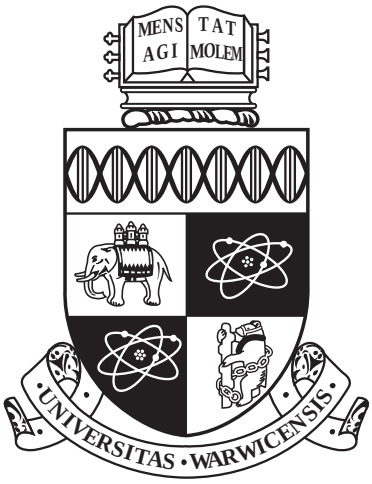

Aspects of Bayesian Inverse Problems

by

Sergios Agapiou

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Mathematics Institute

September 2013

THE UNIVERSITY OF
WARWICK

Contents

Acknowledgments

This thesis would not have been possible without the help and support of the people around me to which I am deeply and truly grateful. This is an attempt to thank at least some of them.

First of all, I would like to thank my supervisor Andrew Stuart for his continuous support and guidance in the last four years. This thesis has only been possible because of his mathematical ideas and I am extremely thankful to him for including me in his broader vision for Bayesian inverse problems. He has been a lot more than a thesis supervisor, a constant source of inspiration and optimism about mathematics and life in general.

I am also grateful to my collaborators Stig Larsson, Omiros Papaspiliopoulos, Johnathan Bardsley and Yuan-Xiang Zhang for introducing me to new mathematical concepts, letting me understand their different points of view and in general for a very fruitful cooperation. I would especially like to thank Omiros for crucial help and guidance which often reached outside the limits of our common work.

I would also like to thank my second supervisor Neil O'Connell and my thesis mentor Martin Hairer for productive conversations and useful advice relating to parts of this work.

I have been very lucky to be a part of a large research group, which has been very beneficial for both this thesis and my life as a postgraduate student in general. I take this opportunity to thank current and old members of Andrew Stuart's research group Masoumeh Dashti, Andrew Duncan, Marco Iglesias, David Kelly, Kody Law, Damon McDougall, Michela Ottobre, Sebastian Vollmer and Hendrik Weber for their friendship and support. I would like to thank Andrew and David especially

for always attempting to answer my often naive mathematical questions.

I would also like to thank my friend and fellow PhD student Dayal Strub for stimulating discussions during lunch breaks and for being a constant source of support throughout my postgraduate studies.

I am deeply grateful to my partner Chara, for her constant personal support in often difficult times.

Finally and above all, I am truly thankful to my parents Dina and Adonis for their unequivocal support and for always encouraging me to chase my dreams. My father was the one to spark my interest in sciences from a very early age; I can only hope that this thesis would have made him proud.

Declarations

This thesis contains original collaborative research conducted during the author's doctoral studies at the University of Warwick. In particular:

i) Chapter 1 is an introduction containing an overview of the thesis and a literature review;

ii) Chapter 2 is work done in collaboration with Stig Larsson (Chalmers University of Technology) and Andrew Stuart (University of Warwick, thesis supervisor) and is published in [3];

iii) Chapter 3 is work done in collaboration with Andrew Stuart and Yuan-Xiang Zhang (Lanzhou University) and has been accepted for publication, [4];

iv) Chapter 4 is work done in collaboration with Johnathan Bardsley (University of Montana), Omiros Papaspiliopoulos (Universitat Pompeu Fabra) and Andrew Stuart and an abridged version of it has been submitted for publication, [2].

For a description of the author's contribution in each of the chapters please refer to Sections 1.2.1, 1.2.2 and 1.3.2 in the introduction. All the figures in this thesis are original and produced by the author.

Abstract

The goal of this thesis is to contribute to the formulation and understanding of the Bayesian approach to inverse problems in function space. To this end we examine two important aspects of this approach: the frequentist asymptotic properties of the posterior, and the extraction of information from the posterior via sampling. We work in a separable Hilbert space setting and consider Gaussian priors on the unknown in conjugate Gaussian models. In the first part of this work we consider linear inverse problems with Gaussian additive noise and study the contraction in the small noise limit of the Gaussian posterior distribution to a Dirac measure centered on the true parameter underlying the data. In a wide range of situations, which include both mildly and severely ill-posed problems, we show how carefully calibrating the scaling of the prior as a function of the size of the noise, based on a priori known information on the regularity of the truth, yields optimal rates of contraction. In the second part we study the implementation in $\mathbb{R}^N$ of hierarchical Bayesian linear inverse problems with Gaussian noise and priors, and with hyper-parameters introduced through the scalings of the prior and noise covariance operators. We use function space intuition to understand the large N behaviour of algorithms designed to sample the posterior and show that the two scaling hyper-parameters evolve under these algorithms in contrasting ways: as N grows the prior scaling slows down while the noise scaling speeds up. We propose a reparametrization of the prior scaling which is robust with respect to the increase in dimension. Our theory on the slowing down of the evolution of the prior scaling extends to hierarchical approaches in more general conjugate Gaussian settings, while our intuition covers other parameters of the prior covariance operator as well. Throughout the thesis we use a blend of results from measure theory and probability theory with tools from the theory of linear partial differential equations and numerical analysis.

Chapter 1

Introduction

1.1 Overview

Inverse problems *are concerned with determining causes for a desired or an observed effect*, [22]. We have an equation of the form

$$y = \mathcal{G}(u), \tag{1.1.1}$$

which we want to solve for the unknown input $u \in \mathcal{X}$ given the observation $y \in \mathcal{Y}$, where $(\mathcal{X}, \|\cdot\|_{\mathcal{X}}), (\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ are Banach spaces. The possibly nonlinear operator $\mathcal{G} : \mathcal{X} \to \mathcal{Y}$ is called the *observation operator* and in general is the composition of a mathematical model, henceforward called the *forward model*, and the observation mechanism. For example, u can be the initial condition of a partial differential equation and the operator $\mathcal{G}$ the solution operator at time T; that is, $\mathcal{G}$ maps the initial condition u to the solution y at time T. The inverse problem would then be to find the initial condition u from an observation of the solution y at time T. Another possibility in this example would be to only observe the solution y at time T on a discrete set of points.

Typically inverse problems are ill-posed in the Hadamard sense: they may not have a solution, the solution may not be unique and may depend sensitively on the observation y. The latter is very important, since typically we have imperfect observations of y, modelled as

$$y = \mathcal{G}(u) + \eta, \tag{1.1.2}$$

where η is an additive noise.

The area of inverse problems has received enormous interest in recent years and during the second half of the last century a classical, deterministic mathematical

theory has been developed for restoring well-posedness. In this classical approach, the lack of existence of solution to the inverse problem is addressed by weakening the notion of solution: for example, instead of trying to find an exact solution, one may look for an approximate solution which is a minimizer of the least squares functional

$$\Phi(u;y) = \frac{1}{2}\big\|y - \mathcal{G}(u)\big\|_{\mathcal{Y}}^2.$$

Furthermore, a criterion for choosing between candidate solutions can be employed to enforce uniqueness, for example one can choose the element $\bar{u} \in \mathcal{X}$ which in addition to minimizing $\Phi(\cdot;y)$, also has the minimal norm among all minimizers; the element $\bar{u} \in \mathcal{X}$ is an example of a *best approximate solution*. However, the existence of minimizers of $\Phi(\cdot,y)$ is not guaranteed in general and even if a best approximate solution exists, it typically depends sensitively on the observation y. These issues can be addressed using regularization techniques which enforce the continuity of the approximate solution with respect to the data by looking for a *regularized approximate solution*. One of the most widely applied classical regularization techniques is the (generalized) Tikhonov-Phillips regularization [81, 61], in which one seeks for a minimizer u_λ of the functional

$$\mathcal{J}(u;y) = \Phi(u;y) + \frac{\lambda}{2}\|u\|_{\mathcal{E}}^2,$$

where $(\mathcal{E}, \|\cdot\|_{\mathcal{E}})$ is a Banach space (often compactly) embedded in $\mathcal{X}$ and $\lambda > 0$ is the *regularization parameter* which determines the relative weight between the fidelity, least squares term and the regularization term.

A typical result of classical regularization assumes the existence of a true parameter $u^\dagger$ underlying the data y and considers a sequence of idealized experiments such that the norm of the noise vanishes, $\big\|\eta\big\|_{\mathcal{Y}} \leq \frac{1}{\sqrt{n}}$ where $n \to \infty$. The objective is then to determine conditions on $\mathcal{G}$ and the underlying truth $u^\dagger$, which imply the existence of a parameter choice rule for the regularization parameter λ as a function of the size of the noise, which needs to be such that the regularization disappears as the noise vanishes ($\lambda \to 0$ as $n \to \infty$), and which secures the convergence of the regularized approximations $u_{\lambda;n}$ to the underlying truth in the small noise limit $n \to \infty$. More sophisticated results allow for data driven choices of λ, that is, they allow λ to also depend on y. Typically rates of convergence are provided depending on the properties of $\mathcal{G}$ and the regularity of the truth. The performance of a regularization method is judged by comparing the convergence rates it achieves to the *worst case error*, which is a lower bound over the best convergence rate that any regularization method can achieve for a given $\mathcal{G}$ and regularity class of the truth.

The literature in this classical approach to the regularization of inverse problems is very rich, see for example the classic books [22, 42, 85].

In this thesis we adopt a Bayesian approach, that is a statistical approach to the regularization of inverse problems. In the Bayesian approach all the variables in problem (1.1.2), that is the unknown u, the noise η and the observation y, are modelled as random. The regularized approximate solution which is the object of interest in the classical approach is replaced by a probability measure. We assume that the noise is distributed according to the probability measure $\mathbb{P}_0$, hence for a fixed u we can find the *data likelihood*, that is the distribution of the observation y given u, by shifting $\mathbb{P}_0$. We then choose a *prior* distribution on the unknown u, encoding any prior information we may have on the unknown, denoted by μ_0. An application of the Bayes' rule gives the *posterior* distribution μ^y, that is the distribution of the unknown u given the data y. The posterior distribution is the object of interest in the Bayesian approach and encompasses our updated beliefs on the unknown after incorporating the observed data.

The books [40, 80] are an excellent introduction to the Bayesian approach to inverse problems, and include many model problems mainly from differential equations as well as case studies of real-world applications. Even though both of these books contain a breadth of Bayesian techniques in the context of inverse problems, they do not attempt to formulate a complete mathematical theory of Bayesian inverse problems.

A more modern and more mathematically structured treatment based on a function space formulation of the Bayesian methodology, is presented in the review article [78]. In the function space setting, Bayes' rule often has the form of the conditioning result [32, Lemma 5.3], which in the context of inverse problems gives the Radon-Nikodym derivative of the posterior with respect to the prior,

$$\frac{d\mu^y}{d\mu_0}(u) = \frac{1}{Z(y)} \exp(-\Phi(u; y)). \tag{1.1.3}$$

Here, $Z(y)$ is a normalization constant securing that the posterior is a probability measure and Φ is a potential which relates to the forward problem and in particular to the properties of the observation operator $\mathcal{G}$. For example, if the noise is Gaussian, the potential Φ can be formally thought of as being a least squares functional (this would be the case if $\mathcal{Y}$ was finite dimensional; in infinite dimensions the situation is a bit more complicated but the intuition is still correct). Candidate solutions with high posterior probability are the ones that compromise between giving low values of $\Phi(\cdot, y)$ and having high prior probability. It is apparent that the classical and

Bayesian approach are linked: the potential Φ acts as a fidelity term while the prior acts as a regularization term.

In [78], several examples of inverse problems in function space are given a Bayesian formulation with Gaussian priors μ_0 and posteriors μ^y which are defined by (1.1.3), and it is shown that they share a common mathematical structure in the sense that they give rise to potentials Φ which satisfy certain conditions. A mathematical theory is then provided, showing that under these conditions, the posterior measure μ^y is well-defined and is Lipschitz continuous in the Hellinger distance of measures with respect to the data. Furthermore, it is shown that these conditions also secure that approximation results on the forward problem translate to approximation results on the corresponding posterior measure. More recently in [19], this mathematical framework has been extended to cover a larger class of function space priors, termed Besov priors, introduced in [49].

The biggest challenge in the Bayesian approach to inverse problems is the one of extracting information from the posterior distribution. This may be done for example by obtaining estimators of the unknown and quantifying the uncertainty around them. Unfortunately, in all but some very special cases such computations are analytically intractable and the posterior distribution is in general an enormously complicated object. In order to exploit the full potential of the Bayesian approach, we need to be able to efficiently sample the posterior distribution which enables the numerical computation of estimators and the corresponding quantification of uncertainty. Typically, Markov chain Monte Carlo (MCMC) methods are used to produce a sequence of samples approximately drawn from the posterior, however in high dimensions the conventional Metropolis Hastings type algorithms are proved to be inadequate and new algorithms need to be used. The general philosophy in [78] is that the function space formulation enables a better understanding of the issues arising in Bayesian inverse problems in finite but large dimensions. For example, even though in practice all the algorithms are implemented in finite dimensions, a sampling algorithm which is not well-defined in the infinite dimensional limit will degenerate as the dimension increases. For this reason it is desirable to design algorithms directly in the infinite dimensional limit.

In this thesis we work in a separable Hilbert space $\mathcal{X}$ and consider Bayesian inverse problems with Gaussian priors which are Gaussian conjugate, that is the data model is such that the posterior is also Gaussian. The main objectives of the two approaches to inverse problems, one being proving convergence results to the underlying truth in the classical approach and the other being probing the posterior distribution in the Bayesian approach, are brought together by the study of the

frequentist asymptotic properties of the posterior distribution in the small noise limit. That is, we assume that we have data produced from an underlying true parameter $u^\dagger$ and study the contraction of the Gaussian posterior distribution to a Dirac distribution centered on $u^\dagger$ as the noise in the data disappears. This is the subject of the first two chapters of the current thesis, briefly introduced in the next section. In particular in Chapters 2 and 3 we study linear inverse problems with Gaussian prior and noise distributions.

It is often desirable to consider prior and noise distributions defined hierarchically, that is having distributions which depend on random hyper-parameters. This may be motivated either by the Bayesian dogma that *if a parameter is not known, then it is a part of the inference*, [40], or through the study of the asymptotic behaviour of the posterior by an analogy to the data driven regularization parameter choice rules often considered in the classical approach. The hierarchical Bayesian approach gives rise to more elaborate posterior distributions simultaneously on the unknown function and the random hyper-parameters given the data, hence sampling is necessary. This is typically done using Metropolis-within-Gibbs algorithms in which a Metropolis-Hastings step is used to generate an instance from the distributions of the unknown parameter and each of the hyper-parameters in turn conditioned on the data and the current values of the other parameters. The situation is simpler if we have conditional conjugacy, in which case a plain Gibbs sampler can be used.

The efficiency of such sampling algorithms when implemented in finite but large dimensions is the topic of Chapter 4 and is briefly introduced in Section 1.3. In particular, in Chapter 4 we consider hierarchical Bayesian linear inverse problems with Gaussian noise and prior distributions and hyper-parameters introduced through the scalings of the covariance operators of the prior and noise distributions. We work with inverse-Gamma hyper-priors which are conditionally conjugate and apply function space intuition to understand the mixing behaviour of the corresponding Gibbs sampler as the discretization level increases. Our intuition carries over to other conjugate Gaussian setups as well; see Chapter 4.5 and 4.8 for details.

1.2 Asymptotic performance - rate of posterior contraction

The study of the asymptotic performance of the Bayesian posterior distribution μ^y from a frequentist point of view is of central importance in Bayesian statistics. As explained in [20], whether one believes in the existence of a true parameter

underlying the data or not, it is always reasonable to ask *what if* the data is produced from an underlying truth. It is then desirable to have that as more informative data comes in, the posterior concentrates around the underlying truth. In Bayesian statistics this is formalized by the notion of *posterior consistency*. Assume that we have a sequence of observations $\{y_n^\dagger\}_{n\in\mathbb{N}}$, generated from an underlying true parameter $u^\dagger \in \mathcal{X}$, where as $n \to \infty$ the information increases. We denote by $\mu^{y_n^\dagger}$ the corresponding posterior distribution and by $\mathbb{E}^{y_n^\dagger}$ the expectation with respect to the distribution of $y_n^\dagger$.

Definition 1.2.1 (Posterior consistency)**.** *The posterior is called consistent (with respect to the metric d) at $u^\dagger$, if for every $\epsilon > 0$*

$$\mathbb{E}^{y_n^\dagger}\mu^{y_n^\dagger}\{u : d(u, u^\dagger) \geq \epsilon\} \to 0,$$

as $n \to \infty$.

Doob in [21] showed that posterior consistency holds under very weak measurability conditions. Essentially Doob's result says that for every prior distribution on the parameter space $\mathcal{X}$, posterior consistency holds for every underlying truth except for a set of truths having prior measure zero. Consistency fails for true parameters $u^\dagger$ in a null set of the prior distribution and Doob's result gives no information about this null set. Even though one may hope that things only go wrong in pathological, special cases, this is not necessarily the case in nonparametric models [26, 20], and we thus need to be careful when choosing the prior distribution. For a given prior and model, it is desirable to be able to secure posterior consistency for an underlying true parameter in a known set. Such posterior consistency results in nonparametric models have been proved in weak metrics in [76] and in stronger metrics in [9, 27].

A more quantitative measure of the asymptotic performance of the posterior distribution is the speed at which the posterior contracts to the truth as captured by the *rate of contraction* defined below.

Definition 1.2.2 (Posterior rate of contraction)**.** *The posterior is said to contract (with respect to the metric d) at $u^\dagger$ with rate ϵ_n, $\epsilon_n \downarrow 0$, if*

$$\mathbb{E}^{y_n^\dagger}\mu^{y_n^\dagger}\{u : d(u, u^\dagger) \geq M_n\epsilon_n\} \to 0,$$

for every sequence $M_n \to \infty$ as $n \to \infty$.

An arsenal of techniques for proving rates of posterior contraction in the Hellinger metric for suitable priors in general contexts has been developed in [28, 77].

More recently, posterior contraction results in general contexts have been proved for example in [83, 82, 84, 37, 29, 74].

In this thesis we study rates of posterior contraction in the small noise limit in linear inverse problems in a separable Hilbert space $(\mathcal{X}, \|\cdot\|)$. In particular, we consider the model

$$y = Ku + \frac{1}{\sqrt{n}}\eta, \tag{1.2.1}$$

where $K : \mathcal{X} \to \mathcal{X}$ is an injective linear bounded operator and $\eta \sim \mathcal{N}(0, \mathcal{C}_1)$ is a Gaussian additive noise. We put a Gaussian prior on the unknown $u \sim \mathcal{N}(0, \tau^2\mathcal{C}_0)$ resulting in a Gaussian posterior distribution $\mu^y = \mathcal{N}(m, \mathcal{C})$ with known mean m and covariance operator $\mathcal{C}$, as proved in [54, 52]. We then consider a sequence of observations of the form $y_n^\dagger = Ku^\dagger + \frac{1}{\sqrt{n}}\eta$, where $u^\dagger$ is an underlying true parameter known to belong to a smoothness class indexed by $\gamma > 0$ and determine rate $\epsilon_n = \epsilon_n(\gamma)$, such that as the noise disappears $(n \to \infty)$

$$\mathbb{E}^{y_n^\dagger}\mu^{y_n^\dagger}\left\{u : \left\|u - u^\dagger\right\| \geq M_n\epsilon_n\right\} \to 0, \quad \forall M_n \to \infty. \tag{1.2.2}$$

The combination of K and $\mathcal{C}_1$ determines the *degree of ill-posedness* of problem (1.2.1). In broad terms, assuming that $\mathcal{C}_1^{-\frac{1}{2}}K$ possesses a discrete set of singular values, we refer to problem (1.2.1) as being *mildly ill-posed* if these singular values decay algebraically, or *severely ill-posed* if the singular values decay exponentially. In an analogous way to the classical approach, to a given smoothness class of the truth and a given degree of ill-posedness, we associate an optimal rate of convergence defined by the *minimax* criterion. The minimax rate is defined as the infimum over all statistical estimators $\hat{u}$, of the supremum over all true parameters $u^\dagger$ in a given smoothness class, of the mean integrated squared error (MISE) of $\hat{u}$, $\mathbb{E}^{y_n^\dagger}\left\|\hat{u} - u^\dagger\right\|^2$, [14]. A prior achieving the minimax rate in a family of smoothness classes indexed by γ, is considered optimal and is called *rate-adaptive* over the particular family: it achieves the optimal rate for a true parameter in the smoothness class γ, without knowledge of γ. It is obvious that such priors are highly desirable, but unfortunately they are very hard to find.

In this thesis the priors will not be adaptive, however we examine how the rate of contraction is affected by the choice of the parameters of the prior and show how a careful calibration of these parameters, based on at least a rough knowledge of the smoothness of the truth, leads to the minimax rates. As hinted earlier, this motivates the use of hierarchical priors, that is priors with hyper-parameters which are part of the inference.

Despite the rich literature in general nonparametric problems, the study of

the frequentist asymptotic properties of the posterior in the small noise limit in the context of inverse problems, has been underdeveloped until recently. In [34, 35], rates of convergence of the posterior distribution $\mu^{y_n^\dagger}$ to a Dirac distribution centered on the underlying true parameter $u^\dagger$ in the Ky-Fan metric are shown in linear inverse problems in finite dimensions with Gaussian prior and noise distributions, while in [58] an attempt to translate these convergence rates to infinite dimensions is undertaken, by first projecting the data to finite dimensions and then letting the dimension increase as the noise disappears in a carefully chosen way. The first paper to study posterior contraction in the sense of (1.2.2) in the infinite dimensional linear conjugate Gaussian setting of (1.2.1) is [44], in which sharp posterior contraction rates are obtained in a mildly ill-posed setup where K^*K and $\mathcal{C}_0$ are diagonalizable in the same eigenbasis with eigenvalues decaying algebraically and where η is Gaussian white noise. Our work in Chapter 2 (published as [3]) sidesteps the simultaneous diagonalizability assumption in [44] and allows for non-white Gaussian noise; when restricting to the mildly ill-posed diagonal case our rates agree (up to $\varepsilon > 0$ arbitrarily small) with the sharp rates in [44] when the truth is in a range of smoothness classes. In [45], sharp rates of posterior contraction are obtained in the linear severely ill-posed inverse problem of the recovery of the initial condition of the heat equation in a simultaneous diagonalizable setup similar to [44], while in our work in Chapter 3 (contained in [4]) we extend these sharp rates to more general severely ill-posed diagonal linear inverse problems. In [43, 79], posterior contraction rates are provided in hierarchical Bayesian diagonal mildly ill-posed linear inverse problems, with hyper-parameters introduced through the regularity and the scaling of the Gaussian prior respectively. Moreover, the methodology of [3] has been extended to prove posterior contraction rates in the problem of nonparametric drift estimation for diffusion processes in the large observation time limit in [64].

All of the results in the previous paragraph rely on the availability of an explicit description of the posterior distribution. The first consistency results based on general techniques in the context of inverse problems have appeared recently. In particular, posterior contraction results for linear inverse problems with Gaussian white noise and non-conjugate priors are considered in [65], while nonlinear inverse problems with a more general class of priors and Gaussian noise are studied in [86]. In both of these papers, when restricting to linear inverse problems with Gaussian priors, the posterior contraction rates obtained are in some cases slower than the ones in [44, 3, 45, 4]. Finally, consistency results in the Ky-Fan metric for linear inverse problems with non-Gaussian noise and priors in finite dimensions, are presented in [12].

1.2.1 Chapter 2 - Posterior contraction rates for the Bayesian approach to linear ill-posed inverse problems*

In this chapter we study the Bayesian linear inverse problem (1.2.1) with $K, \mathcal{C}_0$ and $\mathcal{C}_1$ which are related to each other through certain norm equivalence assumptions which formally express the fact that $K \simeq \mathcal{C}_0^{\ell}$ and $\mathcal{C}_1 \simeq \mathcal{C}_1^{\beta}$ for some $\ell, \beta \geq 0$. In particular, our assumptions allow cases where the three operators defining the problem are *not* simultaneously diagonalizable (see Chapter 2.8)*. While in the simultaneously diagonalizable setting, problem (1.2.1) is reduced to a countable set of uncoupled scalar problems, this is not the case here and the analysis is considerably harder.

We work in a functional setting defined through the prior covariance operator $\mathcal{C}_0$, in particular we work in the Hilbert scale $(X^t)_{t\in\mathbb{R}}$, where the Hilbert spaces X^t are roughly defined as the domains of $\mathcal{C}_0^{-\frac{t}{2}}$, [22]. This choice is natural since both the support and the Cameron-Martin space of the prior are spaces in this Hilbert scale. If the eigenvalues of $\mathcal{C}_0$ decay algebraically, then the spaces X^t can be thought of as rescaled Sobolev-type classes, while for exponential decay the spaces X^t correspond to classes of analytic functions, as defined in [14]. We do not assume that the problem is mildly nor severely ill-posed, however the assumed similarity between K and $\mathcal{C}_1$ to algebraic powers of $\mathcal{C}_0$, suggests that for priors which are supported in Sobolev-type smoothness classes the problem needs to be mildly ill-posed, while for priors which are supported in analytic smoothness classes the problem needs to be severely ill-posed (see Appendix A).

Our first contribution in this chapter is a new method for identifying the posterior which is a generalization to the separable Hilbert space setting of the *completion of the squares* technique natural in finite dimensions. As mentioned earlier, in the linear Gaussian setting of (1.2.1) it has been shown in [54, 52] that the posterior is Gaussian, $\mu^y = \mathcal{N}(m, \mathcal{C})$, and formulae for the posterior mean and covariance were provided. Our method results in alternative formulae for the mean and covariance which are expressed through the unbounded precision operator:

$$\frac{1}{n}\mathcal{C}^{-1} = K^*\mathcal{C}_1^{-1}K + \frac{1}{n\tau^2}\mathcal{C}_0^{-1}, \tag{1.2.3}$$

*Chapter 2 in the current thesis is [3].

*Note that in Chapter 2 we denote the forward operator as $\mathcal{A}^{-1}$ instead of K. This is because we have in mind that the observation operator is the inverse of a Schröndinger-type operator as in the examples in Chapter 2.8. For the same reason, we assume that $\mathcal{A}$ is self-adjoint and positive definite, however an inspection of our proofs suggests that our analysis directly generalizes to the case where the forward operator is bounded and injective.

$$\frac{1}{n}\mathcal{C}^{-1}m = K^*\mathcal{C}_1^{-1}y, \tag{1.2.4}$$

where equation (1.2.4) is interpreted in a weak sense in X^1. Our formulation of the posterior using precision operators has several advantages:

i) First, the precision operator formulation links directly to the theory of calculus of variations and the classical Tikhonov-Phillips regularization method. In particular, equations (1.2.3), (1.2.4) are the Euler-Lagrange equations for minimizing the generalized Tikhonov-Phillips functional

$$\mathcal{J}(u;y) = \frac{1}{2}\left\|\mathcal{C}_1^{-\frac{1}{2}}(y - Ku)\right\|^2 + \frac{\lambda}{2}\left\|\mathcal{C}_0^{-\frac{1}{2}}u\right\|^2,$$

where $\lambda = \frac{1}{n\tau^2}$ acts as the regularization parameter. This suggests that in the small noise limit $n \to \infty$, we need to have $\lambda \to 0$ in order to recover the truth.

ii) The posterior precision operator in (1.2.3) has the simple form of the prior precision operator $\mathcal{C}_0^{-1}$ plus the operator $K^*\mathcal{C}_1^{-1}K$ which in order for the prior to be regularizing we assume to be of lower order. If this is not the case, that is, if $K^*\mathcal{C}_1^{-1}K$ dominates in $\mathcal{C}^{-1}$, then $\frac{1}{n}\mathcal{C}^{-1} \simeq K^*\mathcal{C}_1K$ and the mean equation gives $m \simeq K^{-1}y$, that is, we try to directly invert the data. This simple form is useful because it splits the behaviour of the posterior to a leading order behaviour coming from the prior and a lower order correction due to the model. This splitting is useful when considering noise with non-vanishing size; we use this especially in Chapter 4 both in our intuition and our calculations.

iii) Working with the unbounded precision operators opens the possibility of using methods familiar from the theory of partial differential equations. In particular it enables the use of interpolation techniques to estimate blow-up rates for the operator $\frac{1}{n}\mathcal{C}^{-1}$ in a range of weak spaces X^t, $t \leq 0$ as $\lambda = \frac{1}{n\tau^2} \to 0$. This is useful for obtaining rates of posterior contraction as described below.

iv) Precision matrices corresponding to Markov processes can be sparse hence computation is efficient. Indeed, it is well known that for multivariate Gaussian distributions, the precision matrix relates to the correlation of the corresponding pair of variables given the rest. If a Markov property is assumed in the underlying stochastic process, this translates to conditional independence hence sparse precision matrices.

As part of our method, we formulate the posterior in the form (1.1.3), and show that the corresponding potential $\Phi(u;y)$ satisfies the assumptions of [78] for u and

y in certain spaces of full measure under the assumed prior and model. This secures that the posterior is well-posed and absolutely continuous with respect to the prior; it also secures that the posterior is Lipschitz continuous in the Hellinger metric with respect to the data and with respect to finite dimensional approximations of the forward problem.

Our second contribution in Chapter 2 is that we prove rates of contraction in this more general more difficult non-diagonal setting. We also prove rates of convergence of the mean squared error of the posterior mean in a range of spaces in the Hilbert scale $(X^t)_{t\in\mathbb{R}}$. We now present a brief heuristic for the convergence of the MISE of the posterior mean. We assume that we have data of the form $y_n^\dagger = Ku^\dagger + \frac{1}{\sqrt{n}}\eta$, hence the posterior is $\mu^{y=y_n^\dagger} = \mathcal{N}(m^\dagger, \mathcal{C})$, where $m^\dagger$ satisfies (1.2.4) with $y = y_n^\dagger$ and $\mathcal{C}$ satisfies (1.2.3). That is, we have

$$\frac{1}{n}\mathcal{C}^{-1}m^\dagger = K^*\mathcal{C}_1^{-1}Ku^\dagger + \frac{1}{\sqrt{n}}K^*\mathcal{C}_1^{-1}\eta. \tag{1.2.5}$$

Moreover, by the definition of $\mathcal{C}$, the true solution $u^\dagger$ satisfies the equation

$$\frac{1}{n}\mathcal{C}^{-1}u^\dagger = K^*\mathcal{C}_1^{-1}Ku^\dagger + \frac{1}{n\tau^2}\mathcal{C}_0^{-1}u^\dagger, \tag{1.2.6}$$

hence subtracting we get the equation for the error $e = m^\dagger - u^\dagger$,

$$e = n\mathcal{C}\Big(\frac{1}{\sqrt{n}}K^*\mathcal{C}_1^{-1}\eta - \frac{1}{n\tau^2}\mathcal{C}_0^{-1}u^\dagger\Big). \tag{1.2.7}$$

As $n \to \infty$ and for $\lambda = \frac{1}{n\tau^2} \to 0$, the expectations of sufficiently weak norms of the two terms in the parenthesis go to zero. On the other hand using interpolation techniques we can estimate the blow-up rates of the $n\mathcal{C}$ as $\lambda = \frac{1}{n\tau^2} \to 0$ as a linear operator mapping these spaces to $\mathcal{X}$. Comparing the decay and the blow-up rates yields the rate of convergence of the MISE of the posterior mean which has two contributions: one because of the presence of the noise which, for fixed n, blows-up as $\lambda \to 0$ and one because of the regularizing effect of the prior which, for fixed n, goes to zero as $\lambda \to 0$. We optimize the rate by balancing the two contributions by choosing λ (and hence τ) as an appropriate function of n.

The work in this chapter is in collaboration with Stig Larsson (Chalmers University of Technology) and my PhD supervisor Andrew Stuart (University of Warwick) and is published in [3]. The problem formulation and methodology were developed in conjunction with my two co-authors, while almost all of the technical analysis was carried out by myself.

1.2.2 Chapter 3 - Bayesian posterior contraction rates for linear severely ill-posed inverse problems*

In this chapter we study the Bayesian linear inverse problem (1.2.1), in a setting in which the operators $K^*K, \mathcal{C}_0$ and $\mathcal{C}_1$ are simultaneously diagonalizable and have eigenvalues which decay like $\exp(-2sj^b), j^{-2\alpha}$ and $j^{-2\beta}$, respectively, for $s, b > 0$, $\alpha > \frac{1}{2}$ and $\beta \geq 0$. That is, we consider a family of severely ill-posed linear inverse problems which includes for example the problem of the recovery of the initial condition of the heat equation considered in [45] (corresponding to $b = 2$), or the Cauchy problem for the Helmholtz equation (corresponding to $b = 1$, see Chapter 3.5).

Because of the simultaneous diagonalizability assumption, this problem is reduced to countable uncoupled scalar inverse problems; this enables the sharp calculation of posterior contraction rates as in [45]. Both the problem formulation and the convergence analysis are inspired by [45], however this more general setup leads to some technical improvements in the proofs. Furthermore, we provide new results on the absolute continuity of the posterior with respect to the prior. In particular, we show that in this severely ill-posed case the posterior is absolutely continuous with respect to the prior almost surely with respect to the joint distribution of (u, y), independently of the particular values of α, β, b and s. This is not a trivial statement; we demonstrate this by showing that in the mildly ill-posed case where the eigenvalues of K^*K decay as $j^{-4\ell}$ for $\ell \geq 0$, there are combinations of α, β and ℓ which are such that the (Gaussian) posterior is mutually singular with the (Gaussian) prior independently of the data; in particular this happens if the prior is not sufficiently regular. Finally, we include a numerical simulation of the Cauchy problem for the Helmholtz equation.

The work in this chapter is in collaboration with my PhD supervisor Andrew Stuart (University of Warwick) and Yuan-Xiang Zhang (Lanzhou University). My main contribution in this chapter was to use my experience from Chapter 2 to guide Yuan-Xiang in carrying out the technical convergence analysis, while I was more actively involved in the measure-theoretic considerations and the implementation of the numerical example.

1.3 Sampling the posterior

We now turn our attention to the problem of efficiently sampling the posterior. Suppose we have a possibly unnormalized finite measure π on a space $\mathcal{W}$. The

*Chapter 3 in the current thesis is [4].

most widely applied method for drawing samples from π is to run a MCMC algorithm which produces a Markov chain $\{w^{(1)}, ..., w^{(k)}\}$ constructed to have π as its stationary distribution. Provided the algorithm satisfies certain assumptions, we have that when started from π-almost any point $w \in \mathcal{W}$ the Markov chain converges to its stationary distribution; we then say that the algorithm is *ergodic*, [71]. Ergodicity guarantees that if we run the algorithm for long enough, the samples $w^{(i)}$ are drawn approximately from the target distribution π. Naturally, different algorithms have different convergence properties and in particular different speeds of convergence, which are quantified using various notions of ergodicity [57, 71, 68]. Furthermore, even after the algorithm reaches stationarity, it is desirable that the produced Markov chain explores the target distribution π as quickly as possible, so that less samples are required to extract accurate information; this relates to the rate of decorrelation between samples.

1.3.1 Sampling in function space - diffusion limits and spectral gaps

In our context of Bayesian inverse problems in function space, the target measure π is the posterior measure μ^y on the function space $\mathcal{X}$ which we assume to have a density with respect to a Gaussian prior measure μ_0 as expressed in (1.1.3). In practice the problem is discretized and the MCMC algorithm is implemented in $\mathbb{R}^N$; we are interested in understanding the behaviour of the algorithm as we refine the model, $N \to \infty$. In this context, two theories have been developed in order to understand the advantages of different algorithms: the first is by proving diffusion (scaling) limits of the algorithms, while the second is to establish spectral gaps.

In the diffusion limit approach, it is shown that an appropriately scaled continuous time interpolant of the Markov chain converges weakly, in the infinite dimensional limit $N \to \infty$, to the solution to a stochastic partial differential equation (SPDE) which has π as stationary distribution. Since the SPDE needs a finite time T to explore its stationary distribution, this suggests that the time the Markov chain needs to explore π (or more accurately an approximation of π in $\mathbb{R}^N$) is inversely proportional to the time-step required to get the diffusion limit; the bigger the required time-step the faster the algorithm explores the target distribution with the algorithms which are optimal in this sense being the ones which require no scaling at all. This technique was pioneered by [67, 69, 70] for target distributions of product form and recently extended to targets of the type (1.1.3) in [11, 55, 63]. The intuition obtained from the study of diffusion limits led to the design of algorithms which are well-defined in the infinite dimensional limit and hence are robust with respect to the increase in dimension, [10, 62, 15].

Spectral gaps relate to the geometric ergodicity of Markov chains, [68], and hence to the speed of convergence of an algorithm to the target distribution. In [31], the large dimensional behaviour of the Random Walk Metropolis (RWM) and preconditioned Crank Nicolson (pCN) algorithms was investigated using this theory. It was shown that the RWM algorithm, which is not well-defined in the function space setting, has a dimension-dependent spectral gap implying the deterioration of its rate of convergence to equilibrium, while the pCN algorithm which is well-defined in the infinite dimensional limit, has a dimension independent spectral gap and hence its rate of convergence to equilibrium is robust with respect to the increase in dimension. The application of the spectral gap theory was made possible through recent developments in the theory of Markov chains in infinite dimensions, [30], and it is expected that this method will also be useful for analyzing the large dimensional behaviour of other MCMC algorithms in the near future.

In this thesis we are interested in understanding the large dimensional behaviour of Gibbs samplers naturally arising in hierarchical Bayesian inverse problems in conjugate Gaussian settings; this is the topic of Chapter 4 introduced in the next subsection. We apply infinite dimensional intuition and use a variant of the diffusion limit approach.

1.3.2 Chapter 4 - Dimension dependence of sampling algorithms in hierarchical Bayesian inverse problems

In this chapter we consider a hierarchical variant of the Bayesian linear inverse problem (1.2.1), with hyper-priors introduced through the scalings of the noise and prior covariance operators. In particular, we consider the model

$$y = Ku + \eta, \tag{1.3.1}$$

where $\eta|\sigma \sim \mathcal{N}(0, \sigma^{-1}\mathcal{C}_1)$ and $\sigma \sim \text{Gamma}(\alpha_1, \beta_1)$. We put a mixture prior on u, $u|\delta \sim \mathcal{N}(0, \delta^{-1}\mathcal{C}_0)$ where $\delta \sim \text{Gamma}(\alpha_0, \beta_0)$.

In practice the problem is discretized and the above setup is implemented in $\mathbb{R}^N$, where it is desirable to refine the model by letting $N \to \infty$, [7]. To this end, we assume that we have a way of computing discretizations of the data $\boldsymbol{y_N} \in \mathbb{R}^N$ and we replace the operators $K, \mathcal{C}_0, \mathcal{C}_1$ by $N \times N$ matrices arising from consistent discretizations of the corresponding operators in the underlying Hilbert space $\mathcal{X}$. It is well known that this hierarchical Bayesian model is conditionally conjugate, that is $u|\boldsymbol{y_N}, \delta, \sigma$ is Gaussian and $\delta|\boldsymbol{y_N}, u, \sigma$ and $\sigma|\boldsymbol{y_N}, u, \delta$ are Gamma; this makes natural the use of a Gibbs sampler, which draws from the three conditional distributions

in turn, in order to sample the full posterior on $u, \delta, \sigma | \boldsymbol{y_N}$ in $\mathbb{R}^N \times \mathbb{R} \times \mathbb{R}$. We are interested in understanding the large N behaviour of this Gibbs sampler.

We remark that the use of consistent discretizations secures the interpretability of the scaling parameters δ^{-1} and σ^{-1} in the limit $N \to \infty$, as the scalings of the corresponding covariance operators in the underlying infinite dimensional setup (1.3.1). If as in [7] we do not use consistent discretizations, we cannot make meaningful comparisons of δ and σ across different discretization levels, and in particular it is not natural to examine their limiting behaviour as $N \to \infty$. This is very important especially in applied problems where often simulations are performed at different discretization levels in order to check the robustness to the numerical approximations used.

Our main results in this chapter suggest that under natural assumptions, as the dimension N grows, the behaviour of the algorithm has two scales: an increasingly fast scale on which the scaling of the noise evolves under the Gibbs sampler, and an increasingly slow scale on which the scaling of the prior evolves under the Gibbs sampler.

We now briefly describe our intuition based on infinite dimensional arguments and in particular the following two properties of Gaussian measures in separable Hilbert spaces:

- two centered Gaussian measures with covariance operators which are proportional to each other are mutually singular unless the constant of proportionality is equal to one;

- a complete path drawn from a Gaussian measure $\mathcal{N}(0, a\Sigma)$ contains full information about the scaling parameter a.

We work under the natural assumption that in the infinite dimensional underlying model the conditional posterior on $u|y, \delta, \sigma$ is absolutely continuous with respect to the prior $u|\delta$. By the above properties of infinite dimensional Gaussian measures this assumption suggests that a draw from $u|y, \delta, \sigma$ contains full information on the value of δ, hence as $N \to \infty$ there is a strong dependence between $u|\boldsymbol{y_N}, \delta, \sigma$ and $\delta|\boldsymbol{y_N}, u, \sigma$ leading to strong dependence between successive δ-draws. Indeed, we show under assumptions securing the reasonable behaviour of the discretizations used, that as the dimension N grows, the δ-chain makes moves which on average are of order N^{-1} with fluctuations of order $N^{-\frac{1}{2}}$. This implies that it takes $\mathcal{O}(N)$ steps for the δ-chain to move $\mathcal{O}(1)$ distance and in turn suggests that it takes $\mathcal{O}(N)$ steps for the Gibbs sampler to sample the posterior on $u, \delta, \sigma | \boldsymbol{y_N}$.

In order to analyze the behaviour of the σ-chain, we assume that we have data produced as perturbations of a sufficiently smooth $z \in \mathcal{X}$ by a realization of the noise distribution with a fixed scaling $\bar{\sigma}$; that is, we assume

$$y = z + \bar{\sigma}^{-\frac{1}{2}} \mathcal{C}_1^{\frac{1}{2}} \xi,$$

where ξ is a realization of a Gaussian white noise. The assumption on the regularity of z suggests that the data producing measure $\mathcal{N}(z, \bar{\sigma}^{-1}\mathcal{C}_1)$ is absolutely continuous with respect to the noise distribution $\eta|\bar{\sigma}$. Again by the properties of infinite dimensional Gaussian measures we have that the data contain full information on the value of $\bar{\sigma}$, hence as $N \to \infty$ there is a strong dependence between the σ-draw and the data leading to instant identification of $\bar{\sigma}$ by the σ-chain. Indeed, we show that as the dimension N grows, the σ-chain makes moves which on average are within order N^{-1} distance from $\bar{\sigma}$ with fluctuations of order $N^{-\frac{1}{2}}$.

Based on intuition from [72, 60], we propose a reparametrization of the prior scaling in problem (1.3.1) in order to alleviate the poor mixing of the δ-chain, in which the two components on the unknown and the prior scaling are a priori independent. That is, instead of assuming $u|\delta \sim \mathcal{N}(0, \delta^{-1}\mathcal{C}_0)$ where $\delta \sim \text{Gamma}(\alpha_0, \beta_0)$, we write $u = \tau v$, where $v \sim \mathcal{N}(0, \mathcal{C}_0)$ and $\tau \sim \mathcal{N}(r_0, q_0^2)$. This setup is again conditionally conjugate, hence we again use a Gibbs sampler to sample the full posterior on $u, \tau, \sigma|\boldsymbol{y_N}$. While the reparametrized algorithm is robust with respect to the increase in dimension, it deteriorates in the small noise limit. The reason is that as the noise disappears, v and τ are a posteriori increasingly dependent because they both need to explain the data; more research is required in this small noise limit.

We also extend our results on the slowing down of the δ-chain to other Gaussian conjugate settings, such as the setting of nonparametric drift estimation of SDE's considered in [64, 59, 56]. Furthermore, our theory generalizes to cases where the discretization level of the unknown is different to the discretization level on the data; in this case the slowing down of the δ-chain occurs as the discretization level of the *unknown* increases, while the speeding up of the σ-chain occurs as the discretization level of the *data* increases.

We provide numerical simulations in several linear inverse problems settings which support our theory regarding the behaviour of both the standard hierarchical algorithm and the proposed reparametrization. Finally, we remark that our intuition also applies when attempting to learn other parameters of the noise and prior distributions, as for example in [43].

The work in this chapter is in collaboration with Johnathan Bardsley (Uni-

versity of Montana), Omiros Papaspiliopoulos (Universitat Pompeu Fabra) and my PhD supervisor Andrew Stuart (University of Warwick). The problem formulation and methodology were developed in conjunction with my three co-authors, while almost all of the technical analysis was carried out by myself. All the numerical simulations were performed by myself based on modifications of the code used by Johnathan Bardsley in [7].

Chapter 2

Posterior Contraction Rates for the Bayesian Approach to Linear Ill-Posed Inverse Problems

2.1 Introduction

The solution of inverse problems provides a rich source of applications of the Bayesian nonparametric methodology. It encompasses a broad range of applications from partial differential equations (PDEs) [6], where there is a well-developed theory of classical, non-statistical, regularization [22]. On the other hand, the area of nonparametric Bayesian statistical estimation and in particular the problem of posterior consistency has attracted a lot of interest in recent years; see for instance [28, 77, 74, 83, 82, 29, 20]. Despite this, the formulation of many of these PDE inverse problems using the Bayesian approach is in its infancy [78]. Furthermore, the development of a theory of Bayesian posterior consistency, analogous to the theory for classical regularization, is under-developed with the primary contribution being the recent paper [44]. This recent paper provides a roadmap for what is to be expected regarding Bayesian posterior consistency, but is limited in terms of applicability by the assumption of simultaneous diagonalizability of the three linear operators required to define Bayesian inversion. Our aim in this chapter is to make a significant step in the theory of Bayesian posterior consistency for linear inverse problems by developing a methodology which sidesteps the need for simultaneous diagonalizability. The central idea underlying the analysis is to work with precision

operators rather than covariance operators, and thereby to enable use of powerful tools from PDE theory to facilitate the analysis.

Let $\mathcal{X}$ be a separable Hilbert space, with norm $\|\cdot\|$ and inner product $\langle\cdot,\cdot\rangle$, and let $\mathcal{A}\colon \mathcal{D}(\mathcal{A}) \subset \mathcal{X} \to \mathcal{X}$ be a known self-adjoint and positive-definite linear operator with bounded inverse*. We consider the inverse problem to find u from y, where y is a noisy observation of $\mathcal{A}^{-1}u$. We assume the model,

$$y = \mathcal{A}^{-1}u + \frac{1}{\sqrt{n}}\eta, \tag{2.1.1}$$

where $\frac{1}{\sqrt{n}}\eta$ is an additive noise. We will be particularly interested in the small noise limit where $n \to \infty$.

A popular method in the deterministic approach to inverse problems is the generalized Tikhonov-Phillips regularization method in which u is approximated by the minimizer of a regularized least squares functional: define the Tikhonov-Phillips functional

$$\mathcal{J}_0(u) := \frac{1}{2}\left\|\mathcal{C}_1^{-\frac{1}{2}}(y - \mathcal{A}^{-1}u)\right\|^2 + \frac{\lambda}{2}\left\|\mathcal{C}_0^{-\frac{1}{2}}u\right\|^2, \tag{2.1.2}$$

where $\mathcal{C}_i\colon \mathcal{X} \to \mathcal{X},\ i = 0, 1$, are bounded, possibly compact, self-adjoint positive-definite linear operators. The parameter λ is called the regularization parameter, and in the classical non-probabilistic approach the general practice is to choose it as an appropriate function of the noise size $n^{-\frac{1}{2}}$, which shrinks to zero as $n \to \infty$, in order to recover the unknown parameter u [22].

In this chapter we adopt a Bayesian approach for the solution of problem (2.1.1), which will be linked to the minimization of $\mathcal{J}_0$ via the posterior mean. We assume that the prior distribution is Gaussian, $u \sim \mu_0 = \mathcal{N}(0, \tau^2\mathcal{C}_0)$, where $\tau > 0$ and $\mathcal{C}_0$ is a self-adjoint, positive-definite, trace class, linear operator on $\mathcal{X}$. We also assume that the noise is Gaussian, $\eta \sim \mathcal{N}(0, \mathcal{C}_1)$, where $\mathcal{C}_1$ is a self-adjoint positive-definite, bounded, but not necessarily trace class, linear operator; this allows us to include the case of white observational noise. We assume that the, generally unbounded, operators $\mathcal{C}_0^{-1}$ and $\mathcal{C}_1^{-1}$, have been maximally extended to self-adjoint positive-definite operators on appropriate domains. The unknown parameter and the noise are considered to be independent, thus the conditional distribution of the observation given the unknown parameter u (termed the likelihood) is also Gaussian with distribution $y|u \sim \mathcal{N}(\mathcal{A}^{-1}u, \frac{1}{n}\mathcal{C}_1)$.

*In fact our analysis directly generalizes to the case where the forward operator is bounded and injective, see Section 2.10.

Define $\lambda = \frac{1}{n\tau^2}$ and let

$$\mathcal{J}(u) = n\mathcal{J}_0(u) = \frac{n}{2}\left\|\mathcal{C}_1^{-\frac{1}{2}}(y - \mathcal{A}^{-1}u)\right\|^2 + \frac{1}{2\tau^2}\left\|\mathcal{C}_0^{-\frac{1}{2}}u\right\|^2. \tag{2.1.3}$$

In finite dimensions the probability density of the posterior distribution, that is, the distribution of the unknown given the observation, with respect to the Lebesgue measure is proportional to $\exp(-\mathcal{J}(u))$. This suggests that, in the infinite-dimensional setting, the posterior is Gaussian $\mu^y = \mathcal{N}(m, \mathcal{C})$, where we can identify the posterior covariance and mean by the equations

$$\mathcal{C}^{-1} = n\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1} + \frac{1}{\tau^2}\mathcal{C}_0^{-1} \tag{2.1.4}$$

and

$$\frac{1}{n}\mathcal{C}^{-1}m = \mathcal{A}^{-1}\mathcal{C}_1^{-1}y, \tag{2.1.5}$$

obtained by completing the square. We present a method of justifying these expressions in Section 2.5. We define

$$\mathcal{B}_\lambda = \frac{1}{n}\mathcal{C}^{-1} = \mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1} + \lambda\mathcal{C}_0^{-1} \tag{2.1.6}$$

and observe that the dependence of $\mathcal{B}_\lambda$ on n and τ is only through λ. Since

$$\mathcal{B}_\lambda m = \mathcal{A}^{-1}\mathcal{C}_1^{-1}y, \tag{2.1.7}$$

the linear operator mapping the data to the posterior mean also depends only on λ: $m(y) = m_\lambda(y)$. This is not the case for the posterior covariance $\mathcal{C}$, since it depends on n and τ separately: $\mathcal{C} = \mathcal{C}_{\lambda,n}$. In the following, we suppress the dependence of the posterior mean on λ and the posterior covariance on λ and n and write just m and $\mathcal{C}$.

Observe that the posterior mean is the minimizer of the functional $\mathcal{J}$, hence also of $\mathcal{J}_0$, that is, the posterior mean is the Tikhonov-Phillips regularized approximate solution of problem (2.1.1), for the functional $\mathcal{J}_0$ with $\lambda = \frac{1}{n\tau^2}$.

In [54] and [52], formulae for the posterior covariance and mean are identified in the infinite-dimensional setting, which avoid using any of the inverses of the prior, posterior or noise covariance operators. They obtain

$$\mathcal{C} = \tau^2\mathcal{C}_0 - \tau^2\mathcal{C}_0\mathcal{A}^{-1}(\mathcal{A}^{-1}\mathcal{C}_0\mathcal{A}^{-1} + \lambda\mathcal{C}_1)^{-1}\mathcal{A}^{-1}\mathcal{C}_0 \tag{2.1.8}$$

and

$$m = \mathcal{C}_0\mathcal{A}^{-1}(\mathcal{A}^{-1}\mathcal{C}_0\mathcal{A}^{-1} + \lambda\mathcal{C}_1)^{-1}y, \tag{2.1.9}$$

which are consistent with formulae (2.1.4) and (2.1.7) for the finite-dimensional case. In [54] this is done only for $\mathcal{C}_1$ of trace class while in [52] the case of white observational noise was included. We will work in an infinite-dimensional setting where the formulae (2.1.4), (2.1.7) for the posterior covariance and mean can be justified. Working with the unbounded operator $\mathcal{B}_\lambda$ opens the possibility of using tools of analysis, and also numerical analysis, familiar from the theory of partial differential equations.

In our analysis we always assume that $\mathcal{C}_0^{-1}$ is regularizing, that is, we assume that $\mathcal{C}_0^{-1}$ dominates $\mathcal{B}_\lambda$ in the sense that it induces stronger norms than $\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}$. This is a reasonable assumption since otherwise we would have $\mathcal{B}_\lambda \simeq \mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}$ (here $\simeq$ is used loosely to indicate two operators which induce equivalent norms; we will make this notion precise in due course). This would imply that the posterior mean is $m \simeq \mathcal{A}y$, meaning that we attempt to invert the data by applying the, generally discontinuous, operator $\mathcal{A}$ [22, Proposition 2.7].

We study the consistency of the posterior μ^y in the frequentist setting. To this end, we consider data $y = y_n^\dagger$ which is a realization of

$$y_n^\dagger = \mathcal{A}^{-1}u^\dagger + \frac{1}{\sqrt{n}}\eta, \quad \eta \sim \mathcal{N}(0, \mathcal{C}_1), \tag{2.1.10}$$

where $u^\dagger$ is a fixed element of $\mathcal{X}$; that is, we consider a sequence of observations $\{y_n^\dagger\}$, where $y_n^\dagger$ is a perturbation of the image of a fixed true solution $u^\dagger$ by an additive noise η, scaled by $\frac{1}{\sqrt{n}}$. Since the posterior depends through its mean on the data and also through its covariance operator on the scaling of the noise and the prior, this choice of data model gives as posterior distribution the Gaussian measure $\mu_{\lambda,n}^{y_n^\dagger} = \mathcal{N}(m^\dagger, \mathcal{C})$, where $\mathcal{C}$ is given by (2.1.4) and

$$\mathcal{B}_\lambda m^\dagger = \mathcal{A}^{-1}\mathcal{C}_1^{-1}y_n^\dagger. \tag{2.1.11}$$

Note that the posterior mean $m^\dagger$ now also depends on n through the assumed data. We study the behaviour of the posterior $\mu_{\lambda,n}^{y_n^\dagger}$ as the noise disappears ($n \to \infty$). Our aim is to show that it contracts to a Dirac measure centered on the fixed true solution $u^\dagger$. In particular, we aim to determine ϵ_n such that

$$\mathbb{E}^{y_n^\dagger}\mu_{\lambda,n}^{y_n^\dagger}\left\{u : \left\|u - u^\dagger\right\| \geq M_n\epsilon_n\right\} \to 0, \quad \forall M_n \to \infty, \tag{2.1.12}$$

where the expectation is with respect to the random variable $y_n^\dagger$ distributed according to the data producing measure $\mathcal{N}(\mathcal{A}^{-1}u^\dagger, \frac{1}{n}\mathcal{C}_1)$.

As in the deterministic theory of inverse problems, in order to get convergence in the small noise limit, we let the regularization disappear in a carefully chosen way, that is, we will choose $\lambda = \lambda(n)$ such that $\lambda \to 0$ as $n \to \infty$. The assumption that $\mathcal{C}_0^{-1}$ dominates $\mathcal{B}_\lambda$, shows that $\mathcal{B}_\lambda$ is a singularly perturbed unbounded (usually differential) operator, with an inverse which blows-up in the limit $\lambda \to 0$. This together with equation (2.1.7), opens up the possibility of using the analysis of such singular limits to study posterior contraction: on the one hand, as $\lambda \to 0$, $\mathcal{B}_\lambda^{-1}$ becomes unbounded; on the other hand, as $n \to \infty$, we have more accurate data, suggesting that for the appropriate choice of $\lambda = \lambda(n)$ we can get $m^\dagger \simeq u^\dagger$. In particular, we will choose τ as a function of the scaling of the noise, $\tau = \tau(n)$, under the restriction that the induced choice of $\lambda = \lambda(n) = \frac{1}{n\tau(n)^2}$, is such that $\lambda \to 0$ as $n \to \infty$. The last choice will be made in a way which optimizes the rate of posterior contraction ϵ_n, defined in (2.1.12). In general there are three possible asymptotic behaviours of the scaling of the prior τ^2 as $n \to \infty$, [82, 44]:

i) $\tau^2 \to \infty$; we increase the prior spread, if we know that draws from the prior are more regular than $u^\dagger$;

ii) τ^2 fixed; draws from the prior have the same regularity as $u^\dagger$;

iii) $\tau^2 \to 0$ at a rate slower than $\frac{1}{n}$; we shrink the prior spread, when we know that draws from the prior are less regular than $u^\dagger$.

The problem of posterior contraction in this context is also investigated in [44] and [24]. In [44], sharp convergence rates are obtained in the case where $\mathcal{C}_0, \mathcal{C}_1$ and $\mathcal{A}^{-1}$ are simultaneously diagonalizable, with eigenvalues decaying algebraically, and in particular $\mathcal{C}_1 = I$, that is, the data are polluted by white noise. In this chapter we relax the assumptions on the relations between the operators $\mathcal{C}_0, \mathcal{C}_1$ and $\mathcal{A}^{-1}$, by assuming that appropriate powers of them induce comparable norms (see Section 2.3). In [24], the non-diagonal case is also examined; the three operators involved are related through domain inclusion assumptions. The assumptions made in [24] can be quite restrictive in practice; our assumptions include settings not covered in [24], and in particular the case of white observational noise.

2.1.1 Outline of the rest of the chapter

In the following section we present our main results which concern the identification of the posterior (Theorem 2.2.1) and the posterior contraction (Theorems 2.2.2

and 2.2.3). In Section 2.3 we present our assumptions and their implications. The proofs of the main results are built in a series of intermediate results contained in Sections 2.4-2.7. In Section 2.4, we reformulate equation (2.1.7) as a weak equation in an infinite-dimensional space. In Section 2.5, we present a new method of identifying the posterior distribution: we first characterize it through its Radon-Nikodym derivative with respect to the prior (Theorem 2.5.1) and then justify the formulae (2.1.4), (2.1.7) for the posterior covariance and mean (proof of Theorem 2.2.1). In Section 2.6, we present operator norm bounds for $\mathcal{B}_\lambda^{-1}$ in terms of the singular parameter λ, which are the key to the posterior contraction results contained in Section 2.7 and their corollaries in Section 2.2 (Theorems 2.7.1, 2.7.2 and 2.2.2, 2.2.3). In Section 2.8, we present some nontrivial examples satisfying our assumptions and provide the corresponding rates of convergence. In Section 2.9, we compare our results to known minimax rates of convergence in the case where $\mathcal{C}_0, \mathcal{C}_1$ and $\mathcal{A}^{-1}$ are all diagonalizable in the same eigenbasis and have eigenvalues that decay algebraically. Finally, Section 2.10 is a short conclusion.

