

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

A State-Space Framework for Modeling Seasonally-Varying Functional Time Series of Stratospheric Ozone Profiles

Permalink

<https://escholarship.org/uc/item/62q393bn>

ISBN

9798186382621

Author

Xiong, Moya

Publication Date

2026-06-10

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

**A State-Space Framework for Modeling
Seasonally-Varying Functional Time Series of
Stratospheric Ozone Profiles**

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Statistics and Applied Probability

by

Moya Xiong

Committee in charge:

Professor Wendy Meiring, Chair
Professor Yuedong Wang
Professor Emeritus John Hsu

June 2026

The Dissertation of Moya Xiong is approved.

Professor Yuedong Wang

Professor Emeritus John Hsu

Professor Wendy Meiring, Committee Chair

June, 2026

A State-Space Framework for Modeling Seasonally-Varying Functional Time Series of
Stratospheric Ozone Profiles

Copyright © 2026

by

Moya Xiong

To my family.

Acknowledgements

I would like to express my deepest gratitude to my advisor, Professor Wendy Meiring, for her guidance, support, and encouragement throughout my Ph.D. study. Her thoughtful advice, careful feedback, and patience have shaped both this dissertation and my development as a researcher. I am especially grateful for the time and effort she devoted to discussing research ideas, reading drafts, and helping me navigate the many stages of this work.

I am also sincerely grateful to my committee members, Professors John Hsu and Yuedong Wang, for their time, valuable comments, and support. Their suggestions helped improve the quality and clarity of this dissertation. I would also like to thank the faculty members and staff in the Department of Statistics and Applied Probability for providing a supportive academic environment throughout my graduate study.

I gratefully acknowledge the World Ozone and Ultraviolet Radiation Data Centre (WOUDC), which is part of the World Meteorological Organization Global Atmosphere Watch (WMO/GAW) programme, and the contributing data providers for making the ozonesonde observations used in this dissertation available.

I would also like to thank my fellow graduate students, classmates, and friends for their friendship, encouragement, and many helpful conversations. Their support made the challenges of graduate school more manageable and the experience more enjoyable. I am especially grateful to Tiansi Li for his support throughout my Ph.D. journey. I would also like to express my sincere thanks to Hong Xu. I am extremely grateful for the kindness and support I received from all of them.

Finally, I am deeply grateful to my family for their unconditional love, patience, and support. Their encouragement has been a source of strength throughout my studies. This dissertation would not have been possible without them.

Curriculum Vitæ

Moya Xiong

Education

- | | |
|------|--|
| 2026 | Ph.D. in Statistics and Applied Probability (Expected), University of California, Santa Barbara. |
| 2019 | A.M. in Statistics, Washington University in St. Louis. |
| 2017 | B.A. in Economics, Wuhan University. |
| 2017 | B.S. in Mathematics, Wuhan University. |

Experience

- | | |
|-----------|--|
| 2020-2026 | Teaching Assistant, Department of Statistics and Applied Probability, University of California, Santa Barbara. |
|-----------|--|

Abstract

A State-Space Framework for Modeling Seasonally-Varying Functional Time Series of Stratospheric Ozone Profiles

by

Moya Xiong

This dissertation develops a state-space modeling framework for analyzing stratospheric ozone concentration profiles over altitude and time. Ozone profiles exhibit smooth vertical structure and temporal variation, including seasonal patterns whose strength and timing can differ across altitude. To capture these features, the proposed model represents the vertical profile using cubic B-spline basis functions and models seasonal temporal variation through harmonic state components. The seasonal effects are allowed to vary smoothly across altitude through B-spline loading functions in the observation equation, which keeps the state dimension relatively low while maintaining flexibility.

Model parameters are estimated by maximum likelihood using the Kalman Filter likelihood, and latent states are estimated through Kalman filtering and smoothing. An innovation-based bootstrap is used to quantify parameter uncertainty. Identifiability issues in the seasonal component are also examined, leading to identifying constraints on the seasonal innovation variance and harmonic initial states.

Simulation studies show that the proposed model can recover latent state structure and key scalar parameters reasonably well under several seasonal settings. The bootstrap procedure provides useful uncertainty quantification, although seasonal loading functions become more variable when the seasonal dynamics are stochastic or when additional harmonic components are included.

The method is applied to Hohenpeissenberg ozonesonde profile data from 1979 to

2024. The fitted model captures the main vertical structure of the ozone profiles and reveals altitude-dependent seasonal variation. Diagnostic and bootstrap results suggest that the model provides a useful and interpretable representation of ozone profile dynamics, while also indicating areas where further model refinements may be beneficial.

Contents

Curriculum Vitae	vi
Abstract	vii
1 Introduction	1
1.1 Scientific Background and Motivation	2
1.2 Ozone Profile Data and Modeling Challenges	3
1.3 Existing Approaches for Ozone and Atmospheric Profile Modeling	4
1.4 Thesis Objective	5
1.5 Organization of the Thesis	6
2 Model Overview	7
2.1 State-Space Model	7
2.2 Modeling Motivation	8
2.3 Notation and Setup	10
2.4 State Equation Decomposition and Evolving	12
2.5 Observation equation components and evolution	18
2.6 Model Identifiability	24
3 Related Work and Contributions	32
3.1 Spline-based Joint Modeling Across Altitude	32
3.2 Reduced Fourier-Form Dynamic Linear Models	35
3.3 State-Space Models At Individual Altitudes	40
3.4 Functional Modeling of Entire Ozone Profiles	45
3.5 Functional Mixed Models for Ozone Trends	49
3.6 Positioning of This Work	55
4 Maximum Likelihood Estimation in Linear Gaussian State-Space Models	56
4.1 Maximum Likelihood Estimation	56
4.2 Kalman Filter and State Estimation	58
4.3 Likelihood Evaluation via the Kalman Filter	64

4.4	Innovation-Based Bootstrap for Parameter Inference	65
4.5	Numerical Optimization	68
5	Simulation Study	70
5.1	Parameter Specification	71
5.2	Data Generation	75
5.3	Model Fitting	77
5.4	Simulation Bootstrap Implementation	79
5.5	Simulation Results	80
6	Application to Hohenpeissenberg Ozonesonde Profiles	102
6.1	WOUDC Hohenpeissenberg Ozonesonde Data	102
6.2	Data Preprocessing	104
6.3	Exploratory Analysis	106
6.4	Model Specification for the Real Data	110
6.5	Fitted Model Results	116
6.6	Model Diagnostics	129
6.7	Bootstrap Uncertainty Quantification	132
6.8	Discussion and Summary	136
7	Summary, Discussion, and Future Work	140
7.1	Summary of the Thesis	140
7.2	Main Contributions	142
7.3	Discussion of Findings	145
7.4	Limitations	147
7.5	Future Work	148
7.6	Concluding Remarks	151
A	Cubic B-spline Basis Functions	152
A.1	B-spline Basis Functions	152
A.2	Cubic B-splines	153
A.3	Knots and Degrees of Freedom	154
A.4	B-spline Design Matrix	154
B	Details for Numerical Optimization	156
B.1	Initialization Setup	156
B.2	Optimization Settings	157
B.3	Convergence and Practical Considerations	159
	Bibliography	160

Chapter 1

Introduction

Stratospheric ozone plays an important role in atmospheric chemistry and climate, and its temporal variation has been the subject of extensive scientific study. Although total column ozone provides an important summary of atmospheric ozone abundance, it does not describe how ozone is distributed vertically through the atmosphere. This dissertation therefore focuses on ozone vertical profiles observed repeatedly over time. This profile-valued structure contains both smooth altitude-dependent variation and temporal variation, including seasonal patterns whose magnitude and timing may differ across altitude. These features motivate statistical models that can represent the ozone profile as a smooth function of altitude and describe how this profile evolves over time.

This dissertation develops a state-space modeling framework for analyzing lower-to-middle stratospheric ozone concentration profiles over altitude and time. The proposed model represents vertical structure using B-spline basis functions and models seasonal temporal variation through harmonic state components with altitude-dependent loading functions. The goal is to provide a flexible, interpretable, and computationally feasible framework for studying ozone profile dynamics.

1.1 Scientific Background and Motivation

Stratospheric ozone is a key component of the atmosphere because it absorbs a substantial portion of incoming solar ultraviolet radiation. Changes in stratospheric ozone therefore have implications for atmospheric chemistry, radiative processes, and surface ultraviolet exposure (see [1, 2, 3]). The discovery of ozone depletion and the subsequent monitoring of ozone recovery have made long-term ozone observations an important part of atmospheric science (see [4, 5, 6]).

Much of the historical discussion of ozone depletion and recovery is based on total column ozone, which provides a useful integrated measure of atmospheric ozone abundance. However, total column ozone does not describe the vertical distribution of ozone. Since chemical, radiative, and dynamical processes can differ by altitude, the evolution of vertical ozone profiles across time provides additional information about where and how ozone varies in the atmosphere (see [7, 8, 9]).

Ozone vertical profiles are observed by multiple instrument-types such as ozonesondes, lidar, Umkehr measurements, and satellite sensors, providing data across different spatial and temporal time scales ([3, 10]). In particular, ozonesondes provide long-term vertically resolved measurements extending into the stratosphere and are widely used for ozone trend studies and satellite validation ([11]). Ozonesondes are launched from the ground, and measure ozone partial pressures as they ascend into the stratosphere, carried by balloons. Each ozonesonde flight provides data of the ozonesonde partial pressures as a function of altitude for that flight-date. A sequence of ozonesonde flights from a specific location produce a time series of profile data, representing the evolution of ozone partial pressures across both altitude and time. Such data are scientifically informative but statistically challenging: the observed ozone partial pressure profiles are functions of altitude, observations are temporally dependent, and atmospheric dynamics including

the seasonal cycle differ across altitude levels.

These features motivate the development of statistical models that treat the vertical ozone profile as a structured functional object that evolves over time, rather than as a collection of unrelated single-altitude time series.

1.2 Ozone Profile Data and Modeling Challenges

Ozonesonde profile observations can be viewed as a time series of vertical curves. At each observation time, the data describe ozone concentration over altitude, so the time series of observations contains both vertical structure and temporal structure. In practice, raw profile measurements may be observed from ozonesonde flights at irregular launch dates, at different altitudes ([12]), with complex dependencies across space and time. At midlatitude sites, the seasonal cycle dominates stratospheric ozone variation, with both the magnitude and timing differing smoothly across altitude levels. Modeling each altitude level separately may therefore ignore useful information contained in the vertical structure of the profile. In contrast, modeling the time-evolving vertical ozone profiles using functional data analysis enables the borrowing of strength across altitudes in estimating complex altitude-dependent temporal evolution, including seasonality.

The complexities of the spatio-temporal scales of ozonesonde observations, and of stratospheric ozone dynamics, require a trade-off between flexibility, interpretability, and computational feasibility. A model that treats all altitude levels as unrelated time series may be simple but does not fully use the profile structure. A fully flexible multivariate dynamic model over all altitude levels may be difficult to estimate and computationally expensive. This motivates a reduced-dimensional dynamic representation in which vertical structure is represented through smooth basis functions and temporal evolution is modeled through a lower-dimensional latent state process, within a functional data

analysis state space framework.

1.3 Existing Approaches for Ozone and Atmospheric Profile Modeling

Many previous works have studied stratospheric ozone variability and recovery using linear regression and linear time-series regression models. These approaches often represent ozone as a function of time, seasonal terms, and atmospheric explanatory variables such as solar activity, the quasi-biennial oscillation, El Niño–Southern Oscillation, volcanic aerosol loading, and measures of anthropogenic halogen loading (for example, [13, 14, 15, 16, 17]). Such models have played an important role in assessing ozone depletion and recovery, especially for total column ozone and altitude-specific ozone time series.

State-space models ([18, 19, 20]) and dynamic linear models ([21, 22]) provide another framework for atmospheric time series because they separate latent temporal evolution from measurement error and allow inference through Kalman filtering and smoothing. In ozone applications, these models have been used to estimate time-varying trends, seasonal components, and other latent temporal structures in ozone time series and profile-related settings (for example, see [23, 24, 25, 26]). These methods provide a natural way to describe temporal dependence, but their treatment of vertical profile structure depends on the model formulation and does not necessarily provide a smooth altitude-basis representation with shared seasonal states across the full profile.

Functional data analysis approaches (see [27, 28, 29]) focus more directly on the vertical profile structure by treating each ozone profile as a function observed over altitude or pressure. These methods can represent ozone profiles using smooth basis-function expan-

sions, functional principal components or related low-dimensional functional summaries, thereby borrowing information across neighboring altitude levels. In ozone applications, functional and profile-based approaches have been used to study the vertical structure of ozone variability, nonlinear temporal changes, and the effects of atmospheric covariates (for example, [30, 31, 32]). These methods are useful for representing vertical dependence, although the temporal evolution is not always modeled through an explicit latent dynamic process.

The existing literature therefore highlights a trade-off between modeling vertical structure, describing temporal dependence, preserving interpretability, and maintaining computational feasibility. The approach developed in this dissertation combines these ideas by representing the ozone profile through B-spline basis functions while modeling temporal evolution in a state-space framework. Seasonal variation is represented through shared harmonic state components whose loading functions vary smoothly across altitude, allowing altitude-dependent seasonality without introducing a separate seasonal state process at every altitude level.

1.4 Thesis Objective

The objective of this thesis is to develop a statistical framework for modeling ozone vertical profiles as a dynamic process over altitude and time. The framework is designed for profile-valued observations that exhibit smooth vertical structure, temporal dependence, and seasonal variation whose magnitude and timing may differ across altitude.

Rather than analyzing each altitude level as a separate time series, this thesis treats the vertical ozone profile as a structured object evolving over time. The proposed model, estimation procedure, identifiability discussion, simulation study, and real-data application are developed around this objective.

1.5 Organization of the Thesis

The remainder of this thesis is organized as follows. Chapter 2 introduces the proposed state-space model for ozone profile time series. It presents the state and observation equations, the B-spline representation of the altitude-dependent profile structure at each time, the state evolution of the B-spline coefficients and harmonic seasonal state components, and constraints introduced due to identifiability issues.

Chapter 3 compares the model proposed in Chapter 2 to closely related literature in analyzing stratospheric ozone, and also to reduced harmonic seasonal cycle models in state space models literature.

Chapter 4 describes the maximum likelihood estimation framework for the linear Gaussian state-space model. It reviews Kalman filtering and smoothing, likelihood evaluation through one-step-ahead innovations, innovation-based bootstrap inference, and numerical optimization.

Chapter 5 presents simulation studies under several parameter configurations. The simulations are designed to examine state estimation, scalar parameter estimation, recovery of seasonal loading functions, and bootstrap uncertainty quantification under controlled settings.

Chapter 6 applies the proposed method to Hohenpeissenberg ozonesonde profile data, from a time series of ozonesonde flights. It describes the WOUDC data source ([10, 33]), preprocessing into a regular pentad-level altitude-by-time array, exploratory analysis, model specification, fitted results, diagnostics, and bootstrap uncertainty summaries.

Finally, Chapter 7 summarizes the main findings, discusses limitations of the current framework, and outlines possible directions for future work.

Chapter 2

Model Overview

2.1 State-Space Model

Many time series arising in environmental and atmospheric sciences exhibit complex temporal dynamics across space and time. Observed data are often noisy measurements of an underlying process, itself driven by latent processes that evolve over time. State-space models (SSMs) provide a flexible probabilistic framework to model such systems by introducing latent state variables that capture the unobserved dynamics [19, 20].

A (linear Gaussian) state-space model consists of two equations: an observation equation and a state transition equation. Let $\mathbf{y}_t \in \mathbb{R}^{d_y}$ denote the observed data at time t , and let $\mathbf{x}_t \in \mathbb{R}^{d_x}$ denote the latent (hidden) state vector. Here d_y and d_x denote the dimensions of the observation and state vectors, respectively. In a linear Gaussian state-space model, the observations are linear functions of the latent states plus Gaussian observation noise, while the latent states evolve through a linear transition equation with Gaussian

process noise. The general form is

$$\text{Observation equation:} \quad \mathbf{y}_t = \mathbf{Z}_t \mathbf{x}_t + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_t), \quad (2.1)$$

$$\text{State equation:} \quad \mathbf{x}_t = \mathbf{B}_t \mathbf{x}_{t-1} + \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_t), \quad (2.2)$$

where $\mathbf{Z}_t$ is the observation matrix, $\mathbf{B}_t$ is the state transition matrix, and $\mathbf{v}_t$ and $\mathbf{w}_t$ are Gaussian noise processes, frequently assumed to be mutually independent.

The observation equation relates the latent state to the observed data, while the state equation describes how the latent process evolves over time. This framework naturally separates measurement noise from process variability, making it particularly suitable for modeling dynamical systems with hidden structure [19].

In this work, we adopt a time-invariant linear Gaussian state-space model, where $\mathbf{Z}_t = \mathbf{Z}$, $\mathbf{B}_t = \mathbf{B}$, $\mathbf{R}_t = \mathbf{R}$, and $\mathbf{Q}_t = \mathbf{Q}$ for all t . The specific structure of these matrices will be motivated based on the scientific characteristics of the ozone profile data.

2.2 Modeling Motivation

The objective of this work is to model the evolution of ozone concentration as a function of both altitude and time. The motivating observed ozone profile data, obtained from a time series of balloon-based ozonesonde flights from a midlatitude location, may be considered as data from a time series of profiles (curves), since each ozonesonde flight provides observations of the ozone concentrations on that date as a function of altitude as the sonde ascends [10],[30]. These data exhibit two dominant characteristics: a smooth variation across altitude and a seasonal pattern over time. The proposed model is designed to capture these features in a flexible yet computationally efficient framework.

Ozone concentration varies smoothly with altitude, with localized features such as

peaks and gradients occurring at different altitudes. B-spline basis functions (appendix A) as a function of altitude can provide flexible reduced-dimension approximations for such smooth functions across altitudes. Therefore, we represent the vertical ozone profile at time t as a linear combination of B-spline basis functions, in addition to seasonal harmonic functional features, and model non-harmonic temporal evolution through latent time-evolving B-spline coefficients. Notably, B-splines have compact support, meaning that each basis function is nonzero only over a limited range of altitudes, benefiting modeling and estimation. In particular, the bounded altitude support of each B-spline allows the model to capture variations localized to specific altitude regions, through B-spline coefficient components, without individual coefficients affecting the entire profile. This localized property improves interpretability and numerical stability [34, 35].

Ozone dynamics are strongly influenced by seasonal effects driven by photochemical processes and atmospheric circulation at midlatitudes. Such periodic patterns can be naturally modeled using harmonic (Fourier) components [20]. By introducing pairs of cosine and sine terms into the latent state equation, and including corresponding harmonic loading functions as one component of the observation equation, our model captures both the amplitude and phase of seasonal variation in a parsimonious manner.

The model is formulated within a state-space framework to allow for the dynamic evolution of both the latent state B-spline coefficients and latent state Fourier harmonics. Importantly, we structure the model with static altitude-dependent seasonal loading functions (for the low-dimensional latent state Fourier harmonics) entering through the observation matrix $\mathbf{Z}$, instead of including higher dimensional Fourier coefficients as additional state variables. This design choice reduces the dimension of the latent state vector.

This reduction is critical because the computational cost of standard Kalman filtering and smoothing estimation algorithms scales cubically with the dimension of the state

vector, due to matrix operations such as covariance propagation and inversion [36, 37]. Consequently, increasing the state dimension can lead to a substantial increase in computational burden. By keeping the state dimension relatively low while honoring modeling goals, we achieve improved computational efficiency while maintaining model flexibility.

Overall, this modeling framework balances flexibility, interpretability, and computational tractability, making it well suited for analyzing ozone profile data. The following sections provide the detailed specification of the proposed state-space model (SSM): Section 2.3 introduces notation, dimensions and a broad-overview of proposed SSM structure; Section 2.4 provides additional details about the latent state evolution equation; Section 2.5 similarly expands details about the observation equation. Section 2.6 presents constraints implemented in the estimation, due to identifiability issues.

2.3 Notation and Setup

We consider ozone concentration profiles observed across multiple altitudes over time. Let $a_1, \dots, a_{n_a}$ denote the altitude grid, and suppose observations are available at T evenly spaced time points. The goal is to model the temporal evolution of the ozone profile while capturing both its vertical structure and seasonal variability.

Based on the characteristics of ozone profile data and the modeling considerations discussed above, we now specify the structure of the state vector and the specific form of the state-space model used in this work:

$$\mathbf{y}_t = \tilde{\mathbf{y}}_t + \mathbf{v}_t = \mathbf{Z}\mathbf{x}_t + \mathbf{v}_t \quad \mathbf{v}_t \stackrel{\text{ind}}{\sim} N(\mathbf{0}, \mathbf{R}) \quad t \in \{1, 2, \dots, T\} \quad (2.3)$$

$$\mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{w}_t \quad \mathbf{w}_t \stackrel{\text{ind}}{\sim} N(\mathbf{0}, \mathbf{Q}) \quad t \in \{1, 2, \dots, T\} \quad (2.4)$$

- $\mathbf{y}_t \in \mathbb{R}^{n_a}$ denotes the observed ozone profile across n_a altitude levels at time t .

- The observation equation (equation (2.3)) relates observed variables $\mathbf{y}_t$ to the latent state vector $\mathbf{x}_t$ via the matrix $\mathbf{Z}$, with observation noise $\mathbf{v}_t$. $\tilde{\mathbf{y}}_t$ is the noise-free part of the observation equation. Note that later $\mathbf{y}_t$ may refer to a random variable or its realization, depending on the context.
- The state equation (equation (2.4)) describes the temporal evolution of $\mathbf{x}_t$ governed by transition matrix $\mathbf{B}$, with process noise $\mathbf{w}_t$.
- $\mathbf{v}_t$ and $\mathbf{w}_{t'}$ are independent for all times t and t' .

Since we will discuss details of the state equation in Section 2.4, before we provide details on the observation equation in Section 2.5, we introduce here that the matrix $\mathbf{Z}$ in the observation equation is composed of different components in order to understand the state equation. Specifically, the observation equation will be decomposed as:

$$\mathbf{y}_t = \mathbf{Z}\mathbf{x}_t + \mathbf{v}_t = \mathbf{B}_{sp}\mathbf{x}_{t,c} + \sum_{j=1}^p \mathbf{B}_{sp}\gamma_{c,j}M_{t,j} + \sum_{j=1}^p \mathbf{B}_{sp}\gamma_{s,j}M_{t,j}^* + \mathbf{v}_t, \quad \mathbf{v}_t \stackrel{ind}{\sim} N(\mathbf{0}, \sigma_R^2 \mathbf{I}_{n_a}),$$

where the observation matrix $\mathbf{Z}$ has components:

$$\mathbf{Z} = \begin{pmatrix} \mathbf{B}_{sp} & \mathbf{B}_{sp}\gamma_{c,1} & \mathbf{B}_{sp}\gamma_{s,1} & \cdots & \mathbf{B}_{sp}\gamma_{c,p} & \mathbf{B}_{sp}\gamma_{s,p} \end{pmatrix},$$

and where:

- $\mathbf{B}_{sp} \in \mathbb{R}^{n_a \times k}$ is the cubic B-spline basis matrix with k degrees of freedom evaluated at the altitude grid where n_a is the number of altitudes,
- $\mathbf{x}_{t,c} \in \mathbb{R}^k$ is the time-varying B-spline coefficient vector for the non-harmonic functional component; $\mathbf{x}_{t,c}$ is the first component of the latent state vector,
- $\gamma_{c,j}, \gamma_{s,j} \in \mathbb{R}^k$ are the static B-spline coefficient vectors that determine the altitude-dependent loading functions for the j^{th} harmonic pair, for $j = 1, \dots, p$,

- $M_{t,j}$ and $M_{t,j}^*$ are the harmonic state variables for the j^{th} harmonic pair, for $j = 1, \dots, p$; evolving as the second component of the latent state vector,
- $\mathbf{v}_t$ is the observation noise vector at time point t , with common variance parameter σ_R^2 across altitude levels, assumed independent across time.
- The vector of latent states that evolves in the State Equation is of dimension $k+2p$, and decomposed as

$$\mathbf{x}_t = \begin{pmatrix} \mathbf{x}_{t,c}^\top & M_{t,1} & M_{t,1}^* & \dots & M_{t,p} & M_{t,p}^* \end{pmatrix}^\top \in \mathbb{R}^{k+2p},$$

as detailed further in Section 2.4

We summarize the key variables in the state-space model and their dimensions in the following table. The general dimension is expressed in terms of n_a (number of altitudes) and the dimension of the state vector $(k+2p)$.

Variable	Definition	Dimension
$\mathbf{y}_t$	Observation random variable	$n_a \times 1$
$\mathbf{x}_t$	State vector	$(k+2p) \times 1$
$\mathbf{Z}$	Observation matrix	$n_a \times (k+2p)$
$\mathbf{B}$	Transition matrix	$(k+2p) \times (k+2p)$
$\mathbf{w}_t$	State noise	$(k+2p) \times 1$
$\mathbf{v}_t$	Observation noise	$n_a \times 1$

Table 2.1: Dimensions of components in the state-space model.

2.4 State Equation Decomposition and Evolving

We partition the state vector into a component with time-evolving basis function coefficients (dynamic coefficients for B-splines in our model), and a component containing the harmonic components. Define

$$\mathbf{x}_t = \begin{pmatrix} \mathbf{x}_{t,c} \\ \mathbf{H}_t \end{pmatrix} \quad (2.5)$$

where

$$\mathbf{x}_{t,c} = \begin{pmatrix} x_{t,c_1} & \dots & x_{t,c_k} \end{pmatrix}^\top \quad (2.6)$$

represents the time-varying coefficients associated with the k B-spline basis functions (see appendix A). Here x_{t,c_i} is the i^{th} coefficient for the corresponding B-spline basis function in the non-harmonic component, $i \in \{1, 2, \dots, k\}$, k is the degrees of freedom for the B-spline. And

$$\mathbf{H}_t = \begin{pmatrix} M_{t,1} & M_{t,1}^* & \dots & M_{t,p} & M_{t,p}^* \end{pmatrix}^\top \quad (2.7)$$

denotes the seasonal harmonic state component.

Let $\mathbf{w}_t$ denote the state noise vector. We likewise partition the state noise vector into a spline coefficient component and a harmonic component. Define

$$\mathbf{w}_t = \begin{pmatrix} \mathbf{w}_{t,c} \\ \mathbf{w}_{t,H} \end{pmatrix} \quad (2.8)$$

where

$$\mathbf{w}_{t,c} = \begin{pmatrix} \epsilon_{t,c_1} & \dots & \epsilon_{t,c_k} \end{pmatrix}^\top \quad (2.9)$$

and

$$\mathbf{w}_{t,H} = \begin{pmatrix} \epsilon_{t,1,M} & \epsilon_{t,1,M}^* & \dots & \epsilon_{t,p,M} & \epsilon_{t,p,M}^* \end{pmatrix}^\top \quad (2.10)$$

For $t \in \{1, 2, \dots, T\}$, $i \in \{1, \dots, k\}$, $j \in \{1, \dots, p\}$, we assume

$$\epsilon_{t,c_i} \stackrel{\text{i.i.d}}{\sim} N(0, \sigma_c^2) \quad (2.11)$$

$$\epsilon_{t,j,M}, \epsilon_{t,j,M^*} \stackrel{\text{i.i.d}}{\sim} N(0, \sigma_\omega^2) \quad (2.12)$$

with all components mutually independent across t, i and j .

Consequently, the variance-covariance matrix of the state noise vector $\mathbf{w}_t$, denoted as $\mathbf{Q}$, is given by

$$\mathbf{Q} = \text{Var}(\mathbf{w}_t) = \begin{pmatrix} \mathbf{Q}_c & \mathbf{0}_{k \times 2p} \\ \mathbf{0}_{2p \times k} & \mathbf{Q}_H \end{pmatrix} = \begin{pmatrix} \sigma_c^2 \mathbf{I}_k & \mathbf{0}_{k \times 2p} \\ \mathbf{0}_{2p \times k} & \sigma_\omega^2 \mathbf{I}_{2p} \end{pmatrix} \quad (2.13)$$

In practice, σ_ω^2 must be fixed to avoid the identifiability issue mentioned in Subsection 2.6.1. Without loss of generality, we fix σ_ω^2 to a prespecified constant. In particular, setting $\sigma_\omega^2 = 1$ allows stochastic innovations in the seasonal harmonic components, whereas setting $\sigma_\omega^2 = 0$ yields deterministic seasonal harmonic dynamics.

The seasonal harmonic and conjugate harmonic, with nominal period T_p , evolve according to the following equations:

$$\begin{aligned} M_{t,j} &= d \cos(\omega_j) M_{t-1,j} + d \sin(\omega_j) M_{t-1,j}^* + \epsilon_{t,j,M} \\ M_{t,j}^* &= -d \sin(\omega_j) M_{t-1,j} + d \cos(\omega_j) M_{t-1,j}^* + \epsilon_{t,j,M^*} \end{aligned} \quad (2.14)$$

where ω_j is the Fourier frequency, defined as

$$\omega_j = \frac{2\pi j}{T_p} \quad (2.15)$$

Here, T_p represents the nominal seasonal period. Throughout this thesis, we fix $T_p = 73$, corresponding to a cycle under 5-day pentad aggregation. When $\sigma_\omega^2 = 0$, the j^{th} seasonal harmonic component is deterministic with frequency $\omega_j = \frac{2\pi j}{T_p}$. When $\sigma_\omega^2 > 0$,

stochastic innovations in the harmonic state imply that the resulting trajectories are no longer strictly periodic in finite samples. While the harmonic frequency remains fixed at $\omega_j = 2\pi j/T_p$, the realized seasonal cycle may exhibit stochastic deviations from the exact periodicity.

The transition matrix $\mathbf{B}$ in equation (2.4) is block diagonal under the state partition equation (2.5):

$$\mathbf{B} = \begin{pmatrix} \mathbf{B}_c & \mathbf{0}_{k \times 2} & \cdots & \mathbf{0}_{k \times 2} \\ \mathbf{0}_{2 \times k} & d\mathbf{R}(\omega_1) & \cdots & \mathbf{0}_{2 \times 2} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{2 \times k} & \mathbf{0}_{2 \times 2} & \cdots & d\mathbf{R}(\omega_p) \end{pmatrix} \quad (2.16)$$

where $\mathbf{B}_c$ governs the evolution of the basis function coefficients (coefficients of B-spline basis functions here), and $d\mathbf{R}(\omega_j)$, $j \in \{1, \dots, p\}$ specifies the transition of the j^{th} pair harmonic component. Here, $\mathbf{B}_c \in \mathbb{R}^{k \times k}$ and $\mathbf{R}(\omega_j) \in \mathbb{R}^{2 \times 2}$.

The coefficient block $\mathbf{B}_c$ is specified as a tridiagonal Toeplitz matrix,

$$\mathbf{B}_c = \begin{pmatrix} \alpha_0 & \alpha_1 & 0 & \cdots & 0 \\ \alpha_1 & \alpha_0 & \alpha_1 & \ddots & \vdots \\ 0 & \alpha_1 & \alpha_0 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & \alpha_1 \\ 0 & \cdots & 0 & \alpha_1 & \alpha_0 \end{pmatrix}. \quad (2.17)$$

The 2×2 rotation matrix is

$$\mathbf{R}(\omega_j) = \begin{pmatrix} \cos(\omega_j) & \sin(\omega_j) \\ -\sin(\omega_j) & \cos(\omega_j) \end{pmatrix}. \quad (2.18)$$

The harmonic block $\mathbf{R}(\omega_j)$ is a rotation matrix whose eigenvalues have unit modulus.

