Nonlinear Data Assimilation: towards a prediction of the solar cycle.

Andreas Svedin

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2013

ABSTRACT

Nonlinear Data assimilation: towards a prediction of the solar cycle.

Andreas Svedin

The solar cycle is the cyclic variation of solar activity, with a span of 9-14 years. The prediction of the solar cycle is an important and unsolved problem with implications for communications, aviation and other aspects of our high-tech society. Our interest is *model-based prediction*, and we present a self-consistent procedure for parameter estimation and model state estimation, even when only one of several *model variables* can be observed.

Data assimilation is the art of comparing, combining and transferring observed data into a mathematical model or computer simulation. We use the 3DVAR methodology, based on the notion of least squares, to present an implementation of a traditional data assimilation. Using the Shadowing Filter – a recently developed method for nonlinear data assimilation – we outline a path towards model based prediction of the solar cycle. To achieve this end we solve a number of methodological challenges related to *unobserved variables*. We also provide a new framework for interpretation that can guide future predictions of the Sun and other astrophysical objects.

Contents

List of Figures

List of Tables

ACKNOWLEDGMENTS

This thesis exists due to dedicated support of many marvelous people. This author humbly takes this opportunity to express gratitude and appreciation for their efforts. If a name has been omitted, please know that you are not forgotten, especially to all my Swedish friends still very dear to me.

I first choose to thank my family: my mother Elisabeth, my father Uno and my sister Ingela. You have always told me to walk my own path and have given me the strength, confidence and conviction to pursue it.

Thomas Erlandsson in Uppsala taught me what it means to be a scientist, and Bengt Gustafsson put me on the path of nonlinear science and chaos studies, and finally the world made sense. It was also Bengt's ticket that in 1996 brought me to Copenhagen and the *Nonlinearities - from Turbulent to Magic Conference* where I met Ed Spiegel and had the first conversation about going to the US and Columbia University.

The trip to New York and my time at Columbia were made possible through the Fulbright program, and eventually Gålö-foundation and Sweden-America Foundation (Through the office of the Marshal of the realm of Sweden). This work would not have been completed without the support of the NORDITA, The Columbia Graduate School of Arts and Science, the Columbia Astronomy Department and once again the Gålö-foundation. Without my supervisors, Ed Spiegel and Milena Cuellar I would not have advanced to achieve this important life goal. In Ed I found a person that thinks like me, only faster, deeper and wiser. It has been a privilege to be in the presence of the master.

Milena is the one that have kept me in the present and connected to reality, and she taught me the importance of details. Her perseverance contributed significantly to the completion of this study. I also wish to thank my collaborators, Axel Brandenburg, Lenny Smith and Hailiang Du. Axel has been my mentor away from home, Lenny and Du have guided me through the theoretical and practical aspects of the shadowing filter. Without Lenny I would not have known how to approach the problem, and the performance of my code was made possible by Du's expertise. The proofreading and comments by Christine, Colin, Erika, Uno, Åsa, Lia, Maureen, Jan, El-ad and Linda have been invaluable. The submission would not have been possible without the LaTeX support of Pehr Söderman. Thanks to you all.

My strongest supporter and champion has been Mildred Garcia in the astronomy department office, Millie, you are a star. With the help of Ayoune, she made this academic journey a smooth ride. Thank you Greg for helping setting up the defense. Nelson, the ever-present sage of Pupin hall has been a great source of wisdom. Thank you for telling me that time is not scarce, time is just something that happens now. Time is never running out, we are the ones that run out.

All my classmates that helped me through coursework and my 5 years at Columbia: Dave S, Andrew, Cameron, Erika, Christine, Jennifer, Yuan, Lia, Josh, Emily, Maureen, Jeff, Loren, Jenna, Jana and Colin to name a few. My Columbia friends outside of the department: Ian, Iva, Dionisios, Cherie, Valentine, Mi, Timor, Whitney, Andrew, Emily, Jose, Liya, Spike, Ryan, Monica, Ruairidh, Ricco, Bruce, Michel, Diana, Julian, Eszter, Nika, Christian, Max and Ben. I specially want to thank Monica, Rayat, Tao and Antonios

for showing me the high towers of the Ivy League.

Other people that have been important to my being here in New York are Illion, Cortney, Doug, Richard, Millie, Robert, Phillip, Ana, Katharina, Megan, Bach, Bob, Kalvin, Salim, Matt, Liza, Marchal, Megan, Joakim, Susanne, Linda, Ted and Adella. Special thanks to El-ad, Eynat and Jan for being a family away from home and the rest of the Siwanowicz clan for adopting me.

My darling wife Åsa has been my safe haven during this whole adventure in hidden dimension and strange worlds. She has inspired, encouraged and motivated me even when she just wanted me to come home. If my work will nudge any mountain it is because of her, giving me a stable point to stand. Thank you, thank you more than I can say.

In conclusion, to quote the namesake of my high school Thomas Thorild and my undergraduate alma mater:

Tänka fritt är stort

men tänka rätt är större.

Chapter 1

Introduction

A snapshot of the night sky reveals little of the vast story of the universe that changes and evolves on all time scales. Only through studying the dynamical evolution of nature can we come to understand her past, present and future behavior.

1.1 Background and context

To understand observed *systems* scientists and mathematicians construct mathematical *models* that describe and express some of the observed behaviors. Using these models we can make a *prediction* of future behavior that later can be tested with observations. This study focuses on the application – and some methodological aspects – of *data assimilation* and prediction of nonlinear systems associated with the dynamics of the *solar cycle*. The solar cycle is introduced and discussed in Section 1.2.

Data assimilation is a general term for a procedure designed to match and transfer information from an observed system to a dynamical model. It allows one to find an

approximate solution to the inverse problem, i.e. finding the underlying signal or the process responsible for model outputs or a set of observations. Wunsch (1996) uses the following introduction to data assimilation in his book:

> "*Assimilation* is a term widely used in numerical weather prediction (NWP) to describe the process of combining a forecast with current observations in order to update a dynamical model in preparation for another forecast. The term has more recently come to describe loosely any data/model combination [...] Much of NWP can be regarded as a form of engineering-in which the forecaster's major problem is to find something that works, and only secondary to understand exactly why."(Wunsch 1996)

This study will explore a specific setup where the model used for data assimilation has a mathematical structure that is equivalent to the observed system, i.e. it is possible to find a set of model variables such that the model and systems are identical. The conditions whereby model state estimation is possible in an *ideal case*[1] will be shown even for nonlinear or chaotic systems (McSharry & Smith 1999; Judd & Smith 2001).

In Chapter 4 we will extrapolate what we learned from the ideal case and discuss the application of the methodology to a *real case*: a prediction of the solar cycle using a low-dimensional model. In the real case, any model will be an approximation of the behavior of the observed system and it is necessary to account for this difference in mathematical structure when model state estimates are made.

This work assumes that we have a model, though we may not know a correct set

[1]This corresponds to the *perfect model scenario* discussed in literature (Judd & Smith 2001).

of model parameters. One challenge is to determine what parameter values generate model trajectories that reflect the observed behavior and provide accurate predictions. The parameters found could also eventually be used to infer physical properties of the observed system.

In general, standard interpolation or least square fits can be used to assimilate data into estimates – in order to compare analytic expressions or numerical outputs – if the main behavior of an observed system can be considered to be linear. These common activities – not typically identified as data assimilation – do not require additional theoretical considerations, in the context of this work. When the model is high-dimensional or the behavior of the system is complex and nonlinear then the translation of the observations into consistent model states requires more elaborate methodologies, such as the ones presented and used in this work.

Data assimilation has been successfully developed throughout the last decades, especially regarding meteorology and oceanography (Wunsch 2006). One goal of geophysics is to provide real-time, model-based predictions of very complex and high-dimensional systems such as weather, climate, or large-scale oceanographic flow patterns. There are no fundamental differences between geophysical, astrophysical, or heliophysical systems in terms of the mathematics used to model, analyze, and describe them. The methods applied and developed for terrestrial phenomena and the methodological results obtained should thus be applicable to the general case.

Much effort has been devoted to the determination, identification and separation of multi-periodic astrophysical signals (Carroll & Ostlie 1996). Irregular or semi-regular

signals have always existed that appear to resist these decompositions into superimposed periodic components (Buchler & Kolláth 2001). Classic Miras, RV-Tauri and different types of novae are examples of systems where non-linear or even stochastic behavior might be prevalent (Buchler et al. 2002). In such systems, standard methods of data assimilation and prediction, designed for linear or multi-period signals, could encounter difficulties (Berliner 1991).

Time Domain Astrophysics has been identified as one of the main frontiers for discovery by 2010 National Research Council's Decadal Survey of Astronomy and Astrophysics (National Research Council 2010). In this approach the sky is monitored continuously or recurrently for transient and variable signals of astrophysical origin. The search for exo-planets and the early detection of supernovae and other transient objects have been driving forces in time domain exploration and for the deployment of space missions, e.g., CoRoT (Baglin 2003), Kepler (Borucki et al. 2010) and to some extent Gaia (Perryman et al. 2001) to support such investigations. There are also ground based surveys that will provide immense catalogs of astrophysical time-series. Examples include the Palomar Transient Factory (PTF[2]) and eventually the Large Synoptic Survey Telescope (LSST[3]) and the Square Kilometer Array (SKA[4]). All these endeavors in time domain astrophysics can to some extent be regarded as aspects of data assimilation. Different kinds of models and the possibility of making predictions will prove interesting for most of these applications. The methods and discussions of this work could, for example, be applicable to the exploration of irregular or semi-regular systems and signals.

[2]Palomar Transient Factory, *http* : *//www.astro.caltech.edu/ptf*
[3]Large Synoptic Sky Survey, *http* : *//lsst.org*
[4]Square Kilometer Array *http* : *//skatelescope.org*

First, general data assimilation in the context of simplified solar convection models is presented (Chapter 2). It uses the classic *3DVAR* data assimilation scheme (Lorenc 1986; Daley 1993). The chapter illustrates how traditional data assimilation operates, explains the problem it tries to address, and presents required assumptions and approximations. 3DVAR is based on statistical arguments and all uncertainties and correlations for the model components and the data have to be explicitly known or pre-defined (Lorenc 1986). We use a two-dimensional model of stratified convection, but any model of similar complexity could have been used to illustrate the procedure.

Second, *nonlinear data assimilation* is presented. Specific tools are applied to determine model states and parameters, especially when it is not possible to observe all of the model variables (Chapter 3). Here we use the *Shadowing Filter* method to compare different sets of the model parameters (Smith et al. 2010). This method presents a new approach to data assimilation recently developed in the context of nonlinear systems (Judd & Smith 2001). The use of the word filter is connected to its origin in noise reduction (e.g. Hammel 1990; Davies 1993, 1994). The Shadowing Filter focuses on the dynamics and geometry of the chosen model (Judd et al. 2008). The agreement with the observations that represent the system of interest is also important.

Third, once the methodologies to select model parameters and to produce model-based predictions are implemented, we test a solar cycle model (Chapter 4). The main results of Chapter 4 are presented in Section 4.5. We also look towards a nonlinear prediction of the solar cycle, under real conditions, i.e. using real data.

Finally, all these components are integrated into a larger framework and future pos-

sible work is outlined and recommended (Chapter 5).

In Appendix A, a comparison between the notation used in the context of 3DVAR (Chapter 2) and the Shadowing Filter (Chapter 3 and 4) is made. In Appendix B, notes about the numerical implementation of the Shadowing Filter are given. And in Appendix C a modified solar cycle model is sketched.

1.2 The solar cycle

Solar activity, as manifested in the number of spots visible on the solar surface, varies cyclically on a time scale of 9-14 years, with a mean period, over the last 23 cycles, of 131 ± 14 months or about 11 years (Hathaway et al. 2002; Hathaway 2010). This is referred to as *the solar cycle*. Sunspots are produced when magnetic flux tubes protrude through the solar photosphere and locally inhibit convection (Thomas & Weiss 2008). Photographs, made with suitable exposures for the general intensity of the photosphere, show darker spots in these regions of inhibited convection, though they are only about 20% cooler than the ambient medium. The *relative Sunspot number* R was defined by Rudolf Wolf in the mid 19th century (Usoskin 2008) as the combination of the number of Sunspots N and Sunspot groups G:

$$R = k(10G + N) \tag{1.1}$$

where k is a normalizing factor based on the observer (Usoskin 2008). This has given a framework where multiple sources can be combined into a unified record that stretches

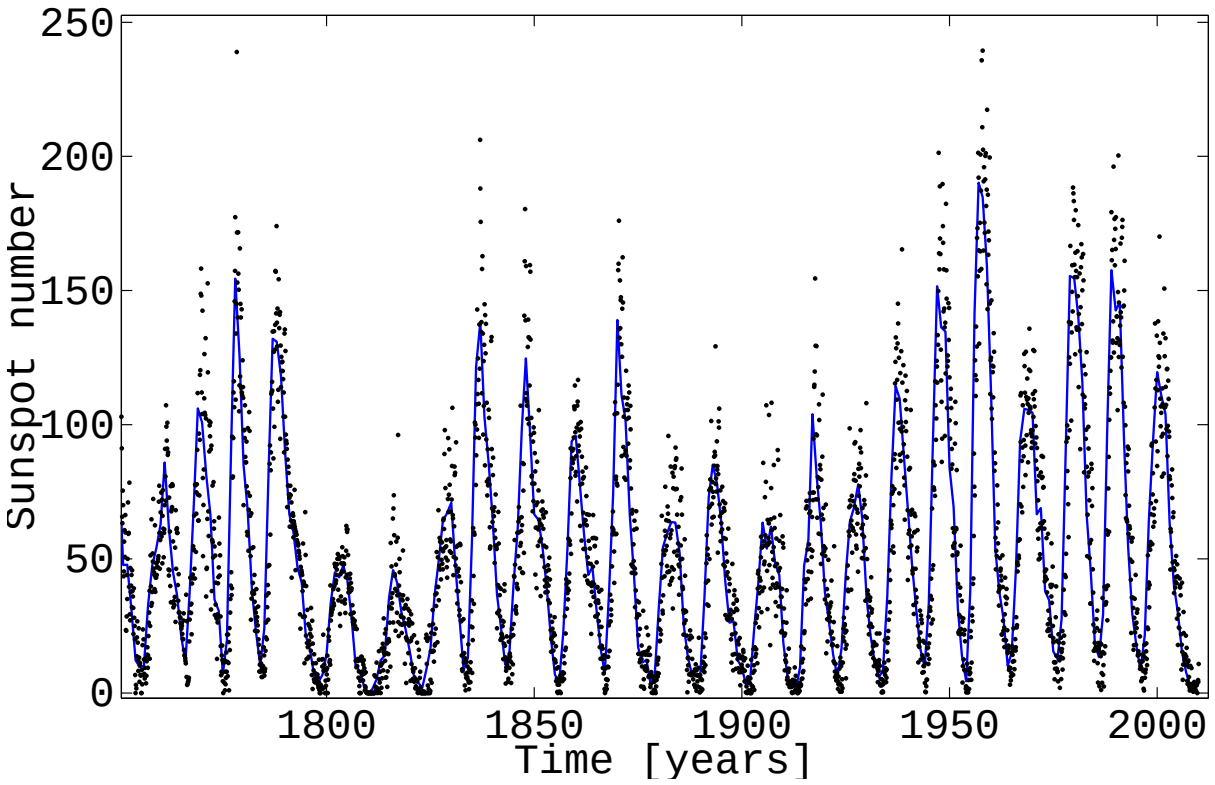

Figure 1.1 The Sunspot number since 1755. Yearly averages (blue line) and monthly averages (black dots).

back about 300 years (Usoskin 2008). This is the quantity that we are going to use throughout this work.

Figure 1.1 shows the yearly (blue line) and monthly (black dots) averaged Sunspot number since 1755, where the average is the accumulated Sunspot number divided by the number of days. It represents one of the longest time-series available in astrophysics and is possibly one of the most studied (Thomas & Weiss 2008).

The record stretches even further back in time to the beginning of the 17th century. However, this historical record is incomplete and does not provide a continuous sequence. It is also difficult to assess how accurate these older estimates are, since few or no in-

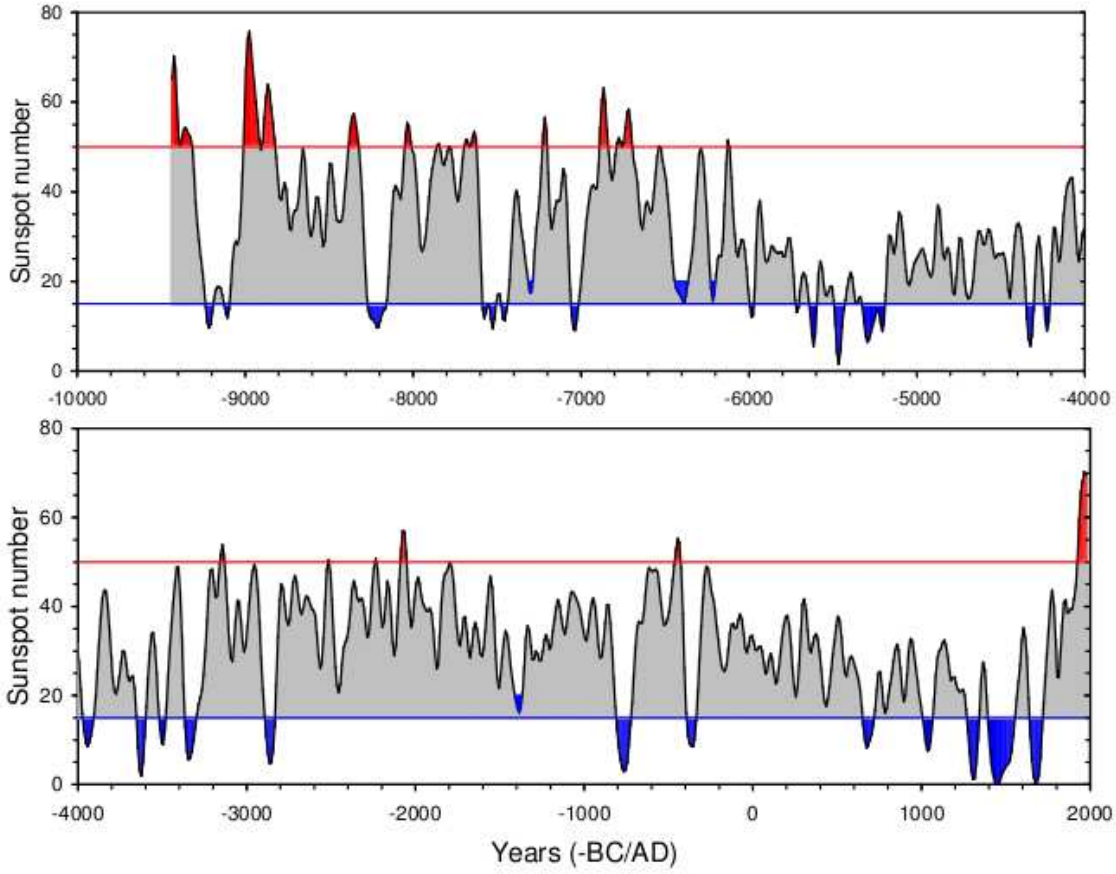

Figure 1.2 The long term variation of solar activity using geomagnetic ^{15}C data of Yang et al. (2000). Reconstructed Sunspots (decades) are smoothed with a 1-2-2-2-1 filter. Red denotes grand maxima and blue denotes grand minima. From Usoskin (2008).

dependent observations exist. There are also other indirect proxies of solar activity that indicate its long term variation. An example, using geomagnetic data, is seen in Figure 1.2.

Although the number of Sunspots could suggest a crude measure of solar activity, it is correlated to other physical quantities associated with solar processes. This is graphically represented in Figure 1.3, where these correlations over the last decades are plotted. The Figure shows the total solar irradiance (C) the 10.7 cm cm radio flux (B) the number of

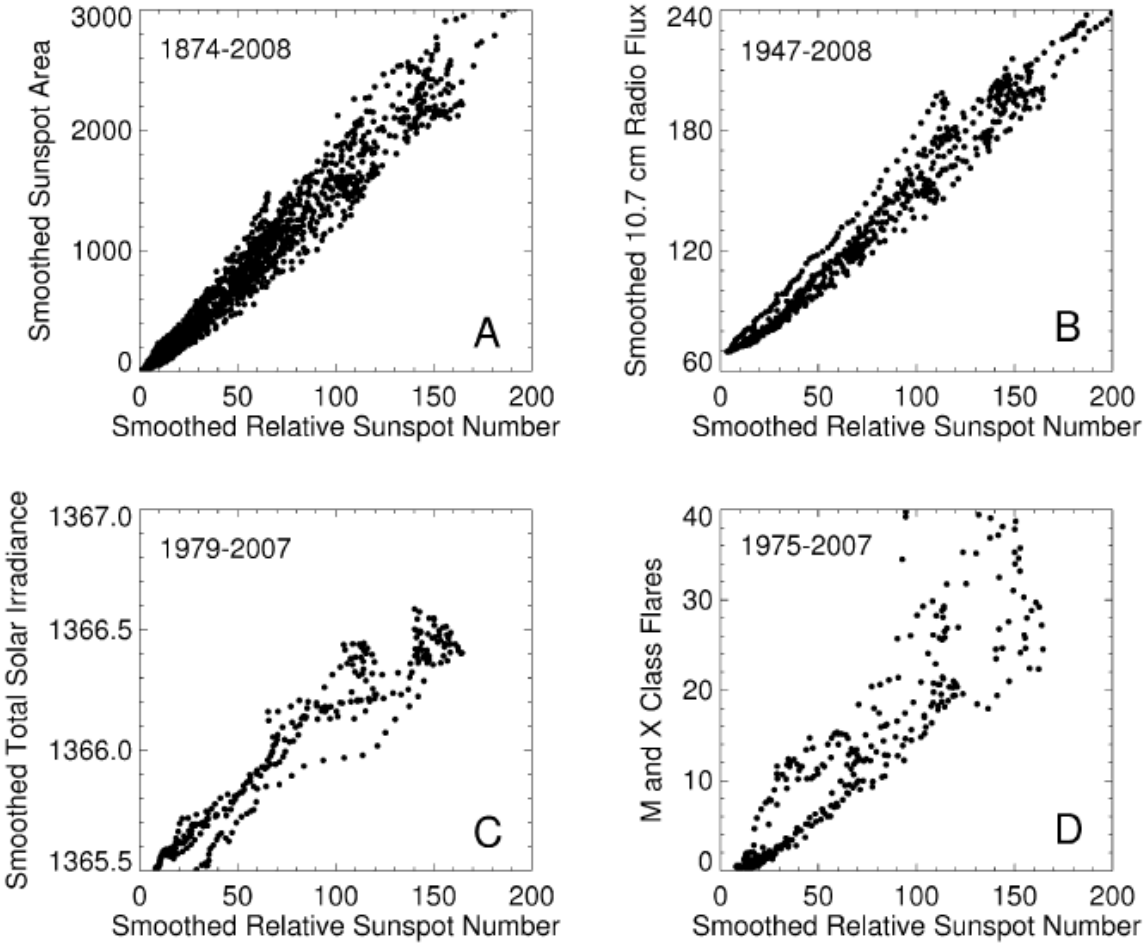

Figure 1.3 Correlation between the relative Sunspot number and other quantities of interest: (A) The Sunspot Area, (B) The 10.7 cm Radio Flux, (C) The Total Solar Irradiance and finally (D) The Number of Solar Flares. From Hathaway (2009).

solar flares (D) and the Sunspot area (A).

In addition to the average eleven-year variation, other time scales are revealed by the long-term solar record. Most notable among these is the *grand minima* occurring approximately every 200 years (see Figure 1.2), when the level of solar activity becomes very low and the detection of a significant spot may be an occasion for publication announcing its observation (Ribes & Nesme-Ribes 1993; Eddy 1976). A famous example is the *Maunder*

minimum (Maunder 1894; Eddy 1976; Usoskin 2008) that lasted approximately the course of Newton's lifetime (1642-1727). The time of low activity at the beginning of the 19th century has been named the *Dalton minimum* in a similar way. Figure 1.2 indicates the times of low activity in blue, for example the Maunder and Dalton Minima towards the end of the record.

Usoskin et al. (2007) showed – using this record of over 10,000 years – that the Sun actually spends about 17% of its time in these phases of low activity. Looking at a population of solar-type stars, one observes that about 10% are also in such inactive phases (Baliunas et al. 1995). Measurements of the solar magnetic field have shown a sharp decline in the mean Sunspot field over the last two decades (Penn & Livingston 2006; Livingston & Penn 2009). It is conjectured that, if the decline holds, the Sun will produce very few Sunspots in 2015 and for some time thereafter. Abreu et al. (2008) studied the long term variations of grand solar minima and maxima and found that the current level of solar activity has been unusually prolonged. Using a statistical argument based on the frequency and length of previous grand solar maxima, Abreu et al. (2008) state that it would be more typical for the current grand solar maximum to decline than to continue at this level. Lockwood (2009) explains that there is an 8% chance that the Sun will enter a Maunder-type minimum when compared to the decline of previous grand maxima (Barnard et al. 2011).

In light of these findings, a possible model-based prediction of the onset of a coming grand minimum would be of interest. Unfortunately, model-based prediction has not been able to address questions related to these indications since no consistent approach

has been presented in this context. The aim of this work is to present techniques that may be of use in making such a prediction.

In addition the activity of the Sun affects *space weather*, i.e. the level of activity in the solar wind, magnetosphere, ionosphere etc. (Schwenn 2006). It is of importance to the operation of satellites, space travel, and other areas of communication (Barnard et al. 2011). Space weather also affects earth's atmosphere, with possible connections between the solar cycle and climate variations (Lean et al. 1995; Haigh 2007). This makes prediction of the solar cycle a study of multidisciplinary interest. The potential historical correlation between the Maunder minimum and the so-called 'little ice age' – when the surface temperature, at least in northern Europe, was about one degree centigrade lower than average (Lean et al. 1995) – will also make this work of interest to climate studies (Lockwood 2009; Barnard et al. 2011). The cooling expected from lower solar activity due to a grand solar minimum is estimated to be approximately 0.3 degrees and would not be sufficient to counteract the expected effect of anthropogenic greenhouse gases emissions (Barnard et al. 2011).

The amplitude of individual solar cycle peaks does not seem to follow any obvious pattern. Much ingenuity and effort have already been invested into solar cycle prediction to find relevant patterns and relations (see Hathaway 2009, 2010; Petrovay 2010, for reviews). The current cycle (number 24 counted since 1755) has proven extra difficult to predict, see Section 1.3.

The solar behavior is regular but not periodic, indicating that the underlying dynamical mechanism might be nonlinear (Spiegel 1997, 2009), making standard (linear)

methods less applicable. Over the last decades, considerable effort has been invested into the creation of low-dimensional mathematical models of the solar cycle (Weiss (2010) and references therein). There are two main ways to approach the model design process (Wilmot-Smith et al. 2005).

The first way is to derive a set of ordinary differential equations by truncating the full set of magnetohydrodynamic dynamo equations. The difference between the models comes from the choice of terms kept in the truncation and what assumptions and simplifications are made. The advantage of this approach is that each term has a clear physical interpretation, but the models obtained might not always be robust (Wilmot-Smith et al. 2005). This approach can be exemplified with the sixth order model by Jones et al. (1985) describing the interaction between the average poloidal and toroidal fields and fluctuations in the differential rotation. In the seventh order model by Feudel et al. (1993) it is the mean magnetic helicity that is explicitly modeled together with the poloidal and toroidal fields.

The second way to design a model is to use mathematical insight and *normal form theory* to look for mathematical structures that generate the general observed behavior (Spiegel 1994; Knobloch et al. 1998; Wilmot-Smith et al. 2005). *Astromathematics*, as this approach has been called (Spiegel 1994), can identify important processes and mechanisms, and guide our understanding of complex systems. An example is the model by Tobias et al. (1995) found from physical arguments where the linear dynamo action leads to exponential growing oscillations that are saturated with a nonlinear Lorenz force. The components of the model correspond to the poloidal and toroidal fields and all other dy-

namical information are contained in a third component of the model. The Tobias et al. (1995) model can also be found from normal form theory where the general behavior of the system can be changed from *periodic*, to *quasi-periodic* and eventually to *chaotic* with the use of a control parameter and a specific bifurcation structure. In the fourth order model by Knobloch & Landsberg (1996) the interaction is between the magnetic dipole and quadrupole moments of the magnetic field. The sixth order Knobloch et al. (1998) model combines the transition from active to inactive phases expressed by the model Tobias et al. (1995) and the regular modulation of the solar cycle is expressed by the Knobloch & Landsberg (1996) model. This creates a unified model of the solar variation. Mathematical insight was also involved in the development of the Platt et al. (1993b) model, another unified model of the solar variation discussed in detail below.

These models capture several qualitative features of the solar cycle. However, quantitative assessment and direct comparison with real data have not previously been possible (Pasquero 1995; Weiss 2010) due to the lack of nonlinear data assimilation techniques.

Many students of the solar cycle now believe that the origin of the solar cycle is in the *solar tachocline*, i.e. the differentially rotating layer at the base of the solar convection zone (Spiegel 1972; Spiegel & Weiss 1980; Spiegel & Zahn 1992; Hughes et al. 2007; Spiegel 2009). The tachocline has also been observed and studied using helioseismology (e.g. Brown et al. 1989; Charbonneau et al. 1999). The origin of the Sunspots and their overall variation is connected to the generally unsolved problem of how the solar magnetic field is generated. It denotes one of the fundamental questions in solar physics (NASA Heliophysics road-map 2009)[5]. Better understanding of the solar variations could guide

[5] $http://sec.gsfc.nasa.gov/sec\_roadmap.htm$ (Accessed August 2012)

the solution of this "dynamo problem". Solar variations should also be used to verify any proposed dynamo theory.

One model of particular interest for the solar cycle has been outlined and explored by Platt, Spiegel and Tresser (Platt et al. 1993b), referred to as *PST*. The model was originally made by Spiegel (1980). This model treats the convection zone dynamo and the tachocline dynamo as a chaotic system coupled to a nonlinear oscillator. Its temporal behavior is reminiscent of the observed variation of the Sunspot number (Platt et al. 1993b; Pasquero 1995). The amplitude of the nonlinear oscillator is modulated by the chaotic system and the *on-off intermittency mechanism* plays an important role in generating periods of inactivity (Platt et al. 1993a; Spiegel 1994, 1997). Pasquero (1995) compares several models of the solar cycle and finds that the PST model best reproduces the observed behavior of the solar cycle. The PST model can generate cycle amplitude variations and times of suppressed solar activity reminiscent of the observed solar behavior. This is the model we study in the context of solar cycle predictions.

1.3 Predictions of the solar cycle

A variety of methods has been applied to the solar cycle prediction problem. It is possible, 2-3 years into a cycle, to anticipate the shape of the current solar cycle by applying templates and curve fitting (e.g. McNish & Lincoln 1949; Hathaway et al. 1994). Since the current cycle, number 24, started a few years ago (November 2009) these methods have predicted[6] the cycle to peak in the spring of 2013 with an amplitude of about 60.

[6] $http://solarscience.msfc.nasa.gov/predict.shtml$

A similar estimate gives a peak value of 65-75 in the middle or end of 2012 (Owens et al. 2011). This would correspond to an atypical cycle with a late onset and low peak amplitude (Lockwood et al. 2012).

Analysis of the solar cycle as a multi-periodic signal can be used to study the overall trend (e.g. Gleissberg 1939; Gnevyshev & Ohl 1948; Ahluwalia 1998) but these methods provide approximately the same accuracy as predicting that the next cycle will correspond to the mean peak value of all the previous cycles (Hathaway 2009). Precursor methods using measurement of variation magnetic fields on Earth due to solar wind activity (Ohl 1966; Feynman 1982; Thompson 1993; Hathaway et al. 1999) or direct measurements of solar magnetic activity (Schatten et al. 1978; Schatten & Sofia 1987; Svalgaard et al. 2005) give an estimate of the peak solar cycle amplitude a factor of two more accurate than predicting the mean peak value (Hathaway 2009).

Using dynamo models and feeding these with observational data offer some promise (Dikpati & Charbonneau 1999; Dikpati et al. 2008). Choudhuri et al. (2007) predict that the current cycle (number 24) should be a weak and early solar cycle with a maximum of 80 Sunspots, whereas Dikpati et al. (2006) predict a strong and late cycle with 150 Sunspots. This can be compared to the mean solar cycle peak of 114 ± 40 (Usoskin 2008). Both predictions incur problems since the current cycle appears to be late and low – disagreeing in some aspect with both prognoses (Hathaway 2009).

Precursor methods and the dynamo models that are fed with data look only at the previous cycle or the last few cycles. They will not, in their current form, be able to capture large scale trends or the possibility of grand minima. There is always a limit

to the precision of any measurement; furthermore, any error in the initial conditions will grow in time and eventually become larger than any signal, making prediction far into the future very difficult. Errors in initial conditions, inaccuracy of parameter estimates, and differences between the mathematical structure of the model and the observations, are effects that may lower predictability. These are shortcomings that need to be addressed to facilitate consistent and reliable predictions. Such limitations have been used to argue against feeding data directly into mean-field dynamo models (Tobias et al. 2006; Cameron & Schüssler 2007; Bushby & Tobias 2007).

The application of data assimilation methods for solar weather and solar cycle prediction has been envisioned (Brun 2007) and explored (Kitiashvili & Kosovichev 2008) using the *Ensemble Kalman Filter*. The Ensemble Kalman Filter (EKF) is a data assimilation method that uses multiple predictions (an ensemble) to generate an empirical expression for the uncertainty (Kalnay 2003; Evensen 2009). The estimate corresponds to the mean value of the ensemble. For a comparison between the Shadowing Filter and the Ensemble Kalman Filter see Judd (2003). Parameter estimation has also been attempted for a dynamo model (Jouve et al. 2011) but not in the context of prediction.

The result of the data assimilation methods is very dependent on the accuracy of the model chosen to represent the system of interest. The challenge is to construct a dynamo model that exhibits and reproduces the most interesting features of the solar cycle. The occurrence of grand minima and a realistic variation of peak amplitude are two such important features. Our work represents the Sun with a five-dimensional model where only one of the variables is directly observed.

To date it is unclear how to assimilate *unobserved* variables into a dynamically consistent way, without using the methods outlined in Chapter 3. In addition the issue that the mathematical structure of the model might be different from the system, has not been fully addressed before, see discussed in Chapter 4.

Most prediction methods estimate whether a coming cycle will have higher or lower activity than the previous cycle and do not include the possibility of predicting a grand minimum (Hathaway 2009). The challenge raised by the nonlinear aspect of the solar cycle is hard to avoid, especially for methods constructed for linear systems. However, in the context of the Shadowing Filter (outlined in Chapter 3 and 4), the geometry of chaotic systems is used explicitly to our advantage, during all stages of the procedure.

1.4 The problem

To understand future activity of the Sun, one must analyze its past behavior. This study will use the yearly Sunspot number since 1755, obtained from the Solar Influence Data Center (SIDC) at the Royal Observatory of Brussels (Clette et al. 2007). Data assimilation will be used to translate observations into model states. At times, data assimilation will also be applied to select optimal sets of model parameter values.

Two different data assimilation methods are implemented: 3DVAR (Chapter 2) and the Shadowing Filter (Chapters 3 and 4). Once the corresponding current state of the system is understood in terms of the model, a model-based prediction of the future state of the Sun (Chapter 4) can be explored and discussed.

For a prediction, the key components are the choice of a model, the determination

of what parameter values to use, and an estimate or an ensemble of estimates of the current state of the system. In this work, prediction is strictly considered as an initial value problem. The best model estimate of the current system state is propagated into the future with a set of determined parameters values. The accuracy of any model-based prediction is dependent on the quality of the model. Outside the ideal case, in the real world, discrepancies will always occur between the model and the system.

When the mathematical structure of the model and the observed system are different, notions, including: "best set of parameters", "optimal estimate", "noise", or even "the true state" do not necessary connote the same meaning as in the ideal case. They are not even well defined.

In general, it is more important for models used for prediction of real systems to be reasonably realistic in order to make accurate predictions rather than just being close to the observations. For nonlinear astrophysical systems that operate on time scales from seconds to years, data assimilation will be of great importance, especially when quantitative agreement between model and observations is assessed. The ultimate verification of a model is its ability to offer reliable prediction, for which data assimilation is necessary.

1.5 Hypothesis

Our work is based on the hypothesis that some deterministic phenomenon in the Sun is expressed through the manifestation of Sunspots. We model the process with a low-dimensional chaotic system. The simplest mathematical structure that can successfully predict the solar cycle is the "optimal" phenomenological model. A model-based pre-

diction would give estimates with an accuracy at least comparable to the other solar cycle predictions schemes would be considered successful. The components of a successful model and their interaction should help us to identify how different processes in the Sun interact to generate the solar cycle. In the end, the model components should also be related to specific solar dynamo processes to connect the model and the system to the mathematical structure derived from physical first principles. This study focuses on methodological aspects of data assimilation and prediction to facilitate further inquiries into the solar cycle and similar phenomena.

1.6 Technical statement of the problem

The time evolution of the model $F : \mathbf{K} \rightarrow \mathbf{K}$ is described by a trajectory in a state space $\mathbf{K} \subseteq \mathbb{R}^d$ spanned by the d variables of the model. A *sequence space* of dimension $\mathbb{R}^{N \times d}$ will also be explored whereby every point is a sequence or set of states over a fixed time interval $n = 1, ..., N$. Dynamics and geometric relations will be discussed in relation to these spaces.

The deterministic model $F(X_n)$ generates a unique image at time $n + 1$, given a general model state $X_n \in \mathbb{R}^d$ at time n. In 3DVAR an estimate $Z_N \in \mathbb{R}^d$ is made at a specific time N. This estimate is a compromise, between the model *background state* $F(X_{(N-1)})$ (the prediction from a previous state $X_{(N-1)}$) and the observation $S_N \in \mathbb{R}^k$, weighted by the relative uncertainty of these two components. The estimate is found by minimizing the weighted squared error of the background state and the observation and will be closer to the more reliable or less uncertain of the two. For this statistical

approach, the distribution of the uncertainty is assumed for the model estimate and for the observations. If the uncertainties are Gaussian, then the estimate is a Maximum Likelihood estimate (Lorenc 1986).

If the behavior of the system is considered to be *deterministic*, the sequence of system states sought should be *causally connected*.

Starting with a sequence of observations $\{S_n\}$ at times $n = 1, 2, ..., N$, the Shadowing Filter look in sequence space for a dynamically consistent set of estimates $\{Z_n\}$ also at times $n = 1, 2, ..., N$. Specifically we want to find a set $\{Z_n\}$ such that $F(Z_n)$ is close to Z_{n+1} for all n. This approach rests only on dynamical properties of the model and the system. It does not require any assumptions about the statistical properties of the observations or the model error (Judd et al. 2008).

In practice, two limitations exist that cannot be avoided in real world applications and should thus be addressed by any data assimilation approach. The first limitation is that all variables might not be observed, that is the observables in $\mathbb{R}^k$ are fewer than the model variables in $\mathbb{R}^d$. The second limitation is that the mathematical structure of the chosen model most likely differs from the mathematical structure of the observed system and may not capture some of the interesting features of the system, that is, the model is not ideal. The Shadowing Filter can account for these limitations without significant changes in the methodology (Judd & Smith 2004; Judd 2008), but 3DVAR or other data assimilation methods would require major changes.

This study also prefers the Shadowing Filter over 3DVAR since the former does not require any initial assumptions about the statistical properties of the observations or of

the model behavior. The Shadowing Filter delivers estimates of optimal parameter values and system states from which measures of the uncertainties related to the model and observations can be calculated.

1.7 General framework of the thesis

The question to be answered by the data assimilation method is where in relation to the observations the estimate should be. 3DVAR makes a weighted compromise in state space between the model background state and the observations. The Shadowing Filter has a similar compromise in sequence space between the set of observations and a trajectory of the model.

In the ideal case – where the mathematical structure of the model and system are equivalent – the goal of state estimation is to find a model trajectory. The main justification for leaving the observations behind is that, in the ideal case, the model and system are, by definition, the same. In this case – where the mathematical structure of the model and system is equivalent – the model and the system would generate identical output, given the same parameter values and initial conditions.

In real world applications – where the model and system might be structurally different – more weight needs to be put on the observations. In this case model and systems would not necessarily generate the same output even with identical initial conditions. In the specific case of unobserved variables, outlined in Section 3.6, the basic approach is to keep the observations fixed and only find dynamically consistent estimates for the unobserved variables.

With the question – where in relation to the observations the estimate should be – in mind one can make estimates and know whether the trust has been implicitly put in the observations or in the model. Knowing the level of trust put in the model and observation should guide our choice of data assimilation method and help us adapt the approach to specific applications.

Chapter 2

Data assimilation for stratified convection[1]

2.1 Abstract

We show how the 3DVAR data assimilation methodology can be used in the astrophysical context of a two-dimensional convection flow. We study the way this variational approach finds best estimates of the current state of the flow from a weighted average of model states and observations. We use numerical simulations to generate synthetic observations of a vertical two-dimensional slice of the outer part of the solar convection zone for varying noise levels and implement 3DVAR when the covariance matrices are scalar. Our simulation results demonstrate the capability of 3DVAR to produce error estimates of system states up to three orders of magnitude below the original noise level present

[1]This chapter is a reformatted and updated version of an article by the same name by A. Svedin, M. Cuellar and A. Brandenburg. Submitted to The Monthly notices to the Royal Astronomical Society.

in the observations. This work exemplifies the importance of applying data assimilation techniques in simulations of stratified convection.

2.2 Introduction

When using models to describe the temporal evolution of observed complex systems we are confronted with a number of challenges. An immediate difficulty in dealing with this question is that we generally do not know in all detail the current state of the system or the initial condition that is to be used. Lacking such information prevents us from keeping a model-based simulation in step with the behavior of the observed system. *Data assimilation* techniques offer means to address such challenges for complex systems by keeping a computer simulation (i.e. model) in synchronization with observations of the system it represents. It provides a general framework for simultaneously comparing, combining, and evaluating observations of physical systems and output from computer simulations.

The methods used in data assimilation have been developed over several decades, primarily in meteorology and oceanography for the prediction of future behavior. Data assimilation is used daily in operational weather prediction (Bengtsson et al. 1981), in climate forecast (Palmer & Hagedorn 2006) and it was even used to correct the path of the Apollo spacecraft during the first moon landings (Cipra 1993). There is a large and growing body of literature including several monographs (Daley 1993; Kalnay 2003; Wunsch 2006) and work discussing its theoretical foundations (Lorenc 1981, 1986; Le Dimet & Talagrand 1986; Ghil 1989). Astrophysical data assimilation has recently been discussed

by (Brun 2007), both in the context of space weather and in solar cycle prediction (Dikpati 2007; Choudhuri et al. 2007; Kitiashvili & Kosovichev 2008) as well as in the context of dynamo models (Jouve et al. 2011).