The entire chapter rests on a rich set of connections between the theory of stochastic processes and various aspects of the theory of linear partial differential equations. In particular, since the Green's function of the precision operator of a Gaussian measure corresponds to its covariance function, our formulation and analysis of the inverse problem via precision operators is very natural. Furthermore, estimates on the inverse of singular limits of these precisions, which have direct implications for localization of the Green's functions, play a key role in the analysis of posterior consistency.

2.2 Main Results

In this section we present our main results. We postpone the rigorous presentation of our assumptions to the next section and the proofs and technical lemmas are presented together with intermediate results of independent interest in Sections 2.4 - 2.7. Recall that we assume a Gaussian prior $\mu_0 = \mathcal{N}(0, \tau^2\mathcal{C}_0)$ and a Gaussian noise distribution $\mathcal{N}(0, \mathcal{C}_1)$. Our first assumption concerns the decay of the eigenvalues of the prior covariance operator and enables us to quantify the regularity of draws from the prior. This is encoded in the parameter $s_0 \in [0, 1)$; smaller s_0 implies more regular draws from the prior. We also assume that $\mathcal{C}_1 \simeq \mathcal{C}_0^\beta$ and $\mathcal{A}^{-1} \simeq \mathcal{C}_0^\ell$, for some $\beta, \ell \geq 0$, where $\simeq$ is used in the manner outlined in Section 2.1, and defined in detail in Section 2.3. Finally, we assume that the problem is sufficiently ill-posed with respect to the prior. This is quantified by the parameter $\Delta := 2\ell - \beta + 1$ which

we assume to be larger than $2s_0$; for a fixed prior, the larger Δ is, the more ill-posed the problem.

2.2.1 Posterior Identification

Our first main theorem identifies the posterior measure as Gaussian and justifies formulae (2.1.4) and (2.1.7). This reformulation of the posterior in terms of the precision operator is key to our method of analysis of posterior consistency and opens the route to using methods from the study of partial differential equations (PDEs). These methods will also be useful for the development of numerical methods for the inverse problem.

Theorem 2.2.1. *Under the Assumptions 2.3.1, the posterior measure $\mu^y(du)$ is Gaussian $\mu^y = \mathcal{N}(m, \mathcal{C})$, where $\mathcal{C}$ is given by (2.1.4) and m is a weak solution of (2.1.7).* □

2.2.2 Posterior Contraction

We now present our results concerning frequentist posterior consistency of the Bayesian solution to the inverse problem. We assume to have data $y = y_n^\dagger$ as in (2.1.10), and examine the behaviour of the posterior $\mu_{\lambda,n}^{y_n^\dagger} = \mathcal{N}(m^\dagger, \mathcal{C})$, where $m^\dagger$ is given by (2.1.11), as the noise disappears ($n \to \infty$). The first convergence result concerns the convergence of the posterior mean $m^\dagger$ to the true solution $u^\dagger$ in a range of weighted norms $\|\cdot\|_\kappa$ induced by powers of the prior covariance operator $\mathcal{C}_0$. The spaces $(X^\kappa, \|\cdot\|_\kappa)$ are rigorously defined in the following section. The second result provides rates of posterior contraction of the posterior measure to a Dirac measure centered on the true solution as described in (2.1.12). In both results, we assume a priori known regularity of the true solution $u^\dagger \in X^\gamma$ and give the convergence rates as functions of γ.

Theorem 2.2.2. *Assume $u^\dagger \in X^\gamma$, where $\gamma \geq 1$ and let $\kappa = (1-\theta)(\beta - 2\ell) + \theta$, where $\theta \in [0,1]$. Under the Assumptions 2.3.1, we have the following optimized rates of convergence, where $\varepsilon > 0$ is arbitrarily small:*

i) if $\gamma \in (1, \Delta + 1]$, for $\tau = \tau(n) = n^{-\frac{\gamma - 1 + s_0 + \varepsilon}{2(\Delta + \gamma - 1 + s_0 + \varepsilon)}}$

$$\mathbb{E}^{y_n^\dagger} \left\| m^\dagger - u^\dagger \right\|_\kappa^2 \leq c n^{-\frac{\Delta + \gamma - 1 - \theta\Delta}{\Delta + \gamma - 1 + s_0 + \varepsilon}};$$

ii) if $\gamma > \Delta + 1$, for $\tau = \tau(n) = n^{-\frac{\Delta + s_0 + \varepsilon}{2(2\Delta + s_0 + \varepsilon)}}$

$$\mathbb{E}^{y_n^\dagger} \left\| m^\dagger - u^\dagger \right\|_\kappa^2 \leq c n^{-\frac{(2-\theta)\Delta}{2\Delta + s_0 + \varepsilon}};$$

iii) if $\gamma = 1$ *and* $\theta \in [0,1)$ *for* $\tau = \tau(n) = n^{-\frac{s_0+\varepsilon}{2(\Delta+s_0+\varepsilon)}}$

$$\mathbb{E}^{y_n^\dagger}\left\|m^\dagger - u^\dagger\right\|_\kappa^2 \leq cn^{-\frac{(1-\theta)\Delta}{\Delta+s_0+\varepsilon}}.$$

If $\gamma = 1$ *and* $\theta = 1$ *then the method does not give convergence.*

□

Theorem 2.2.3. *Assume* $u^\dagger \in X^\gamma$*, where* $\gamma \geq 1$*. Under the Assumptions 2.3.1, we have the following optimized rates for the convergence in (2.1.12), where* $\varepsilon > 0$ *is arbitrarily small:*

i) if $\gamma \in [1, \Delta+1]$ *for* $\tau = \tau(n) = n^{-\frac{\gamma-1+s_0+\varepsilon}{2(\Delta+\gamma-1+s_0+\varepsilon)}}$

$$\epsilon_n = \begin{cases} n^{-\frac{\gamma}{2(\Delta+\gamma-1+s_0+\varepsilon)}}, & if\ \beta - 2\ell \leq 0 \\ n^{-\frac{\Delta+\gamma-1}{2(\Delta+\gamma-1+s_0+\varepsilon)}}, & otherwise; \end{cases}$$

ii) if $\gamma > \Delta+1$ *for* $\tau = \tau(n) = n^{-\frac{\Delta+s_0+\varepsilon}{2(2\Delta+s_0+\varepsilon)}}$

$$\epsilon_n = \begin{cases} n^{-\frac{\Delta+1}{2(2\Delta+s_0+\varepsilon)}}, & if\ \beta - 2\ell \leq 0 \\ n^{-\frac{\Delta}{2\Delta+s_0+\varepsilon}}, & otherwise. \end{cases}$$

□

To summarize, provided the problem is sufficiently ill-posed and the true solution $u^\dagger$ is sufficiently regular we get the convergence in (2.1.12) for

$$\epsilon_n = n^{-\frac{\gamma\wedge(\Delta+1)}{2(\Delta+\gamma\wedge(\Delta+1)-1+s_0+\varepsilon)}}.$$

Our rates of convergence agree, up to $\varepsilon > 0$ arbitrarily small, with the sharp convergence rates obtained in the diagonal case in [44] across a wide range of regularity assumptions on the true solution (Figure 2.1); yet, our rates cover a much more applicable range of non-simultaneously diagonalizable problems. (The reason for the appearance of ε is that in the assumed non-diagonal setting we can only use information about the regularity of the noise as expressed in terms of the spaces X^ρ (cf. Lemma 2.3.5), rather than the explicit representation of the noise.)

The rates we obtain are not as strong as in the simultaneously diagonalizable case when the true solution is too regular; in particular our rates saturate earlier as a function of increasing regularity, and we require a certain degree of regularity

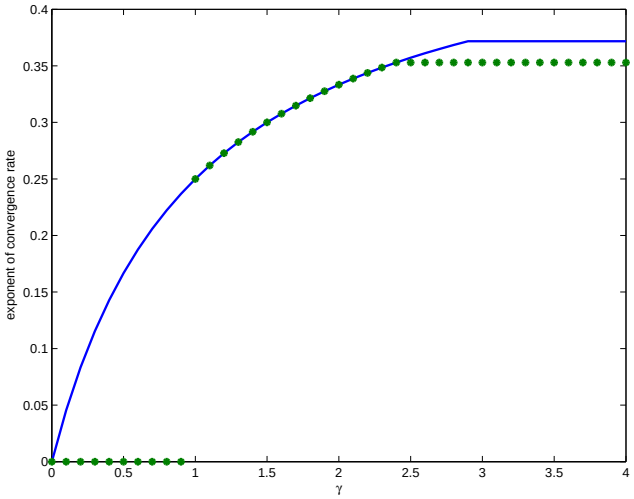

Figure 2.1: Exponents of rates of contraction plotted against the regularity of the true solution, γ. In blue are the sharp convergence rates obtained in the diagonal case in [44], while in green the rates predicted by our method, which applies to the more general non-diagonal case

of the true solution in order to secure convergence. It is not known if our results can be improved but it would be interesting to try. Both of the two discrepancies are attributed to the fact that our method relies on interpolating between rates in a strong and a weak norm of the error $e = m^\dagger - u^\dagger$; on the one hand the rate of the error in the weak norm saturates earlier, and on the other hand the error in the strong norm requires additional regularity in order to converge (cf. Section 2.9).

2.3 The Setting

In this section we present the setting in which we formulate our results. First, we define the spaces in which we work, in particular, we define the Hilbert scale induced by the prior covariance operator $\mathcal{C}_0$. Then we define the probability measures relevant to our analysis. Furthermore, we state our main assumptions, which concern the decay of the eigenvalues of $\mathcal{C}_0$ and the connections between the operators $\mathcal{C}_0$, $\mathcal{C}_1$ and $\mathcal{A}^{-1}$, and present regularity results for draws from the prior, μ_0, and the noise distribution, $\mathcal{N}(0, \mathcal{C}_1)$. Finally we briefly overview the way in which the Hilbert scale defined in terms of the prior covariance operator $\mathcal{C}_0$, which is natural for our analysis, links to scales of spaces defined independently of any prior model.

We start by defining the Hilbert scale which we will use in our analysis.

Recall that $\mathcal{X}$ is an infinite-dimensional separable Hilbert space and $\mathcal{C}_0 \colon \mathcal{X} \to \mathcal{X}$ is a self-adjoint, positive-definite, trace class, linear operator. Since $\mathcal{C}_0 \colon \mathcal{X} \to \mathcal{X}$ is injective and self-adjoint we have that $\mathcal{X} = \overline{\mathcal{R}(\mathcal{C}_0)} \oplus \mathcal{R}(\mathcal{C}_0)^\perp = \overline{\mathcal{R}(\mathcal{C}_0)}$. This means that $\mathcal{C}_0^{-1} \colon \mathcal{R}(\mathcal{C}_0) \to \mathcal{X}$ is a densely defined, unbounded, symmetric, positive-definite, linear operator in $\mathcal{X}$. Hence it can be extended to a self-adjoint operator with domain $\mathcal{D}(\mathcal{C}_0^{-1}) := \{u \in \mathcal{X} : \mathcal{C}_0^{-1} u \in \mathcal{X}\}$; this is the Friedrichs extension [50]. Thus, we can define the Hilbert scale $(X^t)_{t \in \mathbb{R}}$, with $X^t := \overline{\mathcal{M}}^{\|\cdot\|_t}$ [22], where

$$\mathcal{M} := \bigcap_{l=0}^{\infty} \mathcal{D}(\mathcal{C}_0^{-l}), \ \left\langle u, v \right\rangle_t := \left\langle \mathcal{C}_0^{-\frac{t}{2}} u, \mathcal{C}_0^{-\frac{t}{2}} v \right\rangle \quad \text{and} \quad \|u\|_t := \left\| \mathcal{C}_0^{-\frac{t}{2}} u \right\|.$$

The bounded linear operator $\mathcal{C}_1 \colon \mathcal{X} \to \mathcal{X}$ is assumed to be self-adjoint, positive-definite (but not necessarily trace class); thus $\mathcal{C}_1^{-1} \colon \mathcal{R}(\mathcal{C}_1) \to \mathcal{X}$ can be extended in the same way to a self-adjoint operator with domain $\mathcal{D}(\mathcal{C}_1^{-1}) := \{u \in \mathcal{X} : \mathcal{C}_1^{-1} u \in \mathcal{X}\}$. Finally, recall that we assume that $\mathcal{A} \colon \mathcal{D}(\mathcal{A}) \to \mathcal{X}$ is a self-adjoint and positive-definite, linear operator with bounded inverse, $\mathcal{A}^{-1} \colon \mathcal{X} \to \mathcal{X}$.

We assume that we have a probability space $(\Omega, \mathcal{F}, \mathcal{P})$. The expected value is denoted by $\mathbb{E}$ and $\eta \sim \mu$ means that the law of the random variable η is the measure μ.

Let $\mu_0 := \mathcal{N}(0, \tau^2 \mathcal{C}_0)$ and $\mathbb{P}_0 := \mathcal{N}(0, \frac{1}{n} \mathcal{C}_1)$ be the prior and noise distributions respectively. Furthermore, let $\nu(du, dy)$ denote the joint distribution of $u \sim \mu_0$ and $y|u \sim \mathcal{N}(\mathcal{A}^{-1} u, \frac{1}{n} \mathcal{C}_1)$:

$$\nu(du, dy) = \mathbb{P}(dy|u) \mu_0(du),$$

where $\mathbb{P} := \mathcal{N}(\mathcal{A}^{-1} u, \frac{1}{n} \mathcal{C}_1)$. We denote by $\nu_0(du, dy)$ the measure constructed by taking u and y as independent Gaussian random variables $\mathcal{N}(0, \tau^2 \mathcal{C}_0)$ and $\mathcal{N}(0, \frac{1}{n} \mathcal{C}_1)$ respectively:

$$\nu_0(du, dy) = \mathbb{P}_0(dy) \otimes \mu_0(du).$$

Let $\{\lambda_j, \phi_j\}_{j=1}^{\infty}$ be orthonormal eigenpairs of $\mathcal{C}_0$ in $\mathcal{X}$. Thus, $\{\lambda_j\}_{j=1}^{\infty}$ are the eigenvalues, which are positive since $\mathcal{C}_0$ is positive definite, and $\{\phi_j\}_{j=1}^{\infty}$ an orthonormal eigenbasis. Since $\mathcal{C}_0$ is trace class we have that $\sum_{j=1}^{\infty} \lambda_j < \infty$. In fact we require a slightly stronger assumption see Assumption 2.3.1(1) below.

2.3.1 Assumptions

We are now ready to present our assumptions. The first assumption enables us to quantify the regularity of draws from the prior whereas the rest of the assumptions regard interrelations between the three operators $\mathcal{C}_0$, $\mathcal{C}_1$ and $\mathcal{A}^{-1}$; these assumptions

reflect the idea that

$$\mathcal{C}_1 \simeq \mathcal{C}_0^{\beta} \quad \text{and} \quad \mathcal{A}^{-1} \simeq \mathcal{C}_0^{\ell},$$

for some $\beta \geq 0, \ell \geq 0$, where $\simeq$ is used in the same manner as in Section 2.1. This is made precise by the inequalities presented in the following assumption, where the notation $a \asymp b$ means that there exist constants $c, c' > 0$ such that $ca \leq b \leq c'a$.

Assumption 2.3.1. *Suppose there exist $s_0 \in [0,1)$, $\beta \geq 0$, $\ell \geq 0$ and constants $c_i > 0, i = 1,..,4$ such that*

1. *$\mathcal{C}_0^s$ is trace class for all $s > s_0$;*
2. *$\Delta > 2s_0$, where $\Delta := 2\ell - \beta + 1$;*
3. *$\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\| \asymp \left\|\mathcal{C}_0^{\ell-\frac{\beta}{2}}u\right\|, \quad \forall u \in X^{\beta-2\ell}$;*
4. *$\left\|\mathcal{C}_0^{-\frac{\rho}{2}}\mathcal{C}_1^{\frac{1}{2}}u\right\| \leq c_1\left\|\mathcal{C}_0^{\frac{\beta-\rho}{2}}u\right\|, \quad \forall u \in X^{\rho-\beta}, \forall \rho \in [\lceil \beta - s_0 - 1 \rceil, \beta - s_0)$;*
5. *$\left\|\mathcal{C}_0^{\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}u\right\| \leq c_2\left\|\mathcal{C}_0^{\frac{s-\beta}{2}}u\right\|, \quad \forall u \in X^{\beta-s}, \forall s \in (s_0, 1]$;*
6. *$\left\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\| \leq c_3\left\|\mathcal{C}_0^{\frac{2\ell-\beta-s}{2}}u\right\|, \quad \forall u \in X^{s+\beta-2\ell}, \forall s \in (s_0, 1]$;*
7. *$\left\|\mathcal{C}_0^{\frac{\kappa}{2}}\mathcal{A}^{-1}\mathcal{C}_1^{-1}u\right\| \leq c_4\left\|\mathcal{C}_0^{\frac{\kappa}{2}+\ell-\beta}u\right\|, \quad \forall u \in X^{2\beta-2\ell-\kappa}, \forall \kappa \in [\beta - 2\ell, 1]$.*

Notice that, by Assumption 2.3.1(2) we have $2\ell - \beta > -1$ which, in combination with Assumption 2.3.1(3), implies that

$$\left\langle \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u, \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\rangle + \lambda\left\langle \mathcal{C}_0^{-\frac{1}{2}}u, \mathcal{C}_0^{-\frac{1}{2}}u\right\rangle \leq c\left\langle \mathcal{C}_0^{-\frac{1}{2}}u, \mathcal{C}_0^{-\frac{1}{2}}u\right\rangle, \quad \forall u \in X^1,$$

capturing the idea that the regularization through $\mathcal{C}_0$ is indeed a regularization. In fact the assumption $\Delta > 2s_0$ connects the ill-posedness of the problem to the regularity of the prior. We exhibit this connection in the following example:

Example 2.3.2. *Assume $\mathcal{A}, \mathcal{C}_1$ and $\mathcal{C}_0$ are simultaneously diagonalizable, with eigenvalues having algebraic decay $j^{2\hat{\ell}}$, $j^{-2\hat{\beta}}$ and $j^{-2\alpha}$, respectively, for $\hat{\ell}, \hat{\beta} \geq 0$ and $\alpha > \frac{1}{2}$ so that $\mathcal{C}_0$ is trace class. Then Assumptions (1),(3)-(7) are trivially satisfied with $\ell = \frac{\hat{\ell}}{\alpha}$, $\beta = \frac{\hat{\beta}}{\alpha}$ and $s_0 = \frac{1}{2\alpha}$. The Assumption (2) $\Delta > 2s_0$ is then equivalent to $\alpha > 1 + \hat{\beta} - 2\hat{\ell}$. That is, for a certain degree of ill-posedness (encoded in the difference $2\hat{\ell} - \hat{\beta}$) we have a minimum requirement on the regularity of the prior (encoded in α). Put differently, for a certain prior, we require a minimum degree of ill-posedness.*

We refer the reader to Section 2.8 for nontrivial examples satisfying Assumptions 2.3.1.

In the following, we exploit the regularity properties of a white noise to determine the regularity of draws from the prior and the noise distributions using Assumption 2.3.1(1). We consider a white noise to be a draw from $\mathcal{N}(0, I)$, that is a random variable $\zeta \sim \mathcal{N}(0, I)$. Even though the identity operator is not trace class in $\mathcal{X}$, it is trace class in a bigger space X^{-s}, where $s > 0$ is sufficiently large.

Lemma 2.3.3. *Under the Assumption 2.3.1(1) we have:*

i) Let ζ be a white noise. Then $\mathbb{E}\left\|\mathcal{C}_0^{\frac{s}{2}}\zeta\right\|^2 < \infty$ for all $s > s_0$.

ii) Let $u \sim \mu_0$. Then $u \in X^{1-s}$ μ_0-a.s. for every $s > s_0$.

Proof.

i) We have that $\mathcal{C}_0^{\frac{s}{2}}\zeta \sim \mathcal{N}(0, \mathcal{C}_0^s)$, thus $\mathbb{E}\left\|\mathcal{C}_0^{\frac{s}{2}}\zeta\right\|^2 < \infty$ is equivalent to $\mathcal{C}_0^s$ being of trace class. By the Assumption 2.3.1(1) it suffices to have $s > s_0$.

ii) We have $\mathbb{E}\left\|\mathcal{C}_0^{\frac{s-1}{2}}u\right\|^2 = \mathbb{E}\left\|\mathcal{C}_0^{\frac{s}{2}}\mathcal{C}_0^{-\frac{1}{2}}u\right\|^2 = \mathbb{E}\left\|\mathcal{C}_0^{\frac{s}{2}}\zeta\right\|^2$, where ζ is a white noise, therefore using part (i) we get the result.

□

Remark 2.3.4. *Note that as s_0 changes, both the Hilbert scale and the decay of the coefficients of a draw from μ_0 change. The norms $\|\cdot\|_t$ are defined through powers of the eigenvalues λ_j. If $s_0 > 0$, then $\mathcal{C}_0$ has eigenvalues that decay like $j^{-\frac{1}{s_0}}$, thus an element $u \in X^t$ has coefficients $\langle u, \phi_j \rangle$, that decay faster than $j^{-\frac{1}{2}-\frac{t}{2s_0}}$. As s_0 gets closer to zero, the space X^t for fixed $t > 0$, corresponds to a faster decay rate of the coefficients. At the same time, by the last lemma, draws from $\mu_0 = \mathcal{N}(0, \mathcal{C}_0)$ belong to X^{1-s} for all $s > s_0$. Consequently, as s_0 gets smaller, not only do draws from μ_0 belong to X^{1-s} for smaller s, but also the spaces X^{1-s} for fixed s reflect faster decay rates of the coefficients. The case $s_0 = 0$ corresponds to $\mathcal{C}_0$ having eigenvalues that decay faster than any negative power of j. A draw from μ_0 in that case has coefficients that decay faster than any negative power of j (we consider such priors in Appendix A).*

In the next lemma, we use the interrelations between the operators $\mathcal{C}_0, \mathcal{C}_1, \mathcal{A}^{-1}$ to obtain additional regularity properties of draws from the prior, and also determine the regularity of draws from the noise distribution and the joint distribution of the unknown and the data.

Lemma 2.3.5. *Under the Assumptions 2.3.1 we have:*

i) $u \in X^{s_0+\beta-2\ell+\varepsilon}$ μ_0*-a.s. for all* $0 < \varepsilon < (\Delta - 2s_0) \wedge (1 - s_0)$;

ii) $\mathcal{A}^{-1}u \in \mathcal{D}(\mathcal{C}_1^{-\frac{1}{2}})$ μ_0*-a.s.*;

iii) $\eta \in X^\rho$ $\mathbb{P}_0$*-a.s. for all* $\rho < \beta - s_0$;

iv) $y \in X^\rho$ ν*-a.s. for all* $\rho < \beta - s_0$.

Proof.

i) We can choose an ε as in the statement by the Assumption 2.3.1(2). By Lemma 2.3.3(ii), it suffices to show that $s_0 + \beta - 2\ell + \varepsilon < 1 - s_0$. Indeed, $s_0 + \beta - 2\ell + \varepsilon = s_0 + 1 - \Delta + \varepsilon < 1 - s_0$.

ii) Under Assumption 2.3.1(3) it suffices to show that $u \in X^{\beta-2\ell}$. Indeed, by Lemma 2.3.3(ii), we need to show that $\beta - 2\ell < 1 - s_0$, which is true since $s_0 \in [0, 1)$ and we assume $\Delta > 2s_0 \geq s_0$, thus $2\ell - \beta + 1 > s_0$.

iii) It suffices to show it for any $\rho \in [\lceil \beta - s_0 - 1 \rceil, \beta - s_0)$. Noting that $\zeta = \mathcal{C}_1^{-\frac{1}{2}}\eta$ is a white noise, using Assumption 2.3.1(4), we have by Lemma 2.3.3(i)

$$\mathbb{E}\|\eta\|_\rho^2 = \mathbb{E}\big\|\mathcal{C}_0^{-\frac{\rho}{2}}\mathcal{C}_1^{\frac{1}{2}}\mathcal{C}_1^{-\frac{1}{2}}\eta\big\|^2 \leq c\mathbb{E}\big\|\mathcal{C}_0^{\frac{\beta-\rho}{2}}\zeta\big\|^2 < \infty,$$

since $\beta - \rho > s_0$.

iv) By (ii) we have that $\mathcal{A}^{-1}u$ is μ_0-a.s. in the Cameron-Martin space of the Gaussian measures $\mathbb{P}$ and $\mathbb{P}_0$, thus the measures $\mathbb{P}$ and $\mathbb{P}_0$ are μ_0-a.s. equivalent [17, Theorem 2.8] and (iii) gives the result.

□

2.3.2 Guidelines for applying the theory

The theory is naturally developed in the scale of Hilbert spaces defined via the prior. However application of the theory may be more natural in a different functional setting. We explain how the two may be connected. Let $\{\psi_j\}_{j\in\mathbb{N}}$ be an orthonormal basis of the separable Hilbert space $\mathcal{X}$. We define the spaces $\mathcal{H}^t$, $t \in \mathbb{R}$ as follows: for $t > 0$ we set

$$\mathcal{H}^t := \{u \in \mathcal{X} : \sum_{j=1}^{\infty} j^{2t}\langle u, \psi_j\rangle^2 < \infty\}$$

and the spaces $\mathcal{H}^{-t}$, $t > 0$ are defined by duality, $\mathcal{H}^{-t} := (\mathcal{H}^t)^*$.

For example, if we restrict ourselves to functions on a periodic domain $D = [0, L]^d$ and assume that $\{\psi_j\}_{j\in\mathbb{N}}$ is the Fourier basis of $\mathcal{X} = L^2(D)$, then the spaces $\mathcal{H}^t$ can be identified with the Sobolev spaces of periodic functions H^t, by rescaling: $H^t = \mathcal{H}^{\frac{t}{d}}$ [73, Proposition 5.39].

In the case $s_0 > 0$, as explained in Remark 2.3.4 we have algebraic decay of the eigenvalues of $\mathcal{C}_0$ and in particular λ_j decay like $j^{-\frac{1}{s_0}}$. If $\mathcal{C}_0$ is diagonalizable in the basis $\{\psi_j\}_{j\in\mathbb{N}}$, that is, if $\phi_j = \psi_j,\ j \in \mathbb{N}$, then it is straightforward to identify the spaces X^t with the spaces $\mathcal{H}^{\frac{t}{2s_0}}$. The advantage of this identification is that the spaces $\mathcal{H}^t$ do not depend on the prior so one can use them as a fixed reference point for expressing regularity, for example of the true solution.

In our subsequent analysis, we will require that the true solution lives in the Cameron-Martin space of the prior X^1, which in different choices of the prior (different s_0) is a different space. Furthermore, we will assume that the true solution lives in X^γ for some $\gamma \geq 1$ and provide the convergence rate depending on the parameters γ, s_0, β, ℓ. The identification $X^\gamma = \mathcal{H}^{\frac{\gamma}{2s_0}}$ and the intuitive relation between the spaces $\mathcal{H}^t$ and the Sobolev spaces, enable us to understand the meaning of the assumptions on the true solution.

We can now formulate the following guidelines for applying the theory presented in the present chapter: we work in a separable Hilbert space $\mathcal{X}$ with an orthonormal basis $\{\psi_j\}_{j\in\mathbb{N}}$ and we have some prior knowledge about the true solution $u^\dagger$ which can be expressed in terms of the spaces $\mathcal{H}^t$. The noise is assumed to be Gaussian $\mathcal{N}(0, \mathcal{C}_1)$, and the forward operator is known; that is, $\mathcal{C}_1$ and $\mathcal{A}^{-1}$ are known. We choose the prior $\mathcal{N}(0, \mathcal{C}_0)$, that is, we choose the covariance operator $\mathcal{C}_0$, and we can determine the value of s_0. If the operator $\mathcal{C}_0$ is chosen to be diagonal in the basis $\{\psi_j\}_{j\in\mathbb{N}}$ then we can find the regularity of the true solution in terms of the spaces X^t, that is, the value of γ such that $u^\dagger \in X^\gamma$, and check that $\gamma \geq 1$ which is necessary for our theory to work. We then find the values of β and ℓ and calculate the value of Δ appearing in Assumption 2.3.1, checking that our choice of the prior is such that $\Delta > 2s_0$. We now have all the necessary information required for applying the Theorems 2.2.2 and 2.2.3 presented in Section 2.2 to get the rate of convergence.

Remark 2.3.6. *Observe that in the above mentioned example of periodic functions, we have the identification $X^1 = H^{\frac{d}{2s_0}}$, thus since $s_0 < 1$ we have that the assumption $u^\dagger \in X^1$ implies that $u^\dagger \in H^t$, for $t > \frac{d}{2}$. By the Sobolev embedding theorem [73, Theorem 5.31], this implies that the true solution is always assumed to be continuous. However, this is not a disadvantage of our method, since in many cases a Gaussian measure which charges $L^2(D)$ with probability one, can be shown to also charge the*

space of continuous functions with probability one [78, Lemma 6.25]

2.4 Properties of the Posterior Mean and Covariance

We now make sense of equation (2.1.7) weakly in the space X^1, under the assumptions presented in the previous section. To do so, we define the operator $\mathcal{B}_\lambda$ from (2.1.6) in X^1 and examine its properties. In Section 2.5 we demonstrate that (2.1.4) and (2.1.7) do indeed correspond to the posterior covariance and mean.

Consider the equation

$$\mathcal{B}_\lambda w = r, \tag{2.4.1}$$

where

$$\mathcal{B}_\lambda = \mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1} + \lambda\mathcal{C}_0^{-1}.$$

Define the bilinear form $B\colon X^1 \times X^1 \to \mathbb{R}$,

$$B(u,v) := \langle \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u, \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}v\rangle + \lambda\langle \mathcal{C}_0^{-\frac{1}{2}}u, \mathcal{C}_0^{-\frac{1}{2}}v\rangle, \ \ \forall u,v \in X^1.$$

Definition 2.4.1. *Let $r \in X^{-1}$. An element $w \in X^1$ is called a weak solution of (2.4.1), if*

$$B(w,v) = \langle r, v\rangle, \ \forall v \in X^1.$$

Proposition 2.4.2. *Under the Assumptions 2.3.1(2) and (3), for any $r \in X^{-1}$, there exists a unique weak solution $w \in X^1$ of (2.4.1).*

Proof. We use the Lax-Milgram theorem in the Hilbert space X^1, since $r \in X^{-1} = (X^1)^*$.

i) $B\colon X^1 \times X^1 \to \mathbb{R}$ is coercive:

$$B(u,u) = \left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\|^2 + \lambda\left\|\mathcal{C}_0^{-\frac{1}{2}}u\right\|^2 \geq \lambda\|u\|_1^2, \ \forall u \in X^1.$$

ii) $B\colon X^1 \times X^1 \to \mathbb{R}$ is continuous: indeed by the Cauchy-Schwarz inequality and the Assumptions 2.3.1(2) and (3),

$$\begin{aligned}|B(u,v)| &\leq \left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\|\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}v\right\| + \lambda\left\|\mathcal{C}_0^{-\frac{1}{2}}u\right\|\left\|\mathcal{C}_0^{-\frac{1}{2}}v\right\| \\ &\leq c\|u\|_{\beta-2\ell}\|v\|_{\beta-2\ell} + \lambda\|u\|_1\|v\|_1 \leq c'\|u\|_1\|v\|_1, \ \forall u,v \in X^1.\end{aligned}$$

□

Remark 2.4.3. *The Lax-Milgram theorem defines a bounded operator $\mathcal{S} : X^{-1} \to X^1$, such that $B(\mathcal{S}r, v) = \langle r, v \rangle$ for all $v \in X^1$, which has a bounded inverse $\mathcal{S}^{-1} : X^1 \to X^{-1}$ such that $B(w, v) = \langle \mathcal{S}^{-1}w, v \rangle$ for all $v \in X^1$, [87, Section III.7]. Henceforward, we identify $\mathcal{B}_\lambda \equiv \mathcal{S}^{-1}$ and $\mathcal{B}_\lambda^{-1} \equiv \mathcal{S}$. Furthermore, note that in Proposition 2.4.2, Lemma 2.4.4 below, and the two propositions in Section 2.6, we only require $\Delta > 0$ and not the stronger assumption $\Delta > 2s_0$. However, in all our other results we actually need $\Delta > 2s_0$.*

Lemma 2.4.4. *Suppose the Assumptions 2.3.1(2) and (3) hold. Then the operator $\mathcal{S}^{-1} = \mathcal{B}_\lambda : X^1 \to X^{-1}$ is identical to the operator $\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1} + \lambda\mathcal{C}_0^{-1} : X^1 \to X^{-1}$, where $\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}$ is defined weakly in $X^{\beta-2\ell}$.*

Proof. The Lax-Milgram theorem implies that $\mathcal{B}_\lambda : X^1 \to X^{-1}$ is bounded. Moreover, $\mathcal{C}_0^{-1} : X^1 \to X^{-1}$ is bounded, thus the operator $\mathcal{K} := \mathcal{B}_\lambda - \lambda\mathcal{C}_0^{-1} : X^1 \to X^{-1}$ is also bounded and satisfies

$$\langle \mathcal{K}u, v \rangle = \langle \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u, \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}v \rangle, \ \forall u, v \in X^1. \tag{2.4.2}$$

Define $\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}$ weakly in $X^{\beta-2\ell}$, by the bilinear form $A : X^{\beta-2\ell} \times X^{\beta-2\ell} \to \mathbb{R}$ given by

$$A(u, v) = \langle \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u, \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}v \rangle, \ \forall u, v \in X^{\beta-2\ell}.$$

By Assumption 2.3.1(3), A is coercive and continuous in $X^{\beta-2\ell}$, thus by the Lax-Milgram theorem, there exists a uniquely defined, boundedly invertible, operator $T : X^{2\ell-\beta} \to X^{\beta-2\ell}$ such that $A(u, v) = \langle T^{-1}u, v \rangle$ for all $v \in X^{\beta-2\ell}$. We identify $\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}$ with the bounded operator $T^{-1} : X^{\beta-2\ell} \to X^{2\ell-\beta}$. By Assumption 2.3.1(2) we have $\Delta > 0$ hence

$$\|\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}u\|_{-1} \leq c\|\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}u\|_{2\ell-\beta} \leq c\|u\|_{\beta-2\ell} \leq c\|u\|_1, \ \forall u \in X^1,$$

that is, $\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1} : X^1 \to X^{-1}$ is bounded. By the definition of $T^{-1} = \mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}$ and (2.4.2), this implies that $\mathcal{K} = \mathcal{B}_\lambda - \lambda\mathcal{C}_0^{-1} = \mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1}$. □

Proposition 2.4.5. *Under the Assumptions 2.3.1(1),(2),(3),(4),(7), there exists a unique weak solution, $m \in X^1$ of equation (2.1.7), $\nu(du, dy)$-almost surely.*

Proof. It suffices to show that $\mathcal{A}^{-1}\mathcal{C}_1^{-1}y \in X^{-1}$, $\nu(du, dy)$-almost surely. Indeed, by Lemma 2.3.5(iv) we have that $y \in X^\rho$ $\nu(du, dy)$-a.s. for all $\rho < \beta - s_0$, thus by the Assumption 2.3.1(7)

$$\|\mathcal{C}_0^{\frac{1}{2}}\mathcal{A}^{-1}\mathcal{C}_1^{-1}y\| \leq c\|\mathcal{C}_0^{\frac{1}{2}+\ell-\beta}y\| < \infty,$$

since $2\beta - 2\ell - 1 < \beta - s_0$, which holds by the Assumption 2.3.1(2). □

2.5 Characterization of the Posterior using Precision Operators

Suppose that in the problem (2.1.1) we have $u \sim \mu_0 = \mathcal{N}(0, \mathcal{C}_0)$ and $\eta \sim \mathcal{N}(0, \mathcal{C}_1)$, where u is independent of η. Then we have that $y|u \sim \mathbb{P} = \mathcal{N}(\mathcal{A}^{-1}u, \frac{1}{n}\mathcal{C}_1)$. Let μ^y be the posterior measure on $u|y$.

In this section we prove a number of facts concerning the posterior measure μ^y for $u|y$. First, in Theorem 2.5.1 we prove that this measure has density with respect to the prior measure μ_0, identify this density and show that μ^y is Lipschitz in y, with respect to the Hellinger metric. Continuity in y will require the introduction of the space $X^{s+\beta-2\ell}$, to which u drawn from μ_0 belongs almost surely. Secondly, we prove Theorem 2.2.1, where we show that μ^y is Gaussian and identify the covariance and mean via equations (2.1.4) and (2.1.7). This identification will form the basis for our analysis of posterior contraction in the following section.

Theorem 2.5.1. *Under the Assumptions 2.3.1(1),(2),(3),(4),(5),(6), the posterior measure μ^y is absolutely continuous with respect to μ_0 and*

$$\frac{d\mu^y}{d\mu_0}(u) = \frac{1}{Z(y)} \exp(-\Phi(u, y)), \tag{2.5.1}$$

where

$$\Phi(u, y) := \frac{n}{2}\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\|^2 - n\left\langle \mathcal{C}_1^{-\frac{1}{2}}y, \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\rangle \tag{2.5.2}$$

and $Z(y) \in (0, \infty)$ is the normalizing constant. Furthermore, the map $y \mapsto \mu^y$ is Lipschitz continuous, with respect to the Hellinger metric: let $s = s_0 + \varepsilon$, $0 < \varepsilon < (\Delta - 2s_0) \wedge (1 - s_0)$; then there exists $c = c(r)$ such that for all $y, y' \in X^{\beta-s}$ with $\|y\|_{\beta-s}, \|y'\|_{\beta-s} \leq r$

$$d_{\text{Hell}}(\mu^y, \mu^{y'}) \leq c\|y - y'\|_{\beta-s}.$$

Consequently, the μ^y-expectation of any polynomially bounded function $f\colon X^{s+\beta-2\ell} \to E$, where $(E, \|\cdot\|_E)$ is a Banach space, is locally Lipschitz continuous in y. In particular, the posterior mean is locally Lipschitz continuous in y as a function $X^{\beta-s} \to X^{s+\beta-2\ell}$.

The proofs of Theorem 2.5.1 and Theorem 2.2.1 are presented in the next two subsections. Each proof is based on a series of lemmas.

2.5.1 Proof of Theorem 2.5.1

In this subsection we prove Theorem 2.5.1. We first prove several useful estimates regarding Φ defined in (2.5.2), for $u \in X^{s+\beta-2\ell}$ and $y \in X^{\beta-s}$, where $s \in (s_0, 1]$. Observe that, under the Assumptions 2.3.1(1),(2),(3),(4), for $s = s_0 + \varepsilon$ where $\varepsilon > 0$ sufficiently small, the Lemma 2.3.5 implies on the one hand that $u \in X^{s+\beta-2\ell}$ $\mu_0(du)$-almost surely and on the other hand that $y \in X^{\beta-s}$ $\nu(du, dy)$-almost surely.

Lemma 2.5.2. *Under the Assumptions 2.3.1(1),(3),(5),(6), for any $s \in (s_0, 1]$, the potential Φ given by (2.5.2) satisfies:*

i) for every $\delta > 0$ and $r > 0$, there exists an $M = M(\delta, r) \in \mathbb{R}$, such that for all $u \in X^{s+\beta-2\ell}$ and all $y \in X^{\beta-s}$ with $\|y\|_{\beta-s} \leq r$,

$$\Phi(u, y) \geq M - \delta \|u\|^2_{s+\beta-2\ell};$$

ii) for every $r > 0$, there exists an $R = R(r) > 0$, such that for all $u \in X^{s+\beta-2\ell}$ and $y \in X^{\beta-s}$ with $\|u\|_{s+\beta-2\ell}, \|y\|_{\beta-s} \leq r$,

$$\Phi(u, y) \leq R;$$

iii) for every $r > 0$, there exists an $L = L(r) > 0$, such that for all $u_1, u_2 \in X^{s+\beta-2\ell}$ and $y \in X^{\beta-s}$ with $\|u_1\|_{s+\beta-2\ell}, \|u_2\|_{s+\beta-2\ell}, \|y\|_{\beta-s} \leq r$,

$$|\Phi(u_1, y) - \Phi(u_2, y)| \leq L \|u_1 - u_2\|_{s+\beta-2\ell};$$

iv) for every $\delta > 0$ and $r > 0$, there exists a $c = c(\delta, r) \in \mathbb{R}$, such that for all $y_1, y_2 \in X^{\beta-s}$ with $\|y_1\|_{\beta-s}, \|y_2\|_{\beta-s} \leq r$ and for all $u \in X^{s+\beta-2\ell}$,

$$|\Phi(u, y_1) - \Phi(u, y_2)| \leq \exp\left(\delta \|u\|^2_{s+\beta-2\ell} + c\right) \|y_1 - y_2\|_{\beta-s}.$$

Proof.

i) By first using the Cauchy-Schwarz inequality, then the Assumptions 2.3.1 (5) and (6), and then the Cauchy with δ' inequality for $\delta' > 0$ sufficiently small,

we have

$$\begin{aligned}\Phi(u,y) &= \frac{n}{2}\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\|^2 - n\left\langle \mathcal{C}_0^{\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}y, \mathcal{C}_0^{-\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\rangle \\ &\geq -n\left\|\mathcal{C}_0^{\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}y\right\|\left\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\| \geq -cn\|y\|_{\beta-s}\|u\|_{s+\beta-2\ell} \\ &\geq -\frac{cn}{4\delta'}\|y\|_{\beta-s}^2 - cn\delta'\|u\|_{s+\beta-2\ell}^2 \geq M(r,\delta) - \delta\|u\|_{s+\beta-2\ell}^2.\end{aligned}$$

ii) By the Cauchy-Schwarz inequality and the Assumptions 2.3.1(3),(5) and (6), we have since $s > s_0 \geq 0$

$$\begin{aligned}\Phi(u,y) &\leq \frac{n}{2}\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\|^2 + n\left\|\mathcal{C}_0^{\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}y\right\|\left\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\| \\ &\leq c\frac{n}{2}\|u\|_{\beta-2\ell}^2 + cn\|y\|_{\beta-s}\|u\|_{s+\beta-2\ell} \leq R(r).\end{aligned}$$

iii) By first using the Assumptions 2.3.1 (5) and (6) and the triangle inequality, and then the Assumption 2.3.1(3) and the reverse triangle inequality, we have since $s > s_0 \geq 0$

$$\begin{aligned}&|\Phi(u_1,y) - \Phi(u_2,y)| = \\ &\frac{n}{2}\left|\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u_1\right\|^2 - \left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u_2\right\|^2 + 2\left\langle \mathcal{C}_0^{\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}y, \mathcal{C}_0^{-\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}(u_2-u_1)\right\rangle\right| \\ &\leq \frac{n}{2}\left|\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u_1\right\|^2 - \left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u_2\right\|^2\right| + cn\|y\|_{\beta-s}\|u_1-u_2\|_{s+\beta-2\ell} \\ &\leq cn\|u_1-u_2\|_{\beta-2\ell}\left(\|u_1\|_{\beta-2\ell} + \|u_2\|_{\beta-2\ell}\right) + cnr\|u_1-u_2\|_{s+\beta-2\ell} \\ &\leq L(r)\|u_1-u_2\|_{s+\beta-2\ell}.\end{aligned}$$

iv) By first using the Cauchy-Schwarz inequality and then the Assumptions 2.3.1(5) and (6), we have

$$\begin{aligned}|\Phi(u,y_1) - \Phi(u,y_2)| &= n\left|\left\langle \mathcal{C}_0^{\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}(y_1-y_2), \mathcal{C}_0^{-\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\rangle\right| \\ &\leq n\left\|\mathcal{C}_0^{\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}(y_1-y_2)\right\|\left\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}u\right\| \\ &\leq cn\|y_1-y_2\|_{\beta-s}\|u\|_{s+\beta-2\ell} \\ &\leq \exp\left(\delta\|u\|_{s+\beta-2\ell}^2 + c\right)\|y_1-y_2\|_{\beta-s}.\end{aligned}$$

□

Corollary 2.5.3. *Under the Assumptions 2.3.1(1),(2),(3),(5),(6)*

$$Z(y) := \int_{\mathcal{X}} \exp(-\Phi(u,y))\mu_0(du) > 0,$$

for all $y \in X^{\beta-s}, s = s_0 + \varepsilon$ *where* $0 < \varepsilon < (\Delta - 2s_0) \wedge (1 - s_0)$. *In particular, if in addition the Assumption 2.3.1(4) holds, then* $Z(y) > 0$ ν*-almost surely.*

Proof. Fix $y \in X^{\beta-s}$ and set $r = \|y\|_{\beta-s}$. Gaussian measures on separable Hilbert spaces are full [17, Proposition 1.25], hence since by Lemma 2.3.5(i) $\mu_0(X^{s+\beta-2\ell}) = 1$, we have that $\mu_0(B_{X^{s+\beta-2\ell}}(r)) > 0$. By Lemma 2.5.2(ii), there exists $R(r) > 0$ such that

$$\begin{aligned}\int_{\mathcal{X}} \exp(-\Phi(u,y))\mu_0(du) &\geq \int_{B_{X^{s+\beta-2\ell}}(r)} \exp(-\Phi(u,y))\mu_0(du) \\ &\geq \int_{B_{X^{s+\beta-2\ell}}(r)} \exp(-R(r))\mu_0(du) > 0.\end{aligned}$$

Recalling that, under the additional Assumption 2.3.1(4), by Lemma 2.3.5(iv) we have $y \in X^{\beta-s}$ ν-almost surely for all $s > s_0$, completes the proof. □

We are now ready to prove Theorem 2.5.1:

Proof of Theorem 2.5.1. Recall that $\nu_0 = \mathbb{P}_0(dy) \otimes \mu_0(du)$ and $\nu = \mathbb{P}(dy|u)\mu_0(du)$. By the Cameron-Martin formula [13, Corollary 2.4.3], since by Lemma 2.3.5(ii) we have $\mathcal{A}^{-1}u \in \mathcal{D}(\mathcal{C}_1^{-\frac{1}{2}})$ μ_0-a.s., we get for μ_0-almost all u

$$\frac{d\mathbb{P}}{d\mathbb{P}_0}(y|u) = \exp(-\Phi(u,y)),$$

thus we have for μ_0-almost all u

$$\frac{d\nu}{d\nu_0}(y,u) = \exp(-\Phi(u,y)).$$

By [32, Lemma 5.3] and Corollary 2.5.3 we have the relation (2.5.1).
For the proof of the Lipschitz continuity of the posterior measure in y, with respect to the Hellinger distance, we apply [78, Theorem 4.2] for $Y = X^{\beta-s}, X = X^{s+\beta-2\ell}$, using Lemma 2.5.2 and the fact that $\mu_0(X^{s+\beta-2\ell}) = 1$, by Lemma 2.3.5(i). □

2.5.2 Proof of Theorem 2.2.1

We first give an overview of the proof of Theorem 2.2.1. Let $y|u \sim \mathbb{P} = \mathcal{N}(\mathcal{A}^{-1}u, \frac{1}{n}\mathcal{C}_1)$ and $u \sim \mu_0$. Then by Proposition 2.4.5, there exists a unique weak solution, $m \in X^1$,

of (2.1.7), $\nu(du, dy)$-almost surely. That is, with $\nu(du, dy)$-probability equal to one, there exists an $m = m(y) \in X^1$ such that

$$B(m, v) = b^y(v), \quad \forall v \in X^1,$$

where the bilinear form B is defined in Section 2.4, and $b^y(v) = \left\langle \mathcal{A}^{-1}\mathcal{C}_1^{-1}y, v \right\rangle$. In the following we show that $\mu^y = \mathcal{N}(m, \mathcal{C})$, where

$$\mathcal{C}^{-1} = n\mathcal{A}^{-1}\mathcal{C}_1^{-1}\mathcal{A}^{-1} + \frac{1}{\tau^2}\mathcal{C}_0^{-1}.$$

The proof has the same structure as the proof for the identification of the posterior in [64]. We define the Gaussian measure $\mathcal{N}(m^N, \mathcal{C}^N)$, which is the independent product of a measure identical to $\mathcal{N}(m, \mathcal{C})$ in the finite-dimensional space $\mathcal{X}^N$ spanned by the first N eigenfunctions of $\mathcal{C}_0$, and a measure identical to μ_0 in $(\mathcal{X}^N)^\perp$. We next show that $\mathcal{N}(m^N, \mathcal{C}^N)$ converges weakly to the measure μ^y which as a weak limit of Gaussian measures has to be Gaussian $\mu^y = \mathcal{N}(\overline{m}, \overline{\mathcal{C}})$, and we then identify $\overline{m}$ and $\overline{\mathcal{C}}$ with m, $\mathcal{C}$ respectively.

Fix y drawn from ν and let P^N be the orthogonal projection of $\mathcal{X}$ to the finite-dimensional space $\text{span}\{\phi_1, ..., \phi_N\} := \mathcal{X}^N$, where as in Section 2.3, $\{\phi_j\}_{j=1}^\infty$ is an orthonormal eigenbasis of $\mathcal{C}_0$ in $\mathcal{X}$. Let $Q^N = I - P^N$. We define $\mu^{N,y}$ by

$$\frac{d\mu^{N,y}}{d\mu_0}(u) = \frac{1}{Z^N(y)} \exp(-\Phi^N(u, y)) \tag{2.5.3}$$

where $\Phi^N(u, y) := \Phi(P^N u, y)$ and

$$Z^N(y) := \int_{\mathcal{X}} \exp(-\Phi^N(u, y))\mu_0(du).$$

Lemma 2.5.4. *We have $\mu^{N,y} = \mathcal{N}(m^N, \mathcal{C}^N)$, where*

$$P^N\mathcal{C}^{-1}P^N m^N = nP^N\mathcal{A}^{-1}\mathcal{C}_1^{-1}y,$$

$$P^N\mathcal{C}^N P^N = P^N\mathcal{C}P^N, \; Q^N\mathcal{C}^N Q^N = \tau^2 Q^N\mathcal{C}_0 Q^N$$

and $P^N\mathcal{C}^N Q^N = Q^N\mathcal{C}^N P^N = 0$.

Proof. Let $u \in \mathcal{X}^N$. Since $u = P^N u$ we have by (2.5.3)

$$d\mu^{N,y}(P^N u) \propto \exp\left(-\Phi(P^N u; y)\right) d\mu_0(P^N u).$$

The right hand side is N-dimensional Gaussian with density proportional to the

exponential of the following expression

$$-\frac{n}{2}\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}P^N u\right\|^2 + n\langle \mathcal{C}_1^{-\frac{1}{2}} y, \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}P^N u\rangle - \frac{1}{2\tau^2}\left\|\mathcal{C}_0^{-\frac{1}{2}}P^N u\right\|^2, \tag{2.5.4}$$

which by completing the square we can write as

$$-\frac{1}{2}\left\|(\tilde{\mathcal{C}}^N)^{-\frac{1}{2}}(u-\tilde{m}^N)\right\|^2 + c(y),$$

where $\tilde{\mathcal{C}}^N$ is the covariance matrix and $\tilde{m}^N$ the mean. By equating with expression (2.5.4), we find that $(\tilde{\mathcal{C}}^N)^{-1} = P^N\mathcal{C}^{-1}P^N$ and $(\tilde{\mathcal{C}}^N)^{-1}\tilde{m}^N = nP^N\mathcal{A}^{-1}\mathcal{C}_1^{-1}y$, thus on $\mathcal{X}^N$ we have that $\mu^{N,y} = \mathcal{N}(\tilde{m}^N, \tilde{\mathcal{C}}^N)$. On $(\mathcal{X}^N)^\perp$, the Radon-Nikodym derivative in (2.5.3) is equal to 1, hence $\mu^{N,y} = \mu_0 = \mathcal{N}(0, \tau^2\mathcal{C}_0)$. □

Proposition 2.5.5. *Under the Assumptions 2.3.1(1),(2),(3),(4),(5),(6), for all $y \in X^{\beta-s}$, $s = s_0 + \varepsilon$, where $0 < \varepsilon < (\Delta - 2s_0) \wedge (1 - s_0)$, the measures $\mu^{N,y}$ converge weakly in $\mathcal{X}$ to μ^y, where μ^y is defined in Theorem 2.5.1. In particular, $\mu^{N,y}$ converge weakly in $\mathcal{X}$ to μ^y ν-almost surely.*

Proof. Fix $y \in X^{\beta-s}$. Let $f : \mathcal{X} \to \mathbb{R}$ be continuous and bounded. Then by (2.5.1), (2.5.3) and Lemma 2.3.5(i), we have that

$$\int_{\mathcal{X}} f(u)\mu^{N,y}(du) = \frac{1}{Z^N}\int_{X^{s+\beta-2\ell}} f(u)e^{-\Phi^N(u,y)}\mu_0(du)$$

and

$$\int_{\mathcal{X}} f(u)\mu^{y}(du) = \frac{1}{Z}\int_{X^{s+\beta-2\ell}} f(u)e^{-\Phi(u,y)}\mu_0(du).$$

Let $u \in X^{s+\beta-2\ell}$ and set $r_1 = \max\{\|u\|_{s+\beta-2\ell}, \|y\|_{\beta-s}\}$ to get by Lemma 2.5.2(iii) that $\Phi^N(u,y) \to \Phi(u,y)$ as $N \to \infty$, since $\left\|P^N u\right\|_{s+\beta-2\ell} \leq \|u\|_{s+\beta-2\ell} \leq r_1$. By Lemma 2.5.2(i), for any $\delta > 0$, for $r_2 = \left\|y\right\|_{\beta-s}$, there exists $M(\delta, r_2) \in \mathbb{R}$ such that

$$\left|f(u)e^{-\Phi^N(u,y)}\right| \leq \left\|f\right\|_\infty e^{\delta\|u\|^2_{s+\beta-2\ell} - M(\delta,r_2)},\ \forall u \in X^{s+\beta-2\ell},$$

where the right hand side is μ_0-integrable for δ sufficiently small by the Fernique Theorem [13, Theorem 2.8.5]. Hence, by the Dominated Convergence Theorem, we have that $\int_{\mathcal{X}} f(u)\mu^{N,y}(du) \to \int_{\mathcal{X}} f(u)\mu^y(du)$, as $N \to \infty$, where we get the convergence of the constants $Z^N \to Z$ by choosing $f \equiv 1$. Thus we have $\mu^{N,y} \Rightarrow \mu^y$. Recalling, that $y \in X^{\beta-s}$ ν-almost surely completes the proof. □

We are now ready to prove Theorem 2.2.1:

Proof of Theorem 2.2.1. By Proposition 2.5.5 we have that $\mu^{N,y}$ converge weakly in $\mathcal{X}$ to the measure μ^y, ν-almost surely. Since by Lemma 2.5.4, the measures $\mu^{N,y}$ are Gaussian, the limiting measure μ^y is also Gaussian. To see this we argue as follows. The weak convergence of measures implies the pointwise convergence of the Fourier transforms of the measures, thus by Levy's continuity theorem [41, Theorem 4.3] all the one dimensional projections of $\mu^{N,y}$, which are Gaussian, converge weakly to the corresponding one dimensional projections of μ^y. By the fact that the class of Gaussian distributions in $\mathbb{R}$ is closed under weak convergence [41, Chapter 4, Exercise 2], we get that all the one dimensional projections of the μ^y are Gaussian, thus μ^y is a Gaussian measure in $\mathcal{X}$, $\mu^y = \mathcal{N}(\overline{m}, \overline{\mathcal{C}})$ for some $\overline{m} \in \mathcal{X}$ and a self-adjoint, positive semi definite, trace class linear operator $\overline{\mathcal{C}}$. It suffices to show that $\overline{m} = m$ and $\overline{\mathcal{C}} = \mathcal{C}$.

We use the standard Galerkin method to show that $m^N \to m$ in $\mathcal{X}$. Indeed, since by their definition m^N solve (2.1.7) in the N-dimensional spaces $\mathcal{X}^N$, for $e = m - m^N$, we have that $B(e, v) = 0$, $\forall v \in \mathcal{X}^N$. By the coercivity and the continuity of B (see Proposition 2.4.2)

$$\|e\|_1^2 \leq cB(e,e) = cB(e, m-z) \leq c\|e\|_1 \|m-z\|_1, \ \forall z \in \mathcal{X}^N.$$

Choose $z = P^N m$ to obtain

$$\|m - m^N\| \leq c\|m - P^N m\|_1,$$

where as $N \to \infty$ the right hand side converges to zero since $m \in X^1$. On the other hand, by [13, Example 3.8.15], we have that $m^N \to \overline{m}$ in $\mathcal{X}$, hence we conclude that $\overline{m} = m$, as required.

For the identification of the covariance operator, note that by the definition of $\mathcal{C}^N$ we have

$$\mathcal{C}^N = P^N \mathcal{C} P^N + (I - P^N)\mathcal{C}_0(I - P^N).$$

Recall that $\{\phi_j\}_{j=1}^{\infty}$ are the eigenfunctions of $\mathcal{C}_0$ and fix $j \in \mathbb{N}$. Then, for $N > j$ and any $w \in \mathcal{X}$, we have that

$$\begin{aligned} \left|\langle w, \mathcal{C}^N \phi_j \rangle - \langle w, \mathcal{C}\phi_j \rangle\right| &= \left|\langle w, (P^N - I)\mathcal{C}\phi_j \rangle\right| \\ &\leq \|(P^N - I)w\| \|\mathcal{C}\phi_j\|, \end{aligned}$$

where the right hand side converges to zero as $N \to \infty$, since $w \in \mathcal{X}$. This implies that $\mathcal{C}^N \phi_j$ converges to $\mathcal{C}\phi_j$ weakly in $\mathcal{X}$, as $N \to \infty$ and this holds for any $j \in \mathbb{N}$.