Multiplying by the damping factor $d < 1$ scales the eigenvalues of the harmonic block $d\mathbf{R}(\omega_j)$ so that their moduli are strictly less than one (see [19]).

The coefficient block $\mathbf{B}_c$ is a symmetric tridiagonal Toeplitz matrix. Closed-form eigenvalues and eigenvectors for tridiagonal Toeplitz matrices are well known; see, for example, Noschese et al. [38]. For the present symmetric case, where the diagonal entries are α_0 and the subdiagonal and superdiagonal entries are both α_1 , the eigenvalues of $\mathbf{B}_c$ are

$$\lambda_i = \alpha_0 + 2\alpha_1 \cos\left(\frac{i\pi}{k+1}\right), \quad i \in \{1, 2, \dots, k\} \quad (2.19)$$

with corresponding eigenvectors $\mathbf{v}_i \in \mathbb{R}^k$:

$$\mathbf{v}_i = \left(\sin\left(\frac{i\pi}{k+1}\right), \dots, \sin\left(\frac{ki\pi}{k+1}\right) \right)^\top. \quad (2.20)$$

For completeness, we verify this result for the present matrix. Let $v_{i,l}$ be the l^{th} element of $\mathbf{v}_i$, so

$$v_{i,l} = \sin\left(\frac{il\pi}{k+1}\right), \quad l = 1, \dots, k.$$

For notational convenience, define $v_{i,0} = 0$ and $v_{i,k+1} = 0$. This boundary convention is consistent with the sine form of the eigenvector, since

$$v_{i,k+1} = \sin\left(\frac{i(k+1)\pi}{k+1}\right) = \sin(i\pi) = 0.$$

Then, for each $l = 1, \dots, k$, the l th element of $\mathbf{B}_c \mathbf{v}_i$ can be written as

$$(\mathbf{B}_c \mathbf{v}_i)_l = \alpha_1 v_{i,l-1} + \alpha_0 v_{i,l} + \alpha_1 v_{i,l+1}.$$

It follows that

$$\begin{aligned}
 (\mathbf{B}_c \mathbf{v}_i)_l &= \alpha_0 \sin\left(\frac{il\pi}{k+1}\right) + \alpha_1 \left[\sin\left(\frac{i(l-1)\pi}{k+1}\right) + \sin\left(\frac{i(l+1)\pi}{k+1}\right) \right] \\
 &= \left[\alpha_0 + 2\alpha_1 \cos\left(\frac{i\pi}{k+1}\right) \right] \sin\left(\frac{il\pi}{k+1}\right) \\
 &= \lambda_i v_{i,l}
 \end{aligned}$$

$$\text{since} \quad \sin\left(\frac{i(l-1)\pi}{k+1}\right) + \sin\left(\frac{i(l+1)\pi}{k+1}\right) = 2 \cos\left(\frac{i\pi}{k+1}\right) \sin\left(\frac{il\pi}{k+1}\right).$$

This verifies established equations (2.19) and (2.20), by showing

$$\mathbf{B}_c \mathbf{v}_i = \lambda_i \mathbf{v}_i. \quad (2.21)$$

The spectral radius of $\mathbf{B}$, denoted $\rho(\mathbf{B})$, is defined as the maximum modulus of its eigenvalues in [39]. Since the transition matrix $\mathbf{B}$ is block diagonal with coefficient block $\mathbf{B}_c$ and harmonic block $d\mathbf{R}(\omega_j)$, its eigenvalues are the union of the eigenvalues of these blocks. The rotation block $d\mathbf{R}(\omega_j)$ has spectral radius d , since $\mathbf{R}(\omega_j)$ is orthogonal and its eigenvalues have modulus one (see [39, 40]). Therefore, using the eigenvalues verified above,

$$\rho(\mathbf{B}) = \max\{\rho(\mathbf{B}_c), d\}. \quad (2.22)$$

where

$$\rho(\mathbf{B}_c) = \max_{1 \leq i \leq k} \left| \alpha_0 + 2\alpha_1 \cos\left(\frac{i\pi}{k+1}\right) \right|. \quad (2.23)$$

If $\sigma_\omega^2 > 0$, setting $\rho(\mathbf{B}) < 1$ gives a stable state process $\mathbf{x}_t$ (see [41]). If $\sigma_\omega^2 = 0$, we set $d = 1$, to avoid the seasonal cycle decreasing in magnitude across time.

2.5 Observation equation components and evolution

Consistent with the modeling strategy discussed in Section 2.2, the altitude-dependent loading functions associated with the seasonal harmonic states are incorporated through the observation matrix, rather than by introducing additional state variables. This allows flexible representation of vertical structure while keeping the state dimension low.

For each $t \in \{1, \dots, T\}$, let

$$\tilde{\mathbf{y}}_t = \begin{pmatrix} \tilde{y}_{t,a_1} & \cdots & \tilde{y}_{t,a_{n_a}} \end{pmatrix}^\top \quad (2.24)$$

denote the noise-free ozone profile evaluated at the n_a altitude levels. The corresponding observation matrix in equation (2.3) is given by

$$\mathbf{Z} = \begin{pmatrix} \mathbf{B}_{sp} & \mathbf{Z}_H \end{pmatrix} \quad \mathbf{Z}_H = \begin{pmatrix} \mathbf{Z}_{c,1} & \mathbf{Z}_{s,1} & \cdots & \mathbf{Z}_{c,p} & \mathbf{Z}_{s,p} \end{pmatrix} \quad (2.25)$$

where $\mathbf{B}_{sp} \in \mathbb{R}^{n_a \times k}$ denotes the matrix of cubic B-spline basis functions evaluated at the altitude grid $\{a_1, \dots, a_{n_a}\}$, where k is the number of basis functions determined by the chosen knot sequence (see appendix A), and $\mathbf{Z}_{c,j} = \mathbf{B}_{sp} \boldsymbol{\gamma}_{c,j}$, $\mathbf{Z}_{s,j} = \mathbf{B}_{sp} \boldsymbol{\gamma}_{s,j}$ are smooth altitude-dependent loading functions.

Figure 2.1 and 2.2 illustrate a representative set of cubic B-spline basis functions evaluated over the altitude grid. Due to their compact support, each basis function affects only a local altitude region, allowing the model to flexibly capture localized vertical structure while maintaining numerical stability. Smaller values of k produce smoother and more global representations, while larger values allow finer altitude-dependent variation to be captured.




Figure 2.1: Cubic B-spline basis functions evaluated over the scaled altitude grid for $k = 7$. Each basis function has compact support and contributes locally to the representation of the vertical ozone structure.




Figure 2.2: Cubic B-spline basis functions evaluated over the scaled altitude grid for $k = 13$. Each basis function has compact support and contributes locally to the representation of the vertical ozone structure.

The noise-free part of the observation equation is given by

$$\tilde{\mathbf{y}}_t = \mathbf{Z}\mathbf{x}_t = \mathbf{B}_{sp}\mathbf{x}_{t,c} + \sum_{j=1}^p \mathbf{Z}_{c,j}M_{t,j} + \sum_{j=1}^p \mathbf{Z}_{s,j}M_{t,j}^* \quad (2.26)$$

Here, $\mathbf{B}_{sp}$ denotes the matrix of cubic B-spline basis functions evaluated at the altitude grid $\{a_1, \dots, a_{n_a}\}$, where k is the number of basis functions determined by the chosen knot sequence (see appendix A). The vectors $\boldsymbol{\gamma}_{c,j}$ and $\boldsymbol{\gamma}_{s,j}$ are coefficient vectors of length k that determine the vertical structure of the cosine and sine components, respectively. Consequently, $\mathbf{Z}_{c,j}$ and $\mathbf{Z}_{s,j}$ are smooth functions of altitude, each parameterized by k unknown coefficients, providing a parsimonious representation of the seasonal vertical structure. This formulation allows the altitude-dependent seasonal loadings to be represented through a low-dimensional parameterization via $\boldsymbol{\gamma}_{c,j}$ and $\boldsymbol{\gamma}_{s,j}$, while avoiding an increase in the state dimension.

The observation noise is assumed to be Gaussian and temporally independent,

$$\mathbf{v}_t = \mathbf{y}_t - \tilde{\mathbf{y}}_t \stackrel{\text{ind}}{\sim} N(\mathbf{0}, \mathbf{R}), \quad \mathbf{R} = \sigma_R^2 \mathbf{I}_{n_a}, \quad (2.27)$$

where σ_R^2 is an unknown variance parameter.

Let c_{ij} and b_{ij} denote the i th elements of column vectors $\mathbf{Z}_{c,j} = \mathbf{B}_{sp}\boldsymbol{\gamma}_{c,j}$ and $\mathbf{Z}_{s,j} = \mathbf{B}_{sp}\boldsymbol{\gamma}_{s,j}$ for $i = 1, \dots, n_a$, respectively, representing the altitude-dependent cosine and sine loadings in the j^{th} harmonic component. Specifically, denote components of these loadings as

$$\mathbf{Z}_{c,j} = \mathbf{B}_{sp}\boldsymbol{\gamma}_{c,j} = \begin{pmatrix} c_{1j} & \dots & c_{n_a j} \end{pmatrix}^\top \quad \mathbf{Z}_{s,j} = \mathbf{B}_{sp}\boldsymbol{\gamma}_{s,j} = \begin{pmatrix} b_{1j} & \dots & b_{n_a j} \end{pmatrix}^\top$$

The harmonic component could be represented through a cosine-sine pair, so we can write $M_{t,j} = \cos(\psi_{t,j})$ and $M_{t,j}^* = \sin(\psi_{t,j})$ for some angle $\psi_{t,j}$. Then the total seasonal

contribution of all harmonic components to the observation equation (2.3) for i^{th} altitude at time t , denoted as $S_t(a_i)$, would be

$$S_t(a_i) = \sum_{j=1}^p (c_{ij}M_{t,j} + b_{ij}M_{t,j}^*) \quad (2.28)$$

For each harmonic component j , the contribution of the j^{th} harmonic component at altitude a_i can be written as

$$\begin{aligned} c_{ij}M_{t,j} + b_{ij}M_{t,j}^* &= \sqrt{c_{ij}^2 + b_{ij}^2} \left(\frac{c_{ij}}{\sqrt{c_{ij}^2 + b_{ij}^2}} \cos(\psi_{t,j}) + \frac{b_{ij}}{\sqrt{c_{ij}^2 + b_{ij}^2}} \sin(\psi_{t,j}) \right) \\ &= \sqrt{c_{ij}^2 + b_{ij}^2} (\cos(\phi_{ij}) \cos(\psi_{t,j}) + \sin(\phi_{ij}) \sin(\psi_{t,j})) \\ &= \sqrt{c_{ij}^2 + b_{ij}^2} \cos(\psi_{t,j} - \phi_{ij}) \end{aligned} \quad (2.29)$$

where

$$\phi_{ij} = \text{atan2}(b_{ij}, c_{ij}) \quad (2.30)$$

denotes the phase shift associated with altitude level i and harmonic component j . Here, $\text{atan2}(y, x)$ is the two-argument arctangent function, which returns the angle $\phi \in (-\pi, \pi]$ satisfying $\cos \phi = x/\sqrt{x^2 + y^2}$ and $\sin \phi = y/\sqrt{x^2 + y^2}$. Using atan2 ensures that the phase ϕ_{ij} is uniquely defined for all $(c_{ij}, b_{ij}) \neq (0, 0)$ and correctly reflects the quadrant in which (c_{ij}, b_{ij}) lies.

Consequently, the total seasonal contribution at altitude a_i can be written as

$$S_t(a_i) = \sum_{j=1}^p \sqrt{c_{ij}^2 + b_{ij}^2} \cos(\psi_{t,j} - \phi_{ij}). \quad (2.31)$$

Thus, each harmonic component has its own altitude-dependent amplitude $\sqrt{c_{ij}^2 + b_{ij}^2}$ and phase shift ϕ_{ij} . When $p = 1$, the seasonal contribution at altitude a_i can be summarized

by a single amplitude and phase shift. When $p > 1$, the total seasonal effect is the sum of these harmonic contributions rather than a single cosine term with one common phase shift. Since the harmonic components have different frequencies $\omega_j = \frac{2\pi j}{T_p}$, the total seasonal contribution is a Fourier-type summation rather than a single shifted cosine. Therefore, for $p > 1$, it cannot in general be further reduced to one altitude-specific amplitude and one phase shift.

Using the harmonic state update equation (2.14), for each $j = 1, \dots, p$, we have

$$\begin{aligned}
& c_{ij}M_{t+1,j} + b_{ij}M_{t+1,j}^* \\
&= c_{ij} \left(d \cos(\omega_j) M_{t,j} + d \sin(\omega_j) M_{t,j}^* + \epsilon_{t+1,j,M} \right) \\
&\quad + b_{ij} \left(-d \sin(\omega_j) M_{t,j} + d \cos(\omega_j) M_{t,j}^* + \epsilon_{t+1,j,M^*} \right) \\
&= d (c_{ij} \cos(\omega_j) - b_{ij} \sin(\omega_j)) M_{t,j} + d (c_{ij} \sin(\omega_j) + b_{ij} \cos(\omega_j)) M_{t,j}^* \\
&\quad + c_{ij} \epsilon_{t+1,j,M} + b_{ij} \epsilon_{t+1,j,M^*} \\
&= d \sqrt{c_{ij}^2 + b_{ij}^2} \cos(\psi_{t,j} - (\phi_{ij} + \omega_j)) + c_{ij} \epsilon_{t+1,j,M} + b_{ij} \epsilon_{t+1,j,M^*} \tag{2.32}
\end{aligned}$$

$$c_{ij} \epsilon_{t+1,j,M} + b_{ij} \epsilon_{t+1,j,M^*} \sim N(0, (c_{ij}^2 + b_{ij}^2) \sigma_\omega^2) \quad \text{for } t \in \{1, \dots, T\} \tag{2.33}$$

For fixed altitude level i and harmonic component j , this innovation term is independent and identically distributed over time t . However, this identical-distribution statement does not generally hold across different altitude levels or harmonic components, since the variance depends on c_{ij} and b_{ij} .

Therefore, the total seasonal contribution at altitude a_i evolves as

$$S_{t+1}(a_i) = \sum_{j=1}^p \left[d \sqrt{c_{ij}^2 + b_{ij}^2} \cos(\psi_{t,j} - (\phi_{ij} + \omega_j)) + c_{ij} \epsilon_{t+1,j,M} + b_{ij} \epsilon_{t+1,j,M^*} \right] \tag{2.34}$$

Since the harmonic innovations are assumed mutually independent across j , the added

noise term in equation (2.34) satisfies

$$\sum_{j=1}^p (c_{ij}\epsilon_{t+1,j,M} + b_{ij}\epsilon_{t+1,j,M^*}) \sim N \left(0, \sigma_\omega^2 \sum_{j=1}^p (c_{ij}^2 + b_{ij}^2) \right) \quad (2.35)$$

For fixed altitude level i , the summed innovation term is independent and identically distributed over time t . However, its variance depends on the altitude-specific loadings $\{c_{ij}, b_{ij}\}_{j=1}^p$, so the identical-distribution statement does not generally hold across different altitude levels.

This shows that, under the harmonic rotation update equation (2.14), the contribution of each seasonal component remains a cosine function, with amplitude scaled by d and phase advanced by ω_j , up to stochastic innovation. For $p > 1$, the full seasonal effect at each altitude is the summation of these component-wise seasonal effects.

Equivalent reduced Fourier-form representations of harmonic seasonal components are also commonly used in dynamic linear models; for example, Johnson et al. [42] used such a formulation, which is reviewed in Section 3.2.

2.6 Model Identifiability

Following [43], a parametric model is said to be identifiable if distinct parameter values correspond to distinct probability distributions for the observable data. More formally, let $\boldsymbol{\theta}$ denote the vector of model parameters and let $p(\mathbf{y}|\boldsymbol{\theta})$ denote the probability distribution of the observable process. The model is identifiable if

$$p(\mathbf{y}|\boldsymbol{\theta}_1) = p(\mathbf{y}|\boldsymbol{\theta}_2) \quad \text{for all } \mathbf{y} \implies \boldsymbol{\theta}_1 = \boldsymbol{\theta}_2. \quad (2.36)$$

Here the notations are the same as in equation (4.1). If two distinct parameter values produce the same distribution of the observable process, then the model parameters are

not identifiable. In the following sections we examine identifiability issues arising in the proposed state-space model.

We now discuss model identifiability under the specification $p = 2$. In this case, the seasonal harmonic state vector is

$$\mathbf{H}_t = \begin{pmatrix} M_{t,1} & M_{t,1}^* & M_{t,2} & M_{t,2}^* \end{pmatrix}^\top. \quad (2.37)$$

The purpose of this chapter is not to establish a complete set of sufficient conditions for global identifiability of the full state-space model. Such a result is difficult for structured state-space models with latent states and constrained matrix parameterizations. Instead, we focus on the main observational equivalences that arise from the seasonal harmonic component of the proposed model. These equivalences motivate the identifying constraints used in the simulation and real-data analyses.

2.6.1 Identifiability between σ_ω^2 and $\gamma_{c,j}, \gamma_{s,j}$

We investigate the identifiability between the seasonal innovation variance σ_ω^2 and the seasonal loading vectors $\gamma_{c,j}, \gamma_{s,j}$.

Recall that the harmonic innovation covariance matrix is given by

$$\mathbf{Q}_H = \sigma_\omega^2 \mathbf{I}_4, \quad (2.38)$$

and the seasonal loading component of the observation matrix can be written as

$$\mathbf{Z}_H = \begin{pmatrix} \mathbf{Z}_{c,1} & \mathbf{Z}_{s,1} & \mathbf{Z}_{c,2} & \mathbf{Z}_{s,2} \end{pmatrix}, \quad \mathbf{Z}_{c,j} = \mathbf{B}_{sp} \gamma_{c,j}, \quad \mathbf{Z}_{s,j} = \mathbf{B}_{sp} \gamma_{s,j}. \quad (2.39)$$

The issue is that the observable seasonal contribution depends on the product of the seasonal loading matrix and the harmonic states, so the scale of the harmonic states and

the scale of the seasonal loading functions cannot be separately determined without an additional constraint.

Theorem 1 *Let $c > 0$ be a constant. Define the parameter transformation*

$$\sigma_\omega^2 \mapsto c^2 \sigma_\omega^2, \quad (\gamma_{c,j}, \gamma_{s,j}) \mapsto \left(\frac{1}{c} \gamma_{c,j}, \frac{1}{c} \gamma_{s,j} \right), \quad j = 1, 2. \quad (2.40)$$

If the harmonic initial state is transformed consistently, then the resulting state-space model induces the same distribution for $\{y_t\}_{t=1}^T$ as the original model.

Consequently, σ_ω^2 and $(\gamma_{c,j}, \gamma_{s,j})$ are not jointly identifiable.

Proof: Consider the linear transformation of the state vector

$$\mathbf{x}'_t = \mathbf{C} \mathbf{x}_t, \quad \mathbf{C} = \begin{pmatrix} \mathbf{I}_k & \mathbf{0} \\ \mathbf{0} & c \mathbf{I}_4 \end{pmatrix}. \quad (2.41)$$

Define the transformed system matrices

$$\mathbf{B}' = \mathbf{C} \mathbf{B} \mathbf{C}^{-1}, \quad \mathbf{Q}' = \mathbf{C} \mathbf{Q} \mathbf{C}^\top, \quad \mathbf{Z}' = \mathbf{Z} \mathbf{C}^{-1}. \quad (2.42)$$

Then the transformed state equation becomes

$$\mathbf{x}'_t = \mathbf{B}' \mathbf{x}'_{t-1} + \mathbf{w}'_t, \quad \mathbf{w}'_t \sim N(\mathbf{0}, \mathbf{Q}'). \quad (2.43)$$

From $\mathbf{Q}' = \mathbf{C} \mathbf{Q} \mathbf{C}^\top$, the harmonic block transforms as

$$\mathbf{Q}'_H = c^2 \sigma_\omega^2 \mathbf{I}_4, \quad (2.44)$$

which corresponds to the parameter mapping

$$\sigma_\omega^2 \mapsto c^2 \sigma_\omega^2. \quad (2.45)$$

Next, since $\mathbf{Z}' = \mathbf{Z}\mathbf{C}^{-1}$ rescales the harmonic components by $1/c$, we obtain

$$(\mathbf{Z}_{c,j}, \mathbf{Z}_{s,j}) \mapsto \left(\frac{1}{c} \mathbf{Z}_{c,j}, \frac{1}{c} \mathbf{Z}_{s,j} \right), \quad (2.46)$$

which is equivalent to

$$(\gamma_{c,j}, \gamma_{s,j}) \mapsto \left(\frac{1}{c} \gamma_{c,j}, \frac{1}{c} \gamma_{s,j} \right). \quad (2.47)$$

Finally, the observation equation satisfies

$$\mathbf{Z}\mathbf{x}_t = \mathbf{Z}\mathbf{C}^{-1}\mathbf{C}\mathbf{x}_t = \mathbf{Z}'\mathbf{x}'_t. \quad (2.48)$$

Thus, the noise-free observations are identical:

$$\tilde{\mathbf{y}}_t = \mathbf{Z}\mathbf{x}_t = \mathbf{Z}'\mathbf{x}'_t =: \tilde{\mathbf{y}}'_t. \quad (2.49)$$

With the observation noise unchanged, the two parameter sets induce the same distribution for $\{\mathbf{y}_t\}_{t=1}^T$. Therefore, the likelihood functions coincide, implying non-identifiability. ■

Therefore, σ_ω^2 and the seasonal loading vectors are not jointly identifiable in the unconstrained seasonal representation. To remove this scale ambiguity, we fix σ_ω^2 rather than estimate it jointly with the loading vectors.

The numerical value chosen for σ_ω^2 should be interpreted as a modeling choice as well as an identifying constraint. Larger fixed values allow stronger stochastic deviations of the harmonic states from the deterministic rotation, whereas smaller fixed values make

the seasonal dynamics closer to a deterministic harmonic cycle. The special case $\sigma_\omega^2 = 0$ corresponds to fully deterministic seasonal harmonic dynamics. Thus, once σ_ω^2 is fixed, the seasonal loading functions are interpreted relative to that fixed state scale.

2.6.2 Identifiability between $\gamma_{c,j}$, $\gamma_{s,j}$ and $M_{0,j}$, $M_{0,j}^*$

We next consider a rotational ambiguity between the seasonal loading vectors and the phase of the harmonic states. This ambiguity is easiest to see in the deterministic case, where the harmonic states evolve only through rotation. A similar rotational invariance also applies when $Q_H = \sigma_\omega^2 I_{2p}$, since block-wise rotations preserve this isotropic covariance structure.

Theorem 2 *Let $\tilde{\theta}_1, \tilde{\theta}_2 \in \mathbb{R}$ and define the block-diagonal rotation matrix*

$$\mathbf{T} = \begin{pmatrix} \mathbf{T}(\tilde{\theta}_1) & \mathbf{0} \\ \mathbf{0} & \mathbf{T}(\tilde{\theta}_2) \end{pmatrix}, \quad (2.50)$$

where for $j = 1, 2$,

$$\mathbf{T}(\tilde{\theta}_j) = \begin{pmatrix} \cos(\tilde{\theta}_j) & \sin(\tilde{\theta}_j) \\ -\sin(\tilde{\theta}_j) & \cos(\tilde{\theta}_j) \end{pmatrix}. \quad (2.51)$$

Define the transformations

$$\mathbf{H}_t \mapsto \mathbf{T}\mathbf{H}_t, \quad (2.52)$$

$$\mathbf{Z}_H \mapsto \mathbf{Z}_H \mathbf{T}^\top, \quad (2.53)$$

where $\mathbf{Z}_H$ is defined in equation (2.39).

Then the resulting model produces the same deterministic seasonal component

$$\mathbf{Z}_H \mathbf{H}_t. \quad (2.54)$$

Consequently, $(\gamma_{c,j}, \gamma_{s,j})$ and $(M_{0,j}, M_{0,j}^*)$, for $j = 1, 2$, are not jointly identifiable.

Proof: For each harmonic $j = 1, 2$, the deterministic seasonal dynamics satisfy

$$\begin{pmatrix} M_{t+1,j} \\ M_{t+1,j}^* \end{pmatrix} = \mathbf{R}(\omega_j) \begin{pmatrix} M_{t,j} \\ M_{t,j}^* \end{pmatrix}. \quad (2.55)$$

Since both $\mathbf{R}(\omega_j)$ and $\mathbf{T}(\tilde{\theta}_j)$ are rotation matrices, they commute:

$$\mathbf{T}(\tilde{\theta}_j) \mathbf{R}(\omega_j) = \mathbf{R}(\omega_j) \mathbf{T}(\tilde{\theta}_j). \quad (2.56)$$

Thus, for each j ,

$$\mathbf{T}(\tilde{\theta}_j) \begin{pmatrix} M_{t+1,j} \\ M_{t+1,j}^* \end{pmatrix} = \mathbf{R}(\omega_j) \mathbf{T}(\tilde{\theta}_j) \begin{pmatrix} M_{t,j} \\ M_{t,j}^* \end{pmatrix}. \quad (2.57)$$

Stacking the two harmonics yields

$$\mathbf{T} \mathbf{H}_{t+1} = \mathbf{R} \mathbf{T} \mathbf{H}_t, \quad (2.58)$$

where $\mathbf{R}$ is the block-diagonal transition matrix.

Hence, the transformed states follow the same deterministic dynamics.

Next, consider the seasonal component in the observation equation:

$$\mathbf{Z}_H \mathbf{H}_t = \mathbf{Z}_H \mathbf{T}^\top \mathbf{T} \mathbf{H}_t \quad (2.59)$$

$$= (\mathbf{Z}_H \mathbf{T}^\top)(\mathbf{T} \mathbf{H}_t). \quad (2.60)$$

Thus, the transformed parameters and states produce the same seasonal component for all t . This means that changing the phase convention of the harmonic states can be exactly offset by an opposite change in the seasonal loading matrix.

Therefore, the two parameterizations are observationally equivalent, implying that $(\gamma_{c,j}, \gamma_{s,j})$ and $(M_{0,j}, M_{0,j}^*)$ are not jointly identifiable. ■

To remove this rotational ambiguity, we impose the constraint

$$\begin{pmatrix} M_{0,j} \\ M_{0,j}^* \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad j = 1, 2. \quad (2.61)$$

2.6.3 Summary of Identifying Constraints

The discussion above shows that the seasonal harmonic component contains both scale and phase ambiguities if no identifying constraints are imposed. The scaling argument in Subsection 2.6.1 applies to the unconstrained seasonal representation, where the harmonic initial states may also be rescaled. The rotational argument in Subsection 2.6.2 shows that the phase of the harmonic states and the phase encoded in the loading vectors cannot be separated without fixing a phase convention. In the model fitted in this thesis, we therefore fix the harmonic initial states at $(M_{0,j}, M_{0,j}^*) = (1, 0)$ and fix the seasonal state noise variance σ_ω^2 . These restrictions define a stable and interpretable parameterization of the seasonal component. The fixed value of σ_ω^2 should be understood as a modeling choice that determines the stochastic scale of the harmonic states, rather than as a quantity

estimated from the data.

Chapter 3

Related Work and Contributions

Several approaches have been proposed in the literature to model ozone profiles across altitude and time. These methods differ in how they handle vertical dependence, temporal dynamics, and uncertainty quantification. Since the models reviewed below use different formulations and notation, we describe them using notation that is as consistent as possible with the notation introduced in this thesis. This notation is used only to facilitate comparison; it does not imply that the models are identical in structure.

3.1 Spline-based Joint Modeling Across Altitude

Chang et al. [32] proposed a statistical regularization method for detecting long-term trends from sparse tropospheric ozone profile data. Their approach treats the profile time series as a collection of correlated time series observed on discrete pressure surfaces or altitude levels and fits these series jointly using a smoothing spline ANOVA framework with roughness penalties.

Let $y_t(a_i)$ denote the ozone observation at time t and altitude level a_i , where $i = 1, \dots, n_a$. In Chang et al. [32], the observations are first aggregated to monthly means.

For a single altitude or pressure level, the basic time series decomposition is

$$y_t(a_i) = \text{seasonal cycle} + \text{trend} + \text{residual}. \quad (3.1)$$

To model the full vertical profile jointly, they decompose the time index into a monthly seasonal index and an interannual index. Let $q(t) \in \{1, \dots, 12\}$ denote the month corresponding to time t , and let $r(t) \in \{1, \dots, n_y\}$ denote the year or interannual index. Their two-dimensional profile model can then be written as

$$y_t(a_i) = f_{\text{seasonal}}\{q(t), a_i\} + f_{\text{interannual}}\{r(t), a_i\} + \epsilon_t(a_i), \quad (3.2)$$

where the first term $f_{\text{seasonal}}\{q(t), a_i\}$ represents the altitude-dependent seasonal cycle, $f_{\text{interannual}}\{r(t), a_i\}$ represents the de-seasonalized interannual variation, and $\epsilon_t(a_i)$ is the residual term. In vector form, this can be written as

$$\mathbf{y}_t = \mathbf{f}_{\text{seasonal}}\{q(t)\} + \mathbf{f}_{\text{interannual}}\{r(t)\} + \boldsymbol{\epsilon}_t, \quad (3.3)$$

where

$$\mathbf{f}_{\text{seasonal}}\{q(t)\} = \begin{pmatrix} f_{\text{seasonal}}\{q(t), a_1\} \\ \vdots \\ f_{\text{seasonal}}\{q(t), a_{n_a}\} \end{pmatrix}, \quad (3.4)$$

$$\mathbf{f}_{\text{interannual}}\{r(t)\} = \begin{pmatrix} f_{\text{interannual}}\{r(t), a_1\} \\ \vdots \\ f_{\text{interannual}}\{r(t), a_{n_a}\} \end{pmatrix}. \quad (3.5)$$

Both components are represented using two-dimensional penalized regression spline basis expansions. In notation adapted to the altitude grid $a_1, \dots, a_{n_a}$, their smooth

components can be written as

$$f_{\text{seasonal}}(q, a) = \sum_{\ell=1}^{K_1} c_{\ell}^{(s)} \psi_{\ell}^{(s)}(q, a), \quad (3.6)$$

and

$$f_{\text{interannual}}(r, a) = \sum_{\ell=1}^{K_2} c_{\ell}^{(I)} \psi_{\ell}^{(I)}(r, a), \quad (3.7)$$

where $\psi_{\ell}^{(s)}(q, a)$ and $\psi_{\ell}^{(I)}(r, a)$ are two-dimensional smoothing spline basis functions evaluated over the seasonal-altitude and interannual-altitude grids, respectively. The coefficients are estimated by minimizing a penalized least squares criterion of the form

$$\sum_{t,i} [y_t(a_i) - f_{\text{seasonal}}\{q(t), a_i\} - f_{\text{interannual}}\{r(t), a_i\}]^2 + D^2 F, \quad (3.8)$$

where the norm is taken over the observed time and altitude grid, and $D^2 F$ is a roughness penalty involving second-order partial derivatives of the fitted smooth surfaces. This penalty discourages overly rough seasonal and interannual surfaces and allows information to be borrowed across neighboring altitude or pressure levels.

