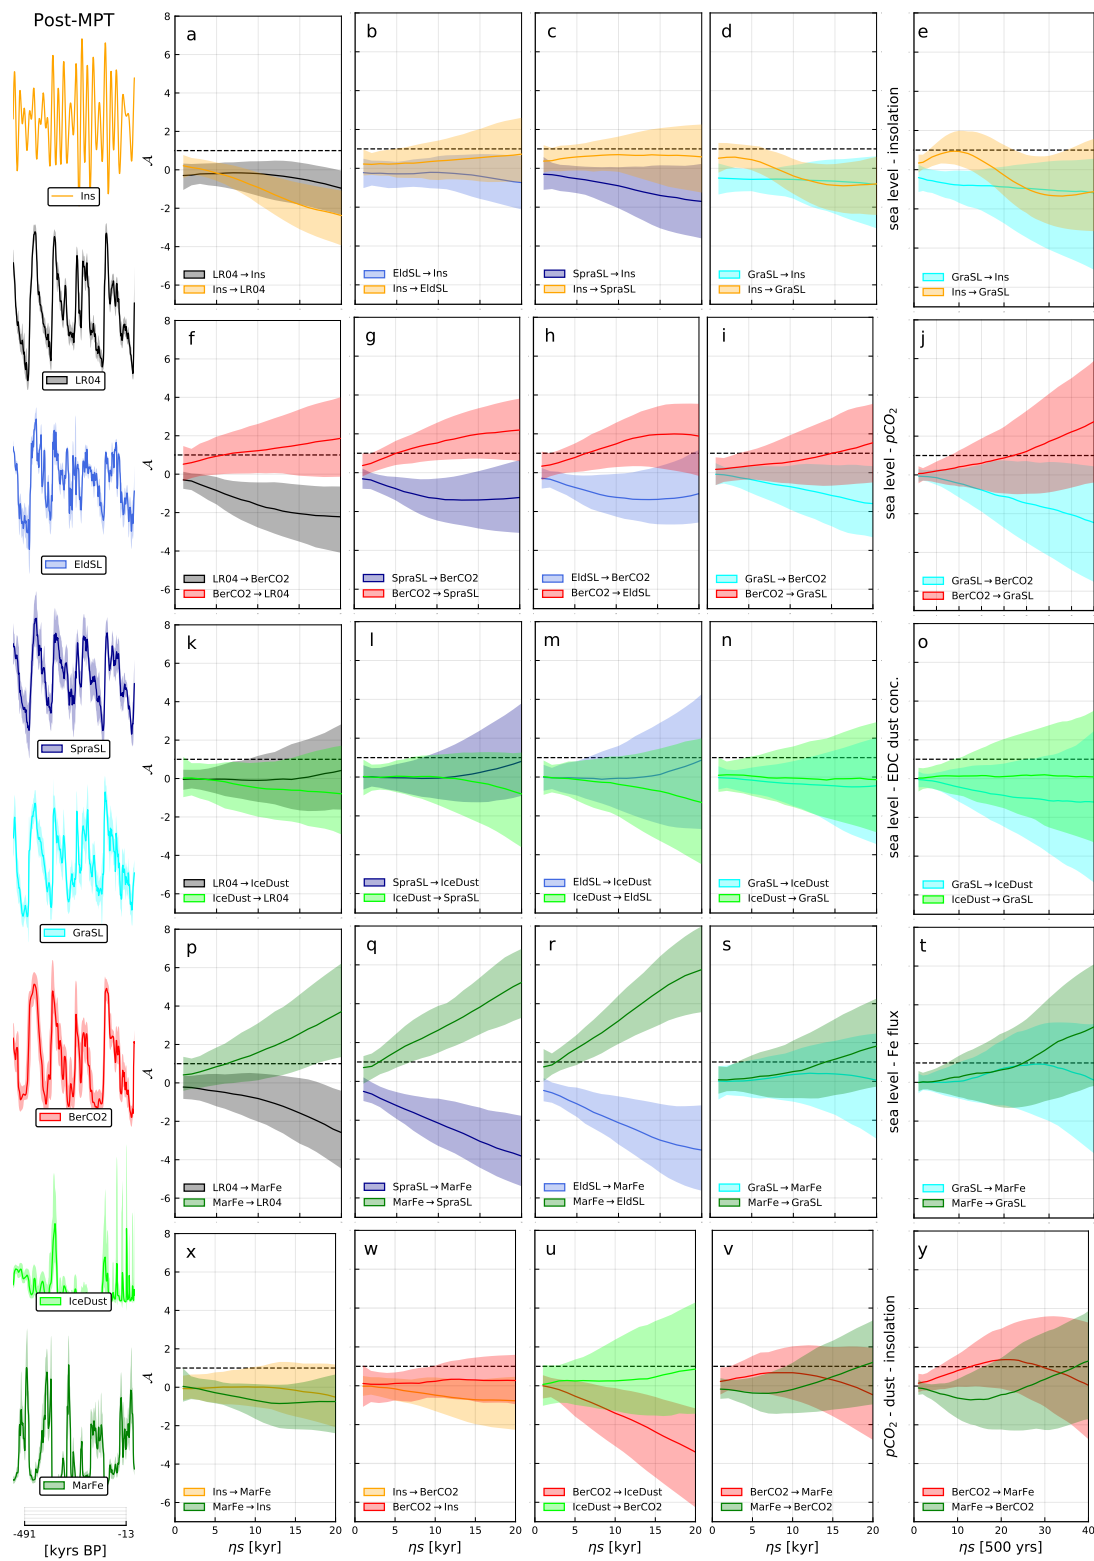


Figure A.4: Predictive asymmetry results for the longer MPT control window (1250-700 ka BP).

### **A.3 Predictive asymmetry results for the post-MPT windows**



**Figure A.5:** Predictive asymmetry results for the post-MPT time window (492-13 ka BP). This figure is included here to facilitate comparison with the corresponding analysis over the longer control window.