On the other hand by [13, Example 3.8.15], we have that $\mathcal{C}^N\phi_j \to \overline{\mathcal{C}}\phi_j$ in $\mathcal{X}$, as $N \to \infty$, for all $j \in \mathbb{N}$. It follows that $\overline{\mathcal{C}}\phi_j = \mathcal{C}\phi_j$, for every j and since $\{\phi_j\}_{j=1}^{\infty}$ is an orthonormal basis of $\mathcal{X}$, we have that $\overline{\mathcal{C}} = \mathcal{C}$. □

2.6 Operator norm bounds on $\mathcal{B}_\lambda^{-1}$

The following propositions contain several operator norm estimates on the inverse of $\mathcal{B}_\lambda$ and related quantities, and in particular estimates on the singular dependence of this operator as $\lambda \to 0$. These are the key tools used in Section 2.7 to obtain posterior contraction results. In all of them we make use of the interpolation inequality in Hilbert scales, [22, Proposition 8.19]. Recall that we consider $\mathcal{B}_\lambda$ defined on X^1, as explained in Remark 2.4.3.

Proposition 2.6.1. *Let $\kappa = (1-\theta)(\beta-2\ell)+\theta$, where $\theta \in [0,1]$. Under the Assumption 2.3.1(3) the following operator norm bounds hold: there is $c > 0$ independent of θ such that*

$$\left\|\mathcal{B}_\lambda^{-1}\right\|_{\mathcal{L}(X^{-\kappa},X^{\beta-2\ell})} \leq c\lambda^{-\frac{\theta}{2}}$$

and

$$\left\|\mathcal{B}_\lambda^{-1}\right\|_{\mathcal{L}(X^{-\kappa},X^{1})} \leq c\lambda^{-\frac{\theta+1}{2}}.$$

In particular, if $\beta - 2\ell \leq 0$, interpolation of the two bounds gives

$$\left\|\mathcal{B}_\lambda^{-1}\right\|_{\mathcal{L}(X^{-\kappa},\mathcal{X})} \leq c\lambda^{-\frac{\theta+\theta_0}{2}},$$

where $\theta_0 = \frac{2\ell-\beta}{\Delta} \in [0,1]$.

Proof. Let $h \in X^{-\kappa} \subset X^{-1}$. Then by Proposition 2.4.2 for $r = h$, there exists a unique weak solution of (2.4.1), $z \in X^1$. By Definition 2.4.1, for $v = z \in X^1$, we get

$$\left\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}z\right\|^2 + \lambda\left\|\mathcal{C}_0^{-\frac{1}{2}}z\right\|^2 = \langle \mathcal{C}_0^{\frac{\kappa}{2}}h, \mathcal{C}_0^{-\frac{\kappa}{2}}z\rangle.$$

Using the Assumption 2.3.1(3), and the Cauchy-Schwarz inequality, we get

$$\|z\|_{\beta-2\ell}^2 + \lambda\|z\|_1^2 \leq c\|\mathcal{C}_0^{\frac{\kappa}{2}}h\|\|z\|_\kappa.$$

We interpolate the norm on z appearing on the right hand side between the norms on z appearing on the left hand side, then use the Cauchy with ε inequality, and then Young's inequality for $p = \frac{1}{1-\theta}, q = \frac{1}{\theta}$, to get successively, for $c > 0$ a changing constant

$$\|z\|_{\beta-2\ell}^2 + \lambda\|z\|_1^2 \leq c\|\mathcal{C}_0^{\frac{\kappa}{2}}h\|\|z\|_{\beta-2\ell}^{1-\theta}\lambda^{-\frac{\theta}{2}}\left(\lambda^{\frac{1}{2}}\|z\|_1\right)^\theta$$

$$\leq \frac{c}{2\varepsilon}\left(\lambda^{-\theta}\|\mathcal{C}_0^{\frac{\kappa}{2}}h\|^2\right) + \frac{c\varepsilon}{2}\left(\|z\|_{\beta-2\ell}^{2(1-\theta)}\left(\lambda^{\frac{1}{2}}\|z\|_1\right)^{2\theta}\right)$$

$$\leq \frac{c}{2\varepsilon}\left(\lambda^{-\theta}\|\mathcal{C}_0^{\frac{\kappa}{2}}h\|^2\right) + \frac{c\varepsilon}{2}\left((1-\theta)\|z\|_{\beta-2\ell}^2 + \theta\lambda\|z\|_1^2\right).$$

By choosing $\varepsilon > 0$ small enough we get, for $c > 0$ independent of θ, λ,

$$\|z\|_{\beta-2\ell} \leq c\lambda^{-\frac{\theta}{2}}\|\mathcal{C}_0^{\frac{\kappa}{2}}h\| \quad \text{and} \quad \|z\|_1 \leq c\lambda^{-\frac{\theta+1}{2}}\|\mathcal{C}_0^{\frac{\kappa}{2}}h\|.$$

Replacing $z = \mathcal{B}_\lambda^{-1}h$ gives the result. □

Proposition 2.6.2. *Let $\kappa = (1-\theta)(\beta-2\ell-s)+\theta(1-s)$, where $\theta \in [0,1]$ and $s \in (s_0, 1]$, where $s_0 \in [0,1)$ as defined in Assumption 2.3.1(1). Under the Assumptions 2.3.1(2) and (3), the following norm bounds hold: there is $c > 0$ independent of θ such that*

$$\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}\|_{\mathcal{L}(X^{-\kappa}, X^{\beta-2\ell-s})} \leq c\lambda^{-\frac{\theta}{2}}$$

and

$$\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}\|_{\mathcal{L}(X^{-\kappa}, X^{1-s})} \leq c\lambda^{-\frac{\theta+1}{2}}.$$

In particular,

$$\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}\|_{\mathcal{L}(X)} \leq c\lambda^{-\frac{2\ell-\beta+s}{\Delta}}, \quad \forall s \in (\{\beta-2\ell\} \vee s_0, 1].$$

Proof. Let $h \in X^{-\kappa} = X^{(1-\theta)\Delta+s-1}$. Then $h \in X^{s-1}$, since $\Delta > 0$, thus $\mathcal{C}_0^{-\frac{s}{2}}h \in X^{-1}$. By Proposition 2.4.2 for $r = \mathcal{C}_0^{-\frac{s}{2}}h$, there exists a unique weak solution of (2.4.1), $z' \in X^1$. Since for $v \in X^{1-s}$ we have that $\mathcal{C}_0^{\frac{s}{2}}v \in X^1$, we conclude that for any $v \in X^{1-s}$

$$\langle \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}\mathcal{C}_0^{\frac{s}{2}}z, \mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}\mathcal{C}_0^{\frac{s}{2}}v\rangle + \lambda\langle \mathcal{C}_0^{\frac{s-1}{2}}z, \mathcal{C}_0^{\frac{s-1}{2}}v\rangle = \langle \mathcal{C}_0^{-\frac{s}{2}}h, \mathcal{C}_0^{\frac{s}{2}}v\rangle,$$

where $z = \mathcal{C}_0^{-\frac{s}{2}}z' \in X^{1-s}$. Choosing $v = z \in X^{1-s}$, we get

$$\|\mathcal{C}_1^{-\frac{1}{2}}\mathcal{A}^{-1}\mathcal{C}_0^{\frac{s}{2}}z\|^2 + \lambda\|\mathcal{C}_0^{\frac{s-1}{2}}z\|^2 = \langle h, z\rangle.$$

By the Assumption 2.3.1(3) and the Cauchy-Schwarz inequality, we have

$$\|z\|_{\beta-2\ell-s}^2 + \lambda\|z\|_{1-s}^2 \leq c\|h\|_{-\kappa}\|z\|_\kappa.$$

We interpolate the norm of z appearing on the right hand side between the norms of z appearing on the left hand side, to get as in the proof of Proposition

2.6.1, for $c > 0$ independent of θ, λ and s

$$\|z\|_{\beta-2\ell-s} \leq c\lambda^{-\frac{\theta}{2}}\|h\|_{-\kappa} \quad \text{and} \quad \|z\|_{1-s} \leq c\lambda^{-\frac{\theta+1}{2}}\|h\|_{-\kappa}.$$

Replacing $z = \mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}h$ gives the first two rates.
For the last claim, note that we can always choose $\{\beta-2\ell\}\vee\{s_0\} < s \leq 1$, since $s_0 < 1$ and $\Delta > 0$. Using the first two estimates, for $\kappa = (1-\theta')(\beta-2\ell-s)+\theta'(1-s) = 0$, that is $\theta' = \frac{2\ell-\beta+s}{\Delta} \in [0,1]$, we have that

$$\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}h\|_{\beta-2\ell-s} \leq c\lambda^{-\frac{\theta'}{2}}\|h\|$$

and

$$\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}h\|_{1-s} \leq c\lambda^{-\frac{\theta'+1}{2}}\|h\|.$$

Using the interpolation inequality we then get the claim, since

$$\begin{aligned}\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}h\| &\leq \|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}h\|_{\beta-2\ell-s}^{1-\theta'}\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}h\|_{1-s}^{\theta'} \\ &\leq c\lambda^{-\theta'}\|h\|.\end{aligned}$$

□

2.7 Posterior Contraction

In this section we employ the developments of the preceding sections to study the posterior consistency of the Bayesian solution to the inverse problem. That is, we consider a family of data sets $y = y_n^\dagger$ given by (2.1.10) and study the limiting behaviour of the posterior measure $\mu_{\lambda,n}^{y_n^\dagger} = \mathcal{N}(m^\dagger, \mathcal{C})$ as $n \to \infty$. Intuitively we would hope to recover a measure which concentrates near the true solution $u^\dagger$ in this limit. Following the approach in [44], [28], [83] and [24], we quantify this idea as in (2.1.12). By the Markov inequality we have

$$\mathbb{E}^{y_n^\dagger}\mu_{\lambda,n}^{y_n^\dagger}\left\{u : \|u-u^\dagger\| \geq M_n\epsilon_n\right\} \leq \frac{1}{M_n^2\epsilon_n^2}\mathbb{E}^{y_n^\dagger}\int \|u-u^\dagger\|^2\mu_{\lambda,n}^{y_n^\dagger}(du),$$

so that it suffices to show that

$$\mathbb{E}^{y_n^\dagger}\int \|u-u^\dagger\|^2\mu_{\lambda,n}^{y_n^\dagger}(du) \leq c\epsilon_n^2. \tag{2.7.1}$$

In addition to n^{-1}, there is a second small parameter in the problem, namely the regularization parameter, $\lambda = \frac{1}{n\tau^2}$, and we will choose a relationship between n and

λ in order to optimize the convergence rates ϵ_n. We will show that determination of optimal convergence rates follows directly from the operator norm bounds on $\mathcal{B}_\lambda^{-1}$ derived in the previous section, which concern only λ dependence; relating n to λ then follows as a trivial optimization. Thus, the λ dependence of the operator norm bounds in the previous section forms the heart of the posterior contraction analysis. The relationship between λ and n will induce a relationship between τ and n, where τ being the scaling parameter in the prior covariance is the relevant parameter in the current Bayesian framework.

We now present our convergence results. In Theorem 2.7.1 we study the convergence of the posterior mean to the true solution in a range of norms, while in Theorem 2.7.2 we study the concentration of the posterior near the true solution as described in (2.1.12). The proofs of Theorems 2.7.1 and 2.7.2 are provided later in the current section. The two main convergence results, Theorems 2.2.2 and 2.2.3 follow as direct corollaries of Remark 2.7.3 and Theorems 2.7.1 and 2.7.2 respectively.

Theorem 2.7.1. *Let $u^\dagger \in X^1$. Under the Assumptions 2.3.1, we have that, for the choice $\tau = \tau(n) = n^{\frac{\theta_2-\theta_1-1}{2(\theta_1-\theta_2+2)}}$ and for any $\theta \in [0,1]$*

$$\mathbb{E}^{y_n^\dagger}\left\|m^\dagger - u^\dagger\right\|_\kappa^2 \leq cn^{\frac{\theta+\theta_2-2}{\theta_1-\theta_2+2}},$$

where $\kappa = (1-\theta)(\beta-2\ell)+\theta$. The result holds for any $\theta_1, \theta_2 \in [0,1]$, chosen so that $\mathbb{E}(\iota^2) < \infty$, for $\iota = \max\left\{\left\|\eta\right\|_{2\beta-2\ell-\kappa_1}, \left\|u^\dagger\right\|_{2-\kappa_2}\right\}$, where $\kappa_i = (1-\theta_i)(\beta-2\ell) + \theta_i$, $i = 1,2$.

Theorem 2.7.2. *Let $u^\dagger \in X^1$. Under the Assumptions 2.3.1, we have that, for $\tau = \tau(n) = n^{\frac{\theta_2-\theta_1-1}{2(\theta_1-\theta_2+2)}}$, the convergence in (2.1.12) holds with*

$$\epsilon_n = n^{\frac{\theta_0+\theta_2-2}{2(\theta_1-\theta_2+2)}}, \quad \theta_0 = \begin{cases} \frac{2\ell-\beta}{\Delta}, & if\ \beta-2\ell \leq 0 \\ 0, & otherwise. \end{cases}$$

The result holds for any $\theta_1, \theta_2 \in [0,1]$, chosen so that $\mathbb{E}(\iota^2) < \infty$, for $\iota = \max\left\{\left\|\eta\right\|_{2\beta-2\ell-\kappa_1}, \left\|u^\dagger\right\|_{2-\kappa_2}\right\}$, where $\kappa_i = (1-\theta_i)(\beta-2\ell)+\theta_i$, $i = 1,2$.

Remark 2.7.3. *i) To get convergence in the PDE method we need $\mathbb{E}\left\|u^\dagger\right\|_{2-\kappa_2}^2 < \infty$ for a $\theta_2 \leq 1$. Under the a priori information that $u^\dagger \in X^\gamma$, we need $\gamma \geq 2-\kappa_2 = 1+(1-\theta_2)\Delta$ for some $\theta_2 \in [0,1]$. Thus the minimum requirement for convergence is $\gamma = 1$ in agreement to our assumption $u^\dagger \in X^1$. On the other hand, to obtain the optimal rate (which corresponds to choosing θ_2 as small as possible) we need to choose $\theta_2 = \frac{\Delta+1-\gamma}{\Delta}$. If $\gamma > 1+\Delta$ then the right*

hand side is negative so we have to choose $\theta_2 = 0$, hence we cannot achieve the optimal rate. We say that the method saturates at $\gamma = 1+\Delta$ which reflects the fact that the true solution has more regularity than the method allows us to exploit to obtain faster convergence rates.

ii) To get convergence we also need $\mathbb{E}\|\eta\|^2_{2\beta-2\ell-\kappa_1} < \infty$ for a $\theta_1 \leq 1$. By Lemma 2.3.5(iii), it suffices to have $\theta_1 > \frac{s_0}{\Delta}$. This means that we need $\Delta > s_0$, which holds by the Assumption 2.3.1(2), in order to be able to choose $\theta_1 \leq 1$. On the other hand, since $\Delta > 0$ and $s_0 \geq 0$, we have that $\frac{s_0}{\Delta} \geq 0$ thus we can always choose θ_1 in an optimal way, that is, we can always choose $\theta_1 = \frac{s_0+\varepsilon}{\Delta}$ where $\varepsilon > 0$ is arbitrarily small.

iii) If we want draws from μ_0 to be in X^γ then by Lemma 2.3.3(ii) we need $1-s_0 > \gamma$. Since the requirement for the method to give convergence is $\gamma \geq 1$ while $1-s_0 \leq 1$, we can never have draws exactly matching the regularity of the prior. On the other hand if we want an undersmoothing prior (which according to [44] in the diagonal case gives asymptotic coverage equal to 1) we need $1 - s_0 \leq \gamma$, which we always have. This, as discussed in Section 2.1, gives an explanation to the observation that in both of the above theorems we always have $\tau \to 0$ as $n \to \infty$.

iv) When $\beta - 2\ell > 0$, in Theorem 2.7.2 and in Theorem 2.2.3 below, we get suboptimal rates. The reason is that our analysis to obtain the error in the $\mathcal{X}$-norm is based on interpolating between the error in the $X^{\beta-2\ell}$-norm and the error in the X^1-norm. When $\beta - 2\ell > 0$, interpolation is not possible since the $\mathcal{X}$-norm is now weaker than the $X^{\beta-2\ell}$-norm. However, we can at least bound the error in the $\mathcal{X}$-norm by the error in the $X^{\beta-2\ell}$-norm, thus obtaining a suboptimal rate. Note, that the case $\beta-2\ell > 0$ does not necessarily correspond to the well posed case: by Lemma 2.3.5 we can only guarantee that a draw from the noise distribution lives in X^ρ, $\rho < \beta - s_0$, while the range of $\mathcal{A}^{-1}$ is formally $X^{2\ell}$. Hence, in order to have a well posed problem we need $\beta - s_0 > 2\ell$, or equivalently $\Delta < 1 - s_0$. This can happen despite our assumption $\Delta > 2s_0$, when $s_0 < 1/3$ and for appropriate choice of ℓ and β. In this case, regularization is unnecessary.

Note that, since the posterior is Gaussian, the left hand side in (2.7.1) is the Square Posterior Contraction

$$SPC = \mathbb{E}^{y_n^\dagger}\left\|m^\dagger - u^\dagger\right\|^2 + \mathrm{Tr}(\mathcal{C}), \tag{2.7.2}$$

which is the sum of the mean integrated squared error (MISE) of the posterior mean and the posterior spread. Let $u^\dagger \in X^1$. By Lemma 2.4.4, the relationship (2.1.10) between $u^\dagger$ and $y_n^\dagger$ and the equation (2.1.11) for $m^\dagger$, we obtain

$$\mathcal{B}_\lambda m^\dagger = \mathcal{A}^{-1}\mathcal{C}_1^{-1}y_n^\dagger = \mathcal{A}^{-1}\mathcal{C}^{-1}\mathcal{A}^{-1}u^\dagger + \frac{1}{\sqrt{n}}\mathcal{A}^{-1}\mathcal{C}^{-1}\eta$$
$$\text{and} \quad \mathcal{B}_\lambda u^\dagger = \mathcal{A}^{-1}\mathcal{C}^{-1}\mathcal{A}^{-1}u^\dagger + \lambda\mathcal{C}_0^{-1}u^\dagger,$$

where the equations hold in X^{-1}, since by a similar argument to the proof of Proposition 2.4.5 we have $m^\dagger \in X^1$. By subtraction we get

$$\mathcal{B}_\lambda(m^\dagger - u^\dagger) = \frac{1}{\sqrt{n}}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta - \lambda\mathcal{C}_0^{-1}u^\dagger.$$

Therefore

$$m^\dagger - u^\dagger = \mathcal{B}_\lambda^{-1}\left(\frac{1}{\sqrt{n}}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta - \lambda\mathcal{C}_0^{-1}u^\dagger\right), \tag{2.7.3}$$

as an equation in X^1. Using the fact that the noise has mean zero and the relation (2.1.6), equation (2.7.3) implies that we can split the square posterior contraction into three terms

$$SPC = \left\|\lambda\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-1}u^\dagger\right\|^2 + \mathbb{E}\left\|\frac{1}{\sqrt{n}}\mathcal{B}_\lambda^{-1}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta\right\|^2 + \frac{1}{n}\mathrm{Tr}(B_\lambda^{-1}), \tag{2.7.4}$$

provided the right hand side is finite. A consequence of the proof of Theorem 2.2.1 is that $\mathcal{B}_\lambda^{-1}$ is trace class. Note that for ζ a white noise, we have that

$$\mathrm{Tr}(\mathcal{B}_\lambda^{-1}) = \mathbb{E}\left\|\mathcal{B}_\lambda^{-\frac{1}{2}}\zeta\right\|^2 = \mathbb{E}\langle\zeta, \mathcal{B}_\lambda^{-1}\zeta\rangle = \mathbb{E}\langle\mathcal{C}_0^{\frac{s}{2}}\zeta, \mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}\mathcal{C}_0^{\frac{s}{2}}\zeta\rangle$$
$$\leq \left\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}\right\|_{\mathcal{L}(\mathcal{X})}\mathbb{E}\left\|\mathcal{C}_0^{\frac{s}{2}}\zeta\right\|^2,$$

which for $s > s_0$ since by Lemma 2.3.3 we have that $\mathbb{E}\left\|\mathcal{C}_0^{\frac{s}{2}}\zeta\right\|^2 < \infty$, provides the bound

$$\mathrm{Tr}(\mathcal{B}_\lambda^{-1}) \leq c\left\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}\right\|_{\mathcal{L}(\mathcal{X})}, \tag{2.7.5}$$

where $c > 0$ is independent of λ. If q is chosen sufficiently small and r sufficiently large so that $\left\|\mathcal{C}_0^{-\frac{q}{2}-1}u^\dagger\right\| < \infty$ and $\mathbb{E}\left\|\mathcal{C}_0^{\frac{r}{2}}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta\right\|^2 < \infty$ then we see that

$$SPC \leq c\left(\lambda^2\left\|\mathcal{B}_\lambda^{-1}\right\|^2_{\mathcal{L}(X^q,\mathcal{X})} + \frac{1}{n}\left\|\mathcal{B}_\lambda^{-1}\right\|^2_{\mathcal{L}(X^{-r},\mathcal{X})} + \frac{1}{n}\left\|\mathcal{C}_0^{-\frac{s}{2}}\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-\frac{s}{2}}\right\|_{\mathcal{L}(\mathcal{X})}\right), \tag{2.7.6}$$

where $c > 0$ is independent of λ and n. Thus identifying ϵ_n in (2.1.12) can be achieved simply through properties of the inverse of $\mathcal{B}_\lambda$ and its parametric dependence on λ.

In the following, we are going to study convergence rates for the square posterior contraction, (2.7.4), which by the previous analysis will secure that

$$\mathbb{E}^{y_n^\dagger}\mu_{\lambda,n}^{y_n^\dagger}\left\{u\colon \left\|u-u^\dagger\right\| \geq \epsilon_n\right\} \to 0,$$

for $\epsilon_n^2 \to 0$ at a rate almost as fast as the square posterior contraction. This suggests that the error is determined by the MISE of the posterior mean and the trace of the posterior covariance, thus we optimize our analysis with respect to these two quantities. In [44] the situation where $\mathcal{C}_0, \mathcal{C}_1$ and $\mathcal{A}$ are diagonalizable in the same eigenbasis is studied, and it is shown that the third term in equation (2.7.4) is bounded by the second term in terms of their parametric dependence on λ. The same idea is used in the proof of Theorem 2.7.2.

We now provide the proofs of Theorem 2.7.1 and Theorem 2.7.2.

Proof of Theorem 2.7.1. Since η has zero mean, we have by (2.7.3)

$$\mathbb{E}\left\|m^\dagger - u^\dagger\right\|_{\beta-2\ell}^2 = \lambda^2\left\|\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-1}u^\dagger\right\|_{\beta-2\ell}^2 + \frac{1}{n}\mathbb{E}\left\|\mathcal{B}_\lambda^{-1}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta\right\|_{\beta-2\ell}^2$$

and

$$\mathbb{E}\left\|m^\dagger - u^\dagger\right\|_1^2 = \lambda^2\left\|\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-1}u^\dagger\right\|_1^2 + \frac{1}{n}\mathbb{E}\left\|\mathcal{B}_\lambda^{-1}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta\right\|_1^2.$$

Using Proposition 2.6.1 and Assumption 2.3.1(7), we get

$$\mathbb{E}\left\|m^\dagger - u^\dagger\right\|_{\beta-2\ell}^2 \leq c\mathbb{E}(\iota^2)(\lambda^{2-\theta_2} + \frac{1}{n}\lambda^{-\theta_1}) = c\mathbb{E}(\iota^2)(n^{\theta_2-2}\tau^{2\theta_2-4} + n^{\theta_1-1}\tau^{2\theta_1})$$

and

$$\mathbb{E}\left\|m^\dagger - u^\dagger\right\|_1^2 \leq c\mathbb{E}(\iota^2)(\lambda^{1-\theta_2} + \frac{1}{n}\lambda^{-\theta_1-1}) = \frac{c\mathbb{E}(\iota^2)}{\lambda}(n^{\theta_2-2}\tau^{2\theta_2-4} + n^{\theta_1-1}\tau^{2\theta_1}).$$

Since the common parenthesis term, consists of a decreasing and an increasing term in τ, we optimize the rate by choosing $\tau = \tau(n) = n^p$ such that the two terms become equal, that is, $p = \frac{\theta_2-\theta_1-1}{2(\theta_1-\theta_2+2)}$. We obtain,

$$\mathbb{E}\left\|m^\dagger - u^\dagger\right\|_{\beta-2\ell}^2 \leq c\mathbb{E}(\iota^2)n^{\frac{\theta_2-2}{\theta_1-\theta_2+2}} \quad \text{and} \quad \mathbb{E}\left\|m^\dagger - u^\dagger\right\|_1^2 \leq c\mathbb{E}(\iota^2)n^{\frac{\theta_2-1}{\theta_1-\theta_2+2}}.$$

By interpolating between the two last estimates we obtain the claimed rate. □

Proof of Theorem 2.7.2. Recall equation (2.7.4)

$$SPC = \left\|\lambda\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-1}u^\dagger\right\|^2 + \mathbb{E}\left\|\frac{1}{\sqrt{n}}\mathcal{B}_\lambda^{-1}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta\right\|^2 + \frac{1}{n}\mathrm{Tr}(B_\lambda^{-1}).$$

The idea is that the third term is always dominated by the second term. Combining equation (2.7.5) with Proposition 2.6.2, we have that

$$\frac{1}{n}\mathrm{Tr}(\mathcal{B}_\lambda^{-1}) \leq c\frac{1}{n}\lambda^{-\frac{2\ell-\beta+s}{\Delta}},\ \forall s \in (\{\beta-2\ell\}\vee\{s_0\}, 1].$$

i) Suppose $\beta - 2\ell \leq 0$, so that by Proposition 2.6.1 we have, where $\theta_0 = \frac{2\ell-\beta}{\Delta} \in [0,1]$, using Assumption 2.3.1(7)

$$\mathbb{E}\left\|\frac{1}{\sqrt{n}}\mathcal{B}_\lambda^{-1}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta\right\|^2 \leq c\frac{1}{n}\mathbb{E}\left\|\eta\right\|^2_{2\beta-2\ell-\kappa_1}\lambda^{-\theta_1-\theta_0}$$

and

$$\left\|\lambda\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-1}u^\dagger\right\|^2 \leq c\left\|u^\dagger\right\|^2_{2-\kappa_2}\lambda^{2-\theta_2-\theta_0}.$$

Note that θ_1 is chosen so that $\mathbb{E}\left\|\eta\right\|^2_{2\beta-2\ell-\kappa_1} < \infty$, that is, by Lemma 2.3.5(iii), it suffices to have $\theta_1 > \frac{s_0}{\Delta}$. Noticing that by choosing s arbitrarily close to s_0, we can have $\frac{2\ell-\beta+s}{\Delta}$ arbitrarily close to $\frac{2\ell-\beta+s_0}{\Delta}$, and since $\theta_1+\theta_0 > \frac{2\ell-\beta+s_0}{\Delta}$, we deduce that the third term in equation (2.7.4) is always dominated by the second term. Combining, we have that

$$SPC \leq \frac{c\mathbb{E}(\iota^2)}{\lambda^{\theta_0}}(\lambda^{2-\theta_2} + \frac{1}{n}\lambda^{-\theta_1}) = \frac{c\mathbb{E}(\iota^2)}{\lambda^{\theta_0}}(n^{\theta_2-2}\tau^{2\theta_2-4} + n^{\theta_1-1}\tau^{2\theta_1}).$$

ii) Suppose $\beta - 2\ell > 0$. Using Proposition 2.6.1 and Assumption 2.3.1(7) we have

$$\left\|\lambda\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-1}u^\dagger\right\|^2 \leq c\left\|\lambda\mathcal{B}_\lambda^{-1}\mathcal{C}_0^{-1}u^\dagger\right\|^2_{\beta-2\ell} \leq c\left\|u^\dagger\right\|^2_{2-\kappa_2}\lambda^{2-\theta_2}$$

and

$$\begin{aligned}\mathbb{E}\left\|\frac{1}{\sqrt{n}}\mathcal{B}_\lambda^{-1}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta\right\|^2 &\leq c\mathbb{E}\left\|\frac{1}{\sqrt{n}}\mathcal{B}_\lambda^{-1}\mathcal{A}^{-1}\mathcal{C}_1^{-1}\eta\right\|^2_{\beta-2\ell} \\ &\leq c\frac{1}{n}\mathbb{E}\left\|\eta\right\|^2_{2\beta-2\ell-\kappa_1}\lambda^{-\theta_1},\end{aligned}$$

where as before $\theta_1 > \frac{s_0}{\Delta}$. The third term in equation (2.7.4) is again dominated by the second term, since on the one hand $\theta_1 > \frac{s_0}{\Delta}$ and on the other hand, since $\beta - 2\ell > 0$, we can always choose $\{\beta-2\ell\}\vee\{s_0\} < s \leq 1\wedge\{s_0+\beta-2\ell\}$

to get $\frac{2\ell-\beta+s}{\Delta} \leq \frac{s_0}{\Delta}$. Combining the three estimates we have that

$$SPC \leq c\mathbb{E}(\iota^2)(n^{\theta_2-2}\tau^{2\theta_2-4} + n^{\theta_1-1}\tau^{2\theta_1}).$$

In both cases, the common term in the parenthesis consists of a decreasing and an increasing term in τ, thus we can optimize by choosing $\tau = \tau(n) = n^p$ making the two terms equal, that is, $p = \frac{\theta_1-\theta_2+1}{2\theta_2-2\theta_1-4}$, to get the claimed rates. □

2.8 Examples

We now present some nontrivial examples satisfying Assumptions 2.3.1.

Let $\Omega \subset \mathbb{R}^d$, $d = 1, 2, 3$, be a bounded and open set. We define $\mathcal{A}_0 := -\boldsymbol{\Delta}$, where $\boldsymbol{\Delta}$ is the Dirichlet Laplacian which is the Friedrichs extension of the classical Laplacian defined on $C_0^2(\Omega)$, that is, $\mathcal{A}_0$ is a self-adjoint operator with a domain $\mathcal{D}(\mathcal{A}_0)$ dense in $\mathcal{X} := L^2(\Omega)$ [50]. For $\partial\Omega$ sufficiently smooth we have $\mathcal{D}(\mathcal{A}_0) = H^2(\Omega) \cap H_0^1(\Omega)$. It is well known that $\mathcal{A}_0$ has a compact inverse and that it possesses an eigensystem $\{\rho_j, e_j\}_{j=1}^{\infty}$, where the eigenfunctions $\{e_j\}$ form a complete orthonormal basis of $\mathcal{X}$ and the eigenvalues ρ_j are positive and behave asymptotically like $j^{\frac{2}{d}}$ [5].

In Subsections 2.8.1 and 2.8.2, we consider the inverse problem to find u from y, where

$$y = z + \frac{1}{\sqrt{n}}\eta,$$

for z solving the partial differential equation

$$\begin{aligned} -\boldsymbol{\Delta}z + qz &= u \quad \text{in} \quad \Omega, \\ z &= 0 \quad \text{on} \quad \partial\Omega, \end{aligned}$$

that is, $\mathcal{A}_0 z + qz = u$, where q is a nonnegative real function of certain regularity. We choose prior and noise distributions with covariance operators which are not simultaneously diagonalizable with the forward operator. Later on, in Subsection 2.8.3, we consider more complicated examples and in particular, we consider fractional powers of the Dirichlet Laplacian in the forward operator, as well as more general choices of prior and noise covariance operators.

Our general strategy for proving the validity of our norm equivalence assumptions is:

i) if needed, use Proposition 2.8.6 below to reduce the range of spaces required to check an assumption's validity to a finite set of spaces;

ii) reformulate the assumptions as statements regarding the boundedness of operators of the form considered in Lemma 2.8.7 below.

The statement of Proposition 2.8.6, which is a well known result from interpolation theory, and the statement and proof of Lemma 2.8.7 are postponed to Subsection 2.8.4.

2.8.1 Example 1 - Non-diagonal forward operator

We study the Bayesian inversion of the operator $\mathcal{A}^{-1} := (\mathcal{A}_0 + \mathcal{M}_q)^{-1}$ where $\mathcal{M}_q \colon L^2(\Omega) \to L^2(\Omega)$ is the multiplication operator by a nonnegative function $q \in W^{2,\infty}(\Omega)$. We assume that the observational noise is white, so that $\mathcal{C}_1 = I$, and we set the prior covariance operator to be $\mathcal{C}_0 = \mathcal{A}_0^{-2}$.

The operator $\mathcal{C}_0$ is trace class. Indeed, let $\lambda_j = \rho_j^{-2}$ be its eigenvalues. Then they behave asymptotically like $j^{-\frac{4}{d}}$ and $\sum_{j=1}^{\infty} j^{-\frac{4}{d}} < \infty$ for $d < 4$. Furthermore, we have that $\sum_{j=1}^{\infty} \lambda_j^s \leq c \sum_{j=1}^{\infty} j^{-\frac{4s}{d}} < \infty$, provided $s > \frac{d}{4}$, that is, the Assumption 2.3.1(1) is satisfied with

$$s_0 = \begin{cases} 1/4, & d = 1, \\ 1/2, & d = 2, \\ 3/4, & d = 3. \end{cases}$$

We define the Hilbert scale induced by $\mathcal{C}_0 = \mathcal{A}_0^{-2}$, that is, $(X^s)_{s\in\mathbb{R}}$, for $X^s := \overline{\mathcal{M}}^{\|\cdot\|_s}$, where

$$\mathcal{M} = \bigcap_{l=0}^{\infty} \mathcal{D}(\mathcal{A}_0^{2l}), \ \langle u, v \rangle_s := \langle \mathcal{A}_0^s u, \mathcal{A}_0^s v \rangle \quad \text{and} \quad \|u\|_s := \|\mathcal{A}_0^s u\|.$$

Observe, $X^0 = \mathcal{X} = L^2(\Omega)$.

Our aim is to show that $\mathcal{C}_1 \simeq \mathcal{C}_0^{\beta}$ and $\mathcal{A}^{-1} \simeq \mathcal{C}_0^{\ell}$, where $\beta = 0$ and $\ell = \frac{1}{2}$, in the sense of the Assumptions 2.3.1. We have $\Delta = 2\ell - \beta + 1 = 2$. Since for $d = 1, 2, 3$ we have $0 < s_0 < 1$, the Assumption 2.3.1(2) is satisfied. Moreover, note that since $\mathcal{C}_1 = I$ the Assumptions 2.3.1(4) and (5) are trivially satisfied.

We now show that Assumptions 2.3.1 (3), (6), (7) are also satisfied. In this example the three assumptions have the form

3. $\|(\mathcal{A}_0 + \mathcal{M}_q)^{-1} u\| \asymp \|\mathcal{A}_0^{-1} u\|, \ \forall u \in X^{-1}$;

6. $\|\mathcal{A}_0^s (\mathcal{A}_0 + \mathcal{M}_q)^{-1} u\| \leq c_3 \|\mathcal{A}_0^{s-1} u\|, \ \forall u \in X^{s-1}, \ \forall s \in (s_0, 1]$;

7. $\|\mathcal{A}_0^{-\kappa} (\mathcal{A}_0 + \mathcal{M}_q)^{-1} u\| \leq c_4 \|\mathcal{A}_0^{-\kappa-1} u\|, \ \forall u \in X^{-\kappa-1}, \ \forall \kappa \in [-1, 1]$.

Observe that Assumption (6) is implied by Assumption (7).

Proposition 2.8.1. *The Assumptions 2.3.1 are satisfied in this example.*

Proof. We only need to show that Assumptions (3) and (7) hold.

3. The assumption is equivalent to $\mathcal{T} := (\mathcal{A}_0 + \mathcal{M}_q)^{-1}\mathcal{A}_0$ and $\mathcal{T}^{-1} = \mathcal{A}_0^{-1}(\mathcal{A}_0 + \mathcal{M}_q)$ being bounded in $\mathcal{X}$. Since $\mathcal{T}^{-1} = I + \mathcal{A}_0^{-1}\mathcal{M}_q$ which is bounded in $\mathcal{X}$, we only need to show that $\mathcal{T}$ is bounded. Indeed, $(\mathcal{A}_0 + \mathcal{M}_q)^{-1}\mathcal{A}_0 = (I + \mathcal{A}_0^{-1}\mathcal{M}_q)^{-1}$, which is bounded by Lemma 2.8.7 applied for $t = -1, s = 1$.

7. By Proposition 2.8.6, it suffices to show $\mathcal{T} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^1) \cap \mathcal{L}(X^{-1})$. We have already shown that $\mathcal{T} \in \mathcal{L}(\mathcal{X})$. For $\mathcal{T} \in \mathcal{L}(X^1)$, note that it is equivalent to $\mathcal{A}_0\mathcal{T}\mathcal{A}_0^{-1} = (I + \mathcal{M}_q\mathcal{A}_0^{-1})^{-1} \in \mathcal{L}(\mathcal{X})$, which holds by Lemma 2.8.7 applied for $t = s = 1$. Finally, for $\mathcal{T} \in \mathcal{L}(X^{-1})$, note that it is equivalent to $\mathcal{A}_0^{-1}\mathcal{T}\mathcal{A}_0 = (I + \mathcal{A}_0^{-2}\mathcal{M}_q\mathcal{A}_0)^{-1} \in \mathcal{L}(\mathcal{X})$, which holds by Lemma 2.8.7 applied for $t = -1, s = 1$.

□

We can now apply Theorem 2.2.2 and Theorem 2.2.3 to get the following convergence result.

Theorem 2.8.2. *Let $u^\dagger \in X^\gamma, \gamma \geq 1$. Then, for $\tau = \tau(n) = n^{\frac{4-d-4(\gamma\wedge 3)-\varepsilon}{8(\gamma\wedge 3)+8+2d+2\varepsilon}}$, the convergence in (2.1.12) holds with $\epsilon_n = n^{-e}$, where*

$$e = \begin{cases} \frac{2\gamma}{4+d+4\gamma+2\varepsilon}, & if\ \gamma < 3 \\ \frac{6}{16+d+2\varepsilon}, & if\ \gamma \geq 3, \end{cases}$$

for $\varepsilon > 0$ arbitrarily small and where $d = 1, 2, 3$, is the dimension. Furthermore, for $t \in [-1, 1)$, for the same choice of τ, we have $\mathbb{E}\left\|m^\dagger - u^\dagger\right\|_t^2 \leq cn^{-h}$, where

$$h = \begin{cases} \frac{4\gamma-4t}{4+d+4\gamma+2\varepsilon}, & if\ \gamma < 3 \\ \frac{12-4t}{16+d+2\varepsilon}, & if\ \gamma \geq 3. \end{cases}$$

For $t = 1$ the above rate holds provided $\gamma > 1$.

2.8.2 Example 2 - A fully non-diagonal example

As in Example 2.8.1, we study the Bayesian inversion of the operator $\mathcal{A}^{-1} = (\mathcal{A}_0 + \mathcal{M}_q)^{-1}$, where $\mathcal{M}_q : L^2(\Omega) \to L^2(\Omega)$ is the multiplication operator by a nonnegative function $q \in W^{2,\infty}(\Omega)$. We assume that the observational noise is Gaussian with covariance operator $\mathcal{C}_1 := (\mathcal{A}_0^{\frac{1}{4}} + \mathcal{M}_r)^{-2}$, where $\mathcal{M}_r : L^2(\Omega) \to L^2(\Omega)$ is the multiplication operator by another nonnegative function $r \in W^{4,\infty}(\Omega)$. As

before, we set the prior covariance operator to be $\mathcal{C}_0 = \mathcal{A}_0^{-2}$, thus the Assumption 2.3.1(1) is satisfied with the same s_0 and we work in the same Hilbert scale $(X^s)_{s\in\mathbb{R}}$.

We show that $\mathcal{C}_1 \simeq \mathcal{C}_0^{\beta}$ and $\mathcal{A}^{-1} \simeq \mathcal{C}_0^{\ell}$, where $\beta = \frac{1}{4}$ and $\ell = \frac{1}{2}$, in the sense of the Assumptions 2.3.1(3)-(7). First note that we have $\Delta = 2\ell - \beta + 1 = \frac{7}{4} > 2s_0$ for $d = 1, 2, 3$, so that the Assumption 2.3.1(2) is satisfied. The rest of the assumptions have the form

3. $\big\|(\mathcal{A}_0^{\frac{1}{4}} + \mathcal{M}_r)(\mathcal{A}_0 + \mathcal{M}_q)^{-1}u\big\| \asymp \big\|\mathcal{A}_0^{-\frac{3}{4}}u\big\|,\ \forall u \in X^{-\frac{3}{4}}$;

4. $\big\|\mathcal{A}_0^{\rho}(\mathcal{A}_0^{\frac{1}{4}} + \mathcal{M}_r)^{-1}u\big\| \le c_1\big\|\mathcal{A}_0^{\rho-\frac{1}{4}}u\big\|,\ \forall u \in X^{\rho-\frac{1}{4}}, \forall \rho \in [\lceil -s_0 - \frac{3}{4}\rceil, \frac{1}{4} - s_0)$;

5. $\big\|\mathcal{A}_0^{-s}(\mathcal{A}_0^{\frac{1}{4}} + \mathcal{M}_r)u\big\| \le c_2\big\|\mathcal{A}_0^{\frac{1}{4}-s}\big\|,\ \forall u \in X^{\frac{1}{4}-s}, \forall s \in (s_0, 1]$;

6. $\big\|\mathcal{A}_0^{s}(\mathcal{A}_0^{\frac{1}{4}} + \mathcal{M}_r)(\mathcal{A}_0 + \mathcal{M}_q)^{-1}u\big\| \le c_3\big\|\mathcal{A}_0^{s-\frac{3}{4}}u\big\|,\ \forall u \in X^{s-\frac{3}{4}},\ \forall s \in (s_0, 1]$;

7. $\big\|\mathcal{A}_0^{-\kappa}(\mathcal{A}_0+\mathcal{M}_q)^{-1}(\mathcal{A}_0^{\frac{1}{4}}+\mathcal{M}_r)^2u\big\| \le c_4\big\|\mathcal{A}_0^{-\kappa-\frac{1}{2}}u\big\|,\ \forall u \in X^{-\kappa-\frac{1}{2}},\ \forall \kappa \in [-\frac{3}{4}, 1]$.

Proposition 2.8.3. *The Assumptions 2.3.1 are satisfied in this example.*

Proof. We have already seen that the first two assumptions are satisfied.

3. We need to show that $\mathcal{T} := (\mathcal{A}_0^{\frac{1}{4}} + \mathcal{M}_r)(\mathcal{A}_0 + \mathcal{M}_q)^{-1}\mathcal{A}_0^{\frac{3}{4}}$ and $\mathcal{T}^{-1}$ are bounded operators in $\mathcal{X}$. Indeed, $\mathcal{T} = (I + \mathcal{M}_r\mathcal{A}_0^{-\frac{1}{4}})(I + \mathcal{A}_0^{-\frac{3}{4}}\mathcal{M}_q\mathcal{A}_0^{-\frac{1}{4}})^{-1}$ which is bounded by Lemma 2.8.7 applied for $t = s = \frac{1}{4}$ and $t = \frac{1}{4}, s = 1$. For $\mathcal{T}^{-1}$ we have, $\mathcal{T}^{-1} = (I + \mathcal{A}_0^{-\frac{3}{4}}\mathcal{M}_q\mathcal{A}_0^{-\frac{1}{4}})(I + \mathcal{M}_r\mathcal{A}_0^{-\frac{1}{4}})^{-1}$, which again by Lemma 2.8.7 is the composition of two bounded operators.

4. Since $\frac{1}{4} - s_0 = 0, -\frac{1}{4}, -\frac{1}{2}$ for $d = 1, 2, 3$ respectively, it suffices to show that it holds for all $\rho \in [-1, 0]$. By Proposition 2.8.6 it suffices to show that $\mathcal{T} := (\mathcal{A}_0^{\frac{1}{4}} + \mathcal{M}_r)^{-1}\mathcal{A}_0^{\frac{1}{4}} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^{-1})$. This is equivalent to showing that $\mathcal{T} = (I + \mathcal{A}_0^{-\frac{1}{4}}\mathcal{M}_r)^{-1}$ and $\mathcal{A}_0^{-1}\mathcal{T}\mathcal{A}_0 = (I + \mathcal{A}_0^{-\frac{5}{4}}\mathcal{M}_r\mathcal{A}_0)^{-1}$ are bounded in $\mathcal{X}$, which holds by Lemma 2.8.7.

5. By Proposition 2.8.6 it suffices to show that $\mathcal{T} := (\mathcal{A}_0^{\frac{1}{4}} + \mathcal{M}_r)\mathcal{A}_0^{-\frac{1}{4}} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^{-1})$. Indeed, $\mathcal{T} = I + \mathcal{M}_r\mathcal{A}_0^{-\frac{1}{4}} \in \mathcal{L}(\mathcal{X})$. On the other hand, to show $\mathcal{T} \in \mathcal{L}(X^{-1})$ it is equivalent to show that $\mathcal{A}_0^{-1}\mathcal{T}\mathcal{A}_0 \in \mathcal{L}(X)$. Indeed, $\mathcal{A}_0^{-1}\mathcal{T}\mathcal{A}_0 = I + \mathcal{A}_0^{-1}\mathcal{M}_r\mathcal{A}_0^{\frac{3}{4}}$ which is bounded by Lemma 2.8.7.

6. By Proposition 2.8.6 it suffices to show that $\mathcal{T} := (\mathcal{A}_0^{\frac{1}{4}}+\mathcal{M}_r)(\mathcal{A}_0+\mathcal{M}_q)^{-1}\mathcal{A}_0^{\frac{3}{4}} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^1)$. Indeed, we have already shown in part (3) of the current proof that $\mathcal{T} \in \mathcal{L}(\mathcal{X})$. To show $\mathcal{T} \in \mathcal{L}(X^1)$ it is equivalent to show that

$\mathcal{A}_0\mathcal{T}\mathcal{A}_0^{-1} \in \mathcal{L}(\mathcal{X})$. Indeed, $\mathcal{A}_0\mathcal{T}\mathcal{A}_0^{-1} = (I + \mathcal{A}_0\mathcal{M}_r\mathcal{A}_0^{-\frac{5}{4}})(I + \mathcal{A}_0^{\frac{1}{4}}\mathcal{M}_q\mathcal{A}_0^{-\frac{5}{4}})^{-1}$ which by Lemma 2.8.7 is the composition of two bounded operators in $\mathcal{X}$..

7. By Proposition 2.8.6 it suffices to show that $\mathcal{T} := (\mathcal{A}_0+\mathcal{M}_q)^{-1}(\mathcal{A}_0^{\frac{1}{4}}+\mathcal{M}_r)^2\mathcal{A}_0^{\frac{1}{2}} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^{-1}) \cap \mathcal{L}(X^1)$. We start by showing $\mathcal{T} \in \mathcal{L}(\mathcal{X})$. Indeed, we have $\mathcal{T} = (I+\mathcal{A}_0^{-1}\mathcal{M}_q)^{-1}(I+\mathcal{A}_0^{-1}\mathcal{M}_r\mathcal{A}_0^{\frac{3}{4}})(I+\mathcal{A}_0^{-\frac{3}{4}}\mathcal{M}_r\mathcal{A}_0^{\frac{1}{2}})$, which by Lemma 2.8.7, is the composition of three bounded operators. For showing $\mathcal{T} \in \mathcal{L}(X^{-1})$ it is equivalent to show that $\mathcal{A}_0^{-1}\mathcal{T}\mathcal{A}_0 \in \mathcal{L}(\mathcal{X})$. Indeed, $\mathcal{A}_0^{-1}\mathcal{T}\mathcal{A}_0 = (I + \mathcal{A}_0^{-2}\mathcal{M}_q\mathcal{A}_0)^{-1}(I+\mathcal{A}_0^{-2}\mathcal{M}_r\mathcal{A}_0^{\frac{7}{4}})(I+\mathcal{A}_0^{-\frac{7}{4}}\mathcal{M}_r\mathcal{A}_0^{\frac{3}{2}})$, which by Lemma 2.8.7, is the composition of three bounded operators. Finally, we show that $\mathcal{T} \in \mathcal{L}(X^1)$ or equivalently $\mathcal{A}_0\mathcal{T}\mathcal{A}_0^{-1} \in \mathcal{L}(\mathcal{X})$. Indeed we have $\mathcal{A}_0\mathcal{T}\mathcal{A}_0^{-1} = (I+\mathcal{M}_q\mathcal{A}_0^{-1})^{-1}(I+\mathcal{M}_r\mathcal{A}_0^{-\frac{1}{4}})(I + \mathcal{A}_0^{\frac{1}{4}}\mathcal{M}_r\mathcal{A}_0^{-\frac{1}{2}})$, which again by Lemma 2.8.7, is the composition of three bounded operators.

□

We can now apply Theorem 2.2.2 and Theorem 2.2.3 to get the following convergence result.

Theorem 2.8.4. *Let $u^\dagger \in X^\gamma, \gamma \geq 1$. Then, for $\tau = \tau(n) = n^{\frac{4-d-(4\gamma\wedge 11)-\varepsilon}{(8\gamma\wedge 22)+6+2d+2\varepsilon}}$, the convergence in (2.1.12) holds with $\epsilon_n = n^{-e}$, where*

$$e = \begin{cases} \frac{2\gamma}{3+d+4\gamma+2\varepsilon}, & if\ \gamma < \frac{11}{4} \\ \frac{11}{28+2d+2\varepsilon}, & if\ \gamma \geq \frac{11}{4}, \end{cases}$$

for $\varepsilon > 0$ arbitrarily small and where $d = 1, 2, 3$, is the dimension. Furthermore, for $t \in [-\frac{3}{4}, 1)$, for the same choice of τ, we have $\mathbb{E}\left\|m^\dagger - u^\dagger\right\|_t^2 \leq cn^{-h}$, where

$$h = \begin{cases} \frac{4\gamma-4t}{3+d+4\gamma+2\varepsilon}, & if\ \gamma < \frac{11}{4} \\ \frac{22-8t}{28+2d+2\varepsilon}, & if\ \gamma \geq \frac{11}{4}. \end{cases}$$

For $t = 1$ the above rate holds provided $\gamma > 1$.

2.8.3 Example 3 - More general lower order perturbations case

We now study the Bayesian inversion of the operator $\mathcal{A} = \mathcal{A}_0^{\ell\alpha}+\mathcal{M}_q$ for $\ell > 0, \alpha > \frac{d}{2}$ and where $\mathcal{M}_q : L^2(\Omega) \to L^2(\Omega)$ is the multiplication operator by a nonnegative function $q \in W^{a_q,\infty}(\Omega)$ where $a_q > 0$ is sufficiently large. We assume that the observational noise is Gaussian with covariance operator $\mathcal{C}_1 := (\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)^{-2}$, where $\beta > 0$ and $\mathcal{M}_r : L^2(\Omega) \to L^2(\Omega)$ is the multiplication operator by another

nonnegative function $r \in W^{a_r,\infty}(\Omega)$ and a_r is sufficiently large. We set the prior covariance operator to be $\mathcal{C}_0 = \mathcal{A}_0^{-\alpha}$, hence since $\alpha > \frac{d}{2}$ we have that $\mathcal{C}_0$ is trace class and Assumption 2.3.1(1) is satisfied with $s_0 = \frac{d}{2\alpha}$. We work in the Hilbert scale $(X^s)_{s\in\mathbb{R}}$, induced by the operator $\mathcal{A}_0^{-\frac{\alpha}{2}}$.

We show that $\mathcal{C}_1 \simeq \mathcal{C}_0^{\beta}$ and $\mathcal{A}^{-1} \simeq \mathcal{C}_0^{\ell}$, in the sense of the Assumptions 2.3.1. First note that the Assumption 2.3.1(2) is satisfied provided $\Delta = 1 + 2\ell - \beta > \frac{d}{\alpha}$. The rest of the assumptions have the form

3. $\big\|(\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)(\mathcal{A}_0^{\ell\alpha} + \mathcal{M}_q)^{-1}u\big\| \asymp \big\|\mathcal{A}_0^{\alpha\frac{\beta-2\ell}{2}}u\big\|,\ \forall u \in X^{\beta-2\ell}$;
4. $\big\|\mathcal{A}_0^{\frac{\rho\alpha}{2}}(\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)^{-1}u\big\| \le c_1\big\|\mathcal{A}_0^{\alpha\frac{\rho-\beta}{2}}u\big\|,\ \forall u \in X^{\rho-\beta}, \forall \rho \in [\lceil \beta - \frac{d}{2\alpha} - 1\rceil, \beta - \frac{d}{2\alpha})$;
5. $\big\|\mathcal{A}_0^{-\frac{s\alpha}{2}}(\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)u\big\| \le c_2\big\|\mathcal{A}_0^{\alpha\frac{\beta-s}{2}}\big\|,\ \forall u \in X^{\beta-s}, \forall s \in (\frac{d}{2\alpha}, 1]$;
6. $\big\|\mathcal{A}_0^{\frac{s\alpha}{2}}(\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)(\mathcal{A}_0^{\ell\alpha} + \mathcal{M}_q)^{-1}u\big\| \le c_3\big\|\mathcal{A}_0^{\alpha\frac{s+\beta-2\ell}{2}}u\big\|,\ \forall u \in X^{s+\beta-2\ell},\ \forall s \in (\frac{d}{2\alpha}, 1]$;
7. $\big\|\mathcal{A}_0^{-\frac{\eta\alpha}{2}}(\mathcal{A}_0^{\ell\alpha} + \mathcal{M}_q)^{-1}(\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)^2u\big\| \le c_4\big\|\mathcal{A}_0^{\alpha\frac{2\beta-2\ell-\eta}{2}}u\big\|,\ \forall u \in X^{\beta-\ell-\eta},\ \forall \eta \in [\beta - 2\ell, 1]$.

Proposition 2.8.5. *The Assumptions 2.3.1 are satisfied in this example for $a_q = a_q(\alpha, \beta, \ell)$ and $a_r = a_r(\alpha, \beta, \ell)$ sufficiently large, provided $1 + 2\ell - \beta > \frac{d}{\alpha}$.*

Proof. We have already discussed the first two assumptions.

3. We need to show that $\mathcal{T} := (\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)(\mathcal{A}_0^{\ell\alpha} + \mathcal{M}_q)^{-1}\mathcal{A}_0^{\alpha\frac{2\ell-\beta}{2}}$ and $\mathcal{T}^{-1}$ are bounded operators in $\mathcal{X}$. Indeed, $\mathcal{T} = (I + \mathcal{M}_r\mathcal{A}_0^{-\frac{\beta\alpha}{2}})(I + \mathcal{A}_0^{\alpha\frac{\beta-2\ell}{2}}\mathcal{M}_q\mathcal{A}_0^{-\frac{\beta\alpha}{2}})^{-1}$, which is bounded by Lemma 2.8.7 applied for $t = s = \frac{\beta\alpha}{2}$ and $t = \frac{\beta\alpha}{2}, s = \ell\alpha$. For $\mathcal{T}^{-1}$ we have, $\mathcal{T}^{-1} = (I + \mathcal{A}_0^{\alpha\frac{\beta-2\ell}{2}}\mathcal{M}_q\mathcal{A}_0^{-\frac{\beta\alpha}{2}})(I + \mathcal{M}_r\mathcal{A}_0^{-\frac{\beta\alpha}{2}})^{-1}$, which again by Lemma 2.8.7 is the composition of two bounded operators.

4. Depending on the value of $\beta - \frac{d}{2\alpha}$ we determine integer σ such that by interpolation theory it suffices to show that $\mathcal{T} := (\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)^{-1}\mathcal{A}_0^{\frac{\beta\alpha}{2}} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^{\sigma})$. This is equivalent to showing that $\mathcal{T} = (I + \mathcal{A}_0^{-\frac{\beta\alpha}{2}}\mathcal{M}_r)^{-1}$ and $\mathcal{A}_0^{\frac{\sigma\alpha}{2}}\mathcal{T}\mathcal{A}_0^{-\frac{\sigma\alpha}{2}} = (I + \mathcal{A}_0^{\alpha\frac{\sigma-\beta}{2}}\mathcal{M}_r\mathcal{A}_0^{-\frac{\sigma\alpha}{2}})^{-1}$ are bounded in $\mathcal{X}$, which holds by Lemma 2.8.7.

5. By Proposition 2.8.6, it suffices to show that $\mathcal{T} := (\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)\mathcal{A}_0^{-\frac{\beta\alpha}{2}} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^{-1})$. Indeed, $\mathcal{T} = I + \mathcal{M}_r\mathcal{A}_0^{-\frac{\beta\alpha}{2}} \in \mathcal{L}(\mathcal{X})$. On the other hand, to show $\mathcal{T} \in \mathcal{L}(X^{-1})$ it is equivalent to show that $\mathcal{A}_0^{-\frac{\alpha}{2}}\mathcal{T}\mathcal{A}_0^{\frac{\alpha}{2}} \in \mathcal{L}(X)$. Indeed, $\mathcal{A}_0^{-\frac{\alpha}{2}}\mathcal{T}\mathcal{A}_0^{\frac{\alpha}{2}} = I + \mathcal{A}_0^{-\frac{\alpha}{2}}\mathcal{M}_r\mathcal{A}_0^{\frac{\alpha}{2}}\mathcal{A}_0^{-\frac{\beta\alpha}{2}}$, which is bounded by Lemma 2.8.7.

6. By Proposition 2.8.6, it is sufficient to show that $\mathcal{T} := (\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)(\mathcal{A}_0^{\ell\alpha} + \mathcal{M}_q)^{-1}\mathcal{A}_0^{\alpha\frac{2\ell-\beta}{2}} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^1)$. Indeed, we have already shown in part (3) of the current proof that $\mathcal{T} \in \mathcal{L}(X)$. To show $\mathcal{T} \in \mathcal{L}(X^1)$ it is equivalent to show that $\mathcal{A}_0^{\frac{\alpha}{2}}\mathcal{T}\mathcal{A}_0^{-\frac{\alpha}{2}} \in \mathcal{L}(\mathcal{X})$. Indeed, $\mathcal{A}_0^{\frac{\alpha}{2}}\mathcal{T}\mathcal{A}_0^{-\frac{\alpha}{2}} = (I + \mathcal{A}_0^{\frac{\alpha}{2}}\mathcal{M}_r\mathcal{A}_0^{-\alpha\frac{1+\beta}{2}})(I + \mathcal{A}_0^{\alpha\frac{\beta+1-2\ell}{2}}\mathcal{M}_q\mathcal{A}_0^{-\alpha\frac{1+\beta}{2}})^{-1}$ which by Lemma 2.8.7 is the composition of two bounded operators in $\mathcal{X}$.