Here we focus on the three-dimensional variational (3DVAR) data assimilation technique, also known as sequential approach (Daley 1993), which produces updates of the current state of a model simulation at times when system observations are available. Propagation of model states between times when the system is observed is done with free simulations of the model initiated at the latest state estimate. An extension of 3DVAR to implicitly incorporate dynamical information is known as four-dimensional variational (4DVAR) data assimilation. 3DVAR dynamically evolves the mean state whereas 4DVAR also evolves other statistical properties of the model dynamics.

State estimates produced by 3DVAR are optimal provided that the model is linear and the uncertainties are Gaussian. In other words, 3DVAR states are Best Linear Unbiased Estimates, where *best* and *optimal* refer to the lowest possible mean squared error of the estimate (Kalman 1960; Talagrand 1997).

For nonlinear models, error statistics may become non-Gaussian even when the initial distribution is Normal and 3DVAR (or 4DVAR) estimates are not longer unbiased. In this case, data assimilation techniques are challenged by the fact that actual applications are typically based on nonlinear processes (Pires et al. 1996). Specifically, that the states exhibited by real systems under observation will diverge from those predicted by a model simulation is clear and this is principally owing to two causes (Palmer & Hagedorn 2006): observational error and sensitivity to initial conditions. The first of these is a result of

what may be called noise. Since its statistical character may not be known, we may need to make some assumptions about its properties. The second source of error occurs on many complex systems and is referred to as chaotic behavior. This has been known for some time, but only in recent decades serious progress in its understanding has been possible. Sensitivity of the model to initial conditions limits how far into the future predictions can be made (Lorenz 1993). Despite challenges and some open questions, 3DVAR is widely used in the oceanographic and meteorological communities, and would make a good candidate method to be explored in the context of astrophysical flows.

2.3 Stratified Convection Model

We are motivated to use data assimilation techniques in the context of stratified convection as a path to obtain predictions of solar subsurface weather events, i.e. the flow structure beneath the surface. Anticipating the possibility of violent events on the solar surface such as coronal mass ejections that affect the space weather and the dynamics of the Earth's magnetosphere is important, (see Ilonidis et al. 2011). The idea is to use a model of solar subsurface convection, ultimately involving the magnetic fields that give rise to surface activity such as coronal mass ejections, although current attempts in that direction is still at a preliminary stage (Warnecke et al. 2011). However, once such models are able to reproduce sufficient details of solar activity, it will be important to synchronize the model with daily observations to be able to use it for predictions.

As a proof of concept, we design a data assimilation experiment to test the implementation of 3DVAR for the Pencil Code, a public domain code of high-order (sixth order

in space and third order in time) for solving the hydrodynamic equations (Brandenburg & Dobler 2002)[2]. We consider here a simple two-dimensional convection model representing the turbulent flows of stars with outer convection zones. In our experiment, *synthetic observations* are generated by adding noise to the output from our model. These observations are then processed by 3DVAR to produce an *analysis*. An analysis is an estimation of the *unknown* state of a system in terms of model variables (Lorenc 1986; Talagrand 1997).

Our implementation is general and can in the future be used for other problems that can be addressed with the Pencil Code. In this work we assume the model to be *ideal* and reproduce the same features present in the observations. In real world applications, the models are far from ideal, and imperfections and uncertainties related to the model are always present. Ideally, we would like to be able to account for some portion of those unknowns by using data assimilation techniques.

We use the sample `2d/conv-slab-MLT` of the Pencil Code (revision r14696 and later). This sample simulates a vertical two-dimensional slice of the outer part of a stellar convection zone. In particular, we use it to simulate convection at low resolution, 64×64, at a Rayleigh number of 8×10^5 (Dintrans et al. 2005), and a Reynolds number of approximately 30. The basic setup is similar to that described in (Brandenburg et al. 2005) and other models before them (Hurlburt et al. 1986; Brandenburg et al. 1996), consisting of a convectively unstable layer sandwiched between two stable layers.

The simulated vertical two-dimensional slice of the solar convection zone has a mean field velocity of $u_{\mathrm{rms}} = 0.08$, the wavenumber of the energy-carrying

[2] `http://pencil-code.googlecode.com/`

eddies is $k_{\rm f} = 2\pi d$ for a depth of the unstable layer $d = 1$, therefore the correlation time $\tau_{\rm cor} = (u_{\rm rms} k_{\rm f})^{-1}$ is approximately 2. Starting from an initial velocity field of perturbations with amplitude of 3×10^{-4} times the average sound speed, convective motion is generated without having to introduce any stochastic elements. This model is chosen to illustrate 3DVAR in an astrophysical context for its sufficiently complex behavior without having any stochastic elements.

2.4 Data assimilation setup

The 3DVAR scheme was developed in the meteorological community to improve model-based weather prediction in spite of observational and modeling uncertainties. It was formulated in a unified Bayesian framework by (Lorenc 1986). 3DVAR produces updates of the current state of the system at times when observations are available, which in turn can be used as a new initial condition to be propagated forward to the time when the next observation is available.

We can use 3DVAR as a black box along with a low resolution simulation to assimilate many data points at low computational cost on a laptop computer. For example, a typical model run for a 64×64 two-dimensional convection field over a time interval of 300 time units takes about 15 minutes on a laptop computer.

3DVAR minimizes the sum of the squared differences between both the model background state and the observations to find a solution that is a compromise between these two estimates of the true state.

It is important to realize that in real problems, the true state is available only through

noisy observations of the system. Therefore, it is impossible to tell how close our model output is to the current and future true states of the system. In our *twin-experiment* (Bengtsson et al. 1981) we select two different initial fields to run the Pencil Code simulation. One of these initial fields represents the unknown true initial state of the system. The other initial field represents what might be, in practice, a good approximation or guess, of the initial state of the system.

As mentioned before, we are set at the ideal case where there is no model uncertainty and the only source of uncertainty is in the observations. In this way, we can assess how far/close the model state is to the true state of the system. The key is to generate a known true state against which the estimated state obtained via data assimilation can be verified.

The experiment is setup as follows: the initial field chosen to represent the true initial state of the system initializes a model run considered to be the original state of the system to be used as reference trajectory or *control*. The other initial field is used to initialize two different runs: one a *free* model run and another that will become the assimilated trajectory or *analysis*. The analysis is a collection of segments of model trajectories initialized at the 3DVAR corrections made at all times where the observations are available. The model (analysis) states at the time when an observation is available are known as the *background states*.

Comparing the free and the control run gives us a measure of the sensitivity to initial conditions of the model, i.e. it shows how similar initial conditions diverge in time. Similarly, comparing the free run and the analysis represents the effect of the data

assimilation procedure over a trajectory starting at the same initial condition field. If the assimilation of the second set of initial conditions is effective it will bring the analysis "closer" to the control run, and further away from the free run.

We generate synthetic observations by adding independent and identically distributed noise to the horizontal velocity field in all grid-points to the control run. These synthetic data are considered to be our experimental observations which in turn will be used to update the analysis. As is explained in the next section, 3DVAR requires both a model and observation state to update the analysis at each given time in the assimilation window.

2.5 3DVAR and the weight factor

The 3DVAR technique finds a model state $\boldsymbol{X}$ that agrees with the current state of the system and the information available in the observations and the model. Specifically, we minimize the weighted average of the residues for both observations $\boldsymbol{Y}_0$ and model states $\boldsymbol{X}_b$ at time t to find an optimal solution. This is expressed in the 3DVAR cost function (Lorenc 1986)

$$\begin{aligned} J(\boldsymbol{X}) =& \frac{1}{2}\left[\boldsymbol{X}-\boldsymbol{X}_b\right]^T \mathbf{B}^{-1}\left[\boldsymbol{X}-\boldsymbol{X}_b\right] \\ &+\frac{1}{2}\left[\boldsymbol{Y}_0-\mathbf{H}(\boldsymbol{X})\right]^T \mathbf{R}^{-1}\left[\boldsymbol{Y}_0-\mathbf{H}(\boldsymbol{X})\right], \end{aligned} \tag{2.1}$$

where $\boldsymbol{X}_b$ is the model state–traditionally called the *background state* and $\boldsymbol{Y}_0$ is the observed state. The background covariance matrix is defined as:

$$\mathbf{B} = \langle(\delta \mathbf{u}_x^A - \langle\delta \mathbf{u}_x^A\rangle)(\delta \mathbf{u}_x^A - \langle\delta \mathbf{u}_x^A\rangle)^T\rangle \tag{2.2}$$

with averages denoted by $\langle\cdot\rangle$ and $\delta\mathbf{u}_x^A$ defined as the difference between the background state $\boldsymbol{X}_b$ and the reference state $\boldsymbol{X}_t$:

$$\delta\mathbf{u}_x^A = \boldsymbol{X}_b - \boldsymbol{X}_t. \tag{2.3}$$

The observational covariance matrix is in a similar way defined as:

$$\mathbf{R} = \langle(\delta\mathbf{u}_x^0 - \langle\delta\mathbf{u}_x^0\rangle)(\delta\mathbf{u}_x^0 - \langle\delta\mathbf{u}_x^0\rangle)^T\rangle \tag{2.4}$$

with $\delta\mathbf{u}_x^0$ defined as the difference between the observations $\boldsymbol{Y}_0$ and the corresponding reference state $\mathbf{H}(\boldsymbol{X}_t)$:

$$\delta\mathbf{u}_x^0 = \boldsymbol{Y}_0 - \mathbf{H}(\boldsymbol{X}_t) \tag{2.5}$$

where the observation operator $\mathbf{H}$ has been used to project the model variables onto observables.

The analysis is $\boldsymbol{X}_a$, minimizes $J(\boldsymbol{X})$ and corresponds to the best estimate of the current state of the system. After the estimate is generated, the model is integrated using as initial condition the analysis up to the next time an observation is available. Synthetic observations are denoted by $\boldsymbol{Y}_0$, these observations contain noise of amplitude, σ_{R}, proportional to the maximum amplitude of the full 2D vertical velocity field. For example, a noise level of 1% corresponds to $\sigma_{\mathrm{R}} = 0.01$ for a normalized field or $\sigma_{\mathrm{R}} = 6 \times 10^{-3}$ for an unnormalized field.

The selection and construction of the observational and background covariance matrices ($\mathbf{R}$ and $\mathbf{B}$) is of great interest in data assimilation (Bannister 2008a,b). In our

case, the observational noise is not correlated in space, and we neglect spatial correlations between model states, making the off-diagonal components of the $\mathbf{B}$ and $\mathbf{R}$ matrices vanish. By doing this, we can set these matrices to be scalars, $\mathsf{R}_{ij} = \delta_{ij}\sigma_{\mathrm{R}}^2$ and $\mathsf{B}_{ij} = \delta_{ij}\sigma_{\mathrm{B}}^2$. Without these spatial correlations 3DVAR generates an analysis that is, in general, less smooth over the two-dimensional domain. In more sophisticated formulations of equation (2.1), the form of $\mathbf{B}$ can also include physical constraints to processes not resolved in the model (Dobricic & Pinardi 2008).

In turn, we set the observation operator to be $\mathsf{H}_{ij} = \delta_{ij}$. This means that we assume observations cover all grid-points in the model domain, i.e. the observables and model variables belong to the same space. In other words, system and model are the same. In more realistic implementations of 3DVAR, $\mathbf{H}$ is typically a computer algorithm that cannot be expressed explicitly as a matrix due to its nonlinear nature (Dobricic & Pinardi 2008). For example the transformation between observables and variable might require the modeling of dynamical processes or making averages.

After these assumptions we can translate the cost function (2.1) to:

$$J(\boldsymbol{X}) = w\,(\boldsymbol{X} - \boldsymbol{X}_b)^2 + (\boldsymbol{X} - \boldsymbol{Y}_0)^2\,, \tag{2.6}$$

where w is the ratio of the scalar variances corresponding to the observed and background states

$$w = (\sigma_{\mathrm{R}}/\sigma_{\mathrm{B}})^2. \tag{2.7}$$

In this setting, to find the state vector $\boldsymbol{X}_a$ that minimizes (2.6), we use Powell mini-

mization (Press et al. 1992). The coefficient w in equation (2.6) behaves as a weight factor in the optimization process and will be referred to as such. Solving $\boldsymbol{\nabla} J(\boldsymbol{X}) = \mathbf{0}$ yields

$$w(\boldsymbol{X} - \boldsymbol{X}_b) + (\boldsymbol{X} - \boldsymbol{Y}_0) = \mathbf{0}, \tag{2.8}$$

and the optimal state of the model that represents the systems given the current observations, or analysis $\boldsymbol{X}_a \equiv \boldsymbol{X}$, is given by:

$$\boldsymbol{X}_a = \frac{1}{w+1}\boldsymbol{Y}_0 + \frac{w}{w+1}\boldsymbol{X}_b. \tag{2.9}$$

Equation (2.9) unveils how the contribution of the background states $\boldsymbol{X}_b$ and observations $\boldsymbol{Y}_0$ affects the analysis $\boldsymbol{X}_a$ in terms of the weight factor w. Figure 2.1 clearly shows the result of this process at times where observations (grey $\diamond$) are available, 3DVAR performs a correction given by equation (2.6) to a value referred to as analysis (black $+$), and from which a segment of background states (dotted curve) is initialized and run up to the next assimilation time. The analysis is the union of the background states for times different from the assimilation times, and the corrected values obtained at assimilation times.

Given a model and a fixed set of observations, equation (2.9) help us understand the effects of the weight factor in the resulting analysis $\boldsymbol{X}_a$, as it is detailed in the following paragraphs. For $w < 1$ or $\sigma_R < \sigma_B$, is interpreted as the case where we assumed that the observational uncertainty is smaller than the model uncertainty, weight is given to the observations since small w will allow the distance $(\boldsymbol{X}_a - \boldsymbol{X}_b)^2$ to grow without making large contributions to the cost function. In contrast, having a $w > 1$ will favor model

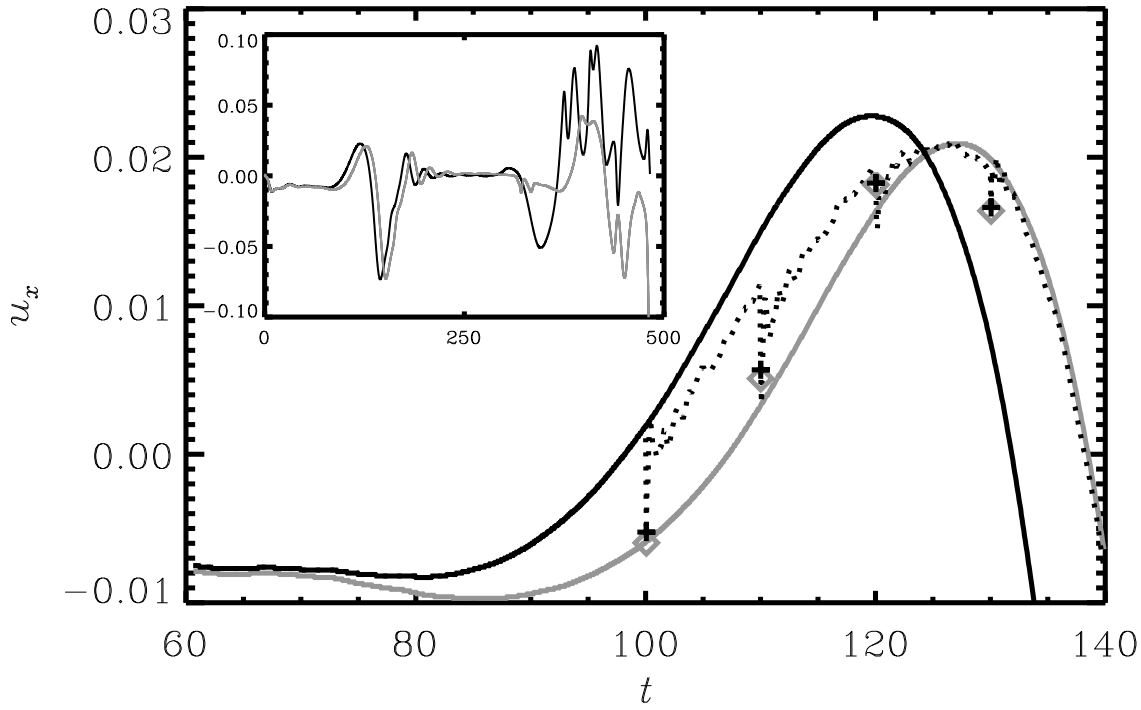

Figure 2.1 Temporal evolution of the horizontal velocity at a certain midpoint of a 2D convection field for $t \in [0, 500]$. Control, free run, and analysis correspond to the grey, black, and dotted curves. A grey '$\diamond$' represent observations, and a '+' corrections ($\boldsymbol{X}_a$) made by 3DVAR, both at assimilation time.

states reflecting our assumption that there is more uncertainty related to the observations than that of the model, or $\sigma_R > \sigma_B$.

The next section presents and describes the results of our numerical experiments. We measure the "quality" of the obtained analyses using 3DVAR by varying values of the weight factor w as well as a couple of sets of observations with different noise levels.

The correlation time was found to be approximately two from simulation parameters (see Section 2.3). This determines the relevant time-scales of the convection features of our simulation and then we choose the most appropriate assimilation time. Consequently, our

choice of assimilation time is large enough to let the oscillations propagate over the two-dimensional field but still small enough to be able to capture the smaller scale dynamics. In each 3DVAR experiment, data assimilation corrections are made each 10 time units (sound-travel-time units) and found no fundamental difference when using assimilation times between 5 and 20 time units.

2.6 Results

We generate analyses using 3DVAR for each of the weight factors $w = \{0, 0.1, 0.5, 1, 10\}$. For a fixed value of w we generate an analysis which is the result of assimilating one set of observations with either 1% or 2% noise level for the same set of two initial field conditions. The resulting horizontal velocity, u_x, at the midpoint of the upper right quadrant of the two-dimensional domain, is plotted in Figure 2.2 for $w = 0$ in the upper panel and $w = 0.5$ in the lower panel, and in Figure 2.3 for $w = 1$ in the upper panel and $w = 10$ in the lower panel.

Note that grey and dashed lines are the same in all panels since they represent the reference states of the system (control) and the corresponding free run of the model. Observations are plotted with grey '$\diamond$' symbols. Black '+' marks are used for the corrections calculated at assimilation time by minimizing expression (2.6). Between assimilations, the analysis is composed by segments of model trajectories (dotted segments) initialized at the corrected state '+' as seen in detail at all insets in Figures 2.2 and 2.3. From equation (2.9) and these insets, we clearly illustrate the amplitude of the correction made in each case for each value of the weight factor w. This amplitude is measured by the gap

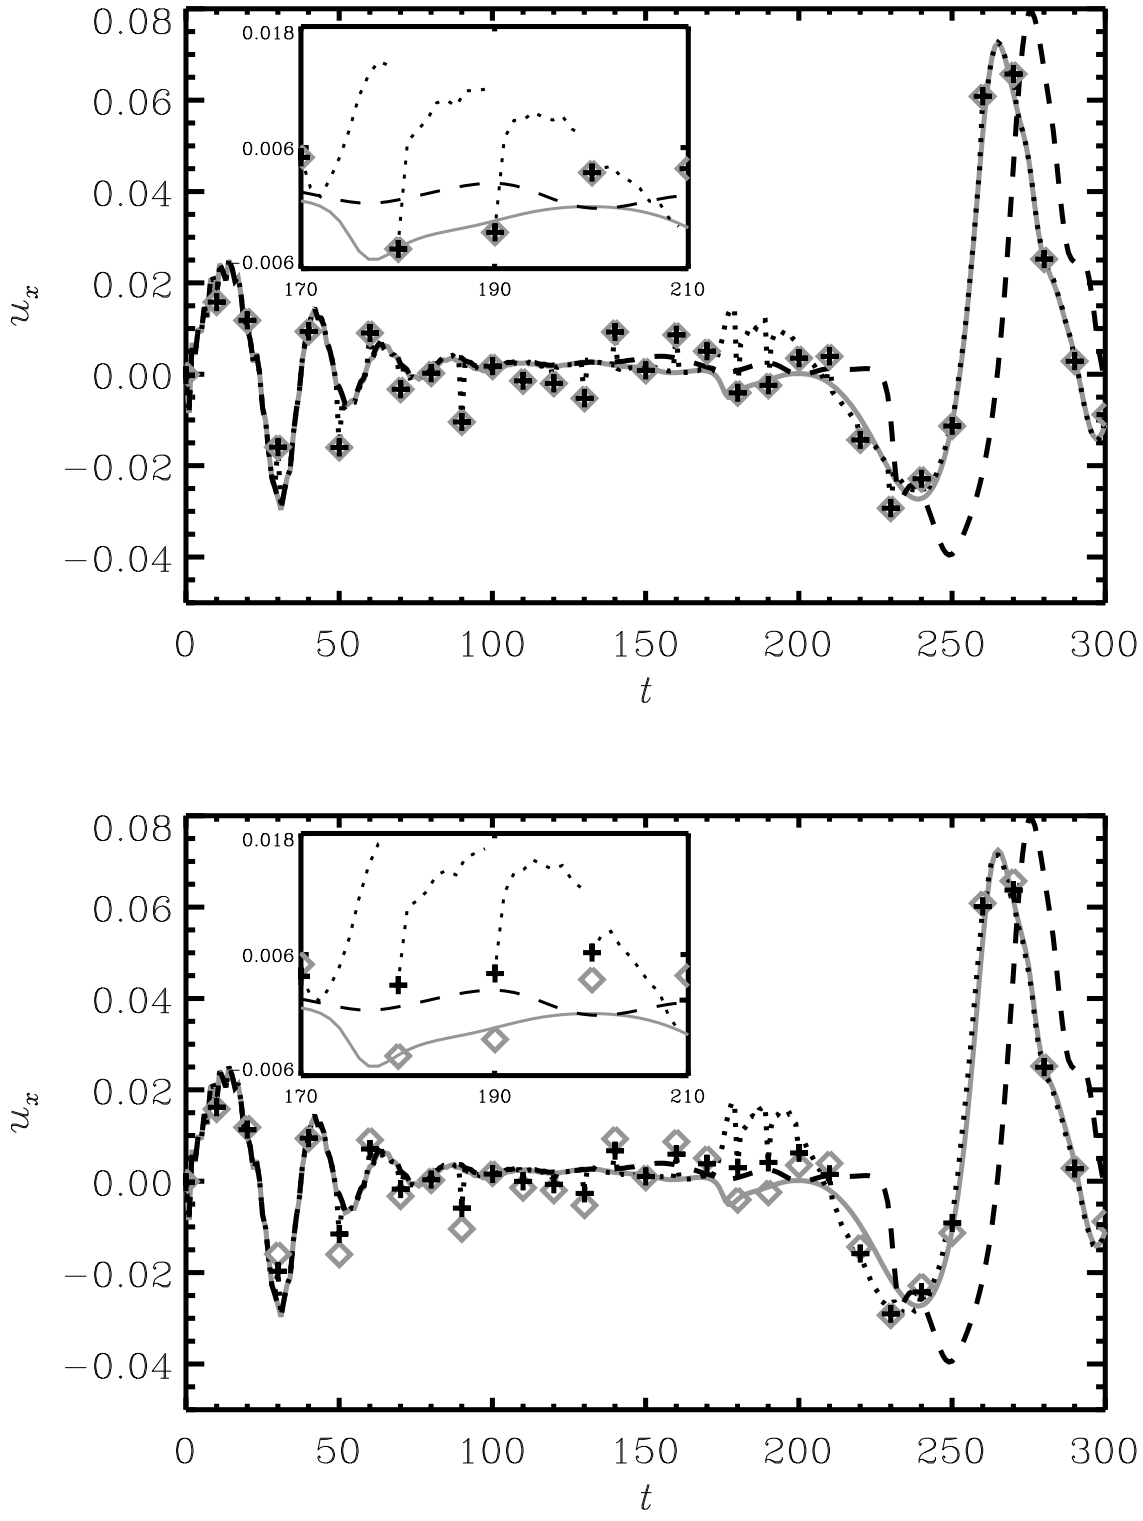

Figure 2.2 Data assimilation run over observations marked with grey '$\diamond$' with 1% noise for $w = 0$ (upper panel) and $w = 0.5$ (lower panel). Black '+' marked the 3DVAR correction at assimilation time. Grey, dotted, and dashed curves correspond to the control, analysis, and free trajectories, respectively. Both panels include zoom-ins for $170 \leq t \leq 210$.

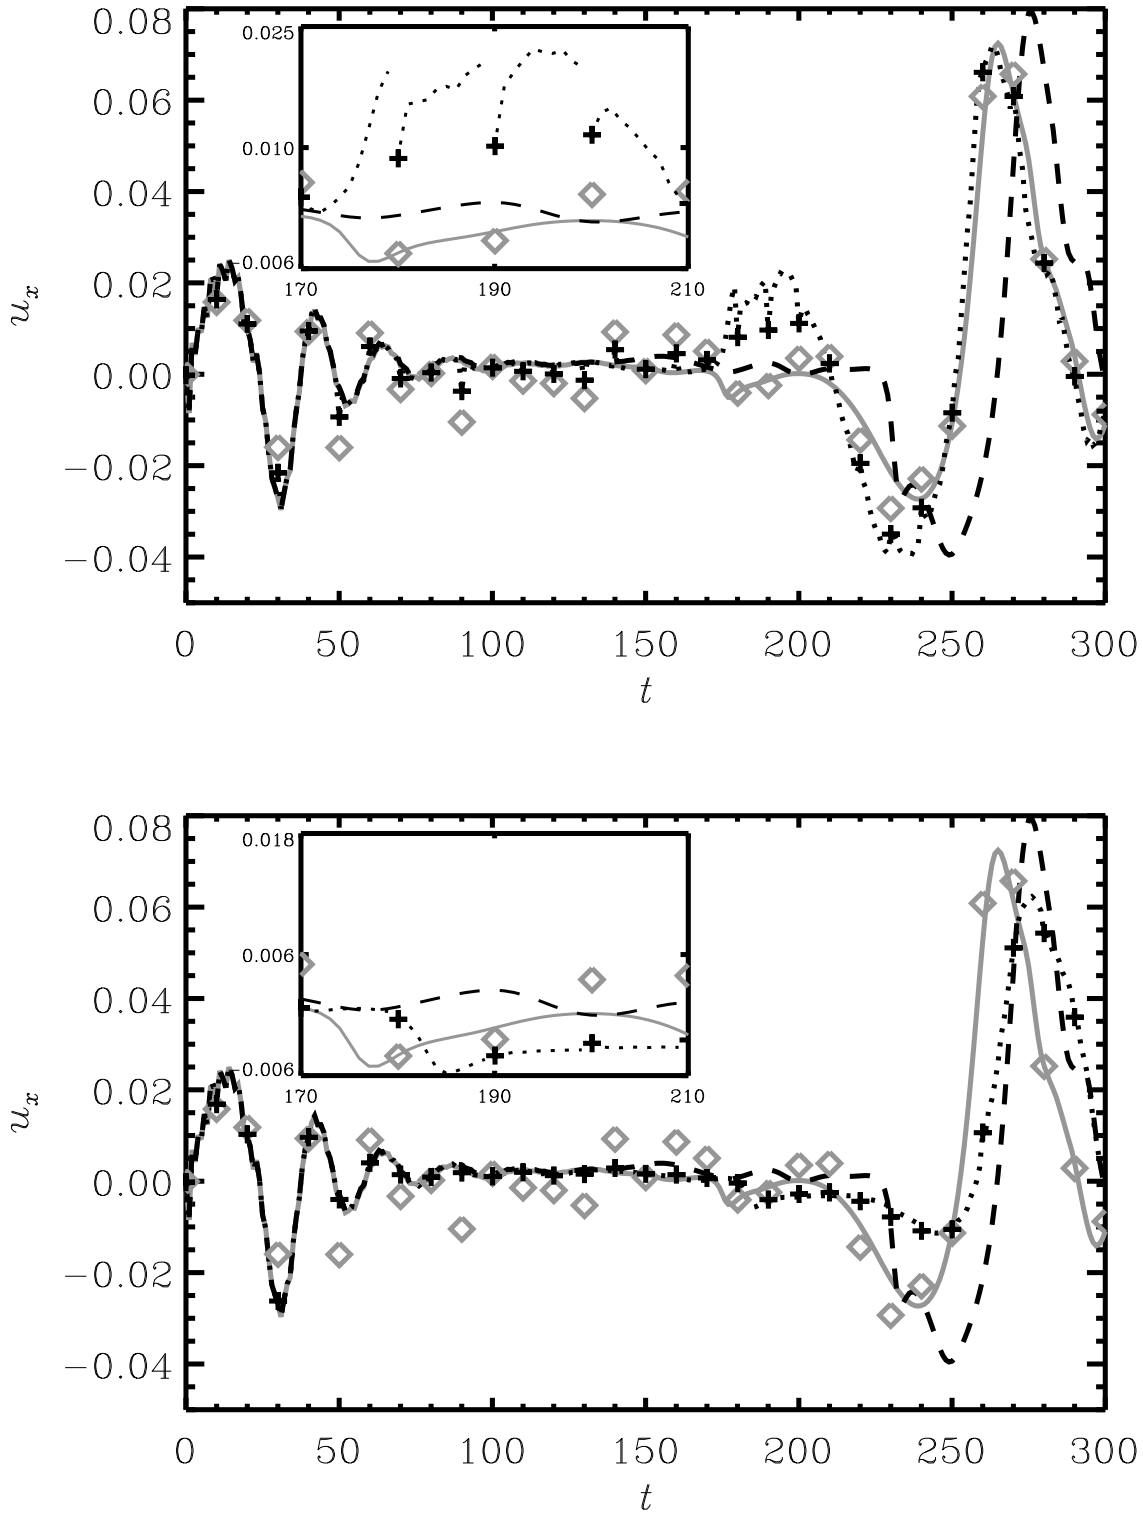

Figure 2.3 Data assimilation run over observations marked with grey '$\diamond$' with 1% noise for $w = 1$ (upper panel) and $w = 10$ (lower panel). Black '+' marked the 3DVAR correction at assimilation time. Grey, dotted, and dashed curves correspond to control, analysis, and free trajectories, respectively. Both panels include zoom-ins for $170 \leq t \leq 210$.

between the '+' and the last dot of the previous dotted segment (background states).

For the trivial case of setting $w = 0$, the first term in (2.6) is neglected and the best estimate of the current state of the system is given by $\boldsymbol{X}_a = \boldsymbol{Y}_0$ from in (2.9), and as seen in the upper panel of Figure 2.2. The correction ('+') is "pulled" from the background state (dotted curve) to the observation ('$\diamond$') at assimilation time. The background states are just segments of transient trajectories initialized at the observations.

For any other value of w the starting point of the correction is in between the observation ('$\diamond$') and the end of the previous background states segment (dotted lines); see insets in Figures 2.2 and 2.3. For $0 < w < 1$, the optimal value of the cost-function (2.6) is by a factor $1/w$ closer to the observations, $\boldsymbol{Y}_o$, than to the background state $\boldsymbol{X}_b$; see Equation (2.9). The lower panel of Figures 2.2 shows results for $w = 0.5$ where the corrections ('+') fall closer to the observations than to the end of the last segment of background states. This figure shows that the analysis follows the control trajectory closely (solid grey line). Note that even if the corrections are large, e.g. at $t = 50$ or $t = 90$, the analysis quickly relaxes to the control state. In the case $w = 1$, equal weights are given to model states and observations. The optimal value of $\boldsymbol{X}_a$ is the average of $\boldsymbol{Y}_0$ and $\boldsymbol{X}_b$, from (2.9). No preference is given to any estimate and the midpoint is the optimal choice for $\boldsymbol{X}_a$ as seen in the upper panel of Figure 2.3.

When $w > 1$, the optimization of equation (2.6) will favor model states rather than observations as follows from equation (2.9). The analysis at assimilation times is by a factor w closer to the model states than to the observations. In the lower panel of Figure 2.3 we plot the resulting trajectories for this case.

For $w = 10$, we observe from this plot and from results at other locations of the two-dimensional domain that for $w > 1$ the estimates of the original state of the system are biased towards the background states.

When the model is linear and ideal, the analysis produced by 3DVAR will be an unbiased estimate of the true state of the system. There, information about the system is only contained in the experimental observations. When the model is neither ideal nor linear information about the known unknown processes important to represent the system of interest, but not part of the model, could only be included via the w. In more general terms the covariance matrices $\mathbf{R}$ and $\mathbf{B}$ informs the 3DVAR procedure about known model and observational uncertainties.

In this simplified experiment, we can see the key importance of the weight factor w in 3DVAR. It points out how crucial the construction of the covariance matrices $\mathbf{B}$ and $\mathbf{R}$ are for optimal results when using variational approaches of data assimilation. One of the motivations for our choice of $\mathbf{R}$ and $\mathbf{B}$ to be scalars is to set a baseline from which we can illustrate, in a simplified way, the inner workings of 3DVAR. It can be hard to see how the different components interact to create a result when more sophisticated choices of $\mathbf{R}$ and $\mathbf{B}$ are used.

We would like to note that the zoom-ins are chosen to be in the interval $t \in [170, 210]$, in all panels of Figures 2.2 and 2.3, as an example of an interval where 3DVAR does not perform very well, and systematically pulls the analysis away from the control trajectory – considered here as the original system trajectory. The reason for this behavior calls for further studies but it is worth noticing that around $t = 200$ the performance of the

assimilation returns to its previous level. The case $w \geq 10$, presented in the lower panel of figure 3, actually performs better during this interval.

Consistently, we observe that the analysis is on average closer to the control trajectory than the observations for all values considered for $w \leq 1$. As noted, exceptions are observed for larger values of the weight factor and during the interval shown in the insets of Figures 2 and 3. Tables 2.1 and 2.2 present several measures of variation of the output from the simulations of our twin experiment. For each simulation, we calculate the variance of the distances between the control and the observations (Table 2.1) or the control and the analysis (Table 2.2) over the data assimilation window.

Specifically, Table 2.1 shows the variance of the distance between the control trajectory and the free trajectory (first row), and the noisy observations for 1% (second row) and 2% (third row) noise levels at the midpoint of the field (second column) denoted by $\langle\langle(\delta u_x^0)^2\rangle\rangle$, this column measures the variability of the local behavior. In addition, the averaged variance over the whole vertical two-dimensional field for the first and the second half of the assimilation window (third and forth column), denoted by $\langle\langle(\delta^O \mathbf{u}_x)^2\rangle\rangle_T$ where the subscript $T = 1, 2$ refers to averages over the first or second half of the assimilation window, $t \in [1, 150]$ or $t \in [151, 300]$ respectively. Note that from the values in Table 2.1, the free run is one or two orders of magnitude further away from the original state of the system (control) than the 1% and 2% noisy observations. Large difference between the second and third columns reflect how free run is diverging from the control run over two different time intervals.

Table 2.1 is the baseline from which we measure the performance of 3DVAR when

Noise Level	$\langle(\delta u_x^O)^2\rangle$	$\langle\langle(\delta \mathbf{u}_x^O)^2\rangle\rangle_1$	$\langle\langle(\delta \mathbf{u}_x^O)^2\rangle\rangle_2$
free	1100	530	3000
1%	27	40	36
2%	110	140	140

Table 2.1 Measures of variability (in 10^{-6}) for the distance between the control and observations with noise levels 1%, 2%, and the free trajectory. See text for notation description.

estimating the system state from noisy observations or the free trajectory–or 0% noise level. Note that the variance of the noise is of the order of 10^{-6}. When assessing the performance of 3DVAR for different values of w and noise levels, we look for variability measures between the control and the analysis lower than the levels set by the free run and the noisy observations. All variance measures for the free trajectory are larger than the corresponding variances values for all w and noise levels. This means that performing 3DVAR data assimilation is more effective at estimating the original state of the system than just using as estimate the trajectory initialized with a very close initial field.

The variance of the distance between the two initial velocity fields is $\approx 2 \times 10^{-7}$. Table 2.2 shows the average variance of the distance between control run and the analysis at the midpoint of the field (second column) denoted by $\langle\langle(\delta u_x^A)^2\rangle\rangle$, and the averaged variance over the whole vertical two-dimensional field for the first and the second half of the assimilation window (third and forth column), denoted by $\langle\langle(\delta^A \mathbf{u}_x)^2\rangle\rangle_T$, for $T = 1, 2$, in analogy of the notation used in Table 2.1.

From the values presented in both tables, we observe that for both noise levels studied and all cases with $w < 10$, the variances of the distance between the control and the analysis are smaller–on average–than the corresponding observations with the same noise level. This shows that on both local and global scales, over the two-dimensional domain,

w	$\langle(\delta u_x^A)^2\rangle$	$\langle\langle(\delta \mathbf{u}_x^A)^2\rangle\rangle_1$	$\langle\langle(\delta \mathbf{u}_x^A)^2\rangle\rangle_2$
Noise level 1%			
0	0.61	0.34	3.80
0.1	0.67	0.35	4.40
1	3.30	4.10	*130.00*
10	*300.00*	*500.00*	*1600.00*
Noise level 2%			
0	2.40	1.30	19.00
0.1	2.70	1.40	22.00
1	33.00	9.80	*840.00*
10	*950.00*	*430.00*	*2800.00*

Table 2.2 Measures of variability (in 10^{-6}) for the distance between the control and the analysis for several values of the w of the x–component of the velocity, $\mathbf{u}_x$. The italic values are larger then the corresponding noise level. See text for notation description.

3DVAR is effectively reducing noise and accounting for sensitivity to initial conditions i.e. it is estimating a value for the horizontal velocity closer to the original state of the system (control) than to the trajectory generated using a guess of the initial state of the system (free run).

In contrast, for $w = 10$ and both noise levels, the free and the analysis are–on average–consistently farther away from the control than the original noisy observations (corresponding values in Table 2.1) both locally and globally. This is also observed, for example at the specific location of the two-dimensional domain shown in Figures 2.2 and 2.3.

More sophisticated choices for $\mathbf{R}$ and $\mathbf{B}$, which inform more realistically the cost-function (2.6) about spatial correlations over the field, might generate an improved and consistent performance. When comparing in this way the variance of the distance between the control and the free run we write $\langle\langle(\delta^F \mathbf{u}_x)^2\rangle\rangle_T$, for $T = 1, 2$. Furthermore, in Figure 2.4 we plot in a semi-logarithmic scale the variance of the point-by-point distances of the

two-dimensional vertical slices between the control and analysis (grey '•'), the control and the observations (grey '⋄'), and the control and the free run (black '•') for all $t \in [1, 300]$. The black '+' mark the variance of distances between the control and analysis fields at assimilation times, i.e. when the corrections from (2.9) are made.

From top to bottom, plots correspond to $w = 0.1, 1, 10$. Black dotted and grey '⋄′ curves are the same for all panels in the figure. The noise level is constant corresponding to 1%.

It is important to note, that values plotted in Figure 2.4 are not the running variance of the distance between control trajectories and the other relevant trajectories but the variance of the distances between the 2D fields at each time step.

We observe in Figure 2.4 that for all values of w and for $t \in [0, 10]$ the analysis and free runs are good estimators of the control trajectory. In addition, for all w and $t \in [20, 100]$, the free run and the analysis variances with respect to control, are below the noise level, with the grey curve below the black curve, showing that for 1% noise level 3DVAR produces a better estimate than the free run. Only for $w = 0.1$ is this the case all values of $t \in [30, 300]$. Note the increase in the error of the estimate from $t \in [160, 190]$. This is part of the interval shown in the insets for Figures 2.2 and 2.3).

Otherwise, the error of the estimate exceeds the noise-level for $t \in [170, 300]$ for $w = 1$, and for $t \in [190, 230]$ for $w = 10$. Figure 2.4 also shows how far apart–on average– the distance between the control and the free run grows as time increases. The increase in the amplitude of the variance (the black curve) is up to six or seven orders of magnitude over the assimilation window with respect to the same distance at small times. As noted

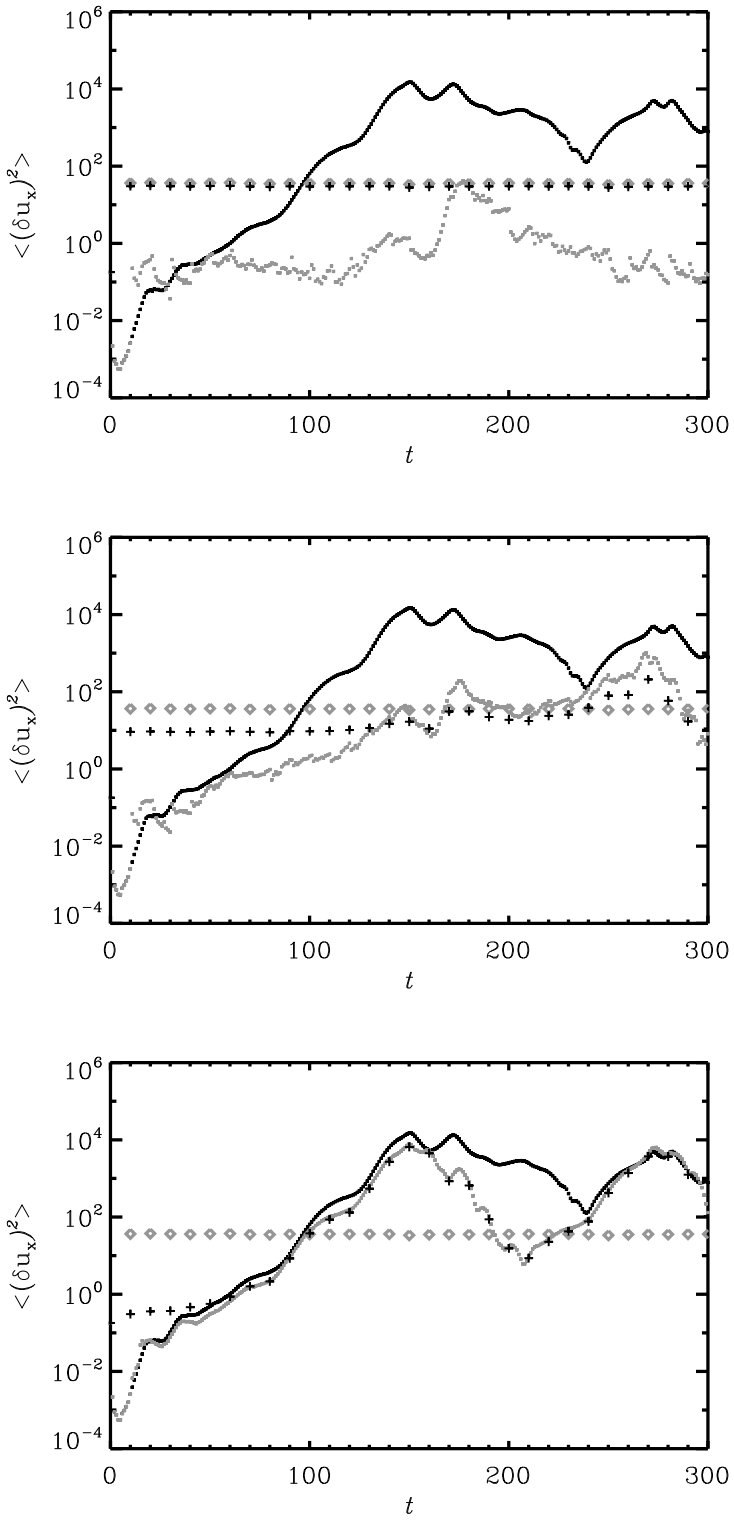

Figure 2.4 Semi-logarithmic plot of the variance of the point-by-point distances of the two-dimensional vertical slices between the control and the analysis in grey '•' ($\langle(\delta^A \mathbf{u}_x)^2\rangle$), the control and the observations in grey '◇' ($\langle(\delta^O \mathbf{u}_x)^2\rangle$), and the control and the free run in black '•' ($\langle(\delta^F \mathbf{u}_x)^2\rangle$) for all $t \in [1, 300]$. The black '+' corresponds to $\langle(\delta^A \mathbf{u}_x)^2\rangle$ at assimilation times, after the correction is made. All quantities are scaled by 10^6.

before, this exhibits the model's sensitivity to initial conditions and it is also a feature observed for all values of the weight factor and noise levels.