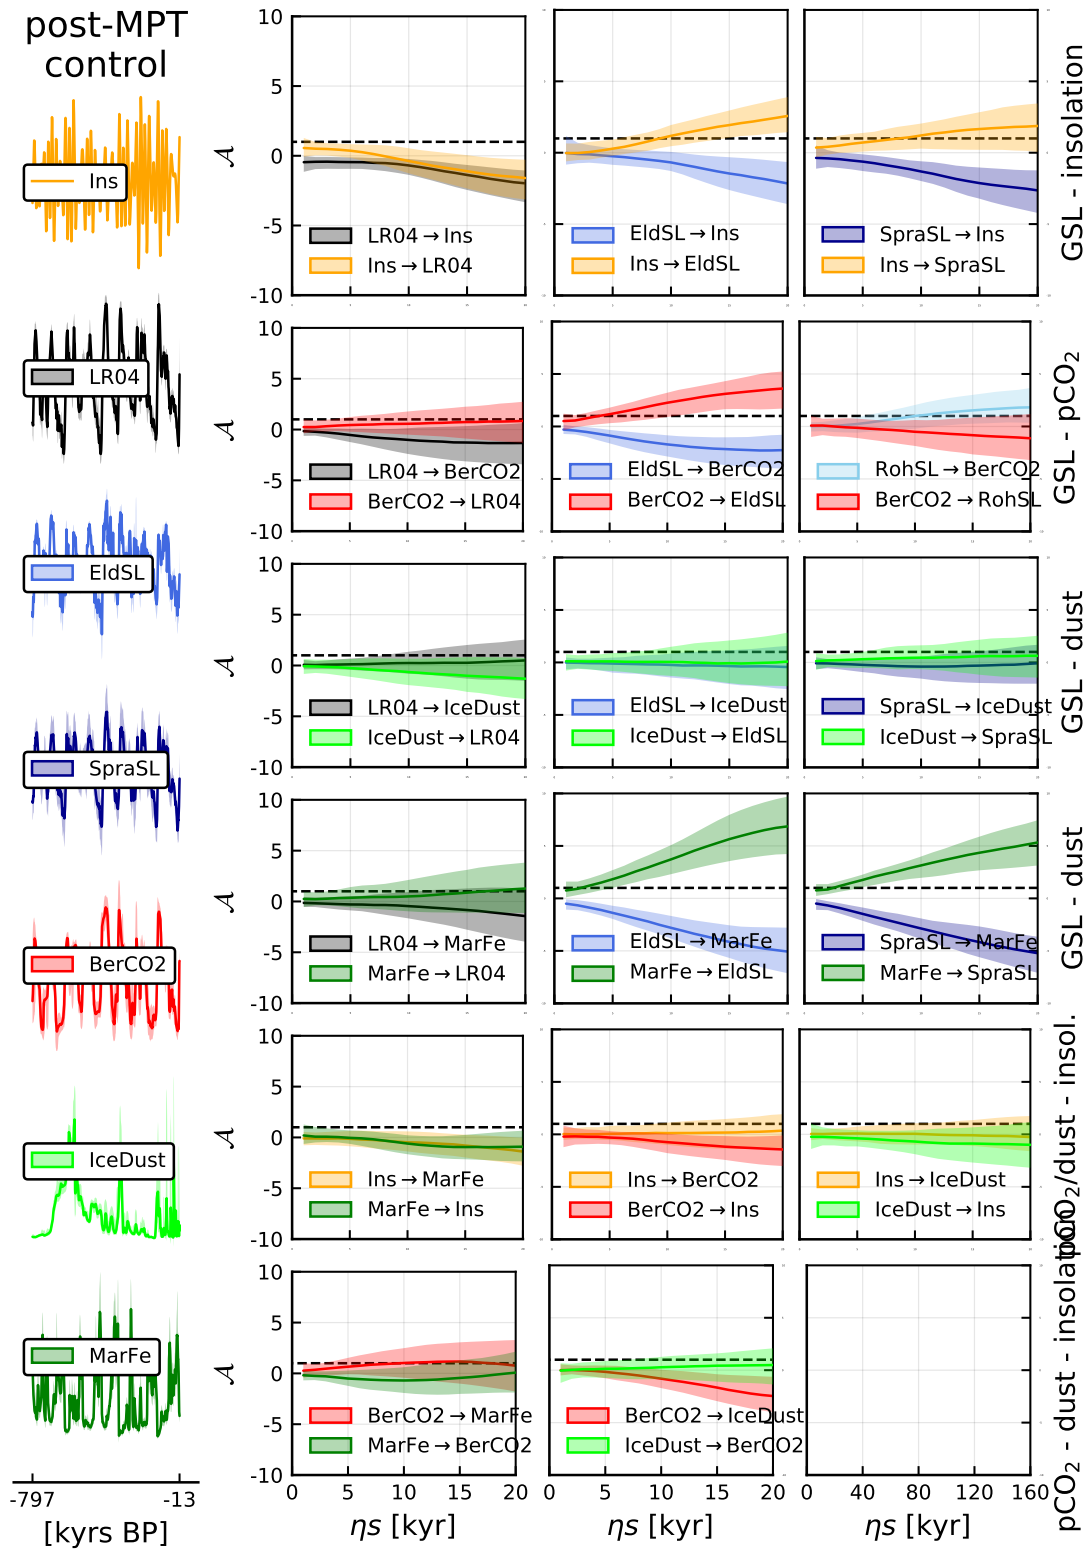