7. By Proposition 2.8.6, it suffices to show that $\mathcal{T} := (\mathcal{A}_0^{\ell\alpha} + \mathcal{M}_q)^{-1}(\mathcal{A}_0^{\frac{\beta\alpha}{2}} + \mathcal{M}_r)^2\mathcal{A}_0^{\alpha(\ell-\beta)} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^\sigma) \cap \mathcal{L}(X^{-1})$ where $\sigma = \lceil 2\ell - \beta \rceil \geq -1$. Note that σ is either equal to zero or a positive integer, by the assumption $\Delta > 0$. If $\sigma = 0$, that is, if $\beta - 2\ell \geq 0$, then we only need to show $\mathcal{T} \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(X^{-1})$. We start by showing $\mathcal{T} \in \mathcal{L}(\mathcal{X})$. Indeed, we have $\mathcal{T} = (I + \mathcal{A}_0^{-\ell\alpha}\mathcal{M}_q)^{-1}(I + \mathcal{A}_0^{-\ell\alpha}\mathcal{M}_r\mathcal{A}_0^{\alpha\frac{2\ell-\beta}{2}})(I + \mathcal{A}_0^{\alpha\frac{\beta-2\ell}{2}}\mathcal{M}_r\mathcal{A}_0^{\alpha(\ell-\beta)})$, which by Lemma 2.8.7 is the composition of three bounded operators. For showing $\mathcal{T} \in \mathcal{L}(X^{-1})$ it is equivalent to show that $\mathcal{A}_0^{-\frac{\alpha}{2}}\mathcal{T}\mathcal{A}_0^{\frac{\alpha}{2}} \in \mathcal{L}(\mathcal{X})$. Indeed, $\mathcal{A}_0^{-\frac{\alpha}{2}}\mathcal{T}\mathcal{A}_0^{\frac{\alpha}{2}} = (I + \mathcal{A}_0^{-\alpha\frac{1+2\ell}{2}}\mathcal{M}_q\mathcal{A}_0^{\frac{\alpha}{2}})^{-1}(I+\mathcal{A}_0^{-\alpha\frac{1+2\ell}{2}}\mathcal{M}_r\mathcal{A}_0^{\alpha\frac{1+2\ell-\beta}{2}})(I+\mathcal{A}_0^{\alpha\frac{\beta-1-2\ell}{2}}\mathcal{M}_r\mathcal{A}_0^{\alpha\frac{1+2\ell-2\beta}{2}})$, which again by Lemma 2.8.7 is the composition of three bounded operators. Finally, if $\sigma > 0$, we need to show that $\mathcal{T} \in \mathcal{L}(X^\sigma)$ or equivalently $\mathcal{A}_0^{\frac{\sigma\alpha}{2}}\mathcal{T}\mathcal{A}_0^{-\frac{\sigma\alpha}{2}} \in \mathcal{L}(\mathcal{X})$. Indeed, $\mathcal{A}_0^{\frac{\sigma\alpha}{2}}\mathcal{T}\mathcal{A}_0^{-\frac{\sigma\alpha}{2}} = (I+\mathcal{A}_0^{-\alpha\frac{2\ell-\sigma}{2}}\mathcal{M}_q\mathcal{A}_0^{-\frac{\sigma\alpha}{2}})^{-1}(I+\mathcal{A}_0^{-\alpha\frac{2\ell-\sigma}{2}}\mathcal{M}_r\mathcal{A}_0^{\alpha\frac{2\ell-\beta-\sigma}{2}})(I+\mathcal{A}_0^{\alpha\frac{\beta+\sigma-2\ell}{2}}\mathcal{M}_r\mathcal{A}_0^{\alpha\frac{2\ell-2\beta-\sigma}{2}})$, which once more by Lemma 2.8.7 is the composition of three bounded operators.

Observe that the application of Lemma 2.8.7 imposes conditions on the values of a_q and a_r which we do not make explicit in this general case. It is straightforward to determine these conditions once the values of α, β and ℓ are known. □

Note that we require $\ell, \beta > 0$ for our compactness arguments to work, however, the cases $\beta = 0$ and/or $\ell = 0$ also work using a slightly modified proof.

As in the previous examples, one can apply Theorem 2.2.2 and Theorem 2.2.3 to get the corresponding convergence results. We do this only for Theorem 2.2.3 and for dimension $d = 1$. In this case we have that the spaces X^t can be identified with $\mathcal{H}^{t\alpha}$, where the spaces $\mathcal{H}^t$ are defined in Subsection 2.3.2 and are often termed as Sobolev-classes in the statistical literature, [14]. Using Theorem 2.2.3 we get the following result which holds for $\mathcal{C}_0 = \mathcal{A}_0^{-\alpha}, \mathcal{A} = \mathcal{A}_0^{\hat{\ell}} + \mathcal{M}_q, \mathcal{C}_1 = (\mathcal{A}_0^{\frac{\hat{\beta}}{2}} + \mathcal{M}_r)^{-2}$, where $\alpha > \frac{1}{2}$ and $\hat{\ell}, \hat{\beta} \geq 0$ such that $\alpha + 2\hat{\ell} - \hat{\beta} > 1$, provided the nonnegative functions q and r are sufficiently regular:

Theorem 2.8.1. *Assume that $u^\dagger \in \mathcal{H}^{\hat{\gamma}}$, where $\hat{\gamma} \geq \alpha$. Then for*

$$\tau = \tau(n) = n^{-\frac{2\hat{\gamma}\wedge(2\alpha+2\hat{\ell}-\hat{\beta})-2\alpha+1+\varepsilon}{8\hat{\ell}-4\hat{\beta}+4(\hat{\gamma}\wedge(2\alpha+2\hat{\ell}-\hat{\beta}))+2+2\varepsilon}},$$

where $\varepsilon > 0$ arbitrarily small, we have that the convergence in (2.1.12) holds with $\epsilon_n = n^{-e}$, where

$$e = \begin{cases} \frac{\hat{\gamma}\wedge(2\alpha+2\hat{\ell}-\hat{\beta})}{4\hat{\ell}-2\hat{\beta}+2(\hat{\gamma}\wedge(2\alpha+2\hat{\ell}-\hat{\beta}))+1+\varepsilon}, & if\ \hat{\beta}-2\hat{\ell} \leq 0 \\ \frac{2\hat{\ell}-\hat{\beta}+\hat{\gamma}\wedge(2\alpha+2\hat{\ell}-\hat{\beta})}{4\hat{\ell}-2\hat{\beta}+2(\hat{\gamma}\wedge(2\alpha+2\hat{\ell}-\hat{\beta}))+1+\varepsilon}, & otherwise. \end{cases}$$

Remark 2.8.2. *In the simultaneously diagonalizable case where the functions q and r are the zero functions (see Example 2.3.2), we have that $\hat{\alpha} = \alpha - \frac{1}{2}$ and $\hat{\beta} - \frac{1}{2}$ are the Sobolev regularities of the prior and noise respectively, and $p = 2\hat{\ell} - \hat{\beta}$ is the degree of ill-posedness in the sense of [14]. One can then use the last theorem to get the contraction rate for a truth $u^\dagger \in \mathcal{H}^{\hat{\gamma}}$: $\epsilon_n = n^{-e}$ where*

$$e = \begin{cases} \frac{\hat{\gamma}\wedge(1+2\hat{\alpha}+p)}{1+2(\hat{\gamma}\wedge(1+2\hat{\alpha}+p))+2p+\varepsilon}, & if\ p \geq 0 \\ \frac{\hat{\gamma}\wedge(1+2\hat{\alpha}+p)+p}{1+2(\hat{\gamma}\wedge(1+2\hat{\alpha}+p))+2p+\varepsilon}, & otherwise. \end{cases}$$

In this form, our rates can be directly compared to the minimax rates contained in [14] as well as the rates obtained in [44]. We elaborate further on this comparison in Section 2.9.

2.8.4 Technical results from interpolation theory

Let $(Y^s)_{s\in\mathbb{R}}$ be the Hilbert scale induced by a self-adjoint positive definite linear operator $\mathcal{Q} \in \mathcal{L}(\mathcal{X})$ (cf. Section 2.3), where $\mathcal{X}$ is a separable Hilbert space. The following result holds [53, Theorems 4.36, 1.18, 1.6]:

Proposition 2.8.6. *For any $t > 0$, the couples $(\mathcal{X}, Y^t)$ and $(\mathcal{X}, Y^{-t})$ are interpolation couples and for every $\theta \in [0,1]$ we have $(\mathcal{X}, Y^t)_{\theta,2} = Y^{\theta t}$ and $(\mathcal{X}, Y^{-t})_{\theta,2} = Y^{-\theta t}$. In particular, for any $s \in \mathbb{R}$, if $T \in \mathcal{L}(\mathcal{X}) \cap \mathcal{L}(Y^s)$ then $T \in \mathcal{L}(Y^{\theta s})$ for any $\theta \in [0,1]$.*

Let $\mathcal{X} = L^2(\Omega)$, where $\Omega \subset \mathbb{R}^d$ bounded open. Let $w \in W^{a_w,\infty}(\Omega)$ be a nonnegative function and define the multiplication operator $\mathcal{M}_w \colon \mathcal{X} \to \mathcal{X}$. Note

that by the Hölder inequality the operator $\mathcal{M}_w$ is bounded. The last proposition, implies the following lemma.

Lemma 2.8.7. *For any $t \in \mathbb{R}$, $\mathcal{A}_0^t \mathcal{M}_w \mathcal{A}_0^{-t}$ is a bounded operator in $\mathcal{X}$, provided $a_w \geq 2\lceil |t| \rceil$. Furthermore, for any $s > 0$ the operators $K_1 := \mathcal{A}_0^t \mathcal{M}_w \mathcal{A}_0^{-t-s}$ and $K_2 := \mathcal{A}_0^{t-s} \mathcal{M}_w \mathcal{A}_0^{-t}$ are compact in $\mathcal{X}$ and $(I + K_i)^{-1}, i = 1, 2$, are bounded in $\mathcal{X}$.*

Proof. We begin by showing that $\mathcal{A}_0^t \mathcal{M}_w \mathcal{A}_0^{-t} \in \mathcal{L}(\mathcal{X})$, for $t \in [-1, 1]$. By the last proposition applied for $\mathcal{Q} = \mathcal{A}_0^{-2}$, $T = \mathcal{M}_w$, and since $\mathcal{M}_w$ is bounded, it suffices to show that $\mathcal{A}_0^{-1} \mathcal{M}_w \mathcal{A}_0$ and $\mathcal{A}_0 \mathcal{M}_w \mathcal{A}_0^{-1}$ are bounded in $\mathcal{X}$. In fact it suffices to show that $\mathcal{A}_0 \mathcal{M}_w \mathcal{A}_0^{-1}$ is bounded since $\|\mathcal{A}_0^{-1} \mathcal{M}_w \mathcal{A}_0\| = \|(\mathcal{A}_0^{-1} \mathcal{M}_w \mathcal{A}_0)^*\| = \|\mathcal{A}_0 \mathcal{M}_w \mathcal{A}_0^{-1}\|$. Indeed, since $\mathcal{A}_0 = -\boldsymbol{\Delta}$,

$$\|\mathcal{A}_0 \mathcal{M}_w \mathcal{A}_0^{-1} \phi\| = \|\boldsymbol{\Delta} \mathcal{M}_w \mathcal{A}_0^{-1} \phi\| = \|(\boldsymbol{\Delta} w) \mathcal{A}_0^{-1} \phi + 2(\nabla w) \cdot (\nabla \mathcal{A}_0^{-1} \phi) + w \boldsymbol{\Delta} \mathcal{A}_0^{-1} \phi\|$$

$$\leq \|w\|_{W^{2,\infty}(\Omega)} (\|\mathcal{A}_0^{-1} \phi\| + \|\nabla \mathcal{A}_0^{-1} \phi\| + \|\phi\|) \leq c \|w\|_{W^{2,\infty}(\Omega)} \|\phi\|.$$

For general $t \in \mathbb{R}$, let $\underline{t} = \lceil |t| \rceil \in \mathbb{N}$, then as before it suffices to show that $\mathcal{A}_0^{\underline{t}} \mathcal{M}_w \mathcal{A}_0^{-\underline{t}}$ is bounded in $\mathcal{X}$. Again, using the fact that $\mathcal{A}_0 = -\boldsymbol{\Delta}$, we have by the product rule for derivatives that $\mathcal{A}_0^{\underline{t}} \mathcal{M}_w \mathcal{A}_0^{-\underline{t}}$ is bounded, provided $w \in W^{2\underline{t},\infty}(\Omega)$.

The operators K_i are compact in $\mathcal{X}$, since they are compositions between the compact operator $\mathcal{A}_0^{-s}$ and the bounded operator $\mathcal{A}_0^t \mathcal{M}_w \mathcal{A}_0^{-t}$. Positivity of the operator $\mathcal{A}_0$ and nonnegativity of the operator $\mathcal{M}_w$ show that -1 cannot be an eigenvalue of K_i, so that by the Fredholm Alternative [33, §27, Theorem 7] we have that $(I + K_i)^{-1}, i = 1, 2$, are bounded in $\mathcal{X}$. □

2.9 The Diagonal Case

In the case where $\mathcal{C}_0, \mathcal{C}_1$ and $\mathcal{A}$, are all diagonalizable in the same eigenbasis our assumptions are trivially satisfied, provided $\Delta > 2s_0$. In [44], sharp convergence rates are obtained for the convergence in (2.1.12), in the case where the three relevant operators are simultaneously diagonalizable and have spectra that decay algebraically; the authors only consider the case $\mathcal{C}_1 = I$ since in this diagonal setting the colored noise problem can be reduced to the white noise one. The rates in [44] agree with the minimax rates provided either the scaling of the prior is optimally chosen or the scaling is fixed and the prior matches exactly the regularity of the truth, [14]. In Figure 2.1 (cf. Section 2.2) we have in green the optimized rates of convergence predicted by Theorem 2.2.3 (see Remark 2.8.2 which contains the rates obtained by our method in this diagonal mildly ill-posed setting) and in blue the sharp conver-

gence rates from [44], plotted against the regularity of the true solution, $u^\dagger \in X^\gamma$, in the case where $\beta = \ell = \frac{1}{2}$ and $\mathcal{C}_0$ has eigenvalues that decay like j^{-2}. In this case $s_0 = \frac{1}{2}$ and $\Delta = \frac{3}{2}$, so that $\Delta > 2s_0$.

As explained in Remark 2.7.3, the minimum regularity for our method to work is $\gamma = 1$ and our rates saturate at $\gamma = 1 + \Delta$, that is, in this example at $\gamma = 2.5$. We note that for $\gamma \in [1, 2.5]$ our rates agree, up to $\varepsilon > 0$ arbitrarily small, with the sharp rates obtained in [44], for $\gamma > 2.5$ our rates are suboptimal and for $\gamma < 1$ the method fails. In [44], the convergence rates are obtained for $\gamma > 0$ and the saturation point is at $\gamma = 2\Delta$, that is, in this example at $\gamma = 3$. In general the PDE method can saturate earlier (if $2\ell - \beta > 0$), at the same time (if $2\ell - \beta = 0$), or later (if $2\ell - \beta < 0$) compared to the diagonal method presented in [44]. However, the case $2\ell - \beta < 0$ in which our method saturates later, is also the case in which our rates are suboptimal, as explained in Remark 2.7.3(iv).

The discrepancies can be explained by the fact that in Proposition 2.6.1, the choice of θ which determines both the minimum requirement on the regularity of $u^\dagger$ and the saturation point, is the same for both of the operator norm bounds. This means that on the one hand to get convergence of the term $\left\|\lambda \mathcal{B}_\lambda^{-1} \mathcal{C}_0^{-1} u^\dagger\right\|$ in equation (2.7.4) in the proof of Theorem 2.7.2, we require conditions which secure the convergence in the stronger X^1-norm and on the other hand the saturation rate for this term is the same as the saturation rate in the weaker $X^{\beta-2\ell}$-norm. For example, when $\beta - 2\ell = 0$ the saturation rate in the PDE method is the rate of the $\mathcal{X}$-norm hence we have the same saturation point as the rates in [44]. In particular, we have agreement of the saturation rate when $\beta = \ell = 0$, which corresponds to the problem where we directly observe the unknown function polluted by white noise (termed the *white noise model*).

Finally, we include in Appendix A another diagonal example which corresponds to the severely ill-posed case with Gaussian priors of analytic regularity.

2.10 Conclusions

We have presented a new method of identifying the posterior distribution in a conjugate Gaussian Bayesian linear inverse problem setting (Section 2.2 and Section 2.5). We used this identification to examine the posterior consistency of the Bayesian approach in a frequentist sense (Section 2.2 and Section 2.7). We provided convergence rates for the convergence of the expectation of the mean error in a range of norms (Theorem 2.7.1, Theorem 2.2.2). We also provided posterior contraction rates (Theorem 2.7.2, Theorem 2.2.3). Our methodology assumed a relation between the

prior covariance, the noise covariance and the forward operator, expressed in the form of norm equivalence relations (Assumptions 2.3.1). We considered Gaussian noise which can be white. In order for our methods to work we required a certain degree of ill-posedness compared to the regularity of the prior (Assumption 2.3.1(2)) and for the convergence rates to be valid a certain degree of regularity of the true solution. In the case where the three involved operators are all diagonalizable in the same eigenbasis, when the problem is mildly ill-posed with a sufficiently large degree of ill-posedness with respect to the prior, and for a range of values of γ, the parameter expressing the regularity of the true solution, our rates agree (up to $\varepsilon > 0$ arbitrarily small) with the sharp (minimax over Sobolev classes) convergence rates obtained in [44] (Section 2.9). Furthermore, again in the simultaneously diagonalizable setting, if the prior has analytic regularity and the problem is severely ill-posed, then our rates agree (up to $\varepsilon > 0$ arbitrarily small) with the minimax rates for severely ill-posed problems over analytic classes (see Appendix A).

Our optimized rates rely on rescaling the prior depending on the size of the noise, achieved by choosing the scaling parameter τ^2 in the prior covariance as an appropriate function of the parameter $n^{-\frac{1}{2}}$ multiplying the noise. However, the relationship between τ and n depends on the unknown regularity of the true solution γ, which raises the question how to optimally choose τ in practice. An attempt to address this question in a similar but more restrictive setting than ours is taken in [24], where an empirical Bayes maximum likelihood based procedure giving a data driven selection of τ is presented. Moreover, in [79] both the hierarchical approach for simultaneous inference on τ and the unknown u and an empirical Bayes approach for estimating τ from the data, are considered in the white noise model with Gaussian priors of Sobolev regularity; it is shown that in both cases the resulting posterior contracts at the minimax rate when the truth is in a range of spaces determined by the regularity of the prior. A different approach is taken in [43] in the more general simultaneously diagonalizable mildly ill-posed case with priors of Sobolev regularity. As discussed in [44], for a fixed value of τ independent of n, the rates are optimal only if the regularity of the prior exactly matches the regularity of the truth. In [43], the hierarchical approach for simultaneous inference on the regularity of the prior and the unknown u and an empirical Bayes method for choosing the regularity of the prior from the data are presented; both methods are shown to give minimax rates over both Sobolev and analytic regularity classes up to slowly varying terms.

A potential extension of our method for obtaining posterior contraction rates to empirical and hierarchical Bayesian approaches of the type discussed in the last

paragraph is limited by our requirement that the truth is in the Cameron-Martin space of the prior; more research is needed in order to sidestep this assumption on the regularity of the truth. In Chapter 4 we investigate a different aspect of hierarchical Bayesian approaches to inverse problems and in particular we study the implementation of hierarchical methods with conjugate priors and hyper-priors for inference simultaneously on the unknown u, the scaling parameter of the prior τ and the noise level $n^{-\frac{1}{2}}$.

The methodology presented in this chapter is extended to drift estimation for diffusion processes in [64]. Future research includes the extension to an abstract setting which includes both the present chapter and [64] as special cases. Other possible directions are the consideration of nonlinear inverse problems, the use of non-Gaussian priors and/or noise and the extension of the credibility analysis presented in [44] to a more general setting.

Finally, we remark that our assumption that the forward operator is of the form $\mathcal{A}^{-1}$ where $\mathcal{A}$ is a self-adjoint positive definite operator with bounded inverse can be relaxed to include forward operators K which are bounded and injective. Such a generalization requires some straightforward modifications to the formulae and norm equivalence assumptions; in particular, whenever the forward operator K appears in front of the noise covariance operator $\mathcal{C}_1$, it needs to be replaced by its adjoint operator K^*. The assumption that $\mathcal{A}$ is self-adjoint and positive definite was made because we had in mind examples of the type considered in Section 2.8; however, an inspection of our proofs shows that it was never used.

Chapter 3

Bayesian Posterior Contraction Rates for Linear Severely Ill-Posed Inverse Problems

3.1 Introduction

Let $\mathcal{X}$ be an infinite dimensional separable Hilbert space and let $K : \mathcal{X} \to \mathcal{X}$ be an injective compact linear operator with non-closed range. We consider the ill-posed inverse problem of finding u from data y, where

$$y = Ku + \eta, \tag{3.1.1}$$

and where η represents noise. The problem (3.1.1) is called mildly or modestly ill-posed if the singular values of the forward mapping K decay algebraically, while it is called severely ill-posed if the singular values of K decay exponentially [22]. Our interest is focussed on the severely ill-posed case, and on the small observational noise limit.

The use of classical (deterministic) regularization methods for (3.1.1), and the small-noise limit in particular, is well-studied in both the mildly ill-posed [22] and severely ill-posed [36] cases; nonlinear inverse problems have also been studied from this perspective [22]. However, if we wish to incorporate information concerning the statistical structure of the noise and solution, then it is natural to adopt a Bayesian perspective. The Bayesian approach to linear ill-posed inverse problems was adopted in [25], in which the severely ill-posed problem of inverting the heat kernel was considered, and then developed systematically in [54, 52]. More recently,

nonlinear inverse problems have been given a Bayesian formulation [46, 78, 47, 48]. However, the study of the small noise limit, known as posterior consistency in the Bayesian context, is an under-developed aspect of the Bayesian methodology for inverse problems. Our work adds to the growing literature in this area.

For mildly ill-posed linear problems, subject to Gaussian observational noise, Bayesian posterior consistency is considered in the recent papers [44, 3]*. In [44], sharp contraction rates are obtained for white observational noise when the forward operator K and the prior covariance operator are simultaneously diagonalizable; this allows the analysis to proceed through the study of an infinite set of uncoupled scalar linear inverse problems. In [3]*, the setting of [44] is generalized to allow for non-white noise and operators which are not simultaneously diagonalizable, using tools from PDE theory. The paper [45] is, to the best of the authors' knowledge, the first to study Bayesian posterior consistency for severely ill-posed problems. It concerns the one-dimensional backward heat equation with white noise, where the j-th eigenvalue of the (self-adjoint) forward mapping decays like e^{-j^2} and works in the simultaneously diagonalizable paradigm of [44]. In this chapter, we generalize the work in [45] by studying Bayesian posterior consistency for a class of severely ill-posed inverse problems in which the j-th singular value of K decays as e^{-sj^b} for arbitrary positive s and b, again working in the simultaneously diagonalizable paradigm of [44]. In addition to the backward heat equation considered in [45] ($b = 2$), there are a variety of ill-posed inverse problems covered by our theory. For instance, the Cauchy problem for the Laplace equation and the Cauchy problem for the Helmholtz equation or the modified Helmholtz equation (see [88] and the references therein): the eigenvalue decay of the forward mapping for these three examples corresponds to $b = 1$. Our analysis is inspired by both the problem and techniques used in [45]; however our generalized setting leads to some technical improvements in the proofs, we discuss new results relating to the equivalence of the prior and posterior and we include a numerical illustration for the Helmholtz equation.

The rest of this chapter is organized as follows. In Section 3.2 we introduce notation and give informal calculations for the posterior mean and covariance operator. In Section 3.3 we characterize the posterior distribution rigorously and show that it is equivalent, in the sense of measures, to the prior – see Theorems 3.3.1 and 3.3.2. In Section 3.4 we present the main results concerning posterior consistency, characterizing the error in the mean in Theorem 3.4.1, the contraction of the posterior covariance in Theorem 3.4.2 and putting these together to estimate posterior

*[3] is Chapter 2 in the current thesis.

contraction rates in Theorem 3.4.3. Some technical lemmas which are essential to the proof of Theorems 3.4.1, 3.4.2 and 3.4.3 are attached at the end of this section. In Section 3.5 contains a simple example for which the theoretical analysis can be applied and includes a numerical experiment which is consistent with the theory. Finally, Section 3.6 contains concluding remarks.

3.2 Notation and Problem Setting

3.2.1 Notation

Throughout the chapter, $\langle \cdot, \cdot \rangle$ and $\| \cdot \|$ denote the inner product and norm of the Hilbert space $\mathcal{X}$. For two sequences k_j and h_j of real numbers, $k_j \asymp h_j$ means that $\frac{|k_j|}{|h_j|}$ is bounded from above and below as $j \to \infty$, $k_j \lesssim h_j$ means that $\frac{k_j}{h_j}$ is bounded from above as $j \to \infty$, and $k_j \sim h_j$ means that $\frac{k_j}{h_j} \to 1$ as $j \to \infty$. We will use M to denote a constant which is different from occurrence to occurrence.

Let $\{\phi_j\}_{j=1}^{\infty}$ denote an orthonormal basis in $\mathcal{X}$. Then we can express an element $u \in \mathcal{X}$ as $u = \sum_{j=1}^{\infty} u_j \phi_j$ where $u_j = \langle u, \phi_j \rangle$. For $\gamma \geq 0$ we define the Sobolev-like spaces

$$\mathcal{H}^{\gamma} = \{u \in \mathcal{X} : \sum_{j=1}^{\infty} j^{2\gamma} u_j^2 < \infty\},$$

with norm $\| \cdot \|_{\mathcal{H}^{\gamma}}$ given by

$$\|u\|_{\mathcal{H}^{\gamma}}^2 := \sum_{j=1}^{\infty} j^{2\gamma} u_j^2.$$

For $\gamma < 0$, we define the spaces $\mathcal{H}^{\gamma}$ by duality: $\mathcal{H}^{\gamma} = (\mathcal{H}^{-\gamma})^*$.

In the following we consider random variables drawn from Gaussian distributions in $\mathcal{X}$, denoted by $\mathcal{N}(\theta, \Sigma)$ where the mean θ is an element of $\mathcal{X}$ and the covariance operator Σ is a positive definite, self-adjoint, trace class, linear operator in $\mathcal{X}$. The operator Σ possesses an infinite set of eigenfunctions $\{\phi_j\}_{j \in \mathbb{N}}$ which correspond to positive eigenvalues $\{\sigma_j\}_{j \in \mathbb{N}}$ and which form an orthonormal basis of $\mathcal{X}$. One can express a draw x from $\mathcal{N}(\theta, \Sigma)$ using the Karhunen-Loeve expansion as

$$x = \theta + \sum_{j=1}^{\infty} \sqrt{\sigma_j} \xi_j \phi_j, \tag{3.2.1}$$

where ξ_j are independent and identically distributed $\mathcal{N}(0, 1)$ real random variables, [17, 78]. In particular, the expansion coefficients $x_j = \theta_j + \sqrt{\sigma_j} \xi_j$ are $\mathcal{N}(\theta_j, \sigma_j)$ real random variables and it is straightforward to see that $\mathbb{E}\|x\|^2 = \|\theta\|^2 + \mathrm{Tr}(\Sigma)$ and

that for any bounded linear operator T in $\mathcal{X}$, Tx is distributed as $\mathcal{N}(T\theta, T\Sigma T^*)$. It is also straightforward to check that if $\theta = 0$ and $\sigma_j = j^{-2r}$ for some $r \in \mathbb{R}$, then $x \in \mathcal{H}^\gamma$ almost surely , for any $\gamma < r - \frac{1}{2}$.

3.2.2 Bayesian setting and informal charaterization of the posterior

In this subsection we describe the assumptions underlying the Bayesian formulation of the linear inverse problem. Furthermore we provide informal calculations which motivate the expressions for the posterior mean and covariance. These will be made precise in Section 3.3.

We place a scaled Gaussian prior on the unknown u of the form $\mu_0 := \mathcal{N}(0, \tau^2\mathcal{C}_0)$, where $\tau > 0$ is a scale parameter and $\mathcal{C}_0$ is a self-adjoint, positive-definite, trace class, linear operator on $\mathcal{X}$. We assume Gaussian observational noise in (3.1.1) which is independent of u. In particular, we model the data as

$$y = Ku + \frac{1}{\sqrt{n}}\eta, \tag{3.2.2}$$

that is we have $\eta = \frac{1}{\sqrt{n}}\eta$ in (3.1.1), where $\frac{1}{\sqrt{n}}$ is a scale parameter modelling the noise level and η is a random variable independent of u and distributed as $\mathcal{N}(0, \mathcal{C}_1)$. The linear operator $\mathcal{C}_1$ is assumed to be self-adjoint, positive-definite, bounded, but not necessarily trace class on $\mathcal{X}$. This allows for the possibility of having irregular noise which is not in $\mathcal{X}$. For example, the case where η is white noise corresponds to $\mathcal{C}_1 = I$, and can be viewed as a Gaussian random variable in $\mathcal{H}^{-r}$ for $r > \frac{1}{2}$. Under these assumptions, the conditional distribution of $y|u$, called the *data likelihood*, is the translation of $\mathcal{N}(0, \mathcal{C}_1)$ by Ku, which is also Gaussian:

$$\mathcal{N}(Ku, \frac{1}{n}\mathcal{C}_1). \tag{3.2.3}$$

In finite dimensions the density of the *posterior* distribution, that is the conditional distribution of $u|y$, is found from Bayes rule to be proportional to $\exp(-\mathcal{J}(u; y))$, where

$$\mathcal{J}(u; y) = \frac{n}{2}\|\mathcal{C}_1^{-\frac{1}{2}}(y - Ku)\|^2 + \frac{1}{2\tau^2}\|\mathcal{C}_0^{-\frac{1}{2}}u\|^2. \tag{3.2.4}$$

This suggests that in our infinite dimensional setting, the posterior distribution is Gaussian, $\mu^y := \mathcal{N}(m, \mathcal{C})$, where the mean m and covariance $\mathcal{C}$ can be informally

derived from (3.2.4) using completion of the square:

$$\mathcal{C}^{-1} = nK^*\mathcal{C}_1^{-1}K + \frac{1}{\tau^2}\mathcal{C}_0^{-1}, \tag{3.2.5}$$

and

$$\frac{1}{n}\mathcal{C}^{-1}m = K^*\mathcal{C}_1^{-1}y. \tag{3.2.6}$$

Observe that the posterior mean m is the minimizer of the functional $\mathcal{J}(u;y)$. If we define $\mathcal{J}_0(u;y) = \frac{1}{n}\mathcal{J}(u;y)$ and denote

$$\lambda := \frac{1}{n\tau^2}, \tag{3.2.7}$$

then m also minimizes the functional $\mathcal{J}_0(u;y)$, that is,

$$m = \arg\min_u \mathcal{J}_0(u;y), \tag{3.2.8}$$

where

$$\mathcal{J}_0(u;y) = \frac{1}{2}\|\mathcal{C}_1^{-\frac{1}{2}}(y - Ku)\|^2 + \frac{\lambda}{2}\|\mathcal{C}_0^{-\frac{1}{2}}u\|^2.$$

Thus the posterior mean is a Tikhonov-Phillips regularized solution in the classical sense (in fact $\mathcal{J}_0$ is almost surely infinite and we should really consider $\Psi_0 = \mathcal{J}_0 - \frac{1}{2}\left\|\mathcal{C}_1^{-\frac{1}{2}}y\right\|^2$ which is finite; the minimizer is unaffected). This reveals the close connection between Bayesian and classical regularization for inverse problems. In the deterministic framework, λ is called the *regularization parameter* which is carefully chosen in order to balance consistency and stability. Similarly, for given inverse noise level n, the scale parameter τ introduced in the prior can be judiciously chosen to guarantee a small error between the posterior mean and the true solution, as we will see in Section 3.4.

Posterior consistency refers, in statistical inverse problems, to studying the relationship between the result of the statistical analysis and the truth which underlies the data in either the small noise or large data limits; we concentrate on the small noise limit. We consider the standard Bayesian variant on frequentist posterior consistency [20, 28] for our severely ill-posed inverse problem. To this end we consider observations which are perturbations of the image of a fixed element $u^\dagger \in \mathcal{X}$ by a scaled Gaussian additive noise, that is, we have data $y = y_n^\dagger$ of the form

$$y_n^\dagger = Ku^\dagger + \frac{1}{\sqrt{n}}\eta \tag{3.2.9}$$

where η is a realization of $\mathcal{N}(0, \mathcal{C}_1)$. This choice of data model gives the posterior

distribution as $\mu_{\lambda,n}^{y_n^\dagger} := \mathcal{N}(m^\dagger, \mathcal{C})$, where $\mathcal{C}$ is given by (3.2.5) and $m^\dagger$ is given by (3.2.6) with $y = y_n^\dagger$. Similar to the practice in the deterministic framework, we assume a-priori known regularity of the true solution and identify contraction rates of the posterior $\mu_{\lambda,n}^{y_n^\dagger}$ to a Dirac measure centered on the true solution, as the noise disappears ($n \to \infty$).

3.2.3 Model assumptions

In this subsection we present our assumptions on the operators appearing in our framework, that is, on the forward operator K, the prior covariance operator $\mathcal{C}_0$ and the noise covariance operator $\mathcal{C}_1$.

Assumption 3.2.1. *The operators K, $\mathcal{C}_0$ and $\mathcal{C}_1$ commute with one another, so that K^*K, $\mathcal{C}_0$ and $\mathcal{C}_1$ have the same eigenfunctions $\{\phi_j\}_{j=1}^\infty$. The corresponding eigenvalues $\{l_j^2\}_{j=1}^\infty$, $\{c_{0j}\}_{j=1}^\infty$ and $\{c_{1j}\}_{j=1}^\infty$ of K^*K, $\mathcal{C}_0$ and $\mathcal{C}_1$ are assumed to satisfy*

$$l_j \asymp \exp(-sj^b), \quad c_{0j} = j^{-2\alpha}, \quad c_{1j} = j^{-2\beta}, \tag{3.2.10}$$

for $s > 0, b > 0, \alpha > \frac{1}{2}, \beta \geq 0$. Furthermore, the fixed true solution $u^\dagger$ belongs to $\mathcal{H}^\gamma$ for some $\gamma > 0$.

Remark 3.2.2. *As is well known in finite dimensions, in the current infinite dimensional separable Hilbert-space setting, if K, $\mathcal{C}_0$ and $\mathcal{C}_1$ commute with one another, then K^*K, $\mathcal{C}_0$ and $\mathcal{C}_1$ have the same eigenfunctions $\{\phi_j\}_{j=1}^\infty$ [51, 75].*

Remark 3.2.3. *One can relax the assumptions on the eigenvalues of $\mathcal{C}_0$ and $\mathcal{C}_1$ to $c_{0j} \asymp j^{-2\alpha}$ and $c_{1j} \asymp j^{-2\beta}$ without affecting any of the subsequent results.*

3.3 Characterization of the Posterior

In [54, 52] it is proved in the infinite dimensional setting that the posterior is Gaussian with covariance and mean given by

$$\mathcal{C} = \tau^2\mathcal{C}_0 - \tau^2\mathcal{C}_0K^*(K\mathcal{C}_0K^* + \lambda\mathcal{C}_1)^{-1}K\mathcal{C}_0 \tag{3.3.1}$$

and

$$m = \mathcal{C}_0K^*(K\mathcal{C}_0K^* + \lambda\mathcal{C}_1)^{-1}y, \tag{3.3.2}$$

respectively. In the simultaneously diagonalizable case considered here, these formulae are equivalent to the formulae (3.2.5) and (3.2.6) [78, Example 6.23]. Furthermore, since K, $\mathcal{C}_0$ and $\mathcal{C}_1$ commute with one another, the equations (3.3.1) and

(3.3.2) can be rewritten as

$$\mathcal{C} = \tau^2 \mathcal{C}_0 - \tau^2 \mathcal{Q} K \mathcal{C}_0 \tag{3.3.3}$$

and

$$m = \mathcal{Q} y, \tag{3.3.4}$$

where $\mathcal{Q} : \mathcal{X} \to \mathcal{X}$ is the continuous linear operator

$$\mathcal{Q} = \mathcal{C}_0^{\frac{1}{2}} \left(\mathcal{C}_0^{\frac{1}{2}} K^* K \mathcal{C}_0^{\frac{1}{2}} + \lambda \mathcal{C}_1 \right)^{-1} \mathcal{C}_0^{\frac{1}{2}} K^* = \mathcal{C}_0 K^* (K \mathcal{C}_0 K^* + \lambda \mathcal{C}_1)^{-1}.$$

In the next two theorems we show that the Gaussian posterior distribution μ^y, with covariance and mean given by (3.3.3) and (3.3.4), is a proper conditional Gaussian distribution on $\mathcal{X}$ and is absolutely continuous with respect to the prior.

Theorem 3.3.1. *Suppose Assumption 3.2.1 holds, then: (i) the covariance operator $\mathcal{C}$ of the conditional distribution μ^y given by (3.3.3) is trace class on $\mathcal{X}$; (ii) the mean m of the conditional posterior distribution given by (3.3.4) is an element of $\mathcal{X}$, almost surely with respect to the joint distribution of (u, y). Thus $\mu^y(\mathcal{X}) = 1$ almost surely with respect to the joint distribution of (u, y).*

Proof. The fact that $\mu^y(\mathcal{X}) = 1$ follows from (i) and (ii) is well-known [17]. We thus prove these two points.

(i) Using the basis $\{\phi_j\}$, by equation (3.3.3) we have that the eigenvalues of $\mathcal{C}$ are given by

$$c_j = \tau^2 c_{0j} - \frac{\tau^2 c_{0j}^2 l_j^2}{c_{0j} l_j^2 + \lambda c_{1j}} = \frac{\tau^2 \lambda c_{0j} c_{1j}}{c_{0j} l_j^2 + \lambda c_{1j}} \leq \tau^2 c_{0j}. \tag{3.3.5}$$

Since $\mathcal{C}_0$ is trace class on $\mathcal{X}$, it follows that $\mathcal{C}$ is trace class on $\mathcal{X}$.

(ii) From (3.3.4) we have that,

$$\begin{aligned} \mathbb{E}\|m\|^2 &= \mathbb{E}\|\mathcal{Q} y\|^2 = \mathbb{E}\|\mathcal{Q} K u + \frac{1}{\sqrt{n}} \mathcal{Q} \eta\|^2 \\ &= \mathbb{E}\|\mathcal{Q} K u\|^2 + \frac{1}{n} \mathbb{E}\|\mathcal{Q} \eta\|^2 \end{aligned} \tag{3.3.6}$$

since η and u are independent and η has mean zero. The distribution of $\mathcal{Q}\eta$ is

$\mathcal{N}(0, \mathcal{Q}\mathcal{C}_1\mathcal{Q}^*)$ and it follows, again working in the basis $\{\phi_j\}_{j=1}^{\infty}$, that

$$\begin{aligned}
\mathbb{E}\|m\|^2 &= \mathbb{E}\|\mathcal{Q}Ku\|^2 + \frac{1}{n}\mathrm{Tr}(\mathcal{Q}\mathcal{C}_1\mathcal{Q}^*) \\
&= \sum_{j=1}^{\infty} \frac{\tau^2 c_{0j}^3 l_j^4}{(l_j^2 c_{0j} + \lambda c_{1j})^2} + \frac{1}{n}\sum_{j=1}^{\infty} \frac{c_{0j}^2 l_j^2 c_{1j}}{(l_j^2 c_{0j} + \lambda c_{1j})^2} \\
&\leq \frac{\tau^2}{\lambda^2}\sum_{j=1}^{\infty} c_{0j}^3 c_{1j}^{-2} l_j^4 + \frac{1}{n\lambda^2}\sum_{j=1}^{\infty} c_{0j}^2 c_{1j}^{-1} l_j^2 \\
&\asymp \frac{\tau^2}{\lambda^2}\sum_{j=1}^{\infty} j^{4\beta-6\alpha}\exp(-4sj^b) + \frac{1}{n\lambda^2}\sum_{j=1}^{\infty} j^{2\beta-4\alpha}\exp(-2sj^b) \\
&< \infty.
\end{aligned}$$

Hence $\|m\|$ is almost surely finite, which completes the proof. □

Theorem 3.3.2. *Suppose Assumption 3.2.1 holds, then the posterior measure $\mu^y = \mathcal{N}(m, \mathcal{C})$ with covariance and mean given by (3.3.3) and (3.3.4), respectively, is equivalent to the prior measure $\mu_0 = \mathcal{N}(0, \tau^2\mathcal{C}_0)$, almost surely with respect to the joint distribution of (u, y).*

Proof. By the Feldman-Hajek theorem [18, Theorem 2.23], to show that the Gaussian measure $\mu^y = \mathcal{N}(m, \mathcal{C})$ is equivalent to $\mu_0 = \mathcal{N}(0, \tau^2\mathcal{C}_0)$, it suffices to show:

(i) The Cameron-Martin spaces associated with μ^y and μ_0 are equal, that is, $\mathcal{D}(\mathcal{C}^{-\frac{1}{2}}) = \mathcal{D}(\mathcal{C}_0^{-\frac{1}{2}}) := E$.

(ii) The posterior mean m lies in the Cameron-Martin space E.

(iii) The operator $T = I - \tau^2\mathcal{C}^{-\frac{1}{2}}\mathcal{C}_0\mathcal{C}^{-\frac{1}{2}}$ is Hilbert-Schmidt.

We now check the validity of the above conditions. For (i) it is equivalent to show that there exists a constant M such that

$$\langle h, \mathcal{C}h\rangle \leq M\langle h, \mathcal{C}_0 h\rangle, \forall h \in \mathcal{X} \tag{3.3.7}$$

and

$$\langle h, \mathcal{C}_0 h\rangle \leq M\langle h, \mathcal{C}h\rangle, \forall h \in \mathcal{X}; \tag{3.3.8}$$

this follows from [78, Lemma 6.15] using [18, Proposition B1]. Using the eigenbasis expansion, these are equivalent to

$$c_j \leq M c_{0j} \tag{3.3.9}$$

and

$$c_{0j} \leq M c_j. \tag{3.3.10}$$

From (3.3.5), we know that (3.3.9) is true with $M = \tau^2$. Again by (3.3.5), we have

$$c_j = \frac{\tau^2 c_{0j}}{1 + \lambda^{-1} l_j^2 c_{0j} c_{1j}^{-1}} \asymp \frac{\tau^2 c_{0j}}{1 + \lambda^{-1} \exp(-2sj^b) j^{2\beta - 2\alpha}} \geq M c_{0j}, \tag{3.3.11}$$

where $M = \frac{\tau^2}{1+\kappa}$ and κ is a constant.

For (ii), it is straightforward to check that $E = \mathcal{D}(\mathcal{C}_0^{-\frac{1}{2}}) = \mathcal{H}^\alpha$. The mean square expectation of the posterior mean m in $\mathcal{H}^\alpha$ can be estimated similarly to (3.3.7):

$$\begin{aligned}
\mathbb{E}\|m\|_{\mathcal{H}^\alpha}^2 &= \mathbb{E}\|\mathcal{C}_0^{-\frac{1}{2}} m\|^2 = \mathbb{E}\|\mathcal{C}_0^{-\frac{1}{2}} \mathcal{Q} y\|^2 \\
&= \mathbb{E}\|\mathcal{C}_0^{-\frac{1}{2}} \mathcal{Q} K u + \frac{1}{\sqrt{n}} \mathcal{C}_0^{-\frac{1}{2}} \mathcal{Q} \eta\|^2 \\
&= \mathbb{E}\|\mathcal{C}_0^{-\frac{1}{2}} \mathcal{Q} K u\|^2 + \frac{1}{n} \mathrm{Tr}(\mathcal{C}_0^{-\frac{1}{2}} \mathcal{Q} \mathcal{C}_1 \mathcal{Q}^* \mathcal{C}_0^{-\frac{1}{2}}) \\
&= \sum_{j=1}^{\infty} \frac{\tau^2 c_{0j}^2 l_j^4}{(l_j^2 c_{0j} + \lambda c_{1j})^2} + \lambda \sum_{j=1}^{\infty} \frac{c_{0j} l_j^2 c_{1j}}{(l_j^2 c_{0j} + \lambda c_{1j})^2} \\
&\leq \frac{\tau^2}{\lambda^2} \sum_{j=1}^{\infty} c_{0j}^2 c_{1j}^{-2} l_j^4 + \frac{1}{\lambda} \sum_{j=1}^{\infty} c_{0j} c_{1j}^{-1} l_j^2 \\
&\asymp \frac{\tau^2}{\lambda^2} \sum_{j=1}^{\infty} j^{4\beta - 4\alpha} \exp(-4sj^b) + \frac{1}{\lambda} \sum_{j=1}^{\infty} j^{2\beta - 2\alpha} \exp(-2sj^b) \\
&< \infty,
\end{aligned} \tag{3.3.12}$$

therefore $m \in E$ almost surely.

For (iii), using (3.3.5) again, we have

$$\sum_{j=1}^{\infty} \left(1 - \frac{\tau^2 c_{0j}}{c_j}\right)^2 = \frac{1}{\lambda^2} \sum_{j=1}^{\infty} c_{0j}^2 l_j^4 c_{1j}^{-2} \asymp \sum_{j=1}^{\infty} \exp(-4sj^b) j^{4\beta - 4\alpha} < \infty, \tag{3.3.13}$$

demonstrating that the operator T is Hilbert-Schmidt. □

The preceding result is interesting because, without the assumption that the inverse problem is severely ill-posed, it is possible to construct linear inverse problems of the form considered in this chapter, but for which the posterior is not absolutely continuous with respect to the prior. For example, suppose that we modify Assumption 3.2.1 so that the forward operator K has singular values that

decay algebraically, $l_j \asymp j^{-2\ell}$, but retain the same assumptions on the prior and noise covariances. Then the posterior is again Gaussian with covariance and mean given by the formulae (3.3.1) and (3.3.2). The following proposition shows that, if the noise is too smooth, then the posterior is not absolutely continuous with respect to the prior:

Proposition 3.3.3. *If $\beta \geq \alpha + 2\ell - \frac{1}{4}$ then the posterior $\mu^y = \mathcal{N}(m, \mathcal{C})$ is not absolutely continuous with respect to the prior $\mathcal{N}(0, \tau^2\mathcal{C}_0)$, independently of the data y.*

Proof. Without lost of generality we assume that $\tau = \lambda = n = 1$. It suffices to show that the third condition of the Feldman-Hajek theorem fails [18, Theorem 2.23]. Indeed, $\mathcal{C}$ is diagonalizable in the basis $\{\phi_j\}_{j\in\mathbb{N}}$ with eigenvalues c_j such that

$$c_j \asymp \frac{j^{-2\alpha-2\beta}}{j^{-2\beta} + j^{-4\ell-2\alpha}}.$$

Thus, the operator $T := I - \mathcal{C}^{-\frac{1}{2}}\mathcal{C}_0\mathcal{C}^{-\frac{1}{2}}$ is also diagonalizable in $\{\phi_j\}_{j\in\mathbb{N}}$ with eigenvalues t_j, where

$$t_j = 1 - \frac{c_{0j}}{c_j} \asymp j^{-2\alpha-4\ell+2\beta}.$$

Hence, the operator T is Hilbert-Schmidt, if and only if the sequence $\{t_j\}$ is square summable, that is, if and only if $\beta < \alpha + 2\ell - \frac{1}{4}$. □

3.4 Posterior Contraction

In this section, we study the limiting behaviour of the posterior distribution $\mu^{y_n^\dagger}_{\lambda,n}$ as the noise disappears, $n \to \infty$. Intuitively, we expect the mass of the posterior to concentrate in a small ball centered on the fixed true solution. As in [3, 44, 45, 64]*, we study this problem by identifying ϵ_n such that, for arbitrary positive numbers $M_n \to \infty$, there holds

$$\mathbb{E}^{y_n^\dagger}\mu^{y_n^\dagger}_{\lambda,n}\{u : \|u - u^\dagger\| \geq M_n\epsilon_n\} \to 0. \tag{3.4.1}$$

Here expectation is with respect to the random variable $y_n^\dagger$, with probability distribution given by the data likelihood $\mathcal{N}(Ku^\dagger, \frac{1}{n}\mathcal{C}_1)$, and ϵ_n is called the contraction rate of the posterior distribution with respect to the $\mathcal{X}$-norm.

*[3] is Chapter 2 in the current thesis.

By the Chebyshev inequality, we have

$$\mathbb{E}^{y_n^\dagger}\mu_{\lambda,n}^{y_n^\dagger}\{u : \|u-u^\dagger\| \geq M_n\epsilon_n\} \leq \frac{1}{M_n^2\epsilon_n^2}\mathbb{E}^{y_n^\dagger}\Big(\int \|u-u^\dagger\|^2\mu_{\lambda,n}^{y_n^\dagger}(du)\Big), \tag{3.4.2}$$

thus if

$$\mathbb{E}^{y_n^\dagger}\Big(\int \|u-u^\dagger\|^2\mu_{\lambda,n}^{y_n^\dagger}(du)\Big) \leq M_0\epsilon_n^2, \tag{3.4.3}$$

where M_0 is a constant, we get $\mathbb{E}^{y_n^\dagger}\mu_{\lambda,n}^{y_n^\dagger}\{u : \|u-u^\dagger\| \geq M_n\epsilon_n\} \to 0$ as $M_n \to \infty$. The left hand side of (3.4.3) is the squared posterior contraction (SPC) which satisfies

$$SPC = \mathbb{E}^{y_n^\dagger}\|m^\dagger - u^\dagger\|^2 + \mathrm{Tr}(\mathcal{C}), \tag{3.4.4}$$

and therefore, it is enough to estimate the mean integrated squared error (MISE) of the posterior mean $\mathbb{E}^{y_n^\dagger}\|m^\dagger - u^\dagger\|^2$ and the trace of the posterior covariance operator $\mathcal{C}$.

By (3.3.4) we have

$$m^\dagger = \mathcal{Q}y_n^\dagger = \mathcal{Q}Ku^\dagger + \frac{1}{\sqrt{n}}\mathcal{Q}\eta.$$

Meanwhile,

$$u^\dagger = \mathcal{Q}Ku^\dagger + (I-\mathcal{Q}K)u^\dagger$$

so that we get the error equation

$$e := m^\dagger - u^\dagger = \frac{1}{\sqrt{n}}\mathcal{Q}\eta + (\mathcal{Q}K - I)u^\dagger.$$

The first part of the error comes from the noise, while the second part comes from the regularization. Note that for $\lambda = 0$ formally we have

$$\mathcal{Q}K = \mathcal{C}_0K^*(K^*)^{-1}\mathcal{C}_0^{-1}K^{-1}K = I,$$

indicating that we can make the error small by ensuring that $\lambda \ll 1$ and $n \gg 1$. Since $\lambda = \frac{1}{n\tau^2}$ this indicates the possibility of an optimal choice of $\tau := \tau(n)$ to ensure that $\lambda = \frac{1}{n\tau(n)^2} \to 0$ as $n \to \infty$ and to balance the two sources of error. In the next three theorems, respectively, we estimate the MISE of the posterior mean, the trace of the covariance and the SPC.

Theorem 3.4.1 (MISE)**.** *Under Assumption 3.2.1 the MISE may be estimated as*

follows

$$\text{MISE}\begin{cases} \asymp \frac{1}{n\lambda}(\ln\lambda^{-\frac{1}{2s}})^{-\frac{2\alpha}{b}} + (\ln\lambda^{-\frac{1}{2s}})^{-\frac{2\gamma}{b}}, & b \geq 1, \\ \lesssim \frac{1}{n\lambda}(\ln\lambda^{-\frac{1}{2s}})^{-\frac{2\alpha+b-1}{b}} + (\ln\lambda^{-\frac{1}{2s}})^{-\frac{2\gamma}{b}}, & b < 1. \end{cases} \tag{3.4.5}$$

Proof. From the expression above for the error e, since η is centred Gaussian, we have

$$\mathbb{E}^{y_n^\dagger}\|m^\dagger - u^\dagger\|^2 = \frac{1}{n}\mathbb{E}^{y_n^\dagger}\|\mathcal{Q}\eta\|^2 + \mathbb{E}^{y_n^\dagger}\|(\mathcal{Q}K - I)u^\dagger\|^2, \tag{3.4.6}$$

from which it follows that

$$\begin{aligned} &\mathbb{E}^{y_n^\dagger}\|m^\dagger - u^\dagger\|^2 = \frac{1}{n}\text{Tr}(\mathcal{Q}\mathcal{C}_1\mathcal{Q}^*) + \|(\mathcal{Q}K - I)u^\dagger\|^2 \\ &= \frac{1}{n}\sum_{j=1}^{\infty}\frac{j^{-4\alpha-2\beta}l_j^2}{(j^{-2\alpha}l_j^2 + \lambda j^{-2\beta})^2} + \sum_{j=1}^{\infty}\frac{\lambda^2 j^{-4\beta}(u_j^\dagger)^2}{(j^{-2\alpha}l_j^2 + \lambda j^{-2\beta})^2} \\ &= \frac{1}{n\lambda^2}\sum_{j=1}^{\infty}\frac{l_j^2 j^{2\beta-4\alpha}}{(1 + \frac{1}{\lambda}l_j^2 j^{2\beta-2\alpha})^2} + \sum_{j=1}^{\infty}\frac{(u_j^\dagger)^2}{(1 + \frac{1}{\lambda}l_j^2 j^{2\beta-2\alpha})^2} \\ &:= \text{I} + \text{II}. \end{aligned} \tag{3.4.7}$$

By Assumption 3.2.1, it follows that

$$\text{I} \asymp \frac{1}{n\lambda^2}\sum_{j=1}^{\infty}\frac{\exp(-2sj^b)j^{2\beta-4\alpha}}{(1 + \frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha})^2},$$

and

$$\text{II} \asymp \sum_{j=1}^{\infty}\frac{(u_j^\dagger)^2}{(1 + \frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha})^2}.$$

To estimate I and II we split the sum according to the dominating term in the denominator. Define

$$F(j;\lambda) := \frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha}$$

and note that $F(1;\lambda) > 1$, for λ sufficiently small. Since we are considering a limit in which $\lambda \to 0$ we assume that $F(1;\lambda) > 1$ henceforth. Let J_λ be the unique solution of the equation $F(j;\lambda) = 1$ which exceeds 1. By Lemma 3.4.5, we have

$$J_\lambda \sim (\ln\lambda^{-\frac{1}{2s}})^{\frac{1}{b}}. \tag{3.4.8}$$

For I, if $1 \le j \le J_\lambda$,

$$\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha} \le 1 + \frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha} \le 2\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha}, \quad (3.4.9)$$

therefore

$$\frac{1}{n\lambda^2}\sum_{j\le J_\lambda}\frac{\exp(-2sj^b)j^{2\beta-4\alpha}}{(1+\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha})^2} \asymp \frac{1}{n}\sum_{j\le J_\lambda}\exp(2sj^b)j^{-2\beta}. \quad (3.4.10)$$

The sum on the right hand side is bounded from above by the integral in the same range, and values at both endpoints. By Lemma 3.4.6, we have

$$\begin{aligned}
&\frac{1}{n}\sum_{j\le J_\lambda}\exp(2sj^b)j^{-2\beta} \\
&\le \frac{1}{n}\exp(2sJ_\lambda^b)J_\lambda^{-2\beta} + \frac{1}{n}\exp(2s) + \frac{1}{n}\int_1^{J_\lambda}\exp(2sx^b)x^{-2\beta}dx \\
&= \frac{1}{n}\exp(2sJ_\lambda^b)J_\lambda^{-2\beta} + \frac{1}{n}\exp(2s) + \frac{M}{n}\exp(2sJ_\lambda^b)J_\lambda^{-2\beta-b+1}(1+o(1)) \\
&= \begin{cases} \frac{M}{n}\exp(2sJ_\lambda^b)J_\lambda^{-2\beta}(1+o(1)), & b \ge 1, \\ \frac{M}{n}\exp(2sJ_\lambda^b)J_\lambda^{-2\beta-b+1}(1+o(1)), & b < 1, \end{cases}
\end{aligned} \quad (3.4.11)$$

Since

$$\frac{1}{n}\sum_{j\le J_\lambda}\exp(2sj^b)j^{-2\beta} \ge \frac{1}{n}\exp(2sJ_\lambda^b)J_\lambda^{-2\beta},$$

we deduce that for, $b \ge 1$,

$$\frac{1}{n}\sum_{j\le J_\lambda}\exp(2sj^b)j^{-2\beta} \asymp \frac{1}{n}\exp(2sJ_\lambda^b)J_\lambda^{-2\beta} = \frac{1}{n\lambda}J_\lambda^{-2\alpha}. \quad (3.4.12)$$

For $0 < b < 1$ we have

$$\frac{1}{n}\sum_{j\le J_\lambda}\exp(2sj^b)j^{-2\beta} \lesssim \frac{1}{n}\exp(2sJ_\lambda^b)J_\lambda^{-2\beta-b+1} = \frac{1}{n\lambda}J_\lambda^{-2\alpha-b+1}. \quad (3.4.13)$$

If $j \ge J_\lambda$, then $1 \le 1 + \frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha} \le 2$, thus we have

$$\frac{1}{n\lambda^2}\sum_{j>J_\lambda}\frac{\exp(-2sj^b)j^{2\beta-4\alpha}}{(1+\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha})^2} \asymp \frac{1}{n\lambda^2}\sum_{j>J_\lambda}\exp(-2sj^b)j^{2\beta-4\alpha}.$$

Under our assumption on λ being sufficiently small, we have that J_λ is large enough so that $\exp(-2sj^b)j^{2\beta-4\alpha}$ is always decreasing with respect to j and hence the sum

on the right hand side is bounded from above by the integral in the same range, and the value at the left endpoint. By Lemma 3.4.7, we have

$$\begin{aligned}
&\frac{1}{n\lambda^2}\sum_{j>J_\lambda}\exp(-2sj^b)j^{2\beta-4\alpha}\\
&\leq \frac{1}{n\lambda^2}\exp(-2sJ_\lambda^b)J_\lambda^{2\beta-4\alpha}+\frac{1}{n\lambda^2}\int_{J_\lambda}^{\infty}\exp(-2sx^b)x^{2\beta-4\alpha}dx\\
&\leq \frac{1}{n\lambda^2}\exp(-2sJ_\lambda^b)J_\lambda^{2\beta-4\alpha}+\frac{M}{n\lambda^2}\exp(-2sJ_\lambda^b)J_\lambda^{2\beta-4\alpha-b+1}(1+o(1))\\
&=\begin{cases}\frac{M}{n\lambda^2}\exp(-2sJ_\lambda^b)J_\lambda^{2\beta-4\alpha}(1+o(1)), & b\geq 1,\\ \frac{M}{n\lambda^2}\exp(-2sJ_\lambda^b)J_\lambda^{2\beta-4\alpha-b+1}(1+o(1)), & b<1.\end{cases}
\end{aligned}\tag{3.4.14}$$

Since $\frac{1}{n\lambda^2}\sum_{j>J_\lambda}\exp(-2sj^b)j^{2\beta-4\alpha}\geq\frac{1}{n\lambda^2}\exp(-2sJ_\lambda^b)J_\lambda^{2\beta-4\alpha}$, for $b\geq 1$, we have

$$\frac{1}{n\lambda^2}\sum_{j>J_\lambda}\exp(-2sj^b)j^{2\beta-4\alpha}\asymp\frac{1}{n\lambda^2}\exp(-2sJ_\lambda^b)J_\lambda^{2\beta-4\alpha}=\frac{1}{n\lambda}J_\lambda^{-2\alpha},$$

and for $0<b<1$,

$$\frac{1}{n\lambda^2}\sum_{j>J_\lambda}\exp(-2sj^b)j^{2\beta-4\alpha}\lesssim\frac{1}{n\lambda^2}\exp(-2sJ_\lambda^b)J_\lambda^{2\beta-4\alpha-b+1}=\frac{1}{n\lambda}J_\lambda^{-2\alpha-b+1}.$$

To estimate II, we employ an analysis similar to that applied to I. We have

$$\begin{aligned}
&\sum_{j\leq J_\lambda}\frac{(u_j^\dagger)^2}{(1+\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha})^2}\asymp\sum_{j\leq J_\lambda}(u_j^\dagger)^2\lambda^2\exp(4sj^b)j^{4\alpha-4\beta}\\
&=\sum_{j\leq J_\lambda}j^{2\gamma}(u_j^\dagger)^2\lambda^2\exp(4sj^b)j^{4\alpha-4\beta-2\gamma}.
\end{aligned}\tag{3.4.15}$$

For λ small enough, the terms $\exp(4sj^b)j^{4\alpha-4\beta-2\gamma}$ for $1\leq j\leq J_\lambda$ are dominated by $\exp(4sJ_\lambda^b)J_\lambda^{4\alpha-4\beta-2\gamma}$, so we have the following upper bound for the sum (3.4.15):

$$\sum_{j\leq J_\lambda}j^{2\gamma}(u_j^\dagger)^2\lambda^2\exp(4sj^b)j^{4\alpha-4\beta-2\gamma}\leq\lambda^2\exp(4sJ_\lambda^b)J_\lambda^{4\alpha-4\beta-2\gamma}\|u^\dagger\|_{\mathcal{H}^\gamma}^2.$$

Furthermore

$$\sum_{j\leq J_\lambda}j^{2\gamma}(u_j^\dagger)^2\lambda^2\exp(4sj^b)j^{4\alpha-4\beta-2\gamma}\geq(u_{J_\lambda}^\dagger)^2\lambda^2\exp(4sJ_\lambda^b)J_\lambda^{4\alpha-4\beta-2\gamma},$$

implying that, since $\gamma > 0$ and $u \in \mathcal{H}^\gamma$,

$$\sum_{j \le J_\lambda} j^{2\gamma}(u_j^\dagger)^2 \lambda^2 \exp(4sj^b) j^{4\alpha-4\beta-2\gamma} \asymp \lambda^2 \exp(4sJ_\lambda^b) J_\lambda^{4\alpha-4\beta-2\gamma} = J_\lambda^{-2\gamma}. \quad (3.4.16)$$

The other part of the sum II satisfies

$$\sum_{j > J_\lambda} \frac{(u_j^\dagger)^2}{(1+\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha})^2} \asymp \sum_{j>J_\lambda} (u_j^\dagger)^2 = \sum_{j>J_\lambda} j^{2\gamma}(u_j^\dagger)^2 j^{-2\gamma}.$$

It follows that

$$\sum_{j>J_\lambda} j^{2\gamma}(u_j^\dagger)^2 j^{-2\gamma} \asymp J_\lambda^{-2\gamma}, \quad (3.4.17)$$

since $u \in \mathcal{H}^\gamma$.