In the top panel of the Figure 2.4, we again see evidence of how the 3DVAR correction (black '+') lands closer to the observations and how the background states start to move towards the control trajectory for lower values of the weight factor, w. The effect of $w < 1$ in this case, is to "pull" the background states back closer to the observations as the grey dotted curve is below the corrections ('+').

We recall that the weight factor is $w = (\sigma_{\mathrm{R}}/\sigma_{\mathrm{B}})^2$, and it is interpreted here as a measure of the relative confidence given to either the observations and the model. High values of w reflect higher trust in the observations than in model representation of the system, and the opposite is valid for low values of w. Note that σ_{B} relates to initial condition sensitivities than with model deficiencies in our simple setting. As commented earlier, the estimate is sensitive to the choices of σ_{B} and a more sophisticated choice can include addition components that might help alleviate some of the model deficiencies.

In conclusion from the results of our experiment, in the particular case where the covariance matrices are assumed to be scalar and for small values of the weight factor w, the effect of trusting the observations more than the model states ($\sigma_R < \sigma_B$), provides a closer estimate of the original state of the system than just generating a trajectory close enough in initial conditions. This means that 3DVAR is successful at finding an optimal state estimator in the limit of small observational noise, model uncertainty related only to sensitivity to initial conditions, and scalar covariance matrices, $\mathbf{R}$ and $\mathbf{B}$. Our results strongly point out the importance of choosing more sophisticated covariance matrices ($\mathbf{R}$

and **B**) to better inform the assimilation procedure about the known uncertainty sources in a problem of interest.

2.7 Discussion

We have presented an idealized case where model and system are the same: A computer simulation is used both to generate synthetic observations and as the model required for the data assimilation procedure. In this way, we can assess how far/close the model state estimates are to the true state of the system. The key is to have access to the true states that we can use to verify and evaluate estimates obtained using data assimilation.

We used a simplified formulation of the 3DVAR data assimilation technique in terms of the weight factor: $w = (\sigma_R/\sigma_B)^2$, that defines the contribution of the model states -that contain propagated information from previous observations- and the current observation to make a state estimate. This formulation of 3DVAR used here is achieved by reducing the covariance matrices, **R** and **B**, to scalars; and the observation operator, **H**, to the identity. This selection corresponds to neglecting all spatial correlations between model states over the two-dimensional domain in addition to one-to-one correspondence between system observables and model variables. In this way, we clearly separate the contribution of observations and model states to the estimated state, as seen in equation (2.6) and (2.9).

It is less direct to see how the different components interact to create an estimate of the original state of the system when more sophisticated choices of covariance matrices, which represent uncertainties and spatial correlations. In that case we would have to think

about the optimal combination in analogy to equation (2.9), in terms of a generalization of the weight factor w as a *weight matrix*, $\mathbf{W} = \mathbf{R}\mathbf{B}^{-1}$. In this analogy, model states and observations will be projected by the matrices $\mathbf{W}[\mathbf{W} + \mathbb{I}]^{-1}$ and $[\mathbf{W} + \mathbb{I}]^{-1}$, respectively, to contribute to the analysis $\boldsymbol{X}_a$. Here, $\mathbb{I}$ is the identity matrix. In general, we can say that to understand the 3DVAR algorithm it is important to look at the weight factor, particularly in the limit where $\mathbf{W}$ is assumed to be the scalar w. Consistently we observe that the error between the state estimate and the original state of the system is below the noise level when more weight is given to the observations then to the model state.

When the contribution from the model state is larger than the contribution from the observations we note that the error eventually becomes larger than the noise level, see case $w = 10$. We note in Figures 2.2, 2.3 and 2.4 that 3DVAR under-performs for $t \in [160, 230]$ both locally and globally for all values of w. Further study of the simulation is needed to account for this atypical behavior. Minute differences in initial conditions generate different time evolution for the different runs, as is expected for nonlinear systems. This is illustrated by the black curves in Figure 2.4 that present how the variance of the distance between the two initial conditions grows over the time interval. It can also be seen in Figure 2.1 where the grey and black curves, which started with close initial condition, are very different at latter times.

On the other hand a large correction made by 3DVAR, for example the black '+' at time 50 and 90 in Figure 2.2, does not appear to have a strong effect. The model run that starts at these far away estimates converges almost instantly back to the control at the same time.

These features, that are reminiscent of chaotic behavior, can be understood in terms of attracting sets where small changes in initial condition generate a different time evolution on the attracting set. A large correction probably takes us outside the attracting set and the solution rapidly falls back when the model is integrated forward.

Another interpretation is that the large correction takes us to states that are not consistent with conservation laws and other physical constraints. The system would then rapidly be forced back onto a more physical state. This dual picture, using both physical and mathematical intuition somewhat clarifies this contradictory behavior.

The 3DVAR methodology is optimal for linear models and Gaussian distributed uncertainties (Lorenc 1986). Very few models in astrophysics have these properties. The validity of variational methods outside the linear or weakly nonlinear case is unclear but if the assimilation is frequent enough the behavior might be closer to linear. Higher assimilation frequency during the interval $t \in [170, 230]$ might for example give a better result.

The indication of chaotic properties and attracting state space sets invites the use of other data assimilation methods that explicitly take these properties into account (Judd & Smith 2001; Judd et al. 2008). They might be more applicable to non-linear astrophysical processes. In this work we have used an ideal model, that is, the system and the model is one and the same. There will always be some limitation to modeling of real system and this might prove problematic. In general it is, in our mind, more important for models used for prediction of real systems to be reasonably realistic rather than just being close to the observations.

For nonlinear astrophysical systems that operate on timescales from seconds to years, data assimilation will be of fundamental importance when quantitative agreement between model and observations is to be assessed. The ultimate verification that a model is correct is its ability to make reliable prediction and for this, data assimilation is necessary.

Chapter 3

Data assimilation and parameter estimation with unobserved variables.

3.1 Introduction

This chapter is looking at *data assimilation* and *parameter estimation* for nonlinear models in the context of *model-based predictions*. In general there is a *system* of interest, i.e. the object we like to study and a mathematical *model* that is used to describe the temporal evolution of the system and eventually make predictions of its future behavior.

A *model based prediction* requires selection of the *model parameters* to use and a *model state estimate* of the current state of the system written in terms of the model variables. This chapter will focus on methodological aspects of how to select model

parameters and states using the Shadowing Filter. It will present in particular how to deal with *unobserved variables*, i.e. model variables that do not correspond to a specific observable.

This chapter will study an ideal case[1] where the model is defined to be *structurally equivalent* to the system, i.e. it is possible to find a set of model variables such that the model and system are identical. This case applies to artificial setups and numerical experiments where both the model and system correspond to the same equations or numerical implementation. The motivation to study the methodology in the ideal case is to test that the method works, before complications are added due to real world concerns. For example, there might be structural differences between model and system that need to be accounted for. This exercise will provide insight on how to eventually apply, explore and interpret the method in a real world case.

This chapter presents a *twin experiment* where "reality" is something explicitly generated by a known mathematical structure and a given set of parameters. *Synthetic observations* are constructed by adding random perturbations to the generated "true" states (or trajectory). The ability to find the known "true" state of the explored system is tested– under different imposed conditions – starting from the set of synthetic observations. The evaluation is possible because in the twin experiment, the "true" state can be used as a reference value to compare with. Without knowing the "truth", it is impossible to exactly verify how close the estimates are to the unknown true values, and it would be difficult to confirm that the method works. In the real case the strongest indication that an estimate, a method or a model is reliable is its ability to make accurate predictions in

[1]This corresponds to the *perfect model scenario* discussed in literature (Judd & Smith 2001).

an out of sample.

The *Shadowing Filter* (Judd & Smith 2001; Stemler & Judd 2009) is the method used in this chapter to find *model state estimates* that best represent the observed system using a set of synthetic observations. The time evolution of any observed *deterministic system*[2] generates causally connected states. The set of observations typically corresponds to these states of the system with some added noise.

The Shadowing Filter approach uses this insight to impose *dynamical consistency* as a constraint to obtain estimates, as outlined in Section 3.2. Given a sequence of observations and a specific mathematical structure or model, the Shadowing Filter method obtains a set of estimates that is close to a dynamically consistent sequence, i.e. *trajectory* of the model.

The Shadowing Filter has been applied to a range of problems: the discrete Ikeda map and Logistic map (Cuéllar 2007); the dynamical flows of the Lorenz'63 (Stemler & Judd 2009) and Lorenz'96 models (Du 2009); and the *Navy Operational Global Atmospheric Prediction System* NOGAPS (Judd et al. 2004a, 2008). Parameter estimation for low-dimensional systems using the Shadowing Filter was explored by Cuéllar (2007) for maps and Du (2009) for flows (see also Smith et al. 2010; Du & Smith 2012). This chapter address the specific case of parameter estimation for unobserved variables mentioned by Smith et al. (2010) in the case of the low-dimensional Lorenz'63 flow.

[2]A deterministic system unlike a stochastic system does not have random components.

3.2 The Shadowing Filter

Given a sequence of observations $\{S_n\}$ at times $n = 1, 2, 3, ..., N$ and a deterministic model F that takes a model state X_n at time n to an image $F(X_n, \theta)$ at time $n+1$, where θ is a fixed set of model parameters, the goal of the Shadowing Filter is to find a set of states $\{Z_n\}$ that is close to a *causally connected sequence*, i.e. where the *mismatch* or 1-step prediction error $\mu_n = [F(Z_n, \theta) - Z_{n+1}]$ is small for all n. See also section (1.6) for a technical statement of the problem.

One can evaluate how close any set of states is to a model trajectory with the *mismatch functional*:

$$\begin{aligned} \mathcal{M}(\{Z_n\}) =& \frac{1}{N-1} \sum_{n=1}^{N-1} [\mu_n]^2 \\ =& \frac{1}{N-1} \sum_{n=1}^{N-1} [F(Z_n) - Z_{n+1}]^2, \end{aligned} \tag{3.1}$$

where the $\mathcal{M}(\cdot)$ is zero only if the set $\{Z_n\}$ is a trajectory of the model $F(\cdot)$. In this case, Equation 3.1 implies that $\mu_n = 0$ for all n (Judd 2008).

We would like to find a sequence of estimates $\{Z_n\}$ that minimizes $\mathcal{M}$ by solving equation:

$$\frac{\partial \mathcal{M}}{\partial Z_n} = 0. \tag{3.2}$$

The gradient can be explicitly written as:

$$\frac{\partial \mathcal{M}}{\partial Z_n} = \frac{2}{N-1} \times \begin{cases} [F(Z_n) - Z_{n+1}]F'(Z_n) & n = 0, \\ [F(Z_n) - Z_{n+1}]F'(Z_n) - [F(Z_{n-1}) - Z_n] & 0 < n < N, \\ -[F(Z_{n-1}) - Z_n] & n = N, \end{cases} \tag{3.3}$$

where $F'(\cdot)$ is the Jacobian matrix of the model $F(\cdot)$. In a numerical implementation F' is called an *adjoint*. For more on adjoints in relation to the Shadowing Filter see Judd et al. (2004b) and Stemler & Judd (2009).

Equation 3.2 is solved, using expression 3.3, to find the lowest value for $\mathcal{M}$. The imposed condition shows that the minimum of the mismatch has to be zero. This is a global minimum and the optimal sequence of state estimates $\{Z_n\}$ are trajectories of the model (Judd & Smith 2001; Judd et al. 2008). Solving this variational problem is in this case equivalent to finding the minimum using gradient descent:

$$\begin{aligned} \dot{Z}_n &= -\nabla \mathcal{M} \\ &= -\frac{\partial \mathcal{M}}{\partial Z_n} \end{aligned} \tag{3.4}$$

for all n, where $\dot{Z}_n$ is the derivative with respect to some *descent time* τ that should not be confused with the time evolution t of the system and the model. Note that τ is a parameter that allows us to move down a hyper-surface in sequence space. The global minimum of this surface is the target we seek. Numerically we can solve 3.4 using an

Euler approximation:

$$\begin{aligned} \dot{Z}_n &= -\frac{\partial \mathcal{M}}{\partial Z_n} \\ \frac{\Delta Z_n}{\Delta \tau} &= -\frac{\partial \mathcal{M}}{\partial Z_n} \\ Z_n^{j+1} &= Z_n^j - \Delta\tau \frac{\partial \mathcal{M}}{\partial Z_n} \end{aligned} \tag{3.5}$$

for iteration j and a step-size $\Delta\tau$. This provides an iterative procedure that can refine some initial estimate $\{Z_n^0\}$ and take the estimates towards a trajectory of the model. Note that the step-size $\Delta\tau$ has to be adjusted to the specific problem. If $\Delta\tau$ is too large it may not be possible to find the minimum. Stemler & Judd (2009) use the following prescription:

$$\frac{2\Delta\tau}{N-1} = 0.1. \tag{3.6}$$

In our work the value is ~ 0.2 or more specifically $\Delta\tau = 52$ for $N = 512$ observations and is adjusted for convergence in other cases. The value sought for $\Delta\tau$ should be large enough to give rapid convergence, but sufficiently small not to generate a divergent result. Later, as a general prescription, we use:

$$\Delta\tau = \frac{N}{10}. \tag{3.7}$$

An interesting set of estimates with which to initialize the Euler-approximation of the gradient descent is the set of observations $\{Z_n^0\} = \{S_n\}$. The method iteratively finds estimates with lower $\mathcal{M}$, i.e. higher dynamical consistency, starting from the set of observations. Note that this process only converges towards a trajectory (Ridout & Judd 2002; Judd 2008).

In practice the set of estimates $\{Z_n^j\}$, after a finite number of iterations, are only members of a *pseudo-trajectory*, i.e. a sequence that has small but nonzero mismatch $\mathcal{M}$. This eventually provides an approximate reference trajectory of the model that corresponds to the set of observations.

The language used in Chaos theory (Strogatz 2001), can be used to describe this process as projecting the observables onto the *attractor* or the *attracting set* of the model – if such an object exists for the chosen model (Judd et al. 2008).

In the ideal case one cannot do better with a single estimate than approximating a trajectory of the model that matches the set of observations (Ridout & Judd 2002) since a trajectory of the model is also, by the definition of the ideal case, a trajectory of the system. One could use a probabilistic or an ensemble description to describe the state and use this to initialize a prediction. For more on the mathematical properties of the mismatch function and the gradient descent see Ridout & Judd (2002) and Judd (2008).

3.3 The Lorenz'63 model

The Lorenz'63 model (Lorenz 1963) will be used in this chapter. It is one of the most studied and cited dissipative nonlinear systems in literature (Strogatz 2001; Hilborn 1994). It represents a system with a temperature gradient generated between two plates. When the gradient is large enough, the system goes from laminar flow to convective motion in

rolls. The Lorenz'63 system:

$$\begin{aligned}\dot{\mathbb{X}} =& s(\mathbb{Y} - \mathbb{X}) \\ \dot{\mathbb{Y}} =& \mathbb{X}(r - \mathbb{Z}) - \mathbb{Y} \\ \dot{\mathbb{Z}} =& \mathbb{X}\mathbb{Y} - b\mathbb{Z}\end{aligned} \tag{3.8}$$

is an adaptation of the Saltzman (1962) model of Rayleigh (1917) convection and exhibits chaotic behavior for parameter values around the standard s, r, $b = [10, 28, 8/3]$.

The s parameter of the Lorenz'63 model corresponds to the *Prandtl number*[3]. The onset of convection is set by $r = 1$ since:

$$r = \frac{R_a}{R_c}, \tag{3.9}$$

is the ratio of the *Rayleigh number*[4] to the *critical Rayleigh number*. The b parameter is a geometric factor that is related to the wavenumber of the rolls. This system exhibits a quasi-periodic behavior for these and other values of the parameters.

Lorenz found that this system was very sensitive to the initial conditions. Any two initial conditions, however close, would rapidly diverge. We can also study the state space divergence of the Lorenz'63 system:

$$\frac{d\dot{\mathbb{X}}}{d\mathbb{X}} + \frac{d\dot{\mathbb{Y}}}{d\mathbb{Y}} + \frac{d\dot{\mathbb{Z}}}{d\mathbb{Z}} = -(s + b + 1) \tag{3.10}$$

and find that it is strictly negative. This means that any set of initial conditions will

[3]The Prandtl number is the dimensionless ratio of the kinematic viscosity to the thermal diffusivity.
[4]The Rayleigh number is a dimensionless measure of the buoyancy force.

converge rapidly[5] towards a zero volume object. To account for the behavior of a system where all initial conditions are converging but nearby states diverge a new kind of object, *a strange attractor* was named by (Ruelle & Takens 1971).

For the Lorenz'63 system the $\mathbb{X}$ variable is proportional to the intensity of the convective motion, $\mathbb{Y}$ relates to the temperature of the ascending and descending fluid currents and $\mathbb{Z}$ is related to the distortion of the laminar vertical temperature profile (Lorenz 1963).

3.4 The Shadowing Filter in practice

Figure 3.1 shows an example for Lorenz'63 where 512 synthetic observations (gray stars) are generated by adding 25% Gaussian noise to the model output (blue line) that is used as the "true" state of the system. It is generated with the standard parameter values [s, r, b] = [10, 28, 8/3] from Lorenz (1963), and the initial conditions [$\mathbb{X}_0$, $\mathbb{Y}_0$, $\mathbb{Z}_0$] = [10, 10, 10]. Using the iterative procedure outlined above, estimates (red circles) are found after 200 iterations of the gradient descent algorithm with $\Delta\tau = 52$, using a 4th order fixed step Runge-Kutta integrator. The error of the estimate is a factor of five to seven times smaller than the observational noise. The convergence of $\mathcal{M}$ as a function of gradient descent iterations is seen in Figure 3.2. The exact number $\Delta\tau = 52$ is not significant; the importance is in its ability to provide a convergent result.

The case of applying the Shadowing Filter to the Lorenz'63 system has been explored in numerical detail by Stemler & Judd (2009). They outline a number of criteria for convergence within 100 iterations. In Table 3.4 we compare their values to the ones we

[5]For every time unit the flow contract 13.7 e-folds or about a factor of a million Tucker (2002).

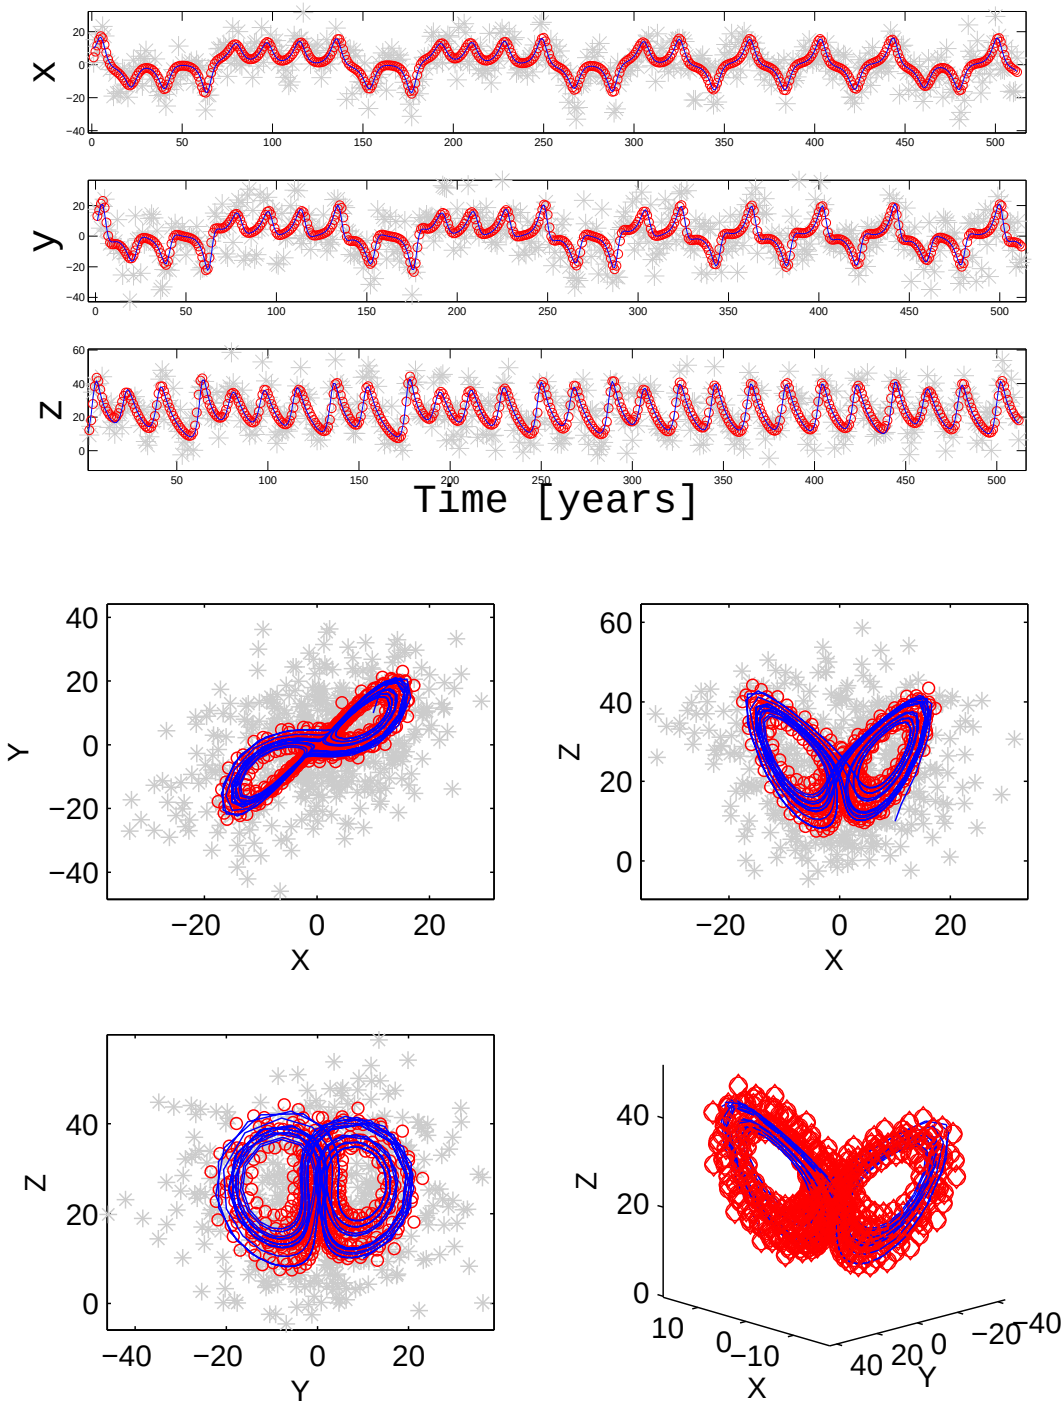

Figure 3.1 From the true model state (blue line) the 512 synthetic observations (gray stars) are generated using the Lorenz'63 system with 25% noise. Following the outline procedure the estimates (red circles) are obtained after 200 gradient descent iterations. In this case the parameters of the model are known. The errors in the final estimate are 3.2%, 3.5% and 5.8% for the three components of the model. The upper set of three figures exhibits the time-series of the components and the lower set of four presents projections of the state space behavior.

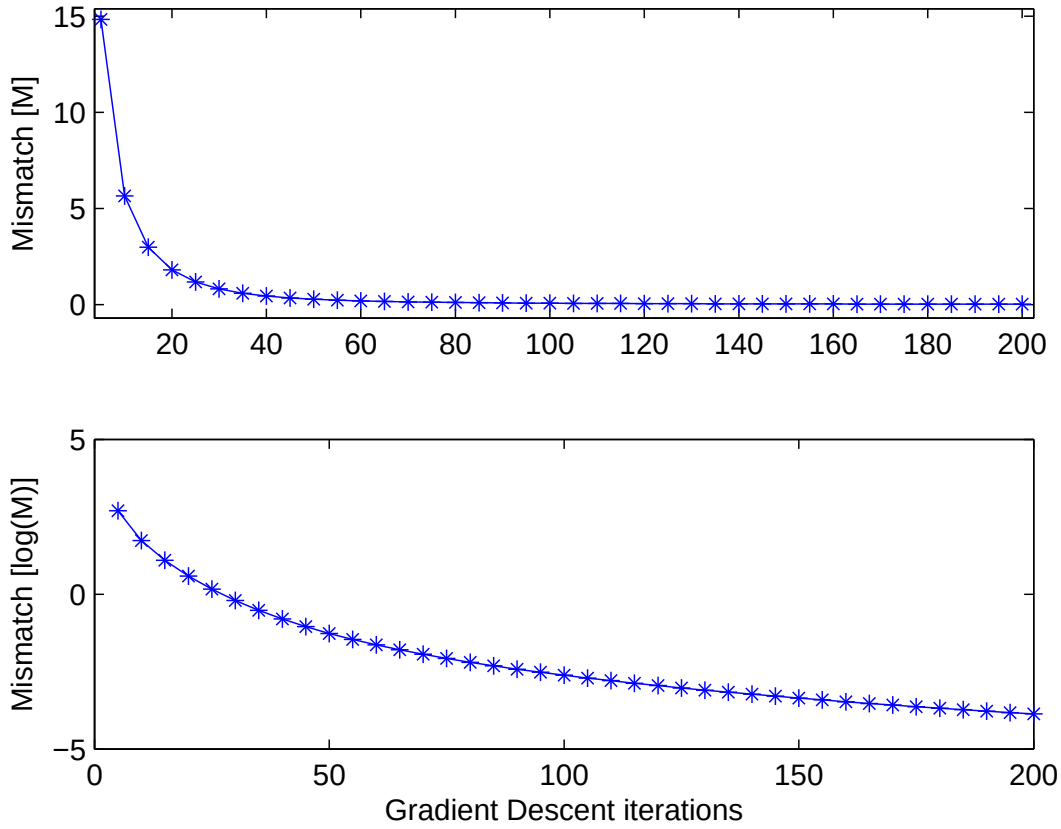

Figure 3.2 The mismatch $\mathcal{M}$ as a function of gradient descent iterations for the Lorenz'63 model (top panel). Lower panel is in logarithmic scale.

have independently found in this study. Their upper bound for Δt and h are higher than our values. Also their lower bound estimate for the number of observations $N > 10$ is in agreement with our value of $N = 512$. The motivation for the large number of observations $N = 512$ used in this work is to include more dynamics of the system to eventually constrain the model parameters.

	Stemler & Judd (2009)	This work
N	> 10	512
Δt	< 0.5	0.15
$h = \Delta t / w$	< 0.05	0.01
w	10	15

Table 3.1 The criteria for convergence within 100 iterations from Stemler & Judd (2009). For the number of observations N, the time interval between observations is Δt (not to be confused with $\Delta\tau$), and the integration step length is h. The number of intermediate integration steps w used in this work is somewhat different, but this should not have any significant effect on the overall outcome. Please note that this w should not be confused with the weight factor used in Chapter 2.

3.5 Parameter estimation

Both the model $F(\cdot)$ and mismatch functional $\mathcal{M}(\cdot)$ are functions of the parameter θ. We evaluate them under identical conditions and compare the result for different values of θ. In this way interesting areas of parameter space, where $\mathcal{M}(\theta)$ is minimal, can be identified,

Three examples of how $\mathcal{M}(\theta)$ changes as a function of different parameters are presented in Figure 3.3. We use the Lorenz'63 system with a noise level of 5% and 80 gradient descent iterations. In this case the lowest $\mathcal{M}$ value, in all figures, also corresponds to the true set of parameters used to generate the system in the first place. In this work we have only shown these examples as illustrative cases. For a thorough exploration of parameters using this and other measures, see Cuéllar (2007); Du (2009); Smith et al. (2010); Du & Smith (2012). The novel contribution of this Chapter – parameter estimation with unobserved variables – is shown in Section 3.7.

Why the $\mathcal{M}$ value would be lower for the "correct" set of parameters is currently not

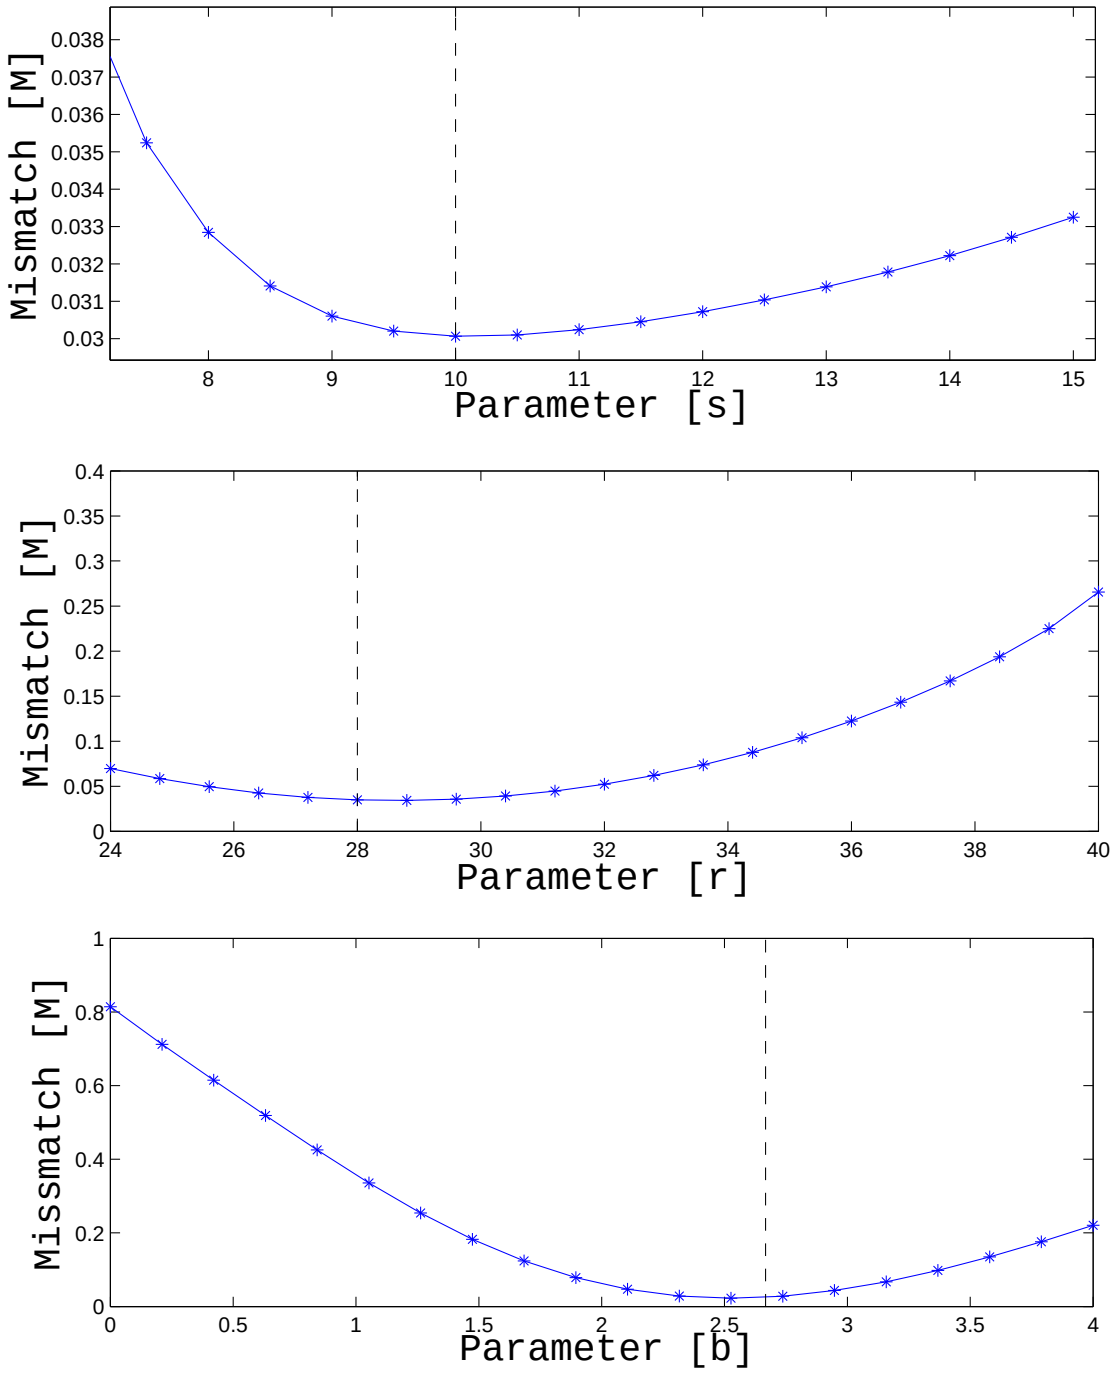

Figure 3.3 Parameter estimation for Lorenz'63 system. Identical state estimates are made and the cost-function $\mathcal{M}(\theta)$ is evaluated for different parameter values θ, after 80 gradient descent iterations at a noise level of 5%. True parameter values (dashed lines) from top to bottom are: $s = 10$, $r = 28$, and $b = 8/3 \approx 2.7$.

fully explained or mathematically proven. One interpretation is that rapid convergence – lower $\mathcal{M}$ after a fixed number of iterations – corresponds to a closer "distance" in sequence space. This would imply that the *projected distance* - to use the language of Judd et al. (2008) - would be closer for "better" sets of parameters. This would make $\mathcal{M}$ an indirect measure of how close in sequence space the dynamics of the system are to the dynamics of the model.

3.6 Unobserved variables

In order to pair model variables and observed quantities, some correspondence is assumed. It is also assumed that some transformation exists between them to enable assimilation and comparison between model estimates and observations. If the number of model variables is larger than the number of observables of the system, then *unobserved variables* occur. In other words, we cannot match model variables with observables.

In such a case the Shadowing Filter – as outlined above – can be applied with slight modifications. One major difference from the previous case is the lack of values for the unobserved variables, and therefore initial guesses are necessary. Our work suggest that the result is not very sensitive to which guesses we choose. We use uniformly distributed random values in typical ranges suitable for the different model components to sample guesses for the unobserved variables.

The case of setting all guesses to zero was also tested, and it does not appear to affect the result significantly.

The observed and the guessed variables make our first estimate: $\{Z_n^0\}$. All the

system information is contained in the observed variables and this information should be propagated to the unobserved variables using the mismatch minimization. The approach mentioned in Smith et al. (2010) and outlined by Du (2009) is to "hold onto" the set of observations when the gradient descent procedure is applied. This procedure has two stages as shown in Figure 3.4. During the first "hold on" stage updates are made only to the unobserved variables whereas the observed variables remain unchanged. It can also be

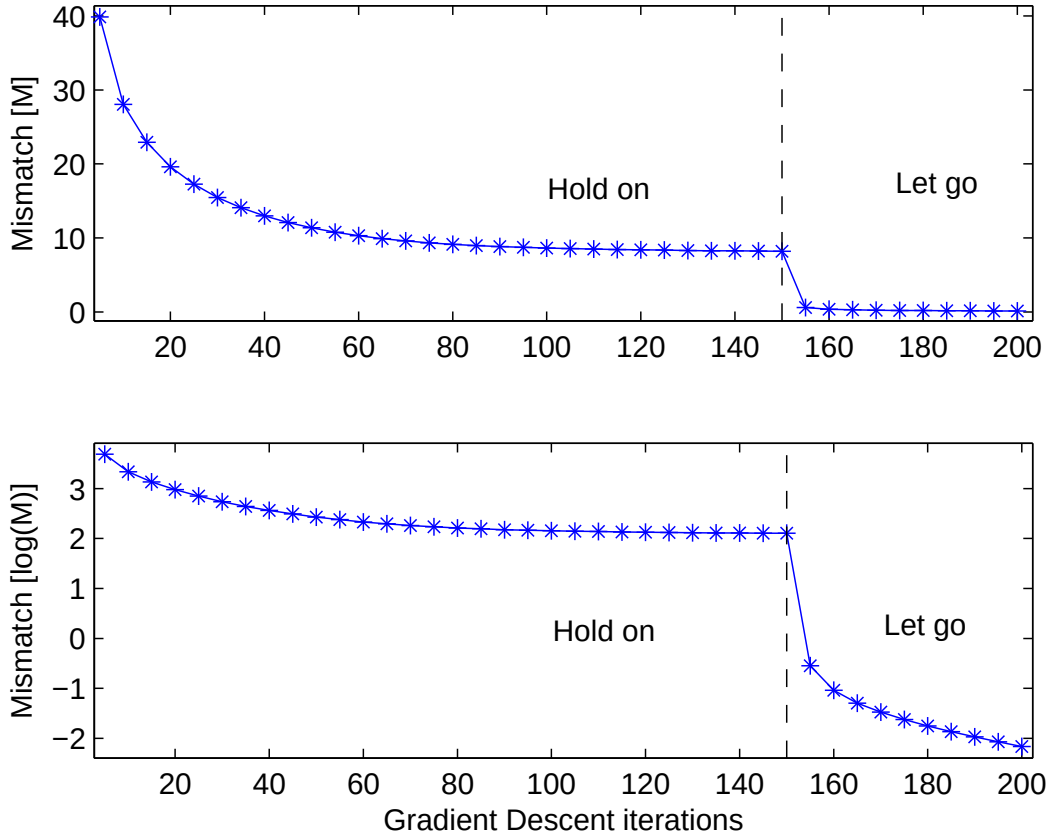

Figure 3.4 The mismatch $\mathcal{M}$ as a function of gradient descent iterations (top panel). Lower panel is in logarithmic scale. Shown are the two stages of the procedure for unobserved variables. First the "hold on" stage where the set of observations are held fixed and then the "let go" stage that operates as the standard procedure, outlined in Section 3.4 and seen in Figure 3.2.

stated that the unobserved variables are updated to lower the $\mathcal{M}$ value, and the observed variables are reset to their original values after each iteration. At the end of the "hold on" stage, the unobserved variables have been taken towards a pseudo-trajectory that is consistent with the set of observations, see Section 3.2.

There will still be some noise in the original set of observations. To make the final estimate one must, during the next stage, "let go" of the data and make a number of gradient descent iterations without resetting the observed variables each time. This second "let go" stage corresponds to the standard gradient descent procedure outlined in Section 3.4 and seen in Figure 3.2. This takes the estimates towards a trajectory of the model that is inferred from the set of observations. At this time, the estimates are not constrained to the set of observations. This is not a problem in the ideal case since a trajectory of the model is by definition the same as a trajectory of the system. However, to "let go" of the data – in a non-ideal or real case – should be done with considerable care since the model and system trajectories are not the same. Leaving the observations in this case does not necessarily take us towards the true state of the system.

Figure 3.5 shows how the Shadowing Filter operates for the Lorenz'63 system with s, r, $b = [10, 28, 8/3]$ when only the $\mathbb{Y}$ component is observed. For the $\mathbb{X}$ and $\mathbb{Z}$ components we generate uniformly distributed guesses (gray stars). The noise level for the $\mathbb{Y}$ component is 5% (also gray stars). The estimates (red circles) obtained after $150 + 50$ iterations of the procedure are close to the generated true state of the system (blue line). The errors in the estimate are 1.7%, 1.9% and 4.5% for the three components of the model.

The result is similar when only the $\mathbb{X}$ variable is observed, see Figure 3.6. The errors

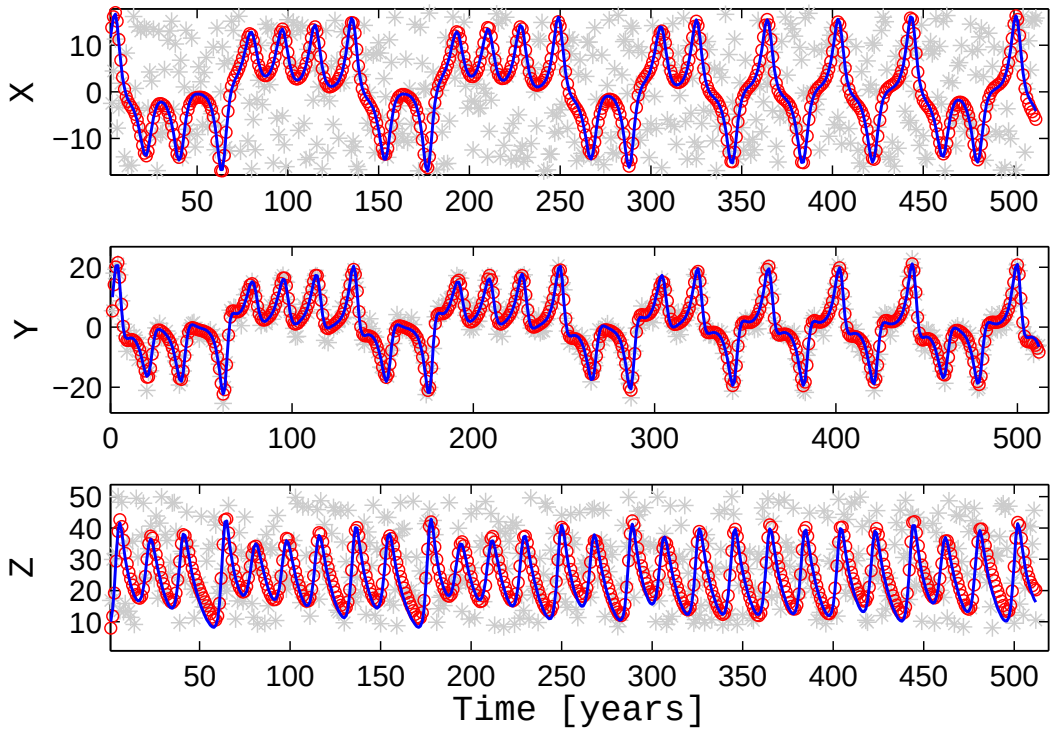

Figure 3.5 State estimation with only one observed variable. The noise level in the $\mathbb{Y}$ component is 5% and the parameters are known. Using the set of observations from the $\mathbb{Y}$ component (gray stars) and random guesses for the $\mathbb{X}$ and $\mathbb{Z}$ components (also gray stars), we obtain estimates (red circles) that are close to the true state (blue line). The errors in the final estimate are 1.7%, 1.9% and 4.5% for the three components.

in this case are: 2.6%, 3.2% and 6.3% for the three model components.