The main similarity between Chang et al. [32] and the present work is that both approaches model ozone profiles jointly across altitude and use smooth basis-function representations to capture vertical dependence. However, the role of the basis functions is different. In Chang et al. [32], the spline basis functions directly define smooth two-dimensional regression surfaces over time and altitude. In our proposed model, the B-spline basis functions are evaluated only over altitude and are used inside the observation matrix of a state-space model. Specifically, the noise-free observation equation in this thesis is described in equation (2.26). Thus, the seasonal variation in the proposed model is represented through latent harmonic states $M_{t,j}$ and $M_{t,j}^*$, while the altitude dependence of the seasonal amplitude and phase is represented through the B-spline

coefficient vectors $\gamma_{c,j}$ and $\gamma_{s,j}$.

The two approaches also differ in their treatment of temporal dependence. Chang et al. [32] model temporal variation through smooth seasonal and interannual functions, with regularization used to reduce the influence of short-lived anomalies. Their method is therefore well suited for estimating smooth long-term trend structure from sparse profile observations. In contrast, our proposed model specifies an explicit stochastic state equation for the latent spline coefficients and harmonic states. Therefore, temporal dependence and uncertainty propagation are handled through the state-space evolution and Kalman filtering framework, rather than through penalized regression surfaces alone.

3.2 Reduced Fourier-Form Dynamic Linear Models

Johnson et al. [42] used a Bayesian reduced Fourier-form dynamic linear model for estimating phenological events from univariate remote sensing time series, commonly used in seasonally varying state space models (e.g., [22]). Although their applications differ from the functional ozone profile setting considered here, our model incorporates closely related harmonic components within a functional state space model. We review their approach and reformulation within the components of our model here.

Consider a univariate time series observed at a single location. Let p denote the number of retained harmonic pairs, corresponding to q in Johnson et al. [42], and let $(S_{t,j}, S_{t,j}^*)$ denote the j th harmonic pair. The reduced Fourier-form observation equation can be written as

$$Y_t = \sum_{j=1}^p S_{t,j} + v_t, \quad (3.9)$$

where only the first component $S_{t,j}$ of each harmonic pair enters the observation equation. The conjugate component $S_{t,j}^*$ is included so that the harmonic pair can evolve through

a rotation matrix:

$$S_{t+1,j} = \cos(\omega_j)S_{t,j} + \sin(\omega_j)S_{t,j}^* + w_{t+1,j}, \quad (3.10)$$

$$S_{t+1,j}^* = -\sin(\omega_j)S_{t,j} + \cos(\omega_j)S_{t,j}^* + w_{t+1,j}^*. \quad (3.11)$$

Equivalently, if

$$\mathbf{S}_{t,j} = \begin{pmatrix} S_{t,j} \\ S_{t,j}^* \end{pmatrix}, \quad \mathbf{w}_{t+1,j} = \begin{pmatrix} w_{t+1,j} \\ w_{t+1,j}^* \end{pmatrix}, \quad (3.12)$$

then the harmonic evolution can be written as

$$\mathbf{S}_{t+1,j} = \mathbf{R}(\omega_j)\mathbf{S}_{t,j} + \mathbf{w}_{t+1,j}, \quad (3.13)$$

where

$$\mathbf{R}(\omega_j) = \begin{pmatrix} \cos(\omega_j) & \sin(\omega_j) \\ -\sin(\omega_j) & \cos(\omega_j) \end{pmatrix}. \quad (3.14)$$

Thus, in the reduced Fourier-form dynamic linear model, the observation equation selects only the first component of each harmonic pair, while the second component is needed for the rotation-based state evolution.

The seasonal component in the present thesis can be viewed as an altitude-profile extension of an equivalent seasonal representation. Recall from equation (2.28) that, for altitude level a_i , the seasonal contribution is

$$S_t(a_i) = \sum_{j=1}^p (c_{ij}M_{t,j} + b_{ij}M_{t,j}^*), \quad (3.15)$$

where c_{ij} and b_{ij} are the i^{th} elements of the loading vectors $\mathbf{Z}_{c,j}$ and $\mathbf{Z}_{s,j}$, respectively. In contrast to equation (3.9), both components $M_{t,j}$ and $M_{t,j}^*$ appear in the observation

equation through the altitude-dependent loadings c_{ij} and b_{ij} .

However, for each fixed altitude level a_i and harmonic component j , the term $c_{ij}M_{t,j} + b_{ij}M_{t,j}^*$ can be rewritten as the first component of a transformed harmonic pair. Define

$$\tilde{S}_{t,j}^{(i)} = c_{ij}M_{t,j} + b_{ij}M_{t,j}^*, \quad (3.16)$$

and define the corresponding conjugate component by

$$\tilde{S}_{t,j}^{*(i)} = -b_{ij}M_{t,j} + c_{ij}M_{t,j}^*. \quad (3.17)$$

In vector form,

$$\begin{pmatrix} \tilde{S}_{t,j}^{(i)} \\ \tilde{S}_{t,j}^{*(i)} \end{pmatrix} = \mathbf{C}_{ij} \begin{pmatrix} M_{t,j} \\ M_{t,j}^* \end{pmatrix}, \quad \text{for} \quad \mathbf{C}_{ij} = \begin{pmatrix} c_{ij} & b_{ij} \\ -b_{ij} & c_{ij} \end{pmatrix}. \quad (3.18)$$

The harmonic pair in our proposed model evolves according to

$$\begin{pmatrix} M_{t+1,j} \\ M_{t+1,j}^* \end{pmatrix} = d\mathbf{R}(\omega_j) \begin{pmatrix} M_{t,j} \\ M_{t,j}^* \end{pmatrix} + \begin{pmatrix} \epsilon_{t+1,j,M} \\ \epsilon_{t+1,j,M^*} \end{pmatrix}, \quad (3.19)$$

where $\mathbf{R}(\omega_j)$ is defined as in equation (3.14), and d is the damping factor. Multiplying both sides of equation (3.19) by $\mathbf{C}_{ij}$ gives

$$\begin{pmatrix} \tilde{S}_{t+1,j}^{(i)} \\ \tilde{S}_{t+1,j}^{*(i)} \end{pmatrix} = d\mathbf{C}_{ij}\mathbf{R}(\omega_j) \begin{pmatrix} M_{t,j} \\ M_{t,j}^* \end{pmatrix} + \mathbf{C}_{ij} \begin{pmatrix} \epsilon_{t+1,j,M} \\ \epsilon_{t+1,j,M^*} \end{pmatrix}. \quad (3.20)$$

Note that:

$$\begin{aligned}
\mathbf{C}_{ij}\mathbf{R}(\omega_j) &= \begin{pmatrix} c_{ij} & b_{ij} \\ -b_{ij} & c_{ij} \end{pmatrix} \begin{pmatrix} \cos(\omega_j) & \sin(\omega_j) \\ -\sin(\omega_j) & \cos(\omega_j) \end{pmatrix} \\
&= \begin{pmatrix} c_{ij}\cos(\omega_j) - b_{ij}\sin(\omega_j) & c_{ij}\sin(\omega_j) + b_{ij}\cos(\omega_j) \\ -c_{ij}\sin(\omega_j) - b_{ij}\cos(\omega_j) & c_{ij}\cos(\omega_j) - b_{ij}\sin(\omega_j) \end{pmatrix} \\
&= \begin{pmatrix} \cos(\omega_j) & \sin(\omega_j) \\ -\sin(\omega_j) & \cos(\omega_j) \end{pmatrix} \begin{pmatrix} c_{ij} & b_{ij} \\ -b_{ij} & c_{ij} \end{pmatrix} \\
&= \mathbf{R}(\omega_j)\mathbf{C}_{ij}.
\end{aligned} \tag{3.21}$$

Using equation (3.21), equation (3.20) becomes

$$\begin{pmatrix} \tilde{S}_{t+1,j}^{(i)} \\ \tilde{S}_{t+1,j}^{*(i)} \end{pmatrix} = d\mathbf{R}(\omega_j) \begin{pmatrix} \tilde{S}_{t,j}^{(i)} \\ \tilde{S}_{t,j}^{*(i)} \end{pmatrix} + \begin{pmatrix} \eta_{t+1,j}^{(i)} \\ \eta_{t+1,j}^{*(i)} \end{pmatrix}, \tag{3.22}$$

where

$$\begin{pmatrix} \eta_{t+1,j}^{(i)} \\ \eta_{t+1,j}^{*(i)} \end{pmatrix} = \mathbf{C}_{ij} \begin{pmatrix} \epsilon_{t+1,j,M} \\ \epsilon_{t+1,j,M^*} \end{pmatrix}. \tag{3.23}$$

Expanding equation (3.22), we obtain

$$\tilde{S}_{t+1,j}^{(i)} = d\cos(\omega_j)\tilde{S}_{t,j}^{(i)} + d\sin(\omega_j)\tilde{S}_{t,j}^{*(i)} + \eta_{t+1,j}^{(i)}, \tag{3.24}$$

$$\tilde{S}_{t+1,j}^{*(i)} = -d\sin(\omega_j)\tilde{S}_{t,j}^{(i)} + d\cos(\omega_j)\tilde{S}_{t,j}^{*(i)} + \eta_{t+1,j}^{*(i)}. \tag{3.25}$$

Therefore, for any fixed altitude level a_i , the transformed harmonic pair $(\tilde{S}_{t,j}^{(i)}, \tilde{S}_{t,j}^{*(i)})$ follows the same rotation-based state evolution as the reduced Fourier-form dynamic linear model, apart from the damping factor d . When $d = 1$, the state evolution has exactly

the same rotation form as equations (3.10)- (3.11); when $d < 1$, it is the corresponding damped version.

Under this reparameterization, the seasonal contribution at altitude a_i becomes

$$S_t(a_i) = \sum_{j=1}^p \tilde{S}_{t,j}^{(i)}. \quad (3.26)$$

This has the same structure as the reduced Fourier-form observation equation (3.9). Hence, for a fixed altitude level, the seasonal component of the proposed model is mathematically equivalent to a reduced Fourier-form dynamic linear model written in terms of the transformed states $\tilde{S}_{t,j}^{(i)}$.

The difference is mainly one of model organization and interpretation. In a univariate reduced Fourier-form dynamic linear model, the amplitude and phase information are absorbed into the harmonic states $S_{t,j}$. In the present multivariate altitude-profile model, the temporal harmonic states $(M_{t,j}, M_{t,j}^*)$ are shared across all altitude levels, while the altitude-dependent amplitude and phase are represented through the loading functions $\mathbf{Z}_{c,j}$ and $\mathbf{Z}_{s,j}$ defined in equation (2.25).

The advantage of our formulation incorporating B-splines in the ozone profile setting is that it avoids introducing a separate harmonic state pair for every altitude level. If separate transformed harmonic states were introduced for each altitude and each harmonic component, the seasonal state dimension would increase to $2pn_a$. In our proposed formulation, the harmonic states are shared across altitudes, so the seasonal state dimension remains $2p$, while the altitude dependence of the seasonal contribution is modeled smoothly through the B-spline loading functions in the observation matrix. This reduces the state dimension and thereby improves computational efficiency [36].

3.3 State-Space Models At Individual Altitudes

Laine et al. [23] proposed a dynamic linear model for analyzing time-varying trends in stratospheric ozone time series. Their study uses combined SAGE II and GOMOS satellite ozone profile observations from 1984 - 2011. The vertical ozone profiles are interpolated onto an altitude grid and monthly averages are formed for each altitude and latitude bin. The state-space model is then applied separately to individual altitude levels or altitude–latitude regions.

Using notation adapted to the present thesis, let $y_t(a_i)$ denote the monthly mean ozone value at time t for a given altitude level or altitude region a_i . Laine et al. [23] write their model using the notation $\mathbf{F}_t$ for the observation matrix and $\mathbf{G}_t$ for the state transition matrix. To make the comparison with the notation in this thesis clearer, we denote the corresponding matrices by $\mathbf{Z}_t^{(L)}$ and $\mathbf{B}^{(L)}(a_i)$, where the superscript (L) indicates the model of Laine et al. For a fixed altitude level or altitude region a_i , their linear Gaussian state-space model can then be written as

$$y_t(a_i) = \mathbf{Z}_t^{(L)} \mathbf{x}_t^{(L)}(a_i) + v_{t,i}, \quad v_{t,i} \sim N(0, \sigma_{\text{obs},t}^2(a_i)), \quad (3.27)$$

$$\mathbf{x}_t^{(L)}(a_i) = \mathbf{B}^{(L)}(a_i) \mathbf{x}_{t-1}^{(L)}(a_i) + \mathbf{w}_t(a_i), \quad \mathbf{w}_t(a_i) \sim N(\mathbf{0}, \mathbf{W}(a_i)), \quad (3.28)$$

where $\mathbf{x}_t^{(L)}(a_i)$ is a latent state vector containing components for the local mean level, local trend, seasonal oscillations, proxy effects, and residual autocorrelation for the ozone time series at altitude level or altitude region a_i .

Specifically, their full state vector for a fixed altitude level or altitude region a_i is

$$\mathbf{x}_t^{(L)}(a_i) = \left(\mu_{t,i} \quad \kappa_{t,i} \quad u_{t,1,i} \quad u_{t,1,i}^* \quad u_{t,2,i} \quad u_{t,2,i}^* \quad \beta_{1,i} \quad \beta_{2,i} \quad \beta_{3,i} \quad \beta_{4,i} \quad \eta_{t,i} \right)^\top. \quad (3.29)$$

Here $\mu_{t,i}$ is the local mean or background ozone level, $\kappa_{t,i}$ is the local trend, $(u_{t,k,i}, u_{t,k,i}^*)$ are seasonal harmonic state pairs, $\beta_{1,i}, \dots, \beta_{4,i}$ are regression coefficients for proxy variables, and $\eta_{t,i}$ is an autoregressive state component of order one.

With this state vector, the observation equation can be written in expanded form as

$$y_t(a_i) = \mu_{t,i} + u_{t,1,i} + u_{t,2,i} + z_{1,t}\beta_{1,i} + z_{2,t}\beta_{2,i} + z_{3,t}\beta_{3,i} + z_{4,t}\beta_{4,i} + \eta_{t,i} + v_{t,i}, \quad (3.30)$$

where $z_{1,t}, \dots, z_{4,t}$ denote the proxy time series used to represent solar activity, QBO, and ENSO effects. Equivalently, the observation matrix is

$$\mathbf{Z}_t^{(L)} = \begin{pmatrix} 1 & 0 & 1 & 0 & 1 & 0 & z_{1,t} & z_{2,t} & z_{3,t} & z_{4,t} & 1 \end{pmatrix}. \quad (3.31)$$

Thus, $\mathbf{Z}_t^{(L)}$ selects the local mean component, the first component of each seasonal harmonic pair, the proxy-variable contributions, and the autoregressive residual component.

The trend component is specified through a local linear trend model. Since Laine et al. [23] fit the model separately for each altitude level or altitude–latitude region, the trend states should be interpreted as specific to the time series being analyzed. For a fixed altitude level or altitude region a_i , this component can be written as

$$\mu_{t,i} = \mu_{t-1,i} + \kappa_{t-1,i} + \epsilon_{\text{level},t,i}, \quad (3.32)$$

$$\kappa_{t,i} = \kappa_{t-1,i} + \epsilon_{\text{trend},t,i}. \quad (3.33)$$

Here $\mu_{t,i}$ represents the slowly varying background ozone level for that altitude level or altitude region, while $\kappa_{t,i}$ represents the local change in that background level from one month to the next. Equivalently, for each fitted univariate time series, the corresponding

transition block is

$$\mathbf{B}_{\text{trend}} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}. \quad (3.34)$$

Thus, the same local linear trend structure is used across altitude levels or altitude regions, but the latent background level and local trend are estimated separately for each individual ozone time series.

Seasonality is modeled using harmonic state pairs. For monthly data, the k^{th} harmonic component is updated by the rotation matrix

$$\mathbf{B}_{\text{seas}}(k) = \begin{pmatrix} \cos(2\pi k/12) & \sin(2\pi k/12) \\ -\sin(2\pi k/12) & \cos(2\pi k/12) \end{pmatrix}, \quad k = 1, 2. \quad (3.35)$$

For the time series at altitude level or altitude region a_i , if $u_{t,k,i}$ and $u_{t,k,i}^*$ denote the two states associated with the k^{th} seasonal harmonic, then

$$\begin{pmatrix} u_{t+1,k,i} \\ u_{t+1,k,i}^* \end{pmatrix} = \mathbf{B}_{\text{seas}}(k) \begin{pmatrix} u_{t,k,i} \\ u_{t,k,i}^* \end{pmatrix}. \quad (3.36)$$

This harmonic representation is closely related to the seasonal rotation structure used in the present thesis.

In addition to trend and seasonal components, Laine et al. [23] include proxy variables for solar activity, QBO, and ENSO through the observation equation. The proxy coefficients are treated as time-invariant, so the proxy transition block is

$$\mathbf{B}_{\text{proxy}} = \mathbf{I}_4. \quad (3.37)$$

For the time series at altitude level or altitude region a_i , the autoregressive residual

component follows an AR(1) transition:

$$\eta_{t,i} = \rho_i \eta_{t-1,i} + \epsilon_{\text{AR},t,i}, \quad (3.38)$$

so that

$$\mathbf{B}_{\text{AR},i} = (\rho_i). \quad (3.39)$$

Combining these components, the full transition matrix is block diagonal:

$$\mathbf{B}^{(L)}(a_i) = \begin{pmatrix} \mathbf{B}_{\text{trend}} & 0 & 0 & 0 & 0 \\ 0 & \mathbf{B}_{\text{seas}}(1) & 0 & 0 & 0 \\ 0 & 0 & \mathbf{B}_{\text{seas}}(2) & 0 & 0 \\ 0 & 0 & 0 & \mathbf{B}_{\text{proxy}} & 0 \\ 0 & 0 & 0 & 0 & \mathbf{B}_{\text{AR},i} \end{pmatrix}. \quad (3.40)$$

The state noise covariance matrix is specified as diagonal:

$$\text{diag}(\mathbf{W}(a_i)) = \begin{pmatrix} 0, & \sigma_{\text{trend},i}^2, & \sigma_{\text{seas},i}^2, & \sigma_{\text{seas},i}^2, & \sigma_{\text{seas},i}^2, & \sigma_{\text{seas},i}^2, & 0, & 0, & 0, & 0, & \sigma_{\text{AR},i}^2 \end{pmatrix}. \quad (3.41)$$

Thus, the local trend, seasonal states, and autoregressive residual component are allowed to evolve stochastically, while the local level innovation and the proxy coefficient innovations are set to zero in their final model specification.

For estimation, Laine et al. [23] use Kalman filtering and smoothing together with Bayesian inference for the auxiliary model parameters. For the time series at altitude level or altitude region a_i , the unknown auxiliary parameters are

$$\boldsymbol{\theta}_i = \begin{pmatrix} \sigma_{\text{trend},i} & \sigma_{\text{seas},i} & \sigma_{\text{AR},i} & \rho_i \end{pmatrix}^\top. \quad (3.42)$$

For fixed θ_i , the Kalman filter provides the likelihood of the observed time series, and the Kalman smoother provides estimates of the latent states. They then use Markov chain Monte Carlo to sample from the posterior distribution of θ_i . Conditional on sampled values of θ_i , a Kalman simulation smoother is used to generate realizations of the latent state sequence. This procedure propagates both latent-state uncertainty and auxiliary-parameter uncertainty into the trend estimates.

The trend is estimated from the posterior distribution of the background level $\mu_{t,i}$. For monthly data, Laine et al. [23] define a 10-year trend at time t as

$$\text{trend}_i(t) = \mu_{t+60,i} - \mu_{t-60,i}, \quad (3.43)$$

where the time index is measured in months. Repeating this calculation over posterior draws of $\mu_{t,i}$ gives posterior summaries and uncertainty intervals for the time-varying trend.

The main similarity between Laine et al. [23] and the present work is that both use a linear Gaussian state-space framework and both represent seasonality through harmonic state pairs evolving by a rotation matrix. Both approaches also use Kalman filtering and smoothing to estimate latent states and quantify uncertainty.

However, the modeling goals and structure differ. Laine et al. [23] focus on estimating time-varying trends in ozone time series at individual altitude levels or altitude–latitude regions. Their state vector is designed to decompose a single ozone time series into background level, trend, seasonal, proxy, and residual components. In contrast, our proposed model treats the full ozone profile at time t as the observation vector $\mathbf{y}_t$ and models all altitude levels jointly, see equation (2.26). The B-spline basis matrix $\mathbf{B}_{sp}$ allows the proposed model to represent smooth vertical dependence directly within the observation equation. Therefore, information is shared across altitudes through the spline

basis structure, rather than by fitting separate state-space models at individual altitude levels.

Another difference is the role of the harmonic states. In Laine et al. [23], the seasonal harmonic states form part of a univariate time-series decomposition for a given altitude or altitude region. In the present model, the harmonic states are shared across the altitude profile, while their altitude-dependent amplitudes and phases are controlled by the loading functions $\mathbf{B}_{sp}\boldsymbol{\gamma}_{c,j}$ and $\mathbf{B}_{sp}\boldsymbol{\gamma}_{s,j}$. This construction preserves the dynamic state-space treatment of seasonal variation while explicitly modeling the vertical structure of the ozone profile.

3.4 Functional Modeling of Entire Ozone Profiles

Meiring [30] proposed a functional data analysis approach to study altitude-dependent patterns of variation in ozonesonde profiles. The analysis focused on Hohenpeissenberg ozonesonde observations over the altitude range from 15.5 km to 30.5 km from January to July of 1967 - 1998. Each ozonesonde flight was treated as an observed curve over altitude, so that a sequence of flights formed a time series of functional observations. This formulation allowed ozone variability to be studied jointly across altitudes, rather than fitting separate regression models at individual altitude levels.

The main goal of the analysis was to estimate altitude-dependent nonlinear time trends and ozone variation associated with the quasi-biennial oscillation (QBO). In contrast to standard multiple regression approaches, where linear or piecewise-linear trends are typically estimated separately by altitude, Meiring [30] used a hierarchical functional modeling strategy. The observed log ozone partial pressure profile was first mean-filtered and standardized. Functional principal components were then used to obtain a low-dimensional basis representation of the standardized profile variability. The approxima-

tion retains the first K functional principal components, ordered by decreasing explained variation. Let $B_k(a)$ denote the continuous basis function representing the k^{th} functional principal component, evaluated at altitude a . Then the standardized ozone profile at time t_j was approximated by

$$y(a, t_j) \approx \sum_{k=1}^K c_k(t_j) B_k(a), \quad (3.44)$$

where the coefficients $c_k(t_j)$ describe the temporal evolution of the dominant modes of vertical profile variability.

The coefficient time series were then modeled as follows. For each month m and k^{th} principal component, a smoothing spline analysis of variance (SSANOVA) model was fitted:

$$c_k(t) = \mu_{k,m} + f_{k,m}(q_t) + g_{k,m}(y_t) + \eta_{k,m}(t), \quad t \text{ within month } m, \quad (3.45)$$

where q_t is a QBO index and y_t denotes calendar year. Here $\mu_{k,m}$ is the month-specific intercept for the k^{th} coefficient model. The function $f_{k,m}(q_t)$ represents the QBO-related component of coefficient variability modeled by a periodic spline, while $g_{k,m}(y_t)$ represents the year effect component modeled using a cubic spline. Then for each k and m , $\mu_{k,m} + g_{k,m}(y_t)$ together represents the year effect, including both the month-specific baseline and the nonlinear temporal departure from that baseline.

These coefficient-level functions were then combined with the functional principal component basis functions and transformed back to the original profile scale to obtain altitude-dependent estimates of QBO effects and long-term time trends, both of which also vary by month. Specifically, the QBO-related profile component and the year-effect

profile component are given by

$$Q(a, q, m) = \sum_{k=1}^K f_{k,m}(q) B_k(a), \quad (3.46)$$

and

$$G(a, y, m) = \sum_{k=1}^K \{\mu_{k,m} + g_{k,m}(y)\} B_k(a), \quad (3.47)$$

respectively. After multiplication by the altitude-dependent standardization curve and addition of the estimated seasonal mean field, these components are transformed back to the original log ozone profile scale. This gives altitude-dependent estimates of QBO effects and long-term time trends. The study found clear altitude-dependent QBO features in some months and provided evidence of nonlinear time variation, including possible approximately 11-year oscillatory behavior in ozone levels.

This work is closely related to the present thesis because both approaches consider ozone profiles as functional observations over altitude and aim to model vertical structure and temporal evolution jointly. However, for computational reasons, the approach in Meiring [30] did not formulate the entire profile evolution in a state-space model. Instead, mean-filtering with standardization and functional principal component analysis were performed in an initial step, followed by a separate step in which serial dependence of dominant modes of variability was represented through smoothing SSANOVA-based evolution of the latent functional principal component coefficients, performed separately by coefficient, relying on approximate orthogonality of the functional principal components. Covariate effects were incorporated within the SSANOVA coefficient models, separately by coefficient, before being combined with the basis functions and transformed back to the original profile scale.

As a result, cross-coefficient temporal evolution dependence is not explicitly modeled

through a stochastic process in Meiring [30]. The coefficient models are fitted separately, so the dependence among different functional principal component coefficients is not represented through a joint transition equation. In contrast, the model developed in this thesis places the latent spline coefficients in a unified state-space model; the temporal dependence among the nonseasonal spline coefficients is modeled through the tridiagonal transition matrix $\mathbf{B}_c$, which allows cross-coefficient dependence to be represented directly in the latent stochastic process.

Another difference is the choice of dimension-reducing basis functions used to approximate the ozonesonde profiles as a function of altitude. Functional principal component basis functions are data-adaptive and efficient for dimension reduction, but they are in general global functions over altitude and are not compactly supported. Consequently, a change in one coefficient can affect the fitted profile over a broad altitude range. In this thesis, compactly supported B-spline basis functions are used to approximate the profile as a function of altitude. This provides a more local representation of altitude dependence and makes the coefficient interpretation more directly connected to neighboring altitude regions. The local support of B-splines also supports the use of structured transition matrices for the latent coefficient process.

The treatment of seasonality also differs. In Meiring [30], seasonal behavior is handled through mean-filtering and month-specific coefficient models, specifically to investigate seasonal variation in QBO associations and time trends. In the model proposed in this thesis, the seasonal component is estimated within the state-space framework. The harmonic seasonal states are included in the latent process, while altitude-dependent seasonal loading functions allow the magnitude and phase of the seasonal cycle to vary smoothly across altitude. This formulation also provides a framework in which the evolution of seasonal states is represented dynamically.

Finally, uncertainty assessment is handled differently. In Meiring [30], uncertainty is

based on SSANOVA Bayesian confidence intervals for the coefficient-level smooth functions, which are then transformed to the profile scale. However, uncertainty in the initial steps, including mean filtering, standardization, and functional principal component estimation, is not included in the final uncertainty assessment. In contrast, this thesis uses a bootstrap-based uncertainty assessment within the fitted state-space modeling framework, to account more directly for uncertainty across the estimated components of the model, including the latent state evolution, basis-function loading parameters, variance parameters, and the fitted seasonal and nonseasonal profile components.

3.5 Functional Mixed Models for Ozone Trends

Park et al. [31] provided a complementary functional data analysis framework for ozone profiles, again by combining functional principal component analysis with additive mixed effects models for the corresponding scores, while including additional covariates. (We note that SSANOVA models used in Meiring ([30]) may be written as additive mixed effect models, linking the two approaches for modeling the principal component scores.) They analyze monthly Umkehr ozone profiles from Boulder, USA, and Arosa, Switzerland, over the period 1978–2011. The profiles were observed over Umkehr layers, corresponding approximately to altitudes from 2 km to 45 km. Since the Umkehr retrievals are naturally ordered over altitude, these authors treat the ozone profiles as smooth functional observations rather than as unrelated measurements at separate altitude levels.

Park et al. [31] proposed a single model that includes seasonality in mean and variance as well as additive effects of additional atmospheric covariates beyond ([30]) within the coefficient models; differing also in steps taken before principal component dimension reduction, and their approach to include seasonally-varying heteroskedasticity, within a

single model including all months. We review their approach here. For each station, let $y_{i,j}$ denote the ozone value corresponding to time i (month and year combination) and layer j . Park et al. [31] first represented each monthly ozone profile using a cubic B-spline basis expansion:

$$y_{ij} = y_i(x_j) + \epsilon_{ij}, \quad (3.48)$$

where

$$y_i(x) = \sum_{k=1}^K c_{ik} \phi_k(x). \quad (3.49)$$

Here $\{\phi_k(x) : k = 1, \dots, K\}$ are cubic B-spline basis functions and c_{ik} are the corresponding basis coefficients for profile at time i . The coefficients were estimated by penalized least squares, with a roughness penalty on the second derivative of the fitted function:

$$\min_{c_{ik}} \sum_j \left\{ y_{ij} - \sum_{k=1}^K c_{ik} \phi_k(x_j) \right\}^2 + \lambda \sum_j \left\{ \sum_{k=1}^K c_{ik} D^2 \phi_k(x_j) \right\}^2. \quad (3.50)$$

The first term measures goodness-of-fit to the observed ozone profile, while the second term penalizes roughness. Thus, the initial spline representation smooths the discrete Umkehr profile observations and provides a functional representation of each monthly ozone profile.

After obtaining smoothed ozone profiles using penalized B-splines, Park et al. [31] used functional principal component analysis for dimension reduction of these smoothed curves. Using the Karhunen–Loève representation, each profile was written as

$$y_i(x) = \bar{y}(x) + \sum_{l=1}^{\infty} \xi_l(x) \theta_{il}, \quad (3.51)$$

where $\bar{y}(x)$ is the mean of the smoothed ozone profiles (already obtained using B-splines), $\xi_l(x)$ is the l^{th} functional principal component, and θ_{il} is the corresponding principal

component score for profile i . In practice, the expansion was truncated to the first d components:

$$y_i(x) \approx \bar{y}(x) + \sum_{l=1}^d \xi_l(x) \theta_{il}. \quad (3.52)$$

Park et al. [31] retained the first $d = 5$ functional principal components, which explained most of the ozone profile variability at both Boulder and Arosa. The score for the l^{th} component was computed by projecting the centered smoothed profile onto the estimated principal component,

$$\theta_{li} = \int \hat{\xi}_l(x) \{ \hat{y}_i(x) - \bar{y}(x) \} dx. \quad (3.53)$$

The time series of scores θ_{li} were then used to represent the dominant temporal evolution of the vertical ozone profile shape.

Each principal component score time series was then modeled using an additive mixed effects model. For each retained component l , the score θ_{li} was modeled as

$$\begin{aligned} \theta_{li} = & c_l + g_{l1}(\text{Month}_i) + g_{l2}(\text{Year}_i) + g_{l3}(\text{QBO1}_i) + g_{l4}(\text{QBO2}_i) \\ & + g_{l5}(\text{solar cycle}_i) + g_{l6}(\text{AO}_i) + g_{l7}(\text{ENSO}_i) + g_{l8}(\text{EP flux}_i) + \epsilon_{li}, \end{aligned} \quad (3.54)$$

where c_l is an intercept, and the functions $g_{lj}(\cdot)$ represent smooth effects of the corresponding covariates on the l^{th} principal component score. Covariates included month, year, two QBO indices, solar cycle, Arctic Oscillation (AO), El Niño–Southern Oscillation (ENSO), and Eliassen–Palm flux. The month effect was modeled using a cyclic cubic spline, while the remaining smooth effects were modeled using penalized regression cubic splines.