Combining (3.4.6) - (3.4.17) completes the proof. □

Theorem 3.4.2 (Trace of $\mathcal{C}$)**.** *Let Assumption 3.2.1 hold and consider the posterior covariance operator $\mathcal{C}$ given by (3.2.5), with λ as in (3.2.7). Then the trace is estimated as*

$$\mathrm{Tr}(\mathcal{C}) \asymp \frac{1}{n\lambda}(\ln \lambda^{-\frac{1}{2s}})^{-\frac{2\alpha-1}{b}}. \quad (3.4.18)$$

Proof. From (3.3.3), we have

$$\mathrm{Tr}(\mathcal{C}) = \sum_{j=1}^{\infty} \frac{\tau^2 \lambda c_{0j} c_{1j}}{c_{0j} l_j^2 + \lambda c_{1j}} \asymp \frac{1}{n\lambda} \sum_{j=1}^{\infty} \frac{j^{-2\alpha}}{1+\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha}}. \quad (3.4.19)$$

As in the proof of Theorem of 3.4.1 we split the sum according to the dominating term in the denominator. For the first part, using equation (3.4.9), we have

$$\frac{1}{n\lambda} \sum_{j \le J_\lambda} \frac{j^{-2\alpha}}{1+\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha}} \asymp \frac{1}{n} \sum_{j \le J_\lambda} \exp(2sj^b) j^{-2\beta}, \quad (3.4.20)$$

where the behaviour of the right hand side is given by (3.4.12) and (3.4.13). The other part of the sum on the right hand side of (3.4.19) satisfies

$$\frac{1}{n\lambda} \sum_{j > J_\lambda} \frac{j^{-2\alpha}}{1+\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha}} \asymp \frac{1}{n\lambda} \sum_{j>J_\lambda} j^{-2\alpha}.$$

By [45, Lemma 6.2], the last sum can be estimated as

$$\sum_{j>J_\lambda} j^{-2\alpha} \asymp J_\lambda^{-2\alpha+1},$$

hence

$$\frac{1}{n\lambda} \sum_{j>J_\lambda} \frac{j^{-2\alpha}}{1+\frac{1}{\lambda}\exp(-2sj^b)j^{2\beta-2\alpha}} \asymp \frac{1}{n\lambda} J_\lambda^{-2\alpha+1}. \tag{3.4.21}$$

Combining (3.4.8), (3.4.19)-(3.4.21) completes the proof. □

We combine the two preceding theorems to determine the overall contraction rate.

Theorem 3.4.3 (Rate of Contraction)**.** *Suppose that Assumption 3.2.1 holds, λ is given by (3.2.7) and $\tau(n) > 0$ satisfies $n\tau^2(n) \to \infty$. Then the posterior distribution $\mu_{\lambda,n}^{y_n^\dagger}$ contracts around the true solution $u^\dagger$ at the rate*

$$\epsilon_n = \big(\ln(n\tau^2)\big)^{-\frac{\gamma}{b}} + \tau\big(\ln(n\tau^2)\big)^{-\frac{\alpha-\frac{1}{2}}{b}}. \tag{3.4.22}$$

In particular, since the rate is undetermined up to a multiplicative constant independent of n, we may take

$$\epsilon_n = \begin{cases} \Big(\ln n\Big)^{-\frac{\gamma\wedge(\alpha-\frac{1}{2})}{b}}, & \tau \equiv 1, \\ \Big(\ln n\Big)^{-\frac{\gamma}{b}}, & n^{-\frac{1}{2}+\sigma} \lesssim \tau \lesssim (\ln n)^{\frac{\alpha-\gamma-\frac{1}{2}}{b}}, \end{cases} \tag{3.4.23}$$

where $\sigma > 0$ is some constant.

Proof. The estimate (3.4.22) follows by combining (3.4.4), Theorem 3.4.1 and Theorem 3.4.2. The rate for $\tau(n) \equiv 1$ follows immediately. In the case of varying $\tau(n)$, observe that in order to balance the contributions of the two terms in (3.4.22), $\tau(n)$ needs to be large enough so that $n\tau^2(n) \to \infty$ as $n \to \infty$, but small enough so that the second term is bounded by the first one. Since the function $\big(\ln(\cdot)\big)^{-\kappa}$, $\kappa > 0$ is decreasing, this can be achieved by choosing $n^{-\frac{1}{2}+\sigma} \lesssim \tau(n) \lesssim (\ln n)^{\frac{\alpha-\gamma-\frac{1}{2}}{b}}$ for some

constant $\sigma > 0$, in which case the rate becomes

$$\begin{aligned}\epsilon_n &\lesssim \big(\ln(n \cdot n^{-1+2\sigma})\big)^{-\frac{\gamma}{b}} + (\ln n)^{\frac{\alpha-\gamma-\frac{1}{2}}{b}} \big(\ln(n \cdot n^{-1+2\sigma})\big)^{-\frac{\alpha-\frac{1}{2}}{b}} \\ &= \big(2\sigma\big)^{-\frac{\gamma}{b}} \big(\ln n\big)^{-\frac{\gamma}{b}} + (2\sigma)^{-\frac{\alpha-\frac{1}{2}}{b}} \big(\ln n\big)^{-\frac{\gamma}{b}} \\ &\lesssim \Big(\ln n\Big)^{-\frac{\gamma}{b}}.\end{aligned}$$

This completes the proof. □

Remark 3.4.4.

i) The rate of the MISE of the posterior mean is determined by the regularity of the prior α, the regularity of the truth γ and the degree of ill-posedness as determined by the power b in the eigenvalues of K (s does not affect the rate). On the other hand, the rate of the trace of the posterior covariance is determined by α and b and has nothing to do with the regularity of the truth γ. Finally the rate of contraction is determined by α, γ and b. Observe that the regularity of the noise, β, does not affect the rate. In the case of mildly ill-posed problems where the singular values of K decay algebraically β does appear in the error estimates, but only through the difference in regularity between the forward operator and the noise covariance [3]. For our severely ill-posed problem this difference may be thought of as being infinite, explaining why β disappears from the error estimates here.*

ii) For fixed $\tau = 1$, the rate of contraction is $\Big(\ln n\Big)^{-\frac{\gamma \wedge (\alpha-\frac{1}{2})}{b}}$, that is, as γ grows the rate improves until $\gamma = \alpha - \frac{1}{2}$, at which point the rate saturates at $\Big(\ln n\Big)^{-\frac{\alpha-\frac{1}{2}}{b}}$. The saturation point $\gamma = \alpha - \frac{1}{2}$ is the crossover point from an undersmoothing prior to an oversmoothing prior. For varying $\tau = \tau(n)$ chosen so that $n^{-\frac{1}{2}+\sigma} \lesssim \tau \lesssim (\ln n)^{\frac{\alpha-\gamma-\frac{1}{2}}{b}}$, the rate is $(\ln n)^{-\frac{\gamma}{b}}$ and never saturates.

iii) For the appropriate choice of $\tau = \tau(n)$, or for fixed $\tau = 1$ and $\alpha \geq \gamma + \frac{1}{2}$, the contraction rate is $\epsilon_n = (\ln n)^{-\frac{\gamma}{b}}$, which for $b = 1$ or 2 is optimal in the minimax sense with L^2-loss [14, 45].

We conclude the section with several technical lemmas used in the proof of the preceding theorems.

Lemma 3.4.5. *Let $a, b > 0$ and $t \in \mathbb{R}$ be constants. For all λ sufficiently small the equation*

$$\frac{1}{\lambda} \exp(-ax^b) x^t = 1, \tag{3.4.24}$$

*[3] is Chapter 2 in the current thesis.

has a unique solution J_λ *in* $\{x \geq 1\}$ *and* $J_\lambda \sim (\ln \lambda^{-\frac{1}{a}})^{\frac{1}{b}}$ *as* $\lambda \to 0$.

Proof. Uniqueness of a root in $\{x \geq 1\}$ follows automatically provided

$$\lambda^{-1} \exp(-a) > 1,$$

since $x \mapsto \exp(-ax^b)x^t$ has at most one maximum in $\{x \geq 0\}$. From (3.4.24), it is straightforward to see that

$$1 = \frac{\ln \lambda^{-\frac{1}{a}}}{J_\lambda^b} + \frac{t}{a}\frac{\ln J_\lambda}{J_\lambda^b}.$$

Since $x \geq 1$, $J_\lambda \to \infty$ as $\lambda \to 0$, thus we have $1 \sim \frac{\ln \lambda^{-\frac{1}{a}}}{J_\lambda^b}$, which completes the proof. □

Lemma 3.4.6. *For* $a > 0, b > 0$ *and* $c \in \mathbb{R}$*, we have as* $J \to \infty$*,*

$$\int_1^J \exp(ax^b)x^c dx \sim \exp(aJ^b)J^{c-b+1}. \tag{3.4.25}$$

Proof. By variable substitution $x^b = y$ and integration by parts, we get

$$\begin{aligned}&\int_1^J \exp(ax^b)x^c dx \\ &= \frac{1}{ab}(\exp(aJ^b)J^{c-b+1} - \exp(a)) - \frac{c-b+1}{ab^2}\int_1^{J^b} \exp(ay)y^{\frac{c-2b+1}{b}} dy,\end{aligned} \tag{3.4.26}$$

thus letting $I(J) := \int_1^{J^b} \exp(ay)y^{\frac{c-2b+1}{b}} dy$, we have that it suffices to show that

$$\lim_{J\to\infty} \frac{I(J)}{\exp(aJ^b)J^{c-b+1}} = 0. \tag{3.4.27}$$

Indeed, if $c - 2b + 1 \geq 0$ then we have

$$\frac{\exp(ay)y^{\frac{c-2b+1}{b}}}{\exp(aJ^b)J^{c-b+1}} \leq \exp(a(y - J^b))J^{-b},$$

and (3.4.27) holds. If $c - 2b + 1 < 0$, we use the variable substitution $e^{ay} = z$ to get that

$$I(J) = \frac{1}{a^{\frac{c-b+1}{b}}}\int_{e^a}^{e^{aJ^b}} (\ln(z))^{\frac{c-2b+1}{b}} dz.$$

By Lemma 3.4.8 below, we then have that

$$I(J) \lesssim \exp(aJ^b)J^{c-2b+1},$$

hence (3.4.27) holds. □

Lemma 3.4.7. *For $J > 0, a > 0, b > 0$ and $c \in \mathbb{R}$ we have*

$$\int_J^\infty \exp(-ax^b)x^c dx \lesssim \exp(-aJ^b)J^{c-b+1}. \tag{3.4.28}$$

Proof. Similar to the proof of Lemma 3.4.6, we have

$$\begin{aligned}&\int_J^\infty \exp(-ax^b)x^c dx\\ &= \frac{1}{ab}\exp(-aJ^b)J^{c-b+1} + \frac{c-b+1}{ab^2}\int_{J^b}^\infty \exp(-ay)y^{\frac{c-2b+1}{b}}dy.\end{aligned}$$

If $\frac{c-b+1}{ab^2} > 0$, then we integrate by parts for n times until $\frac{c-nb+1}{ab^2} < 0$ for the first time. When the constant in front of the integral finally becomes negative we can ignore the integral on the right hand side to get

$$\int_J^\infty \exp(-ax^b)x^c dx \leq \frac{1}{ab}\exp(-aJ^b)(J^{c-b+1}(1+o(1))).$$

□

Lemma 3.4.8. *For any $q, a > 0$ we have as $x \to \infty$*

$$\int_{e^a}^x \frac{dz}{(\ln(z))^q} \leq \frac{x}{(\ln(x))^q}(2+o(1)).$$

Proof. We split the integral as follows

$$\begin{aligned}\int_{e^a}^x \frac{dz}{(\ln(z))^q} &= \int_{e^a}^{e^{2q}} \frac{dz}{(\ln(z))^q} + \int_{e^{2q}}^x \frac{dz}{(\ln(z))^q}\\ &= c(q,a) + \int_{e^{2q}}^x \frac{dz}{(\ln(z))^q},\end{aligned} \tag{3.4.29}$$

where $c(q,a)$ is a real constant. For $z \geq e^{2q}$ it holds

$$\ln(z) \geq 2q,$$

hence dividing by $(\ln(z))^{q+1}$ and rearranging terms we get that

$$\frac{q}{(\ln(z))^{q+1}} \leq \frac{1}{2(\ln(z))^q}. \tag{3.4.30}$$

Integration by parts in the integral on the right hand side of (3.4.29) gives

$$\int_{e^{2q}}^{x} \frac{dz}{(\ln(z))^q} = \frac{x}{(\ln(x))^q} - \frac{e^{2q}}{(2q)^q} + \int_{e^{2q}}^{x} \frac{q}{(\ln(z))^{q+1}} dz,$$

hence using (3.4.30) and rearranging terms, we have

$$\begin{aligned} \int_{e^{2q}}^{x} \frac{dz}{(\ln(z))^q} &\leq 2\frac{x}{(\ln(x))^q} - 2\frac{e^{2q}}{(2q)^q} \\ &= 2\frac{x}{(\ln(x))^q} + \tilde{c}(q). \end{aligned}$$

Concatenating we obtain the result. □

3.5 Example - The Cauchy problem for the Helmholtz equation

In this section, we present the Cauchy problem for the Helmholtz equation as an example to which the theoretical analysis of this chapter can be applied. For simplicity, we only consider the small wave number case ($0 < k < 1$). For more details regarding the more general case, we refer to [88].

Consider the following boundary value problem for the Helmholtz equation:

$$\begin{cases} \Delta v(x_1, x_2) + k^2 v(x_1, x_2) = 0, & (x_1, x_2) \in (0, \pi) \times (0, 1), \\ v_{x_2}(x_1, 0) = 0, & x_1 \in [0, \pi], \\ v(x_1, 1) = u(x_1), & x_1 \in [0, \pi], \\ v(0, x_2) = v(\pi, x_2) = 0, & x_2 \in [0, 1]. \end{cases} \tag{3.5.1}$$

Problem (3.5.1) is well-posed since it corresponds to inversion of a negative-definite elliptic operator with mixed Dirichlet/Neumann data. In fact, by the method of separation of variables, the solution $v(x_1, x_2)$ in the domain $(0, \pi) \times (0, 1)$ can be expressed as

$$v(x_1, x_2) = \sum_{j=1}^{\infty} \frac{\cosh(x_2\sqrt{j^2 - k^2})}{\cosh(\sqrt{j^2 - k^2})} u_j \phi_j(x_1), \tag{3.5.2}$$

where $\phi_j(x_1) = \sqrt{\frac{2}{\pi}} \sin(jx_1)$ and $u_j = \langle u, \phi_j \rangle$.

Define the forward mapping $K : \mathcal{D}(K) \subset L^2(0,\pi) \to L^2(0,\pi)$ by

$$Ku(x_1) = v(x_1, 0) = \sum_{j=1}^{\infty} \frac{1}{\cosh(\sqrt{j^2 - k^2})} u_j \phi_j(x_1),$$

which maps the boundary data of (3.5.1) on $x_2 = 1$ into the solution on $x_2 = 0$. Then K is a self-adjoint, positive-definite, linear operator, with eigenvalues behaving as

$$l_j = \frac{1}{\cosh(\sqrt{j^2 - k^2})} \sim \exp(-j). \tag{3.5.3}$$

The inverse problem is to find the function u, given noisy observations of $v(\cdot, 0)$. More precisely the data y is given by

$$\begin{aligned} y &= v(\cdot, 0) + \frac{1}{\sqrt{n}}\eta, \\ &= Ku + \frac{1}{\sqrt{n}}\eta. \end{aligned}$$

If we place a Gaussian measure $\mathcal{N}(0, \tau^2\mathcal{C}_0)$ as prior on u and assume that η is also Gaussian $\mathcal{N}(0, \mathcal{C}_1)$, then we may apply the theory developed in this chapter. Under Assumption 3.2.1, Theorem 3.4.3 can be applied to this problem with $b = 1$ and $s = 1$ to obtain the contraction rate of the conditional Gaussian posterior distribution.

We now present a numerical simulation for obtaining the rate of the MISE of the posterior mean as the noise disappears ($n \to \infty$), when $\alpha = 2, \gamma = 1$ and we have a fixed $\tau = 1$. In this case, our theory predicts that

$$\text{MISE} \asymp \big(\ln(\sqrt{n})\big)^{-2(\alpha\wedge\gamma)} = \big(\ln(\sqrt{n})\big)^{-2}.$$

To simulate MISE we average the error over a thousand realizations of the noise η, for $n = 10^k$, $k = 1, ..., 100$. We denote the simulated MISE by $\widehat{\text{MISE}}$. The true solution $u^\dagger \in \mathcal{H}^\gamma$ is a fixed draw from a Gaussian measure $\mathcal{N}(0, \Sigma)$, where Σ has eigenvalues $\sigma_j = j^{-2\gamma-1-\varepsilon}$, for $\varepsilon = 10^{-10}$. We use the first 10^5 Fourier modes. In Figure 3.1 we plot $-\frac{1}{2}\ln\big(\widehat{\text{MISE}}\big)$ against $\ln\big(\ln(\sqrt{n})\big)$ in the case $\beta = 0$. The solid line is the relation predicted by Theorem 3.4.1, that is, a line with slope 1. A least squares fit to the simulated points gives a slope of 1.0341 with coefficient of determination 0.9884. In Figure 3.2 we have $\beta = 2$ and all the other parameters the same. The least squares fit gives a slope 0.9723 with coefficient of determination 0.9916, confirming that the regularity of the noise as determined by β does not affect the rate of convergence.

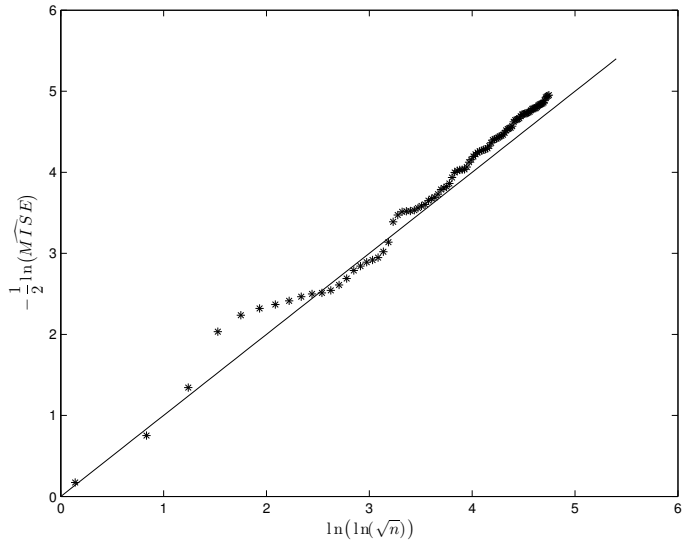

Figure 3.1: $-\frac{1}{2}\ln\left(\widehat{\text{MISE}}\right)$ plotted against $\ln\left(\ln(\sqrt{n})\right)$ for $n = 10^k,\ k = 1, ..., 100$ in the case $b = s = 1, \alpha = 2, \beta = 0, \gamma = 1$, for fixed $\tau = 1$.

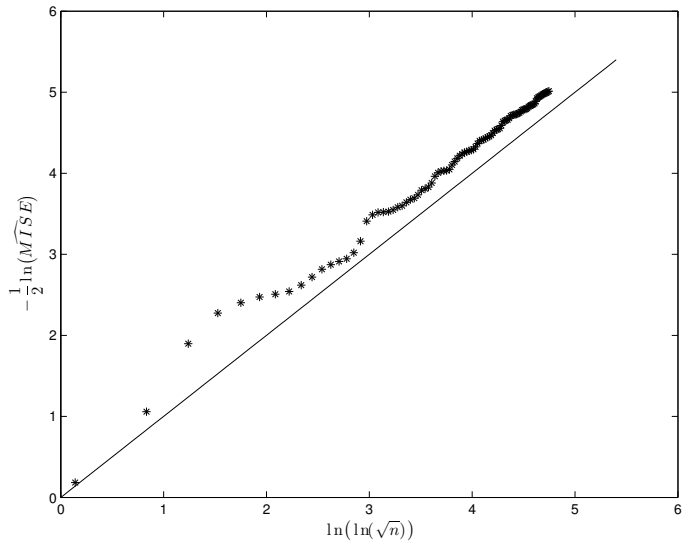

Figure 3.2: $-\frac{1}{2}\ln\left(\widehat{\text{MISE}}\right)$ plotted against $\ln\left(\ln(\sqrt{n})\right)$ for $n = 10^k,\ k = 1, ..., 100$ in the case $b = s = 1, \alpha = 2, \beta = 2, \gamma = 1$, for fixed $\tau = 1$.

3.6 Conclusions

We have considered a class of Bayesian severely ill-posed linear inverse problems with Gaussian additive noise and Gaussian priors in a diagonal setting, that is a setting in which the three operators defining the problem are simultaneously diagonalizable. In particular, we assumed that the forward operator K has singular values which decay like $\exp(-sj^b)$ for $s, b > 0$. In addition to the problem of determining the initial condition of the heat equation considered in [45] ($b = 2$), our theory covers a range of other severely ill-posed inverse problems such as the Cauchy problem for the Helmholtz equation ($b = 1$).

We showed that in our severely ill-posed setting the posterior is absolutely continuous with respect to the prior almost surely with respect to the joint distribution of the unknown and the data (Theorem 3.3.2). This is in contrast to the mildly ill-posed case where it is possible to have that the posterior and the prior are mutually singular independently of the data; this happens if the prior is not sufficiently regularizing (see Proposition 3.3.3).

We also showed rates of posterior contraction in the small noise limit (Theorem 3.4.3) and in particular generalized the sharp rates obtained in [45] for the case $b = 2$ to our generalized setup. Our analysis is inspired by the techniques used in [45], however our more general setting leads to technical improvements in the proofs (for example Lemma 3.4.5). As in [45], we have that the posterior contracts at the minimax rate if either the prior is oversmoothing the truth (in our notation $\alpha \geq \gamma + \frac{1}{2}$) and the scaling of the prior is fixed, or for a prior of any regularity by rescaling it appropriately as the noise disappears.

Finally, we presented a numerical simulation supporting the obtained convergence rate of the mean integrated squared error of the posterior mean, in the case of the Cauchy problem for the Helmholtz equation.

Chapter 4

Dimension Dependence of Sampling Algorithms in Hierarchical Bayesian Inverse Problems

4.1 Introduction

Let $\mathcal{X}$ be an infinite dimensional separable Hilbert space with inner product $\langle \cdot, \cdot \rangle$ and norm $\| \cdot \|$. We denote by $\langle \cdot, \cdot \rangle_{\mathbb{R}^N}$ and $\| \cdot \|_{\mathbb{R}^N}$ the (possibly scaled) Euclidean inner product and norm in $\mathbb{R}^N$ and by $\| \cdot \|_{2,N}$ the corresponding Euclidean norm for $N \times N$ matrices. We assume that this norm and inner product on $\mathbb{R}^N$ are chosen so that, formally, the large N limit recovers the norm on the Hilbert space.

We consider the linear inverse problem of recovering an unknown parameter $u \in \mathcal{X}$ from a blurred noisy observation y. We adopt a Bayesian approach and assume the additive noise model

$$y = Ku + \eta, \tag{4.1.1}$$

where $K : \mathcal{X} \to \mathcal{X}$ is a bounded linear operator, and η is Gaussian noise

$$\eta \sim \mathcal{N}(0, \sigma^{-1}\mathcal{C}_1), \tag{4.1.2}$$

where $\mathcal{C}_1 : \mathcal{X} \to \mathcal{X}$ is a bounded positive definite linear operator; we do not enforce that $\mathcal{C}_1$ is trace-class, thereby allowing the case of Gaussian white noise where it is

the identity. We put a Gaussian prior on the unknown parameter u

$$u \sim \mathcal{N}(0, \delta^{-1}\mathcal{C}_0), \tag{4.1.3}$$

where $\mathcal{C}_0 : \mathcal{X} \to \mathcal{X}$ is a positive definite and trace-class operator. The trace-class assumption ensures that draws from the prior are in $\mathcal{X}$.

For a fixed u the likelihood is Gaussian, $y|u \sim \mathcal{N}(Ku, \sigma^{-1}\mathcal{C}_1)$, and by the linearity of K we have conjugacy, that is, the posterior distribution is also Gaussian

$$u|y \sim \mathcal{N}(m_{\sigma,\delta}(y), \mathcal{C}_{\sigma,\delta});$$

see [54, 52] where formulae for the posterior mean and covariance are provided. Additionally, if we temporarily assume that we are in finite dimensions we have that the log-likelihood is quadratic in u, hence by completing the square we obtain the formulae for the inverse covariance and mean

$$\mathcal{C}_{\sigma,\delta}^{-1} = \sigma K^* \mathcal{C}_1^{-1} K + \delta \mathcal{C}_0^{-1}, \tag{4.1.4}$$

$$\mathcal{C}_{\sigma,\delta}^{-1} m_{\sigma,\delta}(y) = \sigma K^* \mathcal{C}_1^{-1} y. \tag{4.1.5}$$

This can be made rigorous in the separable Hilbert space setting in a range of situations, see for instance [3, 4]*.

In practice the unknown function is discretized and the Bayesian setup is implemented in $\mathbb{R}^N$. It is then of interest to refine the model, that is, to increase N. We study this issue, building on the recent paper [7]. We hence assume to have a way of computing discretizations $\boldsymbol{y_N} \in \mathbb{R}^N$ of the observation y, and we replace the operators $K, \mathcal{C}_0$ and $\mathcal{C}_1$ by $N \times N$ matrices, also denoted by $K, \mathcal{C}_0$ and $\mathcal{C}_1$, which arise from a consistent (in the sense of numerical analysis) family of approximations of the corresponding operators at discretization level N. Since we now work in $\mathbb{R}^N$ the formulae (4.1.4) and (4.1.5) for the posterior covariance and mean are always valid, and in fact are always equivalent to the finite dimensional analogues of the formulae for the posterior covariance and mean obtained in [54, 52], [78, Example 6.23].

We are interested in performing hierarchical inference simultaneously on the unknown u and the hyper-parameters σ^{-1} and δ^{-1}. Note that in the above infinite dimensional setup the hyper-parameters σ^{-1} and δ^{-1} have a clear interpretation as the scalings of the covariance operators of the noise and prior distributions respectively. If consistent approximations of limiting infinite dimensional operators are

*[3, 4] are Chapters 2 and 3 respectively in the current thesis.

not used, as in [7] for example where the Bayesian setup was performed directly on discrete problems, then it is not natural to compare δ^{-1} and σ^{-1} across different discretization levels and, in particular, to consider their limit under model-refinement; in contrast, using consistent approximations leads to interpretable hyper-parameters in the context of the limiting underlying infinite dimensional inverse problem. This is clearly very desirable in the applied context where simulations will routinely be performed at different levels of mesh-refinement.

We put inverse-Gamma hyper-priors on σ^{-1} and δ^{-1}, that is we assume that $\delta \sim \text{Gamma}(\alpha_0, \beta_0)$ and $\sigma \sim \text{Gamma}(\alpha_1, \beta_1)$, where the shape parameters $\alpha_0, \alpha_1 > 0$ and the rate parameters $\beta_0, \beta_1 > 0$ are chosen independently of N. It is well known that inverse-Gamma hyper-priors on the scaling of a Gaussian prior are conditionally conjugate, that is, $\delta|\boldsymbol{y_N}, u, \sigma$ is also Gamma. It is also straightforward to check that we have conditional conjugacy for σ. In particular, we have (see [7]),

$$\delta|\boldsymbol{y_N}, u, \sigma \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}}u\big\|_{\mathbb{R}^N}^2), \tag{4.1.6}$$

$$\sigma|\boldsymbol{y_N}, u, \delta \sim \text{Gamma}(\alpha_1 + \frac{N}{2}, \beta_1 + \frac{1}{2}\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku - \boldsymbol{y_N})\big\|_{\mathbb{R}^N}^2). \tag{4.1.7}$$

We analyze the large N behaviour of algorithms for sampling the full posterior on the discretized unknown u and the hyper-parameters δ and σ given the discretized data $\boldsymbol{y_N}$. Because of the conditional conjugacy of the three components, it is natural to use a Gibbs sampler, where we update the parameters one at a time. The algorithm is described below:

Algorithm 1.

0. Initialize $\delta^{(0)}$ and $\sigma^{(0)}$, and set $k = 0$;
1. Compute $u^{(k)} \sim \mathcal{N}\big(m_{\sigma^{(k)},\delta^{(k)}}(\boldsymbol{y_N}), \mathcal{C}_{\sigma^{(k)},\delta^{(k)}}\big)$;
2. Compute $\delta^{(k+1)} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}}u^{(k)}\big\|_{\mathbb{R}^N}^2)$;
3. Compute $\sigma^{(k+1)} \sim \text{Gamma}(\alpha_1 + \frac{N}{2}, \beta_1 + \frac{1}{2}\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku^{(k)} - \boldsymbol{y_N})\big\|_{\mathbb{R}^N}^2)$;
4. Set $k = k + 1$. If $k < k_{max}$ return to step 1, otherwise stop.

The steps 2 and 3 can be turned on or off depending on whether we want to have fixed δ or σ. Note that the draws $\delta^{(k+1)}$ and $\sigma^{(k+1)}$ appearing in the Gibbs

sampler depend on the dimension N of the discretization of the data y and the unknown parameter u. This dependence will be made explicit in the next section, in which we investigate the behaviour of the Gibbs sampler in the limit $N \to \infty$.

The fact that the Gamma(α, β) distribution has mean and variance $\alpha\beta^{-1}$ and $\alpha\beta^{-2}$ respectively, implies that for any $\mu > 0$, as N grows, the random variable Gamma$(\alpha + \frac{N}{2}, \beta + \mu\frac{N}{2})$ behaves like a Dirac distribution centred on μ^{-1}. Furthermore, we will show that, because of the consistency of the approximation of the operators defining the Bayesian inverse problem, together with scaling of the norms on $\mathbb{R}^N$ to reproduce the Hilbert space norm limit, it is natural to assume that

i) $\left\|\mathcal{C}_0^{-\frac{1}{2}} u^{(k)}\right\|_{\mathbb{R}^N}^2 \simeq (\delta^{(k)})^{-1} N$;

ii) $\left\|\mathcal{C}_1^{-\frac{1}{2}} (K u^{(k)} - \boldsymbol{y_N})\right\|_{\mathbb{R}^N}^2 \simeq \bar{\sigma}^{-1} N$, where $\bar{\sigma}^{-1}$ is the true scaling of the noise in the data, and hence $\boldsymbol{y_N}$.

Using the limiting behaviour of the Gamma distribution described above, this means that as the dimension N increases, on the one hand $\delta^{(k+1)} \simeq \delta^{(k)}$ hence the δ-chain makes very small moves and slows down, while on the other hand $\sigma^{(k+1)} \simeq \bar{\sigma}$ hence the σ-chain goes instantly to the true value of the noise scaling in the data and stays close to it. We will make these ideas precise in what follows.

Our results show that in the context of natural Gibbs sampling algorithms, two seemingly similar choices of hyper-parameterization of the scale in the prior and noise models lead to very different mixing behaviour in the algorithm, in the limit of high dimensional approximations of the inverse problem. To alleviate the effect of poor mixing in the δ-chain, following intuition from [72, 60], we propose a reparametrization of the problem in which the two components on the unknown and the prior scaling are made a priori independent (see Section 4.4). We again have conditional conjugacy, hence use the Gibbs sampler once more. The reparametrized algorithm is robust with respect to the increase in dimension, however, it deteriorates as the quality of the data improves (small observational noise limit); new ideas are required in this situation.

In addition to [7], such conditionally conjugate hierarchical setups have been studied in the case of linear inverse problems in [79] and in the setting of nonparametric drift estimation in [56]. In both [79] and [56] the hierarchical inference is performed only on the prior scaling and is motivated by results on posterior consistency in the frequentist sense, namely the fact that the optimal rates of contraction are in general achieved by appropriately rescaling the prior depending on the quality of the data, [44, 3, 64, 45, 4]*. Formally, the effect of rescaling is to change the reg-

*[3, 4] are Chapters 2 and 3 respectively in the current thesis.

ularity of draws from the prior by changing the spectrum of its covariance operator, in an effort to match the regularity of the unknown. Unfortunately, the optimal relationship between the prior scaling and the quality of the data depends on the regularity of the true solution underlying the data, which is generally not available, hence it is desirable to have a hierarchical setup in order to perform inference simultaneously on the unknown function and the prior scaling parameter.

4.1.1 Notation

In order to avoid heavy notation in the rest of the chapter, we make the dependence on the discretization level N explicit only in the data $\boldsymbol{y_N}$ and the δ and σ-chains. In our subsequent presentation and analysis of the Gibbs sampler we use superscripts to denote the iteration number. We also use the notation $m_{\sigma^{(k)},\delta^{(k)}} := m_{\sigma=\sigma^{(k)},\delta=\delta^{(k)}}$ and similarly for $\mathcal{C}_{\sigma^{(k)},\delta^{(k)}}$. For a random variable x which depends on the mutually independent random variables f_1 and f_2, we use the notation $\mathbb{E}^{f_1}[x]$ to denote the expectation of the random variable x with respect to the random variable f_1 for fixed f_2. We use the notation $x_1 \overset{\mathcal{L}}{=} x_2$ to denote that the random variables x_1 and x_2 have the same law. Finally, for two sequences of positive numbers $\{s_j\}$ and $\{t_j\}$, we use the notation $s_j \asymp t_j$ to mean that s_j/t_j is bounded away from zero and infinity uniformly in j.

4.1.2 Chapter structure

The rest of the chapter proceeds as follows: in the next section we present our main results which hold under certain assumptions on the discrete level which are contained in the same section. We believe that our discrete level assumptions are inherited from certain natural assumptions on the underlying infinite dimensional model whenever consistent numerical discretizations are used. These natural infinite dimensional assumptions are presented in Section 4.3, before providing the proofs of our main results. In Section 4.4 we present our proposed reparametrization of the prior scaling. In Section 4.5 we present extensions of our theory: first we show that our analysis of the behaviour of the δ-chain holds in more general conjugate-Gaussian settings and second we consider cases where the unknown and the data are discretized at different levels. In Section 4.6 we exhibit four classes of inverse problems satisfying our assumptions on the underlying infinite dimensional model. For the first two of these classes, that is a class of mildly ill-posed and a class of severely ill-posed linear inverse problems both in a simultaneously diagonalizable setting, we also explicitly prove that our discrete level assumptions are inherited

from the infinite dimensional assumptions when discretizing via spectral truncation (see Subsections 4.6.1 and 4.6.2). In Section 4.7 we present numerical evidence supporting our theory in a wider class of mildly ill-posed linear inverse problems, using both spectral truncation (Subsection 4.7.1) as well as discretization via finite differences (Subsections 4.7.2 and 4.7.3). The main body of the chapter ends with concluding remarks in Section 4.8, while the Appendix in Section 4.9 contains several technical lemmas.

4.2 Main Results

We now present our main results on the large N behaviour of Algorithm 1. In order to simplify our analysis we examine separately the case where δ is random and to be determined through hierarchical inference while σ is fixed, and the case where σ is random and to be determined through hierarchical inference while δ is fixed. As hinted earlier on, we have two main results regarding the behaviour of the Gibbs sampler in the large N limit. The first one concerns the slowing down of the δ-chain when σ is fixed, while the second one is about the speeding up of the σ-chain when δ is fixed.

In the following, we assume that $\mathcal{C}_0$ and $\mathcal{C}_1$ are positive definite $N \times N$ real matrices which are the discretizations of the positive definite operators $\mathcal{C}_0$ and $\mathcal{C}_1$ respectively, and the $N \times N$ matrix K is the discretization of the bounded operator K. We do not make the dependence of $K, \mathcal{C}_0$ and $\mathcal{C}_1$ on N explicit, however,we do make explicit the dependence on N of the discretized data $\boldsymbol{y_N}$.

4.2.1 Large N behaviour of the δ-chain

In this subsection we present our result regarding the behaviour of the δ-chain for fixed $\sigma > 0$. The Gibbs sampler in this case is as follows:

Algorithm 2.

0. Initialize $\delta^{(0)}$ and set $k = 0$;
1. Compute $u^{(k)} \sim \mathcal{N}\big(m_{\sigma,\delta^{(k)}}(\boldsymbol{y_N}), \mathcal{C}_{\sigma,\delta^{(k)}}\big)$;
2. Compute $\delta^{(k+1)} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}} u^{(k)}\big\|^2_{\mathbb{R}^N})$;
3. Set $k = k + 1$. If $k < k_{max}$ return to step 1, otherwise stop.

Note that the two steps of updating $u|\boldsymbol{y_N}, \delta$ and $\delta|\boldsymbol{y_N}, u$ can be compressed to give one step of updating δ and involving the noise in the u update. Indeed, we denote by $\delta_N^{(k+1)}$ the δ-draw in the $k+1$ iteration of the Gibbs sampler where the problem is discretized in $\mathbb{R}^N$. This draw is made using the previous draw of $u|\boldsymbol{y_N}, \delta$, which assuming that $\delta_N^{(k)} = \delta$, is denoted by $u_\delta^{(k)}$ and can be written as

$$u_\delta^{(k)} = m_{\sigma,\delta}(\boldsymbol{y_N}) + \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta, \tag{4.2.1}$$

where ζ is an N-dimensional Gaussian white noise representing the fluctuation in step 1, and $\mathcal{C}_{\sigma,\delta}, m_{\sigma,\delta}$ are given by the formulae (4.1.4), (4.1.5) respectively. Hence we have

$$\delta_N^{(k+1)} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2}\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2). \tag{4.2.2}$$

Our analysis of the δ-chain is valid under the following assumptions:

Assumptions 4.2.1.

i) For almost all data y, for any $\sigma, \delta > 0$, there exists a constant $c_1 = c_1(y; \sigma, \delta) \geq 0$, independent of N, such that

$$\|\mathcal{C}_0^{-\frac{1}{2}} m_{\sigma,\delta}(\boldsymbol{y_N})\|_{\mathbb{R}^N} \leq c_1;$$

ii) there exists a constant $c_2 \geq 0$, independent of N and y, such that

$$\text{Tr}(\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0 K^* \mathcal{C}_1^{-\frac{1}{2}}) \leq c_2.$$

The above assumptions ensure that the squared norm appearing in (4.2.2) behaves like $\delta^{-1}N$, as assumed in the motivating discussion in the introduction. Combining with the scaling property of the Gamma distribution, we show that as the dimension increases the δ-chain makes smaller and smaller steps, and quantify the scaling of this slowing down.

Theorem 4.2.2. *For fixed $\sigma > 0$, under Assumptions 4.2.1, in the limit $N \to \infty$, we have almost surely with respect to y:*

i) the expected step in the δ-chain scales like $\frac{2}{N}$, that is, for any $\delta > 0$,

$$\frac{N}{2}\mathbb{E}\left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta\right] = (\alpha_0 + 1)\delta - f_N(\delta; \boldsymbol{y_N})\delta^2 + \mathcal{O}(N^{-\frac{1}{2}}),$$

where $f_N(\delta; \boldsymbol{y_N})$ is bounded uniformly in N. In particular, if there exists $f(\delta; y) \in \mathbb{R}$ such that $f_N(\delta; \boldsymbol{y_N}) \to f(\delta; y)$ then

$$\frac{N}{2}\mathbb{E}\left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta\right] = (\alpha_0 + 1)\delta - f(\delta; y)\delta^2 + \mathcal{O}(1);$$

ii) the variance of the step also scales like $\frac{2}{N}$ and in particular for any $\delta > 0$,

$$\begin{aligned}\frac{N}{2}\mathbb{E}\left[\left(\delta_N^{(k+1)} - \delta_N^{(k)} - \mathbb{E}\left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta\right]\right)^2 | \delta_N^{(k)} = \delta\right] \\ = 2\delta^2 + \mathcal{O}(N^{-\frac{1}{2}}).\end{aligned}$$

All the expectations are taken with respect to the randomness in the algorithm.

Remark 4.2.3.

i) The proof of Theorem 4.2.2 is contained in Subsection 4.3.2. Moreover, in Subsection 4.3.1, we have a discussion on our assumptions, and more detailed intuition on the behaviour of the δ-chain based on measure theoretic arguments in the underlying infinite dimensional model. In Subsections 4.6.1 and 4.6.2 we demonstrate two classes of linear inverse problems for which Assumptions 4.2.1 are valid when using discretization via spectral truncation. In Section 4.7 we provide numerical evidence demonstrating the conclusions of Theorem 4.2.2 using both spectral truncation and discretization via finite differences, in several mildly ill-posed linear inverse problem settings.

ii) $f_N(\delta; \boldsymbol{y_N}) := \mathbb{E}^{\zeta}[F_N(\delta; \boldsymbol{y_N})]$, where F_N is defined in the proof of Lemma 4.3.3. The assumption on the convergence of $f_N(\delta; \boldsymbol{y_N})$ is trivially satisfied under Assumptions 4.2.1, if the discretization scheme used is such that if the vector $x \in \mathbb{R}^N$ and the $N \times N$ matrix T are the discretizations at level N of $x \in \mathcal{X}$ and the linear operator T respectively, then $\|Tx\|_{\mathbb{R}^N}$ is a non-decreasing sequence. This is the case for example in spectral truncation discretization methods, when T is diagonalizable in the orthonormal basis used (see Subsections 4.6.1 and 4.6.2).

iii) Theorem 4.2.2 suggests that almost surely with respect to the data, for large N the δ-chain makes moves which on average are of order N^{-1} with fluctuations of order $N^{-\frac{1}{2}}$. As a result, it takes $\mathcal{O}(N)$ steps for the δ-chain to move by $\mathcal{O}(1)$. This is reflected in the numerical simulations in Section 4.7, where we observe that as the dimension N grows, the δ-chain slows down and in particular appears to behave like the solution to a SDE.

iv) Let $\delta_N(t)$ be the path inscribed by the δ-chain in N dimensions, that is, $\delta_N(t)$ is the piecewise constant function defined as $\delta_N^{(k)}$ on $[t_k, t_{k+1})$ for $t_k = k/N$. Formally, in the case where f_N has a limit, we have that almost surely with respect to the data, as $N \to \infty$, $\delta_N(t)$ can be approximated by the solution $\delta = \delta(t)$ of the SDE

$$d\delta = \big(\alpha_0 + 1 - f(\delta; y)\delta\big)\delta dt + \sqrt{2}\delta dW, \tag{4.2.3}$$

where $W = W(t)$ is a standard Brownian motion.

4.2.2 Large N behaviour of the σ-chain

We now present our result about the behaviour of the σ-chain for fixed $\delta > 0$. The Gibbs sampler in this case is as follows:

Algorithm 3.

0. Initialize $\sigma^{(0)}$ and set $k = 0$;
1. Compute $u^{(k)} \sim \mathcal{N}\big(m_{\sigma^{(k)},\delta}(\boldsymbol{y_N}), \mathcal{C}_{\sigma^{(k)},\delta}\big)$;
2. Compute $\sigma^{(k+1)} \sim \text{Gamma}(\alpha_1 + \frac{N}{2}, \beta_1 + \frac{1}{2}\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku^{(k)} - \boldsymbol{y_N})\big\|^2_{\mathbb{R}^N})$;
3. Set $k = k + 1$. If $k < k_{max}$ return to step 1, otherwise stop.

As in Subsection 4.2.1, the two steps of updating $u|\boldsymbol{y_N}, \sigma$ and $\sigma|\boldsymbol{y_N}, u$ can be compressed to give one step of updating σ and involving the noise in the u update. Indeed, we denote by $\sigma_N^{(k+1)}$ the σ-draw in the $k+1$ iteration of the Gibbs sampler where the problem is discretized in $\mathbb{R}^N$. This draw is made using the previous draw of $u|\boldsymbol{y_N}, \sigma$, which assuming that $\sigma_N^{(k)} = \sigma$, is denoted by $u_\sigma^{(k)}$ and can be written as

$$u_\sigma^{(k)} = m_{\sigma,\delta}(\boldsymbol{y_N}) + \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta, \tag{4.2.4}$$

where ζ is an N-dimensional Gaussian white noise representing the fluctuation in step 1, and $\mathcal{C}_{\sigma,\delta}, m_{\sigma,\delta}$ are given by the formulae (4.1.4), (4.1.5) respectively. Hence we have

$$\sigma_N^{(k+1)} \sim \text{Gamma}(\alpha_1 + \frac{N}{2}, \beta_1 + \frac{1}{2}\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku_\sigma^{(k)} - \boldsymbol{y_N})\big\|^2_{\mathbb{R}^N}). \tag{4.2.5}$$

We consider data in $\mathbb{R}^N$ of the form

$$\boldsymbol{y_N} = z + \bar{\sigma}^{-\frac{1}{2}} \mathcal{C}_1^{\frac{1}{2}} \xi, \tag{4.2.6}$$

where z is the discretization in $\mathbb{R}^N$ of an element of the space $\mathcal{X}$, also denoted by z, which satisfies Assumption 4.2.4 below, and ξ is a Gaussian white noise in $\mathbb{R}^N$ which is independent of ζ. In this case, $\bar{\sigma}^{-1}$ is the true value of the scaling σ^{-1}.

Assumption 4.2.4. *There exists $c_3 = c_3(z)$ independent of N and y, such that,*

$$\left\| \mathcal{C}_1^{-\frac{1}{2}} z \right\|_{\mathbb{R}^N} \leq c_3.$$

Assumption 4.2.4 together with Assumption 4.2.1 ensure that the squared norm appearing in (4.2.5) behaves like $\bar{\sigma}^{-1}N$, as assumed in the motivating discussion in the introduction. Combining with the scaling property of the Gamma distribution, we show that as the dimension increases the σ-chain goes to the correct value $\bar{\sigma}$ immediately, and quantify the scaling of this speed up.

Theorem 4.2.5. *For fixed $\delta > 0$, under Assumptions 4.2.1 and 4.2.4, we have that in the limit $N \to \infty$, for any $\sigma, \bar{\sigma} > 0$:*

i)

$$\frac{N}{2} \mathbb{E}\left[\sigma_N^{(k+1)} - \bar{\sigma} | \sigma_N^{(k)} = \sigma \right] = (\alpha_1 + 1)\bar{\sigma} - h_N(\sigma, \bar{\sigma})\bar{\sigma}^2 + \mathcal{O}(N^{-\frac{1}{2}}),$$

where $h_N(\sigma, \bar{\sigma})$ is bounded uniformly in N;

ii)

$$\begin{aligned} &\frac{N}{2} \mathbb{E}\left[\left(\sigma_N^{(k+1)} - \bar{\sigma} - \mathbb{E}\left[\sigma_N^{(k+1)} - \bar{\sigma} | \sigma_N^{(k)} = \sigma \right] \right)^2 | \sigma_N^{(k)} = \sigma \right] \\ &\qquad = 2\bar{\sigma}^2 + \mathcal{O}(N^{-\frac{1}{2}}). \end{aligned}$$

The expectations are taken with respect to both the noise in the data and the randomness in the algorithm.

Remark 4.2.6.

i) The proof of Theorem 4.2.5 can be found in Subsection 4.3.3. Moreover, in Subsection 4.3.1, we have a discussion on our assumptions, and more detailed intuition on the behaviour of the σ-chain based on measure theoretic arguments in the underlying infinite dimensional model. In Subsections 4.6.1 and 4.6.2

we demonstrate two classes of linear inverse problems for which Assumptions 4.2.1 and 4.2.4 are valid when using discretization via spectral truncation. In Section 4.7 we provide numerical evidence demonstrating the conclusions of Theorem 4.2.5 using both spectral truncation and discretization via finite differences, in two mildly ill-posed diagonal linear inverse problem settings.

ii) $h_N(\sigma, \bar{\sigma}) := \mathbb{E}[H_N(\sigma, \bar{\sigma})]$, *where* H_N *is defined in the proof of Lemma 4.2.4.*

iii) Theorem 4.2.5 says that, on average with respect to the data and the algorithmic randomness, the σ*-chain makes moves which are within order* N^{-1} *distance from the true value* $\bar{\sigma}$ *of* σ*, with fluctuations of order* $N^{-\frac{1}{2}}$*. In fact our numerical results in Section 4.7 illustrate that the* σ*-chain exhibits this behaviour almost surely with respect to the data.*

4.3 Proofs of main results

In this section we give the proofs of our two main results presented in Section 4.2. First, in Subsection 4.3.1, we provide some more intuition on our results based on measure theoretic arguments in the underlying infinite dimensional model. Then, in Subsection 4.3.2 we prove Theorem 4.2.2 on the slowing down of the δ-chain when σ is fixed, and in Subsection 4.3.3 we prove Theorem 4.2.5 on the speed up of the σ-chain when δ is fixed. The proofs rely on a series of technical lemmas contained in Section 4.9.

4.3.1 Intuition based on the underlying infinite dimensional model

Starting with the analysis of the δ-chain, we have that on a high level our assumptions express the fact that in the underlying infinite dimensional model, almost surely with respect to the data, the Gaussian conditional posterior on u is absolutely continuous with respect to the prior. In particular, we expect that for a reasonable discretization scheme, Assumptions 4.2.1 are inherited from the following infinite dimensional assumptions.

Assumptions 4.3.1 (Infinite dimensional analogue of Assumptions 4.2.1)**.**

i) For any $\sigma, \delta > 0$*, we have* $m_{\sigma,\delta}(y) \in \mathcal{D}(\mathcal{C}_0^{-\frac{1}{2}})$ *almost surely with respect to* y*; that is, the posterior mean belongs to the Cameron-Martin space of the prior on* $u|\delta$*;*

ii) $\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0 K^* \mathcal{C}_1^{-\frac{1}{2}}$ *is trace-class; that is, the prior is sufficiently regularizing.*

It is a straightforward exercise to check that Assumption 4.3.1(ii) above implies the first and third assumptions of the Feldman-Hajek theorem [18, Theorem 2.23] on the equivalence of two Gaussian measures and in particular the conditional posterior and the prior on u, hence, together with Assumption 4.3.1(i), they imply that the conditional posterior on u is indeed y-almost surely absolutely continuous with respect to the prior on u. In infinite dimensions, the Gaussian measures $\mathcal{N}(0, a_1\Sigma)$ and $\mathcal{N}(0, a_2\Sigma)$ are mutually singular unless $a_1 = a_2$, [17, Remark 2.10]. Hence, $u|y, \sigma, \delta$ is singular with respect to $u|y, \sigma, \delta'$ for any $\delta' \neq \delta$, since each is absolutely continuous to the corresponding prior measure and those are mutually singular. Statistically, this means that a single realization of u from the prior (or by equivalence from the posterior) would perfectly identify the δ that has generated it, that is, $\delta|u$ concentrates on a point pass. However, $\delta|y$ is not concentrated on a point mass, precisely because u is not perfectly identified from y. For discretizations of the model at level N, $\delta|u$ does not concentrate on a point mass but $\mathrm{Var}(\delta|u) = \mathcal{O}(N^{-1})$, whereas $\mathrm{Var}(\delta|\boldsymbol{y_N}) = \mathcal{O}(1)$. Therefore, Algorithm 2 needs to sample a more or less fixed distribution as N increases, while doing smaller and smaller steps, precisely those permitted by $\delta|u$, which are of size $\mathcal{O}(N^{-1/2})$.

For the analysis of the σ-chain, we assume that we have infinite dimensional data of the form

$$y = z + \bar{\sigma}^{-\frac{1}{2}}\mathcal{C}_1^{\frac{1}{2}}\xi,$$

where ξ is Gaussian white noise, $\bar{\sigma}^{-1}$ is the true scaling of the noise in the data, and $z \in \mathcal{X}$. We expect that for a reasonable discretization scheme, Assumption 4.2.4 is inherited from the following regularity assumption on z.

Assumption 4.3.2 (Infinite dimensional analogue of Assumption 4.2.4)**.**
$z \in \mathcal{D}(\mathcal{C}_1^{-\frac{1}{2}})$; that is, for any $\bar{\sigma} > 0$, z belongs to the Cameron-Martin space of the noise measure $\mathcal{N}(0, \bar{\sigma}^{-1}\mathcal{C}_1)$.

The last assumption, secures that the data generating measure $\mathcal{N}(z, \bar{\sigma}^{-1}\mathcal{C}_1)$ is absolutely continuous with respect to the noise measure $\mathcal{N}(0, \bar{\sigma}^{-1}\mathcal{C}_1)$, [17, Theorem 2.8]. Combining again with [17, Remark 2.10] on the singularity of two infinite dimensional Gaussian measures with covariance operators which are proportional to each other, we get that the distribution of the data $\mathcal{N}(z, \bar{\sigma}^{-1}\mathcal{C}_1)$, is absolutely continuous with respect to the conditional noise distribution $\eta|\bar{\sigma}$, but singular to $\eta|\sigma$ for any $\sigma \neq \bar{\sigma}$. Since again, given $\mathcal{C}_1$, the value of σ can be uniquely determined from an infinite dimensional draw of $\eta|\sigma$, we have that the infinite dimensional data contain as much information about $\bar{\sigma}$ as a draw from $\eta|\bar{\sigma}$, which in turn contains full

information about $\bar{\sigma}$. Hence for large N, the draw $\sigma|\boldsymbol{y_N}, u, \delta$ is strongly dependent on the data $\boldsymbol{y_N}$ which, however, encode the true scaling of the noise $\bar{\sigma}^{-1}$, thus σ instantly identifies the true value $\bar{\sigma}$ underlying the data.

In Section 4.6, we present several families of inverse problems satisfying the above underlying infinite dimensional assumptions. For two of these families, we also show that our discrete level assumptions in Section 4.2, are inherited when using spectral truncation (see Subsections 4.6.1 and 4.6.2). Furthermore, in Section 4.7 we provide numerical evidence supporting our theory on the behaviour of the δ and σ-chains, using both spectral truncation as well as finite difference approximation.

4.3.2 Proof of Theorem 4.2.2

We now give the proof of Theorem 4.2.2 under Assumptions 4.2.1. Using the scaling property of the Gamma distribution, $\text{Gamma}(\alpha, \beta) \overset{\mathcal{L}}{=} \beta^{-1}\text{Gamma}(\alpha, 1)$, and multiplying and dividing by $\frac{2}{N}\delta$, we can write the $\delta_N^{(k+1)}$ draw in (4.2.2) as

$$\delta_N^{(k+1)} \overset{\mathcal{L}}{=} \delta \frac{\Gamma_{0,N}}{\frac{2}{N}\delta(\beta_0 + \frac{1}{2}\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2)} \tag{4.3.1}$$

where $\Gamma_{0,N} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \frac{N}{2})$ is independent of $\boldsymbol{y_N}$ and $u_\delta^{(k)}$.