When the Shadowing Filter is applied to the case where only the $\mathbb{Z}$ component is observed for the Lorenz'63 model (shown in Figure 3.7) it is found that the model lacks information about the sign of the $\mathbb{X}$ and $\mathbb{Y}$ components. This generates large errors in the estimate: 32.4%, 28.8% and 1.9% for the 3 components. In the (to us) analogous case of synchronization, a similar behavior has been observed in the context of *Master-Slave synchronization* (Tresser et al. 1995; Balmforth et al. 1997) In the case of synchronization, the $\mathbb{Z}$ component cannot be used as a *synchronizing variable* because it is insensitive to

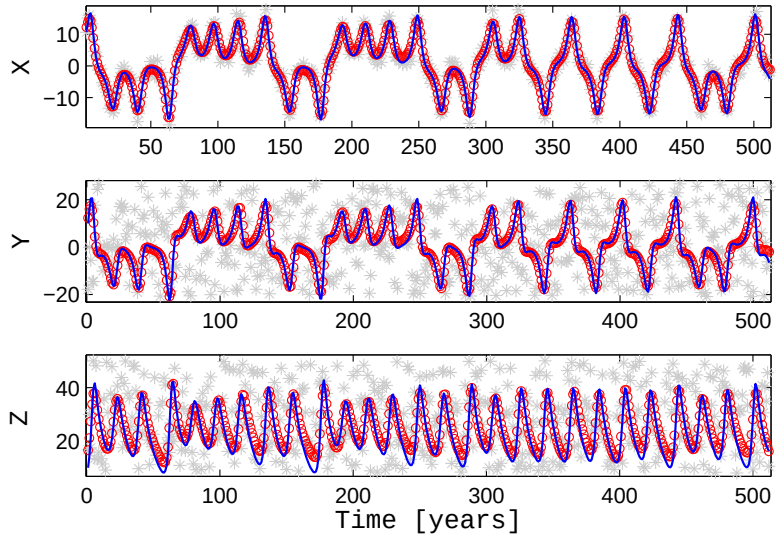

Figure 3.6 State estimation with only one observed variable. The noise level in the $\mathbb{X}$ component is 5% and the parameters are known. Using the set of observations from the $\mathbb{X}$ component (gray stars) and random guesses for the $\mathbb{Y}$ and $\mathbb{Z}$ components (also gray stars), we obtain estimates (red circles) that are close to the true state (blue line). The errors in the final estimate are 2.6%, 3.2% and 6.3% for the three components.

the plus and minus signs of the $\mathbb{X}$ and $\mathbb{Y}$ components.

Note that, despite the sign error, one can still make a very good estimate of the absolute value (see Figure 3.8). In this case the errors of the final estimate for the three components are 2.4%, 2.1% and 1.9%

3.7 Parameter estimation with unobserved variables

The mismatch $\mathcal{M}(\theta)$ is evaluated under identical conditions – same number of iterations etc. – for different parameter values. The smallest value of $\mathcal{M}(\theta)$ corresponds to the optimal value of θ, i.e. parameters that give estimates that are closer to a trajectory are perceived as better.

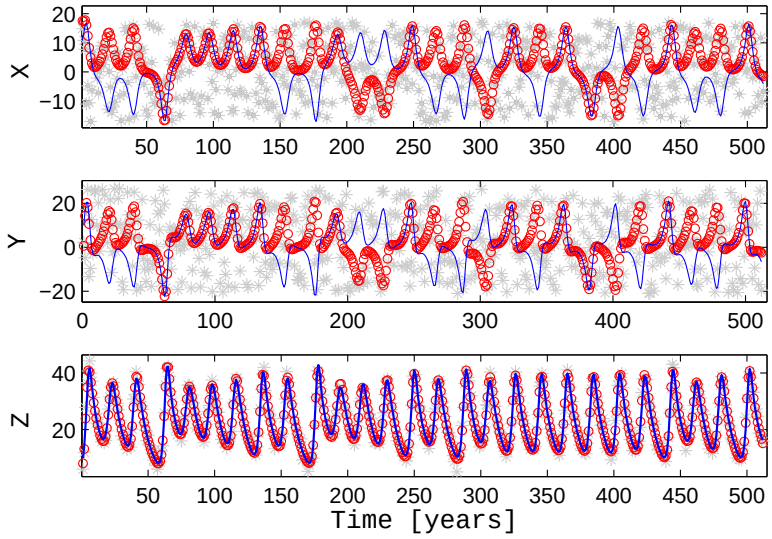

Figure 3.7 State estimation with only one observed variable. The noise level in the $\mathbb{Z}$ component is 5% and the parameters are known. Using the set of observations from the $\mathbb{Z}$ component (gray stars) and random guesses for the $\mathbb{X}$ and $\mathbb{Z}$ components (also gray stars) we obtain estimates (red circles) that are close to the true state (blue line). The errors in the final estimate are 32.4%, 28.8% and 1.9% for the three components.

The mismatch is a measure of how close the sequence of estimates $\{Z_n\}$ is to a trajectory of the model. We would also like to measure how much the noise level has been reduced. For this we use the *implied noise*, which measures the distance between the observations to the estimates (Smith et al. 2010). The implied noise should, in the ideal case, be compared to the noise level as a test for consistency. In cases with unobserved variables the implied noise has to be calculated after we have found a dynamically consistent estimate at the end of the "hold on" stage, as outlined in Section 3.6. It would be difficult to define an implied noise before the end of the "hold on" stage since the observations remain fixed until the "let go" stage and the initial guesses for the unobserved variables can be considered to be arbitrary. In this section, the "hold on" stage consists of 1000 iterations followed by 1000 iterations in the "let go" stage. The implied noise is,

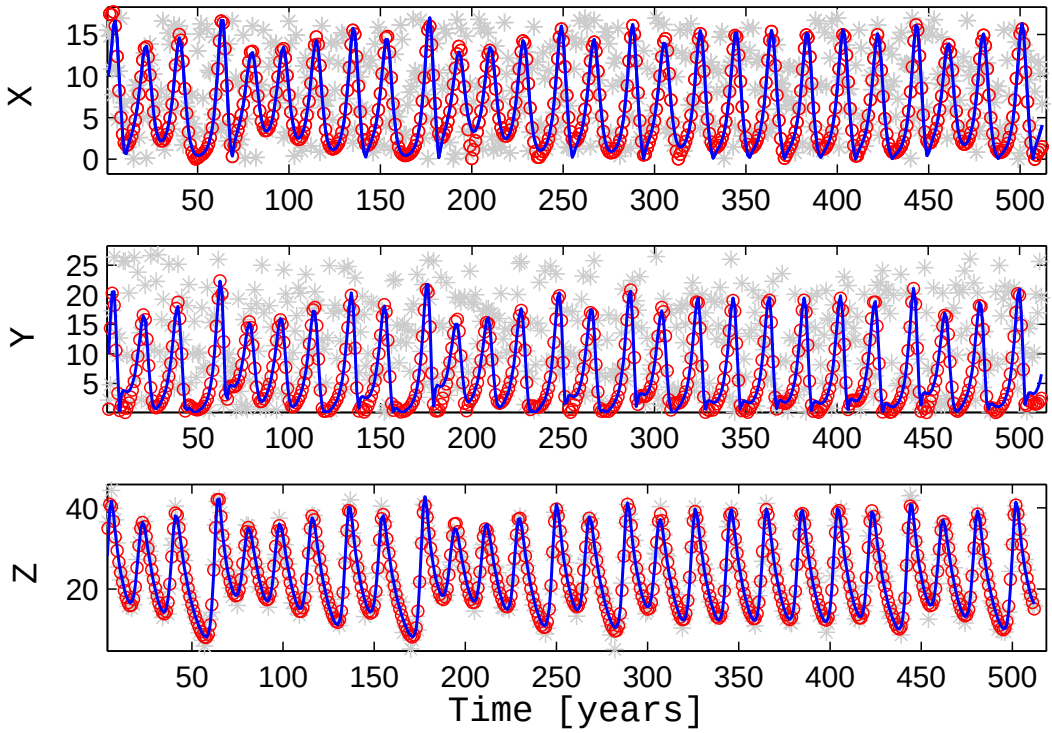

Figure 3.8 State estimation with only one observed variable. Shown is the absolute value, to contrast Figure 3.7. The noise level in the $\mathbb{Z}$ component is 5% and the parameters are known. Using the set of observations from the $\mathbb{Z}$ component (gray stars) and random guesses for the $\mathbb{X}$ and $\mathbb{Z}$ components (also gray stars) we obtain estimates (red circles) that are close to the true state (blue line). The errors in the final estimate are 2.4%, 2.1% and 1.9% for the three components.

in this case, calculated for each component of the model as:

$$\mathcal{I} = \frac{1}{N}\sum_{n=1}^{N}[|Z_n^{1000} - Z_n^{1000+1000}|]. \tag{3.11}$$

Z_n^{1000} is the observation S_n for the observed components. We differentiate between the components because of the different behavior for observed and unobserved variables. In real applications we want the implied noise to be as small as possible, to remain close to

the set of observations.

We explore a two-dimensional slice of parameter space for the r and b parameters. A grid of resolution 50×25 is defined on the $r - b$ parameter space corresponding to a domain of $r \in [24, 48.5]$ and $b \in [0, 4]$ with $s = 10$. There the mismatch and the implied noise are evaluated after a fixed number of iterations.

An explicit true trajectory of the model is generated with a Runge-Kutta of 4th order with $h = 0.01$, and a distance of $w = 4$ integrations steps between each output for each combination of parameters. Three sets of observations of length $N = 512$ for noise levels 5%, 10% and 20% are generated from this trajectory. All variables are simulated, but only one variable is considered to be observed at any specific time. We have also made calculations for two different realizations of the unobserved variables, but only one solution is shown in order to make the diagram more readable. The results of the two cases are not very different.

Figures 3.9, 3.10 and 3.11 show results for the mismatch ($\mathcal{M}$) evaluations over the parameter grid when only $\mathbb{Y}$, $\mathbb{X}$ and $\mathbb{Z}$ are observed, respectively. The mismatch in Figures 3.9 and 3.10 is plotted in logarithmic scale to show higher contrast between values of the functional $\mathcal{M}$ over the parameter space. Note that the range for r in Figure 3.11 is adjusted to [24, 45] for clarity. The $\mathcal{M}$ values for $r > 45$ are very large in this case, affecting the contrast and figure would have been difficult to read without this adjustment.

Two different realizations of the initial guesses for the unobserved variables are shown, labeled as “MismatchA” and “MismatchB”. The difference between the two guesses is only the random seed used in the uniform random sample generated.

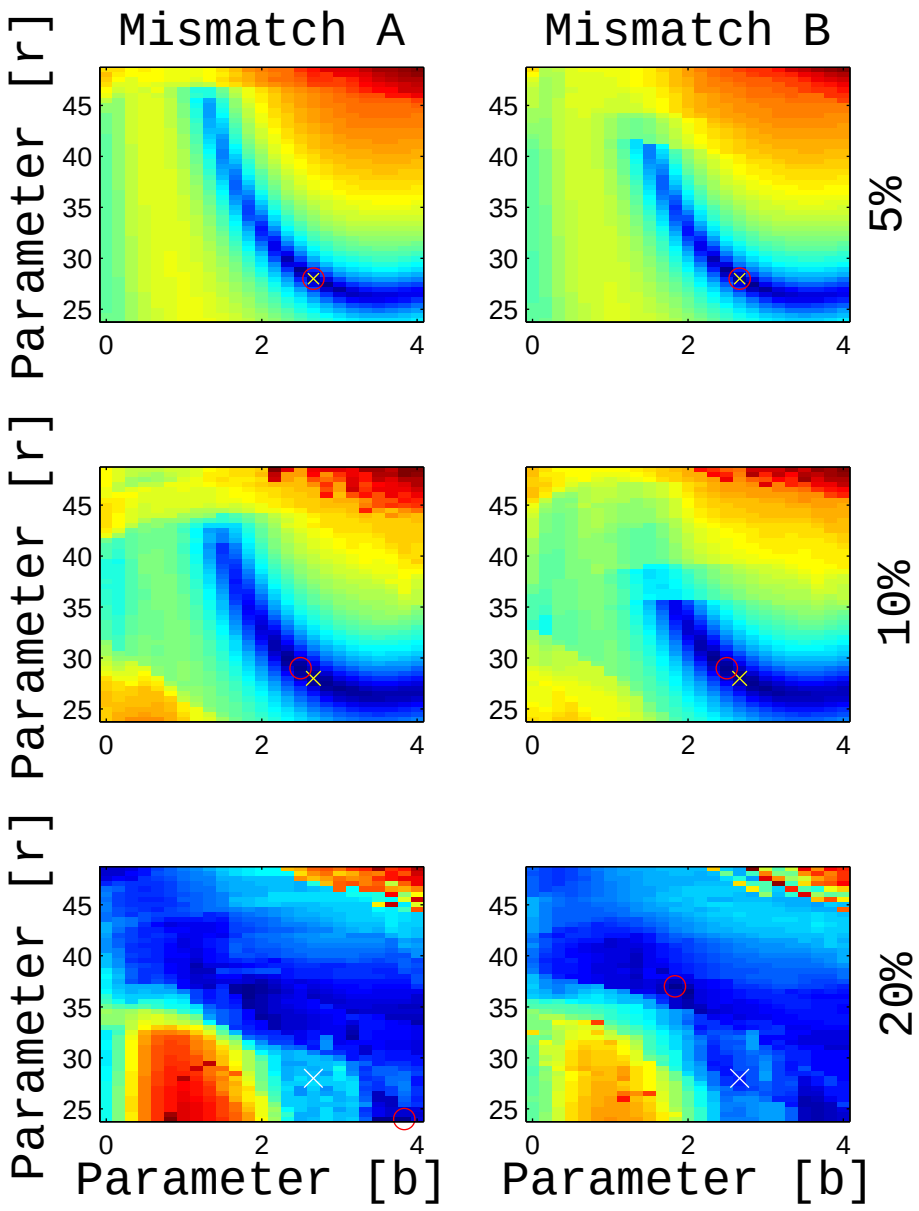

Figure 3.9 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{Y}$ component is observed. The mismatch $\mathcal{M}$ is presented for two different realizations of the unobserved variable guesses and for three noise levels (5%, 10%, 20%). The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

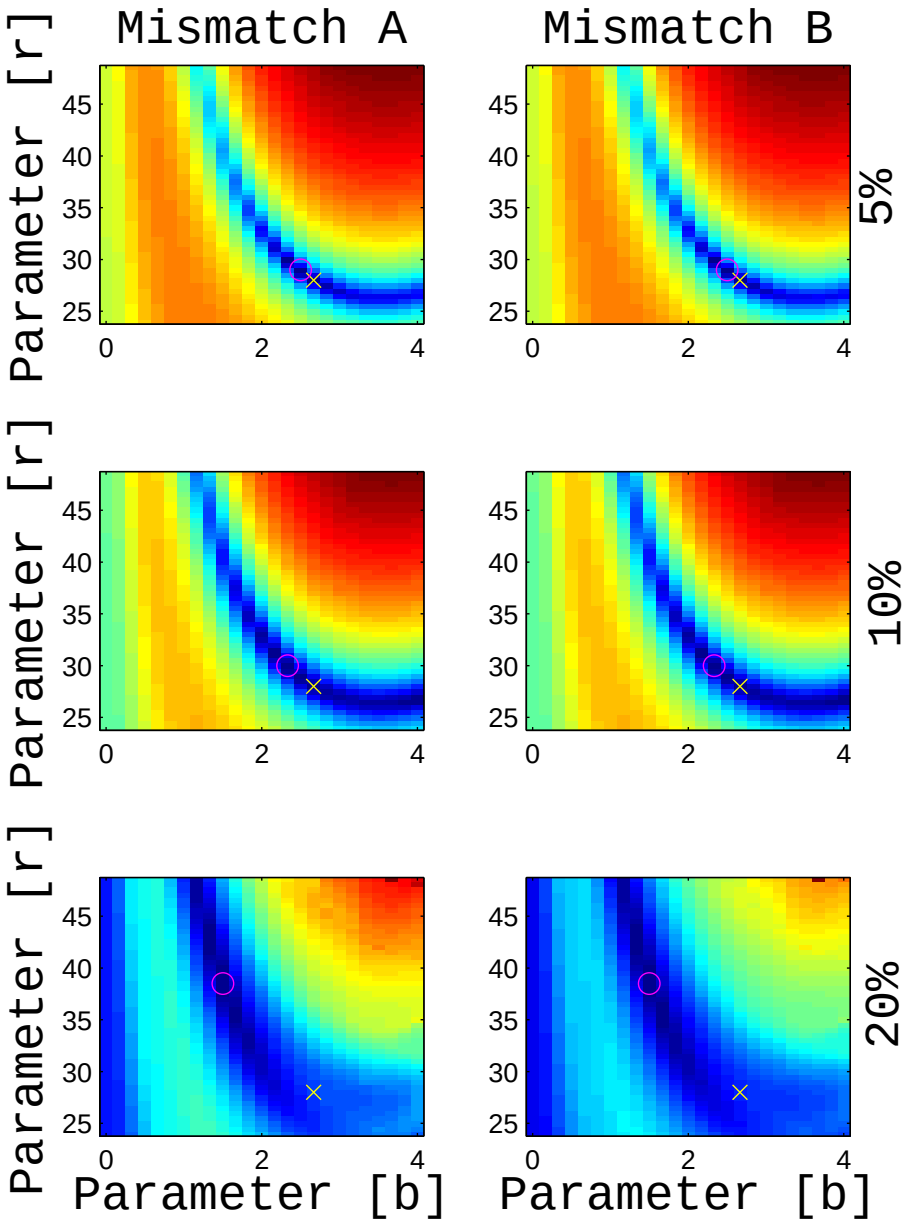

Figure 3.10 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{X}$ component is observed. The mismatch $\mathcal{M}$ is presented for two different realizations for the unobserved variable guesses and for three noise levels (5%, 10%, 20%). The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

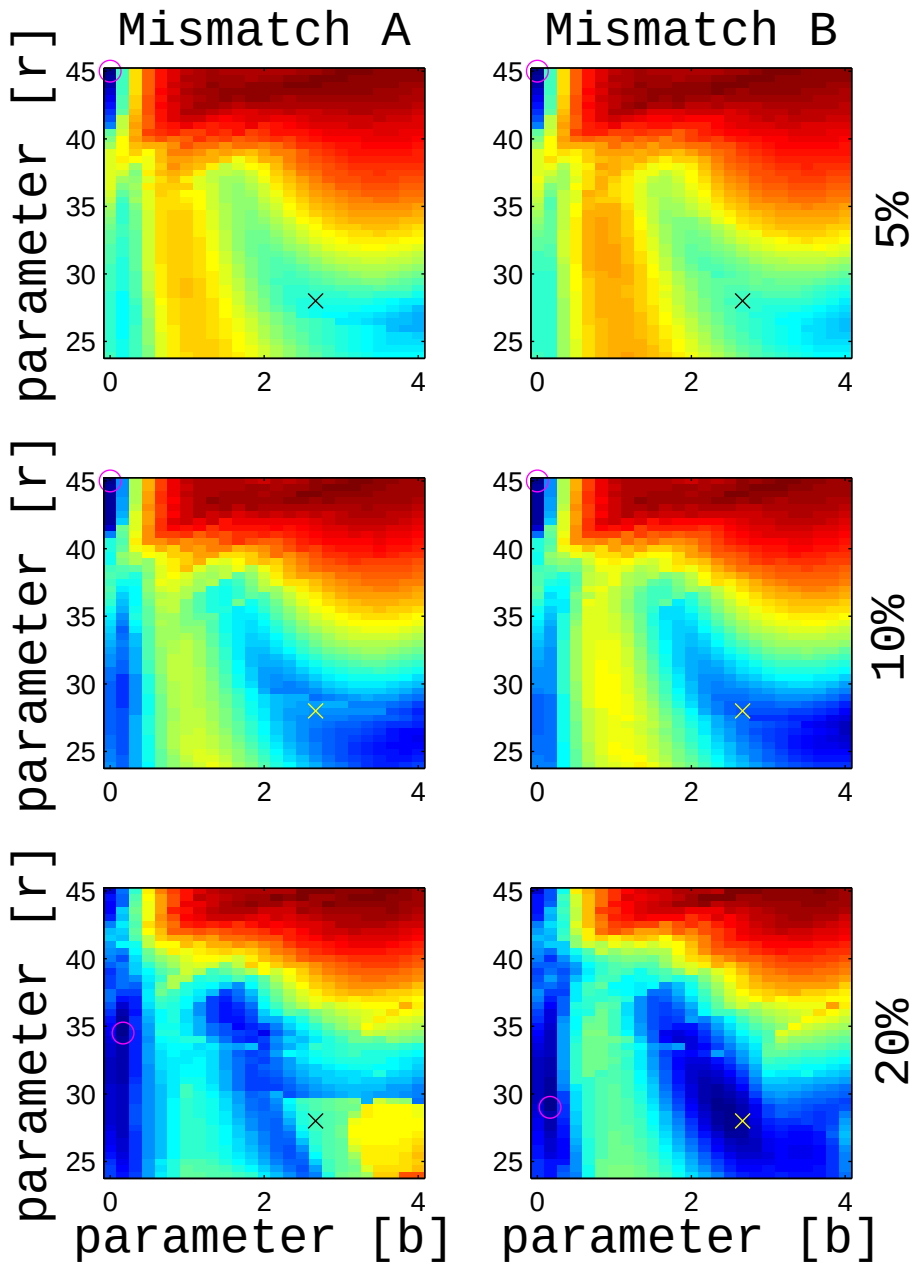

Figure 3.11 Grid representation of parameter values 40x25 in the range $r = [24, 45]$ $b = [0, 4]$ and $s = 10$. The value of $\mathcal{M}$ was very large for $r > 45$ and these values where omitted for readability in this figure. Only the $\mathbb{Z}$ component is observed. The mismatch $\mathcal{M}$ is presented for two different realizations for the unobserved variable guesses and for three noise levels (5%, 10%, 20%). The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

Figures 3.12, 3.13 and 3.14 present the implied noise $\mathcal{I}$ over the parameter grid when only $\mathbb{Y}$, $\mathbb{X}$ and $\mathbb{Z}$ are observed respectively. They are displayed in logarithmic scale for contrast. In all plots, cold dark colors correspond to low values of the measured quantity while warm bright colors correspond to high values of the measures of interest. The true parameter value used to generate the synthetic observations of the system is marked with a cross and the lowest overall in each figure is marked with a circle.

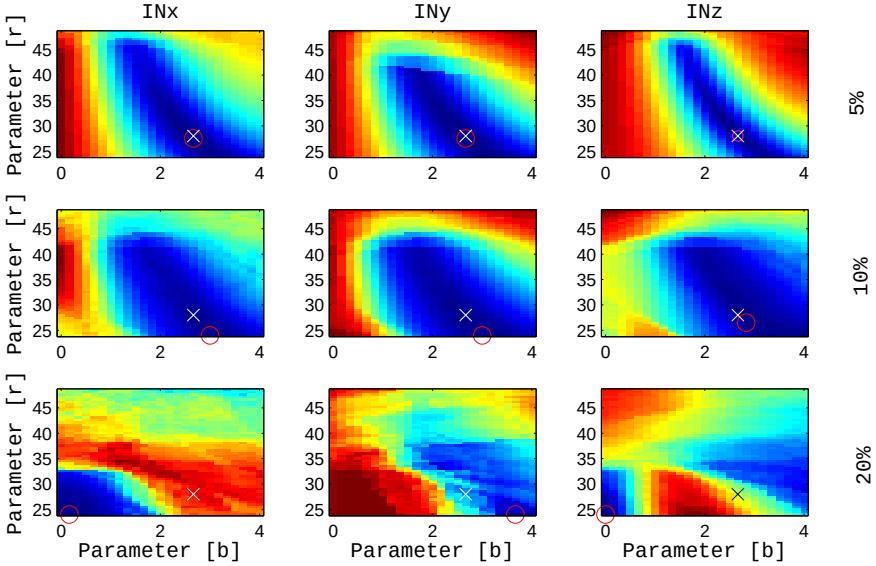

Figure 3.12 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{Y}$ component is observed. The set of nine figures exhibits the implied noise $\mathcal{I}$ shown for the three components of the model (INx, INy, INz) and for three noise levels (5%, 10%, 20%). The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

We have because of the twin experiment construction an explicit true state of the system. Because of this we evaluate the component by component *absolute error* between the true state and the final estimate. This is shown in Figures 3.15, 3.16 and 3.17 over

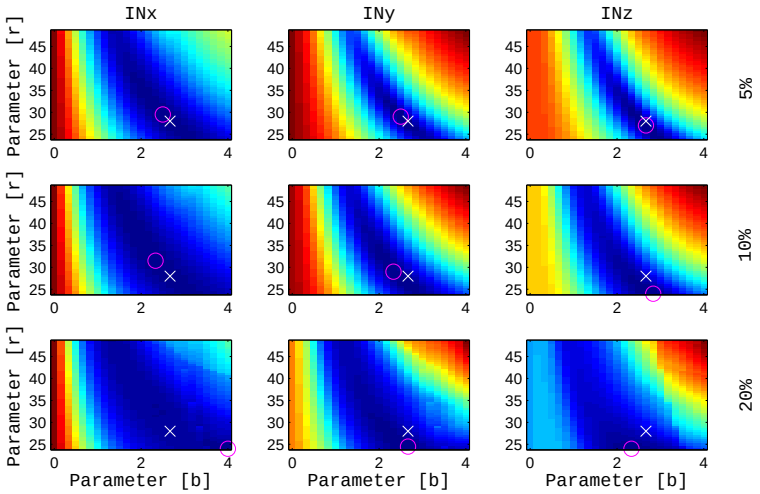

Figure 3.13 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{X}$ component is observed. The nine figures exhibit the implied noise $\mathcal{I}$ shown for the three components of the model (INx, INy, INz) and for three noise levels (5%, 10%, 20%). The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

the parameter grid when only $\mathbb{Y}$, $\mathbb{X}$ and $\mathbb{Z}$ are observed respectively. Of particular interest is Figure 3.17 with its non-regular pattern that does not appear to have a clear correspondence to the estimates in Figure 3.14. As noted before we know that the sign of the estimate is difficult to determine if only the $\mathbb{Z}$ component is observed. This is seen in the absolute error. Figure 3.18 demonstrates the estimate of the amplitude where the ambiguity of the sign has been removed, as was also done in Figure 3.8.

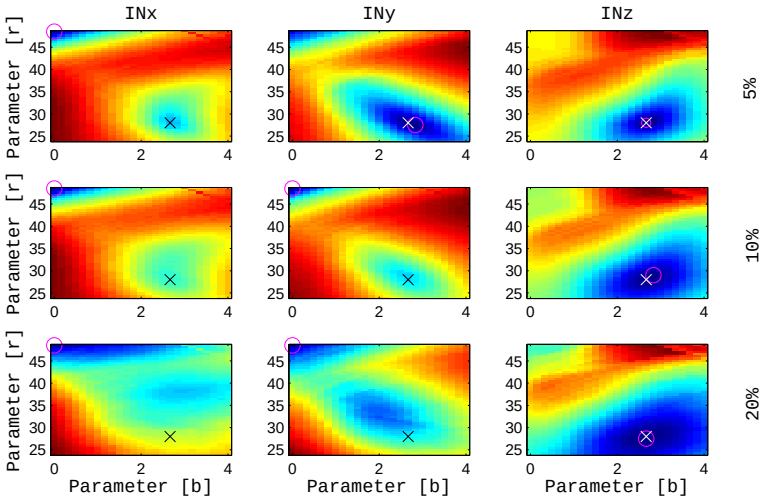

Figure 3.14 Grid representation of parameter values 50×25 in the range $r = [24, 48.5]$, $b = [0, 4]$ and $s = 10$. Only the $\mathbb{Z}$ component is observed. The nine figures exhibits the implied noise $\mathcal{I}$ shown for the three components of the model (INx, INy, INz) and for three noise levels (5%, 10%, 20%).. The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

The *prediction time* is also calculated. For each set of parameters an estimate is made and this estimate is used to make a prediction. The prediction error – distance between the prediction and the known true state – is evaluated. The prediction time is the number of years before the prediction error is larger than the noise level. We can use the absolute error and the prediction time to see what parameters correspond to "close estimates" and "good predictions".

In Figure 3.19 we show the prediction time when only the $\mathbb{Y}$ component is observed. The figure demonstrates the prediction time for grids of 50×25 parameter values and for three noise levels (5%, 10% and 20%) and for the three components of the model (PTx, PTy, PTz). Similar figures for observing the $\mathbb{X}$ component is presented in Figure 3.20,

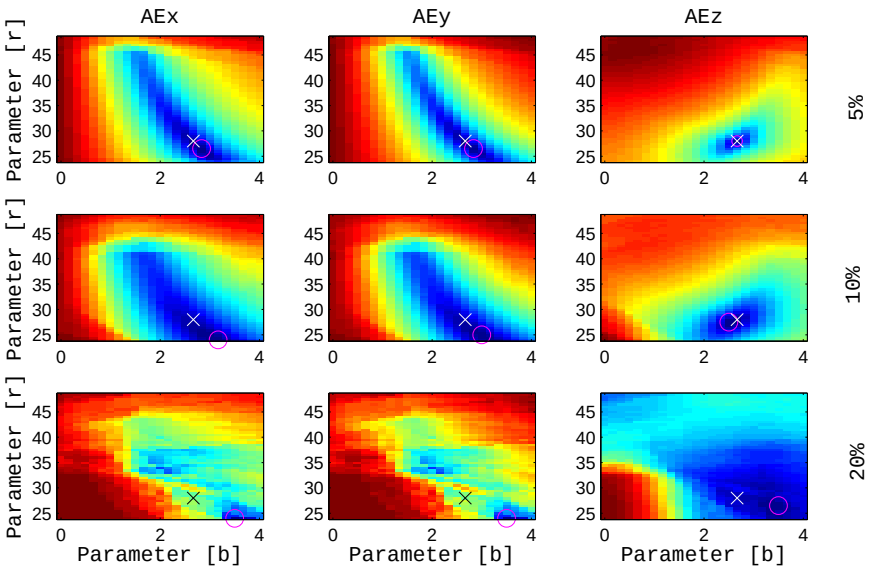

Figure 3.15 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{Y}$ component is observed for three noise levels. The absolute error is shown in the nine figures for the three components (AEx, AEy, AEz) and for the same three noise levels. The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

and the $\mathbb{Z}$ component case is presented in Figure 3.21.

Consistent throughout all these cases is a "valley" of interesting values that starts at the top of the figures and curves down to the lower right. Apparently the dynamics of the model is similar for these combinations of parameters, but it is somewhat unclear if there is another connection. These are values that are indicated as "good" by the mismatch and the implied noise charts. They also give low absolute error and typically long prediction time (high PTx, PTy and PTz values). The exact estimate of the true value becomes harder to get for higher noise levels.

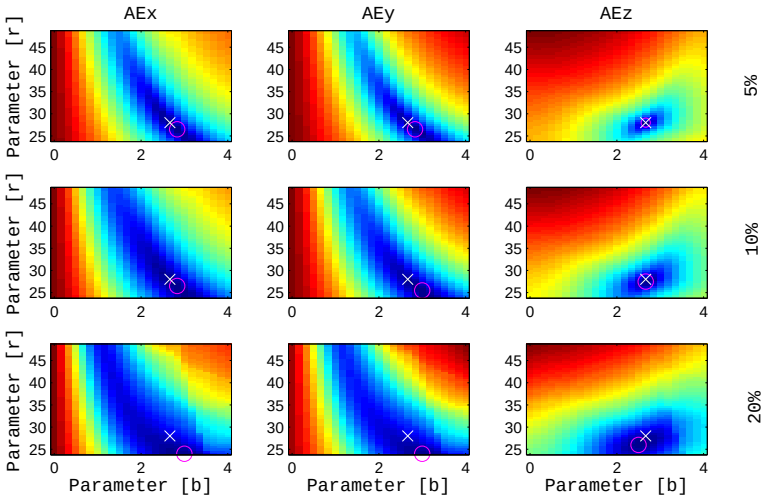

Figure 3.16 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only $\mathbb{X}$ component is observed for three noise levels. The absolute error is shown in the nine figures for the three components (AEx, AEy, AEz) and for three noise levels (5%, 10%, 20%). The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

3.8 Conclusion

In this chapter we have illustrated the study of parameter estimation in the case of unobserved variables for different noise levels for the chaos model of Lorenz (1963). The conclusion is that parameter estimation is possible for moderate noise levels, but an exact estimate of the correct parameters is difficult to achieve if the noise level is high (around and above 20%). The areas of parameter space indicated by low $\mathcal{M}$ and $\mathcal{I}$ values also correspond to close estimates (low absolute error) and good predictions (long prediction time). This is true for a noise level around 10% and below and when the $\mathbb{X}$ or $\mathbb{Y}$ variables are observed. Parameter estimation is possible to some extent even when only the $\mathbb{Z}$ component is observed, but the connection between "good parameters" and

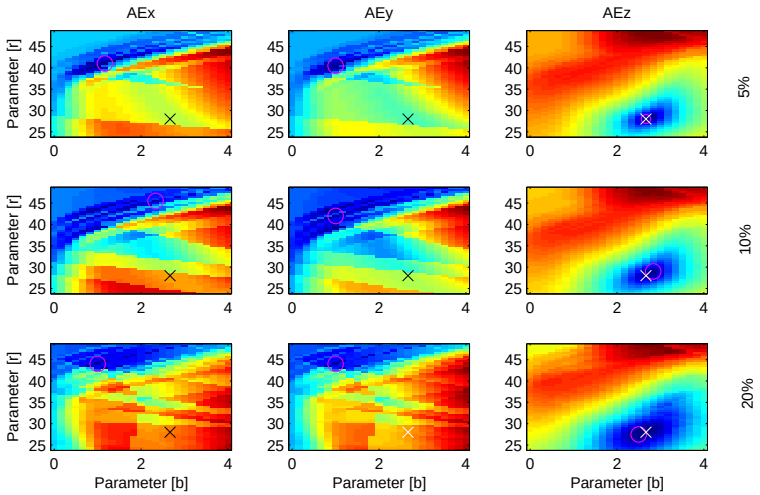

Figure 3.17 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{Z}$ component is observed for three noise levels. The absolute error is shown in the nine figures for the three components (AEx, AEy, AEz) and for three noise levels (5%, 10%, 20%). The color-code of the lower set is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

long prediction is not clear in this case.

The same baseline with 1000+1000 iterations were used: 1000 iterations during the "hold on" stage followed by 1000 iterations for the "let go" stage. The second stage is initialized but not constrained by the set of observations. This might be problematic in a real world application. There might also be better convergence – for the 20% noise case – after additional iterations.

The fact that the r, b and s parameters can be determined, even when the related variable is not observed, provides some confidence that dynamical information has been retrieved from the unobserved variables. The failure of the high noise level cases and when only the $\mathbb{Z}$ component is observed indicates in an interesting way the limits of our

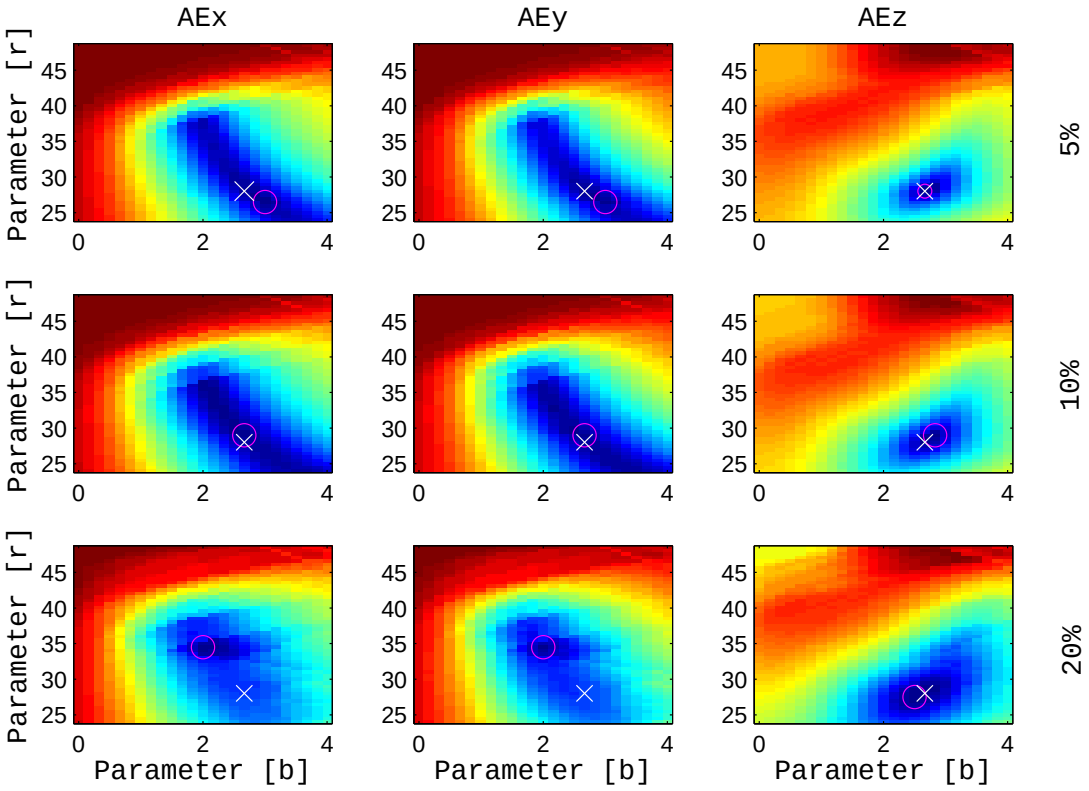

Figure 3.18 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{Y}$ component is observed for three noise levels (5%, 10%, 20%). The three components of the absolute value (AEx, AEy, AEz) are evaluated for the absolute value of the estimate as in Figure 3.8. The color-code is given in logarithmic scale for each individual figure and does not correspond to an absolute scale. The cross indicate the parameter used to generate the observations and the circle indicate the lowest value in the explored parameter grid.

methods. It also tells us that all variables and parameters are not carriers of the same information. Each system has to be explored in depth to determine their specific features and to what extent these methods can be applied.

Unsynchronized variables – in the notion of Master-Slave synchronization – contain limited information about the parameter values of the other components, and a similar behavior is noted here. It is unclear if there is a direct connection between the two

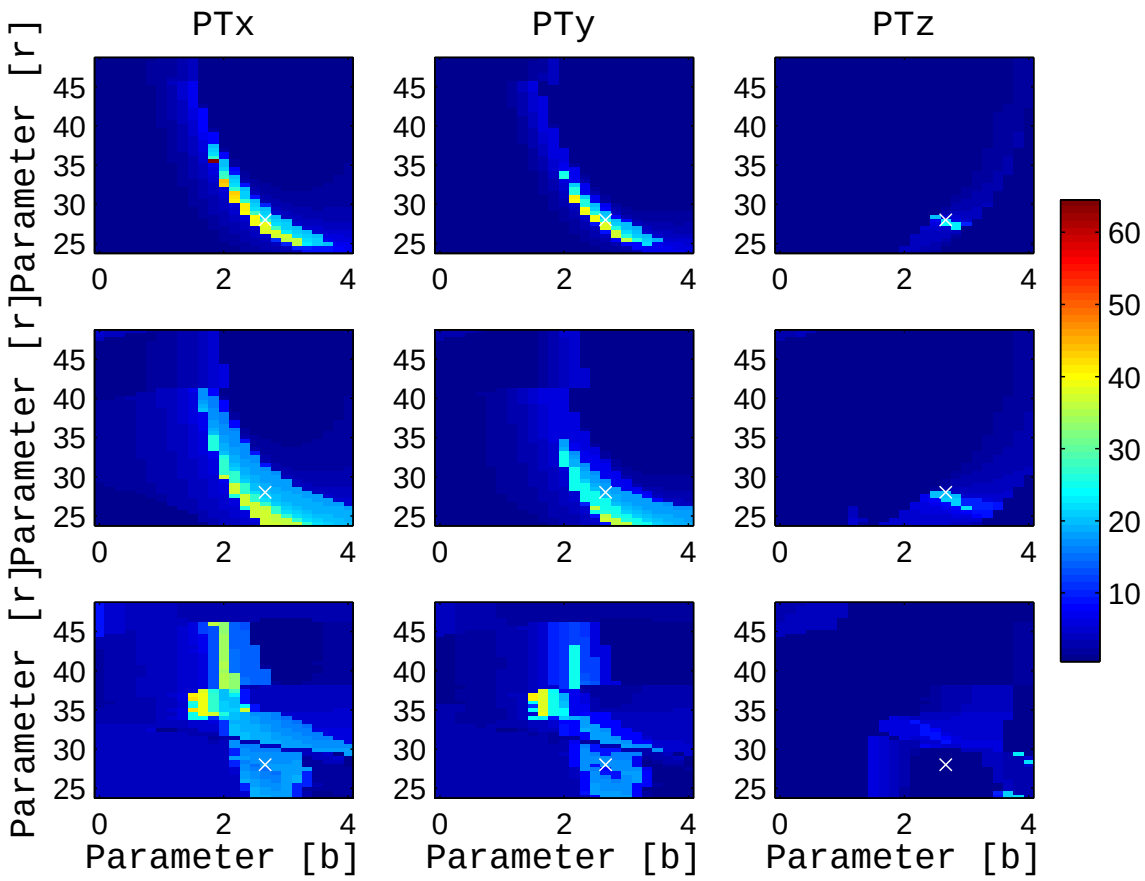

Figure 3.19 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{Y}$ component is observed for three noise levels. Shown is the prediction time for three components of the model (PTx, PTy, PTz) and three noise levels (5%, 10%, 20%). The values of the color code are shown by the color-bar. See Figures 3.9, 3.12 and 3.15 for the corresponding figures.