Each smooth function was represented by a spline expansion,

$$g_{lj}(F_j) = \sum_{k=1}^{K_j} \alpha_{ljk} \phi_{lk}(F_j), \quad (3.55)$$

where F_j denotes the j^{th} covariate, $\phi_{lk}(\cdot)$ are spline basis functions, and α_{ljk} are basis coefficients. The coefficients were estimated by minimizing a penalized least squares criterion,

$$\sum_{i=1}^n \left[\theta_{li} - c_l - \sum_{j=1}^8 g_{lj}(F_{ij}) \right]^2 + \sum_{j=1}^8 \lambda_{lj} \int \{D^2 g_{lj}(F_j)\}^2 dF_j, \quad (3.56)$$

to balance goodness-of-fit with smoothness of the covariate effects. Park et al. [31] formulated the penalized spline model within a mixed effects framework, so that the smoothing parameters could be estimated as variance components. In matrix form, the model for the l^{th} score vector can be written as

$$\boldsymbol{\theta}_l = \mathbf{B}_F \mathbf{b}_l + \mathbf{B}_R \mathbf{u}_l + \boldsymbol{\epsilon}_l, \quad (3.57)$$

where $\mathbf{B}_F \mathbf{b}_l$ contains the fixed-effect part of the spline representation and $\mathbf{B}_R \mathbf{u}_l$ contains the random-effect part. The random effects and errors are assumed to satisfy

$$\mathbf{u}_l \sim N(\mathbf{0}, \boldsymbol{\Omega}_{\lambda_l}), \quad \boldsymbol{\epsilon}_l \sim N(\mathbf{0}, \sigma_l^2 \mathbf{I}_n). \quad (3.58)$$

Under this representation, smoothing is induced through the covariance structure of the random effects, and the smoothing parameters are related to the corresponding variance components.

Park et al. [31] also extended above, including a heteroskedastic error structure in the additive mixed effect models, because diagnostic plots suggested that the residual variability of the principal component score models changed systematically by month.

This enabled the uncertainty assessment for the second-stage score models to reflect seasonal differences in residual variance. However, the uncertainty assessment is still conditional on the earlier dimension-reduction steps, including the initial smoothing of the Umkehr profiles, the estimation of the functional principal components, and the truncation to the first five components.

The fitted score-level smooth functions were then transformed back to the profile scale. In particular, the ozone trend at altitude x and time i was obtained by combining the year-effect terms across the retained functional principal components:

$$O_i(x) = \sum_{l=1}^5 \hat{\xi}_l(x) \hat{g}_{l2}(\text{Year}_i), \quad (3.59)$$

where $\hat{g}_{l2}(\text{Year}_i)$ is the fitted year-effect component from the additive mixed effects model for the l^{th} principal component score, and $\hat{\xi}_l(x)$ is the estimated l^{th} functional principal component evaluated at altitude x . Thus, the altitude-dependent trend estimate is reconstructed by weighting the fitted year effects of the principal component scores by the corresponding estimated functional principal components, giving profile-level estimates of ozone trends after adjusting for month, QBO, solar cycle, AO, ENSO, and EP flux effects.

Compared with Meiring [30], Park et al. [31] modified the multi-stage functional data analysis framework, using additive mixed effects models for the principal component scores with broader set of atmospheric covariates and different heteroskedasticity models, on a different (monthly) time scale. They included all months within a single model, with a more complex heteroskedastic error structure in the score models. Conditional on other covariates in their model, they did not find significant evidence of interactions between several covariates, although their model does not specifically investigate seasonality of the QBO associations after including seasonality (as an effect of month) in the coefficient-

level additive models. They included a four-month lag in the equatorial to mid-latitude impact of QBO. The features in Park et al. [31] provide complementary flexibility to the functional regression approach, enabling the effects of several atmospheric covariates to be studied simultaneously on the monthly time scale, with heteroskedastic covariance.

However, the main distinction of ([31]) from the goals of the present thesis remains similar to that discussed for Meiring [30], except for their study of many additional covariates than we present in this thesis. The approach of ([31]) is still based on separate stages of profile smoothing, functional principal component estimation, and score-level regression modeling. Consequently, the final uncertainty assessment does not fully propagate uncertainty from the first-stage smoothing and dimension reduction into the profile-level trend estimates. In contrast, the model developed in this thesis formulates the latent profile evolution directly within a unified state-space framework, so that the temporal dynamics, seasonal structure, loading parameters, variance parameters, and reconstructed profile components are estimated and assessed within one model-based framework. In this thesis we propose a model including seasonality and local temporal evolution within our unified functional state space model; extensions to include additional atmospheric dynamics remain as future work.

Another important difference is the choice of functional basis. Park et al. [31], like Meiring [30], used functional principal components as data-driven empirical basis functions. These basis functions are efficient for dimension reduction, but they are generally global over altitude. In this thesis, compactly supported B-spline basis functions are used to represent the vertical profile structure within the observation equation of the state space model, with their coefficients being modeled within the latent state evolution equation. This gives the latent coefficients a more local representation of altitude dependence, enabling the simultaneous temporal evolution of all spline coefficients to be modeled directly through the transition matrix in the state-space model.

3.6 Positioning of This Work

The existing approaches highlight a trade-off between capturing vertical dependence, modeling temporal dynamics, and maintaining computational efficiency. State-space models provide temporal dynamics but are often applied independently across altitude levels, without incorporating the vertical dependence that functional representations would provide. Dimension reduction methods frequently improve computational efficiency, but should be included in unified models with explicit modeling of dynamics and improved uncertainty quantification.

The model proposed in this work addresses these limitations by jointly modeling ozone profiles across altitude within a unified state-space framework. Vertical dependence is captured through B-spline basis functions with compact support, allowing localized structure to be represented flexibly, while temporal evolution is governed by a low-dimensional latent state process. In contrast to existing approaches, the proposed model simultaneously incorporates vertical coupling and stochastic temporal dynamics without increasing the state dimension unnecessarily. This is achieved by introducing altitude-dependent seasonal loading functions $\mathbf{Z}_{c,j}$ and $\mathbf{Z}_{s,j}$, defined through B-spline basis expansions in the observation matrix (see equations 2.25 and 2.26), resulting in a parsimonious yet flexible representation.

Chapter 4

Maximum Likelihood Estimation in Linear Gaussian State-Space Models

In this chapter, we describe the maximum likelihood estimation (MLE) framework for the linear Gaussian state-space model introduced in Chapter 2 (see for example [19]). The objective is to estimate the unknown model parameters based on observed data while accounting for the presence of latent state variables.

4.1 Maximum Likelihood Estimation

Recall the linear Gaussian state-space model defined in equation (2.3) - (2.4). Let $\boldsymbol{\theta}$ denote the collection of unknown parameters. In this work, the parameters to be estimated are $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{B}$ and $\mathbf{Z}$, or equivalently, the scalar parameters σ_ω^2 , σ_c^2 , σ_R^2 , α_0 , α_1 , together with the seasonal loading vectors $\{\gamma_{c,1}, \dots, \gamma_{c,p}\}$ and $\{\gamma_{s,1}, \dots, \gamma_{s,p}\}$.

For each time point $t \in \{1, \dots, T\}$, let $\mathcal{Y}_t$ denote the set of all the observations up to and including time point t , namely, $\mathcal{Y}_t = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_t\}$. The likelihood function based on the entire time trajectory is defined as

$$\begin{aligned}
\mathcal{L}(\boldsymbol{\theta}) &= p(\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T \mid \boldsymbol{\theta}) \\
&= p(\mathbf{y}_1 \mid \boldsymbol{\theta}) p(\mathbf{y}_2 \mid \mathbf{y}_1, \boldsymbol{\theta}) \cdots p(\mathbf{y}_T \mid \mathbf{y}_1, \dots, \mathbf{y}_{T-1}, \boldsymbol{\theta}) \\
&= \prod_{t=1}^T p(\mathbf{y}_t \mid \mathcal{Y}_{t-1}, \boldsymbol{\theta}).
\end{aligned} \tag{4.1}$$

where $p(\cdot)$ denotes a probability density function. Here $\mathcal{Y}_0$ denotes the empty information set, so that $p(\mathbf{y}_1 \mid \mathcal{Y}_0, \boldsymbol{\theta}) = p(\mathbf{y}_1 \mid \boldsymbol{\theta})$. This decomposition follows from the chain rule of probability [19].

The maximum likelihood estimator is obtained by solving

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}), \tag{4.2}$$

or equivalently,

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} (-\log \mathcal{L}(\boldsymbol{\theta})). \tag{4.3}$$

A key difficulty in evaluating the likelihood arises from the presence of the latent state sequence $\{\mathbf{x}_t\}$, which is not directly observed. As a result, the likelihood cannot be evaluated directly.

However, under the Gaussian assumption, the likelihood can be evaluated efficiently using the Kalman filter. In particular, the Kalman filter [19] provides a recursive decomposition of the likelihood into a product of one-step-ahead predictive densities $p(\mathbf{y}_t \mid \mathcal{Y}_{t-1}, \boldsymbol{\theta})$, which are Gaussian and can be computed sequentially.

We therefore begin by introducing the Kalman filter and related state estimators, including one-step-ahead forecasting, filtering, and smoothing. The likelihood evaluation is then presented based on these results.

4.2 Kalman Filter and State Estimation

The Kalman Filter recursions and related results presented in this section review established results for linear Gaussian state-space models; see, for example, [19, 20].

The Kalman filter provides a recursive procedure for estimating the latent state vector $\mathbf{x}_t$ based on observed data. At each time point t , the procedure produces different types of state estimates depending on the information set used. The smoothed estimates are obtained through a separate backward recursion after filtering.

In particular, for each time point $t \in \{1, \dots, T\}$, the one-step-ahead predicted state is defined as

$$\hat{\mathbf{x}}_{t|t-1} = \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_{t-1}], \quad (4.4)$$

which represents the expectation of the state at time t given observations up to and including time $t - 1$.

The filtered state estimate is defined as

$$\hat{\mathbf{x}}_{t|t} = \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_t], \quad (4.5)$$

which incorporates information from the current observation $\mathbf{y}_t$ in addition to $\mathcal{Y}_{t-1}$.

Later, we will also consider the smoothed state estimate

$$\hat{\mathbf{x}}_{t|T} = \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_T], \quad (4.6)$$

which is based on the full dataset.

In addition to the state estimates, the Kalman filter tracks the corresponding estima-

tion uncertainty through covariance matrices. These are defined as

$$\mathbf{P}_{t|t-1} = \text{Var}(\mathbf{x}_t \mid \mathcal{Y}_{t-1}), \quad (4.7)$$

$$\mathbf{P}_{t|t} = \text{Var}(\mathbf{x}_t \mid \mathcal{Y}_t), \quad (4.8)$$

$$\mathbf{P}_{t|T} = \text{Var}(\mathbf{x}_t \mid \mathcal{Y}_T). \quad (4.9)$$

The Kalman filter provides recursive expressions for the predicted and filtered state estimates. The smoothing step is then applied as a backward pass using these quantities to incorporate information from future observations.

4.2.1 One-Step-Ahead Prediction

We first consider the one-step-ahead predicted state. Starting from the definition of the one-step-ahead predicted state in equation (4.4) and substituting the state equation (2.4), we take conditional expectation given $\mathcal{Y}_{t-1}$. Using the fact that $\mathbb{E}[\mathbf{w}_t \mid \mathcal{Y}_{t-1}] = \mathbf{0}$, we obtain

$$\begin{aligned} \hat{\mathbf{x}}_{t|t-1} &= \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_{t-1}] \\ &= \mathbb{E}[\mathbf{B}\mathbf{x}_{t-1} + \mathbf{w}_t \mid \mathcal{Y}_{t-1}] \\ &= \mathbf{B}\mathbb{E}[\mathbf{x}_{t-1} \mid \mathcal{Y}_{t-1}] \\ &= \mathbf{B}\hat{\mathbf{x}}_{t-1|t-1}. \end{aligned} \quad (4.10)$$

Similarly, starting from the definition of the prediction covariance in equation (4.7)

and using properties of conditional variance, we obtain

$$\mathbf{P}_{t|t-1} = \text{Var}(\mathbf{B}\mathbf{x}_{t-1} + \mathbf{w}_t \mid \mathcal{Y}_{t-1}) = \mathbf{B}\text{Var}(\mathbf{x}_{t-1} \mid \mathcal{Y}_{t-1})\mathbf{B}^\top + \mathbf{Q} = \mathbf{B}\mathbf{P}_{t-1|t-1}\mathbf{B}^\top + \mathbf{Q}, \quad (4.11)$$

where we have used the independence of $\mathbf{w}_t$ from $\mathcal{Y}_{t-1}$ and equation (2.13).

4.2.2 Observation Prediction and Innovation

We denote the one-step-ahead predicted observation by

$$\hat{\mathbf{y}}_{t|t-1} := \mathbb{E}[\mathbf{y}_t \mid \mathcal{Y}_{t-1}]. \quad (4.12)$$

Using the predicted state $\hat{\mathbf{x}}_{t|t-1}$ in equation (4.10), we can derive the corresponding prediction for the observation. From the observation equation (2.3), taking conditional expectation given $\mathcal{Y}_{t-1}$ yields

$$\begin{aligned} \hat{\mathbf{y}}_{t|t-1} &= \mathbb{E}[\mathbf{y}_t \mid \mathcal{Y}_{t-1}] \\ &= \mathbb{E}[\mathbf{Z}\mathbf{x}_t + \mathbf{v}_t \mid \mathcal{Y}_{t-1}] \\ &= \mathbf{Z}\mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_{t-1}] \\ &= \mathbf{Z}\hat{\mathbf{x}}_{t|t-1}, \end{aligned} \quad (4.13)$$

where we have used the fact that $\mathbb{E}[\mathbf{v}_t \mid \mathcal{Y}_{t-1}] = \mathbf{0}$.

The innovation (or prediction error) is defined as the difference between the observed value and its one-step-ahead prediction:

$$\boldsymbol{\delta}_t = \mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}, \quad (4.14)$$

thereby measuring the discrepancy between the model-based prediction and the actual observation at time t .

We now derive the covariance matrix of the innovation $\boldsymbol{\delta}_t$, denoted as $\mathbf{F}_t$. Starting from the definition of the innovation $\boldsymbol{\delta}_t$ in equation (4.14) and using the observation equation (2.3), we obtain

$$\begin{aligned}
 \mathbf{F}_t &= \text{Var}(\boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1}) \\
 &= \text{Var}(\mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1} \mid \mathcal{Y}_{t-1}) \\
 &= \text{Var}(\mathbf{Z}\mathbf{x}_t + \mathbf{v}_t \mid \mathcal{Y}_{t-1}) \\
 &= \mathbf{Z}\text{Var}(\mathbf{x}_t \mid \mathcal{Y}_{t-1})\mathbf{Z}^\top + \text{Var}(\mathbf{v}_t) \\
 &= \mathbf{Z}\mathbf{P}_{t|t-1}\mathbf{Z}^\top + \mathbf{R},
 \end{aligned} \tag{4.15}$$

where we have used the fact that $\hat{\mathbf{y}}_{t|t-1}$ is deterministic given $\mathcal{Y}_{t-1}$, so subtracting it does not affect the conditional variance. We have also used the independence of $\mathbf{v}_t$ from $\mathcal{Y}_{t-1}$ and $\text{Var}(\mathbf{v}_t) = \mathbf{R}$. The matrix $\mathbf{F}_t$ represents the uncertainty associated with the prediction error at time t .

4.2.3 Filtering

The filtered state $\hat{\mathbf{x}}_{t|t}$ is an update of the one-step-ahead predicted state $\hat{\mathbf{x}}_{t|t-1}$ using the new observation $\mathbf{y}_t$. The filtered state estimate is defined in equation (4.5). The Kalman filter updates the predicted state $\hat{\mathbf{x}}_{t|t-1}$ using the innovation $\boldsymbol{\delta}_t$ derived in (4.14). Under the Gaussian assumption, the conditional expectation of $\mathbf{x}_t$ given $\boldsymbol{\delta}_t$ is linear in $\boldsymbol{\delta}_t$, leading to

$$\begin{aligned}
\hat{\mathbf{x}}_{t|t} &:= \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_t] = \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_{t-1}, \boldsymbol{\delta}_t] \\
&= \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_{t-1}] + \text{Cov}(\mathbf{x}_t, \boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1})[\text{Var}(\boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1})]^{-1} \boldsymbol{\delta}_t \\
&= \hat{\mathbf{x}}_{t|t-1} + \mathbf{P}_{t|t-1} \mathbf{Z}^\top [\text{Var}(\boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1})]^{-1} \boldsymbol{\delta}_t \\
&= \hat{\mathbf{x}}_{t|t-1} + \mathbf{P}_{t|t-1} \mathbf{Z}^\top \mathbf{F}_t^{-1} \boldsymbol{\delta}_t
\end{aligned} \tag{4.16}$$

Again using properties of conditional distributions under Gaussian assumptions, we obtain the corresponding covariance matrix

$$\begin{aligned}
\mathbf{P}_{t|t} &:= \text{Var}(\mathbf{x}_t \mid \mathcal{Y}_t) = \text{Var}(\mathbf{x}_t \mid \mathcal{Y}_{t-1}, \boldsymbol{\delta}_t) \\
&= \text{Var}(\mathbf{x}_t \mid \mathcal{Y}_{t-1}) - \text{Cov}(\mathbf{x}_t, \boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1})[\text{Var}(\boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1})]^{-1} \text{Cov}(\mathbf{x}_t, \boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1})^\top \\
&= \mathbf{P}_{t|t-1} - \mathbf{P}_{t|t-1} \mathbf{Z}^\top \mathbf{F}_t^{-1} \mathbf{Z} \mathbf{P}_{t|t-1}
\end{aligned} \tag{4.17}$$

Both equations (4.16) and (4.17) use the fact that $\text{Cov}(\mathbf{x}_t, \boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1}) = \mathbf{P}_{t|t-1} \mathbf{Z}^\top$, which can be derived as follows, using the fact that $\hat{\mathbf{y}}_{t|t-1}$ is deterministic given $\mathcal{Y}_{t-1}$ and $\text{Cov}(\mathbf{x}_t, \mathbf{v}_t) = 0$:

$$\begin{aligned}
\text{Cov}(\mathbf{x}_t, \boldsymbol{\delta}_t \mid \mathcal{Y}_{t-1}) &= \text{Cov}(\mathbf{x}_t, \mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1} \mid \mathcal{Y}_{t-1}) \\
&= \text{Cov}(\mathbf{x}_t, \mathbf{Z}\mathbf{x}_t + \mathbf{v}_t - \hat{\mathbf{y}}_{t|t-1} \mid \mathcal{Y}_{t-1}) \\
&= \text{Cov}(\mathbf{x}_t, \mathbf{Z}\mathbf{x}_t + \mathbf{v}_t \mid \mathcal{Y}_{t-1}) \\
&= \mathbf{P}_{t|t-1} \mathbf{Z}^\top
\end{aligned} \tag{4.18}$$

The update step combines the model-based prediction with information from the new observation, where the term $\mathbf{P}_{t|t-1} \mathbf{Z}^\top \mathbf{F}_t^{-1}$, in some literature defined as the "Kalman gain", determines the relative weight given to the innovation.

4.2.4 Smoothing

While the filtered state estimate $\hat{\mathbf{x}}_{t|t}$ uses observations up to and including time t , it is often desirable to estimate the latent state using the full set of observations $\mathcal{Y}_T$. This leads to the smoothed state estimate, defined in equation (4.6).

Unlike filtering, which proceeds forward in time, smoothing is obtained through a backward recursion after the filtering step has been completed.

We begin from the definition in equation (4.6):

$$\hat{\mathbf{x}}_{t|T} = \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_T] = \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_t, \mathcal{Y}_{t+1:T}], \quad (4.19)$$

since $\mathcal{Y}_T = \mathcal{Y}_t \cup \mathcal{Y}_{t+1:T}$, using notation $\mathcal{Y}_{t+1:T} = \{\mathbf{y}_{t+1}, \dots, \mathbf{y}_T\}$.

Under the Gaussian assumption, the conditional expectation of $\mathbf{x}_t$ given future observations can be expressed in terms of the next-step smoothed estimate. In particular, we have

$$\begin{aligned} \hat{\mathbf{x}}_{t|T} &= \mathbb{E}[\mathbf{x}_t \mid \mathcal{Y}_t] + \text{Cov}(\mathbf{x}_t, \mathbf{x}_{t+1} \mid \mathcal{Y}_t) [\text{Var}(\mathbf{x}_{t+1} \mid \mathcal{Y}_t)]^{-1} (\hat{\mathbf{x}}_{t+1|T} - \hat{\mathbf{x}}_{t+1|t}) \\ &= \hat{\mathbf{x}}_{t|t} + \mathbf{P}_{t|t} \mathbf{B}^\top \mathbf{P}_{t+1|t}^{-1} (\hat{\mathbf{x}}_{t+1|T} - \hat{\mathbf{x}}_{t+1|t}), \end{aligned} \quad (4.20)$$

Here, we have used the state equation (2.4) and the fact that

$$\text{Cov}(\mathbf{x}_t, \mathbf{x}_{t+1} \mid \mathcal{Y}_t) = \mathbf{P}_{t|t} \mathbf{B}^\top.$$

Now we denote

$$\mathbf{J}_t = \mathbf{P}_{t|t} \mathbf{B}^\top \mathbf{P}_{t+1|t}^{-1}, \quad (4.21)$$

thus equation (4.20) becomes

$$\hat{\mathbf{x}}_{t|T} = \hat{\mathbf{x}}_{t|t} + \mathbf{J}_t (\hat{\mathbf{x}}_{t+1|T} - \hat{\mathbf{x}}_{t+1|t}) \quad (4.22)$$

The matrix $\mathbf{J}_t$ determines how information from future states is incorporated backward to refine the estimate of $\mathbf{x}_t$.

The corresponding covariance matrix is updated as

$$\begin{aligned} \mathbf{P}_{t|T} &= \text{Var}(\mathbf{x}_t \mid \mathcal{Y}_T) \\ &= \mathbf{P}_{t|t} + \mathbf{P}_{t|t} \mathbf{B}^\top \mathbf{P}_{t+1|t}^{-1} (\mathbf{P}_{t+1|T} - \mathbf{P}_{t+1|t}) \mathbf{P}_{t+1|t}^{-1} \mathbf{B} \mathbf{P}_{t|t} \end{aligned} \quad (4.23)$$

The smoothing recursion proceeds backward in time from $t = T - 1$ to $t = 1$, using the filtered estimates $\hat{\mathbf{x}}_{t|t}$ and predicted estimates $\hat{\mathbf{x}}_{t+1|t}$ obtained from the Kalman filter.

Although smoothing is not required for the forward Kalman filtering procedure, it provides improved state estimates by incorporating information from the entire dataset. As such, it is often useful in analysis and inference.

4.3 Likelihood Evaluation via the Kalman Filter

As discussed in Section 4.1, the likelihood function can be written as a product of one-step-ahead predictive densities, as shown in equation (4.1). Under the Gaussian state-space model, each term $p(\mathbf{y}_t \mid \mathcal{Y}_{t-1}, \boldsymbol{\theta})$ follows a multivariate normal distribution. From the results of the Kalman filter, we have

$$\mathbf{y}_t \mid \mathcal{Y}_{t-1}, \boldsymbol{\theta} \sim \mathcal{N}(\hat{\mathbf{y}}_{t|t-1}, \mathbf{F}_t), \quad (4.24)$$

where $\hat{\mathbf{y}}_{t|t-1}$ is the one-step-ahead predicted observation defined in equation (4.12), and $\mathbf{F}_t$ is the innovation covariance matrix defined in (4.15).

Therefore, the likelihood function can be written as

$$\mathcal{L}(\boldsymbol{\theta}) = \prod_{t=1}^T (2\pi)^{-\frac{n_a}{2}} |\mathbf{F}_t|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \boldsymbol{\delta}_t^\top \mathbf{F}_t^{-1} \boldsymbol{\delta}_t\right), \quad (4.25)$$

where $\boldsymbol{\delta}_t = \mathbf{y}_t - \hat{\mathbf{y}}_{t|t-1}$ is the innovation, defined in equation (4.14).

Since our goal is to obtain the maximum likelihood estimator $\hat{\boldsymbol{\theta}}$ which is defined in equation (4.3), we take the negative logarithm of equation (4.25) and obtain

$$-\log \mathcal{L}(\boldsymbol{\theta}) = \frac{1}{2} \sum_{t=1}^T [\log |\mathbf{F}_t| + \boldsymbol{\delta}_t^\top \mathbf{F}_t^{-1} \boldsymbol{\delta}_t + n_a \log(2\pi)], \quad (4.26)$$

where n_a denotes the dimension of the observation vector $\mathbf{y}_t$.

4.4 Innovation-Based Bootstrap for Parameter Inference

While the likelihood framework developed in Section 4.3 provides point estimates of the model parameters, it does not directly quantify their uncertainty. To address this, we employ the innovation-based bootstrap procedure of [44]. This approach resamples standardized innovations from the fitted state-space model and reconstructs bootstrap datasets using the Kalman recursion. The method preserves the temporal dependence structure implied by the fitted model while introducing sampling variability through resampled innovations.

Suppose the model parameters $\boldsymbol{\theta}$ have already been estimated via maximum likelihood as described in Sections 4.1 - 4.3. Recall from Section 4.2 that the innovation (one-step-

ahead prediction error) at time t is defined as $\boldsymbol{\delta}_t$ in equation (4.14), with corresponding variance matrix $\mathbf{F}_t$ given in equation (4.15).

Standardized Innovations

Define the standardized innovations as

$$\mathbf{e}_t = \mathbf{F}_t^{-1/2} \boldsymbol{\delta}_t, \quad (4.27)$$

where $\mathbf{F}_t^{-1/2}$ denotes a symmetric square-root inverse of $\mathbf{F}_t$. Since $\mathbf{F}_t$ is generally a dense $n_a \times n_a$ covariance matrix, the standardization in equation (4.27) simultaneously decorrelates and rescales the entire multivariate innovation vector rather than standardizing each altitude level separately. We therefore refer to elements of $\mathbf{e}_t$ as standardized innovation components in this transformed space, instead of describing them as corresponding to altitude levels $a_1, \dots, a_{n_a}$. This distinction is important when standardized innovations are displayed or used for diagnostic plots in Chapter 6. Under correct model specification and Gaussian assumptions, the standardized innovations are approximately independent and identically distributed with mean zero and identity covariance.

Bootstrap Algorithm

For the b^{th} bootstrap replicate ($b = 1, \dots, n_{boot}$), proceed as follows:

1. Sample with replacement from $\{\mathbf{e}_1, \dots, \mathbf{e}_T\}$ with equal probability, to obtain a bootstrap sequence

$$\{\mathbf{e}_1^{*(b)}, \dots, \mathbf{e}_T^{*(b)}\}. \quad (4.28)$$

2. Construct bootstrap innovations by transforming back to the original scale:

$$\boldsymbol{\delta}_t^{*(b)} = \mathbf{F}_t^{1/2} \mathbf{e}_t^{*(b)} \quad t = 1, \dots, T. \quad (4.29)$$

By construction, $\text{Var}(\boldsymbol{\delta}_t^{*(b)}) = \mathbf{F}_t$.

3. Generate a bootstrap dataset recursively using the Kalman prediction and filter equations, holding the initial state estimate and its covariance fixed at the values obtained from the original fit:

$$\hat{\mathbf{x}}_{1|0}^{*(b)} = \hat{\mathbf{x}}_{1|0}. \quad (4.30)$$

For $t = 1, \dots, T$, compute

$$\mathbf{y}_t^{*(b)} = \mathbf{Z} \hat{\mathbf{x}}_{t|t-1}^{*(b)} + \boldsymbol{\delta}_t^{*(b)}, \quad (4.31)$$

$$\hat{\mathbf{x}}_{t|t}^{*(b)} = \hat{\mathbf{x}}_{t|t-1}^{*(b)} + \mathbf{P}_{t|t-1} \mathbf{Z}^\top \mathbf{F}_t^{-1} \boldsymbol{\delta}_t^{*(b)}, \quad (4.32)$$

$$\hat{\mathbf{x}}_{t+1|t}^{*(b)} = \mathbf{B} \hat{\mathbf{x}}_{t|t}^{*(b)}. \quad (4.33)$$

Since we initialize the one-step-ahead state prediction for $\hat{\mathbf{x}}_{1|0}^{*(b)}$, the algorithm in equations (4.30)-(4.33) matches the order of the Kalman Filter steps in Section 4.2; for $t = T$, the last step in equation (4.33) to obtain the one-step-ahead state prediction $\hat{\mathbf{x}}_{T+1|T}^{*(b)}$ is usually omitted.

For $b = 1, \dots, n_{boot}$, estimate the parameters $\boldsymbol{\theta}$ of the model defined in equations (2.3) - (2.4) from the b^{th} bootstrap dataset

$$\{\mathbf{y}_1^{*(b)}, \dots, \mathbf{y}_T^{*(b)}\}, \quad (4.34)$$

treating this bootstrap dataset as the observations, and estimating model parameters

using the maximum likelihood procedure described in Sections 4.1 - 4.3, yielding the b^{th} bootstrap estimator $\hat{\boldsymbol{\theta}}^{*(b)}$.

Bootstrap Confidence Intervals

After generating n_{boot} bootstrap replicates and estimating $\boldsymbol{\theta}$ from each bootstrap sample, the resulting bootstrap estimators

$$\{\hat{\boldsymbol{\theta}}^{*(1)}, \dots, \hat{\boldsymbol{\theta}}^{*(n_{boot})}\} \quad (4.35)$$

provide an empirical approximation to the sampling distribution of $\hat{\boldsymbol{\theta}}$.

For the j^{th} parameter component θ_j , and any chosen value of α , a $100(1 - \alpha)\%$ bootstrap confidence interval is constructed as

$$\left[\hat{\theta}_j^{*[\alpha/2]}, \hat{\theta}_j^{*[1-\alpha/2]} \right], \quad (4.36)$$

where $\hat{\theta}_j^{*[q]}$ denotes the empirical q^{th} quantile of the bootstrap estimates for this component.

This procedure provides bootstrap-based uncertainty quantification for all model parameters.

4.5 Numerical Optimization

The maximum likelihood estimator is then obtained by minimizing the negative log-likelihood function in equation (4.26). Since this objective function does not have a closed-form solution, numerical optimization methods are required.

In this work, we employ a quasi-Newton method, specifically the BFGS algorithm

[45], to carry out the optimization. This method iteratively updates parameter estimates using gradient information and provides efficient convergence in moderate-dimensional settings. At each iteration of the optimization, evaluation of the objective function requires running the Kalman filter to compute the innovations $\boldsymbol{\delta}_t$ and their covariance matrices $\mathbf{F}_t$. A different estimation approach frequently used for maximum likelihood estimation in state-space models is the expectation - maximization (EM) algorithm (see [46, 47]). In this setting, the E-step uses Kalman filtering and smoothing to compute conditional expectations of the latent states, while the M-step updates the model parameters by maximizing the expected complete-data log-likelihood. However, BFGS is much faster than EM under this framework in general (see [47, 48]). For this reason, we use the BFGS algorithm as the primary optimization method.

The optimization is performed over all model parameters, including the elements of $\mathbf{Q}$, $\mathbf{R}$, $\mathbf{B}$, and $\mathbf{Z}$. Parameter constraints, such as positivity of variance components, are enforced through appropriate parameterization.

The optimization procedure requires specification of initial parameter values. In practice, these are chosen based on simple heuristics or prior knowledge to ensure numerical stability of the optimization.

The numerical optimization is carried out using tools from the MARSS package in R [49], which are adapted to accommodate the specific structure of the proposed model. Further details of initialization setup and optimization settings are provided in Appendix B.