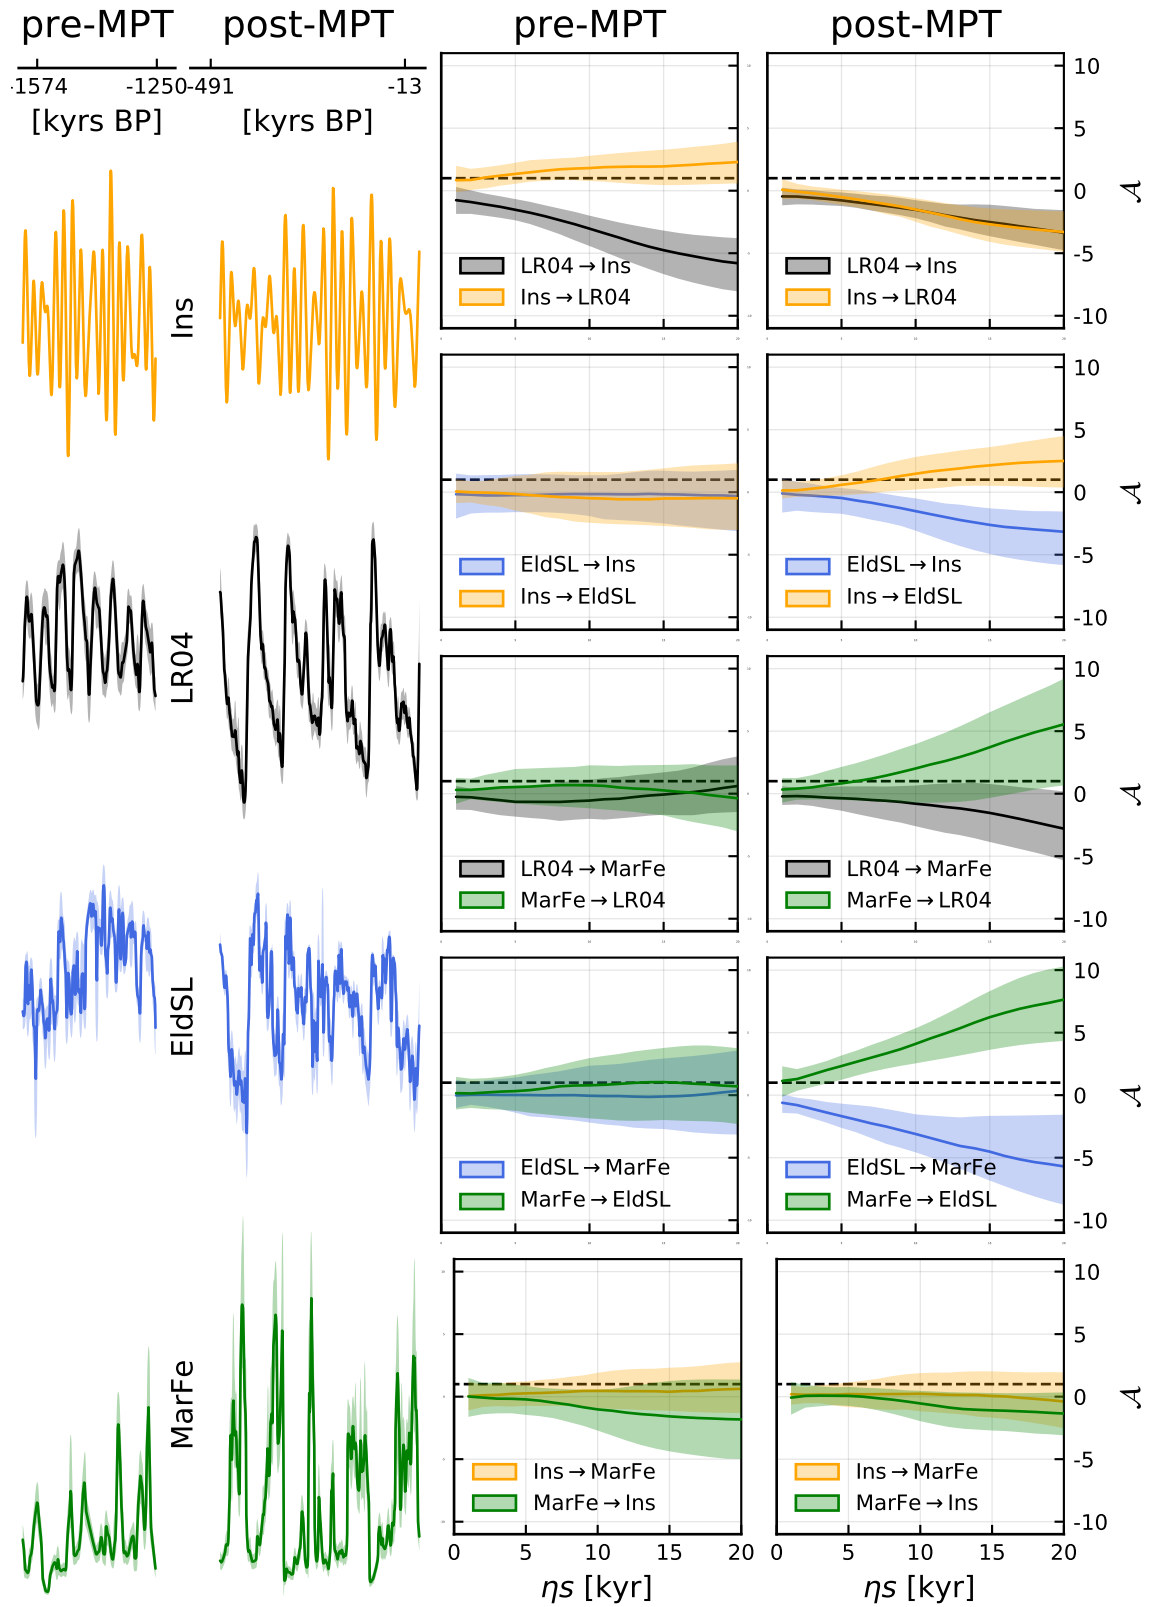