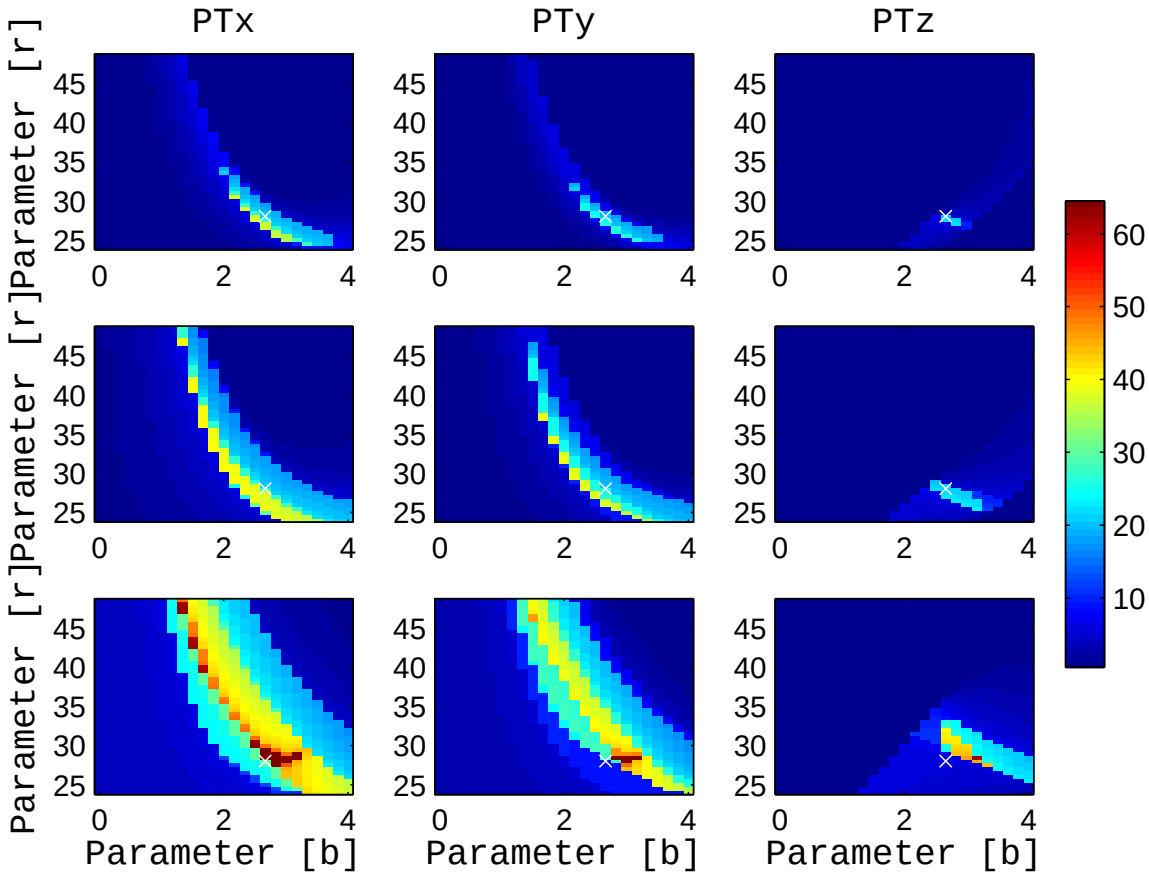

Figure 3.20 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{X}$ component is observed for three noise levels. Shown is the prediction time for three components of the model (PTx, PTy, PTz) and three noise levels (5%, 10%, 20%). The values of the color code are shown by the color-bar. See Figures 3.10, 3.13 and 3.16 for the corresponding figures.

procedures: The Shadowing Filter operates simultaneously over the entire sequence of estimates, whereas Master-Slave synchronization is a process that runs concurrent with the time evolution of the system.

The verification of the procedure was only possible because a true state was generated and synthetic observations were constructed by adding noise to this signal. This work has operated in the ideal case, where the system and the model are one and the same. In

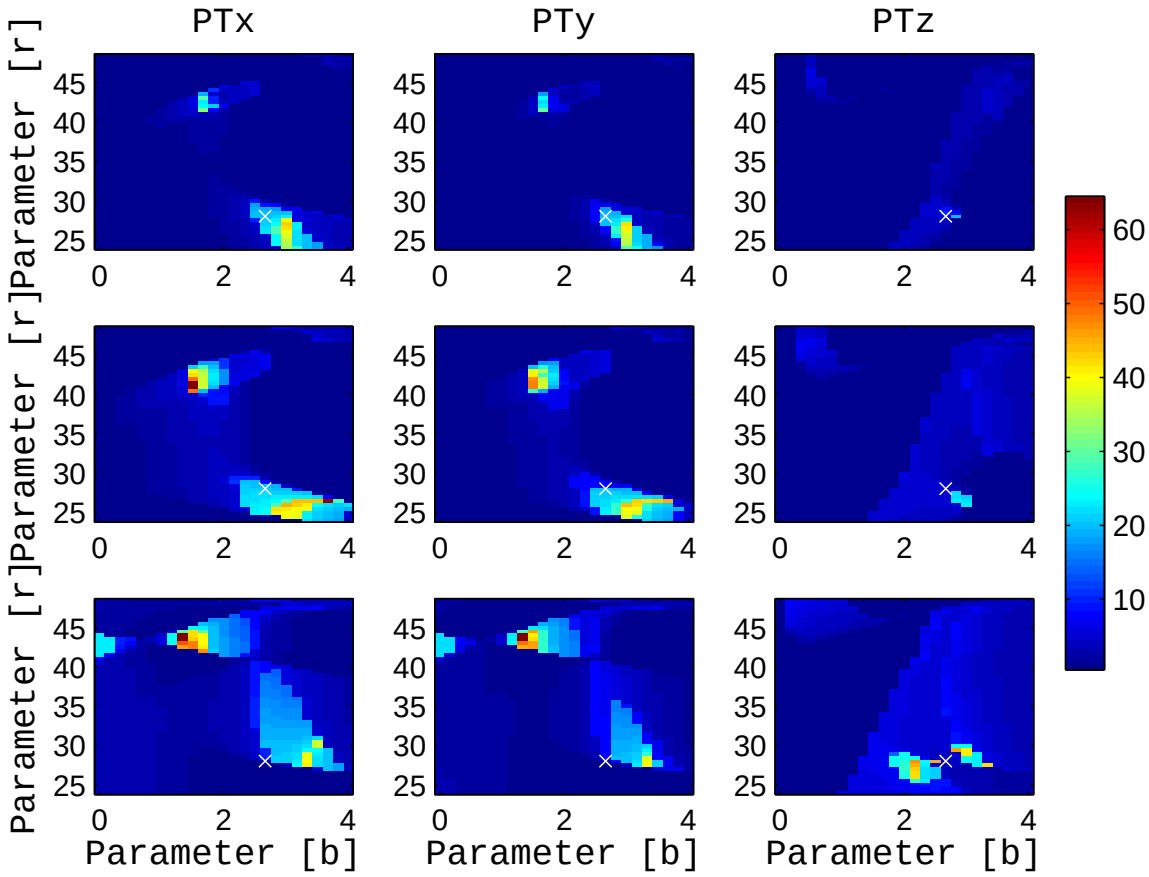

Figure 3.21 Grid representation of 50×25 parameter values in the range $r = [24, 48.5]$, $b = [0, 4]$, $s = 10$. Only the $\mathbb{Z}$ component is observed for three noise levels. Shown is the prediction time for three components of the model (PTx, PTy, PTz) and three noise levels (5%, 10%, 20%). The values of the color code are shown by the color-bar. See Figures 3.11, 3.14, 3.17 and 3.18 for the corresponding figures.

real world application this does not have to be the case and extra considerations have to be made. For example notions such as "finding the true state of the system" might not even make sense in the real world application[6]. In this study we learned that if we want to use the Lorenz'63 system to model a real system one should try to observe at least the $\mathbb{X}$ or $\mathbb{Y}$ components and have a noise level that is expected to be below 10%. This is most

[6]Also referred to as the *imperfect model scenario* in literature (Judd & Smith 2004).

likely not a sufficient condition for correct state or parameter estimation, but it is most certainly a necessary one.

Chapter 4

Data assimilation for the Solar Cycle model

We want to outline a path towards a model-based prediction of the solar cycle. In this chapter we have selected a nonlinear low-dimensional dynamical model to describe the solar cycle, presented in Section 4.1. We want to test the chosen model using the tools developed in Chapter 3 and study under what conditions model state estimation and prediction are possible. The two main features we want from the model are, modulation of the peak amplitude of the oscillation and the ability to enter times of inactivity (see Section 1.2). The chapter is outlined as follows:

First, we qualitatively explore the model and study how changing the model parameters affects the model output (Section 4.1). We find ranges of parameters that generate output reminiscent of the observations.

Second, we explore the solar cycle data and construct a representation of the data

that supports the data assimilation process.

Third, we make a twin experiment – under ideal conditions and with synthetic data – and find that our method does not perform as expected for the chosen model. We identify a specific issue that occurs when the amplitude of one of the model components is significantly lower than the others. An approach using linear scaling is introduced to address this problem in Section 4.4.

Fourth, we make a second twin experiment, where –still under ideal conditions and with synthetic data – we allow only one observed variable (Section 4.5).

This setup closely resembles what would be required for the real case. The main difference from the real case is that our model and system are one and the same, and we know the correct set of parameters. We study how the noise level and the linear scaling affects the error of the estimate of the system state and the ability to predict. We make predictions both at the onset of an inactive phase (called Set A) and at normal activity (called Set B).

Finally, we discuss the implications of our work on the real case (Section 4.6).

4.1 The PST model

There are multiple solar cycle models designed using different mechanisms and assumptions. For a thorough review and comparison between different models see Pasquero (1995); also see Weiss (2010) for an up-to-date discussion. Based on the evaluation of Pasquero (1995) we choose a model developed by Platt, Spiegel and Tresser (Platt et al. 1993b). Throughout this work this model will be referred to as the PST model. The

model is designed to emulate the long term transition from periods of "normal" solar activity to periods of inactivity using the *on-off intermittency* mechanism (Platt et al. 1993a; Spiegel 1997).

It turns out that (see Figure 4.1) the variation in cycle amplitude of the PST model also in a qualitative way resembles the 11-year solar cycle. The model has two components:

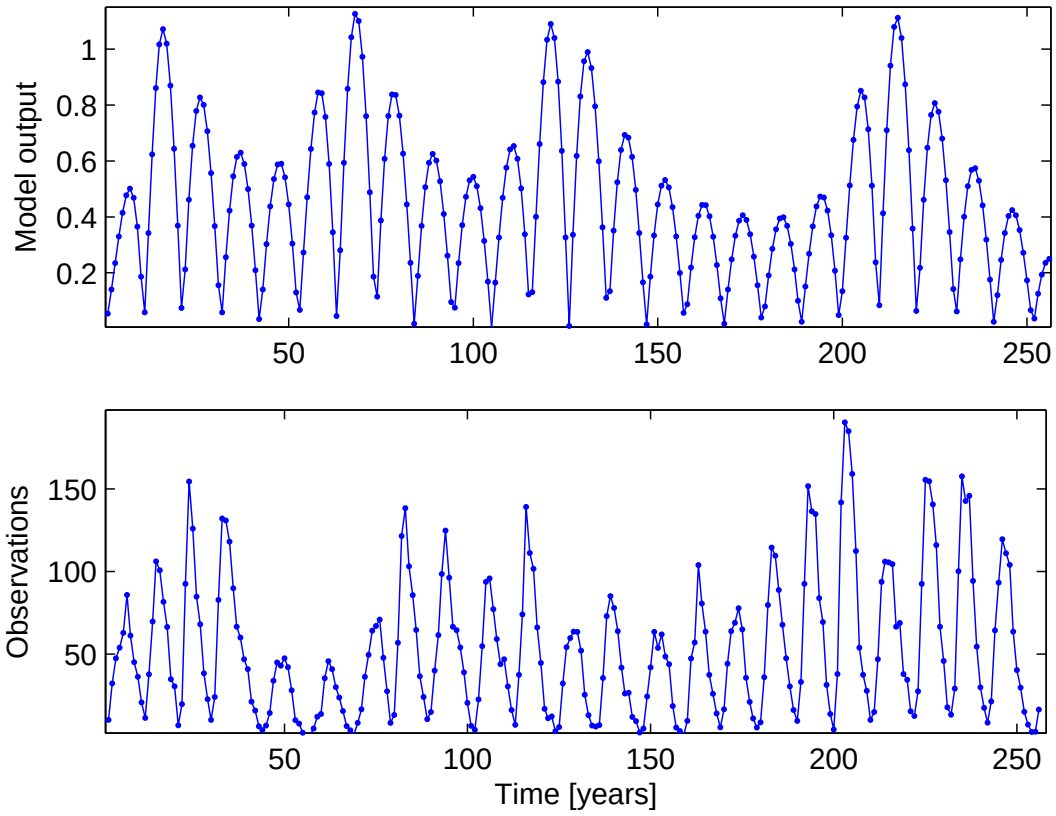

Figure 4.1 Output from the PST model, absolute value of $\mathbb{U}$, for standard parameters values, in the upper panel. The yearly average Sunspot number for the last 24 cycles (1755-2010), in the lower panel.

a nonlinear oscillator and a chaotic system. The oscillator represents a dynamo process

in the Tachocline (Spiegel & Zahn 1992; Spiegel 1997, 2009) with expression:

$$\begin{aligned}\dot{\mathbb{U}} =& -(\mathbb{U}^2+\mathbb{V}^2)\mathbb{U} - \omega\mathbb{V} + \beta(\mathbb{X}-x_0)\mathbb{U} \\ \dot{\mathbb{V}} =& -(\mathbb{U}^2+\mathbb{V}^2)\mathbb{V} + \omega\mathbb{U} + \beta(\mathbb{X}-x_0)\mathbb{V}\end{aligned} \tag{4.1}$$

corresponding to a *supercritical Hopf bifurcation* (Platt et al. 1993b). The amplitude is influenced by the chaotic motion of the solar convection zone:

$$\begin{aligned}\dot{\mathbb{X}} =& \mathbb{Y} \\ \dot{\mathbb{Y}} =& a\mathbb{X} - \mathbb{X}^3 - b\mathbb{Y} + \mathbb{Z} + q\mathbb{U}^2 \\ \dot{\mathbb{Z}} =& -d(\mathbb{Z} - c\mathbb{X}(\mathbb{X}^2-1))\end{aligned} \tag{4.2}$$

reminiscent of the Moore & Spiegel (1966) system. The $\mathbb{U}$ component is related to the solar polodial field, or the Sunspot number, and $\mathbb{X}$ is related to the generation of magnetic activity or "magnetic rainfall" (Platt et al. 1993b).

Output of the PST model, i.e. the absolute value of the $\mathbb{U}$ component, are shown in upper panel of Figure 4.1 together with a time series of the yearly average Sunspot number. There is no real difference between the $\mathbb{U}$ and $\mathbb{V}$ components and we could have used either of them to correspond to the data. The output in Figures 4.1 and 4.2 are both generated from the initial conditions $\mathbb{U}_0$, $\mathbb{V}_0$, $\mathbb{X}_0$, $\mathbb{Y}_0$, $\mathbb{Z}_0$= [−0.0534, 0.2533, 0.6229, 0.4534, 0.0308]. These are the default initial conditions used if nothing else is indicated. The first 250 values of the $\mathbb{U}$ component seen in Figure 4.2 correspond to the upper panel in Figure 4.1. We have also used the standard parameters [β, $x0$, ω, a, b, q, d, c] = [1, −0.15, 2, 0.7, 0, 0, 0.03, −0.5] outlined by Platt et al. (1993b) and explored by Pasquero

(1995).

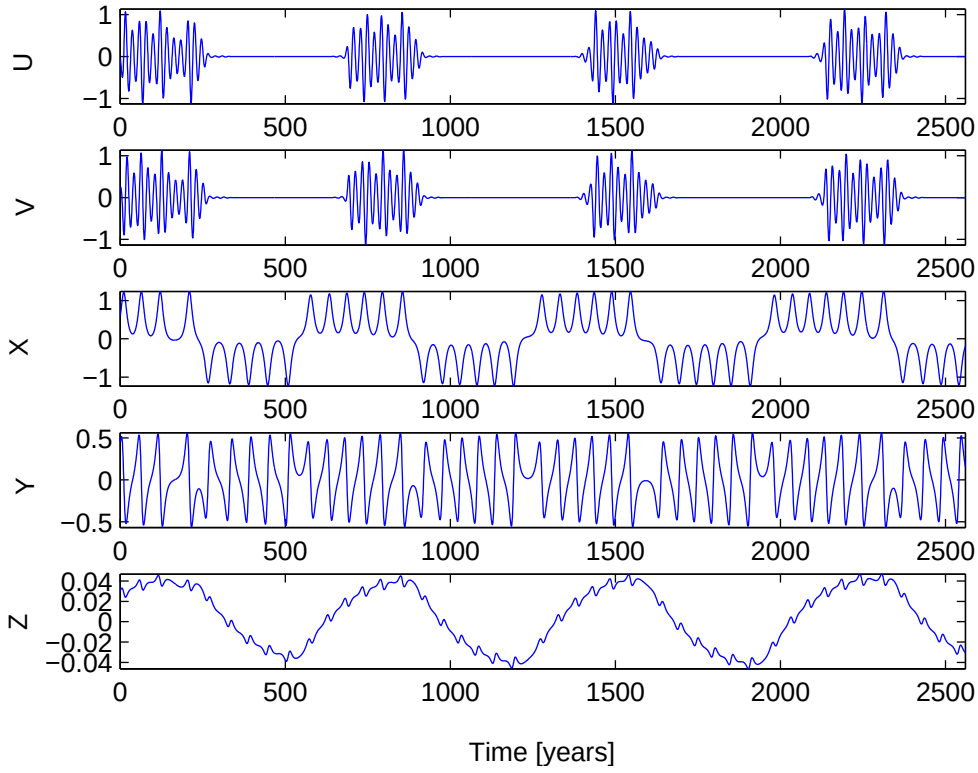

Figure 4.2 Output from the PST model. All five components are generated using standard parameters [β, $x0$, ω, a, b, q, d, c]= [1, -0.15, 2, 0.7, 0, 0, 0.03, -0.5] and default initial conditions: $\mathbb{U}_0$, $\mathbb{V}_0$, $\mathbb{X}_0$, $\mathbb{Y}_0$, $\mathbb{Z}_0$= [-0.0534, 0.2533, 0.6229, 0.4534, 0.0308]. The time units correspond to years.

The observed solar cycle record demonstrates a number of features that are not reproduced by the PST model. Some examples include:

- The cycle length varies from 9-14 years and the PST model has a fixed phase set to 11 years.

- The shape of the cycle is skewed and with fast rise time and slow decay time (Hathaway 2010). The PST output has symmetric cycles.

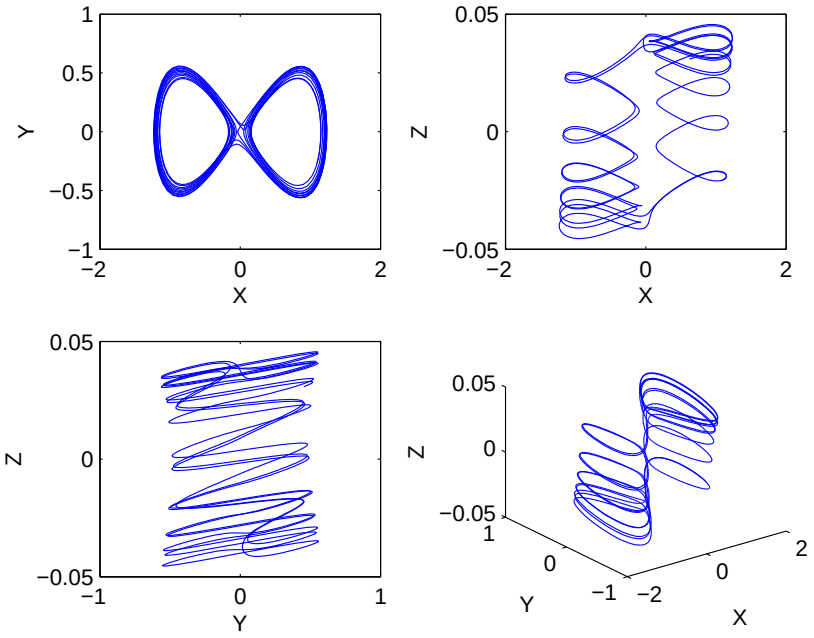

Figure 4.3 Output from the PST model. State space behavior for each combination of the $\mathbb{X}$, $\mathbb{Y}$ and $\mathbb{Z}$ variables. Also seen is the state space behavior for all three components (lower right).

- The envelope of the peaks is smoother for the PST model than for the solar data, see Figure 4.1.

- The transition from activity to inactivity follows the periodic variation of the $\mathbb{Z}$. The Sun does not appear to follow such regular variation.

In Figure 4.2 the long-term variation is shown for all model variables and Figure 4.3 exhibits the corresponding state space behavior for the last three components ($\mathbb{X}$, $\mathbb{Y}$, $\mathbb{Z}$). The peak amplitude of the $\mathbb{U}$ variable – associated with the Sunspot number – is in the [0.6, 1.2] range to be compared to the [50, 190] range for the observed peak Sunspot number.

The $\mathbb{U}$ and $\mathbb{V}$ components of the model vary with a fixed 22 year cycle, where every

11 year cycle has an alternating sign. This behavior and its relation to the solar cycle are presented in Section 4.2. This system also alternates – using standard parameters – between an *on-phase* and an *off-phase*. An off-phase is a time interval when there is no activity neither in the $\mathbb{U}$ nor $\mathbb{V}$ variables. Positive values for $\mathbb{Z}$ appear to correspond to an

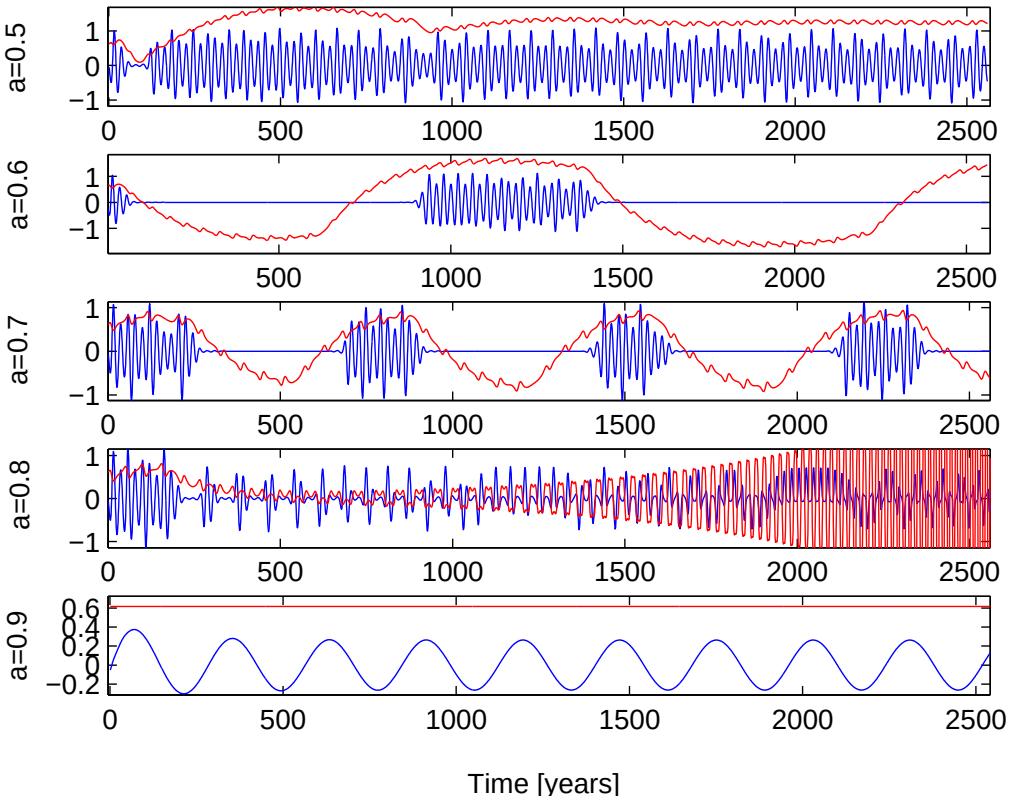

Figure 4.4 Output from PST model, $\mathbb{U}$ is presented in blue and $20 \times \mathbb{Z}$ in red for standard parameters except for the a parameter which is varied here with five cases. From the top: $a = [0.5, 0.6, 0.7, 0.8, 0.9]$. The standard value is $a = 0.7$.

on-phase. This fact could be used to detect a possible transition to a grand minimum, if a negative trend in the $\mathbb{Z}$ variable could be inferred from the Sunspot data. The variation in $\mathbb{Z}$ is cyclic and generates the transitions between activity and inactivity. This transition from periods of activity to inactivity is a key feature in our study. We check ranges of

parameter values to find when the system transitions away from this desired behavior.

The effect of changing parameter a is shown in Figure 4.4. The $\mathbb{U}$ variable (in blue) is shown with the $\mathbb{Z}$ variable scaled with a factor of 20 (in red). The figure suggests the use of values close to $a = 0.7$ to obtain output with the desired behavior.

Parameters c and d determine the length and frequency of the *on-phase* and the *off-phases* as seen in Figure 4.5. If we could determine these parameters accurately, from historical observations of the sun, we could infer an expected length of the current grand maximum.

The amplitude of the cycles is set by the *growth rate* $\beta(\mathbb{X} - x_0)$. For $\mathbb{X} < x_0$ we see a damping and a transition to an "off-phase" occur. To use the language of Platt et al. (1993b): the rate of supply (expressed by the $\mathbb{X}$ variable) is smaller than the demand rate (expressed by the equilibrium parameter value $x0$). The threshold between activity and non-activity is set by $x0$. In the physical world this might be related to the level of the magnetic field strength required to form Sunspots (Penn & Livingston 2006; Livingston & Penn 2009). In Figure 4.6 the effect of changing the $x0$ and β parameters are shown. High values of $x0$ generate a lot of activity in the "off-phase". Both low values and high values of β remove the long-term cyclic behavior.

Higher values for β increase the amplitudes of the $\mathbb{U}$ component but there appears to be a limit to how much the β parameter can enhance the amplitude without affecting the main features of the model.

Parameters b and q are traditionally set to zero. Parameter q acts as a friction term and couples the oscillation to the chaotic system. However, Figure 4.7 shows that even

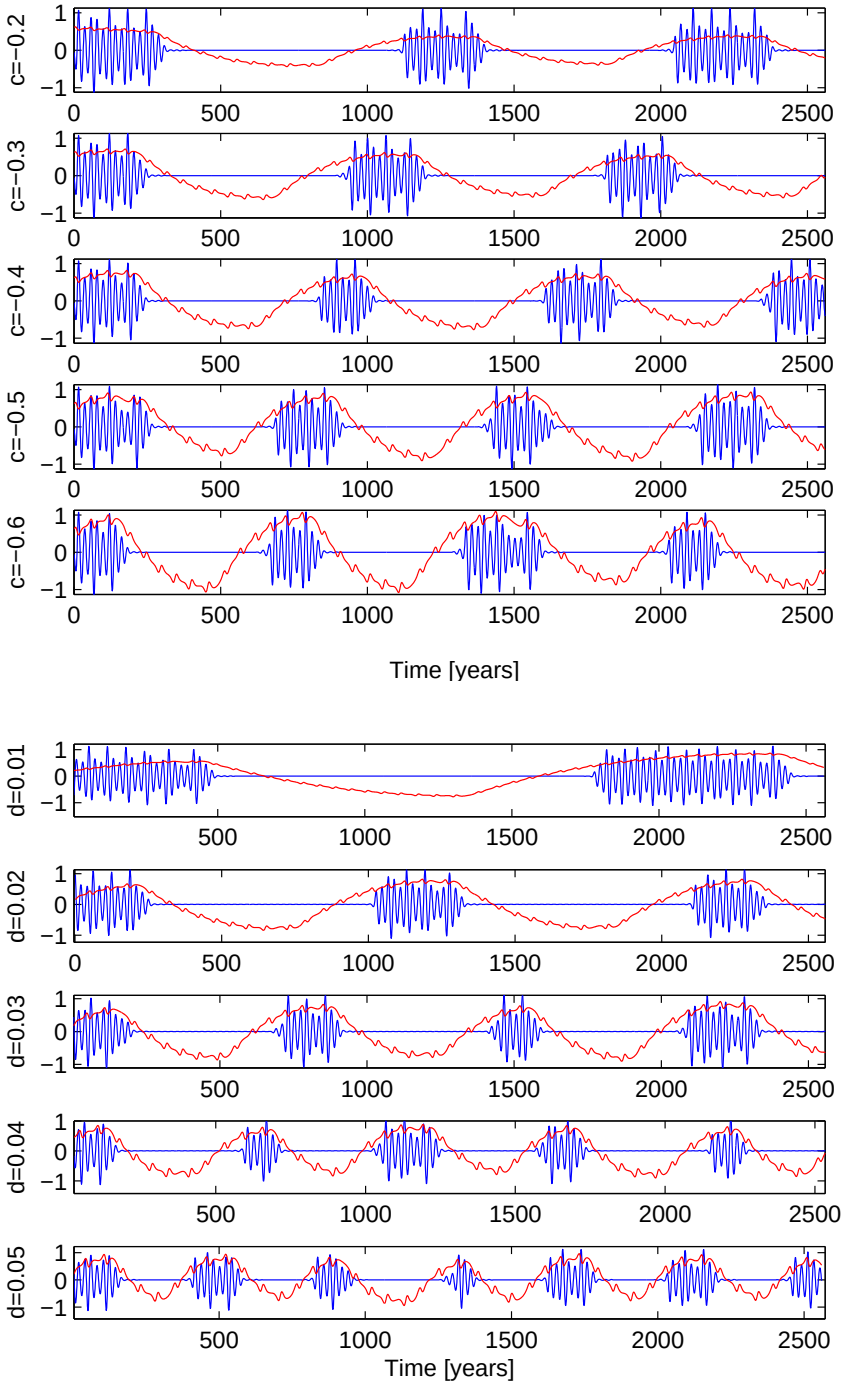

Figure 4.5 Output from PST model. $\mathbb{U}$ is presented in blue and $20 \times \mathbb{Z}$ in red for default initial conditions and standard parameters except for the c and d parameters. In the upper set of five figures the c parameter is varied $c = [-0.2, -0.3, -0.4, -0.5, -0.6]$. In the lower set of five figures the d parameter is varied: $d = [0.01, 0.02, 0.03, 0.04, 0.05]$.

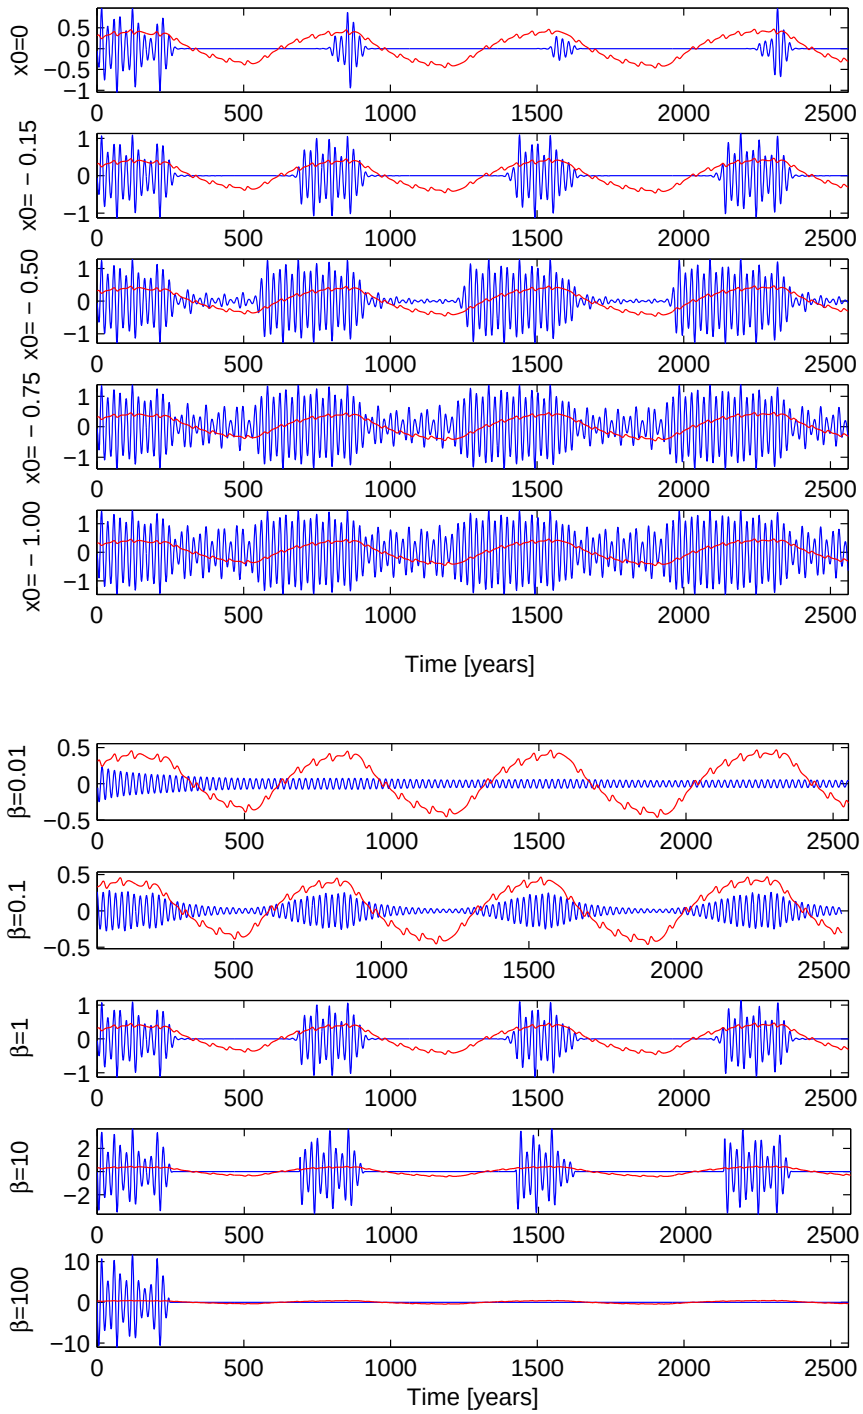

Figure 4.6 Output from PST model. $\mathbb{U}$ is presented in blue and $20 \times \mathbb{Z}$ is shown in red for default initial conditions and standard parameters except for the $x0$ and β parameter. In the upper set of five figures, and the $x0$ parameter is varied $x0$=[0, −0.15, −0.50, −0.75, −1.0]. In the lower set of five figures and the β parameter is varied β=[0.01, 0.1, 1, 10, 100].

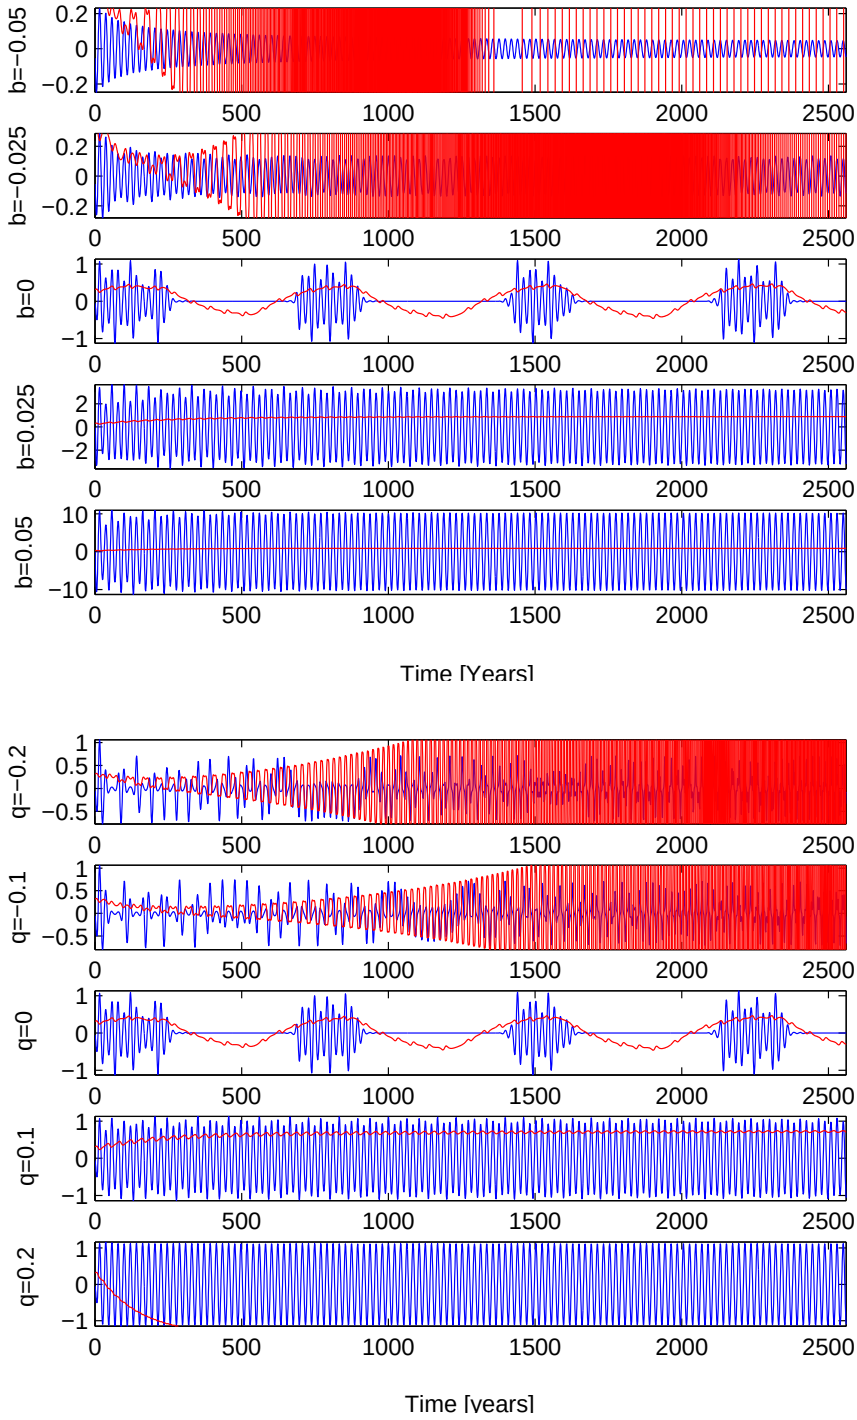

Figure 4.7 Output from PST model. $\mathbb{U}$ is presented in blue and $20 \times \mathbb{Z}$ in red for default initial conditions and standard parameters except for the b and q parameter. In the upper set of five figures the b parameter is varied $b = [-0.05, -0.025, 0.0, 0.025, 0.05]$. In the lower set of five figures the q parameter is varied: $q = [-0.2, -0.1, 0.0, 0.1, 0.2]$.

small variations of the b and q significantly change the output. Sensible output is found only for $b = q = 0$.

Parameter w sets the frequency of the solar cycle and is a fixed value. The period of the model is set to correspond to 11 years. A sketch of a modified PST system with a modulated phase appears in Appendix C.

The ranges of parameters shown in all the figures of this section were chosen to illustrate the transition away from a behavior reminiscent of the solar cycle. This gives some indication on where in parameter space our effort to find optimal parameters for the solar cycle should be focused.

4.2 Data processing

The data used in this work is the yearly average Sunspot number from the (SIDC) *Solar Influences Data Analysis Center* (Clette et al. 2007). The monthly Sunspot number and the 13-month running mean Sunspot number are also provided by SIDC.

A rough estimate of the noise in the data would be the standard deviation of the difference between the unsmoothed and smoothed monthly Sunspot number over the last three centuries. This gives a noise level below 7%. Similar numbers are obtained from the comparison of yearly data and the monthly value from June or July. Comparison between the yearly and daily data would is expected to give a larger number. A conclusion would be that the method should be able to operate when the noise level is at least 10%.

The solar cycle has a mean period of 11 years, but it is possible to define it as 22-year oscillation with positive and negative signs for odd and even cycles. This new period is

related to the 22-year *Hale magnetic polarity cycle* (Hale et al. 1919; Usoskin 2008) and we will refer to this as a *Bracewell construction* (Bracewell 1953). The last 24 cycles start at the cycle that begun in 1755, see Table 4.1.

Cycle #	Start date	End date	Sign	Assign
1	March 1755	June 1766	-	1766: +
2	June 1766	June 1775	+	1775: -
3	June 1775	September 1784	-	1784: -
4	September 1784	May 1798	+	1798: -
5	May 1798	December 1810	-	1810: -
6	December 1810	May 1823	+	1823: -
7	May 1823	November 1833	-	1833: -
8	November 1833	July 1843	+	1843: +
9	July 1843	December 1855	-	1855: -
10	December 1855	March 1867	+	1867: -
11	March 1867	December 1878	-	1878: -
12	December 1878	March 1890	+	1890: -
13	March 1890	February 1902	-	1902: +
14	February 1902	August 1913	+	1913: +
15	August 1913	August 1923	-	1923: -
16	August 1923	September 1933	+	1933: +
17	September 1933	February 1944	-	1944: +
18	February 1944	April 1954	+	1954: -
19	April 1954	October 1964	-	1964: -
20	October 1964	June 1976	+	1976: -
21	June 1976	September 1986	-	1986: -
22	September 1986	May 1996	+	1996: -
23	May 1996	November 2009	-	2009: -
24	November 2009		+	

Table 4.1 Transition between the last 24 solar cycles. The sign of every other cycle is changed to construct a 22-year cycle. The sign assigned to the year of transition is noted in the last column, see text.

The $\mathbb{U}$ and $\mathbb{V}$ components of the PST model have alternating cycles. Note that the model generates a 22-year cycle without any transformation. Thus the $\mathbb{U}$ component should in our case represent the 22-year solar Bracewell cycle.