Chapter 5

Simulation Study

This chapter presents a comprehensive simulation study to evaluate the performance of the proposed state-space model and estimation procedure. The study is organized around three simulation configurations, denoted A, B, and C, which differ in the number of seasonal harmonic pairs and in the seasonal harmonic state noise variance. These configurations allow us to compare model performance under deterministic seasonal dynamics, stochastic seasonal dynamics, and increased seasonal complexity.

The study consists of the following components: parameter specification, data generation, model fitting, bootstrap inference, and performance assessment. We first describe the common data-generating mechanism and the configuration-specific parameter settings. We then describe the estimation procedure and innovation-based bootstrap implementation. The chapter concludes with a comparison of simulation results across the three configurations.

5.1 Parameter Specification

We begin by explicitly specifying the complete state-space model that constitutes the data-generating process for the simulation study. For each time point $t = 1, \dots, T$, the simulated data are generated from the following multivariate autoregressive state space model we introduced in Section 2.3 - 2.5. We first review the observation and state evolution equations in terms of their parameters, before describing data generation with specific simulation parameter settings.

The observation equation is

$$\mathbf{y}_t = \mathbf{Z}\mathbf{x}_t + \mathbf{v}_t = \mathbf{B}_{sp}\mathbf{x}_{t,c} + \sum_{j=1}^p \mathbf{B}_{sp}\gamma_{c,j}M_{t,j} + \sum_{j=1}^p \mathbf{B}_{sp}\gamma_{s,j}M_{t,j}^* + \mathbf{v}_t, \quad \mathbf{v}_t \stackrel{ind}{\sim} N(\mathbf{0}, \sigma_R^2 \mathbf{I}_{n_a}), \quad (5.1)$$

where the observation matrix $\mathbf{Z}$ is defined implicitly through its components as

$$\mathbf{Z} = \begin{pmatrix} \mathbf{B}_{sp} & \mathbf{B}_{sp}\gamma_{c,1} & \mathbf{B}_{sp}\gamma_{s,1} & \dots & \mathbf{B}_{sp}\gamma_{c,p} & \mathbf{B}_{sp}\gamma_{s,p} \end{pmatrix}, \quad (5.2)$$

and where:

- $\mathbf{y}_t \in \mathbb{R}^{n_a}$ denotes the observed ozone profile across n_a altitude levels at time t ,
- $\mathbf{B}_{sp} \in \mathbb{R}^{n_a \times k}$ is the cubic B-spline basis matrix with k degrees of freedom evaluated at the altitude grid where n_a is the number of altitudes,
- $\mathbf{x}_{t,c} \in \mathbb{R}^k$ is the time-varying B-spline coefficient vector for the non-harmonic functional component; $\mathbf{x}_{t,c}$ is the first component of the latent state vector,
- $\gamma_{c,j}, \gamma_{s,j} \in \mathbb{R}^k$ are the static B-spline coefficient vectors that determine the altitude-dependent loading functions for the j^{th} harmonic pair, for $j = 1, \dots, p$,

- $M_{t,j}$ and $M_{t,j}^*$ are the harmonic state variables for the j^{th} harmonic pair, for $j = 1, \dots, p$; evolving as the second component of the latent state vector,
- $\mathbf{v}_t$ is the observation noise vector at time point t , with common variance parameter σ_R^2 across altitude levels, assumed iid across time.

The state evolution equation is

$$\mathbf{x}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{w}_t, \quad \mathbf{w}_t \stackrel{ind}{\sim} N(\mathbf{0}, \mathbf{Q}), \quad (5.3)$$

where the full state vector is defined as

$$\mathbf{x}_t = \left(\mathbf{x}_{t,c}^\top \quad M_{t,1} \quad M_{t,1}^* \quad \dots \quad M_{t,p} \quad M_{t,p}^* \right)^\top \in \mathbb{R}^{k+2p}, \quad (5.4)$$

evolving according to a block-diagonal transition matrix $\mathbf{B}$:

$$\mathbf{B} = \begin{pmatrix} \mathbf{B}_c & \mathbf{0}_{k \times 2} & \dots & \mathbf{0}_{k \times 2} \\ \mathbf{0}_{2 \times k} & d\mathbf{R}(\omega_1) & \dots & \mathbf{0}_{2 \times 2} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}_{2 \times k} & \mathbf{0}_{2 \times 2} & \dots & d\mathbf{R}(\omega_p) \end{pmatrix} \quad (5.5)$$

Within $\mathbf{B}$, the component $\mathbf{B}_c \in \mathbb{R}^{k \times k}$ is a symmetric tridiagonal Toeplitz matrix that controls the evolution of the non-harmonic B-spline coefficients $\mathbf{x}_{t,c}^\top$, specifically

$$\mathbf{B}_c = \begin{pmatrix} \alpha_0 & \alpha_1 & 0 & \cdots & 0 \\ \alpha_1 & \alpha_0 & \alpha_1 & \ddots & \vdots \\ 0 & \alpha_1 & \alpha_0 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & \alpha_1 \\ 0 & \cdots & 0 & \alpha_1 & \alpha_0 \end{pmatrix} \quad (5.6)$$

The evolution of the harmonic block is defined through rotation matrices

$$\mathbf{R}(\omega_j) = \begin{pmatrix} \cos(\omega_j) & \sin(\omega_j) \\ -\sin(\omega_j) & \cos(\omega_j) \end{pmatrix}, \quad \omega_j = \frac{2\pi j}{T_p}, \quad j = 1, \dots, p, \quad (5.7)$$

where T_p denotes the nominal seasonal period and d is a damping factor.

The state noise covariance matrix is given by

$$\mathbf{Q} = \begin{pmatrix} \sigma_c^2 \mathbf{I}_k & \mathbf{0}_{k \times 2p} \\ \mathbf{0}_{2p \times k} & \sigma_\omega^2 \mathbf{I}_{2p} \end{pmatrix}, \quad (5.8)$$

where σ_c^2 controls the variability of the spline coefficient states and σ_ω^2 governs stochastic variation in the seasonal harmonic states. A special case is $\sigma_\omega^2 = 0$, where the seasonal cycle does not change stochastically.

For data generation: The harmonic component is initialized deterministically as

$$(M_{0,1}, M_{0,1}^*, \dots, M_{0,p}, M_{0,p}^*) = (1, 0, \dots, 1, 0), \quad (5.9)$$

and the true initial spline coefficient vector $\mathbf{x}_{0,c}$ is constructed from a smooth ozone profile, as described in Section 5.2. Later in model fitting, the spline component of the initial state is treated as unknown and estimated from the simulated data, while the

harmonic initial states are fixed as above due to the identifiability issues discussed in Section 2.6.

Together, equations (5.1) – (5.8) fully specify the data-generating mechanism used throughout the simulation study. The model used for estimation in Section 5.3 coincides with the data-generating model, so that the simulation study evaluates finite-sample performance under correct model specification.

We now specify the simulation settings for key model parameters, including σ_c^2 , σ_ω^2 , σ_R^2 , α_0 , and α_1 , and the number of replications in our simulation. Table 5.1 summarizes the parameters that are held fixed or varied across the simulation configurations. The three configurations used in the simulation study are then listed in Table 5.2. Specification of the higher-dimensional vector parameters $\gamma_{c,j}$ and $\gamma_{s,j}$, $j \in \{1, \dots, p\}$ is described in Section 5.2.

Notation	Definitions	Simulation Settings
k	B-spline degrees of freedom	7
p	Number of seasonal harmonic pairs	1 and 2
T	Total number of evenly spaced time points	730
n_a	Number of altitudes	30
σ_R^2	Observation noise variance parameter in equation (5.1)	0.06
d	Damping factor in equation (5.5)	1
α_0, α_1	Diagonal and off-diagonal parameters in the spline-state transition block $\mathbf{B}_c$ in equation (5.6)	$\alpha_0 = 0.25, \alpha_1 = 0.1$
σ_c^2	Spline coefficient state noise variance in equation (5.8)	0.01
σ_ω^2	Seasonal harmonic state noise variance in equation (5.8)	0 and 0.1
n_{rep}	Number of replicates for each parameter configuration	100

Table 5.1: Simulation settings: parameters and dimensions in the simulation study. Where 2 numbers are listed for a single parameter under the Simulation Settings column, the simulations are performed for both values of these parameters.

The total number of time points is fixed at $T = 730$. Under 5-day pentad aggregation, one calendar year consists of 73 pentads. For simplicity, the simulation does not account for leap years, and each year is assumed to contain exactly 73 pentads. Thus, $T = 730$ corresponds to 10 years of observations. This length was chosen to provide sufficient temporal information for estimating both state dynamics and seasonal structure while keeping the repeated simulation, fitting, and bootstrap procedures computationally feasible.

The simulation configurations are constructed by varying only a small number of parameters while holding the remaining settings fixed. In particular, the configurations differ in the number of seasonal harmonic pairs p and the seasonal harmonic state noise variance σ_ω^2 . This design allows us to compare deterministic seasonal dynamics, stochastic seasonal dynamics, and increased seasonal complexity in a controlled setting.

For each parameter configuration selected, we generate n_{rep} replicates. The number of replicates n_{rep} is the same across all parameter configurations.

To ensure reproducibility while maintaining independence across replicates within each configuration, we first set a configuration-specific seed value. Using this seed, we generate n_{rep} distinct seed values, each of which is assigned to a specific replicate within this given parameter configuration. Thus, every replicate within each configuration is generated using a distinct, recorded seed value.

5.2 Data Generation

Within the simulation settings in Table 5.1, the configurations differ only in the number of seasonal harmonic pairs p and the seasonal harmonic state noise variance σ_ω^2 . For each selected parameter configuration, n_{rep} replicates are generated. The transition matrix $\mathbf{B}$ is constructed according to equation (5.5), consisting of a tridiagonal Toeplitz

spline-state block and p harmonic rotation blocks with frequencies $\omega_j = j\omega_1$ for $j = 1, \dots, p$.

For each harmonic pair $j = 1, \dots, p$, the seasonal loading coefficient vectors $\gamma_{c,j}$ and $\gamma_{s,j}$ in equation (5.2) are generated independently from the uniform distribution

$$U[\gamma_{\min}, \gamma_{\max}], \quad (5.10)$$

with $\gamma_{\min} = -0.8$ and $\gamma_{\max} = 0.8$ in the simulation study. These loading coefficient vectors are generated once for each parameter configuration and are held fixed across all the n_{rep} replicates within that configuration. The observation matrix $\mathbf{Z}$ is then constructed as in equation (5.2).

The initial state vector $\mathbf{x}_0$ is constructed as follows. To mimic a realistic ozone profile across altitude, we first select the middle altitude $a_{\lfloor n_a/2 \rfloor}$. A baseline value at this altitude is drawn from $N(\mu_0, \sigma_0^2)$, with $\mu_0 = 0.5$ and $\sigma_0 = 0.1$ in the simulation study. Altitude levels above and below the center are then generated by subtracting small independent random increments so that the profile decreases smoothly away from the midpoint. The increments are drawn from

$$U\left[0, \frac{\Delta_{\max}}{\lfloor n_a/2 \rfloor - 1}\right], \quad (5.11)$$

with $\Delta_{\max} = 0.2$ in the simulation study.

The resulting simulated ozone profile is projected onto the B-spline basis, and the estimated spline coefficients form the first k components of $\mathbf{x}_0$. The harmonic components are initialized deterministically as

$$(M_{0,1}, M_{0,1}^*, \dots, M_{0,p}, M_{0,p}^*) = (1, 0, \dots, 1, 0). \quad (5.12)$$

The true initial state vector $\mathbf{x}_0$ is fixed across all the n_{rep} replicates within each parameter

configuration.

For each replicate, the state and observation processes are simulated recursively according to equations (5.3) and (5.1) for $t = 1, \dots, T$. At each time step, the state vector $\mathbf{x}_t$ is updated via the transition equation, and the corresponding observation $\mathbf{y}_t$ is generated from the observation equation.

The simulated state trajectories $\mathbf{x}_t$, noise-free observations $\tilde{\mathbf{y}}_t$, and noisy observations $\mathbf{y}_t$ are stored for subsequent model fitting and performance evaluation.

5.3 Model Fitting

Each simulated dataset is fitted using the state-space model described in Chapter 4. For each parameter configuration, the same initialization scheme is used for all replicates. This ensures that differences in estimation performance are not attributable to differences in the initialization procedure across replicates.

The transition matrix $\mathbf{B}$ contains two free parameters, α_0 and α_1 , governing the spline coefficient block, while the harmonic rotation blocks are fixed and not estimated. The initial values for the free transition parameters are set as

$$\alpha_0^{(0)} \sim \mathcal{N}(0.25, 0.05^2), \quad (5.13)$$

$$\alpha_1^{(0)} \sim \mathcal{N}(0.1, 0.05^2). \quad (5.14)$$

The state covariance matrix $\mathbf{Q}$ is specified as diagonal. The spline coefficient state noise variance σ_c^2 is treated as a free parameter. The seasonal state noise variance σ_ω^2 is fixed at the value specified by the corresponding simulation configuration. The free

spline coefficient state noise variance is initialized as

$$\sigma_c^{2(0)} \sim \mathcal{N}(0.05, 0.01^2). \quad (5.15)$$

The observation matrix $\mathbf{Z}$ consists of a fixed B-spline basis component and free seasonal loading coefficient vectors $\{\boldsymbol{\gamma}_{c,j}, \boldsymbol{\gamma}_{s,j} : j = 1, \dots, p\}$. These coefficients are initialized at zero:

$$\boldsymbol{\gamma}_{c,j}^{(0)} = \mathbf{0}_{k \times 1} \quad j = 1, \dots, p, \quad (5.16)$$

$$\boldsymbol{\gamma}_{s,j}^{(0)} = \mathbf{0}_{k \times 1} \quad j = 1, \dots, p. \quad (5.17)$$

The observation covariance matrix is specified as

$$\mathbf{R} = \sigma_R^2 \mathbf{I}_{n_a}, \quad (5.18)$$

and the common variance parameter is initialized at

$$\sigma_R^{2(0)} = 0.05. \quad (5.19)$$

Collecting the free parameters, the parameter vector estimated in each simulated dataset is

$$\boldsymbol{\theta} = (\alpha_0, \alpha_1, \sigma_c^2, \sigma_R^2, \boldsymbol{\gamma}_{c,1}^\top, \boldsymbol{\gamma}_{s,1}^\top, \dots, \boldsymbol{\gamma}_{c,p}^\top, \boldsymbol{\gamma}_{s,p}^\top, \mathbf{x}_{0,c}^\top)^\top. \quad (5.20)$$

The spline component of the initial state, $\mathbf{x}_{0,c}$, is included in $\boldsymbol{\theta}$ and estimated from the simulated data. The harmonic initial conditions $(M_{0,1}, M_{0,1}^*, \dots, M_{0,p}, M_{0,p}^*)$ are fixed at $(1, 0, \dots, 1, 0)$, and are not estimated. Given these starting values, maximum likelihood estimation proceeds by minimizing the negative log-likelihood function $f(\boldsymbol{\theta}) = -\log \mathcal{L}(\boldsymbol{\theta})$

as described in Chapter 4.

The resulting parameter estimates from all replicates are retained for performance evaluation and for subsequent innovation-based bootstrap inference.

5.4 Simulation Bootstrap Implementation

In this section, we describe how the innovation-based bootstrap procedure introduced in Section 4.4 is implemented within our simulation study.

For each replicate generated under a given parameter configuration, the state-space model is first fitted via maximum likelihood as described in Section 5.3, yielding parameter estimates $\hat{\boldsymbol{\theta}}$. Using the fitted model, we compute the innovations $\{\boldsymbol{\delta}_1, \dots, \boldsymbol{\delta}_T\}$ defined in equation (4.14) and the corresponding innovation covariance matrices $\{\mathbf{F}_1, \dots, \mathbf{F}_T\}$ defined in equation (4.15).

Bootstrap Replication

For each replicate dataset, we generate n_{boot} bootstrap datasets using the innovation-based bootstrap algorithm described in Section 4.4. Specifically, for each bootstrap iteration $b = 1, \dots, n_{\text{boot}}$:

1. The standardized innovations are resampled with replacement.
2. A bootstrap dataset $\{\mathbf{y}_1^{*(b)}, \dots, \mathbf{y}_T^{*(b)}\}$ is reconstructed via the Kalman recursion.
3. The state-space model is re-fitted to the bootstrap dataset, producing a bootstrap estimator $\hat{\boldsymbol{\theta}}^{*(b)}$.

The same model specification and configuration-specific fixed quantities are used in

the bootstrap refits. This yields a collection of bootstrap estimates

$$\{\hat{\boldsymbol{\theta}}^{*(1)}, \dots, \hat{\boldsymbol{\theta}}^{*(n_{\text{boot}})}\} \quad (5.21)$$

for each original replicate.

Bootstrap Confidence Intervals and Coverage

For each estimated parameter component θ_j , a nominal $(1 - \alpha)$ percentile bootstrap confidence interval is constructed using the empirical $\alpha/2$ and $1 - \alpha/2$ quantiles of the bootstrap estimates.

To evaluate finite-sample performance, we compute empirical coverage probabilities across simulation replicates. Specifically, for each parameter θ_j , the coverage rate is defined as the proportion of simulation replicates for which the corresponding bootstrap confidence interval contains the true parameter value used in data generation.

All reported bootstrap confidence intervals and parameter coverage results in Section 5.5 are based on this procedure.

5.5 Simulation Results

We now summarize the simulation results across the three configurations introduced in Sections 5.1 - 5.2. The data are generated according to the state-space model described above, with the common parameter settings given in Table 5.1.

The three configurations are designed to compare the effect of seasonal state noise variation and seasonal model complexity. Configuration A uses one deterministic seasonal harmonic pair. Configuration B keeps one seasonal harmonic pair but introduces stochastic variation in the seasonal harmonic states. Configuration C keeps stochastic

seasonal harmonic states and increases the number of seasonal harmonic pairs from one to two. Thus, the comparison from A to B isolates the effect of allowing stochastic seasonal state variation, while the comparison from B to C examines the effect of increased seasonal complexity.

The configurations considered in this study are summarized in Table 5.2.

Configuration	σ_ω^2	k	p
A	0	7	1
B	0.1	7	1
C	0.1	7	2

Table 5.2: Simulation configurations used in the simulation study. All other parameters are fixed as specified in Table 5.1.

Configuration A represents the baseline setting with deterministic seasonal harmonic dynamics, $p = 1$, and $\sigma_\omega^2 = 0$. Configuration B uses the same number of seasonal harmonic pairs but sets $\sigma_\omega^2 = 0.1$, allowing the seasonal harmonic states to evolve stochastically around the rotation dynamics. Configuration C also uses $\sigma_\omega^2 = 0.1$, but increases the number of seasonal harmonic pairs to $p = 2$. This configuration is used to assess the effect of adding a second seasonal harmonic pair under stochastic seasonal dynamics.

For each configuration, $n_{\text{rep}} = 100$ simulated datasets were generated and fitted. For each fitted dataset, $n_{\text{boot}} = 100$ innovation-bootstrap datasets were generated and refitted using the same model specification. All original model fits and all bootstrap refits converged successfully. Therefore, the parameter estimation summaries and bootstrap coverage results reported below are based on the full set of simulation replicates and bootstrap refits.

5.5.1 Example Simulated Profiles

Before evaluating estimation and bootstrap performance, we first display representative simulated profiles from the three configurations. Figure 5.1 shows the noise-free

profile $\tilde{\mathbf{y}}_t$ and the corresponding observed profile $\mathbf{y}_t$ at time $t = 365$ for Configurations A, B, and C.

The profiles in Figure 5.1 should be interpreted as simulated profile-valued observations generated from the proposed state-space model. They are not intended to reproduce the exact physical shape or scale of observed ozone profiles. The simulated values are interpreted on the model scale; therefore, the data-generating process does not impose positivity constraints on the profile values. Instead, the purpose of the simulation study is to provide a controlled setting in which the true states, loading functions, and model parameters are known. This allows us to evaluate parameter estimation, state estimation, and bootstrap uncertainty quantification under the specified data-generating mechanism.

Across the three panels, the same general state-space construction produces different model-generated profile shapes under different seasonal specifications. Configuration A uses deterministic seasonal dynamics with one harmonic pair. Configurations B and C allow stochastic seasonal state variation, and Configuration C further includes a second harmonic pair. Thus, the example profiles provide a visual summary of the types of simulated observations generated under the three configurations, while the quantitative evaluation is provided in the following subsections.

While Figure 5.1 illustrates the cross-sectional profile shape across altitude at a fixed time point, it is also helpful to examine how the simulated observations vary over time. Figure 5.2 therefore displays illustrative observation series at one selected altitude for each configuration. For visual clarity, each series is split into two time windows, $t = 1, \dots, 365$ and $t = 366, \dots, 730$.

For each configuration, the selected replicate-altitude pair is chosen to have moderate temporal variation in the noise-free signal and a visible one-step-ahead prediction discrepancy. This selection is used only for visualization and is not part of the formal estimation or coverage evaluation.




Figure 5.1: Representative simulated profiles under Configurations A, B, and C at time $t = 365$. The solid curve shows the noise-free profile $\tilde{\mathbf{y}}_t$, and the points show the observed profile $\mathbf{y}_t$. The profiles are model-generated profile-valued observations from the proposed state-space model and are used to illustrate the data-generating process under the three configurations.




Figure 5.2: Illustrative observation series for one selected replicate and one selected altitude under Configurations A, B, and C. Each configuration is split into two time windows to make the temporal variation easier to see. The black points show the observed values $y_t(a_i)$, the blue solid curve shows the noise-free values $\tilde{y}_t(a_i)$, and the red dashed curve shows the one-step-ahead predictions $\hat{y}_{t|t-1}(a_i)$ from the fitted model.

Across the three configurations, the observed series fluctuate around the corresponding noise-free trajectories, while the one-step-ahead predictions track the dominant temporal pattern. Configuration A shows more visibly noisy observations relative to the underlying signal. Configurations B and C exhibit clearer seasonal structure, with Configuration C showing repeated peaks over time, although the peak amplitudes vary across cycles, reflecting the stochastic seasonal dynamics and the inclusion of two seasonal harmonic pairs. Overall, the figure provides a representative visual check that the fitted model captures the main temporal structure in the simulated observations.

5.5.2 State Estimation Performance

We next evaluate the uncertainty of the latent state estimates. We focus on the spline coefficient states because they represent the time-varying altitude-profile component of the model. The harmonic states are shared seasonal states whose initial values are fixed for identifiability, and the uncertainty in the seasonal structure is examined later through the estimated loading functions and bootstrap inference.

Filtered and smoothed state estimators $\hat{\mathbf{x}}_{t|t}$ and $\hat{\mathbf{x}}_{t|T}$, together with their covariance matrices $\mathbf{P}_{t|t}$ and $\mathbf{P}_{t|T}$, are defined in Section 4.2. Under the Gaussian state-space model, $\mathbf{x}_t \mid \mathcal{Y}_T$ is multivariate normal (see [20]). Accordingly, for each simulation replicate $r = 1, \dots, n_{\text{rep}}$, a nominal $100(1 - \alpha)\%$ pointwise confidence interval for the j^{th} state component at time t is constructed using the corresponding filtered or smoothed state estimate and covariance matrix.

For the filtered estimator,

$$\text{CI}_{t,f}^{(j,r)} = \left[\hat{x}_{t|t}^{(j,r)} - z_{1-\alpha/2} \sqrt{P_{t|t}^{(j,j,r)}}, \hat{x}_{t|t}^{(j,r)} + z_{1-\alpha/2} \sqrt{P_{t|t}^{(j,j,r)}} \right]. \quad (5.22)$$

For the smoothed estimator,

$$\text{CI}_{t,s}^{(j,r)} = \left[\hat{x}_{t|T}^{(j,r)} - z_{1-\alpha/2} \sqrt{P_{t|T}^{(j,j,r)}}, \hat{x}_{t|T}^{(j,r)} + z_{1-\alpha/2} \sqrt{P_{t|T}^{(j,j,r)}} \right]. \quad (5.23)$$

In the reported results, we use $\alpha = 0.05$, so that $z_{1-\alpha/2} = 1.96$.

Coverage rates are computed separately for the filtered estimator, denoted by $m = \text{f}$, and the smoothed estimator, denoted by $m = \text{s}$. Let $x_t^{(j,r)}$ denote the true value of the j^{th} spline coefficient latent state at time point t in replicate r . Define the coverage indicator variable

$$I_{t,m}^{(j,r)} = \mathbf{1} \left\{ x_t^{(j,r)} \in \text{CI}_{t,m}^{(j,r)} \right\} \quad m \in \{\text{f}, \text{s}\}. \quad (5.24)$$

The empirical coverage rate for replicate r and state component j is then computed as

$$\widehat{\text{Cover}}_m^{(j,r)} = \frac{1}{T} \sum_{t=1}^T I_{t,m}^{(j,r)} \quad m \in \{\text{f}, \text{s}\}. \quad (5.25)$$

This yields one coverage rate value per replicate for each state component and each estimation method, either filtering or smoothing.

To summarize performance for each latent state across replicates, we further define the pooled mean coverage for state component j as

$$\widehat{\text{Cover}}_m^{(j)} = \frac{1}{n_{\text{rep}}} \sum_{r=1}^{n_{\text{rep}}} \widehat{\text{Cover}}_m^{(j,r)} \quad m \in \{\text{f}, \text{s}\}. \quad (5.26)$$

This pooled mean coverage is used to summarize the calibration of each spline state component.

For comparing the three configurations in a compact way, we further average the

replicate-level state coverage over the k spline coefficient states:

$$\widehat{\text{Cover}}_m^{(r)} = \frac{1}{k} \sum_{j=1}^k \widehat{\text{Cover}}_m^{(j,r)}, \quad m \in \{f, s\}. \quad (5.27)$$

This produces one coverage summary for each replicate, method, and configuration.

Figure 5.3 displays the replicate-level averaged coverage values defined in equation (5.27), where coverage is first evaluated across time for each spline state component and then averaged across the $k = 7$ spline states. Across the three configurations, the state coverage is generally close to the nominal 95% level. For the filtered estimator, the mean replicate-level averaged coverage is 0.9477, 0.9473, and 0.9425 for Configurations A, B, and C, respectively. For the smoothed estimator, the corresponding mean coverage values are 0.9474, 0.9474, and 0.9421. Thus, Configurations A and B are centered very close to the nominal level, while Configuration C shows slightly lower average coverage and more low-coverage outliers. Nevertheless, the mean coverage in Configuration C remains above 0.94 for both filtered and smoothed estimators. Thus, even in the more complex setting with two harmonic pairs, the state intervals remain close to the nominal 95% level.

Figure 5.4 gives a more detailed state-by-state summary based on the replicate-level coverage values in equation (5.25). In Configurations A and B, the state-specific pooled mean coverage values are close to the nominal level for all seven spline states. For Configuration A, the state-specific pooled means range from 0.9457 to 0.9488 for the filtered estimator and from 0.9453 to 0.9486 for the smoothed estimator. For Configuration B, the corresponding ranges are 0.9458 to 0.9491 and 0.9458 to 0.9486. Configuration C has lower state-specific pooled means, ranging from 0.9408 to 0.9448 for the filtered estimator and from 0.9398 to 0.9441 for the smoothed estimator, indicating mild undercoverage but still strong overall calibration across all spline state components.

Overall, the filtered and smoothed intervals exhibit very similar coverage behavior.




Figure 5.3: Empirical coverage rates of nominal 95% pointwise confidence intervals for the spline coefficient states across Configurations A, B, and C. For each replicate, coverage is first averaged over time for each spline state and then averaged across the $k = 7$ spline states. Boxplots summarize the resulting replicate-level coverage values across $n_{\text{rep}} = 100$ replicates. The dashed horizontal line indicates the nominal 95% coverage level.




Figure 5.4: State-specific empirical coverage rates of nominal 95% pointwise confidence intervals for the spline coefficient states. Each panel corresponds to one simulation configuration. For each spline state component, the boxplots summarize the replicate-level coverage values $\widehat{\text{Cover}}_m^{(j,r)}$ across $n_{\text{rep}} = 100$ replicates, separately for the filtered and smoothed estimators. The dashed horizontal line indicates the nominal 95% coverage level.

The mild undercoverage relative to the nominal level is not surprising. The confidence intervals rely on plug-in estimates of model parameters. The intervals in equations (5.22) and (5.23) are constructed from the Kalman filter and smoother covariance matrices computed at the fitted parameter values. Therefore, they describe conditional state uncertainty under the fitted model, but they do not fully incorporate uncertainty from estimating the model parameters (see [49]). In finite samples, this plug-in treatment can lead to slight deviations from nominal coverage. The slightly lower coverage in Configuration C is also consistent with its increased seasonal complexity, since this configuration includes two seasonal harmonic pairs and a larger number of seasonal loading parameters. Nevertheless, the coverage remains broadly close to the nominal level across the three configurations, suggesting that the filtered and smoothed state intervals provide reasonably well-calibrated pointwise uncertainty quantification for the spline coefficient states.

5.5.3 Parameter Estimation Performance

We next evaluate the finite-sample performance of the maximum likelihood estimators for the scalar model parameters. We focus on α_0 , α_1 , σ_c^2 , and σ_R^2 . The seasonal harmonic state noise variance σ_ω^2 is not included in this comparison because it is fixed at the configuration-specific value during estimation to avoid identifiability and scaling issues, as discussed in Section 2.6, rather than as a restriction on the general model formulation.

For each configuration and each simulation replicate $r = 1, \dots, n_{\text{rep}}$, let $\hat{\theta}^{(r)}$ denote the maximum likelihood estimate of a scalar parameter θ_0 . Estimation accuracy is summarized using the estimation error

$$\hat{\theta}^{(r)} - \theta_0. \quad (5.28)$$

Across replicates, we also compute the mean estimate, empirical bias, empirical stan-

dard deviation, and root mean squared error. The empirical bias and root mean squared error are defined as

$$\text{Bias}(\hat{\theta}) = \frac{1}{n_{\text{rep}}} \sum_{r=1}^{n_{\text{rep}}} \left(\hat{\theta}^{(r)} - \theta_0 \right), \quad (5.29)$$

and

$$\text{RMSE}(\hat{\theta}) = \left\{ \frac{1}{n_{\text{rep}}} \sum_{r=1}^{n_{\text{rep}}} \left(\hat{\theta}^{(r)} - \theta_0 \right)^2 \right\}^{1/2}. \quad (5.30)$$

Figure 5.5 shows the estimation errors for the four scalar parameters across Configurations A, B, and C. The variance parameters are estimated accurately across all three configurations. The observation noise variance σ_R^2 is especially stable, with estimation errors tightly centered near zero. The spline state noise variance σ_c^2 also shows small errors, although a mild downward bias is visible, especially in Configuration C.

The transition parameters α_0 and α_1 exhibit greater variability than the variance parameters. This variability is largest in Configuration C, which is consistent with the increased complexity introduced by the second seasonal harmonic pair. Nevertheless, the estimates remain broadly centered near the true values, indicating that the scalar parameters are recovered reasonably well across all three configurations.

Table 5.3 provides numerical summaries of the scalar parameter estimates. The observation noise variance σ_R^2 is recovered very accurately in all three configurations. The mean estimates are essentially equal to the true value 0.06, with RMSE values below 0.001. The spline state noise variance σ_c^2 is also estimated well. The mean estimates are close to the true value 0.01, with small negative biases. The downward bias is slightly larger in Configuration C, but remains small in absolute magnitude.