Figure A.6: Predictive asymmetry results for the longer post-MPT control interval (797-13 ka BP).

## **A.4 Comparative analysis for pre- and post-MPT windows**



**Figure A.7: Sampling-standardized comparison of the pre- and post-MPT dynamics.** The post-MPT time series are constrained to match the number of observations in the pre-MPT window ( $N=325$ ).

## A.5 Sensitivity analyses

### A.5.1 Temporal binning of time series

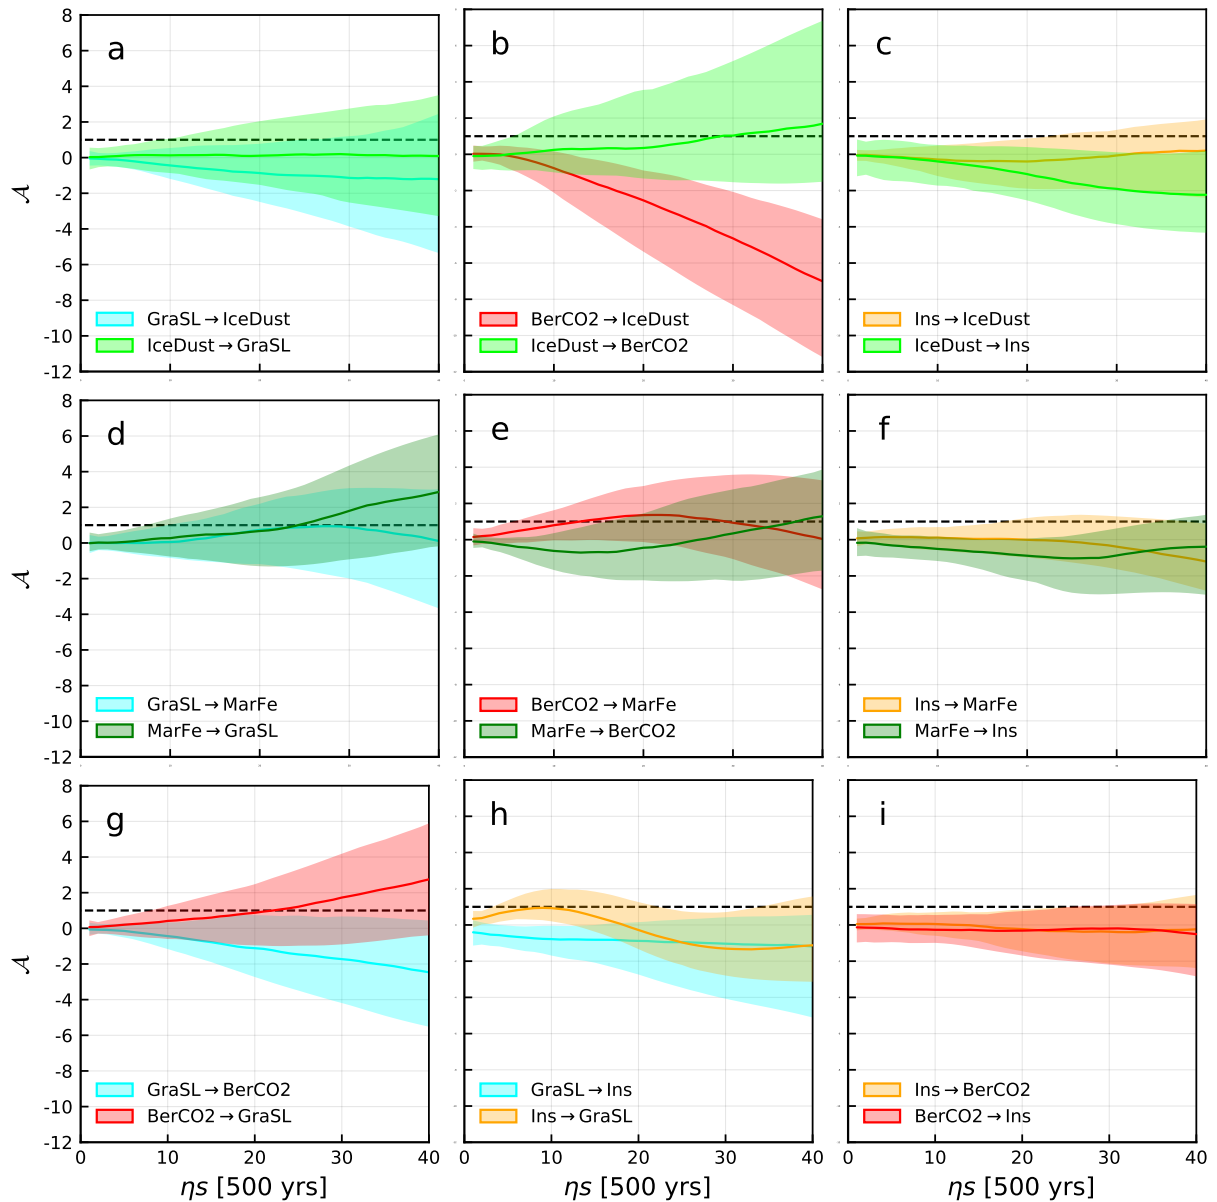




Figure A.8: Predictive asymmetry results for the post-MPT time window (492-13 ka BP) at 500 yr resolution.