Figure 4.8 presents a comparison between the model output and the different representations of the data. The output from the PST model is shown in the top panel. This

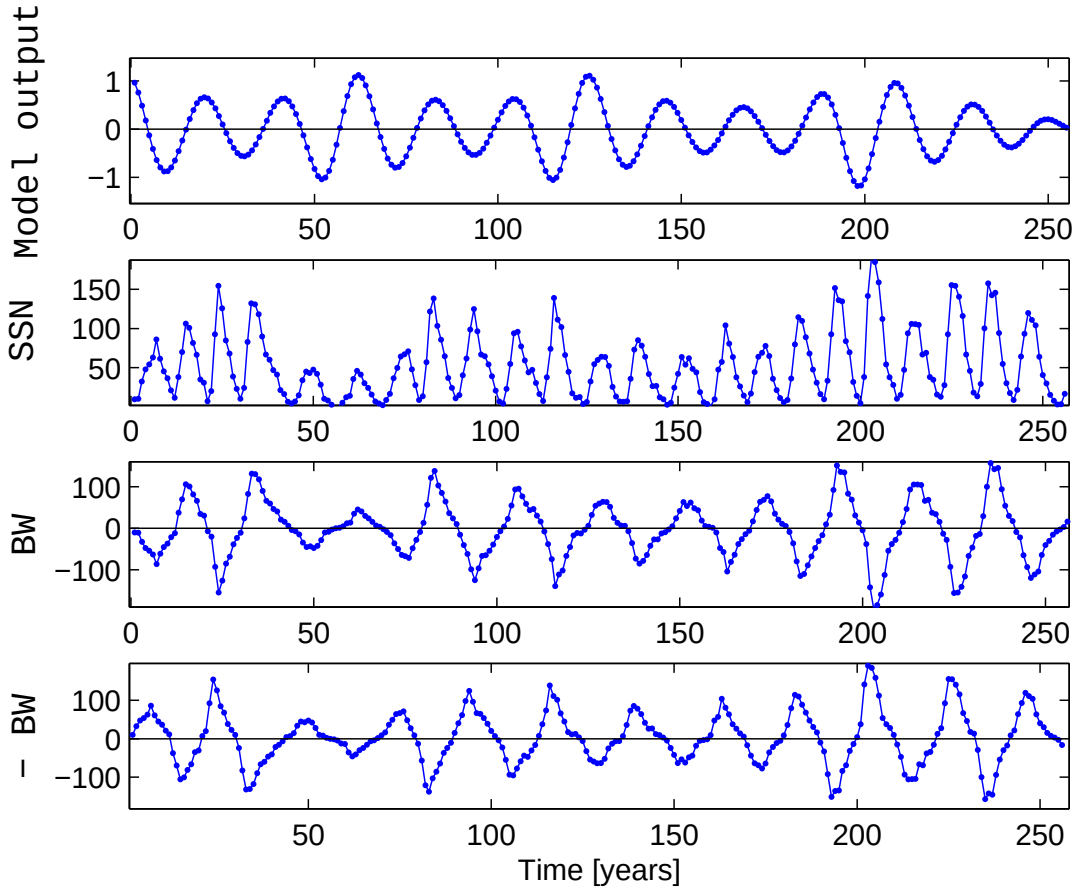

Figure 4.8 Model output $\mathbb{U}$ variable for standard parameters (top panel). Yearly average Sunspot number from 1755-2010 (second panel). Both conventions for alternating signs for odd and even cycles – the Bracewell 22-ear cycle – are shown in the last two panels.

sequence corresponds to the first 256 values used in Figure 4.2. The Sunspot number is shown in the second panel and the last two panels are the 22-year Bracewell cycles with alternating signs. The only difference between the last two data sequences is the sign for odd and even cycles.

In order to create the 22-year cycle we manually change the sign of every other cycle. However this creates some ambiguity for the year of transition between two cycles i.e. when one ends and the other starts. Table 4.1 gives a list of the start and end of the last

24 solar cycles. The official list can be found at:

`ftp://ftp.ngdc.noaa.gov/STP/SOLAR_DATA/SUNSPOT_NUMBERS/maxmin.new.`

In this work we have devised a convention for the sign of the year of transition. It is determined by the cycle that accounts for most months of that year.

For example, cycle 3 ended in September 1784. This provides the longest contribution (9 months) from cycle 3 (an odd cycle) and this in turn gives the year of transition (1784) a negative sign. It would be equally valid to assign positive values to odd cycles and thus we have in fact two different Bracewell cycles generate by this process. Preliminary studies show that the choice of sign convention does not appear to significantly affect the result. As a precaution, estimates should be made with both sequences to confirm consistency.

4.3 Application of Shadowing Filter to the PST model

The first twin experiment of this chapter is similar to the setup outlined in Section 3.4. At this point we generate trajectories of the PST model of length $N = 256$ to match the length of the actual solar cycle record. We use a fixed set of parameter values: $[\beta,x0,w,a,b,q,d,c]$= [1.1, −0.15, 2, 0.7, 0, 0, 0.03, −0.38] and initial conditions: $\mathbb{U}_0$,$\mathbb{V}_0$,$\mathbb{X}_0$, $\mathbb{Y}_0$,$\mathbb{Z}_0$= [0.9614, 0.5235, 0.9292, −0.4947, 0.0079].

These values are slightly different then the standard values but generate similar result. 16 sets of observations are generated for different realizations of the added perturbations. The convergence of $\mathcal{M}$ as a function of gradient descent iterations is seen in

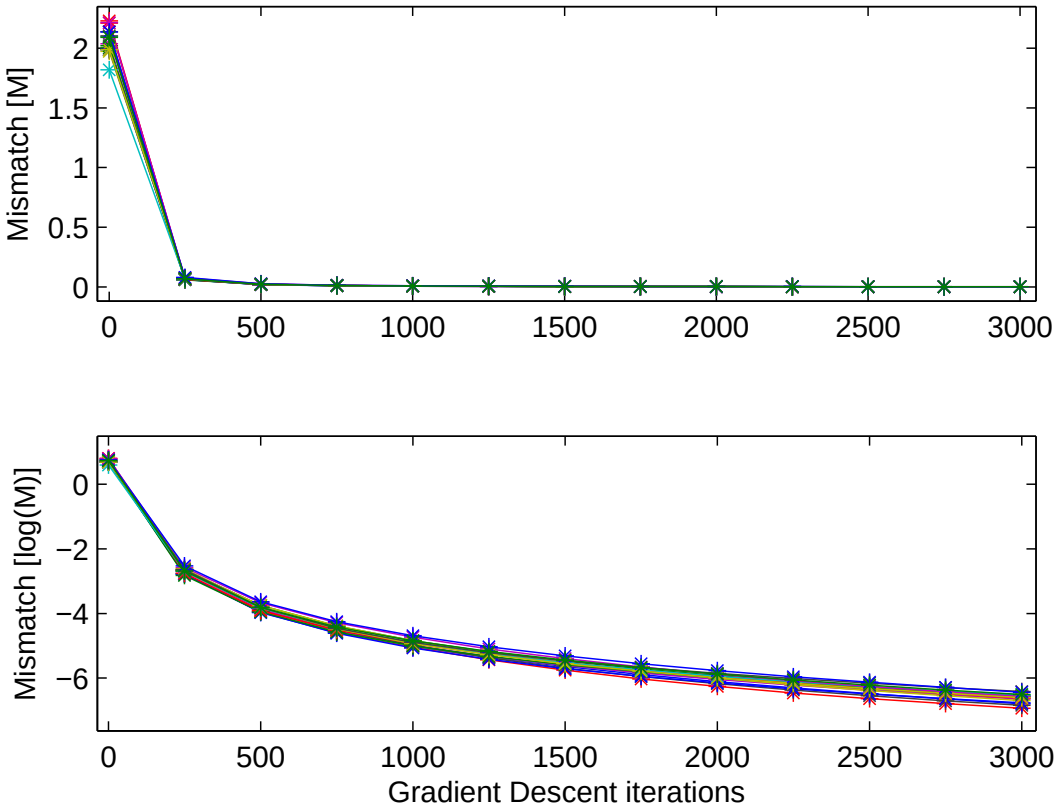

Figure 4.9 The mismatch $\mathcal{M}$ as a function of gradient descent iterations (top panel) for 16 sets of observations generated with different random seeds at a net noise level of 20%. Lower panel is shown in logarithmic scale.

Figure 4.9. The iteration step size for the Euler approximation of the gradient descent is $\Delta\tau = 51.2 \times 10^{-2}$. The value was chosen to guarantee convergence at the expense of effectiveness. The perturbations used to generate the observations have a standard deviation of 25%. The net value of the standard deviation for the 256 generated observations turned out to be closer to 20%.

Figure 4.10 shows the estimate for all of the 16 sets of observations (in red) obtained from the observations (black) generated from the "true state" (blue).

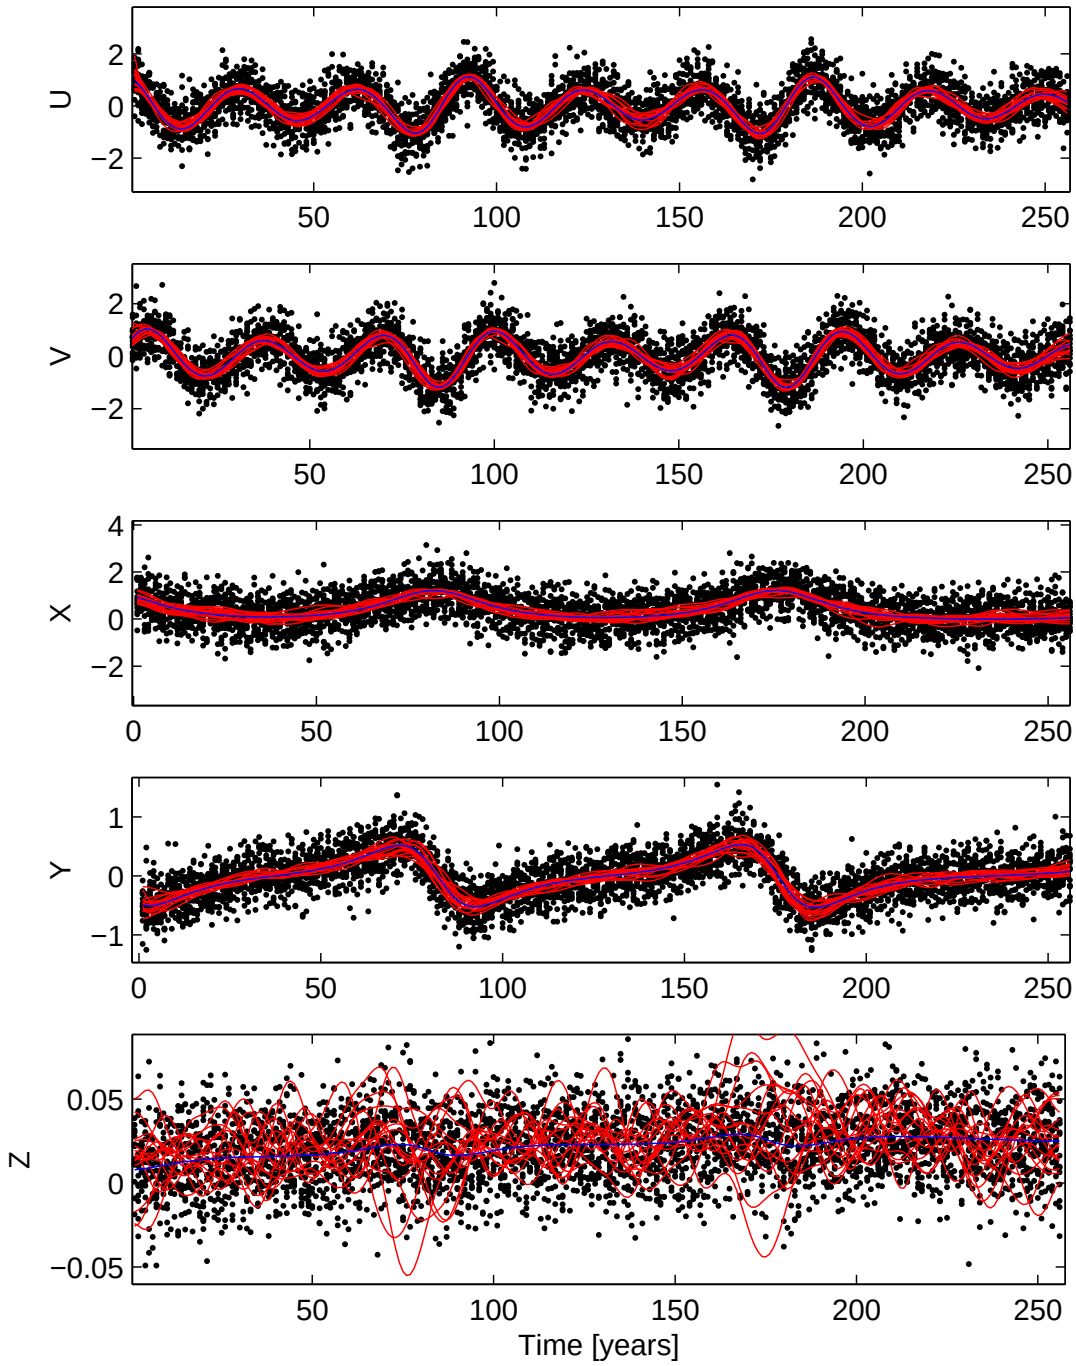

Figure 4.10 From the true model state (blue line) 16 sets of the synthetic observations (black dots) are generated using the PST system with a net noise level of 20%. Following the outlined procedure 16 estimates (red lines) are obtained after 3000 gradient descent iterations. In this case, the parameters of the model are known. Note that the error in the $\mathbb{Z}$ is at times larger than the noise level.

The error (distance between the estimate and the true state) in the $\mathbb{Z}$ component is at time larger than the noise level. The absolute error for the five components is presented in Figure 4.11 as a function of gradient descent iterations. The error for the $\mathbb{Z}$ grows, but remains small in comparison to the other components. Note that the gradient descent is minimizing the total mismatch $\mathcal{M}$ and not component by component. In this case the overall values decrease at the expense of the estimate in the $\mathbb{Z}$ component. This is problematic since we expect that accurate estimate of $\mathbb{Z}$ to be key to the study of the onset of a grand minimum as discussed in Section 4.1.

The absolute error in the $\mathbb{Z}$ component grows for at least 10,000 gradient descent iterations but still remains lower in amplitude than for the other components. After 10,000 iterations the mean error in the $\mathbb{Z}$ component in relation to its own amplitude is larger than the noise level.

Growing error in $\mathbb{Z}$ is also found at the 5% noise level, see Figure 4.12. This indicates that this is not a noise level issue and we explore to what extent the amplitude of the different components has an effect on the answer.

To study the effect of low amplitude of the $\mathbb{Z}$ component we want to introduce a linear scaling. A technical outline of the linear scaling is shown in Section 4.4 and continued exploration of the PST system is found in Section 4.5.

4.4 Linear Scaling of low amplitude component

The long term variation for the five components of the PST system shows that the first three components, $\mathbb{U}$, $\mathbb{V}$ and $\mathbb{X}$ are in the $[-1.2, +1.2]$ range with a signal amplitude close

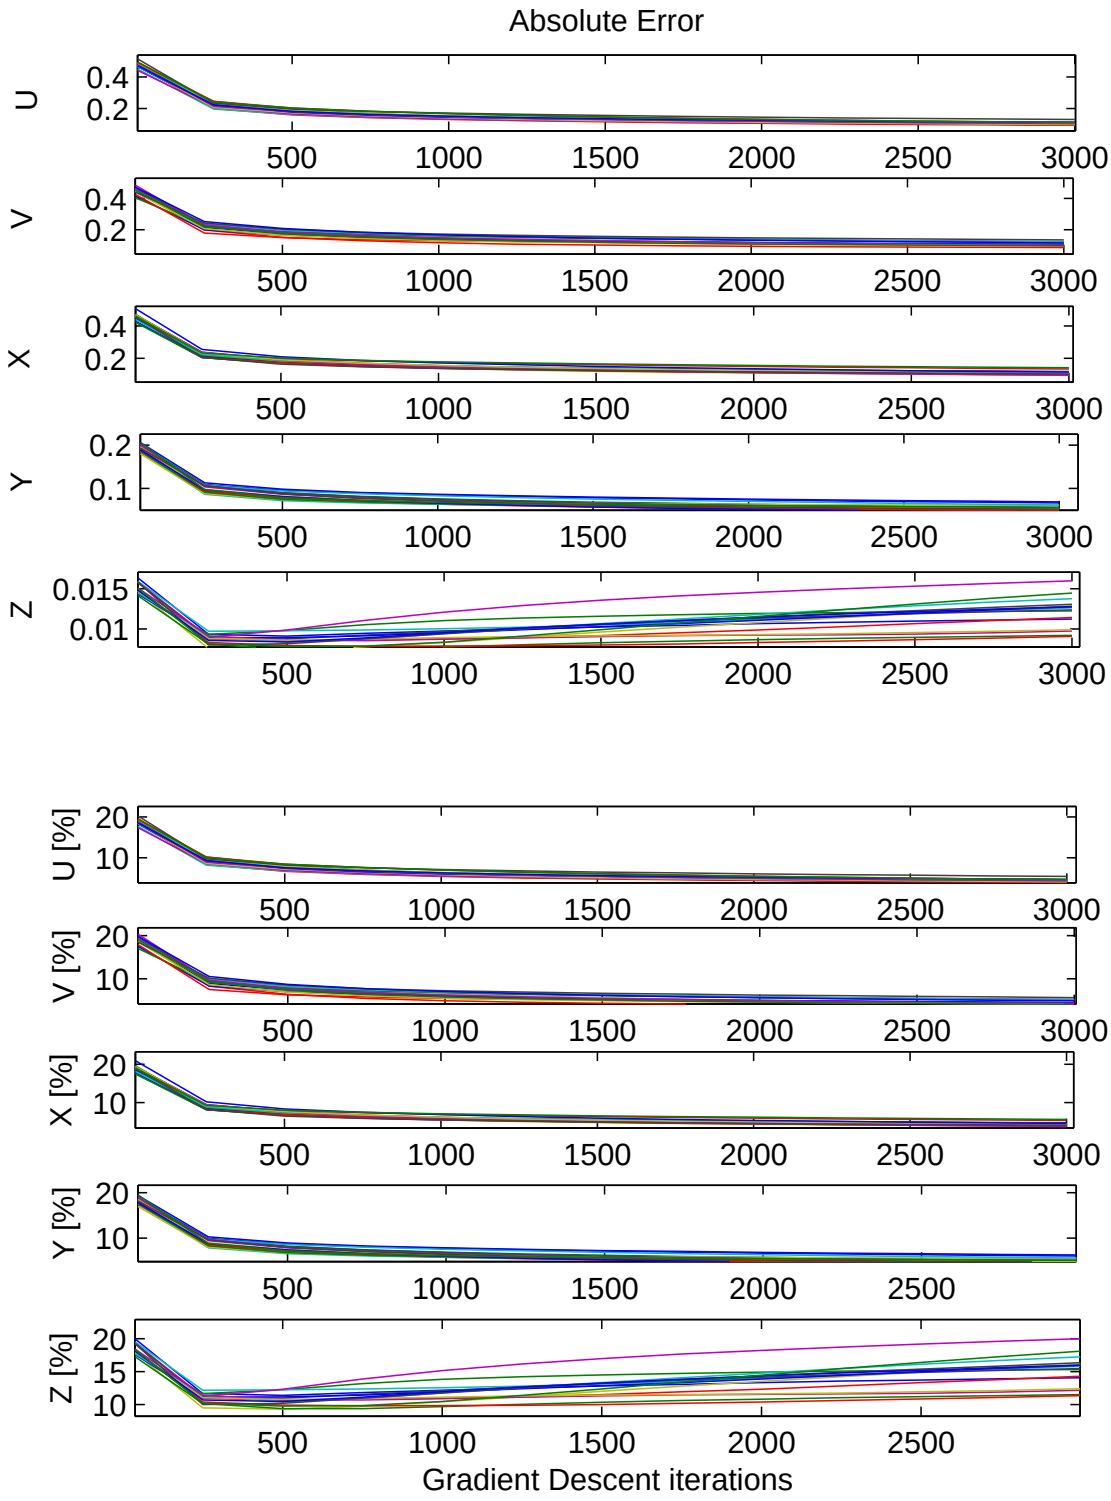

Figure 4.11 Absolute error as a function of gradient descent iterations for 16 realizations of the observations of Figure 4.10. Shown are two sets, each corresponding to the five components of the PST model. The lower set of five is shown in percent. The error of the $\mathbb{Z}$ is small when compared to the other components (see upper set) but very large relative to its own amplitude (see lower set). The initial noise level is 20%.

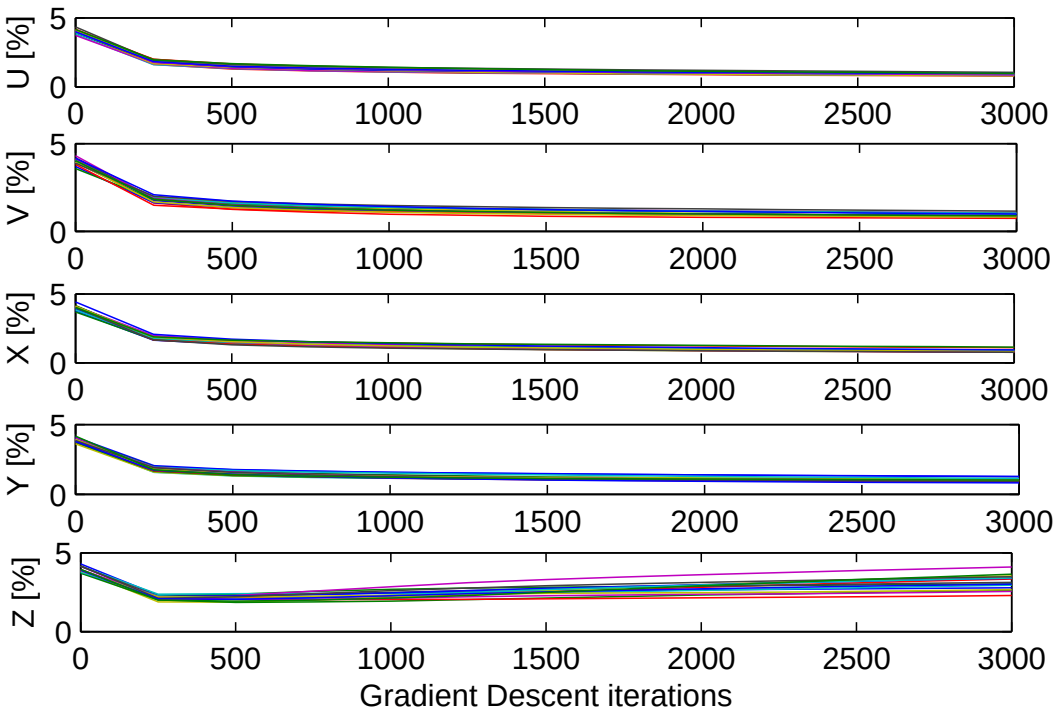

Figure 4.12 Absolute error as a function of gradient descent iterations for 16 realizations of the observations of Figure 4.10. The error in the $\mathbb{Z}$ component keeps growing for at least 10,000 iterations, but remains lower in amplitude than the other components.

to 2.4, see Figure 4.2. The amplitude of the $\mathbb{Y}$ component is about a factor two lower with a range of $[-0.55, 0.55]$ and an amplitude of 1.1. The range for values of the $\mathbb{Z}$ component is $[-0.04, 0.04]$ and an amplitude of 0.08. It is worth noticing that only positive values of $\mathbb{Z}$ correspond to normal solar activity and this would set the effective amplitude lower by a factor two. This would give the $\mathbb{Z}$ component an amplitude of 0.04, a factor 60 lower than the amplitude of the $\mathbb{U}$ component – corresponding to the observations.

We have the PST system:

$$\begin{aligned}
\dot{\mathbb{U}} =& -(\mathbb{U}^2+\mathbb{V}^2)\mathbb{U} - w\mathbb{V} + \beta(\mathbb{X}-x_0)\mathbb{U} \\
\dot{\mathbb{V}} =& -(\mathbb{U}^2+\mathbb{V}^2)\mathbb{V} + w\mathbb{U} + \beta(\mathbb{X}-x_0)\mathbb{V} \\
\dot{\mathbb{X}} =& \mathbb{Y} \\
\dot{\mathbb{Y}} =& a\mathbb{X} - \mathbb{X}^3 - b\mathbb{Y} + \mathbb{Z} + q\mathbb{U}^2 \\
\dot{\mathbb{Z}} =& -d(\mathbb{Z} - c\mathbb{X}(\mathbb{X}^2-1))
\end{aligned} \tag{4.3}$$

and make a linear coordinate transformation:

$$\begin{aligned}
\mathbb{Z}_s =& s_{\mathbb{Z}}\mathbb{Z} \\
\dot{\mathbb{Z}} = \frac{d\mathbb{Z}}{dt} =& \frac{d\mathbb{Z}}{d\mathbb{Z}_s}\frac{d\mathbb{Z}_s}{dt} = \frac{1}{s_{\mathbb{Z}}}\frac{d\mathbb{Z}_s}{dt}
\end{aligned} \tag{4.4}$$

where $\mathbb{U}_s = \mathbb{U}$, $\mathbb{V}_s = \mathbb{V}$,$\mathbb{X}_s = \mathbb{X}$ and $\mathbb{Y}_s = \mathbb{Y}$ remain unscaled for now. This gives us a scaled system:

$$\begin{aligned}
\dot{\mathbb{U}}_s =& -\left(\mathbb{U}_s^2+\mathbb{V}_s^2\right)\mathbb{U}_s - \omega\mathbb{V}_s + \beta_0\left(\mathbb{X}_s - x_0\right)\mathbb{U}_s \\
\dot{\mathbb{V}}_s =& -(\mathbb{U}_s^2+\mathbb{V}_s^2)\mathbb{V}_s + \omega\mathbb{U}_s + \beta_0\left(\mathbb{X}_s - x_0\right)\mathbb{V}_s \\
\dot{\mathbb{X}}_s =& \mathbb{Y}_s \\
\dot{\mathbb{Y}}_s =& a\mathbb{X}_s - \mathbb{X}_s^3 - b\mathbb{Y}_s + \left[\frac{\mathbb{Z}_s}{s_{\mathbb{Z}}}\right] + q\mathbb{U}_s^2 \\
\dot{\mathbb{Z}}_s =& s_{\mathbb{Z}}\left(-d\left(\left[\frac{\mathbb{Z}_s}{s_{\mathbb{Z}}}\right] - c\mathbb{X}_s\left(\mathbb{X}_s^2-1\right)\right)\right).
\end{aligned} \tag{4.5}$$

We generate two data sets of 512 values to study the effect of the scaling. The first set is generated using the unscaled PST system with initial conditions [$\mathbb{U}_0$, $\mathbb{V}_0$, $\mathbb{X}_0$, $\mathbb{Y}_0$, $\mathbb{Z}_0$]= [0.9614, 0.5235, 0.9292, −0.4947, 0.0079], and parameters [a, b, q, d, c, $\beta 0$, x_0, ω]=

[0.7, 0, 0, 0.03, −0.38, 1.1, −0.15, 2]. The second data set is generated with the scaled system 4.5 with the same parameters and initial conditions $[\mathbb{U}_{s,0}, \mathbb{V}_{s,0}, \mathbb{X}_{s,0}, \mathbb{Y}_{s,0}, \mathbb{Z}_{s,0}] = [\mathbb{U}_0, \mathbb{V}_0, \mathbb{X}_0, \mathbb{Y}_0, s_\mathbb{Z} \times \mathbb{Z}_0]$,

where the initial condition $\mathbb{Z}_{s,0}$ for the $\mathbb{Z}_s$ component is scaled with $s_\mathbb{Z}$.

s_z	$\mathbb{U}_s/\mathbb{U}$	$\mathbb{V}_s/\mathbb{V}$	$\mathbb{X}_s/\mathbb{X}$	$\mathbb{Y}_s/\mathbb{Y}$	$\mathbb{Z}_s/\mathbb{Z}$
$1+\epsilon$	$1 \pm 1.3 \times 10^{-13}$	$1 \pm 1.3 \times 10^{-13}$	$1 \pm 1.7 \times 10^{-13}$	$1 \pm 4.0 \times 10^{-12}$	$s_z + 1.1 \times 10^{-12}$
10	$1 \pm 1.4 \times 10^{-13}$	$1 \pm 1.5 \times 10^{-13}$	$1 \pm 1.1 \times 10^{-12}$	$1 \pm 6.4 \times 10^{-12}$	$s_z + 1.7 \times 10^{-11}$
70	$1 \pm 1.4 \times 10^{-13}$	$1 \pm 1.5 \times 10^{-13}$	$1 \pm 1.1 \times 10^{-12}$	$1 \pm 6.4 \times 10^{-12}$	$s_z + 1.7 \times 10^{-11}$
999	$1 \pm 1.6 \times 10^{-12}$	$1 \pm 1.6 \times 10^{-12}$	$1 \pm 1.7 \times 10^{-11}$	$1 \pm 6.3 \times 10^{-11}$	$s_z \pm 1.6 \times 10^{-8}$

Table 4.2 Effect of scaling the $\mathbb{Z}$ component of PST with a factor of $s_\mathbb{Z}$ where $\epsilon = 2.2204 \times 10^{-16}$ is the machine epsilon (smallest available numerical value). Shown is the mean and standard divination of the ratio between the scaled and the unscaled system (see text) for the five components of the model. The dynamics are almost identical for $N = 512$ model output values. The ratio of output for the scaled and unscaled $\mathbb{Z}$ component are close (less than 10^{-8}) to be a factor of $s_\mathbb{Z}$ from each other.

The ratio of the two outputs is seen in Table 4.2. Shown is the mean value and the standard deviation. The first row, $s_\mathbb{Z} = 1 + \epsilon$ exhibits the effect of adding the smallest possible scaling using the *machine epsilon* ϵ. This gives us an estimate of the size of possible numerical effect. The output for the first four components is identical up to numerical effect. The ratio of the outputs from the scale and the unscaled system, for the filth component, is very close ($< 10^{-8}$) to be the scale factor $s_\mathbb{Z}$. The distance, component by component, between the scaled and the unscaled data sets using $s_\mathbb{Z} = 999$ remains close (below 0.1% of the amplitude) for at least $18,000$ years (in units of the model). It can be concluded that the linear scaling does not affect the dynamics of the system, only the amplitude.

The experiment outlined in Section 4.3 is performed again with $s_\mathbb{Z} = 10$ and $s_\mathbb{Z} = 70$.

Figure 4.13 shows the absolute error as a function of gradient descent iterations for the $s_{\mathbb{Z}} = 10$ case. All variables (including $\mathbb{Z}$) have a decreasing absolute error, i.e. all variables are converging towards the "truth". This figure should be compared to previous case Figure 4.11 (corresponding to $s_{\mathbb{Z}} = 1$). Figure 4.14 shows the 16 estimates (red) and 16 sets of observations (black) generated from the "true state" (blue), for $s_{\mathbb{Z}} = 10$. Figure 4.15 show the 16 estimates for three scale-factors $s_{\mathbb{Z}} = [1, 10, 70]$. We find convergence in all variables (including $\mathbb{Z}$) when they are all of comparable amplitude.

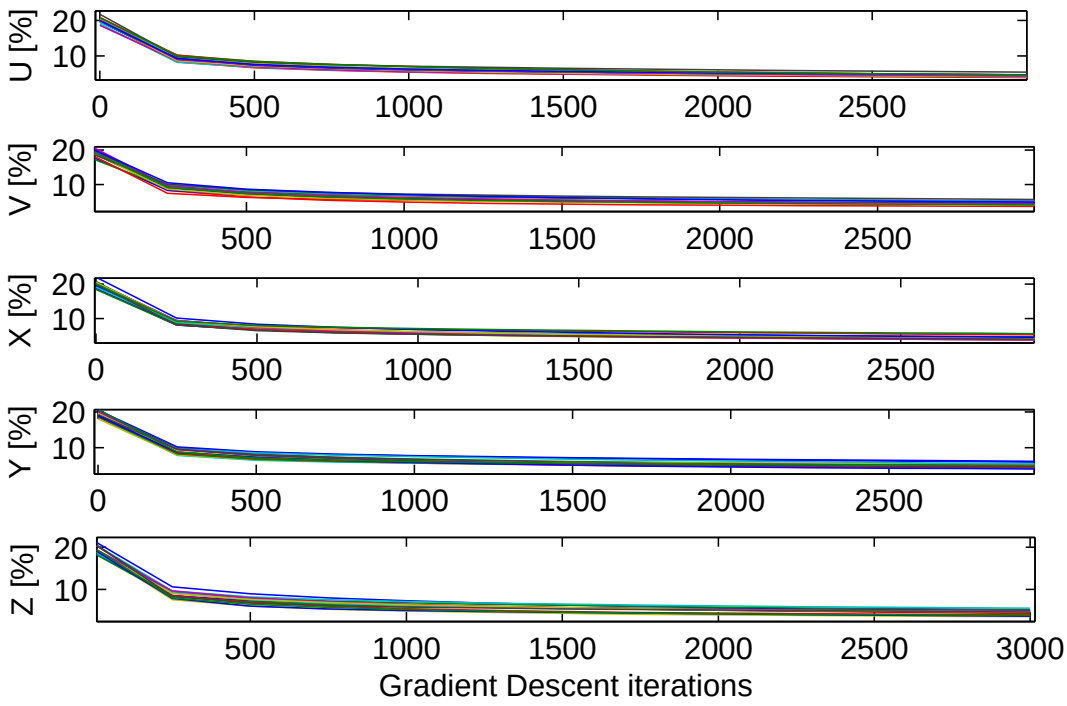

Figure 4.13 Absolute error (in percent) as a function of gradient descent iterations for 16 realizations of the observations of Figure 4.10. The $\mathbb{Z}$ component is scaled by a factor $s_{\mathbb{Z}} = 10$. The procedure reduces the error and the estimate convergence towards the "true" state for all components including $\mathbb{Z}$.

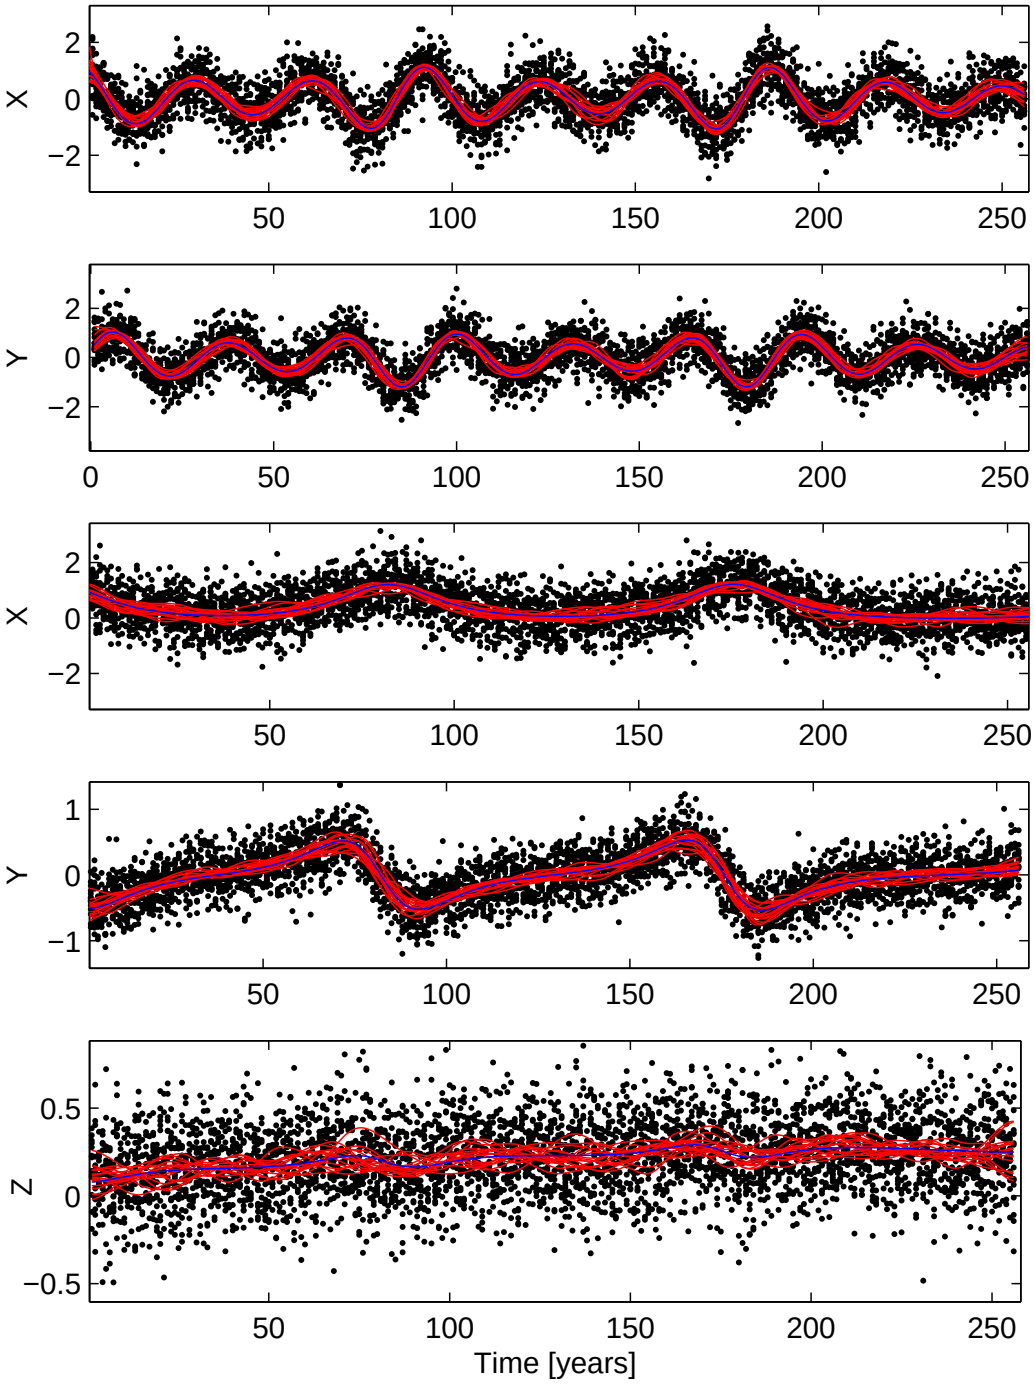

Figure 4.14 From the true model state (blue line) 16 sets of the synthetic observations (black dots) are generated using the PST system with a net noise level of 20%. The $\mathbb{Z}$ component is scaled by a factor of $s_{\mathbb{Z}} = 10$. Following the outline procedure 16 estimates (red lines) are obtained after 3000 gradient descent iterations. In this case, the parameters of the model are known. The scaled system has had its noise reduced in all components. The initial noise level is 20%.

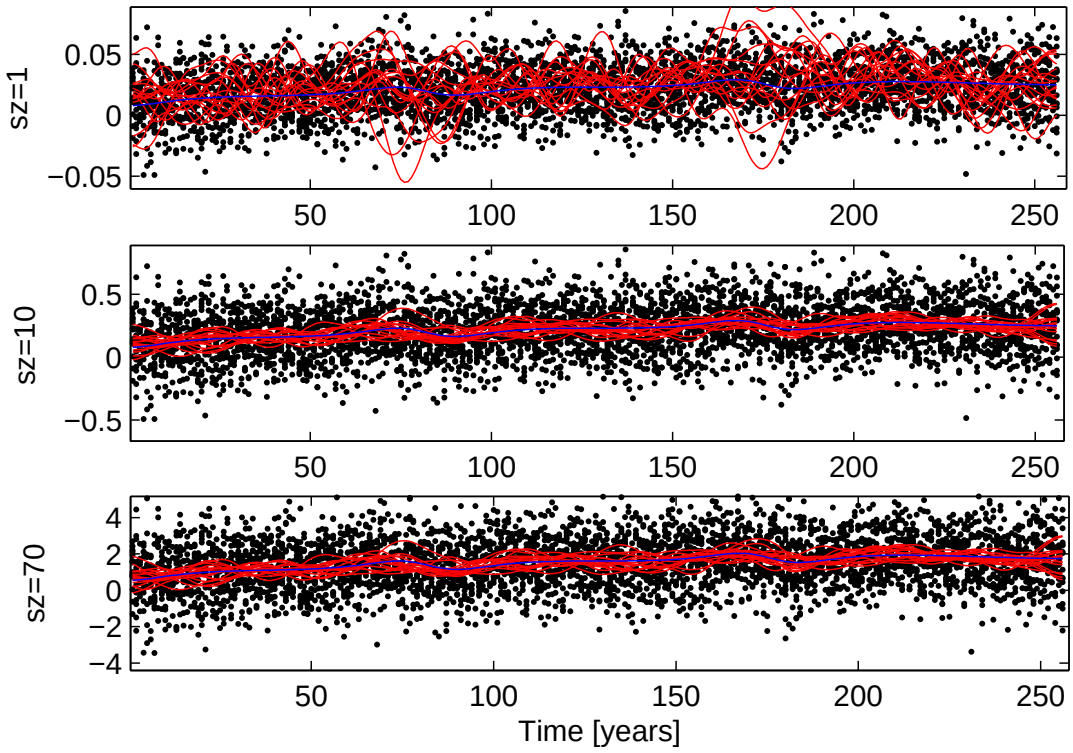

Figure 4.15 From the true model state (blue line) 16 sets of the synthetic observations (black dots) are generated using the PST system with a net noise level of 20%. Shown is the $\mathbb{Z}$ component for scale-factors $s_{\mathbb{Z}} = [1, 10, 70]$. Following the outline procedure 16 estimates (red lines) are obtained after 3000 gradient descent iterations. In this case, the parameters of the model are known.

4.5 Assimilation of unobserved variables for the PST system

The second twin experiment of this chapter is similar to the setup outlined in Section 3.6. A true state of the system is explicitly generated using the scaled PST system of Section 4.4. Perturbations are added to the $\mathbb{U}$ component to generate synthetic observations that

in this case correspond to the modified 22 year Bracewell cycle. The other variables ($\mathbb{V}$, $\mathbb{X}$, $\mathbb{Y}$ and $\mathbb{Z}$) are considered to be *unobserved* and we generate uniformly distributed guesses. We want to determine if it is possible to recover the unobserved variables in an ideal case where the mathematical structure and the parameters of the system are known.

All experiments in this section have the same base setup, outlined in Section 3.6, with 2000+1000 gradient descent iterations where we "hold on" to the observations for 2000 iterations then we let go of the observations for 1000 iteration, see Figure 4.16.

The same set of parameters $[\beta,\ x0,\ w,\ a,\ b,\ q,\ d,\ c] = [1.1,\ -0.15,\ 2,\ 0.7,\ 0,\ 0,\ 0.03,\ -0.38]$ are used for all cases. Two initial conditions are used to generate reference trajectories that we can use to construct synthetic observations: Set A is initialized with: $\mathbb{U}_0, \mathbb{V}_0, \mathbb{X}_0, \mathbb{Y}_0, \mathbb{Z}_0 = [0.9614, 0.5235, 0.9292, -0.4947, s_{\mathbb{Z}} \times 0.0079]$ and Set B with: $[0.2554, 0.8418, 0.5675, 0.0845, s_{\mathbb{Z}} \times 0.1497]$ where the initial condition for $\mathbb{Z}$ is scaled with one of the three linear scale-factor $s_{\mathbb{Z}} = [1, 10, 70]$. See Section 4.4 for more on the linear scaling.

These parameters and initial conditions where chosen to generate 256 observations followed by an inactive phase (Set A) or an active phase (Set B). This was done to evaluate predictions under normal conditions and at the onset of grand minimum, see Figure 4.17.

Figure 4.18 shows the results of only observing the $\mathbb{U}$ using Set A and the standard unscaled PST system ($s_{\mathbb{Z}} = 1$) with a noise level of 1%. The error of the final estimate for the $\mathbb{Z}$ component is very large, about 50% and it does not converge towards the "truth", see Figure 4.19.