For the transition parameters, the mean estimates of α_0 and α_1 are reasonably close to their true values, but their empirical standard deviations and RMSE values are larger than those of the variance parameters. This indicates that the transition parameters are more difficult to estimate in finite samples. Configuration C has the largest RMSE values




Figure 5.5: Estimation errors $\hat{\theta} - \theta_0$ for the scalar parameters α_0 , α_1 , σ_c^2 , and σ_R^2 across Configurations A, B, and C. The dashed horizontal line indicates zero estimation error.

for both α_0 and α_1 , reflecting the additional difficulty introduced by the second seasonal harmonic pair. Even so, the estimates remain within a reasonable range of the true values, and the overall scalar-parameter estimation performance remains satisfactory.

Configuration	Parameter	True	Mean	Bias	SD	RMSE
A	α_0	0.25	0.2270	-0.0230	0.0829	0.0857
B	α_0	0.25	0.2603	0.0103	0.1138	0.1137
C	α_0	0.25	0.2640	0.0140	0.1334	0.1334
A	α_1	0.10	0.1105	0.0105	0.0532	0.0539
B	α_1	0.10	0.0933	-0.0067	0.0668	0.0669
C	α_1	0.10	0.0822	-0.0178	0.0991	0.1002
A	σ_c^2	0.01	0.0098	-0.0002	0.0009	0.0009
B	σ_c^2	0.01	0.0097	-0.0003	0.0012	0.0013
C	σ_c^2	0.01	0.0093	-0.0007	0.0014	0.0016
A	σ_R^2	0.06	0.0600	0.0000	0.0007	0.0007
B	σ_R^2	0.06	0.0599	-0.0001	0.0006	0.0006
C	σ_R^2	0.06	0.0600	0.0000	0.0006	0.0006

Table 5.3: Summary of scalar parameter estimation performance across $n_{\text{rep}} = 100$ simulation replicates. Bias is computed as the mean estimate minus the true value, SD denotes the empirical standard deviation of the estimates across replicates, and RMSE denotes root mean squared error.

5.5.4 Recovery of Seasonal Loading Functions

We next evaluate recovery of the altitude-dependent seasonal loading functions. In the observation equation, the seasonal contribution is represented through the loading vectors

$$\mathbf{Z}_{c,j} = \mathbf{B}_{sp}\boldsymbol{\gamma}_{c,j}, \quad \mathbf{Z}_{s,j} = \mathbf{B}_{sp}\boldsymbol{\gamma}_{s,j}, \quad j = 1, \dots, p. \quad (5.31)$$

These vectors give the altitude-domain loadings for the seasonal components of the j^{th} harmonic pair. We examine these transformed loading functions rather than the raw spline coefficient vectors $\boldsymbol{\gamma}_{c,j}$ and $\boldsymbol{\gamma}_{s,j}$ because $\mathbf{Z}_{c,j}$ and $\mathbf{Z}_{s,j}$ are directly interpretable as

functions over altitude.

For each configuration, the true loading functions are fixed across simulation replicates. After fitting the model to each simulated dataset, we compute the estimated altitude-domain loading functions for the r^{th} replicate within each configuration:

$$\widehat{\mathbf{Z}}_{c,j}^{(r)} = \mathbf{B}_{sp} \widehat{\boldsymbol{\gamma}}_{c,j}^{(r)}, \quad \widehat{\mathbf{Z}}_{s,j}^{(r)} = \mathbf{B}_{sp} \widehat{\boldsymbol{\gamma}}_{s,j}^{(r)}, \quad r = 1, \dots, n_{\text{rep}}. \quad (5.32)$$

Figures 5.6 - 5.8 compare the true loading functions with the mean estimated loading functions across simulation replicates.

It is important to interpret the gray bands carefully. The gray bands show the central 95% empirical range of the estimated loading functions across the $n_{\text{rep}} = 100$ simulation replicates. They are not confidence bands for a single fitted dataset. Instead, they summarize replicate-to-replicate variability in the estimated loading functions across independently generated datasets from the same parameter configuration. Bootstrap-based uncertainty for a single fitted dataset is considered separately in Section 5.5.5.

Configuration A shows a very accurate recovery of the seasonal loading functions. The mean estimated curves are nearly indistinguishable from the true curves, and the empirical ranges across replicates are narrow. This is consistent with Configuration A having deterministic seasonal state dynamics, with $\sigma_{\omega}^2 = 0$, so the seasonal signal is more regular and the loading functions are easier to identify.

In Configuration B, the recovery becomes less stable. The mean estimate of $\mathbf{Z}_{c,1}$ still captures the main shape of the true loading function, but the empirical range across replicates is much wider than in Configuration A. The recovery of $\mathbf{Z}_{s,1}$ is also less stable, and the mean estimated curve deviates more noticeably from the true curve. This indicates that allowing stochastic variation in the seasonal states makes the altitude-dependent loading functions more difficult to estimate precisely.




Figure 5.6: Recovery of the altitude-dependent seasonal loading functions under Configuration A. The black solid curve shows the true loading function, the red dashed curve shows the mean estimated loading function across $n_{\text{rep}} = 100$ simulation replicates, and the gray band shows the central 95% empirical range of the estimated loading functions across replicates.




Figure 5.7: Recovery of the altitude-dependent seasonal loading functions under Configuration B. The black solid curve shows the true loading function, the red dashed curve shows the mean estimated loading function across $n_{\text{rep}} = 100$ simulation replicates, and the gray band shows the central 95% empirical range of the estimated loading functions across replicates.




Figure 5.8: Recovery of the altitude-dependent seasonal loading functions under Configuration C. The black solid curve shows the true loading function, the red dashed curve shows the mean estimated loading function across $n_{\text{rep}} = 100$ simulation replicates, and the gray band shows the central 95% empirical range of the estimated loading functions across replicates.

Configuration C has two harmonic pairs, thus includes four altitude-dependent seasonal loading functions. The mean estimated curves generally follow the main altitude-dependent features of the true loading functions, indicating that the fitted model recovers the overall shapes of these components. The empirical ranges are wider than in Configuration A, reflecting greater replicate-to-replicate variability in the fitted loading functions. The increased variability is expected in this setting because the model must estimate more seasonal loading functions while the seasonal states themselves evolve stochastically. Although the visual agreement between the mean and true curves differs across Configurations B and C, these configurations should not be interpreted as a direct ranking because they are generated under different seasonal structures.

Overall, these results show that the altitude-dependent seasonal loading functions are recovered very well in the deterministic seasonal setting and remain qualitatively recoverable in the more challenging stochastic seasonal settings. The comparison across Configurations A, B, and C indicates that the estimated seasonal loading functions show greater replicate to replicate variability when the seasonal states are stochastic and when additional harmonic components are included. Across several panels, the empirical ranges appear somewhat wider near the boundary altitude levels. This suggests that recovery of the loading functions is less stable at the edges of the altitude grid, which is consistent with the reduced local support of spline-based representations near the boundaries.

5.5.5 Bootstrap Inference for Model Parameters

We next evaluate the empirical coverage of the percentile bootstrap confidence intervals for the scalar model parameters. We focus on the estimated scalar parameters α_0 , α_1 , σ_c^2 and σ_R^2 . The seasonal harmonic state noise variance σ_ω^2 is not included because it is fixed at the configuration-specific value during estimation.

For each simulation replicate $r = 1, \dots, n_{\text{rep}}$, the fitted model is used to generate $n_{\text{boot}} = 100$ bootstrap datasets using the innovation-based bootstrap described in Section 4.4. Let $\hat{\theta}^{*(b,r)}$, $b \in 1, \dots, n_{\text{boot}}$ denote the bootstrap estimates of a scalar parameter θ for the r^{th} replicate within that configuration. A nominal 95% percentile bootstrap confidence interval is constructed as

$$\text{CI}_{\text{boot}}^{(r)}(\theta) = \left[\hat{\theta}^{*[\alpha/2, r]}, \hat{\theta}^{*[1-\alpha/2, r]} \right], \quad (5.33)$$

where $\hat{\theta}^{*[q, r]}$ denotes the empirical q^{th} quantile of $\{\hat{\theta}^{*(1, r)}, \dots, \hat{\theta}^{*(n_{\text{boot}}, r)}\}$. In the reported results, $\alpha = 0.05$.

The empirical coverage rate is then computed across simulation replicates as

$$\widehat{\text{Cover}}_{\text{boot}}(\theta) = \frac{1}{n_{\text{rep}}} \sum_{r=1}^{n_{\text{rep}}} \mathbf{1} \left\{ \theta_{\text{true}} \in \text{CI}_{\text{boot}}^{(r)}(\theta) \right\}, \quad (5.34)$$

where θ_{true} denotes the true parameter value used in data generation. We also record the mean interval width across simulation replicates.

Figure 5.9 summarizes the empirical coverage rates of the nominal 95% percentile bootstrap confidence intervals for the scalar parameters. The points show the empirical coverage rates for Configurations A, B, and C, and the dashed horizontal line indicates the nominal 95% level. Overall, the bootstrap intervals provide reasonably good coverage, although several empirical coverage rates are slightly below the nominal level. The variance parameters σ_c^2 and σ_R^2 show particularly strong performance, with coverage values between 0.92 and 0.95 across the three configurations. The transition parameters α_0 and α_1 show slightly lower coverage, especially for α_1 in Configuration A, where the empirical coverage is 0.87. However, the remaining transition-parameter coverage rates range from 0.90 to 0.92, indicating moderate rather than severe undercoverage.




Figure 5.9: Empirical coverage rates of nominal 95% percentile bootstrap confidence intervals for the scalar parameters α_0 , α_1 , σ_c^2 , and σ_R^2 across Configurations A, B, and C. The points show the empirical coverage rates, and the dashed horizontal line indicates the nominal 95% coverage level.

Table 5.4 provides the corresponding coverage rates and mean interval widths. The interval widths generally increase from Configuration A to Configuration C for α_0 , α_1 , and σ_c^2 . This reflects greater uncertainty in the more complex configurations, especially when stochastic seasonal states are included and when the number of harmonic pairs increases. For σ_R^2 , the mean interval widths are similar across configurations, indicating stable uncertainty quantification for the observation noise variance.

Configuration	Parameter	Coverage	Mean width
A	α_0	0.92	0.3123
B	α_0	0.91	0.4098
C	α_0	0.90	0.5658
A	α_1	0.87	0.1954
B	α_1	0.90	0.2396
C	α_1	0.92	0.3572
A	σ_c^2	0.92	0.0033
B	σ_c^2	0.95	0.0047
C	σ_c^2	0.92	0.0076
A	σ_R^2	0.92	0.0027
B	σ_R^2	0.95	0.0027
C	σ_R^2	0.92	0.0027

Table 5.4: Empirical coverage rates and mean widths of nominal 95% percentile bootstrap confidence intervals for scalar parameters across $n_{\text{rep}} = 100$ simulation replicates.

Overall, the percentile bootstrap intervals provide reasonably good scalar-parameter coverage, with some mild undercoverage. In particular, the transition parameters α_0 and α_1 show greater finite-sample variability than the variance parameters, and their bootstrap intervals tend to have slightly lower coverage. Nevertheless, most coverage values are close to the nominal level, and the bootstrap intervals remain informative for quantifying scalar-parameter uncertainty across the three simulation configurations.

Chapter 6

Application to Hohenpeissenberg Ozonesonde Profiles

In this chapter, we apply the proposed state-space modeling framework to real ozone profile data from the WOUDC Hohenpeissenberg ozonesonde dataset. While the simulation study in Chapter 5 evaluates the performance of the model under controlled settings, the purpose of this chapter is to examine how the method performs on observed real-world atmospheric data, at a midlatitude location (Hohenpeissenberg in Germany). The analysis focuses on modeling the temporal evolution of ozone concentration profiles simultaneously across altitudes, with particular attention to the smooth vertical structure, and altitude-dependent seasonal behavior captured by the fitted model.

6.1 WOUDC Hohenpeissenberg Ozonesonde Data

In this chapter, we analyze ozonesonde profile observations from the Hohenpeissenberg station, obtained from the World Ozone and Ultraviolet Radiation Data Centre (WOUDC) [33]. WOUDC is part of the World Meteorological Organization Global Atmo-

sphere Watch program and serves as an international archive for ozone and ultraviolet radiation observations. The Hohenpeissenberg station was selected because it provides one of the longest ozonesonde records with vertical ozone profile measurements using Brewer-Mast instruments consistently, for estimating temporal variation and altitude-dependent ozone structure. Many analyses of stratospheric ozone trends include data from Hohenpeissenberg analyzed at different spatio-temporal scales (for example, [30, 50, 51]). Our state-space model will produce functional profile estimates on the pentad time scale, based on data interpolated to a regular altitude grid. Extending the model to incorporate responses on differing measurement observations remains a future extension.

The objective of analyzing these data is to examine whether the proposed model can capture the main features of observed ozone profile dynamics, including smooth vertical structure, temporal dependence, and altitude-dependent seasonal variation. Unlike the simulation study in Chapter 5, the true latent states and model parameters are unknown in this real-data setting. Therefore, the analysis focuses on model fit, estimated latent components, seasonal loading functions, and diagnostic assessment. Potential extensions to include additional atmospheric dynamics and chemical impacts (such as associations with the Quasi-Biennial Oscillation, impacts of aerosols from volcanic eruptions) are posed under future directions.

The rest of this chapter is organized as follows. Section 6.2 describes the preprocessing steps used to prepare the raw ozonesonde data for model fitting on a pentad (5-day) time-scale. Section 6.3 presents exploratory summaries/visualizations. Section 6.4 specifies the state-space model used for the Hohenpeissenberg ozone profiles and discusses the model-selection procedure. Section 6.5 presents the fitted model results, including parameter estimates, reconstructed ozone profiles, estimated seasonal loading functions and decomposition of the smoothed reconstruction. Section 6.6 examines model diagnostics based on standardized innovations. Section 6.7 reports bootstrap-based uncertainty

quantification. Finally, Section 6.8 summarizes the main findings from the real-data analysis.

6.2 Data Preprocessing

Each ozonesonde sounding provides measurements of ozone concentration at multiple altitudes; and soundings frequently occur irregularly in time with different altitudes in different soundings. Therefore, before applying the state-space model developed in Chapter 2, we preprocessed the raw profile data into a regular time-by-altitude array. The preprocessing steps included extracting the relevant ozonesonde variables, interpolating profiles onto a common altitude grid, and aggregating observations into 5-day pentads.

The WOUDC Hohenpeissenberg data we downloaded consisted of ozonesonde soundings for each sounding-day, where each file may contain one or more sounding profiles. These raw files are not directly suitable for our state-space model because the profiles are observed at irregular dates, may contain different vertical measurement heights, and may include incomplete or invalid observations. Therefore, the raw data are first processed into a structured altitude-by-time array before model fitting.

The preprocessing procedure consists of three main stages. First, the WOUDC data file URLs were downloaded from the WOUDC data portal for the Hohenpeissenberg station, and the corresponding daily profile files were downloaded on April 28, 2025 ([33]). For each downloaded file, the timestamp information was extracted to identify the date and launch time of each sounding. When a file contained more than one profile on the same day, the profiles were separated according to the timestamp blocks. The profile section of each sounding was then read into R , and rows containing insufficient numerical information were removed.

Second, the relevant profile ozone partial pressure and geopotential height data were

extracted and transformed for each sounding. Geopotential height was converted from meters to kilometers. In cases where the geopotential height variable was entirely missing, altitude was approximated from atmospheric pressure via the empirical formula used in [30]:

$$a = 49.3 - 7.13 \log(P),$$

where P denotes atmospheric pressure (hPa), that approximates the geopotential height formula in [52]. Ozone partial pressure values that were nonpositive were treated as missing.

Third, ozone partial pressures from each sounding were interpolated onto a common altitude grid. Specifically, linear interpolation was used to evaluate the profile variables on the regular grid from 15 km to 30 km with a spacing of 0.5 km. This produced a common set of altitude levels so that each interpolated sounding could be represented as a vector over the same vertical grid. Profiles with insufficient valid observations for a given variable were not interpolated and were instead recorded as missing for that variable. This interpolation step is necessary because the proposed state-space model assumes a fixed observation dimension and uses a fixed B-spline basis matrix evaluated at the common altitude grid.

After interpolation, the ozone profiles were aggregated to the pentad time-scale. The analysis uses observations from January 1, 1979 to December 31, 2024. For each observation date, the day of year was converted to a pentad index by

$$\text{pentad within year} = \min \left\{ 73, \left\lceil \frac{\text{day of year}}{5} \right\rceil \right\}. \quad (6.1)$$

This definition assigns days 1 to 5 of each year to the first pentad, days 6 to 10 to the second pentad, and so on. The use of the minimum with 73 ensures that each calendar year is represented by exactly 73 pentads. In non-leap years, the final pentad contains the

last five days of the year, i.e., days 361 to 365. In leap years, since the year contains 366 days, the final pentad contains six days, i.e., days 361 to 366. This convention preserves a fixed seasonal period of 73 pentads per year, matching the harmonic period $T_p = 73$ used in the state-space model.

The full time index across pentads and years was then constructed as

$$\text{pentad} = 73(\text{year} - 1979) + \text{pentad within year.} \quad (6.2)$$

When more than one profile was available within the same pentad, ozone values were averaged at each altitude level. A complete sequence of pentads from 1979 through 2024 was then created, and the aggregated ozone profiles were joined to this full pentad sequence. For pentads with no available sounding data at specific altitudes, ozone values were filled using the most recent previous available value at the same altitude level. The resulting dataset is a regular pentad-level ozone profile array, with rows corresponding to pentad time points and columns corresponding to altitude levels. This processed dataset is used as the observational input for the functional state-space model in later sections.

6.3 Exploratory Analysis

Before fitting the proposed state-space model, we first examine the main empirical features of the processed Hohenpeissenberg ozone profiles. The purpose of this exploratory analysis is to assess whether the data exhibit the two main structures targeted by the model: smooth variation across altitude and systematic seasonal variation over time. All plots in this section are based on the pentad-level ozone profile array constructed in Section 6.2.

Figure 6.1 displays average ozone profiles for four selected pentads within an annual

cycle, calculating by averaging the altitude-dependent profiles for each selected pentad across all study years. The four selected pentads are chosen to be roughly evenly spaced across the annual cycle, representing different phases of the year. For each selected pentad q , the plotted curve is computed by averaging ozone partial pressure at each altitude over all years in the study period for which the pentad-within-year equals q . Let $y_t(a_i)$ denote the component of the observation vector $\mathbf{y}_t$ corresponding to altitude level a_i . More specifically, for altitude level a_i , the plotted value is

$$\bar{y}_q(a_i) = \frac{1}{N_q} \sum_{\{t: q(t)=q\}} y_t(a_i), \quad (6.3)$$

where $q(t)$ denotes the pentad within year corresponding to time point t , and $N_q = |\{t : q(t) = q\}|$ is the number of time points contributing to pentad q . The profiles vary smoothly with altitude, with a clear ozone maximum in the stratospheric altitude range, while the magnitude and shape of the profiles differ across pentads. This smoothness supports the use of B-spline basis functions to represent the vertical ozone structure and illustrates the systematic seasonal changes in the vertical profile.

Figure 6.2 illustrates the full preprocessed pentad-level ozone profile dataset from 1979 through 2024. In this heatmap, the horizontal axis represents calendar time, indexed by pentad, the vertical axis represents altitude, and the color represents the ozone partial pressure value at that time and altitude. Equivalently, the figure visualizes the entries of the processed data matrix $\{y_t(a_i), \quad t = 1, \dots, T, \quad i = 1, \dots, n_a\}$. The heatmap shows that ozone concentration varies over both time and altitude. The vertical structure is not constant over time, and recurring temporal patterns are visible across years. This motivates the use of a dynamic model rather than a purely static altitude-profile model.

To more directly examine the annual pattern, Figure 6.3 displays the mean seasonal cycle of the ozone profiles. For each pentad within year q and each altitude level a_i , the




Figure 6.1: Averaged ozone profiles at selected pentads within the annual seasonal cycle. For each selected pentad, ozone partial pressure is averaged across all years from 1979 through 2024 at each altitude level separately.




Figure 6.2: Heatmap of the processed pentad-level ozone profiles from 1979 through 2024. The horizontal axis represents calendar time, the vertical axis represents altitude, and color represents ozone partial pressure.

plotted value is the average ozone partial pressure over all years from 1979 through 2024. Thus, the figure visualizes the full collection of mean profile values

$$\bar{y}_q(a_i) = \frac{1}{N_q} \sum_{\{t: q(t)=q\}} y_t(a_i), \quad q = 1, \dots, 73, \quad i = 1, \dots, n_a. \quad (6.4)$$

Unlike Figure 6.2, which shows the full time series, this plot averages across years and therefore emphasizes the typical within-year seasonal pattern. The seasonal variation is clearly altitude-dependent: the timing and magnitude of the annual variation differ across altitude levels. This feature is consistent with the modeling strategy in Chapter 2, where seasonal effects enter through altitude-dependent loading functions rather than through a single common seasonal effect shared by all altitudes.

Together, these exploratory plots indicate that the processed Hohenpeissenberg ozone data exhibit smooth vertical structure, temporal dependence, and a recurring seasonal cycle whose magnitude and timing vary across altitude levels. These features are consistent with the proposed state-space model, in which B-spline basis functions represent smooth altitude dependence and harmonic components capture seasonal temporal variation.

6.4 Model Specification for the Real Data

We fit the state-space model introduced in Chapter 2 to the processed Hohenpeissenberg pentad-level ozone profiles. The observation vector at time t is defined as

$$\mathbf{y}_t = (y_t(a_1), \dots, y_t(a_{n_a}))^\top, \quad (6.5)$$

where $a_1, \dots, a_{n_a}$ denote the common altitude grid. In the real-data analysis, the altitude range is 15 km to 30 km with spacing 0.5 km, so that $n_a = 31$. The time index is measured




Figure 6.3: Mean seasonal cycle of ozone profiles. For each pentad within year and altitude level, ozone partial pressure is averaged across all years from 1979 through 2024.

in pentads from 1979 through 2024, giving

$$T = 46 \times 73 = 3358 \tag{6.6}$$

time points.

To guide the choice of the number of seasonal harmonic pairs p and the number of B-spline basis functions k , we fitted a collection of candidate models and compared their AIC, AICc, and BIC values. The information criteria were computed from the Kalman-filter log-likelihood. The detailed computation of these criteria is given in Section 6.5.1.

Figure 6.4 shows the information criteria for candidate models with different values of p and k . Over the candidate grid considered, larger models generally produce smaller values of AIC, AICc, and BIC. Therefore, the information criteria alone do not select a unique parsimonious model within this range. Instead, they indicate that additional harmonic pairs and additional B-spline basis functions improve the fitted likelihood enough to offset the corresponding penalty terms.

Here we use $p = 2$ harmonic pairs as a balance between model fit, interpretability, and computational feasibility. Under the pentad time scale, the first two harmonic pairs correspond to annual and semiannual seasonal components, which have a clear interpretation for ozone profile seasonality and are consistent with harmonic seasonal representations used in previous ozone state-space analyses [23]. Although models with $p = 3$ or $p = 4$ further reduce the information criteria, they introduce additional seasonal loading functions and make the fitted seasonal structure more difficult to interpret. They also increase the computational cost of repeated model fitting, especially for the bootstrap uncertainty analysis. Therefore, we regard $p = 2$ as a practically interpretable seasonal specification rather than the strict minimizer of the information criteria.

Conditional on $p = 2$, the information criteria decrease substantially when k in-




Figure 6.4: Comparison of AIC, AICc, and BIC for candidate real-data models with different numbers of seasonal harmonic pairs p and B-spline basis functions k . The criteria generally decrease as model complexity increases. The final choice $p=2$ and $k=13$ is made as a practical compromise between model fit, interpretability, and computational feasibility.

creases from 12 to 13, indicating that the additional B-spline flexibility captures important altitude-dependent structure. Further increases from $k = 13$ to $k = 14$ and $k = 15$ continue to reduce the information criteria, but the incremental improvement is smaller than the decrease from $k = 12$ to $k = 13$ and comes with additional parameters and greater computational cost. We therefore use $k = 13$ as a parsimonious choice that captures much of the improvement while keeping the model and subsequent bootstrap analysis computationally feasible.

Based on these considerations, the real-data analysis uses $p = 2$ seasonal harmonic pairs and $k = 13$ cubic B-spline basis functions.

The model fitted to the real data follows the observation and state equations defined in equations (2.3) - (2.4). The state vector is partitioned into the time-varying B-spline coefficient component and the seasonal harmonic component, as described in equation (2.5). The B-spline coefficient vector $\mathbf{x}_{t,c}$ represents the smoothly varying nonseasonal vertical profile component, while the harmonic states represent seasonal variation over the pentad time scale.

For the vertical representation, cubic B-spline basis functions are evaluated on the altitude grid $a_1, \dots, a_{n_a}$. Based on the model-selection comparison above, we use $k = 13$ B-spline basis functions. Let $\mathbf{B}_{sp} \in \mathbb{R}^{n_a \times k}$ denote the corresponding B-spline design matrix. The observation matrix has the same structure as in equation (2.25), where

$$\mathbf{Z}_{c,j} = \mathbf{B}_{sp} \boldsymbol{\gamma}_{c,j}, \quad \mathbf{Z}_{s,j} = \mathbf{B}_{sp} \boldsymbol{\gamma}_{s,j}, \quad j = 1, \dots, p. \quad (6.7)$$

Thus, the seasonal loading functions are smooth functions of altitude, parameterized through the B-spline coefficient vectors $\boldsymbol{\gamma}_{c,j}$ and $\boldsymbol{\gamma}_{s,j}$.

The seasonal period is fixed at $T_p = 73$, corresponding to one annual cycle under the

pentad aggregation used in Section 6.2. The harmonic frequencies are therefore

$$\omega_j = \frac{2\pi j}{73}, \quad j = 1, \dots, p. \quad (6.8)$$

Based on the same comparison, we use $p = 2$ seasonal harmonic pairs. The damping factor is fixed at $d = 0.9999$. The harmonic initial states are fixed as

$$(M_{0,j}, M_{0,j}^*) = (1, 0), \quad j = 1, \dots, p, \quad (6.9)$$

to remove the rotational non-identifiability discussed in Section 2.6.

The state noise covariance matrix is specified in equation (2.13). Following the identifiability discussion in Section 2.6, the seasonal innovation variance σ_ω^2 is fixed at 1, rather than estimated jointly with the seasonal loading vectors. The observation noise covariance is assumed to be diagonal and equal across altitudes, as described in equation (2.27).

The initial values of the spline coefficient states, denoted by $\mathbf{x}_{0,c} \in \mathbb{R}^k$ (for $t = 0$ in equation 2.6), are estimated from the data, while the harmonic initial states are fixed as specified above. The set of unknown parameters estimated from the real data across the state and observation equations, are therefore

$$\boldsymbol{\theta} = (\alpha_0, \alpha_1, \sigma_c^2, \sigma_R^2, \boldsymbol{\gamma}_{c,1}^\top, \boldsymbol{\gamma}_{s,1}^\top, \dots, \boldsymbol{\gamma}_{c,p}^\top, \boldsymbol{\gamma}_{s,p}^\top, \mathbf{x}_{0,c}^\top)^\top, \quad (6.10)$$

with σ_ω^2 , d , T_p , and the harmonic initial states treated as fixed. Maximum likelihood estimation is carried out using the Kalman-filter likelihood and the numerical optimization procedure described in Chapter 4.

6.5 Fitted Model Results

6.5.1 Model Fit Summary and Parameter Estimates

The model described in Section 6.4 was fitted to the processed Hohenpeissenberg pentad-level ozone profiles using maximum likelihood estimation. In this fit, we used $k = 13$ cubic B-spline basis functions, $p = 2$ seasonal harmonic pairs, damping factor $d = 0.9999$, and fixed seasonal innovation variance $\sigma_\omega^2 = 1$. The model was fitted using the BFGS optimization method, within R library MARSS [49]. The optimization converged successfully after 335 iterations.

The maximized log-likelihood was

$$\log \mathcal{L}(\hat{\boldsymbol{\theta}}) = -131569.9 \quad (6.11)$$

For this real-data fit, the number of estimated parameters is

$$n_{\boldsymbol{\theta}} = 2pk + k + 4. \quad (6.12)$$

Here, $2pk$ is the number of seasonal loading coefficients in $\{\boldsymbol{\gamma}_{c,j}, \boldsymbol{\gamma}_{s,j}\}_{j=1}^p$ within observation matrix $\mathbf{Z}$ (equation (2.25)), k is the number of estimated initial B-spline coefficient states in $\mathbf{x}_{0,c}$, and the remaining four parameters are α_0 , α_1 , σ_c^2 , and σ_R^2 . Since $p = 2$ and $k = 13$ in the present fit, this gives

$$n_{\boldsymbol{\theta}} = 2 \times 2 \times 13 + 13 + 4 = 69. \quad (6.13)$$

The information criteria are computed from the maximized log-likelihood and the

number of estimated parameters. The Akaike information criterion is

$$\text{AIC} = -2 \log \mathcal{L}(\hat{\boldsymbol{\theta}}) + 2n_{\boldsymbol{\theta}}. \quad (6.14)$$

For the Bayesian information criterion, we use the number of time points T as the effective sample size since the Kalman-filter likelihood in equation (4.1) is written as a product of T one-step-ahead predictive densities for the multivariate observation vectors $\mathbf{y}_t$. Thus,

$$\text{BIC} = -2 \log \mathcal{L}(\hat{\boldsymbol{\theta}}) + n_{\boldsymbol{\theta}} \log(T). \quad (6.15)$$

The corrected AIC also requires a choice of effective sample size in the correction term. In this analysis, we use T rather than Tn_a as the effective sample size, for the same reason: the altitude components observed at the same time point are modeled jointly as a multivariate observation vector, rather than treated as independent replicates [53]. Therefore, we compute

$$\text{AICc} = \text{AIC} + \frac{2n_{\boldsymbol{\theta}}(n_{\boldsymbol{\theta}} + 1)}{T - n_{\boldsymbol{\theta}} - 1}. \quad (6.16)$$

Using $\log \mathcal{L}(\hat{\boldsymbol{\theta}}) = -131569.9$, $n_{\boldsymbol{\theta}} = 69$, and $T = 3358$, we obtain

$$\text{AIC} = 263277.8, \quad \text{AICc} = 263280.7, \quad \text{BIC} = 263700.0. \quad (6.17)$$

The AICc value is close to the AIC value because the number of time points is large relative to the number of estimated parameters. However, it is not exactly identical to AIC under the T -based effective sample size convention.

The main scalar parameter estimates are summarized in Table 6.1. The estimated transition parameters $\hat{\alpha}_0$ and $\hat{\alpha}_1$ determine the temporal evolution of the B-spline coefficient states through the tridiagonal coefficient block $\mathbf{B}_c$. The estimate $\hat{\sigma}_c^2$ represents

the innovation variance of the spline coefficient states, while $\hat{\sigma}_R^2$ represents the common observation noise variance across altitude levels.

Parameter	Estimate
α_0	0.3627
α_1	0.3252
σ_c^2	1.4253
σ_R^2	0.2657

Table 6.1: Estimated scalar parameters for the real-data state-space model fitted to the Hohenpeissenberg pentad-level ozone profiles.

The seasonal loading coefficient vectors $\gamma_{c,j}$ and $\gamma_{s,j}$ are also estimated as part of the observation matrix $\mathbf{Z}$. Since these coefficients depend on the chosen B-spline basis, their individual numerical values are not directly interpreted. Instead, we interpret the corresponding altitude-dependent loading functions

$$\hat{\mathbf{Z}}_{c,j} = \mathbf{B}_{sp} \hat{\gamma}_{c,j}, \quad \hat{\mathbf{Z}}_{s,j} = \mathbf{B}_{sp} \hat{\gamma}_{s,j}, \quad (6.18)$$

which describe how the seasonal harmonic states contribute to the ozone profile at each altitude. These fitted loading functions are examined in subsection 6.5.3.