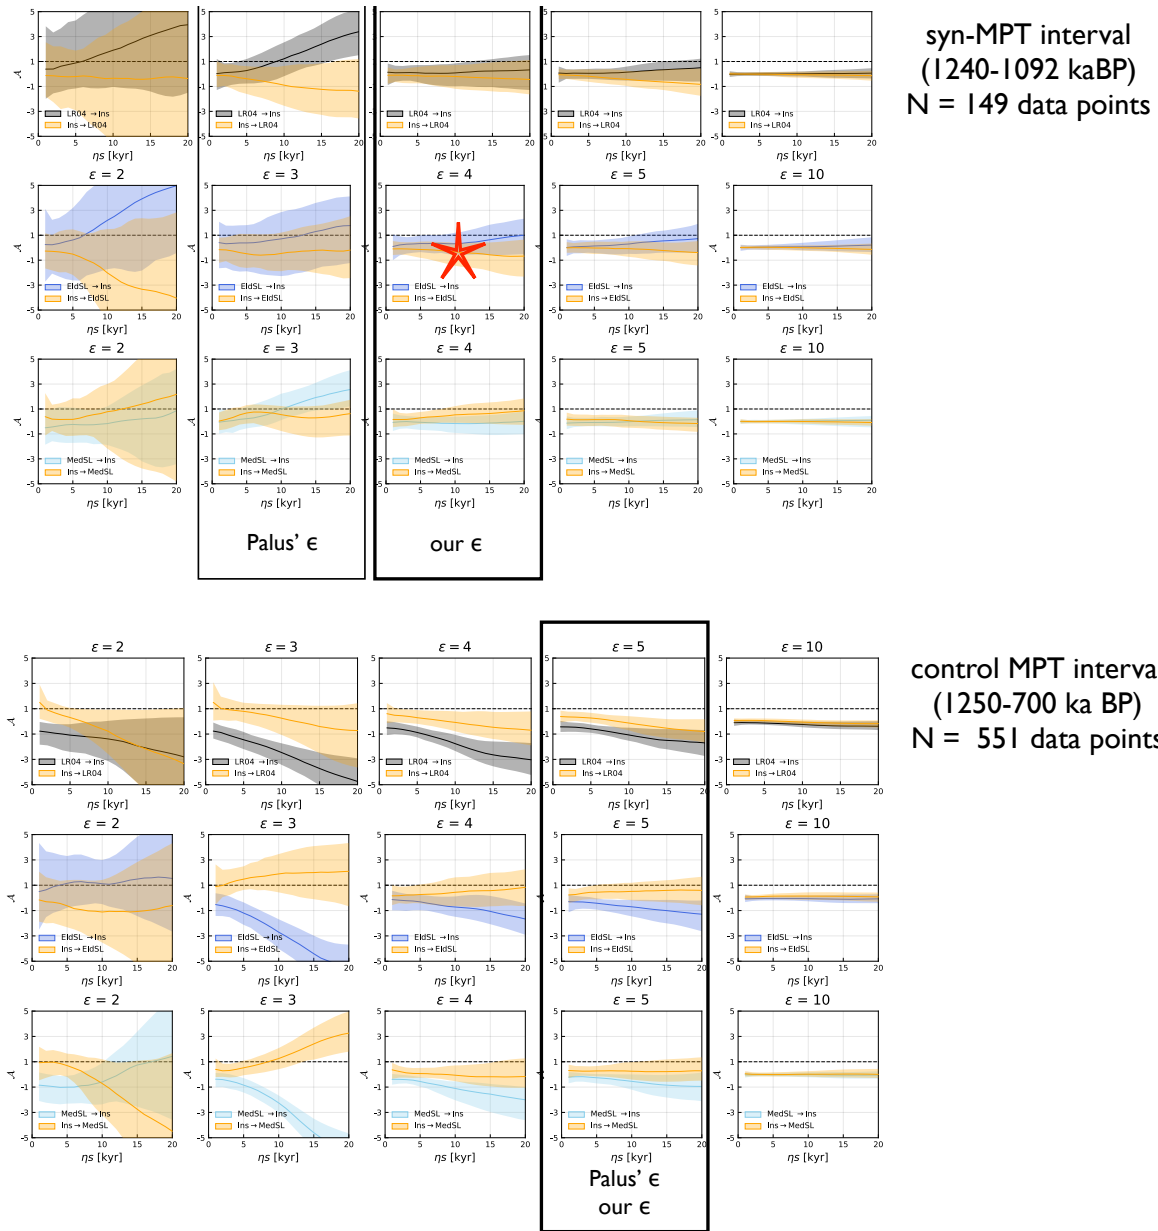
**Figure A.9:** Sensitivity analysis of temporal resolution of over-sampled ChaCO<sub>2</sub> record (1kyr, 500 yrs and 125 yrs) on the syn-MPT time window (1240-1092 ka BP).





**Figure A.10:** Sensitivity analysis of temporal resolution of GraSL record (1kyr, 500 yrs and 125 yrs) on the post-MPT time window (492-13 ka BP).

### A.5.2 Binning of the delay reconstruction ( $\epsilon$ )



**Figure A.11:** Sensitivity analysis on  $\epsilon$  over the synMPT and control MPT time intervals ( $N=149$  and  $N=551$ , respectively).

**A.5.3 Time series length ( $N$ )**