The low amplitude of the $\mathbb{Z}$ has been discussed as a possible issue above and the same

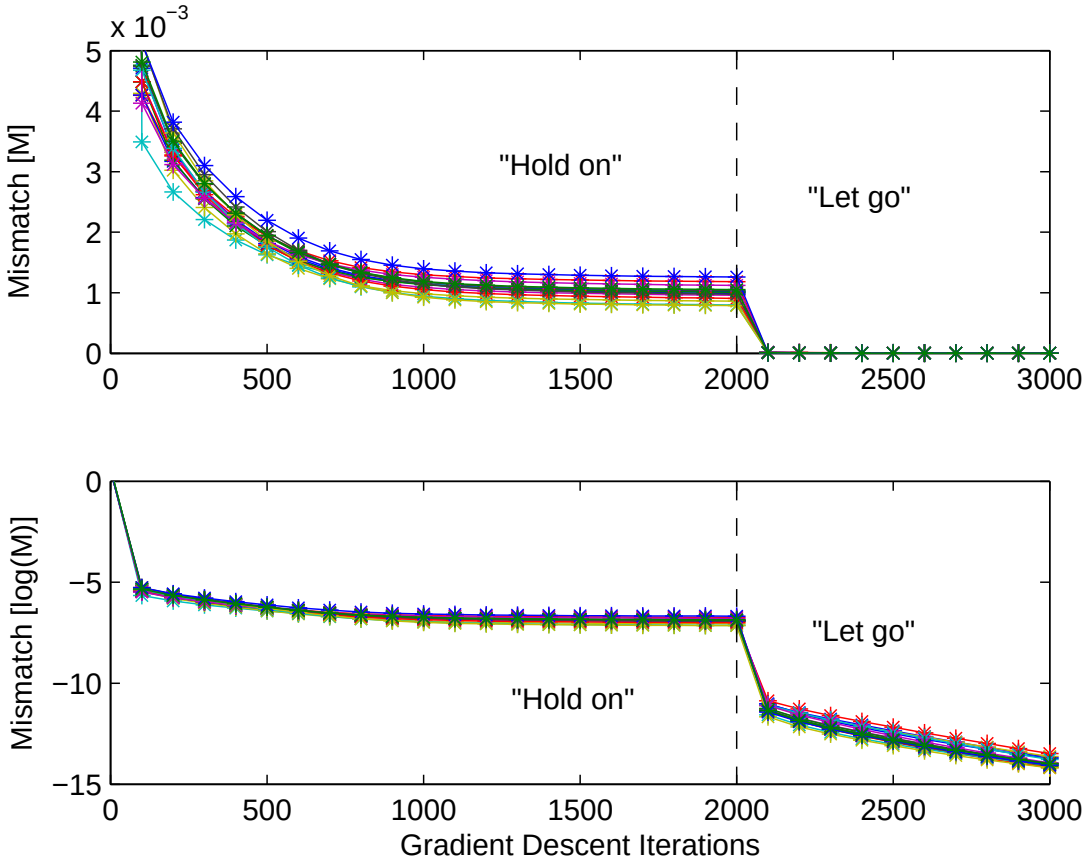

Figure 4.16 The mismatch $\mathcal{M}$ as a function of gradient descent iterations for 16 sets of observations (Set A) and a linear scale-factor $s_{\mathbb{Z}} = 1$. First 2000 iterations in the "hold on" stage, and where the observations are held fixed, followed by 1000 iteration in the "let go" stage where we only minimize $\mathcal{M}$ without any additional constraints. The $\mathbb{U}$ component is observed with a noise level of 1% and four other components are unobserved. Lower panel in shown in natural logarithmic scale.

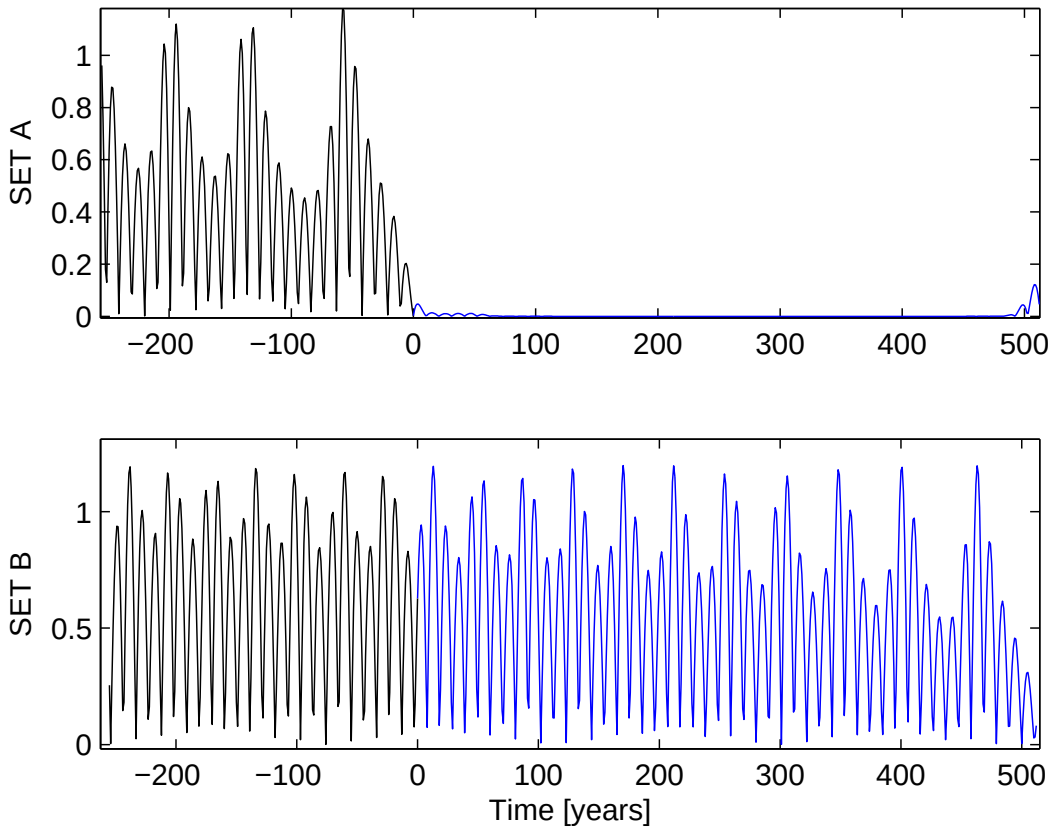

Figure 4.17 The two reference trajectories, with different initial conditions, are used in this chapter. The first 256 values (black) are considered to be "historical values" and the following 512 (blue) are "future values". Using the synthetic observations generated from the first 256 values we want to make predictions starting at time t=0 and evaluate how far prediction is possible, in the ideal case. In the upper panel (Set A) we predict into an *inactive phase*. In the lower panel (Set B) we predict into an *active phase*. The same set of parameter are used for both cases.

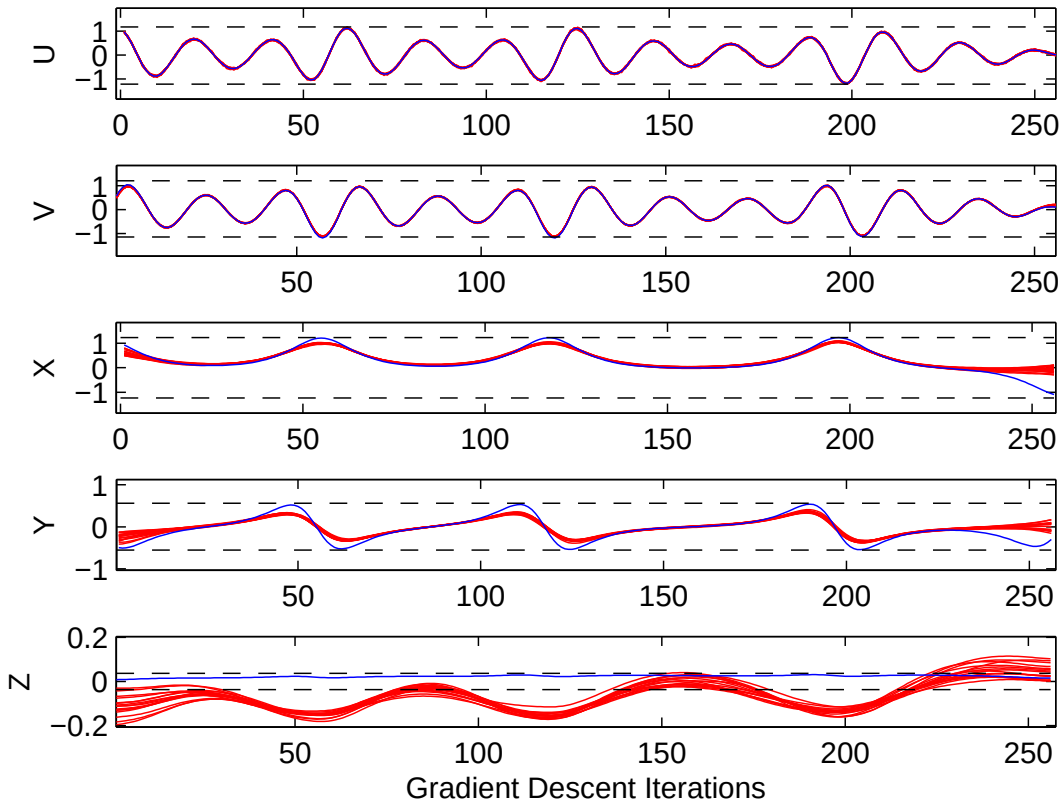

Figure 4.18 State estimation with only one observed variable. From a "true" model state (blue) 16 sets of observations (Set A) for the $\mathbb{Z}$, a 1% noise level for linear scale-factor $s_{\mathbb{Z}} = 1$. The random guesses for the $\mathbb{V}$, $\mathbb{X}$, $\mathbb{Y}$ and $\mathbb{Z}$ components are drawn from a uniform range (dashed lines). Starting from observations, the 16 estimates are made after 2000+1000 gradient descent iterations (see text). Most notable is the error of the $\mathbb{Z}$ component that is larger than the range of the guesses for the unobserved variables.

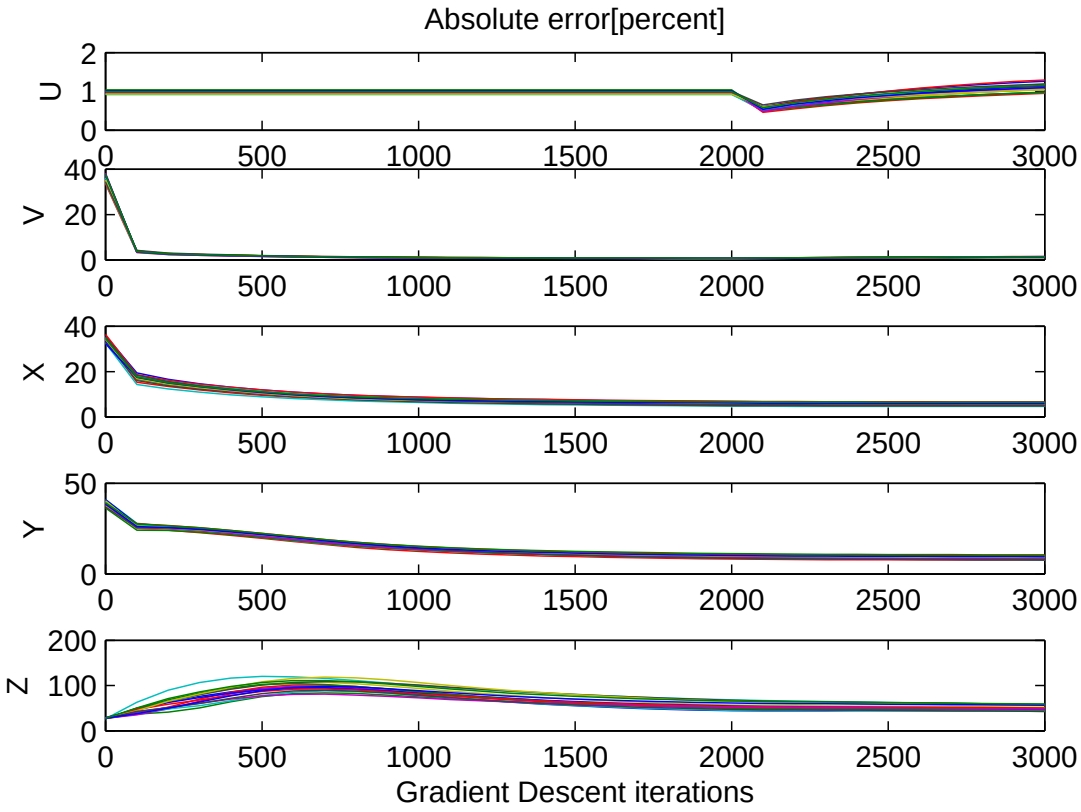

Figure 4.19 Absolute error (in percent) as a function of gradient descent iterations for 16 realizations of observations (Set A) for linear scale-factor $s_{\mathbb{Z}} = 1$ and a noise level of 1%. Note the increase in the error of the $\mathbb{U}$ component during the "let go" stage (from 2000-3000 iterations). The final error for the $\mathbb{Z}$ component is about 50%.

test is performed with $s_{\mathbb{Z}} = 70$. In this case – where the amplitude of the $\mathbb{Z}$ component is comparable to the other variables – the error goes down for the unobserved variables $\mathbb{V}$, $\mathbb{X}$, $\mathbb{Y}$ and notably also for $\mathbb{Z}$. The error grows somewhat for the $\mathbb{U}$ component from 2000-3000 iterations and this clearly illustrates that we are moving away from the observations during the "let go" stage. We will return to the issue of the growing error during the "let go" stage when we discuss Table 4.3 and 4.4. In this case the initial noise level for the $\mathbb{U}$ variable is 1% and grows to about 1.4%. The final estimate for $s_{\mathbb{Z}} = 70$ is shown in

Figure 4.21. The estimates for the $\mathbb{Z}$ component do not converge toward the "truth" (blue line) but the estimate remains within range of the unobserved variables (dashed lines). The final error is about 5%. Note that the ensemble fans out towards the end of the time interval in Figure 4.21. This is most notable in the $\mathbb{Y}$ component from t=200-250. Most values in a sequence are constrained both by past and future values This does not apply to the last values that are only constrained from the past. This might explain this behavior.

We study three scale-factors $s_{\mathbb{Z}}$=$[1, 10, 70]$ and four noise levels nl=[1%, 5%, 10%, 25%] both for Set A and Set B. All these experiments, including the two cases presented above are presented in Table 4.3 and 4.4. The absolute error – ensemble mean and standard deviation – for the 16 estimates are made after both 2000 and 3000 iterations. The first estimate (after 2000 iterations) corresponds to the best estimate that is dynamically consistent with the observations. The last 1000 iterations are in the "let go" stage where we leave the observations behind to go towards a trajectory of the model. The second estimate (after 3000 iterations) corresponds to a dynamically consistent estimate where the constraint on the observations has been removed. Values in bold font, in Table 4.3 and 4.4, indicate instances where the error has grown has grown during the "let go" stage. No clear pattern in found in the distribution of these values. In some instances more iterations during the second stage might eventually give convergence for all variables.

We notice that the error of the estimate of the unobserved variables is larger when the noise level is increased. We also note that the error in the $\mathbb{Z}$ component is very large for the $s_{\mathbb{Z}} = 1$ case, up to about 500%. It is about 1-2 magnitudes lower using $s_{\mathbb{Z}} = [10, 70]$.

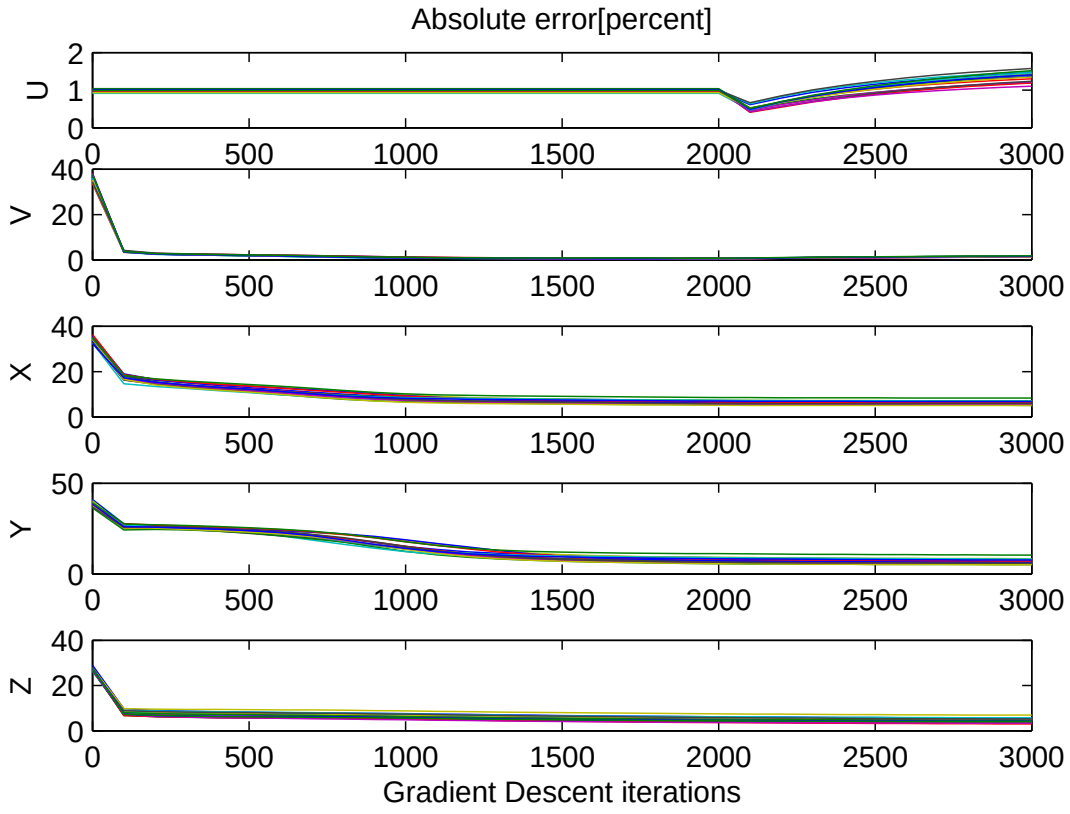

Figure 4.20 Absolute error (in percent) as a function of gradient descent iterations for 16 realizations of observations (Set A) for linear scale-factor $s_{\mathbb{Z}} = 70$ and a noise level of 1%. The errors in the last four components are decreasing. The increase in the error in the $\mathbb{U}$ component to 1.4% is a consequence of "letting go" of the observation when the noise level is low. The final error for the $\mathbb{Z}$ component is about 5%.

The error in the $\mathbb{Y}$ component has large values overall, up to about 70% and we do not see a strong noise reduction during the "let go" stage. It is difficult to differentiate the $s_{\mathbb{Z}} = 10$ and the $s_{\mathbb{Z}} = 70$ case but it is clear that the $s_{\mathbb{Z}} = 1$ case has much larger errors. The error in Set B is somewhat larger than the ones in Set A, especially the error in the $\mathbb{Z}$ component for $s_{\mathbb{Z}} = 10$ and $s_{\mathbb{Z}} = 70$.

The *prediction time*, defined in Section 3.7, for the 16 ensemble members is shown in

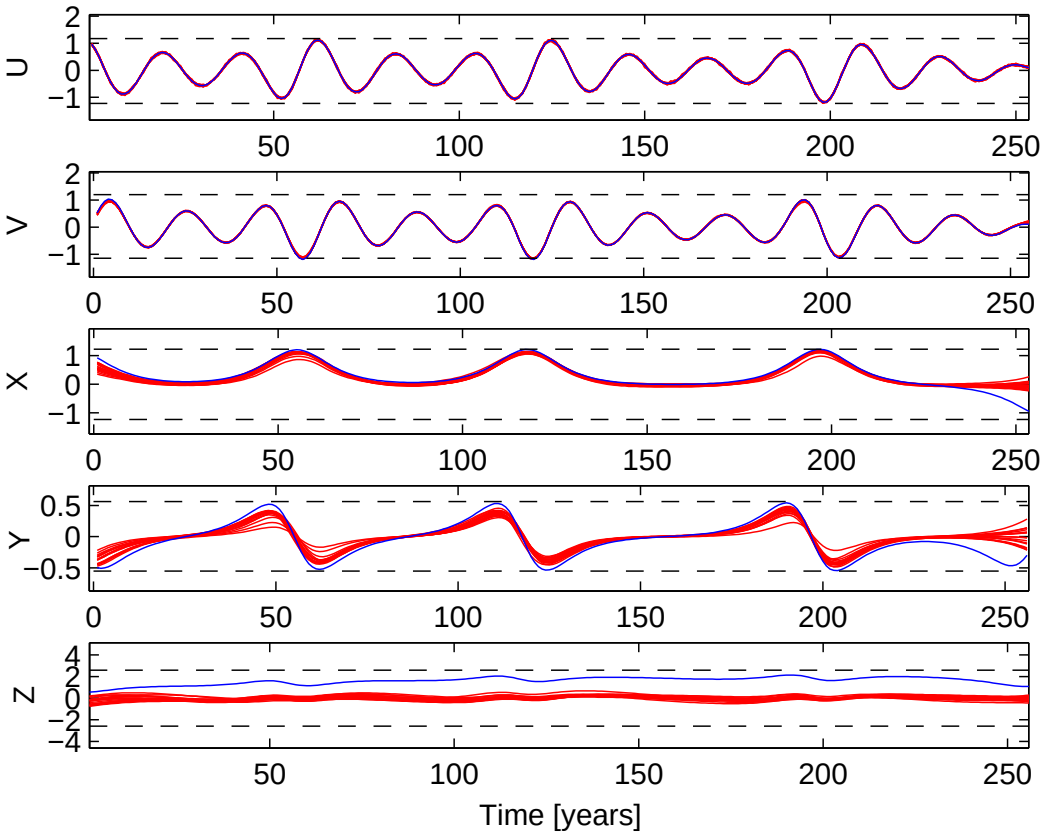

Figure 4.21 State estimation with only one observed variable. From a "true" model state [blue] 16 sets of observations (Set A) for linear scale-factor $s_{\mathbb{Z}} = 70$ and a noise level of 1%. The random guesses for the $\mathbb{V}$, $\mathbb{X}$, $\mathbb{Y}$ and $\mathbb{Z}$ components are drawn from a uniform range (dashed lines). Starting from observations, 16 sets of estimates are made after 2000+1000 gradient descent iterations (see text). The set of 16 estimates forms a tight bundle but fans out towards the end.

$s_{\mathbb{Z}}$	nl	Iterations	$AE_{\mathbb{U}}$	$AE_{\mathbb{V}}$	$AE_{\mathbb{X}}$	$AE_{\mathbb{Y}}$	$AE_{\mathbb{Z}}$
1	1%	2000	0.98 ± 0.04	0.66 ± 0.07	5.88 ± 0.67	$9.53 + 0.90$	$55.13 + 7.98$
1	1%	3000	$\mathbf{1.10 \pm 0.10}$	$\mathbf{1.23 \pm 0.13}$	5.66 ± 0.62	$8.82 + 0.95$	$49.62 + 6.00$
1	5%	2000	4.92 ± 0.18	2.45 ± 0.18	4.88 ± 1.21	$9.35 + 1.78$	$101.74 + 15.72$
1	5%	3000	1.59 ± 0.21	1.61 ± 0.20	4.30 ± 1.23	$8.28 + 1.79$	$75.38 + 15.60$
1	10%	2000	9.84 ± 0.36	4.94 ± 0.29	13.61 ± 1.83	$23.62 + 3.71$	$181.13 + 22.83$
1	10%	3000	7.01 ± 1.26	$\mathbf{7.00 \pm 1.27}$	13.41 ± 1.92	$\mathbf{24.04} + 24.04$	$160.44 + 30.11$
1	25%	2000	24.59 ± 0.89	13.88 ± 2.70	25.88 ± 3.96	$70.57 + 11.13$	$508.05 + 106.06$
1	25%	3000	17.23 ± 1.87	$\mathbf{18.57 \pm 2.77}$	25.05 ± 3.44	$64.56 + 8.13$	$396.76 + 95.94$
10	1%	2000	0.98 ± 0.04	0.58 ± 0.07	5.89 ± 0.70	$7.45 + 1.08$	$4.73 + 0.90$
10	1%	3000	$\mathbf{1.11 \pm 0.09}$	$\mathbf{1.23 \pm 0.13}$	5.70 ± 0.67	$6.81 + 1.03$	$4.15 + 0.76$
10	5%	2000	4.92 ± 0.18	2.45 ± 0.18	4.36 ± 1.46	$6.91 + 1.67$	$5.15 + 1.25$
10	5%	3000	1.30 ± 0.23	1.32 ± 0.27	3.92 ± 1.52	$5.90 + 1.47$	$4.12 + 1.18$
10	10%	2000	9.84 ± 0.36	4.99 ± 0.32	11.97 ± 2.23	$22.75 + 4.46$	$8.62 + 1.84$
10	10%	3000	6.37 ± 2.03	$\mathbf{6.39 + 2.03}$	$\mathbf{13.05 \pm 2.84}$	$\mathbf{23.03 \pm 14.71}$	$7.82 + 2.08$
10	25%	2000	24.59 ± 0.89	13.88 ± 2.74	21.98 ± 4.39	$51.95 + 12.65$	$13.46 + 3.36$
10	25%	3000	17.04 ± 1.70	$\mathbf{18.33 \pm 2.65}$	20.38 ± 3.75	$38.65 + 9.21$	$11.53 + 3.92$
70	1%	2000	0.98 ± 0.04	0.60 ± 0.08	6.44 ± 0.93	$7.66 + 1.50$	$4.99 + 1.07$
70	1%	3000	$\mathbf{1.36 \pm 0.14}$	$\mathbf{1.46 \pm 0.17}$	6.23 ± 0.88	$6.91 + 1.43$	$4.50 + 1.03$
70	5%	2000	4.92 ± 0.18	2.45 ± 0.18	5.35 ± 1.70	$8.19 + 2.24$	$5.33 + 1.12$
70	5%	3000	1.80 ± 0.27	1.82 ± 0.31	4.89 ± 1.72	$6.60 + 1.91$	$4.61 + 1.13$
70	10%	2000	9.82 ± 0.36	4.97 ± 0.32	10.83 ± 2.35	$20.63 + 3.92$	$7.06 + 1.16$
70	10%	3000	$\mathbf{6.49 \pm 1.74}$	$\mathbf{6.49 \pm 1.76}$	$\mathbf{12.34 \pm 2.37}$	$\mathbf{21.48 + 4.50}$	$6.09 + 1.14$
70	25%	2000	24.59 ± 0.89	13.86 ± 2.74	21.39 ± 4.18	$47.32 + 10.37$	$8.83 + 1.20$
70	25%	3000	17.02 ± 1.70	$\mathbf{18.30 \pm 2.66}$	20.43 ± 3.79	$39.10 + 8.45$	$7.80 + 1.20$

Table 4.3 Absolute error in percent – ensemble mean and standard deviation– for the five components of the PST model [$AE_{\mathbb{U}}$, $AE_{\mathbb{V}}$, $AE_{\mathbb{X}}$, $AE_{\mathbb{Y}}$, $AE_{\mathbb{Z}}$]. Shown are four different noise levels nl = [1%, 5%, 10%, 25%] and three different scale factors ($s_{\mathbb{Z}}$)=[1, 10, 70]. The ensemble of 16 estimates (Set A) and the values are shown after both 2000 and 3000 gradient descent iterations. The "hold on" stage ends after 2000 iterations, and the "let go" stage ends after 3000 iterations (see Section 3.6). Values in bold font indicate instances where the error has grown has grown during the "let go" stage.

$s_{\mathbb{Z}}$	nl	Iterations	$AE_{\mathbb{U}}$	$AE_{\mathbb{V}}$	$AE_{\mathbb{X}}$	$AE_{\mathbb{Y}}$	$AE_{\mathbb{Z}}$
1	1%	2000	0.98 ± 0.04	0.76 ± 0.05	2.54 ± 0.16	$9.13 + 0.55$	$24.36 + 6.04$
1	1%	3000	0.97 ± 0.06	$\mathbf{1.00 \pm 0.06}$	$\mathbf{2.91 \pm 0.18}$	$\mathbf{9.19 + 0.57}$	$23.03 + 6.09$
1	5%	2000	4.92 ± 0.18	2.80 ± 0.35	4.27 ± 0.63	$11.95 + 2.18$	$65.17 + 16.63$
1	5%	3000	1.60 ± 0.27	2.04 ± 0.43	4.17 ± 0.74	$11.64 + 2.40$	$56.58 + 15.12$
1	10%	2000	9.84 ± 0.36	6.32 ± 2.00	8.93 ± 1.24	$20.13 + 3.75$	$127.80 + 29.89$
1	10%	3000	3.54 ± 0.78	5.12 ± 2.27	8.12 ± 1.63	$19.32 + 4.28$	$115.13 + 30.47$
1	25%	2000	24.59 ± 0.89	17.54 ± 5.12	34.41 ± 3.67	$76.50 + 10.15$	$501.44 + 99.01$
1	25%	3000	19.68 ± 1.97	$\mathbf{23.19 \pm 4.71}$	33.90 ± 3.37	$\mathbf{70.58 + 10.64}$	$434.99 + 105.15$
10	1%	2000	0.98 ± 0.04	0.68 ± 0.05	2.04 ± 0.12	$7.23 + 0.37$	$26.22 + 1.38$
10	1%	3000	0.95 ± 0.03	$\mathbf{0.96 \pm 0.03}$	$\mathbf{2.73 \pm 0.09}$	$\mathbf{8.88 + 0.24}$	$\mathbf{27.25 + 1.32}$
10	5%	2000	4.92 ± 0.18	2.76 ± 0.34	3.76 ± 0.64	$10.18 + 2.03$	$26.42 + 1.60$
10	5%	3000	1.50 ± 2.21	1.93 ± 0.44	3.73 ± 0.60	$\mathbf{10.63 + 1.66}$	$\mathbf{27.45 + 1.60}$
10	10%	2000	9.84 ± 0.36	6.31 ± 1.98	9.02 ± 1.58	$20.17 + 4.32$	$26.58 + 2.02$
10	10%	3000	3.76 ± 1.21	5.30 ± 2.27	8.44 ± 2.51	$19.14 + 4.79$	$\mathbf{27.68 + 2.40}$
10	25%	2000	24.59 ± 0.89	17.32 ± 5.12	29.94 ± 4.53	$59.19 + 11.79$	$27.10 + 3.64$
10	25%	3000	20.64 ± 2.64	$\mathbf{24.01 \pm 4.92}$	28.19 ± 4.62	$47.38 + 10.49$	$25.76 + 4.33$
70	1%	2000	0.98 ± 0.04	1.56 ± 0.04	1.16 ± 0.19	$3.97 + 0.54$	$26.34 + 1.27$
70	1%	3000	0.65 ± 0.05	0.70 ± 0.07	$\mathbf{1.31 \pm 0.22}$	$\mathbf{3.98 + 0.47}$	$\mathbf{27.02 + 1.24}$
70	5%	2000	4.92 ± 0.18	2.78 ± 0.35	4.02 ± 1.79	$9.96 + 2.62$	$26.39 + 1.27$
70	5%	3000	1.58 ± 0.24	2.01 ± 0.46	3.59 ± 0.83	$8.78 + 2.19$	$\mathbf{27.05 + 1.23}$
70	10%	2000	9.84 ± 0.36	6.48 ± 1.99	11.34 ± 1.89	$22.81 + 4.55$	$26.21 + 1.16$
70	10%	3000	6.66 ± 2.34	$\mathbf{7.85 \pm 2.50}$	$\mathbf{12.82 \pm 3.59}$	$\mathbf{23.19 + 5.75}$	$\mathbf{26.78 + 1.13}$
70	25%	2000	24.59 ± 0.89	17.29 ± 5.16	28.92 ± 4.26	$53.21 + 10.13$	$25.15 + 1.31$
70	25%	3000	20.46 ± 2.49	$\mathbf{23.92 \pm 4.67}$	28.12 ± 3.99	$46.41 + 9.72$	$24.74 + 1.26$

Table 4.4 Absolute error in percent – ensemble mean and standard deviation– for the five components of the PST model [$AE_{\mathbb{U}}$, $AE_{\mathbb{V}}$, $AE_{\mathbb{X}}$, $AE_{\mathbb{Y}}$, $AE_{\mathbb{Z}}$]. Shown are four different noise levels nl = [1%, 5%, 10%, 25%] and three different scale factors ($s_{\mathbb{Z}}$)=[1, 10, 70]. The ensemble of 16 estimates (Set B) and the values are shown after both 2000 and 3000 gradient descent iterations. The "hold on" stage ends after 2000 iterations, and the "let go" stage ends after 3000 iterations (see Section 3.6). Values in bold font indicate instances where the error has grown has grown during the "let go" stage.

$s_{\mathbb{Z}}$	nl	j																
1	1%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1	1%	3000	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
1	5%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1	5%	3000	3	4	4	**509**	509	509	509	509	509	509	509	509	509	509	509	509
1	10%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	**30**	30
1	10%	3000	**23**	370	433	528	528	528	528	528	528	528	528	528	528	528	528	528
1	25%	2000	1	**51**	52	62	453	539	539	539	539	539	539	539	539	539	539	540
1	25%	3000	**109**	392	539	539	539	539	539	539	539	539	539	539	539	539	540	540
10	1%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
10	1%	3000	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
10	5%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
10	5%	3000	3	3	3	5	**268**	340	349	351	415	432	443	444	497	507	509	518
10	10%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	**32**
10	10%	3000	**26**	35	37	46	475	528	528	528	528	528	528	528	528	528	528	528
10	25%	2000	1	1	**30**	32	39	40	52	64	539	539	539	539	539	539	539	539
10	25%	3000	**12**	35	56	67	79	87	97	98	109	109	119	159	539	539	539	548
70	1%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
70	1%	3000	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
70	5%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
70	5%	3000	2	3	3	3	5	**308**	349	422	423	423	432	434	497	497	509	518
70	10%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	**40**
70	10%	3000	**34**	39	45	182	267	351	422	519	528	528	528	528	528	528	528	528
70	25%	2000	1	1	**32**	40	40	40	52	60	63	539	539	539	539	539	539	539
70	25%	3000	**32**	27	48	66	88	88	98	98	99	110	120	130	130	169	539	539

Table 4.5 Prediction time – number of years before the error of the $\mathbb{U}$ component is larger than the noise level – for the five components of the PST model [$AE_{\mathbb{U}}$, $AE_{\mathbb{V}}$, $AE_{\mathbb{X}}$, $AE_{\mathbb{Y}}$, $AE_{\mathbb{Z}}$]. Shown are four different noise levels nl = [1%, 5%, 10%, 25%] and three different scale factors ($s_{\mathbb{Z}}$)=[1, 10, 70]. The ensemble of 16 estimates (Set A) and the values are shown after both 2000 and 3000 gradient descent iterations. The "hold on" stage ends after 2000 iterations, and the "let go" stage ends after 3000 iterations (see Section 3.6). Bold font indicates the first ensemble member with a prediction time longer than 11 years.

$s_{\mathbb{Z}}$	nl	j																
1	1%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1	1%	3000	1	1	1	1	1	1	1	1	1	1	1	1	1	3	5	6
1	5%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1	5%	3000	1	1	1	2	2	2	2	2	3	4	**13**	44	45	75	76	76
1	10%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1	10%	3000	1	1	1	2	2	2	2	3	3	4	**34**	35	46	66	75	76
1	25%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1	25%	3000	1	1	1	2	2	2	2	2	2	2	**12**	13	97	127	443	444
10	1%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
10	1%	3000	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
10	5%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
10	5%	3000	1	1	1	2	3	3	**12**	35	36	44	44	46	75	77	117	160
10	10%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
10	10%	3000	1	1	1	2	2	2	3	4	4	**76**	77	86	117	118	160	200
10	25%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
10	25%	3000	1	1	1	2	2	2	2	2	2	2	**11**	34	43	54	170	496
70	1%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
70	1%	3000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2	**11**
70	5%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
70	5%	3000	1	1	1	1	2	2	2	3	4	4	**45**	45	75	116	117	159
70	10%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
70	10%	3000	1	1	1	1	2	2	2	3	3	3	**44**	77	77	117	159	201
70	25%	2000	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
70	25%	3000	1	1	1	2	2	2	2	2	2	2	**12**	24	34	44	55	494

Table 4.6 Prediction time – number of years before the error of the $\mathbb{U}$ component is larger than the noise level – for the five components of the PST model [$AE_{\mathbb{U}}$, $AE_{\mathbb{V}}$, $AE_{\mathbb{X}}$, $AE_{\mathbb{Y}}$, $AE_{\mathbb{Z}}$]. Shown are four different noise levels nl = [1%, 5%, 10%, 25%] and three different scale factors ($s_{\mathbb{Z}}$)=[1, 10, 70]. The ensemble of 16 estimates (Set B) and the values are shown both after 2000 and 3000 gradient descent iterations. The "hold on" stage ends after 2000 iterations, and the "let go" stage ends after 3000 iterations (see Section 3.6). Bold font indicates the first ensemble member with a prediction time longer than 11 years.

Tables 4.5 and 4.6 for the 12 cases and for 2000 and 2000+1000 iterations. The prediction time is based on free trajectories starting from the 16 estimates at time t=0. Note that the values correspond to the first year when the error of prediction is larger than the noise level. Also note that the distance from the truth allowed is equal to the noise level and this gives longer prediction times for higher noise levels. The prediction time for all ensemble members is shown, sorted in rising order. The definition of prediction time has a strong effect on the 1% noise level making it difficult to evaluate these cases. It is difficult to remain within 1% of the true value. The prediction that was initialized after 2000 iterations has similar issues. The initial condition for the U variable corresponds, in this case, to the last observation. It is difficult to remain below the noise level if the initial condition start at the noise level. The first ensemble member in Tables 4.5 and 4.6 with a prediction time larger than a cycle (11 years) is marked in bold.

We find that for Set A a large majority of the ensemble members correctly predict the onset of an inactive phase. This is even true for the $s_{\mathbb{Z}} = 1$ case where about 80% of the ensemble members enter the grand minimum.

The most astonishing result for Set A is that the estimates for $s_{\mathbb{Z}} = 1$ are superior for 1%, 5% and 10% despite the large error in the $\mathbb{Z}$ component. For Set B the results are comparable with a slight edge to the $s_{\mathbb{Z}} = 10$ case. About 40% of the ensemble members of Set B make an accurate prediction for more than one cycle, and most of these can predict several cycles.

These results would call into question the reasoning that high estimation errors have to imply unreliable predictions. This might bode well for the real world application where

we know that good estimates of all variables might not even be possible.

The ability to predict the onset of the inactive phase is reassuring. It makes it plausible that the PST system and the Shadowing Filter might be applicable to the solar cycle problem. In this section we have not addressed two major challenges. First, we did not determine model parameters and, second, the model and the system were identical. In addition we could also have explored the amount of data needed to make reasonable predictions using, for example, the window test of Stemler & Judd (2009). In the end it is only with real data that we can determine to what extent model-based prediction of the solar cycle is possible.

4.6 Outlook towards a prediction

To make a model-based prediction of the solar cycle we need to determine what model parameters to use and also make an estimate of the current state of the system (the Sun), in terms of the model, using the observed data. We have 24 solar cycles in the record corresponding to about 256 observations. The main difference between the ideal case studied in this work and the real case is that we have to determine to what extent we should "let go" of the data. The basic approach should be to keep the observations fixed and only update the unobserved variables. Unfortunately the result of Section 4.5 did not tell us very much about predictions starting after 2000 iterations at the end of the "hold on" stage.

We have two measures that can be used to determine parameters: the implied noise ($\mathcal{I}$) and the mismatch ($\mathcal{M}$). The implied noise is only defined during the "let go" stage

of the procedure. One interesting use of the implied noise would be to fully let go of the observations and let the gradient descent method converge towards a trajectory of the model. The evaluation of $\mathcal{I}$ at this time would be very informative, as it would give an estimate of the distance between the observations and the trajectories of the models. Low values would indicate a model that is close (in state space) to the observations.

We could call the mismatch $\mathcal{M}$ before we let go of the data $\mathcal{M}^0$, and $\mathcal{I}^M$ would be the implied noise close to a trajectory of the model. Parameter values that give both low $\mathcal{M}^0$ and $\mathcal{I}^M$ would be very interesting. Are the estimates consistent over time or does estimates during the 19th century indicate different optimal parameters than estimates obtained from the 20th century? Also, would estimates using 50 observations be the same as ones obtained from 250 observations?

The c and d parameters of the PST system are of great interest since they are related to the lengths of the active phase of the model. If these parameters can be determined we could infer the length of the current solar cycle. This would give us a rough estimate of when the transition to a grand minimum will occur since we know when it started (early 18th century).

Estimates using ensembles with at least 64 members can be used to study statistical properties. Unfortunately it would be very expensive to sample a large volume of parameter space if each estimate required a full ensemble. A balance has to be struck between these interests for any reasonable progress to be made.

We would also need to check the predictive power of the model using in-sample prediction. An example would be to use the first 101 observations and make a prediction

starting at the beginning of solar cycle number 10 that began in 1855. For how long would the prediction stay close to the observed data? Would the time the prediction remains close to the observations be longer for different parameter values? Here the notion of *shadowing time* (Smith et al. 2010) and *ignorance* (Du & Smith 2012) would be of great interest. We could make predictions in sample for each of the last 14 cycles and that might be enough to estimate the predictive power. One could also make predictions starting from each year in the record and study if the stage of the cycle affects the prediction. There are many tests that we would like to do using the data but we should note that we only have a limited record and the data we use for calibration should not also be used for verification. Real data is precious and we should try not to contaminate it if possible.

Chapter 5

Conclusion and further work

5.1 Summary

There are several indications in literature that the Sun might be entering a grand minimum. Model-based prediction is needed to verify these claims but the methodological situation has been unclear due to the nonlinear nature of solar variation. This work has been aimed at removing methodological obstacles – mostly related to unobserved variables – to facilitate further solar cycle inquires; more specifically the study has outlined a path towards model-based predictions of the solar cycle using low-dimensional chaotic models. We have also found (Section 4.5) that the PST model can, at least in an ideal case, predict the onset of an inactive phase – if nonlinear data assimilation tools are applied.

Data assimilation is a general term for the study of model/system/data integration. It has generated significant progress in neighboring fields like geophysics and meteorology. As the time domain become more popular these methods will become very relevant for

astronomers and astrophysicists that want to explore the semi-regular and irregular signals of the sky.

This work presents both classic and nonlinear data assimilation in a general framework. These approaches share the fundamental choice of whether the trust will be implicitly put in the observations or in the model.