6.5.2 Model-Based Reconstructed Ozone Profiles

The fitted model can be used to construct model-based reconstructions of the ozone profile at each time point. Using the one-step-ahead, filtered, and smoothed state estimates defined in equations (4.4) - (4.6), the corresponding reconstructed profiles are

$$\hat{\mathbf{y}}_{t|t-1} = \hat{\mathbf{Z}} \hat{\mathbf{x}}_{t|t-1} \quad (6.19)$$

$$\hat{\mathbf{y}}_{t|t} = \hat{\mathbf{Z}} \hat{\mathbf{x}}_{t|t} \quad (6.20)$$

$$\hat{\mathbf{y}}_{t|T} = \hat{\mathbf{Z}} \hat{\mathbf{x}}_{t|T}. \quad (6.21)$$

The one-step-ahead predicted profile uses observations only up to and including time $t - 1$, the filtered profile uses observations up to and including time t , and the smoothed profile uses the full observed time series.

Uncertainty in the smoothed reconstructed profile can be obtained from the smoothed state covariance matrix. Let $\mathbf{P}_{t|T}$ denote the covariance matrix of the smoothed state estimate $\hat{\mathbf{x}}_{t|T}$. Conditional on the fitted model parameters, the covariance matrix of the smoothed reconstructed profile is

$$\widehat{\text{Var}}(\hat{\mathbf{y}}_{t|T}) = \hat{\mathbf{Z}}\mathbf{P}_{t|T}\hat{\mathbf{Z}}^\top. \quad (6.22)$$

Therefore, a pointwise approximate 95% confidence interval for the reconstructed ozone value at altitude a_i is given by

$$\hat{y}_{t|T}(a_i) \pm 1.96\sqrt{\left[\hat{\mathbf{Z}}\mathbf{P}_{t|T}\hat{\mathbf{Z}}^\top\right]_{ii}}. \quad (6.23)$$

These intervals describe the uncertainty in the fitted latent signal due to state estimation, after propagating the smoothed state uncertainty through the observation equation. They do not include additional observation error, since the target is the reconstructed latent ozone profile rather than a new noisy observation.

Since this chapter analyzes the data retrospectively, we use the smoothed reconstructed profiles $\hat{\mathbf{y}}_{t|T}$ for visualization and interpretation.

Figure 6.5 compares the observed and smoothed reconstructed ozone profiles at four selected pentads in the year 2000. The year 2000 is chosen as a representative year near the middle of the study period 1979 - 2024, so that the comparison is not based on either the beginning or the end of the record. Within this year, the selected pentads are roughly evenly spaced across the annual cycle and correspond to different phases of




Figure 6.5: Observed and smoothed reconstructed ozone profiles at selected pentads in the year 2000, with pointwise approximate 95% confidence intervals. Observed ozone partial pressure values are shown as points, smoothed reconstructed profiles are shown as solid curves, and confidence intervals are shown as shaded bands. The selected pentads 1, 19, 37, and 55 are roughly evenly spaced across the annual cycle. The reconstructed profiles are computed as $\hat{\mathbf{y}}_{t|T} = \hat{\mathbf{Z}}\hat{\mathbf{x}}_{t|T}$ using the smoothed state estimates, and the pointwise intervals are computed from the diagonal entries of $\hat{\mathbf{Z}}\mathbf{P}_{t|T}\hat{\mathbf{Z}}^\top$.

the seasonal pattern. The observed ozone partial pressure values are shown as points, while the smoothed reconstructed profiles are shown as solid curves. The shaded bands represent pointwise approximate 95% confidence intervals obtained by propagating the smoothed state covariance matrix through the fitted observation matrix.

Across the four selected pentads, the fitted profiles capture the main smooth vertical structure of the observed ozone profiles. In particular, the fitted curves reproduce the broad ozone maximum in the stratospheric altitude range and the decrease toward lower and higher altitudes. Local deviations between the observed points and fitted curves are expected, since the model represents the vertical structure through a low-dimensional B-spline basis and separates the fitted signal from observation noise. The confidence bands are relatively narrow where the latent profile is well informed by the data and the fitted state-space structure, and they provide a visual summary of uncertainty in the reconstructed latent profile. The bands are slightly wider near the lower or upper altitude boundary; this reflects greater uncertainty in the reconstructed profile at the edge of the fitted altitude domain. Overall, the fitted profiles indicate that the model captures the dominant altitude-dependent structure of the observed ozone profiles at representative time points.

6.5.3 Estimated Seasonal Loading Functions

We next examine how the fitted model represents seasonal variation across altitude. In the proposed model, the seasonal harmonic state pairs are shared across all altitude levels, but their contributions to each altitude level are controlled by the loading functions $\mathbf{Z}_{c,j}$ and $\mathbf{Z}_{s,j}$, $j = 1, 2$. Therefore, these loading functions describe how the common seasonal harmonic dynamics are translated into altitude-dependent ozone variation. Since the raw coefficient vectors $\gamma_{c,j}$ and $\gamma_{s,j}$ depend on the chosen B-spline basis, we interpret the

altitude-domain loading functions

$$\hat{\mathbf{Z}}_{c,j} = \mathbf{B}_{sp} \hat{\boldsymbol{\gamma}}_{c,j}, \quad j = 1, 2. \quad (6.24)$$

$$\hat{\mathbf{Z}}_{s,j} = \mathbf{B}_{sp} \hat{\boldsymbol{\gamma}}_{s,j}, \quad j = 1, 2. \quad (6.25)$$

For altitude level a_i , the corresponding entries $\hat{Z}_{c,j}(a_i)$ and $\hat{Z}_{s,j}(a_i)$ determine how the two components of the j^{th} seasonal harmonic state pair contribute to the ozone profile at that altitude.

Figure 6.6 shows the estimated seasonal loading functions for the two harmonic state pairs over altitude. The loading functions vary substantially across altitude, indicating that the fitted seasonal component is not simply a common annual pattern added equally to all altitude levels. Instead, the model estimates altitude-dependent combinations of the harmonic state components. For the first harmonic, $\hat{\mathbf{Z}}_{c,1}$ is mostly positive across the altitude range, while $\hat{\mathbf{Z}}_{s,1}$ is mostly negative except near the lower boundary. For the second harmonic, the loading functions are most pronounced in the lower part of the analyzed altitude range and are closer to zero at higher altitudes. These patterns show that the two harmonic pairs contribute differently across altitude. The estimated loading functions also show some local variation near the lower boundary of the analyzed altitude range. This may reflect both stronger variability in the lower stratospheric portion of the profiles and boundary effects from estimating smooth B-spline loading functions on a finite altitude interval. Therefore, the local behavior near 15 km should be interpreted cautiously.

To summarize the overall magnitude of the fitted seasonal effect for each harmonic, we compute the estimated seasonal amplitude at each altitude level as

$$\hat{A}_j(a_i) = \sqrt{\hat{Z}_{c,j}(a_i)^2 + \hat{Z}_{s,j}(a_i)^2}, \quad j = 1, 2. \quad (6.26)$$




Figure 6.6: Estimated seasonal loading functions for the two harmonic state pairs. For each harmonic j , the solid curve shows $\hat{Z}_{c,j} = \mathbf{B}_{sp}\hat{\gamma}_{c,j}$, and the dashed curve shows $\hat{Z}_{s,j} = \mathbf{B}_{sp}\hat{\gamma}_{s,j}$. Together, these loading functions determine how the seasonal harmonic state pairs contribute to the ozone profile across altitude.

This amplitude combines the two seasonal loading functions for the j^{th} harmonic into a single measure of the magnitude of the annual harmonic contribution at altitude a_i .

Figure 6.7 indicates that the estimated seasonal amplitudes are highly altitude-dependent. The first harmonic has relatively large amplitude over a broad portion of the analyzed altitude range, with larger values in the lower and middle altitudes and somewhat smaller values near the upper part of the range. The second harmonic is most pronounced in the lower stratospheric portion of the profile, especially around 15–18 km, and decreases rapidly with altitude. Above roughly 22 km, the second harmonic amplitude is much smaller.

These results are consistent with the motivation for introducing altitude-dependent seasonal loading functions in the observation equation. Rather than forcing all altitude levels to share the same seasonal amplitude or the same combination of harmonic components, the fitted model allows each harmonic state pair to have different strengths and combinations across altitude. The estimated loading functions and amplitude therefore provide an interpretable summary of the altitude-dependent seasonal structure captured by the model. Model adequacy is assessed separately through the diagnostic analysis in Section 6.6.

6.5.4 Decomposition of the Smoothed Reconstruction

The previous subsection compared observed and smoothed reconstructed ozone profiles at selected time points in Figure 6.5. We now examine the same smoothed reconstruction over the full study period and decompose it into two profile-scale contributions. This provides a more global view of how the fitted model represents the ozone profile dynamics from 1979 through 2024.




Figure 6.7: Estimated seasonal amplitude by altitude for the two harmonic state pairs. The amplitude is computed as $\hat{A}_j(a_i) = \sqrt{\hat{Z}_{c,j}(a_i)^2 + \hat{Z}_{s,j}(a_i)^2}$ and summarizes the strength of that harmonic contribution at each altitude level.

For each time point t , the smoothed reconstructed profile can be written as

$$\hat{\mathbf{y}}_{t|T} = \mathbf{B}_{sp} \hat{\mathbf{x}}_{t,c|T} + \sum_{j=1}^p \left(\hat{\mathbf{Z}}_{c,j} \hat{M}_{t,j|T} + \hat{\mathbf{Z}}_{s,j} \hat{M}_{t,j|T}^* \right). \quad (6.27)$$

We define the smoothed spline contribution as

$$\hat{\mathbf{b}}_{t|T} = \mathbf{B}_{sp} \hat{\mathbf{x}}_{t,c|T}, \quad (6.28)$$

and the smoothed seasonal contribution as

$$\hat{\mathbf{s}}_{t|T} = \sum_{j=1}^p \left(\hat{\mathbf{Z}}_{c,j} \hat{M}_{t,j|T} + \hat{\mathbf{Z}}_{s,j} \hat{M}_{t,j|T}^* \right). \quad (6.29)$$

Thus,

$$\hat{\mathbf{y}}_{t|T} = \hat{\mathbf{b}}_{t|T} + \hat{\mathbf{s}}_{t|T}. \quad (6.30)$$

Here $\hat{\mathbf{b}}_{t|T}$ represents the profile-scale contribution of the time-varying B-spline coefficient states, while $\hat{\mathbf{s}}_{t|T}$ represents the profile-scale seasonal contribution induced by the seasonal harmonic state pairs and the estimated seasonal loading functions.

It is important to note that $\hat{\mathbf{b}}_{t|T}$ should not be interpreted as a purely nonseasonal background component. Rather, it represents the portion of the smoothed reconstruction captured by the time-varying B-spline coefficient states. Therefore, it may include slow temporal variation, interannual variation, and other non-harmonic changes in the ozone profile. In contrast, $\hat{\mathbf{s}}_{t|T}$ represents the portion of the reconstruction associated specifically with the seasonal harmonic state pairs and their altitude-dependent loading functions.

Figure 6.8 shows the smoothed spline contribution over time and altitude. This component captures the broad vertical structure of the ozone profiles, including higher

ozone values in the middle part of the analyzed altitude range and lower values near the lower and upper boundaries. Because this component is constructed through the B-spline basis, it represents the smoothly varying profile-scale contribution estimated from the time-varying B-spline coefficient states.




Figure 6.8: Smoothed spline contribution to the reconstructed ozone profiles. The heatmap shows $\hat{\mathbf{b}}_{t|T} = \mathbf{B}_{sp}\hat{\mathbf{x}}_{t,c|T}$ over time and altitude. This contribution represents the portion of the smoothed reconstruction captured by the time-varying B-spline coefficient states.

Figure 6.9 shows the smoothed seasonal contribution. In contrast to the spline contribution, this quantity represents signed deviations associated with the seasonal harmonic

state pairs. The seasonal contribution shows visible variation over time and altitude, with stronger local variation near the lower part of the analyzed altitude range and a broader positive contribution across much of the upper altitude range. This agrees with the estimated seasonal amplitude in Subsection 6.5.3, where the two harmonic contributions showed different altitude-dependent strengths.




Figure 6.9: Smoothed seasonal contribution to the reconstructed ozone profiles. The heatmap shows $\hat{s}_{t|T} = \sum_{j=1}^p \left(\hat{Z}_{c,j} \hat{M}_{t,j|T} + \hat{Z}_{s,j} \hat{M}_{t,j|T}^* \right)$ over time and altitude. Positive and negative values represent signed deviations associated with the seasonal harmonic state pairs and their altitude-dependent loading functions.

Figure 6.10 shows the total smoothed reconstruction over the full time period. While

Figure 6.5 displayed selected vertical profiles, this heatmap summarizes the reconstructed ozone profile over all pentads from 1979 through 2024. The reconstruction preserves the main altitude-time structure seen in the processed data, while the decomposition provides a model-based separation between the spline contribution and the signed seasonal contribution. This decomposition provides an interpretable summary of the fitted model: the B-spline coefficient states describe the evolving smooth vertical profile, and the seasonal harmonic state pairs contribute altitude-dependent seasonal deviations through the estimated loading functions.

6.6 Model Diagnostics

After examining the fitted reconstruction and its profile-scale contributions, we assess the adequacy of the fitted state-space model using diagnostics based on one-step-ahead innovations. For a state-space model, the primary diagnostic residual is the one-step-ahead prediction error, rather than the residual from the smoothed reconstruction [54]. This is because the Kalman filter likelihood is constructed from the one-step-ahead predictive distribution. Let $\hat{\mathbf{y}}_{t|t-1}$ defined in equation (4.13) denote the one-step-ahead predicted observation vector. The innovation is defined in equation (4.14). The corresponding innovation covariance matrix is defined in equation (4.15). The standardized innovation is then computed as in equation (4.27).

Under a well-specified linear Gaussian state-space model, the standardized innovations should be approximately centered around zero and show little remaining autocorrelation.

Figure 6.11 shows the standardized innovations. As discussed in Section 4.4, since the standardization uses the full innovation covariance matrix $\mathbf{F}_t$, the plotted rows are innovation components rather than direct altitude levels. The standardized innovations are




Figure 6.10: Total smoothed reconstruction of the ozone profiles over the full study period. The heatmap shows $\hat{\mathbf{y}}_{t|T} = \hat{\mathbf{b}}_{t|T} + \hat{\mathbf{s}}_{t|T}$ across time and altitude.

mostly centered around zero and do not show an obvious persistent trend over the study period. This suggests that, at the level of the standardized one-step-ahead innovations, the fitted model removes much of the systematic prediction structure.




Figure 6.11: Standardized innovations for the fitted real-data model, with 73 pentads within each year (horizontal axis). The standardized innovation vector is computed as $\mathbf{e}_t = \mathbf{F}_t^{-1/2} \boldsymbol{\delta}_t$, where $\mathbf{F}_t$ is the innovation covariance matrix for each time point t (i.e., for each pentad). Since the full innovation covariance matrix is used in the standardization, the vertical axis represents entries of the standardized innovation vector rather than specific altitude levels.

For the ACF diagnostics, we display representative interior components of the standardized innovation vector. As discussed in Section 4.4, these components are entries

of the decorrelated and rescaled innovation vector rather than altitude-specific residuals. Therefore, they are selected by their positions in $\mathbf{e}_t$, not by altitude. The displayed components are selected at approximately the 20%, 40%, 60%, and 80% positions of the standardized innovation vector. Boundary components are not shown in the main figure because they can be more sensitive to edge effects from the finite altitude range and from the whitening transformation.

Figure 6.12 shows that the representative interior components have small autocorrelations at most lags, with most values lying within the approximate confidence bounds. This suggests that, for these representative components, the fitted model removes much of the remaining temporal dependence in the standardized one-step-ahead prediction errors. Boundary components were also inspected and showed somewhat stronger short-lag behavior, indicating that diagnostics near the edges of the altitude range should be interpreted with additional caution.

Overall, the diagnostic plots suggest that the fitted model captures important predictable structure in the Hohenpeissenberg ozone profiles. The standardized innovation heatmap does not show a strong large-scale remaining pattern, and the representative interior ACF plots show limited remaining autocorrelation. At the same time, the boundary behavior indicates that the fit near the edges of the altitude range should be interpreted with additional caution.

6.7 Bootstrap Uncertainty Quantification

To assess estimation uncertainty for the real-data fit, we apply the innovation-based bootstrap procedure described in Section 4.4. Starting from the fitted state-space model, bootstrap datasets are generated by resampling standardized innovations and reconstructing bootstrap observation sequences through the fitted Kalman filter recursion.




Figure 6.12: Sample autocorrelation functions of representative interior standardized innovation components. The displayed components are selected at approximately the 20%, 40%, 60%, and 80% positions of the standardized innovation vector $\mathbf{e}_t = \mathbf{F}_t^{-1/2} \boldsymbol{\delta}_t$. Since the full innovation covariance matrix is used in the standardization, these components are entries of the whitened innovation vector rather than specific altitude levels.

The same model specification is then refitted to each bootstrap dataset.

In this analysis, we use 100 bootstrap samples. All bootstrap refits converged numerically. The resulting bootstrap estimates are used to summarize uncertainty in the scalar model parameters and to assess the stability of the fitted seasonal loading estimates.

Figure 6.13 shows the bootstrap distributions of the main scalar parameters. The solid vertical line in each panel indicates the estimate from the original real-data fit, and the dashed vertical lines indicate the 2.5% and 97.5% bootstrap quantiles. The bootstrap distributions for α_0 , α_1 , σ_c^2 , and σ_R^2 are concentrated over relatively narrow ranges, suggesting that these scalar parameters are reasonably stable under the innovation bootstrap.

The bootstrap percentile intervals for the scalar parameters are summarized in Table 6.2. These intervals provide an empirical measure of estimation uncertainty under the fitted model and the innovation bootstrap scheme.

Parameter	Estimate	Bootstrap 2.5%	Bootstrap 97.5%
α_0	0.3627	0.3436	0.3672
α_1	0.3252	0.3240	0.3361
σ_c^2	1.4253	1.3912	1.4794
σ_R^2	0.2657	0.2533	0.2700

Table 6.2: Bootstrap percentile intervals for scalar parameters in the real-data model.

The bootstrap behavior of the seasonal loading structure was also examined on the altitude scale. Since the raw B-spline coefficient vectors $\gamma_{c,j}$ and $\gamma_{s,j}$ are basis-dependent, we focus on the corresponding altitude-domain loading functions

$$\mathbf{Z}_{c,j} = \mathbf{B}_{sp}\gamma_{c,j}, \quad \mathbf{Z}_{s,j} = \mathbf{B}_{sp}\gamma_{s,j}, \quad j = 1, \dots, p. \quad (6.31)$$

These quantities are more directly interpretable because they represent fitted seasonal loading values at specific altitudes. Figure 6.14 shows bootstrap distributions of selected




Figure 6.13: Bootstrap distributions of scalar parameter estimates for the real-data model. The solid vertical line in each panel denotes the estimate from the original fit, and the dashed vertical lines denote the 2.5% and 97.5% bootstrap quantiles.

values of $\hat{Z}_{s,1}(a)$ and $\hat{Z}_{c,2}(a)$ at 15, 22, and 30 km. The solid vertical line in each panel denotes the estimate from the original fit, and the dashed vertical lines denote the bootstrap 2.5% and 97.5% quantiles.

Compared with the scalar parameters in Figure 6.13, these selected loading values generally show wider and less regular bootstrap distributions. In the selected examples shown here, the first-harmonic loading values, represented by $\hat{Z}_{s,1}(a)$, tend to be more dispersed across bootstrap refits. By contrast, the selected second-harmonic loading values, represented by $\hat{Z}_{c,2}(a)$, are often more concentrated near the fitted estimate, although some extreme bootstrap values are still present. Thus, the bootstrap results suggest that uncertainty in the detailed seasonal loading structure is larger than for the scalar transition and variance parameters.

Overall, the bootstrap analysis suggests that the main scalar parameters are relatively stable for the fitted real-data model, whereas uncertainty in the detailed seasonal loading structure is substantially larger. Consequently, the fitted seasonal loading functions in Subsection 6.5.3 should be interpreted as model-based estimates of altitude-dependent seasonal structure, with greater uncertainty than the scalar parameter estimates.

6.8 Discussion and Summary

This chapter applied the proposed B-spline harmonic state-space model to the Hohenpeissenberg ozonesonde profile data. The raw WOUDC profiles were first processed into a regular pentad-level altitude-by-time array over the period 1979–2024. This allowed the model to be fitted to a multivariate time series of ozone profiles on a common altitude grid from 15 km to 30 km.

The fitted model captures several important features of the processed ozone profiles. The smoothed reconstructed profiles reproduce the broad vertical structure of the ob-

Bootstrap Distributions of Selected Altitude–Domain Loading Values




Figure 6.14: Bootstrap distributions of selected altitude-domain seasonal loading values. The panels show selected values of $\hat{Z}_{s,1}(a)$ and $\hat{Z}_{c,2}(a)$ at 15, 22, and 30 km. In each panel, the solid vertical line denotes the estimate from the original fit, and the dashed vertical lines denote the 2.5% and 97.5% bootstrap quantiles.

served profiles, including the main stratospheric ozone maximum. The estimated seasonal loading functions show clear altitude dependence, indicating that the seasonal component is not represented as a single common seasonal pattern added equally to all altitudes. Instead, the harmonic state pairs contribute differently across altitude through the fitted loading functions.

The decomposition of the smoothed reconstruction provides an interpretable model-based summary of the fitted ozone dynamics. The B-spline coefficient states capture the smooth time-varying profile-scale contribution, while the seasonal harmonic state pairs contribute signed seasonal deviations through the altitude-dependent loading functions. This decomposition should not be interpreted as a unique physical separation between background and seasonal ozone variation. In particular, the spline contribution may also contain low-frequency, interannual, and other non-harmonic temporal variation.

The innovation diagnostics suggest that the fitted state-space model captures important predictable structure in the ozone profiles. The standardized innovation heatmap does not show a strong large-scale remaining pattern, and the representative interior ACF plots show limited remaining autocorrelation. However, boundary behavior indicates that the fit near the edges of the altitude range should be interpreted with additional caution.

The bootstrap analysis provides further information about estimation uncertainty. The main scalar parameters, including the transition parameters and variance parameters, show relatively stable bootstrap distributions and reasonably narrow percentile intervals. In contrast, selected altitude-domain seasonal loading values show larger bootstrap variability. This suggests that the detailed seasonal loading structure is more uncertain than the scalar parameters, even though the fitted loading functions still provide a useful model-based description of altitude-dependent seasonality.

Overall, the real-data analysis demonstrates how the proposed state-space model can be used to represent ozone profile dynamics through smooth altitude dependence and

altitude-dependent seasonal structure. At the same time, the diagnostic and bootstrap results indicate that the fitted model should be viewed as a useful but imperfect representation of the Hohenpeissenberg ozone profiles. Further refinements, including alternative choices of the number of B-spline basis functions, the number of harmonic pairs, the seasonal damping factor, and the seasonal innovation variance, may improve model fit and estimation stability. Extensions incorporating additional atmospheric covariates or more flexible error structures may also be useful for future work.

Chapter 7

Summary, Discussion, and Future Work

This chapter concludes the thesis by summarizing the main ideas, findings, and contributions developed in the preceding chapters. The thesis has focused on constructing a functional state-space modeling framework for analyzing time series of ozone vertical profiles, with particular emphasis on representing altitude dependence through B-spline basis functions and seasonal variation through a reduced-dimensional harmonic state-space structure. The chapter first reviews the main components of the proposed framework and the results of simulation and real-data studies. We then discuss the main contributions and limitations of the current work, followed by several possible directions for future research.

7.1 Summary of the Thesis

This thesis develops a functional state-space modeling framework for time series of ozone vertical profiles. The motivating data consist of ozone measurements observed

repeatedly over time and across altitude, where the scientific interest lies not only in temporal variation at individual altitude levels, but also in the smooth vertical structure of the full ozone profile. A model applied separately at each altitude level would not directly use the smoothness of the vertical profile and may become inefficient when the number of altitude levels is large, failing to borrow strength across altitudes. To address this issue, the proposed framework represents the altitude-dependent ozone profile using B-spline basis functions, so that the temporal evolution is modeled through a lower-dimensional set of latent spline coefficients, in addition to seasonally varying functional components.

The model is formulated as a linear Gaussian state-space model. In the observation equation, the ozone profile at each time point is decomposed into a combination of a smooth dynamic component and a seasonal component. The smooth dynamic component is modeled through B-spline basis functions across altitudes, while the seasonal component is represented using harmonic state variables with altitude-dependent loading functions. This formulation allows the seasonal cycle to vary smoothly with altitude while avoiding the need to introduce a separate seasonal state process at every altitude level. As a result, the model provides a reduced-dimensional and interpretable representation of both temporal and vertical variation in ozone profiles.

The estimation procedure is based on maximum likelihood, where the likelihood is evaluated using the Kalman Filter. Unknown parameters in the observation and state equations are estimated by numerical optimization of the log-likelihood. The thesis also discusses several practical modeling issues that arise in this setting, including the choice of B-spline dimension, the number of seasonal harmonics, parameter constraints, and identifiability considerations. In particular, constraints on the seasonal states or innovation variances are needed to avoid scaling ambiguities between the latent seasonal states and their corresponding altitude-dependent loading functions.

Simulation studies were used to evaluate the finite-sample behavior of the proposed model under controlled settings. These studies examined the accuracy of parameter estimation, the recovery of latent components, and the performance of bootstrap uncertainty quantification under different configurations. The simulation results show that the model can recover the main dynamic and seasonal structures, while also illustrating that some parameters are more difficult to estimate than others. The bootstrap results provide useful empirical uncertainty assessments, although the coverage behavior depends on the parameter, model configuration, and finite-sample variability.

The proposed method was then applied to ozonesonde data from Hohenpeissenberg, demonstrating how the model can be used to estimate time-evolving functional variation in the shape of the ozone profile, based on a long-term ozone profile time series interpolated to a common altitude grid. Exploratory plots show strong vertically-dependent seasonal structure in the observed data, motivating the use of a functional state-space approach. The fitted model is used to estimate the profile variation across time, including altitude-dependent seasonal behavior. The diagnostic and bootstrap analyses provide additional information about model fit and uncertainty. Together, the simulation and real-data studies illustrate the usefulness of the proposed framework for modeling ozone profile time series in a way that is both interpretable and computationally feasible, and will naturally lead to future extensions, built on the foundations of this dissertation.

7.2 Main Contributions

The primary contribution of this thesis is the development of a functional state-space modeling framework for analyzing a time series of ozone profiles (each profile being a function of altitude). The motivating data are observed repeatedly over time and across altitudes, so the model must account for both temporal dependence and smooth vertical

structure. The proposed framework addresses this by combining B-spline basis functions with a linear Gaussian state-space model. This allows each ozone profile to be treated as a smooth altitude-dependent function, while the temporal evolution of the profile is modeled through a lower-dimensional latent state process, including latent seasonal harmonics, that drive the evolution of ozone profiles at each time through the observation equation, using B-splines and seasonal harmonic loadings in the observation equation.

A central feature of the proposed model is the use of cubic B-spline basis functions to represent the vertical structure of the ozone profile as a functional object at each time. Instead of modeling the time series independently at each altitude level, the altitude dependence is represented through a fixed set of smooth basis functions, and the time-varying behavior is captured by the corresponding latent spline coefficients and latent seasonal harmonics. This formulation reduces the dimension of the dynamic system and links nearby altitude levels through the shared spline basis functions. It also gives the latent states an interpretable role, since temporal evolution in the latent spline coefficients and seasonal harmonics corresponds to changes in the smooth vertical profile across time.

The thesis develops a reduced-dimensional representation of altitude-dependent seasonal variation in the ozone profile. Seasonal dynamics are described using harmonic state variables, while their effects on the ozone profile are allowed to vary smoothly with altitude through B-spline loading functions. This construction avoids assigning a separate seasonal state process to every altitude level. As a result, the model can describe altitude-dependent seasonal cycles while keeping the state dimension relatively small. This is appropriate for ozone profile data, where seasonal behavior is important but is expected to change smoothly with altitude.

Another contribution is the introduction of an interpretable transition structure for the latent spline coefficients. The state transition matrix allows each spline coefficient at a given time to depend on both its own value and on neighboring coefficient values at

the prior time point, reflecting the local support of the B-spline basis and cross-altitude dependence in the evolution of the smooth vertical ozone profile. This structure provides a simple way to model dependence among nearby altitude regions without losing the reduced-dimensional nature of the state-space formulation.

The thesis further contributes a practical likelihood-based estimation procedure for this model. The likelihood is evaluated using the Kalman Filter, and the unknown parameters are estimated through numerical optimization. In developing this estimation framework, the thesis also addresses several practical issues that arise in fitting the model, including parameter constraints, initialization, choice of model dimension, and identifiability. In particular, the interaction between latent seasonal states and altitude-dependent loading functions creates scaling ambiguities, so appropriate constraints are imposed to remove such ambiguities.

The proposed framework is evaluated through simulation studies and a real-data application. The simulation studies examine the finite-sample behavior of parameter estimation, latent component recovery, and bootstrap uncertainty quantification under several model configurations. These results show that the model can recover the main dynamic and seasonal structures, while also showing that some parameters and derived quantities are more difficult to estimate than others. The thesis then fits the proposed model to long-term ozonesonde data from Hohenpeissenberg and demonstrates how the model can be used to decompose ozone profile variation into smooth dynamic and seasonal components. Together, these studies demonstrate how the proposed framework can be used to model seasonality and local trend structures in ozone profile time series in an interpretable and structured way.

7.3 Discussion of Findings

The results in this thesis show how the proposed functional state-space framework can be used to describe temporal and vertical patterns in ozone profile time series. By representing the vertical ozone profile through B-spline basis functions, the model avoids treating measurements at different altitude levels as unrelated time series. At the same time, the model does not require a separate latent process for every altitude level. This reduced-dimensional structure is important because ozone profile data contain strong dependence across altitude, and modeling this dependence with individual states per altitude level can quickly become computationally demanding. In addition, the true ozone profiles are functions, motivating a functional data analysis approach.

The seasonal component plays an important role in explaining ozone profile variation. The harmonic state-space representation provides a compact way to describe recurring seasonal behavior, while the altitude-dependent loading functions enable the magnitude and shape of the seasonal cycle to vary across the vertical profile. Both the simulation studies and real-data application indicate that this formulation can capture the dominant seasonal structure while maintaining a relatively low-dimensional state process. In particular, the model can capture smooth seasonal structure without substantially increasing the dimension of the latent state vector. This is one of the main practical advantages of sharing the seasonal states across altitude.

The simulation studies also show that the difficulty of estimation is not the same for all model parameters and components. Parameters controlling the main dynamic and seasonal structures can often be recovered reasonably well. Transition parameters and variance parameters may show larger variability, especially under finite-sample settings or more challenging configurations. In many cases, however, the bootstrap confidence intervals reflect this additional uncertainty through wider intervals, so the coverage re-

sults remain reasonably satisfactory even when the point estimates are more variable. This suggests that model interpretation should not rely only on point estimates. Diagnostic plots, bootstrap distributions, and derived altitude-dependent quantities provide important additional information for understanding the fitted model.