The following lemma forms the backbone of our analysis of the δ-chain and in particular of the proof of Theorem 4.2.2.

Lemma 4.3.3. *Under Assumptions 4.2.1, for any $\sigma, \delta > 0$ we have,*

$$\beta_0 + \frac{1}{2}\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2 = \delta^{-1}\frac{N}{2} + \delta^{-1}\sqrt{\frac{N}{2}} W_{1,N} + F_N(\delta), \tag{4.3.2}$$

where i) $W_{1,N}$ only depends on the white noise ζ in (4.2.1), has mean zero and variance one, higher order moments which are bounded uniformly in N, and converges weakly to a standard normal random variable as $N \to \infty$; ii) $F_N(\delta)$ depends on the data $\boldsymbol{y_N}$ and y-almost surely has finite moments of all positive orders uniformly in N (where the expectation is taken with respect to ζ).

The proof of Lemma 4.3.3 is contained in Subsection 4.9.1. Defining

$$W_{2,N} = \frac{\Gamma_{0,N} - 1 - \frac{2\alpha_0}{N}}{\sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}}},$$

we have

$$\Gamma_{0,N} = 1 + \frac{2\alpha_0}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}} W_{2,N},$$

where for every N, the random variable $W_{2,N}$ has mean zero and variance one, third and fourth moments which are bounded uniformly in N (see Lemma 4.9.6), and is independent of the data $\boldsymbol{y_N}$ and ζ, the white noise expressing the fluctuation in $u_\delta^{(k)}$. Concatenating we get

$$\delta_N^{(k+1)} \overset{\mathcal{L}}{=} \delta \frac{1 + \frac{2\alpha_0}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}} W_{2,N}}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2}{N} F_N(\delta)\delta}, \tag{4.3.3}$$

and we are now ready to prove Theorem 4.2.2:

Proof. By the independence of $W_{2,N}$ and ζ and since $\mathbb{E}[W_{2,N}] = 0$, we have

$$\begin{aligned}
\mathbb{E}[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] &= \delta \mathbb{E}\left[\frac{1 + \frac{2\alpha_0}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}} W_{2,N}}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N\delta}{N}} - 1 \right] \\
&= \delta \mathbb{E}^\zeta \left[\frac{1 + \frac{2\alpha_0}{N}}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N\delta}{N}} - 1 \right] \\
&= \delta \mathbb{E}^\zeta \left[\frac{\frac{2\alpha_0}{N} - \sqrt{\frac{2}{N}} W_{1,N} - \frac{2F_N\delta}{N}}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N\delta}{N}} \right].
\end{aligned}$$

Using the identity $\frac{1}{1+x} = 1 - x + \frac{x^2}{1+x}$ we get

$$\begin{aligned}
&\mathbb{E}[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] \\
=&\delta \mathbb{E}^\zeta \left[\left(\frac{2(\alpha_0 - F_N\delta)}{N} - \sqrt{\frac{2}{N}} W_{1,N} \right) \left(1 - \sqrt{\frac{2}{N}} W_{1,N} - \frac{2F_N\delta}{N} \right) \right] + \mathbb{E}^\zeta[e_{1,N}],
\end{aligned}$$

where

$$e_{1,N} = \delta \frac{\left(\frac{2(\alpha_0 - F_N\delta)}{N} - \sqrt{\frac{2}{N}} W_{1,N} \right) \left(\frac{2W_{1,N}^2}{N} + \frac{4F_N^2\delta^2}{N^2} + \frac{4\sqrt{2}F_N W_{1,N}\delta}{N^{\frac{3}{2}}} \right)}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N\delta}{N}}.$$

Using Hölder's inequality and the fact that F_N and $W_{1,N}$ have moments of

all positive orders which are bounded uniformly in N, we get

$$\mathbb{E}[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] = \frac{2}{N}\left((\alpha_0 + 1)\delta - \mathbb{E}^\zeta[F_N]\delta^2\right) + \mathcal{O}(N^{-\frac{3}{2}}) + \mathbb{E}^\zeta[e_{1,N}],$$

almost surely with respect to y. For the residual $e_{1,N}$, by Cauchy-Schwarz inequality and (4.3.2), we have

$$\begin{aligned}
\mathbb{E}^\zeta[e_{1,N}] &= \mathbb{E}^\zeta\Big[\frac{\left(\frac{2(\alpha_0 - F_N\delta)}{N} - \sqrt{\frac{2}{N}}W_{1,N}\right)\left(W_{1,N}^2 + \frac{2}{N}F_N^2\delta^2 + \frac{2\sqrt{2}}{N^{\frac{1}{2}}}F_N W_{1,N}\delta\right)}{\frac{N}{2\delta}(1 + \sqrt{\frac{2}{N}}W_{1,N} + \frac{2F_N\delta}{N})}\Big] \\
&\leq \left(\mathbb{E}\Big[\left(\frac{2(\alpha_0 - F_N\delta)}{N} - \sqrt{\frac{2}{N}}W_{1,N}\right)^2\left(W_{1,N}^2 + \frac{2F_N^2\delta^2}{N} + \frac{2\sqrt{2}F_N W_{1,N}\delta}{N^{\frac{1}{2}}}\right)^2\Big]\right)^{\frac{1}{2}} \\
&\quad \cdot\left(\mathbb{E}\Big[(\beta_0 + \frac{1}{2}\|\mathcal{C}_0^{-\frac{1}{2}}u_\delta^{(k)}\|_{\mathbb{R}^N}^2)^{-2}\Big]\right)^{\frac{1}{2}}.
\end{aligned}$$

The square root of the first expectation on the right hand side of the inequality is of order $N^{-\frac{1}{2}}$, while by Lemma 4.9.1 the square root of the second expectation is of order N^{-1} for almost all y. Combining we get that $\mathbb{E}^\zeta[e_{1,N}] = \mathcal{O}(N^{-\frac{3}{2}})$, almost surely with respect to y, hence

$$\mathbb{E}[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] = \frac{2}{N}\left((1 + \alpha_0)\delta - \mathbb{E}^\zeta[F_N]\delta^2\right) + \mathcal{O}(N^{-\frac{3}{2}}),$$

y-almost surely.

For the expected diffusion, we have

$$\begin{aligned}
&\mathbb{E}\left[\left(\delta_N^{(k+1)} - \delta_N^{(k)} - \mathbb{E}\left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta\right]\right)^2 | \delta_N^{(k)} = \delta\right] \\
=&\mathbb{E}\left[(\delta_N^{(k+1)} - \delta_N^{(k)})^2 | \delta_N^{(k)} = \delta\right] - \mathbb{E}\left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta\right]^2,
\end{aligned}$$

where by the first part of the proof the second term is $\mathcal{O}(N^{-2})$. Thus, we need only consider the first term, which will be shown to be $\mathcal{O}(N^{-1})$. Indeed, for the first term by equation (4.3.3) we have

$$
\begin{aligned}
&\mathbb{E}\left[(\delta_N^{(k+1)}-\delta_N^{(k)})^2|\delta_N^{(k)}=\delta\right]\\
&=\delta^2\mathbb{E}\left[\left(\frac{\frac{2\alpha_0}{N}+\sqrt{\frac{2}{N}+\frac{4\alpha_0}{N^2}}W_{2,N}-\sqrt{\frac{2}{N}}W_{1,N}-\frac{2F_N\delta}{N}}{1+\sqrt{\frac{2}{N}}W_{1,N}+\frac{2F_N\delta}{N}}\right)^2\right]\\
&=\delta^2\mathbb{E}\left[\frac{\frac{2W_{2,N}^2}{N}+\frac{2W_{1,N}^2}{N}+\frac{V_N}{N^{\frac{3}{2}}}}{\left(1+\sqrt{\frac{2}{N}}W_{1,N}+\frac{2F_N\delta}{N}\right)^2}\right],
\end{aligned}
$$

where the random variable V_N depends only on $W_{1,N}$ and F_N and has higher order moments which are bounded uniformly in N, y-almost surely (the dependence on $W_{2,N}$ disappears by the independence of $W_{2,N}$ and ζ and the fact that $W_{2,N}$ has mean zero and variance one). Using the identity $\frac{1}{(1+x)^2}=1-2x+\frac{3x^2+2x^3}{(1+x)^2}$, we get

$$
\begin{aligned}
&\mathbb{E}\left[(\delta_N^{(k+1)}-\delta_N^{(k)})^2|\delta_N^{(k)}=\delta\right]\\
=&\delta^2\mathbb{E}\left[\left(\frac{2W_{2,N}^2}{N}+\frac{2W_{1,N}^2}{N}+\frac{V_N}{N^{\frac{3}{2}}}\right)\left(1-2\sqrt{\frac{2}{N}}W_{1,N}-\frac{4}{N}F_N\delta\right)\right]+\mathbb{E}[e_{2,N}],
\end{aligned}
$$

where

$$
\begin{aligned}
&e_{2,N}\\
=&\delta^2\left(\frac{2W_{2,N}^2}{N}+\frac{2W_{1,N}^2}{N}+\frac{V_N}{N^{\frac{3}{2}}}\right)\frac{3\left(\sqrt{\frac{2}{N}}W_{1,N}+\frac{2F_N\delta}{N}\right)^2+2\left(\sqrt{\frac{2}{N}}W_{1,N}+\frac{2F_N\delta}{N}\right)^3}{\left(1+\sqrt{\frac{2}{N}}W_{1,N}+\frac{2F_N\delta}{N}\right)^2}\\
:=&\frac{E_N\delta^2}{\left(1+\sqrt{\frac{2}{N}}W_{1,N}+\frac{2F_N\delta}{N}\right)^2}.
\end{aligned}
$$

Using the fact that y-almost surely $W_{1,N}$, F_N and V_N have moments of all positive orders which are bounded uniformly in N, by Hölder inequality (we do not need to consider higher order moments for $W_{2,N}$ here, because it is independent with $W_{1,N}$ and F_N, hence bounding terms involving $W_{2,N}$ does not require the use of Hölder's inequality which needs higher moments), we get that

$$
\mathbb{E}[(\delta_N^{(k+1)}-\delta_N^{(k)})^2|\delta_N^{(k)}=\delta]=\frac{2\delta^2}{N}\left(\mathbb{E}[W_{2,N}^2]+\mathbb{E}[W_{1,N}^2]\right)+\mathcal{O}(N^{-\frac{3}{2}})+\mathbb{E}[e_{2,N}],
$$

y-almost surely. For the residual $e_{2,N}$, as before using Cauchy-Schwarz inequality

and (4.3.2), we have

$$\mathbb{E}[e_{2,N}] \leq \frac{N^2}{4}\big(\mathbb{E}[E_N^2]\big)^{\frac{1}{2}}\left(\mathbb{E}[(\beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}}u_\delta^{(k)}\big\|_{\mathbb{R}^N}^2)^{-4}]\right)^{\frac{1}{2}}.$$

Since by Lemma 4.9.6 the first four moments of $W_{2,N}$ are also bounded uniformly in N, the square root of the first expectation on the right hand side is of order N^{-2}, while by Lemma 4.9.1 the square root of the second expectation is of order N^{-2}, for almost all y. Combining we get $\mathbb{E}^\zeta[e_{2,N}] = \mathcal{O}(N^{-2})$, almost surely with respect to y, hence since $\mathbb{E}[W_{1,N}^2] = \mathbb{E}[W_{2,N}^2] = 1$, we have

$$\mathbb{E}[(\delta_N^{(k+1)} - \delta_N^{(k)})^2 | \delta_N^{(k)} = \delta] = \frac{4\delta^2}{N} + \mathcal{O}(N^{-\frac{3}{2}}),$$

y-almost surely. Concatenating, we get the result. □

4.3.3 Proof of Theorem 4.2.5

We now prove Theorem 4.2.5 under Assumptions 4.2.1 and 4.2.4. Again using the scaling property of the Gamma distribution, $\text{Gamma}(\alpha, \beta) \overset{\mathcal{L}}{=} \beta^{-1}\text{Gamma}(\alpha, 1)$, and multiplying and dividing by $\frac{2}{N}\bar{\sigma}$, we can write the $\sigma_N^{(k+1)}$ draw in (4.2.5) as

$$\sigma_N^{(k+1)} \overset{\mathcal{L}}{=} \bar{\sigma}\frac{\Gamma_{1,N}}{\frac{2}{N}\bar{\sigma}(\beta_1 + \frac{1}{2}\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku_\sigma^{(k)} - y)\big\|_{\mathbb{R}^N}^2)} \tag{4.3.4}$$

where $\Gamma_{1,N} \sim \Gamma(\alpha_1 + \frac{N}{2}, \frac{N}{2})$ is independent of $\boldsymbol{y_N}$ and $u_\sigma^{(k)}$.

In the next lemma we expand the norm $\big\|\mathcal{C}_1^{-\frac{1}{2}}(\boldsymbol{y_N} - Ku_\sigma^{(k)})\big\|_{\mathbb{R}^N}^2$ in (4.2.5) for drawing $\sigma_N^{(k+1)}$ given $\sigma_N^{(k)} = \sigma$ and $u_\sigma^{(k)}$. This expansion forms the basis for the analysis of the σ-chain in the linear inverse problem case and in particular for the proof of Theorem 4.2.5.

Lemma 4.3.4. *Under Assumptions 4.2.1 and 4.2.4, for any $\bar{\sigma}, \sigma, \delta > 0$, we have*

$$\beta_1 + \frac{1}{2}\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku_\sigma^{(k)} - \boldsymbol{y_N})\big\|_{\mathbb{R}^N}^2 = \frac{N}{\bar{\sigma}} + \frac{\sqrt{2N}}{\bar{\sigma}}Z_{1,N} + H_N(\sigma, \bar{\sigma}), \tag{4.3.5}$$

where i) $Z_{1,N}$ only depends on the white noise ξ in (4.2.6), has mean zero and variance one, higher order moments which are bounded uniformly in N, and converges weakly to a standard normal random variable as $N \to \infty$; ii) $H_N(\sigma, \bar{\sigma})$ depends on both the white noise ζ in (4.2.4) and ξ, and has bounded moments of all positive orders uniformly in N (where the expectation is taken with respect to ζ and ξ).

The proof of Lemma 4.3.4 is contained in Subsection 4.9.2. Similar to Subsection 4.3.2 we define

$$Z_{2,N} = \frac{\Gamma_{1,N} - 1 - \frac{2\alpha_1}{N}}{\sqrt{\frac{2}{N} + \frac{4\alpha_1}{N^2}}},$$

and we have

$$\Gamma_{1,N} = 1 + \frac{2\alpha_1}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_1}{N^2}} Z_{2,N},$$

where for every N the random variable $Z_{2,N}$ has mean zero and variance one, third and fourth moments which are bounded uniformly in N (see Lemma 4.9.6), and is independent of ξ and ζ. Concatenating we get

$$\sigma_N^{(k+1)} \stackrel{\mathcal{L}}{=} \bar{\sigma} \frac{1 + \frac{2\alpha_1}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_1}{N^2}} Z_{2,N}}{1 + \sqrt{\frac{2}{N}} Z_{1,N} + \frac{2}{N} H_N(\sigma_N^{(k)}, \bar{\sigma})\bar{\sigma}}, \tag{4.3.6}$$

and we are now ready to prove Theorem 4.2.5:

Proof. As in the proof of Theorem 4.2.2, the independence of $Z_{2,N}$ and ξ, ζ, and the fact that $\mathbb{E}[Z_{2,N}] = 0$ give

$$\mathbb{E}[\sigma_N^{(k+1)} - \bar{\sigma} | \sigma_N^{(k)} = \sigma] = \bar{\sigma} \mathbb{E}^{\xi,\zeta} \left[\frac{\frac{2\alpha_1}{N} - \sqrt{\frac{2}{N}} Z_{1,N} - \frac{2H_N\bar{\sigma}}{N}}{1 + \sqrt{\frac{2}{N}} Z_{1,N} + \frac{2H_N\bar{\sigma}}{N}} \right]$$

and using the identity $\frac{1}{1+x} = 1 - x + \frac{x^2}{1+x}$ we have

$$\begin{aligned}
&\mathbb{E}[\sigma_N^{(k+1)} - \bar{\sigma} | \sigma_N^{(k)} = \sigma] \\
&= \bar{\sigma} \mathbb{E}^{\xi,\zeta} \left[\left(\frac{2(\alpha_1 - H_N\bar{\sigma})}{N} - \sqrt{\frac{2}{N}} Z_{1,N} \right) \left(1 - \sqrt{\frac{2}{N}} Z_{1,N} - \frac{2H_N\bar{\sigma}}{N} \right) \right] + \mathbb{E}^{\xi,\zeta}[d_{1,N}],
\end{aligned}$$

where

$$d_{1,N} = \bar{\sigma} \frac{\left(\frac{2(\alpha_1 - H_N\bar{\sigma})}{N} - \sqrt{\frac{2}{N}} Z_{1,N} \right) \left(\frac{2Z_{1,N}^2}{N} + \frac{4H_N^2\bar{\sigma}^2}{N^2} + \frac{4\sqrt{2}H_N Z_{1,N}\bar{\sigma}}{N^{\frac{3}{2}}} \right)}{1 + \sqrt{\frac{2}{N}} Z_{1,N} + \frac{2H_N\bar{\sigma}}{N}}.$$

Using Hölder's inequality and the fact that H_N and $Z_{1,N}$ have bounded

moments of all positive orders uniformly in N, we get that

$$\mathbb{E}[\sigma_N^{(k+1)} - \bar{\sigma}|\sigma_N^{(k)} = \sigma] = \frac{2}{N}\left((\alpha_1+1)\bar{\sigma} - \mathbb{E}^{\xi,\zeta}[H_N]\bar{\sigma}^2\right) + \mathcal{O}(N^{-\frac{3}{2}}) + \mathbb{E}^{\xi,\zeta}[d_{1,N}],$$

where in a similar way to the proof of Theorem 4.2.2, the error term $d_{1,N}$, can be shown using Cauchy-Schwarz and Lemma 4.9.2 to be of order $\mathcal{O}(N^{-\frac{3}{2}})$.

For the expected diffusion, we have

$$\begin{aligned}
&\mathbb{E}\left[\left(\sigma_N^{(k+1)} - \bar{\sigma} - \mathbb{E}\left[\sigma_N^{(k+1)} - \bar{\sigma}|\sigma_N^{(k)} = \bar{\sigma}\right]\right)^2 |\sigma_N^{(k)} = \bar{\sigma}\right] \\
=&\mathbb{E}\left[(\sigma_N^{(k+1)} - \bar{\sigma})^2|\sigma_N^{(k)} = \bar{\sigma}\right] - \mathbb{E}\left[\sigma_N^{(k+1)} - \bar{\sigma}|\sigma_N^{(k)} = \bar{\sigma}\right]^2
\end{aligned}$$

where by the first part of the proof the second term is $\mathcal{O}(N^{-2})$. Thus, we only consider the first term, which will be shown to be $\mathcal{O}(N^{-1})$. Indeed, for the first term by equation (4.3.6) we have since $Z_{2,N}$ is independent of $Z_{1,N}$ and H_N and has mean zero,

$$\begin{aligned}
&\mathbb{E}\left[(\sigma_N^{(k+1)} - \bar{\sigma})^2|\sigma_N^{(k)} = \bar{\sigma}\right] \\
&= \bar{\sigma}^2\mathbb{E}\left[\left(\frac{\frac{2\alpha_1}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_1}{N^2}}Z_{2,N} - \sqrt{\frac{2}{N}}Z_{1,N} - \frac{2H_N\bar{\sigma}}{N}}{1 + \sqrt{\frac{2}{N}}Z_{1,N} + \frac{2H_N\bar{\sigma}}{N}}\right)^2\right] \\
&= \bar{\sigma}^2\mathbb{E}^{\xi,\zeta}\left[\frac{\frac{2Z_{2,N}^2}{N} + \frac{2Z_{1,N}^2}{N} + \frac{U_N}{N^{\frac{3}{2}}}}{\left(1 + \sqrt{\frac{2}{N}}Z_{1,N} + \frac{2H_N\bar{\sigma}}{N}\right)^2}\right],
\end{aligned}$$

where the random variable U_N depends on $Z_{1,N}$ and H_N and has moments of every positive order which are bounded uniformly in N (the dependence on $Z_{2,N}$ disappears by independence and the fact that $Z_{2,N}$ is mean zero and variance one). Using the identity $\frac{1}{(1+x)^2} = 1 - 2x + \frac{3x^2+2x^3}{(1+x)^2}$, we get

$$\begin{aligned}
&\mathbb{E}\left[(\sigma_N^{(k+1)} - \bar{\sigma})^2|\sigma_N^{(k)} = \bar{\sigma}\right] \\
=&\bar{\sigma}^2\mathbb{E}\left[\left(\frac{2Z_{2,N}^2}{N} + \frac{2Z_{1,N}^2}{N} + \frac{U_N}{N^{\frac{3}{2}}}\right)\left(1 - 2\sqrt{\frac{2}{N}}Z_{1,N} - \frac{4H_N\bar{\sigma}}{N}\right)\right] + \mathbb{E}[d_{2,N}],
\end{aligned}$$

where

$$d_{2,N}$$
$$=\bar{\sigma}^2\left(\frac{2Z_{2,N}^2}{N}+\frac{2Z_{1,N}^2}{N}+\frac{U_N}{N^{\frac{3}{2}}}\right)\frac{3\left(\sqrt{\frac{2}{N}}Z_{1,N}+\frac{2H_N\bar{\sigma}}{N}\right)^2+2\left(\sqrt{\frac{2}{N}}Z_{1,N}+\frac{2H_N\bar{\sigma}}{N}\right)^3}{\left(1+\sqrt{\frac{2}{N}}Z_{1,N}+\frac{2H_N\bar{\sigma}}{N}\right)^2}.$$

Using the fact that $Z_{1,N}, H_N$ and U_N have bounded moments of all positive orders, by Hölder inequality and independence (we do not need to consider higher order moments for $Z_{2,N}$ because it is independent with $Z_{1,N}$ and H_N, hence bounding terms involving $Z_{2,N}$ does not require the use of Hölder's inequality which needs higher moments), and since $\mathbb{E}[Z_{1,N}^2]=\mathbb{E}[Z_{2,N}^2]=1$, we get that

$$\mathbb{E}[(\sigma_N^{(k+1)}-\bar{\sigma})^2|\sigma_N^{(k)}=\sigma]=\frac{4\bar{\sigma}^2}{N}+\mathcal{O}(N^{-\frac{3}{2}})+\mathbb{E}[d_{2,N}].$$

In a similar way to the proof of Theorem 4.2.2, the error term $d_{2,N}$ can be shown using the Cauchy-Schwarz inequality and Lemma 4.9.2 to be of order $\mathcal{O}(N^{-2})$. Concatenating we get the result.

□

4.4 Reparametrization

In Subsection 4.2.1 we have stated our result on the slowing down of the δ-chain. As explained in Subsection 4.3.1 this slowing down arises due to the strong dependence between the u and δ components of the Gibbs sampler as the dimension N gets larger. Drawing intuition from [72, 60], we propose a reparametrization of the problem in which we make the unknown and the prior scaling a priori independent, thus alleviating the undesirable effects of the dependency on the mixing of the Gibbs sampler. Instead of having a prior $u|\delta\sim\mathcal{N}(0,\delta^{-1}\mathcal{C}_0)$ where $\delta\sim\text{Gamma}(\alpha_0,\beta_0)$, we write $u=\tau v$ and put Gaussian priors on $\tau\in\mathbb{R}$ and $v\in\mathbb{R}^N$ which are mutually independent. We hence have the reparametrized model

$$\boldsymbol{y_N}=\tau Kv+\eta \tag{4.4.1}$$

where $v\sim\mathcal{N}(0,\mathcal{C}_0)$, $\tau\sim\mathcal{N}(r_0,q_0^2)$ restricted to be positive and $\eta|\sigma\sim\mathcal{N}(0,\sigma^{-1}\mathcal{C}_1)$, $\sigma\sim\text{Gamma}(\alpha_1,\beta_1)$, and where v,τ and $\eta|\sigma$ are mutually independent. Then, the

log-likelihood is quadratic in v, hence

$$v|\boldsymbol{y_N}, \sigma, \tau \sim \mathcal{N}(m_{\sigma,\tau}(\boldsymbol{y_N}), \mathcal{C}_{\sigma,\tau}),$$

where by completing the square we get the formulae for the inverse covariance and mean

$$\mathcal{C}_{\sigma,\tau}^{-1} = \sigma\tau^2 K^* \mathcal{C}_1^{-1} K + \mathcal{C}_0^{-1}, \tag{4.4.2}$$

$$\mathcal{C}_{\sigma,\tau}^{-1} m_{\sigma,\tau} = \sigma\tau K^* \mathcal{C}_1^{-1} \boldsymbol{y_N}. \tag{4.4.3}$$

Furthermore, the log-likelihood is quadratic also in τ, hence $\tau|\boldsymbol{y_N}, v, \sigma \sim \mathcal{N}(r_{\sigma,v}, q_{\sigma,v}^2)$, where again by completing the square we obtain the formulae for the inverse covariance and mean

$$\frac{1}{q_{\sigma,v}^2} = \sigma \left\| \mathcal{C}_1^{-\frac{1}{2}} K v \right\|_{\mathbb{R}^N}^2 + \frac{1}{q_0^2} \tag{4.4.4}$$

and

$$\frac{r_{\sigma,v}}{q_{\sigma,v}^2} = \sigma \left\langle K^* \mathcal{C}_1^{-1} \boldsymbol{y_N}, v \right\rangle_{\mathbb{R}^N} + \frac{r_0}{q_0^2}. \tag{4.4.5}$$

Again we have the restriction that $\tau|\boldsymbol{y_N}, v, \sigma > 0$. As in the original parametrization, it is straightforward to check that we have conditional conjugacy for σ

$$\sigma|\boldsymbol{y_N}, u, \tau \sim \text{Gamma}(\alpha_1 + \frac{N}{2}, \beta_1 + \frac{1}{2}\left\| \mathcal{C}_1^{-\frac{1}{2}} (\tau K v - \boldsymbol{y_N}) \right\|_{\mathbb{R}^N}^2).$$

We again use the notation $m_{\sigma^{(k)},\tau^{(k)}} := m_{\sigma=\sigma^{(k)},\tau=\tau^{(k)}}$ and in a similar way for $\mathcal{C}_{\sigma^{(k)},\tau^{(k)}}, r_{\sigma^{(k)},v^{(k)}}, q_{\sigma^{(k)},v^{(k)}}$. Because of the conditional conjugacy, it is once more natural to use the Gibbs sampler:

Algorithm 4.

0) Initialize $\tau^{(0)}$ and $\sigma^{(0)}$, and set $k = 0$;

1) Compute $v^{(k)} \sim \mathcal{N}\big(m_{\sigma^{(k)},\tau^{(k)}}(\boldsymbol{y_N}), \mathcal{C}_{\sigma^{(k)},\tau^{(k)}}\big)$;

2) Compute $\tau \sim \mathcal{N}(r_{\sigma^{(k)},v^{(k)}}, q_{\sigma^{(k)},v^{(k)}}^2)$.
Reject if $\tau \leq 0$ otherwise $\tau^{(k+1)} = \tau$;

3) Compute $\sigma^{(k+1)} \sim \text{Gamma}(\alpha_1 + \frac{N}{2}, \beta_1 + \frac{1}{2}\left\| \mathcal{C}_1^{-\frac{1}{2}} (\tau^{(k)} K v^{(k)} - \boldsymbol{y_N}) \right\|_{\mathbb{R}^N}^2)$;

4) Set $k = k + 1$. If $k < k_{max}$ return to step 1, otherwise stop.

As our numerical results show, this reparametrization is indeed robust with respect to the increase in dimension (see Section 4.7), however, the τ-chain slows down in the small noise limit. This is because even though v and τ are a priori independent, they both need to explain the data hence they are a posteriori dependent. If the noise is small, this dependence becomes stronger since τKv needs to be very close to the observed data $\boldsymbol{y_N}$, hence τ and v concentrate near a lower dimensional manifold and the chain mixes poorly. Hence we need new ideas to study the Bayesian linear inverse problem with hyper-parameters in the small observational noise limit.

4.5 Extensions

We now present two important generalizations of the theory, arising from the form of the inverse problem considered. The first concerns nonlinear inverse problems where the likelihood is still conjugate to a Gaussian prior, resulting in a posterior distribution which is also Gaussian; however, in contrast to the linear inverse problem, the posterior covariance can depend on the observed data y. This setting occurs, for example, in the problem of nonparametric drift estimation in diffusions. Here only the hierarchical parameter δ is relevant and we study the Gibbs sampler arising for sampling (u, δ). The second generalization concerns problems where the data is discretized at a different level from the unknown function, and indeed the case where the data is of a fixed finite dimension.

4.5.1 General conjugate-Gaussian setting

Our analysis of the behaviour of δ, generalizes to nonlinear inverse problems in the separable Hilbert space $\mathcal{X}$, with Gaussian priors $u \sim \mathcal{N}(0, \delta^{-1}\mathcal{C}_0)$ and conditionally Gaussian posterior. We formalize this setup by assuming that almost surely with respect to the data

$$u|y \sim \mathcal{N}(m_\delta(y), \mathcal{C}_\delta(y)), \tag{4.5.1}$$

where the inverse covariance operator and mean have the general form

$$\mathcal{C}_\delta^{-1}(y) = R^*(y)R(y) + \delta\mathcal{C}_0^{-1}, \tag{4.5.2}$$

$$\mathcal{C}_\delta^{-1}(y)m_\delta(y) = l(y), \tag{4.5.3}$$

where $l(y) \in \mathcal{D}(\mathcal{C}_0^{\frac{1}{2}})$ and $R(y) : \mathcal{D}(R(y)) \to \mathcal{X}$ is a possibly unbounded linear operator and both $l(y)$ and $R(y)$ are independent of δ. Such models arise for example

in the nonparametric drift estimation of SDE's, [64, 59] (see Section 4.6).

As in the linear inverse problems case, in practice we discretize the unknown function and we implement the Bayesian setup in $\mathbb{R}^N$. We assume that we have a way of computing discretizations $\boldsymbol{y_N} \in \mathbb{R}^N$ of the observation y, and we replace the positive definite operator $\mathcal{C}_0$ by an $N \times N$ positive definite matrix also denoted by $\mathcal{C}_0$, the operator $R(y)$ by an $N \times N$ matrix denoted by $R(\boldsymbol{y_N})$ and the functional $l(y)$ by a vector in $\mathbb{R}^N$ denoted by $l(\boldsymbol{y_N})$. As before, we consider consistent discretizations, so that δ^{-1} retains its meaning as the scaling of the prior at different discretization levels N. We consider the same Gamma hyper-prior on δ, and again we have conditional conjugacy and in particular $\delta|\boldsymbol{y_N}, u$ is given by (4.1.6). We can thus sample the full posterior on u and δ given the data using a Gibbs sampler which is essentially identical to Algorithm 2 in Subsection 4.2.1:

Algorithm 5.

0. Initialize $\delta^{(0)}$ and set $k = 0$;
1. Compute $u^{(k)} \sim \mathcal{N}\big(m_{\delta^{(k)}}(\boldsymbol{y_N}), \mathcal{C}_{\delta^{(k)}}(\boldsymbol{y_N})\big)$;
2. Compute $\delta^{(k+1)} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}} u^{(k)}\big\|^2_{\mathbb{R}^N})$;
3. Set $k = k + 1$. If $k < k_{max}$ return to step 1, otherwise stop.

As before, we denote by $\delta_N^{(k+1)}$ the δ-draw in the $k+1$ iteration of the Gibbs sampler where the problem is discretized in $\mathbb{R}^N$. This draw is made using the previous draw of $u|\boldsymbol{y_N}, \delta$, which assuming that $\delta_N^{(k)} = \delta$, is denoted by $u_\delta^{(k)}$ and can be written as

$$u_\delta^{(k)} = m_\delta(\boldsymbol{y_N}) + \mathcal{C}_\delta^{\frac{1}{2}}(\boldsymbol{y_N})\zeta, \tag{4.5.4}$$

where ζ is an N-dimensional Gaussian white noise representing the fluctuation in step 1, and $\mathcal{C}_\delta, m_\delta$ are given by the formulae (4.5.3), (4.5.2) respectively. Hence we have

$$\delta_N^{(k+1)} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\big\|^2_{\mathbb{R}^N}). \tag{4.5.5}$$

The following assumptions on the discrete level generalize Assumptions 4.2.1:

Assumptions 4.5.1 (δ-chain)**.** *For almost all data y:*

i) for any $\delta > 0$ there exists a constant $c_1' = c_1'(y;\delta)$, independent of N, such that $\left\|\mathcal{C}_0^{-\frac{1}{2}} m_\delta(\boldsymbol{y_N})\right\|_{\mathbb{R}^N} \leq c_1'$;

ii) there exists a constant $c_2' = c_2'(y)$, independent of N, such that $\mathrm{Tr}(R(\boldsymbol{y_N})\mathcal{C}_0 R(\boldsymbol{y_N})^*) \leq c_2'$.

We again expect that for a reasonable discretization scheme, Assumptions 4.5.1 are inherited from the following infinite dimensional assumptions which generalize Assumptions 4.3.1.

Assumptions 4.5.2 (Infinite dimensional analogue of Assumptions 4.5.1)**.** *Almost surely with respect to the data y the following two assumptions hold:*

i) for any $\delta > 0$, we have $m_\delta(y) \in \mathcal{D}(\mathcal{C}_0^{-\frac{1}{2}})$; that is, the posterior mean belongs to the Cameron-Martin space of the prior on $u|\delta$;

ii) $R(y)\mathcal{C}_0 R(y)^*$ is trace-class; that is, the prior is sufficiently regularizing.

As in Subsection 4.3.1, it is a straightforward exercise to check that Assumption 4.5.2(ii) above implies the first and third assumptions of the Feldman-Hajek theorem [18, Theorem 2.23] on the equivalence of two Gaussian measures and in particular the conditional posterior and the prior on u, hence, together with Assumption 4.5.2(i), they imply that the conditional posterior on u is y-almost surely absolutely continuous with respect to the prior. This suggests thats our intuition on the behaviour of the δ-chain contained in Subsection 4.3.1 is still valid, since it relies on the absolute continuity of the posterior with respect to the prior.

Indeed, on the discrete level, Assumptions 4.5.1 enable us to show that the squared norm appearing in (4.5.5) behaves like $\delta^{-1}N$, hence combining with the scaling property of the Gamma distribution we can show the following result which is the generalization of Theorem 4.2.2:

Theorem 4.5.3. *Under Assumptions 4.5.1, in the limit $N \to \infty$ we have almost surely with respect to y:*

i) the expected step in the δ-chain scales like $\frac{2}{N}$, that is, for any $\delta > 0$,

$$\frac{N}{2}\mathbb{E}\left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta\right] = (\alpha_0 + 1)\delta - \tilde{f}_N(\delta; \boldsymbol{y_N})\delta^2 + \mathcal{O}(N^{-\frac{1}{2}}),$$

where $\tilde{f}_N(\delta; \boldsymbol{y_N})$ is bounded uniformly in N. In particular, if there exists $\tilde{f}(\delta; y) \in \mathbb{R}$ such that $\tilde{f}_N(\delta; \boldsymbol{y_N}) \to \tilde{f}(\delta; y)$ then

$$\frac{N}{2}\mathbb{E}\left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta\right] = (\alpha_0 + 1)\delta - \tilde{f}(\delta; y)\delta^2 + \mathcal{o}(1);$$

ii) the variance of the step also scales like $\frac{2}{N}$ and in particular for any $\delta > 0$,

$$\frac{N}{2}\mathbb{E}\left[\left(\delta_N^{(k+1)} - \delta_N^{(k)} - \mathbb{E}\left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta\right]\right)^2 | \delta_N^{(k)} = \delta\right] = 2\delta^2 + \mathcal{O}(N^{-\frac{1}{2}}).$$

All the expectations are taken with respect to the randomness in the algorithm.

Remark 4.5.4.

i) The proof of Theorem 4.5.3 proceeds exactly as the proof of Theorem 4.2.2 in Subsection 4.3.2, the only difference being that instead of Lemma 4.3.3 and Lemma 4.9.1, we now use Lemma 4.9.9 and Lemma 4.9.10 to analyse the behaviour of the rate parameter in the δ-draw (4.5.5). The statements and proofs of the Lemmas 4.9.9 and 4.9.10 are in fact essentially identical to those of Lemmas 4.3.3 and 4.9.1 respectively, the only difference being that instead of Assumptions 4.2.1 we now use Assumptions 4.5.1.

ii) $\tilde{f}_N(\delta; \boldsymbol{y_N}) := \mathbb{E}^{\zeta}[\tilde{F}_N(\delta; \boldsymbol{y_N})]$, where $\tilde{F}_N$ is a problem specific term defined in Lemma 4.9.9.

iii) Parts (iii)-(iv) of Remark 4.2.3 are also valid here with obvious adjustments.

iv) In Section 4.6 we show that the nonlinear Bayesian inverse problem of non-parametric drift estimation of SDE's studied in [64, 59], satisfies Assumptions 4.5.2 on the underlying infinite dimensional model.

v) Since the source of the slowing down of the δ-chain is again the strong dependence between the two components of Algorithm 5, we can reparametrize as in Section 4.4 by making the unknown and the scaling a priori independent. We write $u = \tau v$ and put Gaussian priors on $\tau \in \mathbb{R}$ and $v \in \mathbb{R}^N$ which are mutually independent. It is straightforward to check that we again have conditional conjugacy, making the use of a Gibbs sampler natural. As in Section 4.4, we expect the resulting algorithm to be robust with respect to model-refinement but to deteriorate as the quality of the data improves.

4.5.2 Differing discretization levels of data and unknown

It is often of interest to consider situations where the observation y and the unknown parameter u live in possibly different Hilbert spaces $\mathcal{X}_1$ and $\mathcal{X}_2$ and have different discretization levels M and N respectively. In the linear inverse problem case this

means that the forward operator K maps $\mathcal{X}_1$ to $\mathcal{X}_2$ and its discretization is now an $M \times N$ matrix. In the general conjugate setting considered in Subsection 4.5.1, it is also reasonable to have that the possibly unbounded operator R maps $\mathcal{X}_1$ to another Hilbert space $\mathcal{X}_3$ and its discretization is an $L \times N$ matrix. An inspection of our proofs suggests that: i) Theorems 4.2.2 and 4.5.3 on the slowing-down of the δ-chain generalize to the case where the dimension N of the discretization of u goes to infinity, for any values of M and in the general conjugate setting L; ii) Theorem 4.2.5 on the speeding-up of the σ-chain in the linear inverse problems setting generalizes to the case where the dimension M of the discretization of y goes to infinity, for any value of N. The only difference is that f_N in Theorem 4.2.2 and h_N in Theorem 4.2.5, now also depend on M and are uniformly bounded in both N and M, and $\tilde{f}_N$ in Theorem 4.5.3 now also depends on M and L and is uniformly bounded in N, M and L.

One should note that it is possible to have an unknown parameter u which is a field, but an observation y which is of a fixed finite dimension, that is $\mathcal{X}_2 = \mathbb{R}^M$. According to the considerations above, our results on the slowing down of the δ-chain as we refine the discretization of u, that is as $N \to \infty$, also generalize to this case. In particular, in Section 4.7 we show numerical results demonstrating the slowing down of the δ-chain as N grows for M fixed, in a setup where we discretize the unknown parameter u on a grid of N points and observe a blurred version of it at only M of these points subject to additive white noise.

4.6 Examples satisfying underlying infinite dimensional model assumptions

In this section we present several examples satisfying our assumptions on the underlying infinite dimensional Bayesian inverse problem.

We first present three instances of linear inverse problems satisfying Assumptions 4.3.1 and 4.3.2: a family of mildly ill-posed linear inverse problems, where the operators defining the problem are simultaneously diagonalizable, [44]; a family of severely ill-posed inverse problems again in a diagonal setting, [45, 4]*; a family of mildly ill-posed inverse problems in a nondiagonal setting, [3]*. We expect that Assumptions 4.2.1 and 4.2.4, will be satisfied by consistent (in the numerical analysis sense) discretizations of these infinite dimensional models. Indeed, we show that our discrete level assumptions are satisfied if we discretize the two diagonal examples using spectral truncation. Furthermore, in Section 4.7 we provide numerical

*[3, 4] are Chapters 2 and 3 respectively in the current thesis.

evidence that our ideas also apply when using other discretization schemes, and in particular discretization via finite difference approximations.

We also present a nonlinear inverse problem satisfying our more general assumptions in Section 4.5. In particular, we show that the inverse problem of nonparametric drift estimation of SDE's using local time, considered in [64, 59], satisfies Assumptions 4.5.2. Again we expect that Assumptions 4.5.1 on the discrete level, will be satisfied by consistent discretizations.

4.6.1 Linear mildly ill-posed simultaneously diagonalizable inverse problem

We consider the linear inverse problem setting (4.1.1)-(4.1.3), where $K, \mathcal{C}_0$ and $\mathcal{C}_1$ commute with each other and $K^*K, \mathcal{C}_0$ and $\mathcal{C}_1$ are simultaneously diagonalizable with common complete orthonormal eigenbasis $\{e_j\}_{j\in\mathbb{N}}$. Note, that we do not assume that K is compact, but we do assume that $K^\star K$ is diagonalizable in $\{e_j\}_{j\in\mathbb{N}}$; in particular, we allow for K to be the identity. For any $w \in \mathcal{X}$, let $w_j := \langle w, e_j \rangle$. Let Σ be a positive definite and trace class operator in $\mathcal{X}$ which is diagonalizable in the orthonormal basis $\{e_j\}_{j\in\mathbb{N}}$, with eigenvalues $\{\lambda_j^\Sigma\}_{j\in\mathbb{N}}$. Then for any $\rho \in \mathcal{X}$, we can write a draw $x \sim \mathcal{N}(\rho, \Sigma)$ as

$$x = \rho + \sum_{j=1}^{\infty} \sqrt{\lambda_j^\Sigma}\gamma_j e_j,$$

where γ_j are independent standard normal random variables in $\mathbb{R}$; this is the Karhunen-Loeve expansion [1]. This expansion suggests that since we are in a simultaneously diagonalizable setting we can use the Parseval identity and work entirely in the frequency domain. Indeed, we identify an element $w \in \mathcal{X}$ with the square summable sequence of its coefficients $\{w_j\}_{j\in\mathbb{N}}$, and the norm and inner product in $\mathcal{X}$ with the ℓ^2-norm and inner product. Furthermore, we identify the operators $\mathcal{C}_0, \mathcal{C}_1$ and K with the sequences of their eigenvalues $\{\lambda_j^{\mathcal{C}_0}\}_{j\in\mathbb{N}}, \{\lambda_j^{\mathcal{C}_1}\}_{j\in\mathbb{N}}$ and $\{\lambda_j^K\}_{j\in\mathbb{N}}$ respectively. Algebraic operations on the operators $\mathcal{C}_0, \mathcal{C}_1, K$ are defined through the corresponding operations on the respective sequences.

We make the following assumptions on the decay of the spectrum of $K, \mathcal{C}_0$ and $\mathcal{C}_1$:

Assumptions 4.6.1. *The eigenvalues of $K^*K, \mathcal{C}_0$ and $\mathcal{C}_1$, denoted by $(\lambda_j^K)^2, \lambda_j^{\mathcal{C}_0}, \lambda_j^{\mathcal{C}_1}$, respectively, satisfy*

- $(\lambda_j^K)^2 \asymp j^{-4\ell}$, $\ell \geq 0$;

- $\lambda_j^{\mathcal{C}_0} \asymp j^{-2\alpha},\ \alpha > \frac{1}{2}$*;

- $\lambda_j^{\mathcal{C}_1} \asymp j^{-2\beta},\ \beta \geq 0$*.

Let ν be the joint distribution of y and u, where $u|\delta \sim \mathcal{N}(0, \delta^{-1}\mathcal{C}_0)$ and $y|u, \sigma, \delta \sim \mathcal{N}(Ku, \sigma^{-1}\mathcal{C}_1)$. Then in this diagonal case, it is straightforward to show in the infinite dimensional setting that the conditional posterior $u|y, \sigma, \delta$ is ν-almost surely Gaussian, $\mathcal{N}(m_{\sigma,\delta}(y), \mathcal{C}_{\sigma,\delta})$, where $\mathcal{C}_{\sigma,\delta}$ and $m_{\sigma,\delta}(y)$ satisfy (4.1.4) and (4.1.5) respectively. We make the following additional assumption:

Assumption 4.6.2. *The parameters α, β, ℓ in Assumptions 4.6.1 satisfy $2\alpha + 4\ell - 2\beta > 1$.*

We show that under Assumptions 4.6.1 and 4.6.2, the Assumptions 4.3.1 on the underlying infinite dimensional model are satisfied ν-almost surely. Without loss of generality we assume that $\delta = \sigma = 1$. For Assumption 4.3.1(i), we have using the Karhunen-Loeve expansion and Assumption 4.6.1,

$$\mathbb{E}^\nu \left\| \mathcal{C}_0^{-\frac{1}{2}} m(y) \right\|^2 \leq c\mathbb{E}^\nu \sum_{j=1}^{\infty} \frac{j^{2\alpha-4\ell+4b}}{(j^{-4\ell+2\beta} + j^{2\alpha})^2} (j^{-2\ell-\alpha}\zeta_j + j^{-\beta}\xi_j)^2,$$

where $\{\zeta_j\}_{j\in\mathbb{N}}, \{\xi_j\}_{j\in\mathbb{N}}$ are two independent sequences of independent standard normal random variables. The assumption $2\alpha + 4\ell - 2\beta > 1$ secures that the right hand side is finite, hence $m(y) \in \mathcal{D}(\mathcal{C}_0^{-\frac{1}{2}})$ ν-almost surely. For Assumption 4.3.1(ii), the operator $\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0 K^* \mathcal{C}_1^{-\frac{1}{2}}$ has eigenvalues which decay like $j^{-2\alpha-4\ell+2\beta}$ hence are summable by Assumption 4.6.2.

We consider the additional Assumption 4.3.2 on the data required for the analysis of the σ-chain. We first define the Sobolev-like spaces $\mathcal{H}^t, t \in \mathbb{R}$: for $t \geq 0$, we define

$$\mathcal{H}^t := \{u \in \mathcal{X} : \|u\|_{\mathcal{H}^t} := \sum_{j=1}^{\infty} j^{2t} \langle u_j, e_j \rangle^2 < \infty\},$$

and for $t < 0$, $\mathcal{H}^t := (\mathcal{H}^{-t})^*$.

We assume data of the form (4.2.6), where z is the image under K of an underlying, sufficiently regular, true solution $\bar{u}$:

Assumption 4.6.3. *$y = K\bar{u} + \bar{\sigma}^{-\frac{1}{2}} \mathcal{C}_1^{\frac{1}{2}} \xi$, where $\bar{u} \in \mathcal{H}^{\beta-2\ell}$ and ξ is a Gaussian white noise.*

*α, β not to be confused with $\mathfrak{a}, \mathfrak{b}$ used respectively as shape and rate parameters of the Gamma distribution.

Under Assumptions 4.6.1 and 4.6.3, we have that Assumption 4.3.2 is satisfied, since

$$\left\|\mathcal{C}_1^{-\frac{1}{2}} z\right\|^2 \leq c \sum_{j=1}^{\infty} j^{2\beta-4\ell} \bar{u}_j^2 = c\left\|\bar{u}\right\|_{\mathcal{H}^{\beta-2\ell}}^2.$$

Note that under Assumptions 4.6.1, 4.6.2 and 4.6.3, it is straightforward to check that Assumption 4.3.1(i) is also satisfied ξ-almost surely. Indeed, using the Karhunen-Loeve expansion we have,

$$\mathbb{E}\left\|\mathcal{C}_0^{-\frac{1}{2}} m(y)\right\|^2 \leq c\mathbb{E} \sum_{j=1}^{\infty} \frac{j^{2\alpha-4\ell+4b}}{(j^{-4\ell+2\beta}+j^{2\alpha})^2} (j^{-2\ell}\bar{u}_j^2 + \bar{\sigma}^{-\frac{1}{2}} j^{-\beta}\xi_j)^2,$$

where $\{\xi_j\}_{j\in\mathbb{N}}$ is a sequence of independent standard normal random variables. The assumption $2\alpha + 4\ell - 2\beta > 1$ together with the assumption that $\bar{u} \in \mathcal{H}^{\beta-2\ell}$ secure that the right hand side is finite. Assumption 4.3.1(ii) is independent of y, hence also holds by our previous considerations.

A natural way to discretize this setup is to truncate the Karhunen-Loeve expansion. We define the $N \times N$ matrices $\mathcal{C}_0, \mathcal{C}_1$ and K by identifying them with the truncated sequences $\{\lambda_1^{\mathcal{C}_0}, ..., \lambda_N^{\mathcal{C}_0}\}, \{\lambda_1^{\mathcal{C}_1}, ..., \lambda_N^{\mathcal{C}_1}\}$ and $\{\lambda_1^K, ..., \lambda_N^K\}$ respectively. As before, algebraic operations on the matrices $\mathcal{C}_0, \mathcal{C}_1$ and K are defined through the corresponding operations on the corresponding truncated sequences. Furthermore, we identify y with the sequence $\{y_j\}_{j\in\mathbb{N}}$ and its discretization $\boldsymbol{y_N}$ at level N with the truncated sequence $\{y_1, ..., y_N\}$. In general, we consider the discretization of an element $w \in \mathcal{X}$ at level N, by identifying it with the truncated sequence $\{w_1, ..., w_N\} \in \mathbb{R}^N$. The Hilbert space norm of w is then replaced by the Euclidean norm of the vector $\{w_1, ..., w_N\} \in \mathbb{R}^N$ and the Hilbert space inner product of w and v by the Euclidean inner product of the vectors $\{w_1, ..., w_N\}$ and $\{v_1, ..., v_N\}$. We consider discrete data $\boldsymbol{y_N}$ which arise by discretizing y in Assumption 4.6.3. That is, we assume

$$\boldsymbol{y_N} = K\bar{u} + \bar{\sigma}^{-\frac{1}{2}} \mathcal{C}_1^{\frac{1}{2}} \xi,$$

where $K, \mathcal{C}_1, \bar{u}$ and ξ are discretized as explained above. We show that Assumptions 4.2.1 and 4.2.4, are satisfied under Assumptions 4.6.1, 4.6.2 and 4.6.3 for data of this form and for this discretization scheme.

By Assumption 4.6.1, we have that there exists a constant $c \geq 0$ independent

of N, such that

$$\mathbb{E}\left\|\mathcal{C}_0^{-\frac{1}{2}} m(\boldsymbol{y_N})\right\|_{\mathbb{R}^N}^2 \leq c\mathbb{E}\sum_{j=1}^{N} \frac{j^{2\alpha-4\ell+4b}}{(j^{-4\ell+2\beta}+j^{2\alpha})^2}(j^{-2\ell}\bar{u}_j + j^{-\beta}\xi_j)^2,$$

where the right hand side is bounded uniformly in N, since we are summing non-negative numbers and we have seen that under Assumptions 4.6.2 and 4.6.3 the corresponding infinite series is summable. Furthermore, again by Assumption 4.6.1, there exists another constant $c \geq 0$ independent of N, such that

$$\mathrm{Tr}(\mathcal{C}_1^{-\frac{1}{2}} K\mathcal{C}_0 K^* \mathcal{C}_1^{-\frac{1}{2}}) \leq c\sum_{j=1}^{N} j^{-2\alpha-4\ell+2\beta},$$

where the right hand side is bounded uniformly in N, since we have seen that under Assumption 4.6.2 the corresponding infinite series is summable. Finally, under Assumption 4.6.1 there exists a constant $c \geq 0$ independent of N, such that

$$\left\|\mathcal{C}_1^{-\frac{1}{2}} K\bar{u}\right\|_{\mathbb{R}^N}^2 \leq c\sum_{j=1}^{N} j^{2\beta-4\ell}\bar{u}_j^2,$$

where the right hand side is bounded uniformly in N, since by Assumption 4.6.3 the corresponding infinite series is summable.

In Section 4.7 below we show numerical results for both the slowing down of the δ-chain and the speeding up of the σ-chain, in this simultaneously diagonalizable linear inverse problem setting. We first show results using the discretization via truncation of the Karhunen-Loeve expansion presented here, and then we show results using discretization via finite differences. We do not prove that discretization via finite differences satisfies our discrete level assumptions, however, we expect this to be true; our belief is supported by our numerical results.

4.6.2 Linear severely ill-posed simultaneously diagonalizable inverse problem

We consider the setting of [45, 4]*, that is, a similar situation with the previous example, where instead of having $(\lambda_j^K)^2 \asymp j^{-4\ell}$ we now have $(\lambda_j^K)^2 \asymp e^{-2sj^b}$, for $b, s > 0$. The proof of the validity of Assumptions 4.3.1 ν-almost surely is identical to the proof in the previous example, where we now have the added advantage of the exponential decay of the eigenvalues of K^*K. Furthermore, Assumption 4.3.2

*[4] is Chapter 3 in the current thesis.

holds if we consider $z = K\bar{u}$ where $\bar{u} \in \mathcal{X}$. We can again prove that for data of the form (4.2.6) where $z = K\bar{u}$, Assumption 4.3.1 is satisfied ξ-almost surely. Finally, in a similar way to the previous example, Assumptions 4.2.1 and 4.2.4 are valid if we discretize this setup by truncating the Karhunen-Loeve expansion.

4.6.3 Nondiagonal linear inverse problem

We consider the setting of [3]*, that is the linear inverse problem setting where $K^*K, \mathcal{C}_0$ and $\mathcal{C}_1$ are not necessarily simultaneously diagonalizable but they are related to each other via the norm equivalence assumptions 2.3.1. As before let ν be the joint distribution of y and u, where $u|\delta \sim \mathcal{N}(0, \delta^{-1}\mathcal{C}_0)$ and $y|u, \sigma, \delta \sim \mathcal{N}(Ku, \sigma^{-1}\mathcal{C}_1)$. Then as in the simultaneously diagonalizable case examined above, we have that the conditional posterior $u|y, \sigma, \delta$ is ν-almost surely $\mathcal{N}(m_{\sigma,\delta}(y), \mathcal{C}_{\sigma,\delta})$, where $\mathcal{C}_{\sigma,\delta}$ and $m_{\sigma,\delta}(y)$ satisfy (4.1.4) and (4.1.5) respectively (see Theorem 2.2.1). It is implicit in Theorem 2.2.1 that $m_{\sigma,\delta}(y) \in \mathcal{D}(\mathcal{C}_0^{-\frac{1}{2}})$ ν-almost surely, hence Assumption 4.3.1(i) holds ν-almost surely. Assumption 4.3.1(ii) also holds ν-almost surely since if $\{\phi_j\}_{j\in\mathbb{N}}$ is a complete orthonormal system of eigenfunctions of $\mathcal{C}_0$ and $\{\lambda_j^{\mathcal{C}_0}\}_{j\in\mathbb{N}}$ the corresponding eigenvalues, by Assumption 2.3.1(3) we have $\left\|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0^{\frac{1}{2}} \phi_j\right\|^2 \leq c\left\|\mathcal{C}_0^{-\frac{\beta}{2}+\ell+\frac{1}{2}} \phi_j\right\|^2 = c(\lambda_j^{\mathcal{C}_0})^{-\beta+2\ell+1}$ which is summable by Assumption 2.3.1(1) and (2). Hence, we have that $\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0^{\frac{1}{2}}$ is Hilbert-Schmidt which in turn implies that $\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0 K^* \mathcal{C}_1^{-\frac{1}{2}}$ is trace-class.

Furthermore, by Assumption 2.3.1(3), we have that Assumption 4.3.2 is satisfied if we consider $z = K\bar{u}$, for $\bar{u} \in X^{\beta-2\ell}$, where the spaces X^t are defined in Section 2.3 of Chapter 2.

4.6.4 Nonparametric drift estimation

We consider the setting presented in [59] and analyzed in [64]. In particular, we consider the problem of nonparametrically estimating the drift function u from an observation of a single path up to time T of the solution of the SDE

$$dy_t = u(y_t)dt + dW_t,$$

where W_t is a Brownian motion and where we assume that the path lives on the circle $\mathbb{T} = [0, 1)$. We work in $\mathcal{X} = L^2(\mathbb{T})$, and put a Gaussian prior on u, $u|\delta \sim \mathcal{N}(0, \delta^{-1}\mathcal{C}_0)$, where $\mathcal{C}_0^{-1} = (-\frac{d^2}{dx^2})^p + I$. Let $\nu(y, u)$ be the joint density of the path y and the drift u where $y|u$ is given by the Girsanov theorem and u is drawn from

*[3] is Chapter 2 in the current thesis.

the prior. In [64], it is shown that $u|y,\delta \sim \mathcal{N}(m(y),\mathcal{C}_{\sigma,\delta}(y))$ ν-almost surely, for

$$\mathcal{C}_{\sigma,\delta}(y)^{-1} = L_T(.;y) + \delta\mathcal{C}_0^{-1},$$

$$\mathcal{C}_{\sigma,\delta}(y)^{-1}m(y) = \frac{1}{2}L_T(.;y)' + \chi_T(.;y),$$

where L_T is the semimartingale local time and χ_T is a function which is of $\mathcal{O}(1)$ with respect to T. In particular, we are in the setting described in (4.5.1)-(4.5.2) where $R(y)$ is the multiplication operator $L_T^{\frac{1}{2}}(.;y)$. We check the validity of Assumptions 4.5.2. Indeed, it is shown in [64] that $m(y) \in \mathcal{D}(\mathcal{C}_0^{-\frac{1}{2}})$ ν-almost surely hence Assumption 4.5.2(i) holds ν-almost surely. On the other hand for Assumption 4.5.2(ii), the Hölder inequality implies that the operator $R(y)$ is bounded in $\mathcal{X}$ ν-almost surely, hence, since the family of trace class operators is a $*$-ideal in the space of bounded operators in $\mathcal{X}$ [66, Theorem VI.19], and since $\mathcal{C}_0$ is trace class, we have that $R(y)\mathcal{C}_0R^*(y)$ is also trace class ν-almost surely.