The 3DVAR methodology (discussed in Chapter 2) explicitly makes a compromise, balancing the uncertainty of the model estimate and the uncertainty of the observations. Imagine the model estimate and the observation as two points in *state space*. The compromise is determined by the weight factor w. For example, $w < 1$ implies that the uncertainty of the model state is larger than the uncertainty of the observations, and thus our estimate should be closer to the observation. A value for the weight factor of $w = 1$ indicates that the expected variance of the observations is the same as for the model state. In this case the optimal compromise is the mean value. For weight factors $w > 1$, it is the model state that is more reliable, i.e. the estimate is a factor w closer to the model state than the observations. In the worst case the compromise gives us an estimate that is neither close to the observations nor consistent with the model. The result of our study (Chapter 2) show that estimates are closer to the true state of the system when weight is given to the observations ($w < 1$). We also have seen that it is possible to reduce the noise level, below that of the observations.

In the case of the Shadowing Filter approach (discussed in Chapter 3 and 4), the full sequence of observations corresponds to a specific point in *sequence space*. As stated before, every model trajectory and pseudo-trajectory of the model are represented by a

point in this space. The *gradient descent* method takes an initial set of estimates (typically the observations) and refines the estimate towards one of these trajectories. In the ideal case, in which the model and the system are the same, the true states of the system are a causally connected sequence. It is reasonable, in the *ideal case*, to place full faith in the model and "let go" of the data, as a trajectory for the model is also – by definition – a trajectory for the system.

Beyond the ideal case we should put our trust in the data as much as possible. There are several concepts like 'the true state,' 'the true parameters' or 'the best estimate' that are very unclear beyond the ideal case. It is important to urge caution in the real case because both the model and the observations are uncertain. We cannot put all of our trust in the model as we did in the ideal case. An estimate that is close to a trajectory of the model may not necessarily be close to the trajectory of the system. Actually, it is only through observations that we can know the state of the system. The "true state" might be beyond our reach. For the unobserved variable case the basic approach should be to "hold on" to the data. The only reason to "let go" of the data would be to improve the predictive power of the procedure. This would also imply that we shift some of our trust away from the observation onto the model, and this should be done with great care.

We have explored two specific challenges. First we address the question "how do we treat *unobserved variables* that do not correspond to some specific observable?" This has been outlined in Section 3.6. It is difficult to find the set of parameters used to generate that signal, for high noise levels. What we find (Section 3.7) is that areas of interesting parameter values (low $\mathcal{I}$ and $\mathcal{M}$) in parameter space correspond to close model state

estimate (low absolute error) and better prediction (longer prediction time).

The second issue is "how do we determine model parameters?" The general procedure is presented in Section 3.5. Parameter estimation for unobserved variables for the Lorenz'63 system is shown in Section 3.7. The PST system studied in Chapter 4 has eight parameters and at least one scale-factor and we have not found an efficient way to explore such high-dimensional parameter space. The focus of this work has been to outline the procedure and measures that can be used to explore parameter space and compare different parameter sets.

Our results for the Shadowing Filter (Sections 3.7, 4.3 and 4.5) show us that – in cases where the true state is known – the estimates are close to the true state. The estimate is closer for lower noise levels. For a fixed number of gradient descent iterations of the Shadowing Filter, the closeness to a trajectory (mismatch $\mathcal{M}$) is evaluated and the distance from the observation (implied noise $\mathcal{I}$) is calculated. "Good" parameter values correspond to low values of $\mathcal{M}$ and $\mathcal{I}$.

Why a low mismatch value corresponds to a better set of parameters can be understood in terms of relative distances in sequence space between the initial and final estimate. If the trajectory of the model is "further away" from the initial estimate when a specific set of parameters is used, then it will require more iterations to converge towards the trajectory. The convergence would be faster for what we may call a closer distance and the final mismatch value would be lower given a fixed number of iterations. This would connect "good" parameters to "close" distances and lower mismatch values.

Not all model variables contain the same information. For example, the signs of the

$\mathbb{X}$ and $\mathbb{Y}$ in the Lorenz'63 system (Equation 3.8) are impossible to obtain from $\mathbb{Z}$ using our method, though our method can deliver a good estimate of the absolute value, and in some cases this might be enough.

There is also the issue of the $\mathbb{Z}$ parameter of the PST system (Equation 4.3). The amplitude of this variable is much lower than the amplitude of the other components and the procedure does not, in relative terms, make a very good estimate. We obtain better estimates if a linear scaling of the low amplitude variable is used.

In the synthetic prediction of the solar cycle where we only observe the $\mathbb{X}$ component we find that prediction into an "on phase" of the system (set B) is somewhat better using a scaled system. About 6/16 ensemble members can predict more than one cycle and out of these most can predict much longer, up to ten or more cycles. Prediction into an "off phase" of the system (set A) is better using an unscaled system despite the very large error (up to over 500%) in the $\mathbb{Z}$ component. This would be a case where being close to the observations does not imply a good prediction. The majority (about 80%) of the ensemble members enters the off phase of the system and correctly predicts the continuation of the grand minimum for several hundred years. It is good news that the method can catch the onset of the "off stage" but it does not guarantee that it is going to work in the real case.

5.2 Further work

This work has outlined a path towards a model-based prediction of the solar cycle. The main hurdle not overcome in this work is an efficient exploration of parameter space for

the PST system. Not only do we need to visualize a 5+ dimensional space, we also have to deal with the computational cost. This should only be a problem of implementation and finding a sufficiently large computer. One intermediate step would be to explore the window test of Stemler & Judd (2009). This would give us an indication of how much data is needed to make a prediction of a particular length. Here we have used 256 or 512 observations and it might be possible to use much less. Parameter estimation might complicate the picture by requiring more data necessitating us to have to strike a balance.

There are also other sets of data that could be used. The monthly average Sunspot number from SIDC, – both smoothed and not smoothed – might provide more information useful to make a accurate prediction. It might also introduce more noise into the problem and the information gained might not be useful for long term prediction. Other solar records like the group Sunspot number (Hoyt & Schatten 1998; Hathaway et al. 2002) might also be of interest.

Another lane of exploration is the comparison of different solar cycle models starting with the Knobloch et al. (1998) model. Here the individual models would first have to be explored in an ideal setting and than applied to the real case. It is expected that the different models would have different behaviors and the method would have to be adjusted to account for these. In the end it would put the proposed models to the ultimate test, prediction. This would be a great tool for model development where different features and hypotheses could be tested against nature.

The PST system can be seen as an engine that generates the main behavior of the solar cycle. There are a number of features that could be added that might generate

a better prediction. In Appendix C we show an easy example of how to generate a dynamically modulated phase. Other features to add include asymmetrical cycle shape and more varied transitions from low cycles to high cycles.

Data assimilation for solar sub-surface weather would also be of interest. Here it might be of interest to look not at 3DVAR but rather at the Shadowing Filter. The satellite observation we have of the solar surface is typically 2D and the depth information would need to be inferred if we are to use 3D. The unobserved lower layers can be seen as unobserved variables and the methods discussed in Chapter 3 would in principle be applicable to this case.

5.3 Conclusion

With the outlined procedure for model-based prediction it is now possible to engage the questions about the solar cycle using models of different levels of complexity. It will be possible to evaluate, compare and validate different models by using data assimilation and by making predictions. It remains to be seen to what extent phenomenological models can give accurate predictions of the solar cycle. If prediction is possible it might indicate that there could be some low-dimensional process in the Sun responsible for the observed variation. The possible structural similarity between a successful model and the Sun would be a great discovery. This would invite an interesting study where components and parameters of the phenomenological models would be connected to equation and theory derived from physical first principles. We might be able to better understand the past, present and future behavior the Sun by using the tools outlined in this work and study

the dynamical evolution of the solar cycle.

Chapter 6

Epilogue

6.1 A preliminary prediction of the solar cycle

In Chapter 4 we outline how to make a model based prediction of the solar cycle using for example, the PST model (presented in Section 4.1). Independently of what model we use, we only have one observable provided by the solar cycle record leaving a number of model variable unobserved. We need to obtain dynamically consistent estimates, for all model parameters and variables, observed and unobserved, see Sections 3.6, 3.7, 4.4 and 4.5.

The 5−dimensional PST model has a high dimensional parameter space making model evaluation computationally expensive and difficult to visualize, see Sections 4.1 and 4.6. This is a standard optimization problem without major conceptual challenges but a full scale implementation that will require the use of suitable parallelization and visualization tools.

The ΛCDM model of cosmology has six free parameters and the method of *Monte Carlo Markov Chain* was used with great success to find the best fit to the *WMAP* cosmic background radiation data (Spergel et al. 2003). In the case of ΛCDM the specific parameter estimation approach was appropriate since the data was assumed to be Gaussian and random. It is unclear to what extent this approach would be applicable to the solar cycle data. For more on *Monte Carlo Markov Chain* in the context deterministic models and parameter estimation see Cuéllar (2007).

In order to generate a first preliminary model based prediction of the solar cycle, we use real solar cycle data (Clette et al. 2007) and what we call the PST standard parameters, discussed in Section 4.2: $[\beta, x0, \omega, a, b, q, d, c] = [1, -0.15, 2, 0.7, 0, 0, 0.03, -0.5]$. For a physical interpretation of some of these parameters see Section 4.1.

The general setup is:

- Take all available data starting at 1755 up until *the time of prediction*. These observations correspond to the $\mathbb{U}$ variable of the model.

- Transform the solar cycle data with alternate cycles using the *Bracewell construction* as outlined in section 4.2.

- Scale the data to a range corresponding to the $\mathbb{U}$ variable using a *conversion factor*, see below.

- Make 64 sets of random uniform guesses for the four unobserved variables $\mathbb{V}$, $\mathbb{X}$, $\mathbb{Y}$, $\mathbb{Z}$ of the scaled PST model (Sections 4.4 and 4.5) to create an initial condition ensemble in sequence space. The ranges of the guesses for the maximum and minimum values

of the unobserved variables are obtained from 10.000 years of model output.

- Apply 2,000 gradient descent iterations to each of the 64 ensemble members in sequence space. Recall that the estimate obtained is a pseudo- trajectory of the model updated 2,000 times such that each time the one-step prediction error or mismatch is minimized, see Section 3.6. After this procedure is applied we obtain 64 estimates of the state of the system at the time of prediction.

- Use the set of 64 estimates at the *time of prediction* as an initial condition ensemble to either make an in-sample prediction or predict the future behavior of the observed system.

Figure 6.1 shows ensemble predictions (blue) starting in 1954, 1964, 1976, 1986, 1996 and 2009 at the onset of the last six cycles using the general setup outlined above. The predictions that are made, 44 years in-sample and out-of-sample, and qualitative comparison with the observation observations (red) are made. For Figure 6.1 we use a conversion factor of 0.98458/150 to transform the solar cycle record to the range of the $\mathbb{U}$ variables. This sets the peak value for typical model output 0.98458 to correspond to a sunspot number of 150. Other conversion factors where tested but they did not improved the qualitative agreement of the in-sample prediction to the observation over the last few cycle. A scale factor $s_{\mathbb{Z}} = 10$ was also used, see Sections 4.4 and 4.5. The estimates used for prediction are obtained after 2000 iterations, at the end of the *hold on* phase. This is the basic approach where the observations have been held fixed and dynamically consistent estimates for the unobserved variables have been found, see discussion in Sections 1.7 and 4.6. The agreement looks very good at times, for example the prediction that started in

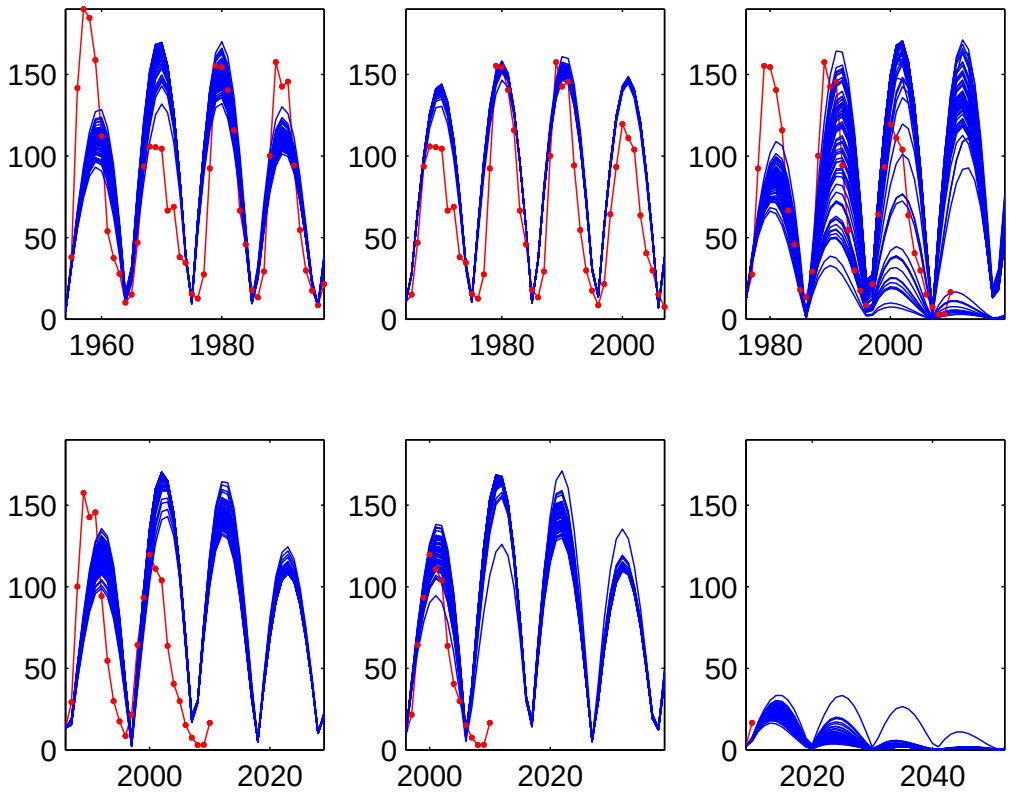

Figure 6.1 Ensemble prediction (blue) using the scaled PST model (Sections 4.1, 4.4 and 4.5) made at the onset of the last six cycles: 1954, 1964, 1976, 1986, 1996 and 2009. The estimates used to initialize the prediction are made after 2000 iterations where the observations are held fixed. In this case the estimates are constrained by the observations and we never "let go" of the data at any point. We use a scale factor $s_{\mathbb{Z}} = 10$, a conversion factor of 0.98458/150 and standard model parameters. The different between the 64 ensemble members is the realization of the guesses for the four unobserved variables. The observations are shown in red for reference. Predictions are made using the data available up until 2009 (last figure on second row) predict the onset of an inactive phase.

1996 (shown in center lower panel in Figure 6.1). The prediction that started in 2009 (last panel in Figure 6.1) indicate that we will be entering an inactive phase.

The set of 15 predictions is presented in Figure 6.2 where the last six sets of predictions are the same as in Figure 6.1. The predictions before 1948, corresponding to the first 8 panels in Figure 6.2, do not agree with the observations. In particular the model predicts a Grand minimum during the beginning of the 20th century that did not occur. Reasons for these behaviors are under study but the fact that more data is used for the later predictions might explain the better predictions after 1948. We also know that the daily sunspot number is available from 1849 and this might give more reliable yearly average values to the second half of the dataset.

The study of why the model predicted a grand minimum at the beginning of the 20th century might help us understand how the PST model operates and how to improve our ability to make more accurate predictions of the onset of an inactive phase. We should also more exhaustively explore parameter space of the PST model.

With this model-based prediction of the solar cycle we have archived the main goal of this work: to present a self-consistent procedure that generate dynamically consistent results. The performance of this particular prediction is promising, as this is only the first model and set of parameter tested. We expect the predictive power of the model to improve for areas of parameter space with lower average one-step prediction error. This approach also allows us to compare and evaluate different models and parameter sets both in qualitative and quantitative ways. We can also look for times when the model predictions do not behave like the Sun and this might help us identify what aspect of the

model to adjust. This is not the end of the line; this is the beginning of something new.

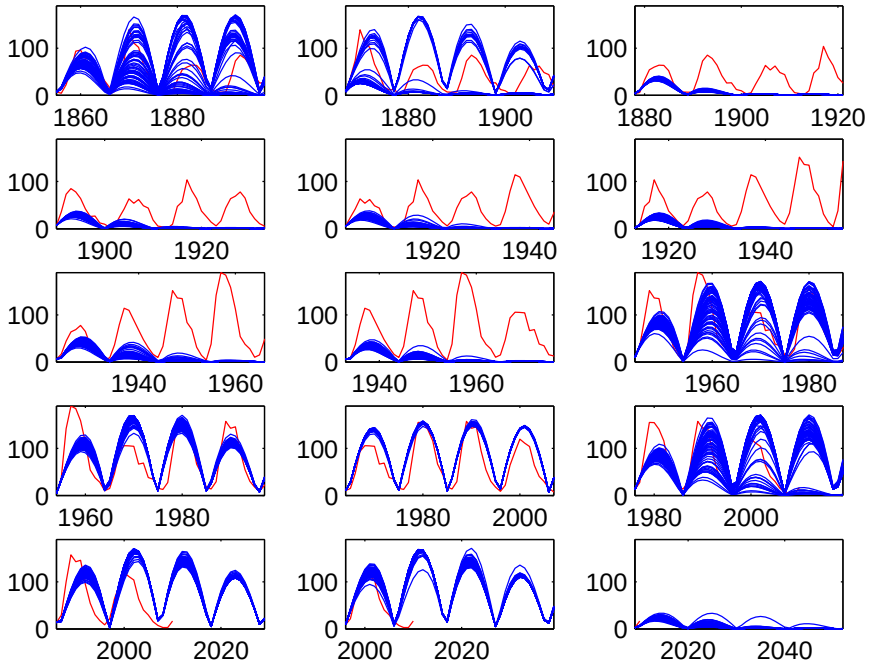

Figure 6.2 Ensemble predictions (blue) using the scaled PST model made at the onset of the last 14 cycles see Table 4.1. The estimates are made after 2000 iterations where the observations are held fixed. We use a scale factor $s_{\mathbb{Z}} = 10$, a conversion factor of 0.98458/150 and standard model parameters. The difference between the 64 ensemble members is the realization of the guesses for the four unobserved variables. The observations are shown in red for reference. Note that the prediction made using the data available up until 2009 predict the onset of an inactive phase. Also note that the predictions made using the data available at the beginning of the 20th century indicated an onset of an inactive phase that did not occur.

Bibliography

Abreu, J., Beer, J., Steinhilber, F., Tobias, S., & Weiss, N. 2008, Geophysical Research Letters, 35, L20109

Ahluwalia, H. S. 1998, Journal of Geophysical Research, 103, 12103

Baglin, A. 2003, Advances in Space Research, 31, 345

Baliunas, S. L., Donahue, R. A., Soon, W. H., Horne, J. H., Frazer, J., Woodard-Eklund, L., Bradford, M., Rao, L. M., Wilson, O. C., Zhang, Q., Bennett, W., Briggs, J., Carroll, S. M., Duncan, D. K., Figueroa, D., Lanning, H. H., Misch, T., Mueller, J., Noyes, R. W., Poppe, D., Porter, A. C., Robinson, C. R., Russell, J., Shelton, J. C., Soyumer, T., Vaughan, A. H., & Whitney, J. H. 1995, The Astrophysical Journal, 438, 269

Balmforth, N., Tresser, C., Worfolk, P., & Wu, C. W. 1997, Chaos, 7, 392

Bannister, R. 2008a, Quart. J. Roy. Meteor. Soc., 134, 1951

—. 2008b, Quart. J. Roy. Meteor. Soc., 134, 1971

Barnard, L., Lockwood, M., Hapgood, M. A., Owens, M. J., Davis, C. J., & Steinhilber, F. 2011, Geophys. Res. Lett., 38, 16103

Bengtsson, L., Ghil, M., & Källén, E. 1981, Dynamic meteorology: data assimilation methods, Vol. 36 (Springer)

Berliner, L. 1991, Journal of the American Statistical Association, 86

Borucki, W. J., Koch, D., Basri, G., Batalha, N., Brown, T., Caldwell, D., Caldwell, J., Christensen-Dalsgaard, J., Cochran, W. D., DeVore, E., Dunham, E. W., Dupree,

A. K., Gautier, T. N., Geary, J. C., Gilliland, R., Gould, A., Howell, S. B., Jenkins, J. M., Kondo, Y., Latham, D. W., Marcy, G. W., Meibom, S., Kjeldsen, H., Lissauer, J. J., Monet, D. G., Morrison, D., Sasselov, D., Tarter, J., Boss, A., Brownlee, D., Owen, T., Buzasi, D., Charbonneau, D., Doyle, L., Fortney, J., Ford, E. B., Holman, M. J., Seager, S., Steffen, J. H., Welsh, W. F., Rowe, J., Anderson, H., Buchhave, L., Ciardi, D., Walkowicz, L., Sherry, W., Horch, E., Isaacson, H., Everett, M. E., Fischer, D., Torres, G., Johnson, J. A., Endl, M., MacQueen, P., Bryson, S. T., Dotson, J., Haas, M., Kolodziejczak, J., Van Cleve, J., Chandrasekaran, H., Twicken, J. D., Quintana, E. V., Clarke, B. D., Allen, C., Li, J., Wu, H., Tenenbaum, P., Verner, E., Bruhweiler, F., Barnes, J., & Prsa, A. 2010, Science, 327, 977

Bracewell, R. N. 1953, Nature, 171, 649

Brandenburg, A., Chan, K. L., Nordlund, Å., & Stein, R. F. 2005, Astron. Nachr., 326, 681

Brandenburg, A., & Dobler, W. 2002, Computer Physics Communications, 147, 471

Brandenburg, A., Jennings, R. L., Nordlund, Å., Rieutord, M., Stein, R. F., & Tuominen, I. 1996, Journal of Fluid Mechanics, 306, 325

Brown, T. M., Christensen-Dalsgaard, J., Dziembowski, W. A., Goode, P., Gough, D. O., & Morrow, C. A. 1989, Astrophysical Journal, 343, 526

Brun, A. S. 2007, Astron. Nachr., 328, 329

Buchler, J. R., & Kolláth, Z. 2001, Nonlinear analysis of irregular variables, ed. M. Takeuti & D. D. Sasselov, 185–213

Buchler, J. R., Kolláth, Z., & Cadmus, R. 2002, in American Institute of Physics Conference Series, Vol. 622, Experimental Chaos, ed. S. Boccaletti, B. J. Gluckman, J. Kurths, L. M. Pecora, & M. L. Spano, 61–73

Bushby, P. J., & Tobias, S. M. 2007, The Astrophysical Journal, 661, 1289

Cameron, R., & Schüssler, M. 2007, The Astrophysical Journal, 659, 801

Carroll, & Ostlie. 1996, An introduction to modern astrophysics (Addison-Wesley)

Charbonneau, P., Christensen-Dalsgaard, J., Henning, R., Larsen, R. M., Schou, J., Thompson, M. J., & Tomczyk, S. 1999, Astrophysical Journal, 527, 445

Choudhuri, A. R., Chatterjee, P., & Jiang, J. 2007, Physical Review Letters, 98, 131103

Cipra, B. 1993, SIAM News, 26, 757

Clette, F., Berghmans, D., Vanlommel, P., van der Linden, R. A. M., Koeckelenbergh, A., & Wauters, L. 2007, Advances in Space Research, 40, 919

Cuéllar, M. C. 2007, PhD thesis, Department of Statistics, London School of Economics

Daley, R. 1993, Atmospheric Data Analysis, ed. R. Daley (Cambridge Univ Press)

Davies, M. 1993, International Journal of Bifurcation and Chaos, 3, 113

—. 1994, Physica D: Nonlinear Phenomena, 79, 174

Dikpati, M. 2007, APS Northwest Section Meeting Abstracts

Dikpati, M., & Charbonneau, P. 1999, The Astrophysical Journal, 518, 508

Dikpati, M., de Toma, G., & Gilman, P. A. 2006, Geophysical Research Letters, 33, 5102

—. 2008, The Astrophysical Journal, 675, 920

Dintrans, B., Brandenburg, A., Nordlund, Å., & Stein, R. F. 2005, Astronomy and Astrophysics, 438, 365

Dobricic, S., & Pinardi, N. 2008, Ocean Modelling, 22, 89

Du, H. 2009, PhD thesis, Department of Statistics, London School of Economics

Du, H., & Smith, L. A. 2012, Phys. Rev. E, 86, 016213

Eddy, J. A. 1976, Science, 192, 1189

Evensen, G. 2009, Data assimilation: The ensemble Kalman filter (Springer Verlag)

Feudel, U., Jansen, W., & Kurths, J. 1993, International Journal of Bifurcation and Chaos, 3, 131

Feynman, J. 1982, Journal of Geophysical Research, 87, 6153

Ghil, M. 1989, Dynamics of Atmospheres and Oceans, 13, 171

Gleissberg, W. 1939, The Observatory, 62, 158

Gnevyshev, M., & Ohl, A. 1948, Astronom. J, 38, 15

Haigh, J. D. 2007, Living Reviews in Solar Physics, 4

Hale, G. E., Ellerman, F., Nicholson, S. B., & Joy, A. H. 1919, Astrophysical Journal, 49, 153

Hammel, S. 1990, Physics letters A, 148, 421

Hathaway, D. H. 2009, Space Science Reviews, 144, 401

—. 2010, Living Reviews in Solar Physics, 7, 1

Hathaway, D. H., Wilson, R. M., & Reichmann, E. J. 1994, Solar Physics, 151, 177

—. 1999, Journal of geophysical research, 104, 22375

—. 2002, Solar Physics, 211, 357

Hilborn, R. C. 1994, Chaos and nonlinear dynamics : an introduction for scientists and engineers (Oxford Univ. Press)

Hoyt, D. V., & Schatten, K. H. 1998, Solar Physics, 181, 491

Hughes, D. W., Rosner, R., & Weiss, N. O., eds. 2007, The Solar Tachocline

Hurlburt, N. E., Toomre, J., & Massaguer, J. M. 1986, The Astrophysical Journal, 311, 563

Ilonidis, S., Zhao, J., & Kosovichev, A. 2011, Science, 333, 993

Jones, C., Weiss, N., & Cattaneo, F. 1985, Physica D: Nonlinear Phenomena, 14, 161

Jouve, L., Brun, A. S., & Talagrand, O. 2011, The Astrophysical Journal, 735, 31

Judd, K. 2003, Physica D: Nonlinear Phenomena, 183, 273

—. 2008, Journal of Nonlinear Science, 18, 57

Judd, K., Reynolds, C. A., & Rosmond, T. 2004a, Naval Research Laboratory Technical Report, NRL/MR/7530-04018

Judd, K., Reynolds, C. A., Rosmond, T. E., & Smith, L. A. 2008, Journal of the atmospheric sciences, 65, 1749

Judd, K., & Smith, L. A. 2001, Physica D Nonlinear Phenomena, 151, 125

—. 2004, Physica D Nonlinear Phenomena, 196, 224

Judd, K., Smith, L. A., & Weisheimer, A. 2004b, Physica D, 190, 153

Kalman, R. 1960, Journal of Basic Engineering, 82, 35

Kalnay, E. 2003, Atmospheric modeling, data assimilation, and predictability (Cambridge Univ Press)

Kitiashvili, I., & Kosovichev, A. 2008, The Astrophysical Journal Letters, 688, L49

Knobloch, E., & Landsberg, A. S. 1996, Monthly Notices of the Royal Astronomical Society, 278, 294

Knobloch, E., Tobias, S. M., & Weiss, N. O. 1998, Monthly Notices of the Royal Astronomical Society, 297, 1123

Le Dimet, F. X., & Talagrand, O. 1986, Tellus A, 38, 97

Lean, J., Beer, J., & Bradley, R. 1995, Geophysical Research Letters, 22, 3195

Livingston, W., & Penn, M. 2009, EOS Transactions, 90, 257

Lockwood, M. 2009, Royal Society of London Proceedings Series A, 466, 303

Lockwood, M., Owens, M., Barnard, L., Davis, C., & Thomas, S. 2012, Astronomy and Geophysics, 53, 030000

Lorenc, A. C. 1981, Monthly Weather Review, 109, 701

—. 1986, Quarterly Journal of the Royal Meteorological Society, 112, 1177

Lorenz, E. N. 1963, Journal of Atmospheric Sciences, 20, 130

—. 1993, The essence of chaos (Seattle: University of Washington Press)

Maunder, R. 1894, Knowledge, 17, 173

McNish, A., & Lincoln, J. 1949, Trans. Am. Geophys. Union, 30, 673

McSharry, P. E., & Smith, L. A. 1999, Physical review letters, 83, 4285

Moore, D. W., & Spiegel, E. A. 1966, Astrophysical Journal, 143, 871

National Research Council. 2010, New Worlds, New Horizons in Astronomy and Astrophysics (The National Academies Press)

Ohl, A. 1966, Solice Danie, 9, 84

Owens, M. J., Lockwood, M., Barnard, L., & Davis, C. J. 2011, Geophysical Research Letters, 38, 19106

Palmer, T., & Hagedorn, R. 2006, Predictability of Weather and Climate (Cambridge University Press)

Pasquero, C. 1995, PhD thesis, Tesi di Laurea, Faculta di Scienza MFN, Universita di Torino

Penn, M., & Livingston, W. 2006, The Astrophysical Journal Letters, 649, 45

Perryman, M. A. C., de Boer, K. S., Gilmore, G., Høg, E., Lattanzi, M. G., Lindegren, L., Luri, X., Mignard, F., Pace, O., & de Zeeuw, P. T. 2001, Astronomy and Astrophysics, 369, 339

Petrovay, K. 2010, Living Reviews in Solar Physics, 7

Pires, C., Vautard, R., & Talagrand, O. 1996, Tellus A, 48, 96

Platt, N., Spiegel, E., & Tresser, C. 1993a, Physical Review Letters, 70, 279

—. 1993b, Geophysical & Astrophysical Fluid Dynamics, 73, 147

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. 1992, Numerical recipes in FORTRAN. The art of scientific computing (Cambridge: University Press)

Rayleigh, L. 1917, Royal Society of London Proceedings Series A, 93, 148

Ribes, J. C., & Nesme-Ribes, E. 1993, Astronomy and Astrophysics, 276, 549

Ridout, D., & Judd, K. 2002, Physica D Nonlinear Phenomena, 165, 26

Ruelle, D., & Takens, F. 1971, Comm. Math. Phys, 20, 167

Saltzman, B. 1962, Journal of Atmospheric Sciences, 19, 329

Schatten, K. H., Scherrer, P. H., Svalgaard, L., & Wilcox, J. M. 1978, Geophysical Research Letters, 5, 411

Schatten, K. H., & Sofia, S. 1987, Geophysical Research Letters, 14, 632

Schwenn, R. 2006, Living Reviews in Solar Physics, 3

Smith, L. A., Cuéllar, M. C., Du, H., & Judd, K. 2010, Physics Letters A, 374, 2618

Spergel, D. N., Verde, L., Peiris, H. V., Komatsu, E., Nolta, M. R., Bennett, C. L., Halpern, M., Hinshaw, G., Jarosik, N., Kogut, A., Limon, M., Meyer, S. S., Page, L., Tucker, G. S., Weiland, J. L., Wollack, E., & Wright, E. L. 2003, The Astrophysical Journal Supplement Series, 148, 175

Spiegel, E. A. 1972, NASA Special Publication, 300, 61

—. 1980, Annals of the New York Academy of Sciences, 357, 305

Spiegel, E. A. 1994, in Lectures on Solar and Planetary Dynamos, ed. M. R. E. Proctor & A. D. Gilbert, 245

Spiegel, E. A. 1997, in Past and present variability of the solar-terrestrial system: measurement, data analysis and theoretical models, ed. G. C. Castignoli & A. Provencale, 311

—. 2009, Space Science Reviews, 144, 25

Spiegel, E. A., & Weiss, N. O. 1980, Nature, 287, 616

Spiegel, E. A., & Zahn, J.-P. 1992, Astronomy and Astrophysics, 265, 106

Stemler, T., & Judd, K. 2009, Physica D Nonlinear Phenomena, 238, 1260

Strogatz, S. 2001, Nonlinear dynamics and chaos: With applications to physics, biology, chemistry, and engineering (Perseus Books)

Svalgaard, L., Cliver, E. W., & Kamide, Y. 2005, Geophysical Research Letters, 32, 1104

Talagrand, O. 1997, J. Meteorol. Soc. Japan, 75, 191

Thomas, J. H., & Weiss, N. O. 2008, Sunspots and Starspots, ed. Thomas, J. H. & Weiss, N. O. (Cambridge University Press)

Thompson, R. J. 1993, Solar Physics, 148, 383

Tobias, S., Hughes, D., & Weiss, N. 2006, Nature, 442, 26

Tobias, S. M., Weiss, N. O., & Kirk, V. 1995, Monthly Notices of the Royal Astronomical Society, 273, 1150

Tresser, C., Worfolk, P. A., & Bass, H. 1995, Chaos, 5, 693

Tucker, W. 2002, Found. Comput. Math, 2, 53

Usoskin, I. G. 2008, Living Reviews in Solar Physics, 5

Usoskin, I. G., Solanki, S. K., & Kovaltsov, G. A. 2007, Astronomy and Astrophysics, 471, 301

Warnecke, J., Brandenburg, A., & Mitra, D. 2011, Astronomy and Astrophysics, 534, A11

Weiss, N. 2010, Astronomy and Geophysics, 51, 030000

Wilmot-Smith, A. L., Martens, P. C. H., Nandy, D., Priest, E. R., & Tobias, S. M. 2005, Monthly Notices of the Royal Astronomical Society, 363, 1167

Wunsch, C. 1996, The Ocean Circulation Inverse Problem (Cambridge University Press)

Wunsch, C. 2006, Discrete inverse and state estimation problems: with geophysical fluid applications (Cambridge Univ Press)

Yang, S., Odah, H., & Shaw, J. 2000, Geophysical Journal International, 140, 158

Appendix A

Notation

Different methods have been developed in different fields and for different purposes and with different notations. The notation for 3DVAR and the Shadowing Filter follows the different conventions in the literature and this makes comparison hard at times. The time component that is an important part of the shadowing-filter does not play any direct role in 3DVAR. The shadowing-filter is looking for a sequence of estimates $\{Z_n\}$ given a sequence of observations $\{S_n\}$ at times $n = 1, 2, ..., N$. 3DVAR is looking for a compromise made at a specific time N. The 3DVAR observation Y_0 is equivalent to S_N at the current time N. The *analysis* X_a is the estimate Z_N at the current time N and the *background state* X_b is the image $F(Z_{N-1})$ from the previous estimate Z_{N-1}. In 3DVAR, Y_0 and X_b are fixed estimates that give a unique X_a. The 3DVAR costfunction:

$$J(X) = \frac{1}{2}\left[X - X_b\right]^T \mathbf{B}^{-1}\left[X - X_b\right] + \frac{1}{2}\left[Y_0 - \mathbf{H}(X)\right]^T \mathbf{R}^{-1}\left[Y_0 - \mathbf{H}(X)\right] \tag{A.1}$$

presented in Chapter 2 would in the standard notation used here be

$$\mathcal{J}(X_N) = \frac{1}{2}[X_N - F(Z_{N-1})]^T\mathbf{B}^{-1}[X_N - F(Z_{N-1})] + \frac{1}{2}[S_N - \mathbf{H}(X_N)]^T\mathbf{R}^{-1}[S_N - \mathbf{H}(X_N)] \tag{A.2}$$

and X_a or Z_N would be the X_N that would minimize $\mathcal{J}$ and Z_{N-1} would be the final estimate at the previous time.

Appendix B

Notes on numerics

For the integration of the model in the shadowing-filter in Chapter 3 and 4 we use a 4th order *Runge Kutta* algorithm:

$$\begin{aligned} X_{n+h} =& X_n + \frac{1}{6}(k_1 + 2k_2 + 2k_3 + k_4) \\ \mathbb{X}_{n\to n+h} =& \mathbb{I} + \frac{1}{6}(\mathbb{K}_1 + 2\mathbb{K}_2 + 2\mathbb{K}_3 + \mathbb{K}_4) \end{aligned} \tag{B.1}$$

that takes some model state X_n and gives us an image X_{n+h}=$F(X_n)$ where $\mathbb{X}_{n\to n+h}$ is the adjoint matrix for the transition n to $n+h$, $\mathbb{I}$ is the identity matrix and the 4 coefficients k_i and matrices $\mathbb{K}_i$ are:

$$
\begin{aligned}
k_1 =& h\mathbf{F}(X_n) \\
k_2 =& h\mathbf{F}(X_n + k_1/2) \\
k_3 =& h\mathbf{F}(X_n + k_2/2) \\
k_4 =& h\mathbf{F}(X_n + k_3) \\
\mathbb{K}_1 =& h\mathbb{F}(X_n) \\
\mathbb{K}_2 =& h\mathbb{F}(X_n + k_1/2)(\mathbb{I} + \mathbb{K}_1/2) \\
\mathbb{K}_3 =& h\mathbb{F}(X_n + k_2/2)(\mathbb{I} + \mathbb{K}_2/2) \\
\mathbb{K}'_4 =& h\mathbb{F}(X_n + k_3)(\mathbb{I} + \mathbb{K}_3)
\end{aligned}
\tag{B.2}
$$

where h is the step-size of the integration, $\mathbf{F}$ is the analytic expression for the model F and $\mathbb{F}$ is the corresponding analytic Jacobian matrix. Typically the integration step is smaller then the distance between observation and we need to take multiple integration steps to arrive at time $n+1$. It is also worth noticing that the adjoint $\mathbb{X}$ is valid only at some specific time and to create the *joint adjoint* for the integration from n to $n+1$ you multiply all the individual adjoints together:

$$
\mathbb{X}_{n \to n+1} = \prod_{j=0}^{w-1} [\mathbb{X}_{(n+h*j)\to(n+h(j+1))}] \tag{B.3}
$$

where it takes w Runge Kutta steps of size h to go from time n to $n+1$. The difference between having the product run from $j=0$ to w or to $w-1$ has been explored, i.e. should the Jacobian matrix of the end point be including in the product or not. No noticeable difference has been found between the two cases.

Appendix C

PST+

The PST model has a fixed phase related to the parameter w. The amplitude of the PST model is set by the growth rate $\beta(\mathbb{X} - x0)$ and w can be changed in a similar fashion:

$$w = w_0 + (\beta(\mathbb{X} - x0))^2 \tag{C.1}$$

to give us a modified PST system:

$$\begin{aligned}
\dot{\mathbb{U}} =& -(\mathbb{U}^2 + \mathbb{V}^2)\mathbb{U} - \left[w_0 + (\beta(\mathbb{X} - x0))^2\right]\mathbb{V} + \beta(\mathbb{X} - x_0)\mathbb{U} \\
\dot{\mathbb{V}} =& -(\mathbb{U}^2 + \mathbb{V}^2)\mathbb{V} + \left[w_0 + (\beta(\mathbb{X} - x0))^2\right]\mathbb{U} + \beta(\mathbb{X} - x_0)\mathbb{V} \\
\dot{\mathbb{X}} =& \mathbb{Y} \\
\dot{\mathbb{Y}} =& a\mathbb{X} - \mathbb{X}^3 - b\mathbb{Y} + \mathbb{Z} + q\mathbb{U}^2 \\
\dot{\mathbb{Z}} =& -d(\mathbb{Z} - c\mathbb{X}(\mathbb{X}^2 - 1))
\end{aligned} \tag{C.2}$$

that has phase that varies in time and is correlated to the peak amplitude of the cycle. A negative sign:

$$w = w_0 - (\beta(\mathbb{X} - x0))^2 \tag{C.3}$$

would have given a positive correlation between amplitude and cycle length. There is a weak empirical inverse relation between peak amplitude and the cycle length (Hathaway 2010). This would imply that a positive sign is preferred.

Figure C.1 show model output for the modified system, for standard parameters where $w_0 = 2$ corresponds to the 11 year cycle. Note that this is a qualitative illustration of a possible way to modify the the system to obtain a nonuniform phase.

Parameter adjustment and other ways to express the time variation of w might give an output more reminiscent of the solar cycle. Unfortunately this updated system does not provide any additional constraint on the unobserved variables $\mathbb{Y}$ and $\mathbb{Z}$.

Additional studies are needed to determine if this model has any improved skill in terms of prediction.

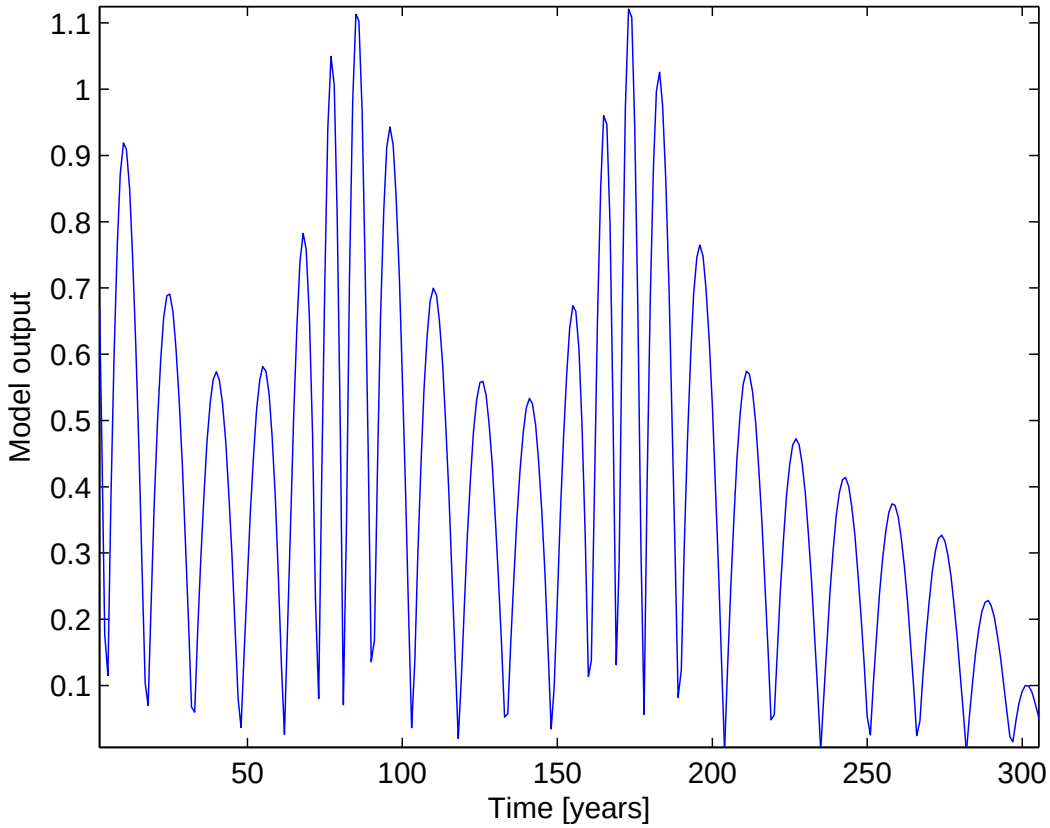

Figure C.1 Output from the modified PST system where the amplitude and phase of the oscillation is modified in a similar fashion. This is a qualitative illustration of a possible mechanism for nonuniform phase.