The bootstrap results provide an empirical assessment of uncertainty for both model parameters and derived quantities. In the simulation study, the confidence intervals are evaluated through coverage rates and interval widths across repeated simulation replicates. The results show that interval behavior depends on the parameter being considered and on the simulation configuration, but the coverage rates are reasonably satisfying in most cases. In the real-data application, the bootstrap procedure is used to assess uncertainty in fitted parameters and selected altitude-dependent model components. Thus, the bootstrap results should be viewed as empirical uncertainty summaries for the fitted model rather than as a complete theoretical guarantee of uncertainty quantification. The latter motivates additional study beyond this thesis.

The real-data application to Hohenpeissenberg ozonesonde profiles demonstrates how the proposed framework can be implemented and interpreted in practice. The exploratory plots show strong vertical structure, temporal dependence, and seasonal variation, all of which motivate the use of a functional state-space model. The fitted model captures the main smooth variation in the ozone profiles and provides an interpretable decomposition into dynamic and seasonal components. At the same time, the diagnostic results show that some residual structure may remain, which is expected given the complexity of long-term atmospheric profile data, and planned extensions to include additionally atmospheric covariates. These findings suggest that the proposed model can serve as a structured baseline framework, while also leaving room for future extensions that include additional atmospheric drivers or more flexible dependence structures.

Overall, the findings demonstrate that this functional state-space formulation pro-

vides an interpretable way to describe ozone profile time series by combining smooth altitude dependence with dynamic temporal evolution. The proposed framework is especially useful when the goal is to describe broad vertical-temporal patterns, seasonal behavior, and uncertainty in fitted components. The results also highlight the importance of combining statistical modeling with careful diagnostics, since no single model summary is sufficient to evaluate all aspects of fit, estimation stability, and scientific interpretation.

7.4 Limitations

Although the proposed framework provides a structured approach for modeling ozone profile time series, several limitations remain. Some modeling choices are guided by practical considerations. In particular, the number of B-spline basis functions, the number of seasonal harmonics, and certain variance constraints are selected based on a combination of model fit, diagnostic checks, identifiability considerations, and computational feasibility. These choices are reasonable for the settings considered in this thesis, but they may need to be reconsidered for other datasets or applications.

The model relies on constraints to address identifiability issues. In particular, the interaction between the latent seasonal states and the altitude-dependent seasonal loading functions can lead to scaling ambiguities if no constraints are imposed. The constraints used in this thesis reflect only one possible way of resolving the ambiguities. In particular, additional study of sensitivity to different choices of the seasonal harmonic evolution variance σ_ω^2 and the damping factor d will be important in practice, as well as studies of sensitivity and estimability when including additional atmospheric covariates with complex interactions.

Another limitation is that the current real-data analysis focuses on a single ozonesonde

station. This allows the model behavior and interpretation to be studied carefully in one long-term profile dataset, but it does not address spatial variation across stations or regions. Extending the framework to multiple stations would be needed for broader spatial comparisons.

Finally, the current model remains a linear Gaussian state-space model. This assumption provides a tractable and interpretable framework for estimation using the Kalman Filter, but it may not fully capture all nonlinear, non-Gaussian, or event-driven features of atmospheric ozone processes. These limitations do not undermine the proposed framework, but they indicate where future extensions may further improve the model.

7.5 Future Work

The current thesis develops and evaluates a functional state-space framework for ozone profile time series. Several extensions could further develop this framework and broaden its scientific use. One natural direction is to incorporate additional atmospheric covariates, such as quasi-biennial oscillation (QBO), the El Niño–Southern Oscillation (ENSO), aerosol loading, water vapor, and temperature. There are different ways to introduce these covariates into the model. If covariates are included in the observation equation, they can be interpreted as having direct effects on the observed ozone profile after accounting for the latent dynamic and seasonal components. Alternatively, if covariates are included in the state equation, they can be viewed as drivers of the latent ozone-profile evolution. The former approach may provide a more direct interpretation at the observation level, whereas the latter may better reflect the idea that atmospheric drivers influence the underlying dynamic process. Comparing these formulations and studying their statistical and scientific implications would be an important direction for future work.

Another extension is to apply the proposed framework to multiple ozonesonde stations. A multi-station analysis would allow comparison of fitted dynamic and seasonal components across different locations and could provide a broader view of vertical ozone variability. However, this extension would also introduce additional challenges. Different stations may have different altitude ranges, missingness patterns, observational histories, and data quality levels. Since the current model uses B-spline basis functions evaluated at on a standardized altitude grid, a multi-station analysis would also require careful consideration of whether to use a common altitude basis across stations or station-specific choices. More comprehensive multi-station modeling may also require station-level or spatial structure, which is beyond the scope of the current thesis.

One possible related direction is to extend the current implementation beyond using the regular time-by-altitude grid. In Chapter 6, the raw ozonesonde profiles are interpolated and aggregated onto a regular pentad-by-altitude array before model fitting. This step makes the observation vectors have a common dimension and allows a fixed B-spline design matrix to be used in the observation equation throughout the analysis. However, the state-space framework itself is not inherently limited to complete observations on a common altitude grid. Missing observations can be handled naturally in Kalman Filtering by modifying the observation equation at each time point to include only the available components of the profile. More generally, one could allow the observation matrix to vary with time, so that the B-spline basis is evaluated at the altitude levels actually observed for each ozonesonde profile. Such an extension would reduce the need for interpolation or imputing missing data (in our case a few pentads were imputed), and would make the method more directly applicable to irregularly sampled ozonesonde records in both altitude and time. Developing and evaluating this irregular-grid version of the model would be a useful direction for future work.

The model could also be refined by allowing more flexible variance and seasonal struc-

tures. For example, the current framework could be extended to allow the observation variance to vary across altitudes, or to allow different state noise variances for boundary and interior spline coefficients, although at the cost of increasing the number of parameters to be estimated in the model. These refinements may be useful when diagnostic results suggest altitude-dependent residual variability or different behavior near the boundaries of the spline basis. Another possible extension is to relax the assumption that the seasonal loading functions remain fixed over the full study period. A practical first step would be to divide the long observational record into several shorter periods and allow the seasonal loading functions to differ across periods. These periods could be chosen based on scientific considerations, such as major atmospheric events, or explored using data-driven approaches in future work. More generally, slowly time-varying seasonal loading functions could be considered, although this would introduce additional identifiability and computational challenges.

Finally, although this thesis focuses on ozone profiles, the proposed framework is not mathematically limited to ozone. The same general modeling idea could be applied to other atmospheric profile variables observed repeatedly over time and altitude, such as water vapor or temperature, or time series of other seasonally varying smooth functions across domains. In such applications, the interpretation of the dynamic and seasonal components would need to be adapted to the scientific context of the variable being studied. Nevertheless, the combination of smooth functional dependence and state-space temporal dynamics provides a general structure that may be useful for other types of atmospheric profile data.

7.6 Concluding Remarks

This thesis develops a functional state-space modeling framework for ozone profile time series by combining smooth altitude dependence with latent temporal dynamics. The proposed model uses B-spline basis functions to represent vertical structure and a reduced-dimensional harmonic formulation to describe altitude-dependent seasonal variation.

The simulation studies and real-data application illustrate how the framework can be estimated, evaluated, and interpreted in practice. They also show the importance of using uncertainty summaries and diagnostic checks when interpreting fitted model components. Overall, this work provides a foundation for further development of functional state-space models for atmospheric profile data, including extensions to additional covariates, multiple stations, more flexible model structures, and other profile variables.

Appendix A

Cubic B-spline Basis Functions

Spline basis functions provide a flexible representation for smooth functions and are widely used in regression and smoothing problems due to their local support and numerical stability [34, 35]. In this part we briefly introduce the cubic B-spline basis functions used to construct the altitude-dependent loading vectors in the proposed model.

A.1 B-spline Basis Functions

Let

$$\tau_1 = \tau_2 = \cdots = \tau_m < \tau_{m+1} \leq \cdots \leq \tau_{K+m} < \tau_{K+m+1} = \cdots = \tau_{K+2m} \quad (\text{A.1})$$

denote a sequence of knots. Here K denotes the number of interior knots. The left boundary point is given by $\tau_1 = \cdots = \tau_m$, and the right boundary point by $\tau_{K+m+1} = \cdots = \tau_{K+2m}$. The knots $\tau_{m+1}, \dots, \tau_{K+m}$ correspond to the K interior knots. The knot sequence determines the support and shape of the B-spline basis functions and thus controls the flexibility of the spline representation used in the model. This structure

corresponds to a clamped knot sequence, which ensures that the spline is well-defined at the boundaries of the domain. The B-spline basis functions are defined recursively using the Cox–de Boor recursion formula [34, 35].

For order $m = 1$, the basis functions are

$$B_{i,1}(x) = \begin{cases} 1 & \tau_i \leq x < \tau_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.2})$$

for $i = 1, \dots, K + 1$. These functions are also known as the Haar basis functions.

For order $m \geq 2$, the recursion is

$$B_{i,m}(x) = \frac{x - \tau_i}{\tau_{i+m-1} - \tau_i} B_{i,m-1}(x) + \frac{\tau_{i+m} - x}{\tau_{i+m} - \tau_{i+1}} B_{i+1,m-1}(x), \quad (\text{A.3})$$

for $i = 1, \dots, K + m$.

Each B-spline basis function has compact support on the interval $[\tau_i, \tau_{i+m}]$, which leads to sparse design matrices and stable numerical computation.

A.2 Cubic B-splines

In this thesis we use cubic B-splines, corresponding to order $m = 4$ (or equivalently, degree 3). Cubic splines are commonly used in statistical modeling because they provide sufficient smoothness while remaining computationally efficient.

A spline function can be expressed as a linear combination of the B-spline basis functions,

$$f(x) = \sum_{j=1}^{K+m} \gamma_j B_{j,m}(x) = \sum_{j=1}^{K+4} \gamma_j B_{j,4}(x), \quad (\text{A.4})$$

where $\{B_{j,m}(x)\}$ denote the B-spline basis functions of order m , and in the cubic case

$\{B_{j,4}(x)\}$ are the cubic B-spline basis functions. The coefficients γ_j are the corresponding spline coefficients.

A.3 Knots and Degrees of Freedom

As defined in section A.1, K denotes the number of interior knots and m denotes the order of the spline. The total number of B-spline basis functions is

$$k = K + m. \quad (\text{A.5})$$

In the proposed state-space model discussed in chapter 2, the notation k is used for this quantity. This is the same quantity as defined in equation (A.5).

To ensure proper boundary behavior, each boundary knot (i.e., the leftmost and rightmost knots) is repeated m times. That is, the first m knots are equal to the left boundary knot, and the last m knots are equal to the right boundary knot. This produces a knot sequence of length $K + 2m$. These are exactly the knots in equation (A.1).

A.4 B-spline Design Matrix

Let $a_1, \dots, a_{n_a}$ denote the altitude grid used in the ozone profile data. Evaluating the cubic B-spline basis functions $\{B_{j,4}(x)\}_{j=1}^k$ at the altitude grid points $a_1, \dots, a_{n_a}$ yields the B-spline design matrix

$$\mathbf{B}_{sp} = \begin{pmatrix} B_{1,4}(a_1) & \cdots & B_{k,4}(a_1) \\ \vdots & & \vdots \\ B_{1,4}(a_{n_a}) & \cdots & B_{k,4}(a_{n_a}) \end{pmatrix}. \quad (\text{A.6})$$

The matrix $\mathbf{B}_{sp} \in \mathbb{R}^{n_a \times k}$ appears as a block of the observation matrix $\mathbf{Z}$ in the

observation equation (see equation (2.25)). At the same time, it provides the evaluated spline basis used to construct the altitude-dependent seasonal loading vectors in $\mathbf{Z}$, as described in equation (2.25).

$$\mathbf{Z}_{c,j} = \mathbf{B}_{sp}\boldsymbol{\gamma}_{c,j}, \quad \mathbf{Z}_{s,j} = \mathbf{B}_{sp}\boldsymbol{\gamma}_{s,j}, \quad (\text{A.7})$$

where $\boldsymbol{\gamma}_{c,j}$ and $\boldsymbol{\gamma}_{s,j}$ are spline coefficient vectors estimated within the state-space model.

Appendix B

Details for Numerical Optimization

This appendix provides additional details for the numerical optimization procedure described in Section 4.5. These details include the initialization setup, optimization settings, and convergence assessment used in fitting the proposed state-space model.

B.1 Initialization Setup

The numerical optimization procedure requires the initial values for all free parameters in the model. These include the free components of the transition matrix $\mathbf{B}$, the state covariance matrix $\mathbf{Q}$, the observation matrix $\mathbf{Z}$, and the observation covariance matrix $\mathbf{R}$.

The transition matrix $\mathbf{B}$ contains free parameters only in the spline coefficient block $\mathbf{B}_c$ (as defined in equation (2.17)). The harmonic block, given by the damped rotation matrix $d\mathbf{R}(\omega_1)$ in equation (2.16), is treated as fixed. In particular, the damping factor d is regarded as known and is not estimated. Consequently, only the spline block $\mathbf{B}_c$ contributes free parameters to $\mathbf{B}$.

The state covariance matrix $\mathbf{Q}$ is specified as diagonal. Variance parameters are

associated with the spline coefficient states. For the harmonic states, the corresponding innovation variance σ_ω^2 is treated as fixed in order to ensure identifiability, as discussed in Chapter 2.6. Depending on the modeling assumption, σ_ω^2 may be set to 0 (deterministic harmonic dynamics) or to other fixed constants (stochastic harmonic innovations), for example, 1 or 0.1.

The observation matrix $\mathbf{Z}$ contains both fixed B-spline basis components and free seasonal loading parameters. Here, the seasonal loading vectors $\{\gamma_{c,1}, \dots, \gamma_{c,p}\}$ and $\{\gamma_{s,1}, \dots, \gamma_{s,p}\}$ described in Section 2.5 are initialized to 0.

The observation covariance matrix is specified as

$$\mathbf{R} = \sigma_R^2 \mathbf{I}_{n_a} \quad (\text{B.1})$$

is parameterized by a single variance parameter. And the initial value of σ_R^2 is set to be 0.05.

The initial state vector is specified based on simple heuristics consistent with the scale of the observed data. In simulation studies, it is constructed from a simulated ozone profile projected onto the B-spline basis, while in real data applications it is initialized using simple deterministic values. In both cases, the initialization is chosen to ensure numerical stability of the optimization.

B.2 Optimization Settings

The negative log-likelihood function is evaluated using the Kalman filter recursion described in Chapter 4. At the i^{th} iteration of the optimization, given the current parameter estimate $\boldsymbol{\theta}_i$, the Kalman filter produces the innovations $\boldsymbol{\delta}_t$ and their covariance matrices $\mathbf{F}_t$, which are used to compute the objective function in (4.26).

The optimization is carried out using the BFGS algorithm. The algorithm iteratively updates the parameter vector in order to reduce the negative log-likelihood. Convergence is assessed based on changes in the objective function and changes in the parameter estimates across iterations.

1. **Log-likelihood evaluation:** Compute the negative log-likelihood

$$f(\boldsymbol{\theta}_i) = -\log \mathcal{L}(\boldsymbol{\theta}_i) \quad (\text{B.2})$$

by running the Kalman filter with the current parameter value $\boldsymbol{\theta}_i$.

2. **Gradient computation:** Let $\nabla f(\boldsymbol{\theta}_i)$ denote the gradient of the objective function.

In practice, this gradient is approximated numerically using finite differences.

3. **Parameter update:** Update the parameter estimate using the BFGS update rule:

$$\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i - \eta_i \mathbf{H}_i \nabla f(\boldsymbol{\theta}_i), \quad (\text{B.3})$$

where $\mathbf{H}_i$ is an approximation to the inverse Hessian of f at iteration i , and η_i is a step size obtained via a line search procedure.

The step size η_i is chosen to satisfy the sufficient decrease (Armijo) condition:

$$f(\boldsymbol{\theta}_i - \eta_i \mathbf{H}_i \nabla f(\boldsymbol{\theta}_i)) \leq f(\boldsymbol{\theta}_i) - c_1 \eta_i \nabla f(\boldsymbol{\theta}_i)^\top \mathbf{H}_i \nabla f(\boldsymbol{\theta}_i), \quad (\text{B.4})$$

where $c_1 \in (0, 1)$ is a small fixed constant.

4. **Inverse Hessian update:** The inverse Hessian approximation $\mathbf{H}_i$ is updated using the BFGS formula:

$$\mathbf{H}_{i+1} = \mathbf{H}_i + \frac{\mathbf{r}_i \mathbf{r}_i^\top}{\mathbf{r}_i^\top \mathbf{s}_i} - \frac{\mathbf{H}_i \mathbf{s}_i \mathbf{s}_i^\top \mathbf{H}_i}{\mathbf{s}_i^\top \mathbf{H}_i \mathbf{s}_i}, \quad (\text{B.5})$$

where

$$\mathbf{s}_i = \boldsymbol{\theta}_{i+1} - \boldsymbol{\theta}_i, \quad (\text{B.6})$$

$$\mathbf{r}_i = \nabla f(\boldsymbol{\theta}_{i+1}) - \nabla f(\boldsymbol{\theta}_i). \quad (\text{B.7})$$

In implementation, the BFGS algorithm is carried out using the `optim()` function in R, where gradients are computed numerically and the line search is performed using default settings.

B.3 Convergence and Practical Considerations

Convergence of the numerical optimization is assessed using the diagnostic information returned by the optimization routine. In particular, convergence is determined based on the stopping criteria of the BFGS algorithm, which typically involve thresholds on changes in the objective function and parameter estimates.

In simulation and bootstrap studies, convergence information is recorded for each model fit. This allows identification of runs in which the optimization procedure may not have converged successfully. Such cases are monitored when evaluating parameter estimates and constructing bootstrap confidence intervals.

When performing repeated model fitting, such as in bootstrap procedures, we carefully manage computational resources and ensure that outputs from different runs are properly organized.

Bibliography

- [1] G. P. Brasseur and S. Solomon, *Aeronomy of the Middle Atmosphere*. Springer Dordrecht, 3 ed., 2005.
- [2] R. P. Wayne, *Chemistry of atmospheres : an introduction to the chemistry of the atmospheres of earth, the planets, and their satellites*. Oxford University Press, 3rd ed. ed., 2000.
- [3] World Meteorological Organization, “Scientific assessment of ozone depletion: 2022.” <https://csl.noaa.gov/assessments/ozone/2022/>, (2022).
- [4] M. J. Molina and F. S. Rowland, *Stratospheric sink for chlorofluoromethanes: Chlorine atom-catalysed destruction of ozone*, *Nature* **249** (1974) 810–812.
- [5] J. C. Farman, B. G. Gardiner, and J. D. Shanklin, *Large losses of total ozone in Antarctica reveal seasonal ClO_x/NO_x interaction*, *Nature* **315** (1985) 207–210.
- [6] S. Solomon, *Stratospheric ozone depletion: A review of concepts and history*, *Reviews of Geophysics* **37** (1999), no. 3 275–316.
- [7] W. J. Randel, R. S. Stolarski, D. M. Cunnold, J. A. Logan, M. J. Newchurch, and J. M. Zawodny, *Trends in the vertical distribution of ozone*, *Science* **285** (1999), no. 5434 1689–1692.
- [8] S. Godin-Beekmann, N. Azouz, V. F. Sofieva, D. Hubert, I. Petropavlovskikh, P. Effertz, G. Ancellet, D. A. Degenstein, D. Zawada, L. Froidevaux, S. Frith, J. Wild, S. Davis, W. Steinbrecht, T. Leblanc, R. Querel, K. Tourpali, R. Damadeo, E. Maillard Barras, R. Stübi, C. Vigouroux, C. Arosio, G. Nedoluha, I. Boyd, R. Van Malderen, E. Mahieu, D. Smale, and R. Sussmann, *Updated trends of the stratospheric ozone vertical distribution in the 60°S – 60°N latitude range based on the LOTUS regression model*, *Atmospheric Chemistry and Physics* **22** (2022), no. 17 11657–11673.
- [9] N. R. P. Harris, B. Hassler, F. Tummon, G. E. Bodeker, D. Hubert, I. Petropavlovskikh, W. Steinbrecht, J. Anderson, P. K. Bhartia, C. D. Boone, A. Bourassa, S. M. Davis, D. Degenstein, A. Delcloo, S. M. Frith, L. Froidevaux, S. Godin-Beekmann, N. Jones, M. J. Kurylo, E. Kyrölä, M. Laine, S. T. Leblanc,

- J.-C. Lambert, B. Liley, E. Mahieu, A. Maycock, M. de Mazière, A. Parrish, R. Querel, K. H. Rosenlof, C. Roth, C. Sioris, J. Staehelin, R. S. Stolarski, R. Stübi, J. Tamminen, C. Vigouroux, K. A. Walker, H. J. Wang, J. Wild, and J. M. Zawodny, *Past changes in the vertical distribution of ozone – part 3: Analysis and interpretation of trends*, *Atmospheric Chemistry and Physics* **15** (2015), no. 17 9965–9982.
- [10] World Ozone and Ultraviolet Radiation Data Centre, “WOUDC contributor guide.” <https://guide.woudc.org/en/>, (2025). Accessed 2026-05-31.
- [11] H. G. J. Smit and ASOPOS, “Quality assurance and quality control for ozonesonde measurements in GAW.” https://library.wmo.int/doc_num.php?explnum_id=7167, (2014). Accessed 2022-08-17.
- [12] B. Hassler, I. Petropavlovskikh, J. Staehelin, T. August, P. K. Bhartia, C. Clerbaux, D. Degenstein, M. D. Mazière, B. M. Dinelli, A. Dudhia, G. Dufour, S. M. Frith, L. Froidevaux, S. Godin-Beekmann, J. Granville, N. R. P. Harris, K. Hoppel, D. Hubert, Y. Kasai, M. J. Kurylo, E. Kyrölä, J.-C. Lambert, P. F. Levelt, C. T. McElroy, R. D. McPeters, R. Munro, H. Nakajima, A. Parrish, P. Raspollini, E. E. Remsberg, K. H. Rosenlof, A. Rozanov, T. Sano, Y. Sasano, M. Shiotani, H. G. J. Smit, G. Stiller, J. Tamminen, D. W. Tarasick, J. Urban, R. J. van der A, J. P. Veefkind, C. Vigouroux, T. von Clarmann, C. von Savigny, K. A. Walker, M. Weber, J. Wild, and J. M. Zawodny, *Past changes in the vertical distribution of ozone - part 1: Measurement techniques, uncertainties and availability*, *Atmospheric Measurement Techniques* **7** (2014), no. 5 1395–1427.
- [13] G. C. Reinsel, A. J. Miller, E. C. Weatherhead, L. E. Flynn, R. M. Nagatani, G. C. Tiao, and D. J. Wuebbles, *Trend analysis of total ozone data for turnaround and dynamical contributions*, *Journal of Geophysical Research: Atmospheres* **110** (2005), no. D16.
- [14] S. Dhomse, M. Weber, I. Wohltmann, M. Rex, and J. P. Burrows, *On the possible causes of recent increases in northern hemispheric total ozone from a statistical analysis of satellite data from 1979 to 2003*, *Atmospheric Chemistry and Physics* **6** (2006), no. 5 1165–1180.
- [15] D. I. Vyushin, V. E. Fioletov, and T. G. Shepherd, *Impact of long-range correlations on trend detection in total ozone*, *Journal of Geophysical Research: Atmospheres* **112** (2007), no. D14.
- [16] M. Weber, C. Arosio, M. Coldewey-Egbers, V. E. Fioletov, S. M. Frith, J. D. Wild, K. Tourpali, J. P. Burrows, and D. Loyola, *Global total ozone recovery trends attributed to ozone-depleting substance (ODS) changes derived from five merged ozone datasets*, *Atmospheric Chemistry and Physics* **22** (2022), no. 10 6843–6859.

- [17] Y. Li, S. S. Dhomse, M. P. Chipperfield, W. Feng, J. Bian, Y. Xia, and D. Guo, *Quantifying stratospheric ozone trends over 1984–2020: a comparison of ordinary and regularized multivariate regression models*, *Atmospheric Chemistry and Physics* **23** (2023), no. 20 13029–13047.
- [18] A. C. Harvey, *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, Cambridge, 1989.
- [19] J. Durbin and S. J. Koopman, *Time Series Analysis by State Space Methods: Second Edition*. OUP Catalogue. Oxford University Press, 2012.
- [20] R. H. Shumway and D. S. Stoffer, *Time Series Analysis and Its Applications*. Springer Cham, 5 ed., 2025.
- [21] M. West and J. Harrison, *Bayesian Forecasting and Dynamic Models*. Springer, New York, 2 ed., 1997.
- [22] G. Petris, S. Petrone, and P. Campagnoli, *Dynamic Linear Models with R*. Use R! Springer, New York, 2009.
- [23] M. Laine, N. Latva-Pukkila, and E. Kyrölä, *Analysing time-varying trends in stratospheric ozone time series using the state space approach*, *Atmospheric Chemistry and Physics* **14** (2014), no. 18 9707–9725.
- [24] E. Maillard Barras, A. Haefele, R. Stübi, A. Jouberton, H. Schill, I. Petropavlovskikh, K. Miyagawa, M. Stanek, and L. Froidevaux, *Dynamical linear modeling estimates of long-term ozone trends from homogenized Dobson Umkehr profiles at Arosa/Davos, Switzerland*, *Atmospheric Chemistry and Physics* **22** (2022), no. 21 14283–14302.
- [25] K. Nilsen, R. Kivi, M. Laine, D. Poyraz, R. Van Malderen, P. von der Gathen, D. W. Tarasick, L. Thölix, and N. Jepsen, *Time-varying trends from Arctic ozonesonde time series in the years 1994–2022*, *Scientific Reports* **14** (2024), no. 1 27683.
- [26] J. Lee and J. O. Berger, *Space–time modeling of vertical ozone profiles*, *Environmetrics* **14** (2003), no. 6 617–639.
- [27] J. O. Ramsay and B. W. Silverman, *Functional Data Analysis*. Springer Series in Statistics. Springer, New York, 2 ed., 2005.
- [28] F. Ferraty and P. Vieu, *Nonparametric Functional Data Analysis: Theory and Practice*. Springer Series in Statistics. Springer, New York, 2006.
- [29] J. S. Morris, *Functional regression*, *Annual Review of Statistics and Its Application* **2** (2015) 321–359.

- [30] W. Meiring, *Oscillations and time trends in stratospheric ozone levels*, *Journal of the American Statistical Association* **102** (2007), no. 479 788–802.
- [31] A. Park, S. Guillas, and I. Petropavlovskikh, *Trends in stratospheric ozone profiles using functional mixed models*, *Atmospheric Chemistry and Physics* **13** (2013), no. 22 11473–11501.
- [32] K.-L. Chang, O. R. Cooper, A. Gaudel, I. Petropavlovskikh, and V. Thouret, *Statistical regularization for trend detection: an integrated approach for detecting long-term trends from sparse tropospheric ozone profiles*, *Atmospheric Chemistry and Physics* **20** (2020), no. 16 9915–9938.
- [33] WMO/GAW Ozone Monitoring Community, “World Meteorological Organization-Global Atmosphere Watch Program (WMO-GAW)/World Ozone and Ultraviolet Radiation Data Centre (WOUDC): OzoneSonde Data.” WOUDC [data set], <https://doi.org/10.14287/10000008>, (2025). Retrieved April 28, 2025.
- [34] C. de Boor, *A Practical Guide to Splines*. Applied Mathematical Sciences. Springer New York, NY, 1 ed., 2001.
- [35] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York, NY, 2 ed., 2009.
- [36] L. Pogorelyuk, C. W. Rowley, and N. J. Kasdin, *An efficient approximation of the Kalman Filter for multiple systems coupled via low-dimensional stochastic input*, *Automatica* **117** (2020) 108972.
- [37] L. Paninski, *Fast Kalman Filtering on quasilinear dendritic trees*, *Journal of Computational Neuroscience* **28** (2010) 211–228.
- [38] S. Noschese, L. Pasquini, and L. Reichel, *Tridiagonal Toeplitz matrices: properties and novel applications*, *Numerical Linear Algebra with Applications* **20** (2013), no. 2 302–326.
- [39] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge University Press, 2 ed., 2012.
- [40] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Johns Hopkins University Press, 4 ed., 2013.
- [41] H. Lütkepohl, *New Introduction to Multiple Time Series Analysis*. Springer Berlin, Heidelberg, 1 ed., 2005.
- [42] M. Johnson, P. C. Caragea, W. Meiring, C. Jeganathan, and P. M. Atkinson, *Bayesian dynamic linear models for estimation of phenological events from remote sensing data*, *Journal of Agricultural, Biological, and Environmental Statistics* **24** (2019), no. 1 pp. 1–25.

- [43] T. J. Rothenberg, *Identification in parametric models*, *Econometrica* **39** (1971), no. 3 577–591.
- [44] D. S. Stoffer and K. D. Wall, *Bootstrapping state-space models: Gaussian maximum likelihood estimation and the Kalman Filter*, *Journal of the American Statistical Association* **86** (1991), no. 416 1024–1033.
- [45] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer, 2 ed., 2006.
- [46] R. H. Shumway and D. S. Stoffer, *An approach to time series smoothing and forecasting using the em algorithm*, *Journal of Time Series Analysis* **3** (1982), no. 4 253–264.
- [47] E. E. Holmes, “Derivation of an EM algorithm for constrained and unconstrained multivariate autoregressive state-space (MARSS) models.” <https://arxiv.org/abs/1302.3919>, 2013. Accessed 2026-06-02.
- [48] E. E. Holmes, M. D. Scheuerell, and E. J. Ward, “Analysis of multivariate time series using the marss package.” <https://cran.r-project.org/web/packages/MARSS/vignettes/UserGuide.pdf>, (2025). Version 3.11.10.
- [49] E. E. Holmes, E. J. Ward, M. D. Scheuerell, and K. Wills, “Marss: Multivariate autoregressive state-space modeling.” <https://CRAN.R-project.org/package=MARSS>, (2024). R package version 3.11.9.
- [50] J. A. Logan, I. A. Megretskaya, A. J. Miller, G. C. Tiao, D. Choi, L. Zhang, R. S. Stolarski, G. J. Labow, S. M. Hollandsworth, G. E. Bodeker, H. Claude, D. De Muer, J. B. Kerr, D. W. Tarasick, S. J. Oltmans, B. Johnson, F. Schmidlin, J. Staehelin, P. Viatte, and O. Uchino, *Trends in the vertical distribution of ozone: A comparison of two analyses of ozonesonde data*, *Journal of Geophysical Research: Atmospheres* **104** (1999), no. D21 26373–26399, [<https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/1999JD900300>].
- [51] W. Steinbrecht, V. A. Velazco, R. Dirksen, L. Doppler, P. Oelsner, R. Van Malderen, H. De Backer, E. M. Barras, R. Stübi, S. Godin-Beekmann, and A. Hauchecorne, *Ground-based monitoring of stratospheric ozone and temperature over germany since the 1960s*, *Earth and Space Science* **12** (2025), no. 3 e2024EA003821, [<https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2024EA003821>]. e2024EA003821 2024EA003821.
- [52] J. R. Holton and G. J. Hakim, *An introduction to dynamic meteorology*. Academic press, 5 ed., 2013.

- [53] C. M. Hurvich and C.-L. Tsai, *A corrected Akaike information criterion for vector autoregressive model selection*, *Journal of Time Series Analysis* **14** (05, 1993) 271 – 279.
- [54] A. C. Harvey and S. J. Koopman, *Diagnostic checking of unobserved-components time series models*, *Journal of Business & Economic Statistics* **10** (1992), no. 4 377–389.