4.7 Numerical Results

In this section we present some numerical simulations supporting our main results contained in Section 4.2, their extensions in Section 4.5 and our intuition on the benefits of the reparametrization introduced in Section 4.4.

We consider the mildly ill-posed diagonal setting presented in Subsection 4.6.1 and examine three instances of the linear inverse problem of determining an unknown function $u \in \mathcal{X} = L^2[0,1]$ from a noisy observation y of a linear transformation of the signal. In Subsection 4.7.1, we discretize using the truncated Karhunen-Loeve expansion as explained in Subsection 4.6.1 and work in the frequency domain. This discretization scheme satisfies our assumptions (see Subsection 4.6.1), and indeed our theory is supported by the numerical results. In Subsection 4.7.2, we discretize the unknown and data on a uniform grid of N points in $[0,1]$, and use finite differences to approximate the operators $K, \mathcal{C}_0$ and $\mathcal{C}_1$ which are assumed to be inverses of differential operators. We do not check the validity of Assumptions 4.2.1 and 4.2.4 in this case, however, the simulation results are again consistent with our theory. Finally, in Subsection 4.7.3, we consider a similar but in fact nondiagonal setting, where again we discretize the unknown on a uniform grid of N points but we assume that we only observe $M < N$ of these points subject to additive white noise. As indicated in Section 4.5 our theory extends to such cases and indeed this is supported by our numerical results.

In the first two instances, in Subsections 4.7.1 and 4.7.2, we consider hier-

archical inference on both the prior and noise scaling parameters, despite the fact that for simplicity our analysis presented in Section 4.2 is performed separately on the behaviour of the δ-chain for fixed σ and on the behaviour of the σ-chain for fixed δ. The simulation results when fixing one of the two hyper-parameters, are qualitatively the same as the ones presented here, indicating that since σ instantly identifies its true value, when considering hierarchical inference simultaneously on both hyper-parameters, the two components of the joint chain (δ, σ) decorrelate instantly.

4.7.1 Signal in white noise model using truncated Karhunen-Loeve expansion

We consider the simultaneously diagonalizable setup described in Subsection 4.6.1, where $\mathcal{X} = L^2[0,1]$ with Dirichlet boundary conditions. We consider the orthonormal basis $e_j(x) = \sqrt{2}\sin(j\pi x),\ x \in [0,1]$, and define the operators $K, \mathcal{C}_0$ and $\mathcal{C}_1$ directly through their eigenvalues $\lambda_j^K = 1, \lambda_j^{\mathcal{C}_0} = j^{-3}$ and $\lambda_j^{\mathcal{C}_1} = 1$, for all $j \in \mathbb{N}$, respectively. In other words, we study the problem of recovering an unknown signal through a direct observation subject to white noise. Defining $\mathcal{A}_0$ to be the negative Laplace operator in $[0,1]$ with Dirichlet boundary conditions, we have that we use a Gaussian prior with covariance operator $\mathcal{C}_0$ which is proportional to $\mathcal{A}_0^{-\frac{3}{2}}$.

In the language of Subsection 4.6.1, we have that Assumptions 4.6.1 are satisfied with $\alpha = 1.5$ and $\beta = \ell = 0$, hence since $2\alpha + 4\ell - 2\beta = 3 > 1$, Assumption 4.6.2 is also satisfied. We assume that we have data produced from the underlying true signal $\bar{u}(x) = \sum_{j=1}^{\infty} \bar{u}_j \sqrt{2}\sin(j\pi x)$, for $x \in [0,1]$, where $\bar{u}_j = j^{-2.25}\sin(10j)$ and $\bar{\sigma} = 200$, and in particular we have that the coefficients of y, are given as

$$y_j = \bar{u}_j + \bar{\sigma}^{-\frac{1}{2}}\xi_j,$$

where ξ_j are standard normal random variables. It is straightforward to check that $\bar{u} \in \mathcal{H}^t$ for any $t < 1.75$, hence we have that Assumption 4.6.3 is also satisfied. According to the considerations in Subsection 4.6.1, we thus have that Assumptions 4.2.1 and 4.2.4 hold when using the truncated Karhunen-Loeve expansion discretization method.

This particular example is studied in [79] using both a hierarchical Bayesian approach and an empirical Bayes approach to perform simultaneous inference on the scaling of the prior covariance and the unknown function. We use the hierarchical setup presented in Section 4.1 and in particular implement Algorithm 1. We also implement the reparametrization presented in Section 4.4 and in particular

Algorithm 4.

In the standard hierarchical algorithm, the values of the shape and rate parameters in the two Gamma hyper-priors on δ and σ, are chosen to be $\alpha_0 = \alpha_1 = 1$ and $\beta_0 = \beta_1 = 10^{-4}$ respectively. We do this in order to have uninformative hyper-priors, since this choice gives that the mean is equal to $\frac{\alpha_0}{\beta_0} = \frac{\alpha_1}{\beta_1} = 10^4$ and the variance is equal to $\frac{\alpha_0}{\beta_0^2} = \frac{\alpha_1}{\beta_1^2} = 10^8$. We choose $\sigma^{(0)} = 1$ and $\delta^{(0)} = 100$. The latter choice is made in order to pronounce the slowness of the δ-chain as the discretization level increases, by starting the chain from a value which is a certain distance apart from the typical value of the chain in stationarity. In the reparametrized algorithm, in order to again have uninformative hyper-priors, we keep the same choice of α_1, β_1 in the Gamma prior on σ, and choose mean $r_0 = 1$ and variance $q_0^2 = 10^4$ in the Gaussian prior on τ. We choose $\sigma^{(0)} = 1$ and $\tau^{(0)} = 1/10$. The latter choice is made in order to be consistent with the fact that τ^2 acts like δ^{-1}. We implement the algorithms at discretization levels $N = 32, 128, 512, 2048$ and 8192 and in each case we use 10^4 iterations of the corresponding Gibbs sampler. In order to have fair comparisons, in the calculation of the sample mean and variance we use a fixed burn-in time of 1000 iterations. We take the viewpoint that we have a fixed computational budget, hence we choose not to increase the burn-in time as N increases as one can do if infinite resources are available. This has a negative effect on the quality of the reconstruction using the standard hierarchical algorithm, since as N increases the burn-in time becomes insufficient for the δ-chain to reach stationarity. This is not to suggest that the slowness of the δ-chain is only in the transition to stationarity. Our results, both numerical presented below and theoretical (Theorem 4.2.2), show that the δ-chain slows down both before and after reaching stationarity.

In Figure 4.1 we have in the left column the true solution (dashed black) and discretized noisy data (blue continuous), and in the middle and right columns the true solution (dashed black), the sample mean (red continuous) and 87.5% credibility bounds (shaded area) using the standard hierarchical algorithm and the reparametrized algorithm respectively, for dimensions $N = 32$ (top) and $N = 8192$ (bottom). We can see that the reconstruction and credibility bounds in the case of the standard hierarchical algorithm deteriorate for large dimensions, while for the reparametrized algorithm appear to be stable.

In Figure 4.2 we have the plots of the σ-chains on the left and the δ-chains on the right, in the standard algorithm, for increasing dimension as we move from top to bottom. As predicted by Theorems 4.2.2 and 4.2.5, the plots show that while in small dimensions both chains appear to have a similar behaviour with a healthy mixing, as N increases the σ-chain moves to the true value $\bar{\sigma} = 200$ and fluctuates

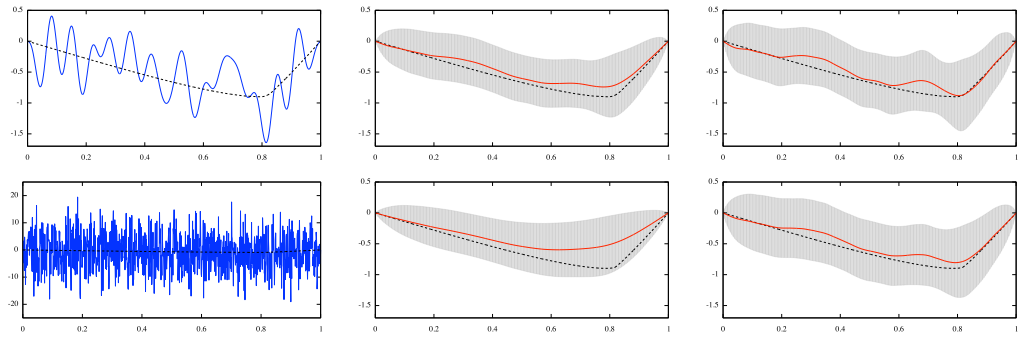

Figure 4.1: Left column: true solution (dashed black) and noisy data (blue continuous). Middle and right columns: true solution (dashed black), sample mean (red continuous) and 87.5% credibility bounds (shaded area) for standard hierarchical (middle) and reparametrized algorithm (right). Dimension is $N = 32$ (top) and $N = 8192$ (bottom).

independently around it with fluctuations which decrease as N increases, while the δ-chain becomes slower and exhibits diffusive behaviour. In Figure 4.3 we have the plots of the σ-chains on the left and the τ^2-chains on the right, in the reparametrized algorithm for increasing dimension top to bottom. As expected, the σ-chain exhibits the same behaviour as in the standard algorithm, but the τ^2-chain appears to be robust with respect to the increase in dimension.

Our observations in Figures 4.2 and 4.3 are also supported by the autocorrelation plots presented in Figure 4.4. We have four panels with the plots of the autocorrelation functions for time lag $1 - 20$ of the four chains at the different discretization levels N. On the left column we have the autocorrelation functions of the σ-chains using the standard algorithm (top) and the reparametrized algorithm (bottom) which are practically the same; in both cases the rate of decay of correlations seems to increase as N increases, and indeed for $N \geq 512$ even consecutive samples are practically independent. On the right column we have the autocorrelation function of the δ-chain in the case of the standard algorithm (top) and of the τ^2-chain in the case of the reparametrized algorithm (bottom). The rate of decay of correlations in the δ-chain appears to decrease as the dimension increases, and in particular for $N = 8192$ the correlations seem not to decay at all. On the contrary, the rate of decay of correlations in the τ^2-chain does not seem to be affected by the increase in dimension.

The fact that in low dimensions the rate of decay of correlations is slower in the τ^2-chain than in the δ-chain, is due to the small noise effect explained in Section 4.4. To highlight this effect, we run the reparametrized algorithm again in

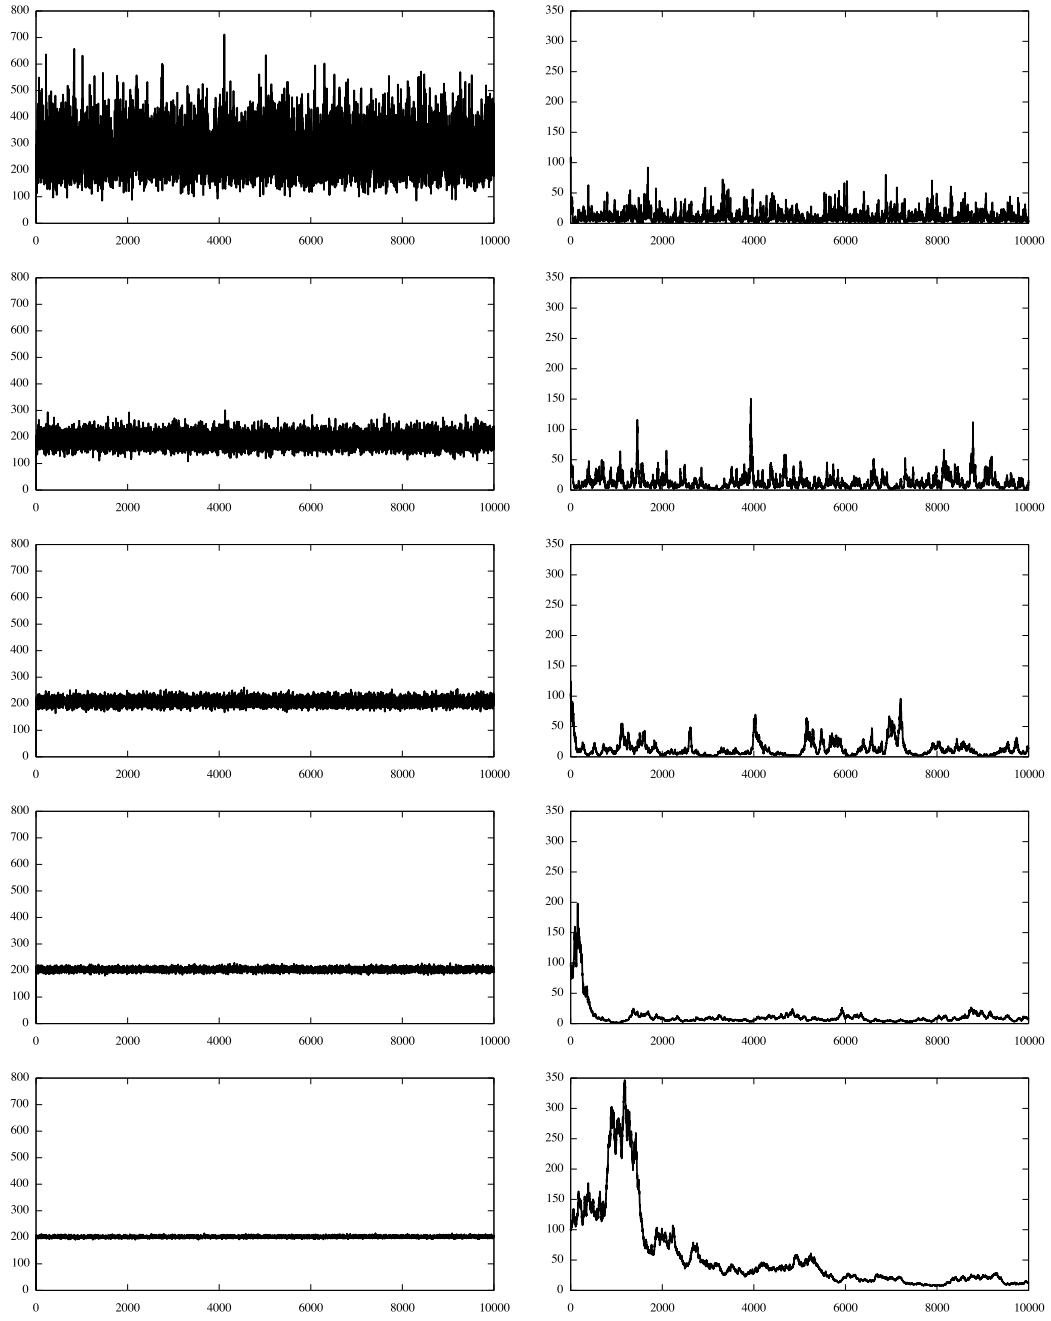

Figure 4.2: Standard algorithm: σ-chains (left column) and δ-chains (right column) for dimensions $N = 32, 128, 512, 2048$ and 8192 top to bottom.

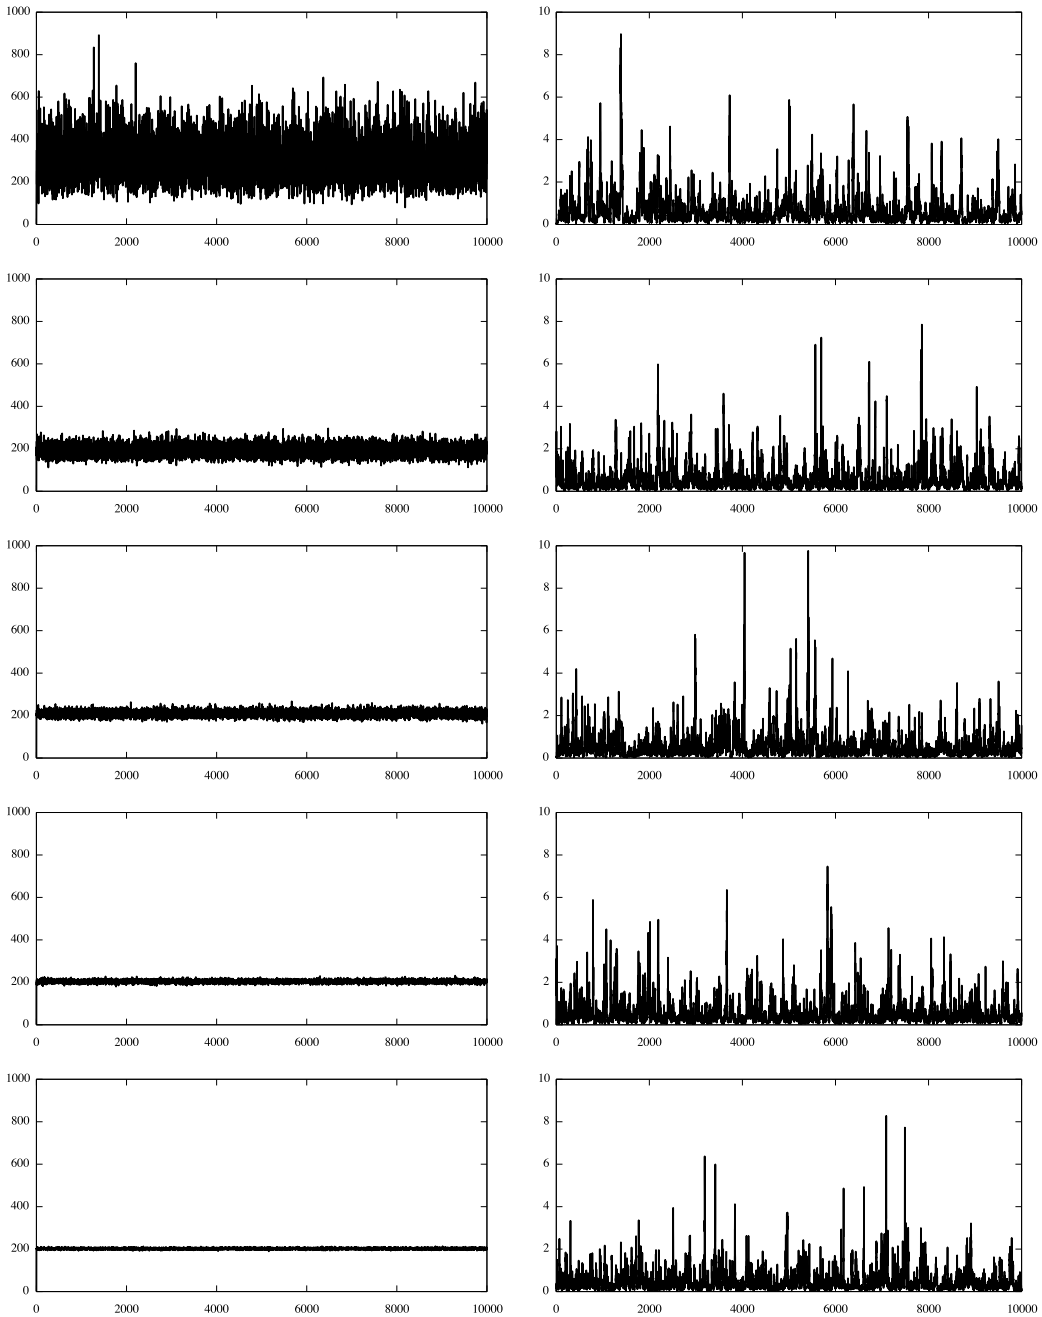

Figure 4.3: reparametrized algorithm: σ-chains (left column) and τ^2-chains (right column) for dimensions $N = 32, 128, 512, 2048$ and 8192 top to bottom.

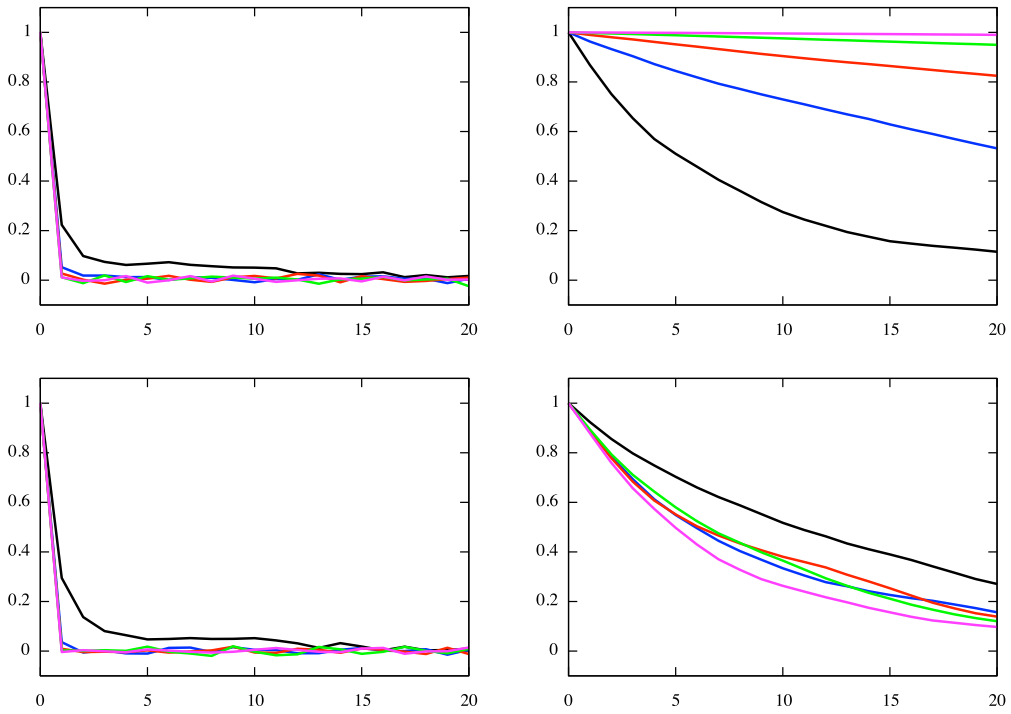

Figure 4.4: Autocorrelation functions for dimensions 32 (black), 128 (blue), 512 (red), 2048 (green) and 8192 (violet). Top left is for σ-chain in standard algorithm, top right for δ-chain in standard algorithm, bottom left for σ-chain in reparametrized algorithm and bottom right for τ^2-chain in reparametrized algorithm.

the case of a much smaller noise, namely $\bar{\sigma} = 200^2$, and plot the σ and τ^2-chains in Figure 4.5. As expected, the σ-chain exhibits the same behaviour as before, but the τ^2-chain mixes very poorly. New work is required to produce effective hierarchical algorithms in this small noise limit.

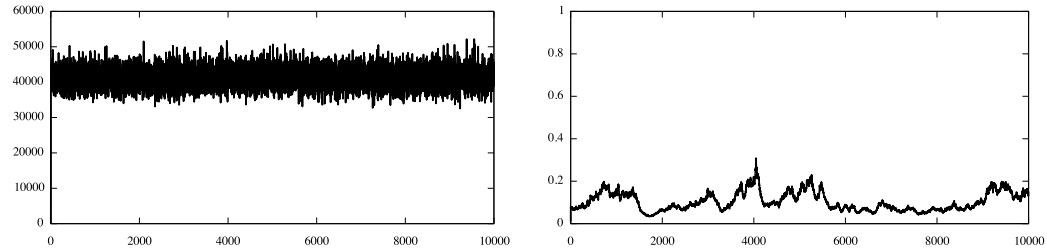

Figure 4.5: reparametrized algorithm for small noise, $\bar{\sigma} = 200^2$: σ-chain (left) and τ^2-chain (right) for dimension $N = 512$.

In conclusion, our numerical simulations support the results on the standard hierarchical algorithm presented in Section 4.2 as well as our intuition on the reparametrized algorithm discussed in Section 4.4. In fact, they suggest that it should be possible to improve Theorem 4.2.5 on the behaviour of the σ-chain to a result formulated almost surely with respect to the data.

4.7.2 Linear Bayesian inverse problem using finite difference discretization

We again consider the simultaneously diagonalizable setup described in Subsection 4.6.1, where $\mathcal{X} = L^2[0,1]$ with Dirichlet boundary conditions. As in the previous subsection we define $\mathcal{A}_0$ to be the negative Laplacian with Dirichlet boundary condtions in $[0,1]$. We consider the case where $K = (I + \frac{1}{10\pi^2}\mathcal{A}_0)^{-1}$, $\mathcal{C}_0 = \mathcal{A}_0^{-1}$ and $\mathcal{C}_1 = \mathcal{A}_0^{-\frac{4}{5}}$. In the language of Subsection 4.6.1, we have that Assumptions 4.6.1 are satisfied with $\alpha = 1, \beta = 4/5$ and $\ell = 1$, hence, since $2\alpha + 4\ell - 2\beta = 22/5 > 1$, Assumption 4.6.2 is also satisfied. We assume that we have data of the form

$$y = K\bar{u} + \bar{\sigma}^{-\frac{1}{2}}\mathcal{C}_1^{\frac{1}{2}}\xi, \tag{4.7.1}$$

where

$$\bar{u}(x) = 0.75 \cdot \mathbb{1}_{[0.1,0.25]}(x) + 0.25 \cdot \mathbb{1}_{[0.35,0.38]} + \sin^4(2\pi x) \cdot \mathbb{1}_{[0.5,1]}(x),\ x \in [0,1],$$

is the true underlying signal, $\bar{\sigma} = 256$ and ξ is a Gaussian white noise. Noticing that $\beta - 2\ell < 0$, we have that $\bar{u} \in \mathcal{H}^{\beta-2\ell}$ hence Assumption 4.6.3 is satisfied. We

hence have that the infinite dimensional assumptions on the underlying model are satisfied and our intuition presented in Subsection 4.3.1 applies.

Instead of working in the frequency domain and truncating the Karhunen-Loeve expansion, we discretize the domain $[0,1]$ using a uniform grid of N points and use finite differences to discretize $\mathcal{A}_0$ hence also $K, \mathcal{C}_0$ and $\mathcal{C}_1$, [7, 8]. In particular, we replace $\mathcal{A}_0$ by the $N \times N$ matrix

$$\mathcal{A}_0 = N^2 \begin{bmatrix} 2 & -1 & 0 & \dots & 0 \\ -1 & 2 & -1 & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & -1 & 2 & -1 \\ 0 & \dots & 0 & -1 & 2 \end{bmatrix},$$

and the operators $K, \mathcal{C}_0$ and $\mathcal{C}_1$ by the corresponding $N \times N$ matrices calculated through the appropriate functions of the matrix $\mathcal{A}_0$. In defining K, we also replace the identity operator by the $N \times N$ identity matrix. We define the inner product and norm in $\mathbb{R}^N$

$$\langle u, v \rangle_{\mathbb{R}^N} = \frac{1}{N} \sum_{j=1}^{N} u_j v_j,$$

and

$$\|u\|_{\mathbb{R}^N} = \left(\frac{1}{N} \sum_{j=1}^{N} u_j^2 \right)^{\frac{1}{2}}.$$

Since we consider u to be discretized on the grid, we have $u_j = u(\frac{j}{N})$, and hence have a discrete approximation of $\mathcal{X}$ with norm and inner product which are the discrete analogues of the L^2-norm and inner product. We do not prove that this discretization scheme satisfies Assumptions 4.2.1 and 4.2.4, however we expect this to be the case. In fact, instead of discretizing y in (4.7.1) by discretizing $\bar{u}$ on the grid and replacing the operators K and $\mathcal{C}_1$ by the corresponding matrices and ξ by a white noise in $\mathbb{R}^N$, we do this only for $N = 8192$ and produce the data at the lower discretization levels $N = 32, 128, 512$ and $N = 2048$ by subsampling. That is we treat the data at level $N = 8192$ as our infinite dimensional data and discretize it by subsampling. This is not exactly what we assume in (4.2.6), however it is very closely related.

As in the previous subsection, we implement the two hierarchical setups

presented in Section 4.1 and Section 4.4, with hyper-parameters $\alpha_0 = \alpha_1 = r_0 = 1, \beta_0 = \beta_1 = q_0^2 = 10^4$ chosen to give uninformative hyper-priors, that is, hyper-priors with variance which is much larger than their mean. We use 10^4 iterations of the two Gibbs samplers and chose $\sigma^{(0)} = 1$ in both cases, $\delta^{(0)} = 10$ and $\tau^{(0)} = 1/\sqrt{10}$. In the calculation of the sample mean and variance of the unknown, we again use a constant burn-in time of 1000 iterations.

In Figure 4.6 we have in the left panel the true solution (dashed black) and discretized noisy data (blue continuous), and in the middle and right panels the true solution (dashed black), the sample mean (red continuous) and 87.5% credibility bounds (shaded area) using the standard hierarchical algorithm and the reparametrized algorithm respectively, for dimension $N = 8192$. The sample means and credibility bounds at other discretization levels are similar, hence omitted.

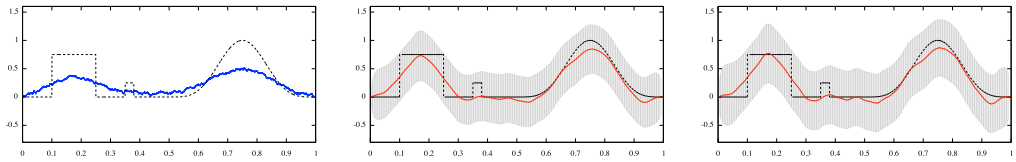

Figure 4.6: Left panel: true solution (dashed black) and blurred noisy data (blue continuous). Middle and right panels: true solution (dashed black), sample mean (red continuous) and 87.5% credibility bounds (shaded area) for standard hierarchical (middle) and reparametrized algorithm (right). Dimension is $N = 8192$.

In Figure 4.7 we have the plots of the σ-chains on the left and the δ-chains on the right, in the standard algorithm, for increasing dimension as we move from top to bottom. As predicted by Theorems 4.2.2 and 4.2.5, the plots show that while in small dimensions both chains appear to have a healthy mixing, as N increases the σ-chain moves to the true value $\bar{\sigma} = 256$ and fluctuates independently around it with fluctuations which decrease as N increases, while the δ-chain becomes slower and exhibits diffusive behaviour. In Figure 4.8 we have the plots of the σ-chains on the left and the τ^2-chains on the right, in the reparametrized algorithm for increasing dimension top to bottom. As expected, the σ-chain exhibits the same behaviour as in the standard algorithm but the τ^2-chain appears to be robust with respect to the increase in dimension.

Our observations in Figures 4.7 and 4.8 are also supported by the autocorrelation plots presented in Figure 4.9. We have four panels with the plots of the autocorrelation functions for time lag $1 - 20$ of the four chains at the different discretization levels N. On the left column we have the autocorrelation functions of the σ-chains using the standard algorithm (top) and the reparametrized algorithm

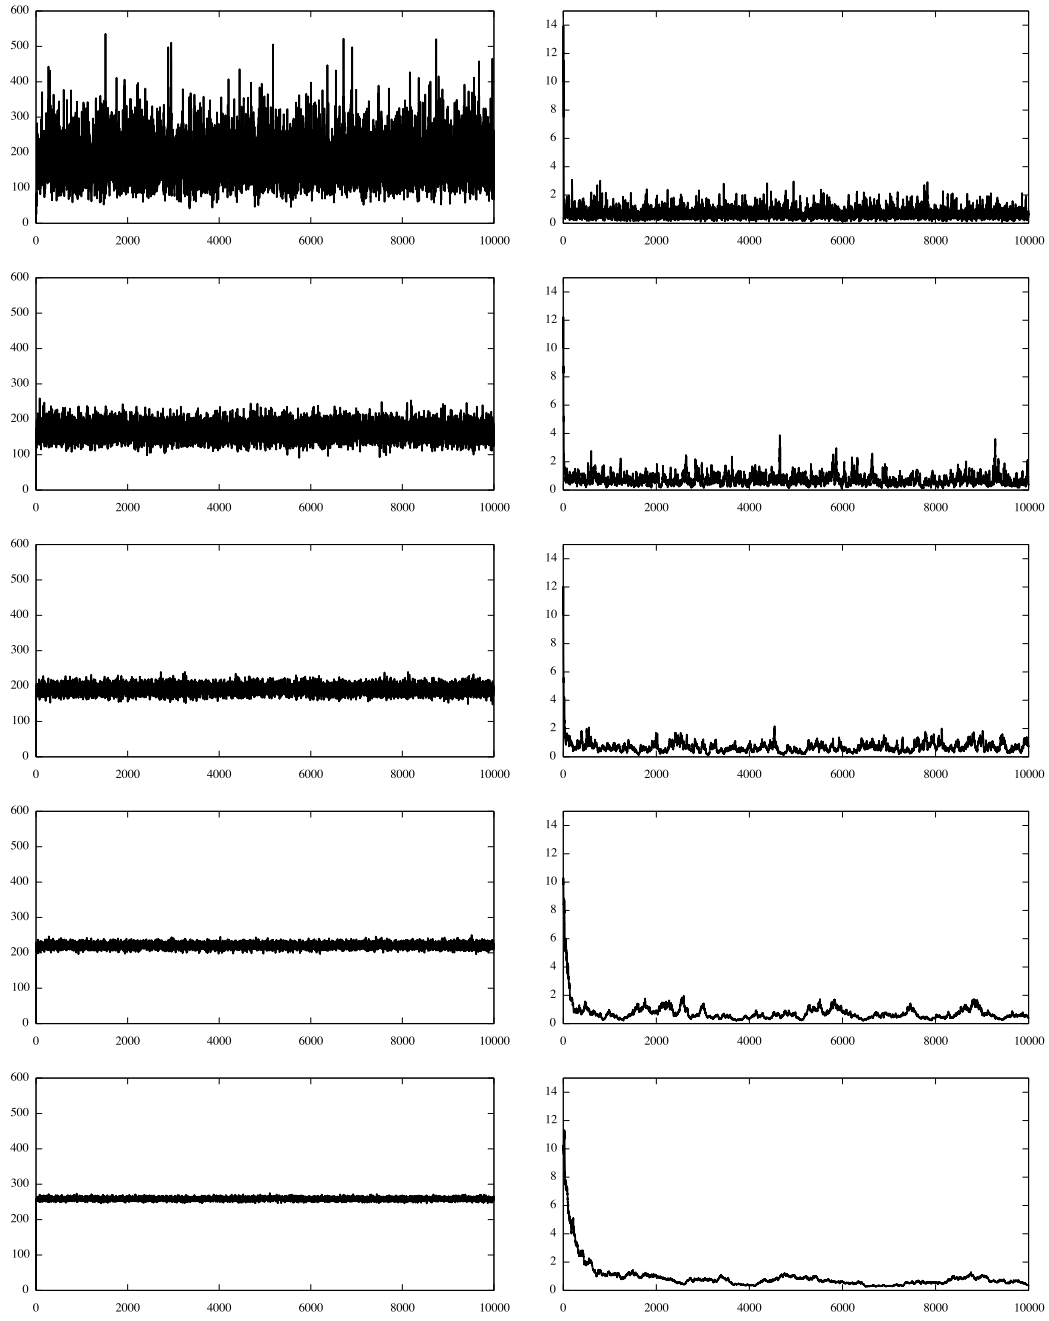

Figure 4.7: Standard algorithm: σ-chains (left column) and δ-chains (right column) for dimensions $N = 32, 128, 512, 2048$ and 8192 top to bottom.

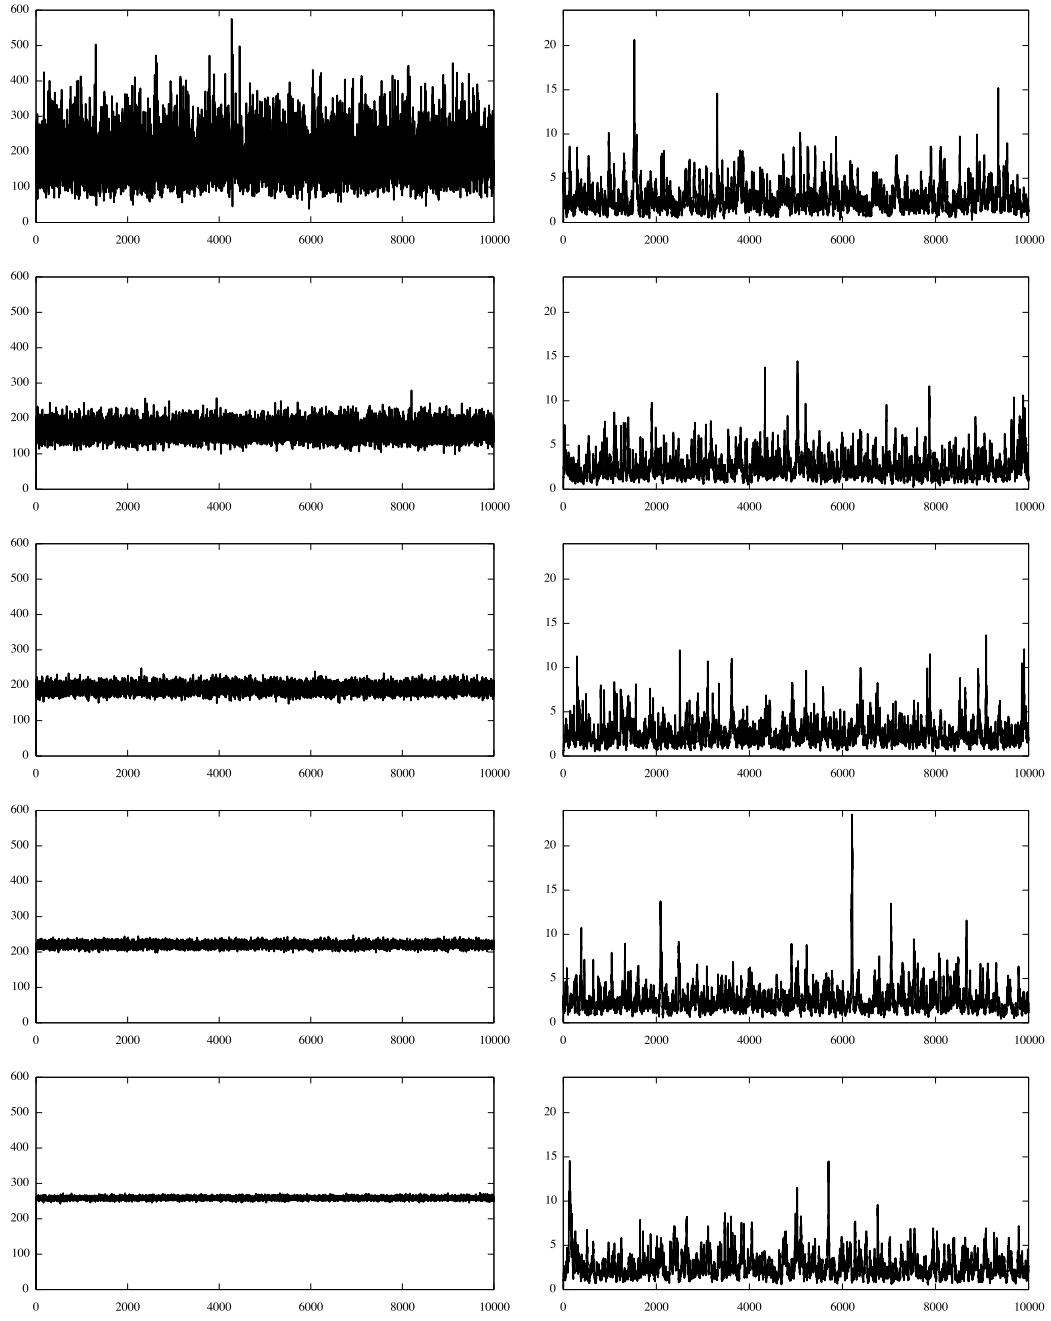

Figure 4.8: reparametrized algorithm: σ-chains (left column) and τ^2-chains (right column) for dimensions $N = 32, 128, 512, 2048$ and 8192 top to bottom.

(bottom) which are practically the same; in both cases the rate of decay of correlations seems to increase as N increases, and indeed for $N \geq 512$ even consecutive samples are practically independent. On the right column we have the autocorrelation function of the δ-chain in the case of the standard algorithm (top) and of the τ^2-chain in the case of the reparametrized algorithm (bottom). The rate of decay of correlations in the δ-chain appears to decrease as the dimension increases, and in particular for large N the correlations seem to decay very slowly. On the contrary, the rate of decay of correlations in the τ^2-chain does not seem to be affected by the increase in dimension.

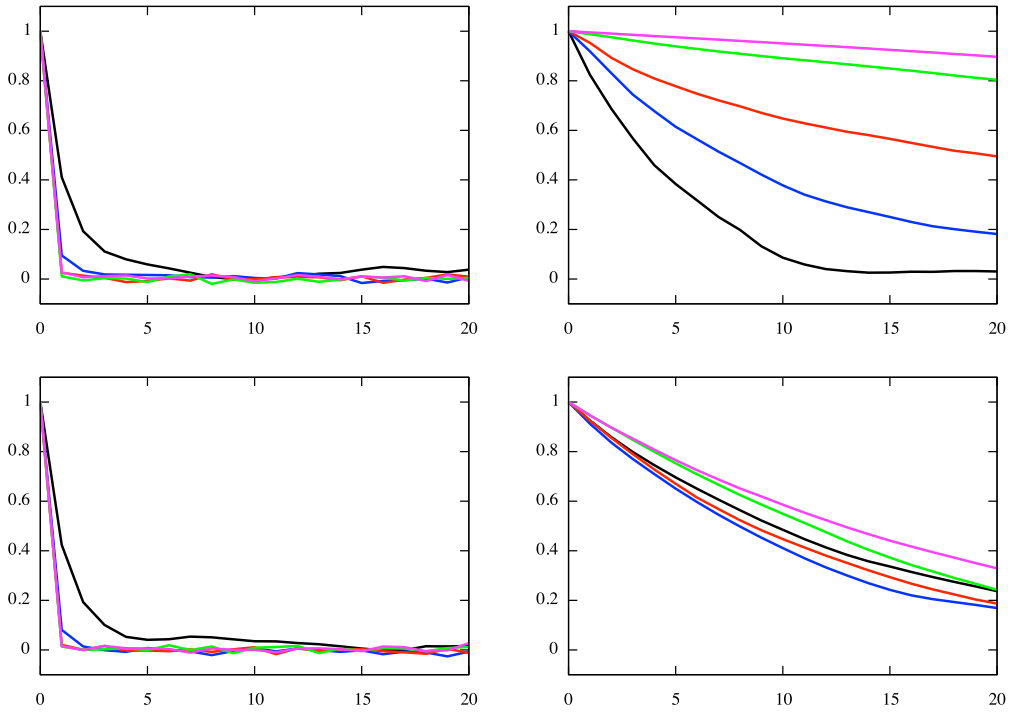

Figure 4.9: Autocorrelation functions for dimensions 32 (black), 128 (blue), 512 (red), 2048 (green) and 8192 (violet). Top left is for σ-chain in standard algorithm, top right for δ-chain in standard algorithm, bottom left for σ-chain in reparametrized algorithm and bottom right for τ^2-chain in reparametrized algorithm.

The fact that in low dimensions the rate of decay of correlations is slower in the τ^2-chain than in the δ-chain, is due to the small noise effect explained in Section 4.4. To highlight this effect, we run the reparametrized algorithm again in the case of a much smaller noise, namely $\bar{\sigma} = 256^2$, and plot the σ and τ^2-chains in Figure 4.10. As expected, the σ-chain exhibits the same behaviour as before, but the τ^2-chain mixes very poorly. We once more highlight that new work is required

to produce effective hierarchical algorithms in this small noise limit.

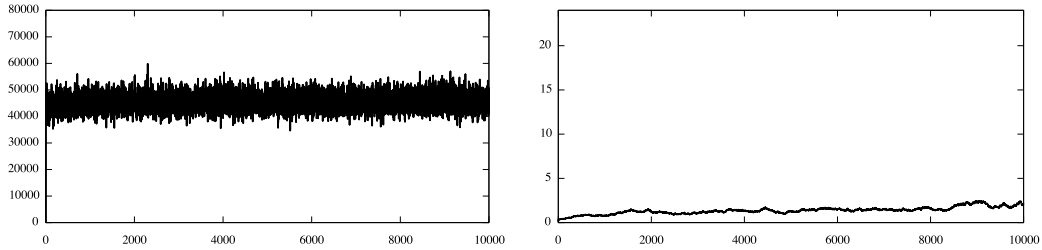

Figure 4.10: reparametrized algorithm for small noise, $\bar{\sigma} = 256^2$: σ-chain (left column) and τ^2-chain (right column) for dimension $N = 512$.

In conclusion, our numerical simulations again support the results on the standard hierarchical algorithm presented in Section 4.2 and our intuition on the reparametrized algorithm discussed in Section 4.4. Once more, they suggest that it should be possible to improve Theorem 4.2.5 on the behaviour of the σ-chain to a result formulated almost surely with respect to the data.

4.7.3 Linear Bayesian inverse problem with coarse data using finite difference discretization

We consider a slight modification of the simultaneously diagonalizable setup described in Subsection 4.6.1, where $\mathcal{X} = L^2[0,1]$ with Dirichlet boundary conditions, and where we allow K to map $\mathcal{X}$ into $\mathbb{R}^M$. In particular, we consider the problem of recovering a true signal $\bar{u}$, by observing a blurred version of it at M uniformly spaced points $\{\frac{1}{M}, ..., \frac{M}{M}\}$, polluted by additive noise of constant variance $\bar{\sigma}^{-1}$. The forward operator K is now defined as the composition of P which is the linear operator of pointwise evaluations at the M observation points, and the blurring operator $\tilde{K} = (I + \frac{1}{100\pi^2}\mathcal{A}_0)^{-1}$, where as in the previous examples $\mathcal{A}_0$ is the Dirichlet Laplacian in $[0,1]$, $K = P\tilde{K}$. We have that $\mathcal{C}_1$ is the $M \times M$ identity matrix and we choose $\mathcal{C}_0 = \mathcal{A}_0^{-1}$. Note, that due to the presence of P, the operator K is not simultaneously diagonalizable with $\mathcal{C}_0$.

As discussed in Section 4.5.2, our theory on the slowing down of the δ-chain extends to cover such settings, and it is straightforward to check that the generalized Assumptions 4.3.1 which allow $K : \mathcal{X} \to \mathbb{R}^M$ and $\mathcal{C}_1 : \mathbb{R}^M \to \mathbb{R}^M$, are satisfied. Indeed, assuming without loss of generality that $\bar{\sigma} = \delta = 1$, we have by [78, Example 6.23] that the posterior covariance and mean satisfy (4.1.4) and (4.1.5), hence $\mathcal{C}_0^{-\frac{1}{2}} m(y) = \mathcal{C}_0^{-\frac{1}{2}}(\mathcal{C}_0^{-1} + K^*K)^{-1}K^*y = (I + \mathcal{C}_0^{\frac{1}{2}}K^*K\mathcal{C}_0^{\frac{1}{2}})^{-1}\mathcal{C}_0^{\frac{1}{2}}K^*y$, where $\mathcal{C}_0^{\frac{1}{2}}K^\star y \in \mathcal{X}$, and $(I + \mathcal{C}_0^{\frac{1}{2}}K^*K\mathcal{C}_0^{\frac{1}{2}})^{-1}$ is bounded in $\mathcal{X}$ by the nonnegativity of

$\mathcal{C}_0^{\frac{1}{2}} K^* K \mathcal{C}_0^{\frac{1}{2}}$. Furthermore, we have that $\mathrm{Tr}(\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0 K^* \mathcal{C}_1^{-\frac{1}{2}}) = \mathrm{Tr}(K \mathcal{C}_0 K^*)$, which is finite since $K\mathcal{C}_0K^*$ is an $M \times M$ matrix. We do not check that the discretization scheme considered satisfies the generalized Assumptions 4.2.1, however, we expect this to be the case.

As in the previous subsection, we discretize the domain $[0, 1]$ using a uniform grid of N points, where now N is chosen such that $N/M \in \mathbb{N}$. As before we use finite differences to discretize $\mathcal{A}_0$, hence also $\tilde{K}$ and $\mathcal{C}_0$. Finally, we discretize P using the $M \times N$ matrix $P = [P_{i,j}]$, where for $i \in \{1, ..., M\}$ and $j \in \{1, ..., N\}$,

$$P_{i,j} = \begin{cases} 1, & if\ j = i\frac{N}{M} \\ 0, & otherwise. \end{cases}$$

We implement this setup for $M = 32$, $N = 32, 128, 512$ and $\bar{\sigma} = 10^3$. We use the hierarchical setups presented in Sections 4.1 and 4.4, and in particular Algorithms 2 and 4, where in Algorithm 4 the third step is switched off. As before, we use $\alpha_0 = r_0 = 1, \beta_0 = q_0^2 = 10^4$ chosen to give uninformative hyper-priors, that is, hyper-priors with variance which is much larger than their mean. As in the previous examples we use 10^4 iterations of the two Gibbs samplers and chose $\delta^{(0)} = 10$ and $\tau^{(0)} = 1/\sqrt{10}$. In the calculation of the sample mean and variance of the unknown, we use a constant burn-in time of 1000 iterations.

In Figure 4.11 we have in the left column the true solution (dashed black) and blurred noisy data (blue dots), and in the middle and right columns the true solution (dashed black), the sample mean (red continuous) and 87.5% credibility bounds (shaded area) using the standard hierarchical algorithm and the reparametrized algorithm respectively, for increasing discretization level of the unknown, N, top to bottom. We can see that as N increases, the sample mean gets worse in both cases, which is reasonable since we observe proportionally less points. However, for large N the standard algorithm appears to perform a lot worse than the reparametrized one, and in particular for $N = 512$ fails completely to produce a decent approximation.

In Figure 4.12 we have the plots of the δ-chains on the left and the τ^2-chains on the right, in the standard hierarchical and reparametrized algorithms respectively, for increasing N as we move from top to bottom. As predicted in Subsection 4.5.2, while for small N the δ-chain mixes well, as N increases it becomes slower and exhibits diffusive behaviour. On the other hand, the τ^2-chain appears to be mixing well independently of N.

Our observations in Figure 4.12 are also supported by the autocorrelation plots presented in Figure 4.13. We have two panels with the plots of the autocorre-

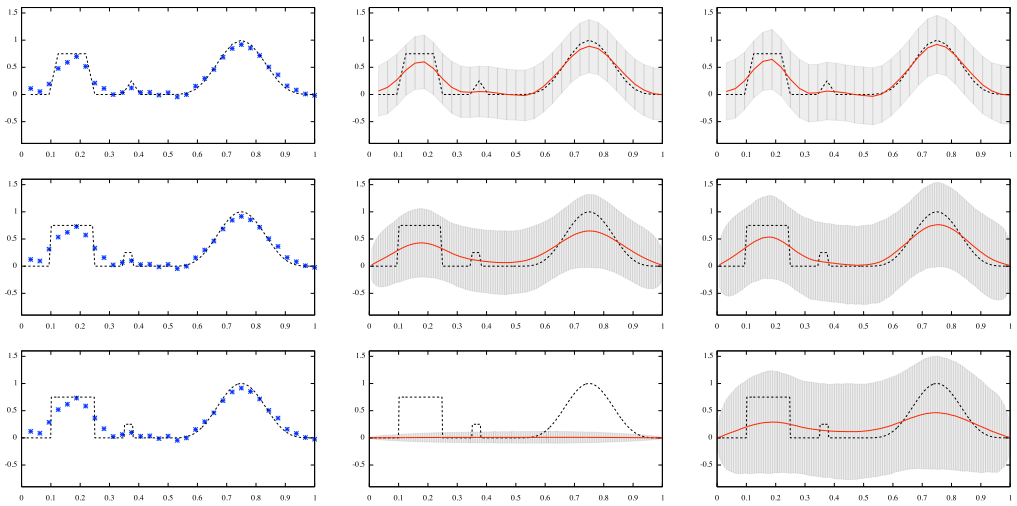

Figure 4.11: Left column: true solution (dashed black) and blurred noisy data (blue dots). Middle and right columns: true solution (dashed black), sample mean (red continuous) and 87.5% credibility bounds (shaded area) for standard hierarchical (middle) and reparametrized algorithm (right). $N = 32, 128, 512$ top to bottom.

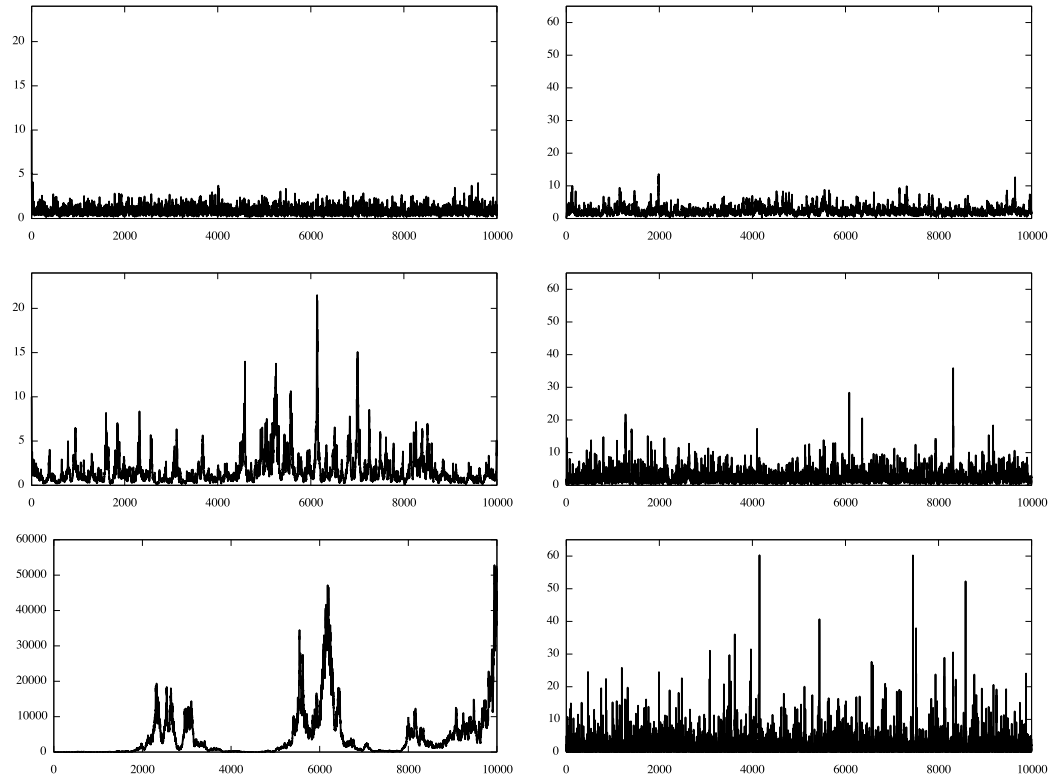

Figure 4.12: δ-chains in standard algorithm (left column) and τ^2-chains in reparametrized algorithm (right column) for $N = 32, 128, 512$ top to bottom.

lation functions for time lag $1 - 20$ of the two chains at the different discretization levels of the unknown, N. On the left panel we have the autocorrelation functions of the δ-chains in the case of the standard algorithm and on the right the autocorrelation functions of the τ^2-chains in the case of the reparametrized algorithm. The rate of decay of correlations in the δ-chain appears to decrease as N increases, and in particular for $N = 512$ the correlations seem to decay very slowly. On the contrary, the rate of decay of correlations in the τ^2-chain actually seems to increase as N increases.

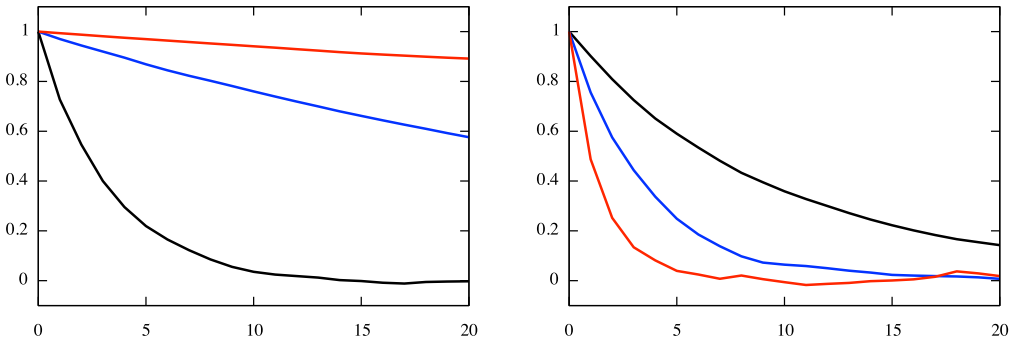

Figure 4.13: Autocorrelation functions for dimensions 32 (black), 128 (blue) and 512 (red). Left is for δ-chain in standard algorithm, right for τ^2-chain in reparametrized algorithm.

The fact that in low dimensions the rate of decay of correlations is slower in the τ^2-chain than in the δ-chain, is due to the small noise effect explained in Section 4.4. To highlight this effect, we run the reparametrized algorithm again in the case of a much smaller noise, namely $\bar{\sigma} = 10^6$, and plot the τ^2-chain in Figure 4.14. As expected, the τ^2-chain mixes very poorly in this case and new work is required to produce effective hierarchical algorithms in this small noise limit.

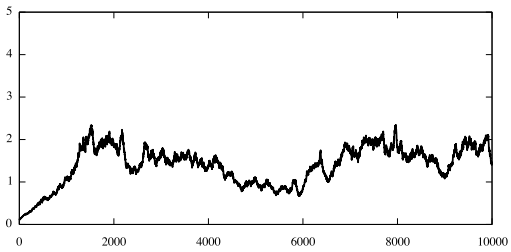

Figure 4.14: τ^2-chain using reparametrized algorithm for small noise $\bar{\sigma} = 10^6$ and $N = 128$.

Note that the small noise effect can also explain the increase in the rate of

decay of correlations in the τ^2-chain when increasing N. Indeed, for fixed M, as N gets larger we observe proportionally less points of the unknown function, hence the restriction $\tau v \simeq \bar{u}$ is on proportionally less points. This suggests that the a posteriori dependence between τ and v is weaker hence the better mixing in the τ^2-chain.

4.8 Conclusions

We have considered a class of Bayesian linear inverse problems with Gaussian additive noise and Gaussian priors, for determining an unknown parameter u from a blurred noisy observation y in a Hilbert space setting. We discretized the setup in $\mathbb{R}^N$ using consistent discretizations of the operators defining the problem. The use of consistent discretizations enabled the comparison of the values of the scalings σ^{-1} and δ^{-1} of the noise and prior covariance operators respectively, across different discretization levels; this is very important in applied problems where it is often of interest to perform simulations at different levels of model-refinement. We studied a standard conditionally conjugate hierarchical setup for simultaneous inference on the unknown u and the two scalings σ^{-1}, δ^{-1}. Conditional conjugacy makes natural the use of a Gibbs sampler to sample the posterior (Algorithm 1). We showed that under assumptions on the discrete level (Assumptions 4.2.1 and 4.2.4) which we believe, and indeed have proved in some specific scenarios, that are inherited by reasonable discretizations from natural assumptions on the underlying infinite dimensional problem (Assumptions 4.3.1 and 4.3.2), the behaviour of the Gibbs sampler has two scales: an increasingly slow scale and an increasingly fast scale on which the δ and σ-chains evolve respectively. We provided both intuition based on the underlying infinite dimensional model (Subsection 4.3.1) as well as rigorous theorems quantifying the slowing down of the δ-chain when σ is fixed (Theorem 4.2.2) and the speed up of the σ-chain when δ is fixed (Theorem 4.2.5).

We proposed a reparametrization of the prior scaling which is robust as the dimension increases (Section 4.4), however, this reparametrization deteriorates as the noise disappears. The frequentist properties of the posterior distribution in the small noise limit using Gaussian priors with a scaling hyper-parameter are studied in [79] in the white noise model (see Subsection 4.7.1) using two methods: i) an empirical Bayes method for estimating the value of the prior scaling from the data; ii) the standard hierarchical Bayesian method considered also in the present chapter (Subsection 4.2.1) for inference simultaneously on the unknown and the scaling of the prior. It is shown that both methods achieve optimal posterior contraction rates over

a range of regularity classes of the true solution. However, as our results suggest, the implementation of the hierarchical Bayesian method in the large dimensional limit is problematic, and in particular the Gibbs sampler naturally used to sample the posterior suffers from increasingly poor mixing as the dimension grows. On the other hand, while the empirical Bayes method is appealing because of the lack of mixing issues, it involves solving an optimization problem which in more complicated models can be computationally demanding, and it does not provide uncertainty quantification of the prior scaling which may be desirable. We believe that more research and new ideas are required in the small noise limit.

Our theory on the slowing down of the δ-chain was extended to cover nonlinear Gaussian-conjugate Bayesian inverse problems (Subsection 4.5.1) and in particular the nonparametric drift estimation in SDE's setting considered in [64, 59, 56]. Again our main result (Theorem 4.5.3) holds under assumptions on the discrete level (Assumptions 4.5.1) which we believe are inherited by reasonable discretizations of natural assumptions on the underlying infinite dimensional model (Assumptions 4.5.2). We also extended our theory to cases where the discretization levels of the data and the unknown differ (Subsection 4.5.2).

We provided four families of inverse problems satisfying our assumptions on the underlying infinite dimensional model (Section 4.6), and for two of them, which are families of mildly and severely ill-posed linear inverse problems in a simultaneously diagonal setting, we also showed that a spectral method based on the common eigenbasis satisfies our discrete level assumptions (Subsections 4.6.1 and 4.6.2). We also provided numerical evidence supporting our theory in more general linear inverse problem settings, using both spectral truncation as well as discretization via finite differences (Section 4.7).

Future directions of interest include a rigorous proof of the diffusion limit (4.2.3) of the δ-chain when σ is fixed, using the standard theory of diffusion approximation of Markov processes [23] in a similar way as in [67, 69] and more recently in [63], as well as exploring the properties of the limiting SDE. Furthermore, our numerical experiments suggest that it may be possible to improve Theorem 4.2.5 and obtain a result on the speed up of the σ-chain formulated almost surely with respect to the data. Our simulations indicate that it may also be possible to prove a result on the joint (δ, σ)-chain, stating that the two components de-correlate instantly, and the δ-chain converges to a limiting SDE while the σ-chain converges to the true value of the noise scaling in the data.

Finally, our infinite dimensional intuition extends to hierarchical setups for inference on other hyper-parameters, for instance the prior and noise regularity

parameters as studied in [43]. The idea is the same as the intuition presented in Subsection 4.3.1, since in infinite dimensions two Gaussian measures $\mathcal{N}(0, \Sigma_1)$ and $\mathcal{N}(0, \Sigma_2)$, where Σ_1 and Σ_2 are simultaneously diagonalizable with eigenvalues $\{j^{-a_1}\}_{j\in\mathbb{N}}$ and $\{j^{-a_2}\}_{j\in\mathbb{N}}$ respectively, are mutually singular unless $a_1 = a_2$. We expect that the chain for estimating the prior regularity slows down as the dimension increases, while the chain for estimating the noise regularity speeds up.

4.9 Appendix

In this section we present several technical results necessary for proving our main theorems presented in Sections 4.2 and 4.5. First, in Subsections 4.9.1 and 4.9.2 we present the proofs of Lemmas 4.3.3 and 4.3.4 which form the basis of the proofs of Theorems 4.2.2 and 4.2.5 on the analysis of the δ and σ-chains respectively, contained in Section 4.3. Then, in Subsection 4.9.3 we state and prove two lemmas on the negative moments of the rate parameters in the δ and σ draws (4.2.2) and (4.2.5) respectively, which allow us to control several lower order terms arising in the proofs of our main results in Section 4.3. In Subsection 4.9.4, we present and prove several technical probability lemmas as well as two linear algebra lemmas, which are useful in our analysis; in particular, Lemma 4.9.8 contains several useful implications of Assumptions 4.2.1. Finally, in Subsection 4.9.5 we present the necessary generalizations of the lemmas contained in the earlier subsections of the appendix, required for proving Theorem 4.5.3 on the analysis of the δ-chain in the general Gaussian-conjugate setting presented in Subsection 4.5.1.

4.9.1 Proof of Lemma 4.3.3

Proof. Let $\{e_j\}_{j=1}^N$ be any orthonormal basis of $\mathbb{R}^N$ (with respect to the possibly scaled norm $\|\cdot\|_{\mathbb{R}^N}$) and for any $w \in \mathbb{R}^N$ write $w_j := \langle w, e_j\rangle_{\mathbb{R}^N}$. We then have that $\zeta = \sum_{j=1}^N \zeta_j e_j$ where $\{\zeta_j\}_{j=1}^N$ is a sequence of independent standard normal random variables.

Using (4.2.1) we have,

$$\begin{aligned}\left\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\right\|_{\mathbb{R}^N}^2 &= \left\|\mathcal{C}_0^{-\frac{1}{2}} m_{\sigma,\delta}(\boldsymbol{y_N})\right\|_{\mathbb{R}^N}^2 + \left\|\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta\right\|_{\mathbb{R}^N}^2 + 2\left\langle \mathcal{C}_0^{-\frac{1}{2}} m_{\sigma,\delta}(\boldsymbol{y_N}), \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta\right\rangle_{\mathbb{R}^N} \\ &:= A_N + B_N + C_N.\end{aligned}$$

Under Assumptions 4.2.1, we can analyze each term as follows:

A) by Assumption 4.2.1(i), for almost all data y, this term and all its positive integer powers are bounded uniformly in N.

B) the second term can be written as

$$\begin{aligned}\big\|\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta\big\|_{\mathbb{R}^N}^2 &= \big\langle \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta, \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta\big\rangle_{\mathbb{R}^N} = \big\langle \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\mathcal{C}_0^{-1}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta, \zeta\big\rangle_{\mathbb{R}^N} \\ &= \delta^{-1}\big\langle \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}(\mathcal{C}_{\sigma,\delta}^{-1} - \sigma K^*\mathcal{C}_1^{-1}K)\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta, \zeta\big\rangle_{\mathbb{R}^N} \\ &= \delta^{-1}\big\langle (I - \sigma\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}K^*\mathcal{C}_1^{-1}K\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}})\zeta, \zeta\big\rangle_{\mathbb{R}^N} \\ &= \delta^{-1}\big\|\zeta\big\|_{\mathbb{R}^N}^2 - \delta^{-1}\sigma\big\|\mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta\big\|_{\mathbb{R}^N}^2 \\ &:= b_{1,N} - b_{2,N},\end{aligned}$$

where

b1) for the first term we have

$$b_{1,N} = \delta^{-1}\big\|\zeta\big\|_{\mathbb{R}^N}^2 = \frac{N}{\delta} + \frac{1}{\delta}\sum_{j=1}^{N}(\zeta_j^2 - 1) := \frac{N}{\delta} + \frac{\sqrt{2N}}{\delta}W_{1,N},$$

where as $N \to \infty$, $W_{1,N} = \frac{1}{\sqrt{2N}}\sum_{j=1}^{N}(\zeta_j^2 - 1)$ converges weakly to a standard normal random variable by the Central Limit Theorem and by Lemma 4.9.5 has moments of every order which are bounded uniformly in N;

b2) for the second term we have by Lemma 4.9.8(ii) that $\mathbb{E}^{\zeta}[b_{2,N}]$ is uniformly bounded in N. In fact using Lemma 4.9.4 together with Lemma 4.9.8(ii), we get that for any $p \in \mathbb{N}$, $\mathbb{E}^{\zeta}[b_{2,N}^p]$ is bounded independently of N.

C) for the third term we have

$$\begin{aligned}\big\langle \mathcal{C}_0^{-\frac{1}{2}}m_{\sigma,\delta}(\boldsymbol{y_N}), \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta\big\rangle_{\mathbb{R}^N} &= \big\langle (\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}})^*\mathcal{C}_0^{-\frac{1}{2}}m_{\sigma,\delta}(\boldsymbol{y_N}), \zeta\big\rangle_{\mathbb{R}^N} \\ &= \sum_{j=1}^{N}((\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}})^*\mathcal{C}_0^{-\frac{1}{2}}m_{\sigma,\delta}(\boldsymbol{y_N}))_j\zeta_j.\end{aligned}$$

It holds that

$$\sum_{j=1}^{N}((\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}})^*\mathcal{C}_0^{-\frac{1}{2}}m_{\sigma,\delta}(\boldsymbol{y_N}))_j^2 = \big\|(\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}})^*\mathcal{C}_0^{-\frac{1}{2}}m_{\sigma,\delta}(\boldsymbol{y_N})\big\|_{\mathbb{R}^N}^2,$$

and we claim that the norm on the right hand side is uniformly bounded in N almost surely with respect to the data. Indeed, by (4.1.4), the Cauchy-Schwarz inequality and the non-negative definiteness of the matrix $\mathcal{C}_0^{\frac{1}{2}}K^*\mathcal{C}_1^{-1}K\mathcal{C}_0^{\frac{1}{2}}$, we

have

$$\begin{aligned}\left\|(\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}})^* u\right\|_{\mathbb{R}^N}^2 &= \left\langle \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\sigma,\delta}\mathcal{C}_0^{-\frac{1}{2}}u, u\right\rangle_{\mathbb{R}^N} \\ &= \left\langle \delta^{-1}(I+\frac{\sigma}{\delta}\mathcal{C}_0^{\frac{1}{2}}K^*\mathcal{C}_1^{-1}K\mathcal{C}_0^{\frac{1}{2}})^{-1}u, u\right\rangle_{\mathbb{R}^N} \\ &\leq \left\|\delta^{-1}(I+\frac{\sigma}{\delta}\mathcal{C}_0^{\frac{1}{2}}K^*\mathcal{C}_1^{-1}K\mathcal{C}_0^{\frac{1}{2}})^{-1}u\right\|_{\mathbb{R}^N}\left\|u\right\|_{\mathbb{R}^N} \\ &\leq \delta^{-1}\left\|u\right\|_{\mathbb{R}^N}^2.\end{aligned}$$

Combining with Assumption 4.2.1(i) we get the claim and therefore by Lemma 4.9.3 below we get that the third term has y-almost surely all even moments uniformly bounded in N.

We define $F_N = \beta_0 + \frac{A_N - b_{2,N} + C_N}{2}$ and observe that since all terms have bounded moments of every order uniformly in N y-almost surely, Hölder's inequality secures that F_N also has bounded moments of every order uniformly in N almost surely with respect to y. □

4.9.2 Proof of Lemma 4.3.4

Proof. Let $\{e_j\}_{j=1}^N$ be any orthonormal basis of $\mathbb{R}^N$ (with respect to the possibly scaled norm $\|\cdot\|_{\mathbb{R}^N}$) and for any $w \in \mathbb{R}^N$ write $w_j := \langle w, e_j\rangle$. We then have that $\zeta = \sum_{j=1}^N \zeta_j e_j$ and $\xi = \sum_{j=1}^N \xi_j e_j$ where $\{\zeta_j\}_{j=1}^N$ and $\{\xi_j\}_{j=1}^N$ are two independent sequences of independent standard normal random variables.

We have

$$\begin{aligned}\left\|\mathcal{C}_1^{-\frac{1}{2}}(Ku_\sigma^{(k)} - \boldsymbol{y_N})\right\|_{\mathbb{R}^N}^2 =& \left\|\mathcal{C}_1^{-\frac{1}{2}}(z - Ku_\sigma^{(k)})\right\|_{\mathbb{R}^N}^2 + \bar{\sigma}^{-1}\left\|\xi\right\|_{\mathbb{R}^N}^2 \\ &+ 2\bar{\sigma}^{-1}\left\langle \mathcal{C}_1^{-\frac{1}{2}}(z - Ku_\sigma^{(k)}), \xi\right\rangle_{\mathbb{R}^N} \\ :=& A_N + B_N + C_N.\end{aligned}$$

Under our assumptions we can analyze each term as follows:

A) we have using (4.1.5), (4.2.6) and the triangle inequality that

$$\begin{aligned}\left\|\mathcal{C}_1^{-\frac{1}{2}}(z - Ku_\sigma^{(k)})\right\|_{\mathbb{R}^N}^2 =& \left\|\mathcal{C}_1^{-\frac{1}{2}}z - \mathcal{C}_1^{-\frac{1}{2}}Km_{\sigma,\delta}(\boldsymbol{y_N}) - \mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta\right\|_{\mathbb{R}^N}^2 \\ \leq& \left\|\mathcal{C}_1^{-\frac{1}{2}}z\right\|_{\mathbb{R}^N}^2 + \sigma\left\|\mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_{\sigma,\delta}K^*\mathcal{C}_1^{-\frac{1}{2}}\mathcal{C}_1^{-\frac{1}{2}}z\right\|_{\mathbb{R}^N}^2 \\ &+ \frac{\sigma}{\bar{\sigma}}\left\|\mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_{\sigma,\delta}K^*\mathcal{C}_1^{-\frac{1}{2}}\xi\right\|_{\mathbb{R}^N}^2 + \left\|\mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\zeta\right\|_{\mathbb{R}^N}^2,\end{aligned}$$

hence it suffices to examine each term separately. By Assumption 4.2.4 and

Lemma 4.9.8(iv), any positive power of the first two terms is bounded uniformly in N. For the third term, since for an $N \times N$ matrix A it holds $\|A\|_{2,N} = \|A^*\|_{2,N} = \sqrt{\|A^*A\|_{2,N}}$, by Lemma 4.9.8(iv) we have

$$\begin{aligned}\|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta} K^* \mathcal{C}_1^{-\frac{1}{2}} \xi\|_{\mathbb{R}^N} &\leq \|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}}\|_{2,N} \|\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} K^* \mathcal{C}_1^{-\frac{1}{2}} \xi\|_{\mathbb{R}^N} \\ &\leq c \|\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} K^* \mathcal{C}_1^{-\frac{1}{2}} \xi\|_{\mathbb{R}^N},\end{aligned}$$

where $c \geq 0$ is independent of N. Using Lemma 4.9.8(ii) and Lemma 4.9.4 we get that the third term has bounded moments of all positive orders uniformly in N. Finally, Lemma 4.9.8(ii) and Lemma 4.9.4 secure that the fourth term also has bounded moments of all positive orders uniformly in N. Hence, we have that for any $p \in \mathbb{N}$, $\mathbb{E}^{\xi,\zeta}[A_N^p]$ is uniformly bounded in N.

B) for the second term we have for any $\varepsilon > 0$

$$\bar{\sigma}^{-1} \|\xi\|_{\mathbb{R}^N}^2 = \frac{N}{\bar{\sigma}} + \frac{1}{\bar{\sigma}} \sum_{j=1}^{N} (\xi_j^2 - 1) := \frac{N}{\bar{\sigma}} + \frac{\sqrt{2N}}{\bar{\sigma}} Z_{1,N},$$

where as $N \to \infty$, $Z_{1,N} = \frac{1}{\sqrt{2N}} \sum_{j=1}^{N} (\xi_j^2 - 1)$ converges weakly to a standard normal random variable by the Central Limit Theorem and by Lemma 4.9.5 has moments of every positive order which are bounded uniformly in N.

C) we expand the third term using expression (4.2.4) for $u_\sigma^{(k)}$,

$$\begin{aligned}&2\bar{\sigma}^{-\frac{1}{2}} \big\langle \mathcal{C}_1^{-\frac{1}{2}} (z - K u_\sigma^{(k)}), \xi \big\rangle_{\mathbb{R}^N} \\ &\quad = 2\bar{\sigma}^{-\frac{1}{2}} \big\langle \mathcal{C}_1^{-\frac{1}{2}} (z - K \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} \zeta), \xi \big\rangle_{\mathbb{R}^N} - 2\bar{\sigma}^{-\frac{1}{2}} \big\langle \mathcal{C}_1^{-\frac{1}{2}} K m_{\sigma,\delta}(\boldsymbol{y_N}), \xi \big\rangle_{\mathbb{R}^N}.\end{aligned}$$

Replacing the assumed data (4.2.6) in the mean equation (4.1.5), we have

$$\begin{aligned}C_N &= 2\bar{\sigma}^{-\frac{1}{2}} \big\langle \mathcal{C}_1^{-\frac{1}{2}} z - \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} \zeta - \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta} K^* \mathcal{C}_1^{-1} z, \xi \big\rangle_{\mathbb{R}^N} \\ &\quad - 2\bar{\sigma}^{-1} \big\langle \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta} K^* \mathcal{C}_1^{-\frac{1}{2}} \xi, \xi \big\rangle_{\mathbb{R}^N} \\ &:= 2\bar{\sigma}^{-\frac{1}{2}} \big\langle x, \xi \big\rangle_{\mathbb{R}^N} - 2\bar{\sigma}^{-1} \|\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} K^* \mathcal{C}_1^{-\frac{1}{2}} \xi\|_{\mathbb{R}^N}^2,\end{aligned}$$

where $x = \mathcal{C}_1^{-\frac{1}{2}} z - \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta} K^* \mathcal{C}_1^{-1} z - \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} \zeta$ is a Gaussian random variable which only depends on ζ. By Lemma 4.9.8(ii) and Lemma 4.9.4 the second term has bounded moments of any positive order. For the first term, we have that $x_j = m_j + t_j \zeta_j$, where m_j and t_j are deterministic real numbers, and

by Assumption 4.2.4, Lemma 4.9.8(ii) and Lemma 4.9.8(iv), $\sum_{j=1}^{N} m_j^2$ and $\sum_{j=1}^{N} t_j^2$ are bounded uniformly in N. We claim that the random sequence $\mathbb{E}^{\xi,\zeta}[\langle x(\zeta), \xi\rangle_{\mathbb{R}^N}^{2p}]$, is bounded uniformly in N for any $p \in \mathbb{N}$. Indeed,

$$\mathbb{E}^{\xi,\zeta}[\langle x(\zeta), \xi\rangle_{\mathbb{R}^N}^{2p}] = \mathbb{E}^{\xi,\zeta}\left[\Big(\sum_{j=1}^{N} x(\zeta)_j \xi_j\Big)^{2p}\right],$$

and by the Minkowski inequality it suffices to separately check the boundedness of $\mathbb{E}^{\xi}\left[\Big(\sum_{j=1}^{N} m_j \xi_j\Big)^{2p}\right]$ and $\mathbb{E}^{\xi,\zeta}\left[\Big(\sum_{j=1}^{N} t_j \zeta_j \xi_j\Big)^{2p}\right]$, both of which can be established using Lemma 4.9.3 below, since m_j and t_j are square summable and both $\{\xi_j\}_{j=1}^{N}$ and $\{\zeta_j \xi_j\}_{j=1}^{N}$ are sequences of independent and identically distributed random variables with finite even moments of any order and zero odd moments. Concatenating, we have that $\mathbb{E}^{\xi,\zeta}[C_N^p]$ is bounded independently of N, for any $p \in \mathbb{N}$.

We define $H_N = \beta_1 + \frac{A_N + C_N}{2}$ and observe that since all terms have bounded moments of every positive order uniformly in N, Hölder's inequality secures that H_N also has bounded moments of every positive order uniformly in N. □

4.9.3 Lemmas on the negative moments of the rate parameters in the δ and σ draws

Lemma 4.9.1. *Let $u_\delta^{(k)}$ be as in (4.2.1), for any $\delta, \sigma > 0$. Under Assumptions 4.2.1, for $i = 1, 2$, we have*

$$\mathbb{E}^{\zeta}\left[(\beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\big\|_{\mathbb{R}^N}^2)^{-2i}\right] = \mathcal{O}(N^{-2i}),$$

as $N \to \infty$, almost surely with respect to y.

Proof. Without loss of generality we consider the case $\delta = \sigma = 1$ and drop the σ and δ dependence in u, m and $\mathcal{C}$. To de-clutter our notation we also drop the dependence of m on the data. Since $\beta_0 \geq 0$ it suffices to show it for $\beta_0 = 0$. Formally, the random variable $\big\|\mathcal{C}_0^{-\frac{1}{2}} u^{(k)}\big\|_{\mathbb{R}^N}^2$ behaves like a chi-squared random variable with N degrees of freedom. We estimate the squared norm by a random variable Y_N of known moment generating function $M_{Y_N}(t)$, and use the following formula from [16] for the calculation of negative moments of nonnegative random variables

$$\mathbb{E}[Y_N^{-l}] = \Gamma(l)^{-1} \int_0^\infty t^{l-1} M_{Y_N}(-t) dt, \; l \in \mathbb{N}. \tag{4.9.1}$$

We begin by showing there exists a constant $c > 0$ independent of N such that $\big\|\mathcal{C}^{-\frac{1}{2}}\mathcal{C}_0^{\frac{1}{2}}v\big\|_{\mathbb{R}^N} \leq c\big\|v\big\|_{\mathbb{R}^N}$ for any $v \in \mathbb{R}^N$. We have,

$$\begin{aligned}\big\|\mathcal{C}^{-\frac{1}{2}}\mathcal{C}_0^{\frac{1}{2}}v\big\|_{\mathbb{R}^N}^2 &= \big\langle \mathcal{C}_0^{\frac{1}{2}}\mathcal{C}^{-1}\mathcal{C}_0^{\frac{1}{2}}v, v\big\rangle_{\mathbb{R}^N} \\ &= \big\langle (I + \mathcal{C}_0^{\frac{1}{2}}K^*\mathcal{C}_1^{-1}K\mathcal{C}_0^{\frac{1}{2}})v, v\big\rangle_{\mathbb{R}^N} \\ &= \big\|v\big\|_{\mathbb{R}^N}^2 + \big\|\mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_0^{\frac{1}{2}}v\big\|_{\mathbb{R}^N}^2 \\ &\leq (1+c_2)\big\|v\big\|_{\mathbb{R}^N}^2,\end{aligned}$$

by Lemma 4.9.8(iii). The proved claim gives the estimate

$$\begin{aligned}\big\|\mathcal{C}_0^{-\frac{1}{2}}u^{(k)}\big\|_{\mathbb{R}^N}^2 = \big\|\mathcal{C}_0^{-\frac{1}{2}}(m + \mathcal{C}^{\frac{1}{2}}\zeta)\big\|_{\mathbb{R}^N}^2 &= \big\|\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}^{\frac{1}{2}}(\mathcal{C}^{-\frac{1}{2}}m + \zeta)\big\|_{\mathbb{R}^N}^2 \\ &\geq c^{-1}\big\|\mathcal{C}^{-\frac{1}{2}}m + \zeta\big\|_{\mathbb{R}^N}^2,\end{aligned}$$

hence it suffices to show that almost surely with respect to y we have $\mathbb{E}^{\zeta}[Y_N^{-2i}] = \mathcal{O}(N^{-2i})$, for $Y_N := \big\|\mathcal{C}^{-\frac{1}{2}}m + \zeta\big\|_{\mathbb{R}^N}^2$. Indeed, let $\{e_j\}_{j=1}^N$ be any orthonormal basis of $\mathbb{R}^N$ (with respect to the possibly scaled norm $\|\cdot\|_{\mathbb{R}^N}$), and define $w_j := \big\langle w, e_j\big\rangle$ for any $w \in \mathbb{R}^N$. Then we have

$$Y_N = \sum_{j=1}^{N}((\mathcal{C}^{-\frac{1}{2}}m)_j + \zeta_j)^2,$$

where $\zeta_j \sim \mathcal{N}(0,1)$ are the mutually independent components of the white noise ζ and $(\mathcal{C}^{-\frac{1}{2}}m)_j$ are independent of ζ, therefore Y_N is a non-central chi-squared random variable with N degrees of freedom and non-centrality parameter $p_N := \sum_{j=1}^{N}(\mathcal{C}^{-\frac{1}{2}}m)_j^2 \geq 0$. The definition and properties of the non-central chi-squared distribution can be found in [39], where in particular, we find that the moment generating function of Y_N is

$$M_{Y_N}(t) = (1-2t)^{-\frac{N}{2}}\exp\big(\frac{p_N t}{1-2t}\big),$$

hence using (4.9.1) we have for $i = 1, 2$,

$$\begin{aligned}\mathbb{E}^{\zeta}[Y_N^{-2i}] &= \Gamma(2i)^{-1}\int_0^{\infty} t^{2i-1}(1+2t)^{-\frac{N}{2}}\exp\big(\frac{-p_N t}{1+2t}\big)dt \\ &\leq c\int_0^{\infty} t^{2i-1}(1+2t)^{-\frac{N}{2}}dt \\ &= \mathcal{O}(N^{-2i}),\end{aligned}$$

provided $N > 4i$, where the last integral can by calculated analytically by integration by parts. □

Lemma 4.9.2. *Let $u_\sigma^{(k)}$ and $\boldsymbol{y_N}$ be as in (4.2.4) and (4.2.6), respectively, for any $\bar{\sigma}, \sigma, \delta > 0$. Then under Assumptions 4.2.1 and 4.2.4, for $i = 1, 2$, we have*

$$\mathbb{E}^{\xi,\zeta}\left[(\beta_1 + \frac{1}{2}\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku_\sigma^{(k)} - \boldsymbol{y_N})\big\|_{\mathbb{R}^N}^2)^{-2i}\right] = \mathcal{O}(N^{-2i}),$$

as $N \to \infty$.

Proof. As in Lemma 4.9.1 without loss of generality we show it for $\bar{\sigma} = \sigma = \delta = 1$ and $\beta_1 = 0$. We once more drop the σ and δ dependence in u, m and $\mathcal{C}$. For fixed ζ, we estimate $\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku_\sigma^{(k)} - \boldsymbol{y_N})\big\|_{\mathbb{R}^N}^2$ by a random variable $Y_N = Y_N(\zeta)$ of known moment generating function $M_{Y_N}(t;\zeta)$ and use (4.9.1) for the calculation of the second and fourth negative moments with respect to ξ. By (4.1.5), (4.2.4) and (4.2.6) we have

$$\begin{aligned}
\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku_\sigma^{(k)} - \boldsymbol{y_N})\big\|_{\mathbb{R}^N}^2 &= \big\|\mathcal{C}_1^{-\frac{1}{2}}(z + \mathcal{C}_1^{\frac{1}{2}}\xi - K\mathcal{C}K^*\mathcal{C}_1^{-1}z - K\mathcal{C}K^*\mathcal{C}_1^{-\frac{1}{2}}\xi - K\mathcal{C}^{\frac{1}{2}}\zeta)\big\|_{\mathbb{R}^N}^2 \\
&= \big\|(I-S)\mathcal{C}_1^{-\frac{1}{2}}z - Q\zeta + (I-S)\xi\big\|_{\mathbb{R}^N}^2 \\
&= \big\|(I-S)(\mathcal{C}_1^{-\frac{1}{2}}z - (I-S)^{-1}Q\zeta + \xi)\big\|_{\mathbb{R}^N}^2
\end{aligned}$$

where $Q := \mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}^{\frac{1}{2}}$ and $S = QQ^*$. Noting that $S = Q_0(I + Q_0^*Q_0)^{-1}Q_0^*$, where $Q_0 := \mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_0^{\frac{1}{2}}$, we have by Lemma 4.9.8(iii) and Lemma 4.9.7(ii) that the eigenvalues of S are bounded above by a non-negative constant $c < 1$ uniformly in N, hence $(I-S)^{-1}$ has eigenvalues which are bounded from above by $\frac{1}{1-c}$ uniformly in N, that is, $\big\|(I-S)^{-1}v\big\|_{\mathbb{R}^N} \leq \frac{1}{1-c}\big\|v\big\|_{\mathbb{R}^N}$ uniformly in N. Moreover, we have

$$\big\|\mathcal{C}_1^{-\frac{1}{2}}(Ku_\sigma^{(k)} - \boldsymbol{y_N})\big\|_{\mathbb{R}^N}^2 \geq (1-c)\big\|\mathcal{C}_1^{-\frac{1}{2}}z - (I-S)^{-1}Q\zeta + \xi\big\|_{\mathbb{R}^N}^2,$$

and it suffices to show that $\mathbb{E}^{\xi,\zeta}[Y_N^{-2i}] = \mathcal{O}(N^{-2i})$, for $Y_N := \big\|\mathcal{C}_1^{-\frac{1}{2}}z - (I-S)^{-1}Q\zeta + \xi\big\|_{\mathbb{R}^N}^2$. Indeed, let $\{e_j\}_{j\in\mathbb{N}}$ be any orthonormal basis of $\mathbb{R}^N$ (with respect to the possibly scaled norm $\|\cdot\|_{\mathbb{R}^N}$), and define $w_j := \big\langle w, e_j\big\rangle_{\mathbb{R}^N}$. Then as in Lemma 4.9.1, since ξ and ζ are independent white noises, we have that if we fix ζ, then Y_N is a non-central chi-squared random variable with N degrees of freedom and non-centrality parameter $p_N(\zeta) = \sum_{j=1}^N(\mathcal{C}_1^{-\frac{1}{2}}z - (I-S)^{-1}Q\zeta)_j^2 \geq 0$. The moment

generating function of Y_N for fixed ζ is

$$M_{Y_N}(t;\zeta) = (1-2t)^{-\frac{N}{2}} \exp(\frac{p_N(\zeta)t}{1-2t}),$$

hence using (4.9.1) we have for $i = 1, 2$,

$$\mathbb{E}^{\xi}[Y_N^{-2i}] \leq c \int_0^{\infty} t^{2i-1}(1+2t)^{-\frac{N}{2}} dt = \mathcal{O}(N^{-2i}),$$

where the constant does not depend on ζ, hence $\mathbb{E}^{\xi,\zeta}[Y_N^{-2i}] \leq \mathcal{O}(N^{-2i})$. □

4.9.4 Technical lemmas

Lemma 4.9.3. *Let $\{X_j\}$ be a sequence of random variables, such that $X_j = c_j Y_j$, where the Y_j, $j \in \mathbb{N}$ are independent and identically distributed random variables with finite even moments up to order $2r \in \mathbb{N}$ and zero odd moments, and the c_j, $j \in \mathbb{N}$ are deterministic real numbers. Then for any $N \in \mathbb{N}$,*

$$\mathbb{E}[(\sum_{j=1}^{N} X_j)^{2r}] \leq \kappa(\sum_{j=1}^{N} c_j^2)^r,$$

where $\kappa = \mathbb{E}[Y_1^{2r}] > 0$ is independent of N.

Proof. Denote by m_n the $2n$-th moment of Y_1, $m_n = \mathbb{E}[Y_1^{2n}]$. Observe that since by Hölder's inequality for $0 < s \leq t$, $\mathbb{E}[|Y_1|^s]^{\frac{1}{s}} \leq \mathbb{E}[|Y_1|^t]^{\frac{1}{t}}$, we have that for $n_1, ..., n_q > 0$ such that $n_1 + ... + n_q = r$

$$m_{n_1}...m_{n_q} \leq \mathbb{E}[Y_1^{2r}]^{\frac{n_1+...+n_q}{r}} = \mathbb{E}[Y_1^{2r}].$$

Using this and the fact that the random variables Y_j are independent with zero odd moments, we get

$$\mathbb{E}[(\sum_{j=1}^{N} X_j)^{2r}] = \sum_{j=1}^{N} c_j^{2r} m_r + \sum_{j_1 \neq j_2}^{N} c_{j_1}^{2(r-1)} m_{r-1} c_{j_2}^2 m_1 + \sum_{j_1 \neq j_2}^{N} c_{j_1}^{2(r-2)} m_{r-2} c_{j_2}^4 m_2$$
$$+ ... + \sum_{j_1 \neq j_2 \neq ... \neq j_r}^{N} c_{j_1}^2 c_{j_2}^2 ... c_{j_r}^2 m_1^r \leq m_r (\sum_{j=1}^{N} c_j^2)^r.$$

□

Lemma 4.9.4. *For any $p \in \mathbb{N}$, there exists a constant $c = c(p) \geq 0$, independent of*

N such that for any centered Gaussian random variable x_N in $\mathbb{R}^N$, it holds

$$\mathbb{E}[\|x_N\|_{\mathbb{R}^N}^{2p}] \leq c(p)(\mathbb{E}[\|x_N\|_{\mathbb{R}^N}^2])^p.$$

Proof. Direct consequence of [18, Corollary 2.17]. □

Lemma 4.9.5. *Let $(\gamma_j)_{j\in\mathbb{N}}$ be a sequence of independent standard normal random variables and define*

$$G_N := \frac{1}{\sqrt{2N}} \sum_{j=1}^{N} (\gamma_j^2 - 1).$$

Then all the integer moments of G_N are bounded uniformly in N.

Proof. For $k \in \mathbb{N}$, we have

$$\mathbb{E}[G_N^k] = \frac{1}{(2N)^{\frac{k}{2}}} \sum_{j_1,\ldots,j_k}^{N} \mathbb{E}[(\gamma_{j_1}^2 - 1)...(\gamma_{j_k}^2 - 1)].$$

Since $\gamma_j^2 - 1$ are independent and identically distributed with finite moments of every order, the sum on the right hand side has a dependence on N which is determined by the total number of non zero terms in the summation. By independence and the fact that $\mathbb{E}[\gamma_j^2 - 1] = 0$, all the terms in the sum which contain a term with an index j_i which occurs only once in the product are equal to zero. We thus have that if k is even the sum on the right hand side is of order $N^{\frac{k}{2}}$, while if k is odd it is of order $N^{\frac{k-1}{2}}$. In both cases the k-th moment of G_N is bounded uniformly in N. □

Lemma 4.9.6. *Let $\Gamma_N \sim \text{Gamma}(\alpha + \frac{N}{2}, \frac{N}{2})$, for $\alpha > 0$, and define*

$$\Theta_N := \frac{\Gamma_N - 1 - \frac{2\alpha}{N}}{\sqrt{\frac{2}{N} + \frac{4\alpha}{N^2}}}.$$

Then the first four moments of Θ_N are bounded uniformly in N.

Proof. By [38] the random variable $\text{Gamma}(a, 1)$ has mean and variance a and third and fourth central moments $2a$ and $3a^2 + 6a$ respectively. Hence by the scaling property of the Gamma distribution, $\Gamma_N \overset{\mathcal{L}}{=} \frac{2}{N}\text{Gamma}(\alpha + \frac{N}{2}, 1)$ has mean $1 + \frac{2\alpha}{N}$, variance $\frac{2}{N} + \frac{4\alpha}{N^2}$, and third and fourth central moments which are both of order N^{-2}. It is thus straightforward to see that Θ_N has mean zero, variance equal to one, and since the denominator in Θ_N is of order $N^{-\frac{1}{2}}$ it has third and fourth moments which are $\mathcal{O}(N^{-\frac{1}{2}})$ and $\mathcal{O}(1)$ respectively. □

Lemma 4.9.7. *For any $M \times N$ matrix A and any $c > 0$, it holds that*

i) $\mathrm{Tr}(A(I + cA^*A)^{-1}A^*) \leq \mathrm{Tr}(A^*A)$;

ii) $\left\|A(I + A^*A)^{-1}A^*\right\|_{2,M} \leq \frac{\left\|A^*A\right\|_{2,N}}{1+c\left\|A^*A\right\|_{2,N}}$.

Proof. By the Singular Value Decomposition, we have that there exist unitary $M \times M$ and $N \times N$ matrices U and V respectively such that $A = U\Sigma V^*$ where Σ is a diagonal $M \times N$ matrix with the square roots of the non-zero eigenvalues $\rho_1^2, ..., \rho_q^2$ of both A^*A and AA^*, where $q = \mathrm{rank}(A)$, and possibly zeros on the diagonal. We thus have that $A(I + cA^*A)^{-1}A^* = U\Sigma(I + c\Sigma^*\Sigma)^{-1}\Sigma^*U^*$ where by the diagonal form of Σ, the $M \times M$ matrix $\Sigma(I + c\Sigma^*\Sigma)^{-1}\Sigma^*$ is also diagonal with $\frac{\rho_1^2}{1+c\rho_1^2}, ..., \frac{\rho_q^2}{1+c\rho_q^2}$ and possibly zeros on the diagonal. Thus the eigenvalues of $A(I + cA^*A)^{-1}A^*$ are $\frac{\rho_1^2}{1+c\rho_1^2}, ..., \frac{\rho_q^2}{1+c\rho_q^2}$ and possibly zero and both claims follow immediately. □

Lemma 4.9.8. *Under Assumptions 4.2.1, we have that for any $\sigma, \delta > 0$,*

i) $\mathrm{Tr}(\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta} K^* \mathcal{C}_1^{-\frac{1}{2}}) \leq c_2 \delta^{-1}$;

ii) $\mathbb{E}^\theta \left\|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} \theta\right\|_{\mathbb{R}^N}^2 = \mathbb{E}^\theta \left\|\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} K^* \mathcal{C}_1^{-\frac{1}{2}} \theta\right\|_{\mathbb{R}^N}^2 \leq c_2 \delta^{-1}$, *where θ is a Gaussian white noise in $\mathbb{R}^N$;*

iii) $\left\|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0^{\frac{1}{2}}\right\|_{2,N} \leq \sqrt{c_2}$;

iv) $\left\|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta} K^* \mathcal{C}_1^{-\frac{1}{2}}\right\|_{2,N} \leq c_2 \delta^{-1}$,

where c_2 is defined in Assumption 4.2.1(ii).

Proof.

i) By (4.1.4), we have

$$\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta} K^* \mathcal{C}_1^{-\frac{1}{2}} = \delta^{-1} \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0^{\frac{1}{2}} (I + \frac{\sigma}{\delta} \mathcal{C}_0^{\frac{1}{2}} K^* \mathcal{C}_1^{-1} K \mathcal{C}_0^{\frac{1}{2}})^{-1} \mathcal{C}_0^{\frac{1}{2}} K^* \mathcal{C}_1^{-\frac{1}{2}},$$

hence Lemma 4.9.7(i) together with Assumption 4.2.1(ii) and the fact that for any matrix A it holds $\mathrm{Tr}(A^*A) = \mathrm{Tr}(AA^*)$ give the claim.

ii) It is well known that for $x \sim \mathcal{N}(0, \Sigma)$, $\mathbb{E}\|x\|_{\mathbb{R}^N}^2 = \mathrm{Tr}(\Sigma)$. Since for $\theta \sim \mathcal{N}(0, I)$ we have $\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} \theta \sim \mathcal{N}(0, \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\sigma,\delta} K^* \mathcal{C}_1^{-\frac{1}{2}})$ and $\mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} K^* \mathcal{C}_1^{-\frac{1}{2}} \theta \sim \mathcal{N}(0, \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}} K^* \mathcal{C}_1^{-1} K \mathcal{C}_{\sigma,\delta}^{\frac{1}{2}})$, the claim follows from part (i), using again the fact that for any matrix A it holds $\mathrm{Tr}(A^*A) = \mathrm{Tr}(AA^*)$.

iii) It is well known that the Euclidean norm of a matrix A, is equal to the Euclidean norm of its adjoint matrix and also equal to the square root of the largest eigenvalue of A^*A, which by the non-negativity of A^*A is smaller than the square root of the trace of A^*A. Hence we have $\big\|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0^{\frac{1}{2}}\big\|_{2,N} \leq \sqrt{\mathrm{Tr}(\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0 K^* \mathcal{C}_1^{-\frac{1}{2}})} \leq \sqrt{c_2}$, by Assumption 4.2.1(ii).

iv) Follows from Lemma 4.9.8(i) since for any matrix A it holds $\big\|A^*A\big\|_{2,N} = \big\|A\big\|_{2,N}^2$, and by applying the same reasoning which gives part (iii) from Assumption 4.2.1(ii).

□

4.9.5 Lemmas for Subsection 4.5.1

Lemma 4.9.9. *Under Assumptions 4.5.1, for any $\delta > 0$, we have*

$$\beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\big\|_{\mathbb{R}^N}^2 = \frac{N}{2\delta} + \sqrt{\frac{N}{2\delta}} W_{1,N} + \tilde{F}_N(\delta), \tag{4.9.2}$$

where i) $W_{1,N}$ only depends on the white noise ζ in (4.5.4), has mean zero and variance one, higher order moments which are bounded uniformly in N, and converges weakly to a standard normal random variable as $N \to \infty$; ii) $\tilde{F}_N(\delta)$ depends on the data and y-almost surely has finite moments of all positive orders uniformly in N (where the expectation is taken with respect to ζ).

Proof. In order to de-clutter our notation we drop the dependence of $m_{\sigma,\delta}, R$ and $\mathcal{C}_{\sigma,\delta}$ on the data. We have,

$$\begin{aligned}\big\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\big\|_{\mathbb{R}^N}^2 &= \big\|\mathcal{C}_0^{-\frac{1}{2}} m_\delta\big\|_{\mathbb{R}^N}^2 + \big\|\mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_\delta^{\frac{1}{2}} \zeta\big\|_{\mathbb{R}^N}^2 + 2\big\langle \mathcal{C}_0^{-\frac{1}{2}} m_\delta, \mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_\delta^{\frac{1}{2}} \zeta\big\rangle_{\mathbb{R}^N} \\ &:= A_N + B_N + C_N.\end{aligned}$$

Under the Assumptions 4.5.1, we can analyze each term as follows:

A) by Assumption 4.5.1(i), for almost all data y, this term and all its positive integer powers are bounded uniformly in N.

B) the second term can be written as

$$\begin{aligned}\big\|\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}}\zeta\big\|_{\mathbb{R}^N}^2 &= \big\langle \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}}\zeta, \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}}\zeta\big\rangle_{\mathbb{R}^N} = \big\langle \mathcal{C}_\delta^{\frac{1}{2}}\mathcal{C}_0^{-1}\mathcal{C}_\delta^{\frac{1}{2}}\zeta, \zeta\big\rangle_{\mathbb{R}^N} \\ &= \delta^{-1}\big\langle \mathcal{C}_\delta^{\frac{1}{2}}(\mathcal{C}_\delta^{-1} - R^*R)\mathcal{C}_\delta^{\frac{1}{2}}\zeta, \zeta\big\rangle_{\mathbb{R}^N} \\ &= \delta^{-1}\big\langle (I - \mathcal{C}_\delta^{\frac{1}{2}}R^*R\mathcal{C}_\delta^{\frac{1}{2}})\zeta, \zeta\big\rangle_{\mathbb{R}^N} \\ &= \delta^{-1}\big\|\zeta\big\|_{\mathbb{R}^N}^2 - \delta^{-1}\big\|R\mathcal{C}_\delta^{\frac{1}{2}}\zeta\big\|_{\mathbb{R}^N}^2 \\ &:= b_{1,N} - b_{2,N},\end{aligned}$$

where

b1) for the first term we have

$$b_{1,N} = \delta^{-1}\big\|\zeta\big\|_{\mathbb{R}^N}^2 = \frac{N}{\delta} + \frac{1}{\delta}\sum_{j=1}^{N}(\zeta_j^2 - 1) := \frac{N}{\delta} + \frac{\sqrt{2N}}{\delta}W_{1,N},$$

where as $N \to \infty$, $W_{1,N} = \frac{1}{\sqrt{2N}}\sum_{j=1}^{N}(\zeta_j^2 - 1)$ converges weakly to a standard normal random variable by the Central Limit Theorem and by Lemma 4.9.5 has moments of every order which are bounded uniformly in N;

b2) for the second term we have by Lemma 4.9.11(ii), that for almost all data $\mathbb{E}^\zeta[b_{2,N}]$ is uniformly bounded in N. In fact using Lemma 4.9.4 together with Lemma 4.9.11(ii), we get that for any $p \in \mathbb{N}$, $\mathbb{E}^\zeta[b_{2,N}^p]$ is bounded independently of N, almost surely with respect to y.

C) for the third term we have

$$\begin{aligned}\big\langle \mathcal{C}_0^{-\frac{1}{2}}m_\delta, \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}}\zeta\big\rangle_{\mathbb{R}^N} &= \big\langle (\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}})^*\mathcal{C}_0^{-\frac{1}{2}}m_\delta, \zeta\big\rangle_{\mathbb{R}^N} \\ &= \sum_{j=1}^{N}((\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}})^*\mathcal{C}_0^{-\frac{1}{2}}m_\delta)_j\zeta_j,\end{aligned}$$

where for any $w \in \mathbb{R}^N$ we write $w_j := \big\langle w, e_j\big\rangle$. It holds that

$$\sum_{j=1}^{N}((\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}})^*\mathcal{C}_0^{-\frac{1}{2}}m_\delta)_j^2 = \big\|(\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}})^*\mathcal{C}_0^{-\frac{1}{2}}m_\delta\big\|_{\mathbb{R}^N}^2,$$

and we claim that the norm on the right hand side is uniformly bounded in N almost surely with respect to the data. Indeed, by (4.5.2), Cauchy-Schwarz and the positive definiteness of the matrix $\mathcal{C}_0^{\frac{1}{2}}R^*R\mathcal{C}_0^{\frac{1}{2}}$, we have almost surely

with respect to y that

$$\begin{aligned}\big\|(\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta^{\frac{1}{2}})^* u\big\|_{\mathbb{R}^N}^2 &= \big\langle \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_\delta \mathcal{C}_0^{-\frac{1}{2}} u, u\big\rangle_{\mathbb{R}^N} \\ &= \big\langle \delta^{-1}(I+\delta^{-1}\mathcal{C}_0^{\frac{1}{2}} R^* R \mathcal{C}_0^{\frac{1}{2}})^{-1} u, u\big\rangle_{\mathbb{R}^N} \\ &\leq \big\|\delta^{-1}(I+\delta^{-1}\mathcal{C}_0^{\frac{1}{2}} R^* R \mathcal{C}_0^{\frac{1}{2}})^{-1} u\big\|_{\mathbb{R}^N} \big\|u\big\|_{\mathbb{R}^N} \\ &\leq \delta^{-1}\big\|u\big\|_{\mathbb{R}^N}^2.\end{aligned}$$

Combining with Assumption 4.5.1(i) we get the claim and therefore by Lemma 4.9.3 we get that the third term has y-almost surely all even moments uniformly bounded in N.

We define $\tilde{F}_N = \beta_0 + \frac{A_N - b_{2,N} + C_N}{2}$ and observe that since all terms have bounded moments of every order uniformly in N y-almost surely, Hölder's inequality secures that $\tilde{F}_N$ also has bounded moments of every order uniformly in N almost surely with respect to y. □

Lemma 4.9.10. *Let $u_\delta^{(k)}$ be as in (4.5.4), for any $\delta > 0$. Under Assumptions 4.5.1, for $i = 1, 2$, we have*

$$\mathbb{E}^\zeta\left[(\beta_0 + \frac{1}{2}\big\|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\big\|_{\mathbb{R}^N}^2)^{-2i}\right] = \mathcal{O}(N^{-2i}),$$

as $N \to \infty$, almost surely with respect to y.

Proof. The proof is almost identical to the proof of Lemma 4.9.10, however we include it for completeness. Without loss of generality we consider the case $\delta = 1$ and drop the δ dependence in u, m and $\mathcal{C}$. To de-clutter our notation we also drop the dependence of m and $\mathcal{C}$ on the data. Since $\beta_0 \geq 0$ it suffices to show it for $\beta_0 = 0$. Formally, the random variable $\big\|\mathcal{C}_0^{-\frac{1}{2}} u^{(k)}\big\|_{\mathbb{R}^N}^2$ behaves like a chi-squared random variable with N degrees of freedom. We estimate the squared norm by a random variable Y_N of known moment generating function $M_{Y_N}(t)$, and use (4.9.1) for the calculation of negative moments of nonnegative random variables.

We begin by showing that for almost all y there exists $c = c(y) > 0$ indepen-

dent of N such that $\|\mathcal{C}^{-\frac{1}{2}}\mathcal{C}_0^{\frac{1}{2}}v\|_{\mathbb{R}^N} \leq c\|v\|_{\mathbb{R}^N}$ for any $v \in \mathbb{R}^N$. We have,

$$\begin{aligned}\|\mathcal{C}^{-\frac{1}{2}}\mathcal{C}_0^{\frac{1}{2}}v\|_{\mathbb{R}^N}^2 &= \langle \mathcal{C}_0^{\frac{1}{2}}\mathcal{C}^{-1}\mathcal{C}_0^{\frac{1}{2}}v, v\rangle_{\mathbb{R}^N} \\ &= \langle (I + \mathcal{C}_0^{\frac{1}{2}}R^*R\mathcal{C}_0^{\frac{1}{2}})v, v\rangle_{\mathbb{R}^N} \\ &= \|v\|_{\mathbb{R}^N}^2 + \|R\mathcal{C}_0^{\frac{1}{2}}v\|_{\mathbb{R}^N}^2 \\ &\leq (1 + c_2')\|v\|_{\mathbb{R}^N}^2,\end{aligned}$$

by Lemma 4.9.11(iii). The proved claim gives the estimate

$$\begin{aligned}\|\mathcal{C}_0^{-\frac{1}{2}}u^{(k)}\|_{\mathbb{R}^N}^2 = \|\mathcal{C}_0^{-\frac{1}{2}}(m + \mathcal{C}^{\frac{1}{2}}\zeta)\|_{\mathbb{R}^N}^2 &= \|\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}^{\frac{1}{2}}(\mathcal{C}^{-\frac{1}{2}}m + \zeta)\|_{\mathbb{R}^N}^2 \\ &\geq c^{-1}\|\mathcal{C}^{-\frac{1}{2}}m + \zeta\|_{\mathbb{R}^N}^2,\end{aligned}$$

hence it suffices to show that almost surely with respect to y we have $\mathbb{E}^\zeta[Y_N^{-2i}] = \mathcal{O}(N^{-2i})$, for $Y_N := \|\mathcal{C}^{-\frac{1}{2}}m + \zeta\|_{\mathbb{R}^N}^2$. Indeed, let $\{e_j\}_{j=1}^N$ be any orthonormal basis of $\mathbb{R}^N$ (with respect to the possibly scaled norm $\|\cdot\|_{\mathbb{R}^N}$), and define $w_j := \langle w, e_j\rangle$ for any $w \in \mathbb{R}^N$. Then we have

$$Y_N = \sum_{j=1}^N ((\mathcal{C}^{-\frac{1}{2}}m)_j + \zeta_j)^2,$$

where $\zeta_j \sim \mathcal{N}(0,1)$ are the mutually independent components of the white noise ζ and $(\mathcal{C}^{-\frac{1}{2}}m)_j$ are independent of ζ, therefore Y_N is a non-central chi-squared random variable with N degrees of freedom and non-centrality parameter $p_N := \sum_{j=1}^N (\mathcal{C}^{-\frac{1}{2}}m)_j^2 \geq 0$. The definition and properties of the non-central chi-squared distribution can be found in [39], where in particular, we find that the moment generating function of Y_N is

$$M_{Y_N}(t) = (1-2t)^{-\frac{N}{2}} \exp\big(\frac{p_N t}{1-2t}\big),$$

hence using (4.9.1) we have for $i = 1, 2$,

$$\begin{aligned}\mathbb{E}^\zeta[Y_N^{-2i}] &= \Gamma(2i)^{-1}\int_0^\infty t^{2i-1}(1+2t)^{-\frac{N}{2}} \exp\big(\frac{-p_N t}{1+2t}\big)dt \\ &\leq c\int_0^\infty t^{2i-1}(1+2t)^{-\frac{N}{2}}dt \\ &= \mathcal{O}(N^{-2i}),\end{aligned}$$

provided $N > 4i$, where the last integral can by calculated analytically by integration

by parts. □

Lemma 4.9.11. *Under Assumptions 4.5.1, we have that for any $\delta > 0$, almost surely with respect to y,*

i) $\mathrm{Tr}(R(\boldsymbol{y_N})\mathcal{C}_\delta(\boldsymbol{y_N})R(\boldsymbol{y_N})^*) \leq c_2'\delta^{-1}$;

ii) $\mathbb{E}^\theta \big\| R(\boldsymbol{y_N})\mathcal{C}_\delta^{\frac{1}{2}}(\boldsymbol{y_N})\zeta \big\|_{\mathbb{R}^N}^2 \leq c_2'\delta^{-1}$, *where θ is Gaussian white noise in $\mathbb{R}^N$*;

iii) $\big\| R(\boldsymbol{y_N})\mathcal{C}_0^{\frac{1}{2}} \big\|_{2,N} \leq \sqrt{c_2'}$;

where $c_2' = c_2'(y)$ is defined in Assumptions 4.5.1(ii).

Proof.

i) By (4.5.2), we have

$$R(\boldsymbol{y_N})\mathcal{C}_\delta R(\boldsymbol{y_N})^* = \delta^{-1} R(\boldsymbol{y_N})\mathcal{C}_0^{\frac{1}{2}} \big(I + \delta^{-1}\mathcal{C}_0^{\frac{1}{2}} R(\boldsymbol{y_N})^* R(\boldsymbol{y_N})\mathcal{C}_0^{\frac{1}{2}}\big)^{-1} \mathcal{C}_0^{\frac{1}{2}} R(\boldsymbol{y_N})^*,$$

hence Lemma 4.9.7(i) together with Assumption 4.5.1(ii) and the fact that for any matrix A it holds $\mathrm{Tr}(A^*A) = \mathrm{Tr}(AA^*)$ give the claim.

ii) It is well known that for $x \sim \mathcal{N}(0, \Sigma)$, $\mathbb{E}\big\|x\big\|_{\mathbb{R}^N}^2 = \mathrm{Tr}(\Sigma)$. Since $R(\boldsymbol{y_N})\mathcal{C}_\delta^{\frac{1}{2}}\theta \sim \mathcal{N}(0, R(\boldsymbol{y_N})\mathcal{C}_\delta R(\boldsymbol{y_N})^*)$, the claim follows from part (i).

iii) It is well known that the Euclidean norm of a matrix A, is equal to the Euclidean norm of its adjoint matrix and also equal to the square root of the largest eigenvalue of A^*A, which by the non-negativity of A^*A is smaller than the square root of the trace of A^*A. Hence we have $\big\| R(\boldsymbol{y_N})\mathcal{C}_0^{\frac{1}{2}} \big\|_{2,N} \leq \sqrt{\mathrm{Tr}(R(\boldsymbol{y_N})\mathcal{C}_0 R(\boldsymbol{y_N})^*)} \leq \sqrt{c_2'}$, by Assumption 4.5.1(ii).

□

Appendix A

Posterior contraction rates for severely ill-posed problems with Gaussian priors of analytic regularity

We apply the theory developed in Chapter 2 to a diagonal setup where $\mathcal{C}_0$ has eigenvalues $\lambda_j \asymp \exp(-2\alpha j)$, $\mathcal{C}_1$ is the identity so that the noise is white, and $\mathcal{A}^{-1} = \mathcal{C}_0^{\ell}$ for some $\ell > 0$. That is, using the terminology of [14], we study a severely ill-posed problem with degree of ill-posedness $\hat{\ell} = 2\alpha\ell$. Let A^t denote the spaces of analytic class t as defined in [14],

$$A^t = \left\{ u \in \mathcal{X} : \sum_{j=1}^{\infty} \exp(2tj) \langle u, \phi_j \rangle^2 < \infty \right\},$$

where we recall that ϕ_j are the eigenfunctions of $\mathcal{C}_0$ which form a complete orthonormal basis in $\mathcal{X}$. We then have that the spaces X^t are identified with the spaces $A^{\alpha t}$. Since the eigenvalues of $\mathcal{C}_0$ decay exponentially, we have that $s_0 = 0$; by Lemma 2.3.3(ii) we have that draws from the prior belong to $A^{\alpha(1-s)}$ with probability one, for any $s > 0$. Furthermore, $\Delta = 1 + 2\ell > 2s_0 = 0$, hence the Assumptions 2.3.1 are satisfied and we can apply Theorem 2.2.3 to get the following convergence result:

Theorem A.0.12. *Assume that $u^\dagger \in A^{\hat{\gamma}}$, where $\hat{\gamma} \geq \alpha$. Then for*

$$\tau = \tau(n) = n^{-\frac{\hat{\gamma} \wedge (2\alpha + \hat{\ell}) - \alpha + \varepsilon}{2\hat{\ell} + 2(\hat{\gamma} \wedge (2\alpha + \hat{\ell})) + 2\varepsilon}},$$

where $\varepsilon > 0$ arbitrarily small, we have that the convergence in (2.1.12) holds with $\epsilon_n = n^{-e}$, where

$$e = \frac{\hat{\gamma} \wedge (2\alpha + \hat{\ell})}{2\hat{\ell} + 2(\hat{\gamma} \wedge (2\alpha + \hat{\ell})) + 2\varepsilon}.$$

In particular, for analytic regularity $\hat{\gamma}$ which is smaller than $2\alpha + \hat{\ell}$, we have that if we rescale the prior appropriately the posterior contracts to the truth in the small noise limit at the optimal minimax rate (up to $\varepsilon > 0$ arbitrarily small), [14].

Bibliography

[1] R. J. Adler. *An introduction to continuity, extrema, and related topics for general Gaussian processes.* Institute of Mathematical Statistics Lecture Notes—Monograph Series, 12. Institute of Mathematical Statistics, Hayward, CA, 1990.

[2] S. Agapiou, J. M. Bardsley, O. Papaspiliopoulos, and A. M. Stuart. Analysis of the Gibbs sampler for hierarchical inverse problems. *Arxiv preprint: 1311.1138.*

[3] S. Agapiou, S. Larsson, and A. M. Stuart. Posterior contraction rates for the Bayesian approach to linear ill-posed inverse problems. *Stoch. Proc. Apl.*, 123(10):3828–3860, 2013.

[4] S. Agapiou, A. M. Stuart, and Y. X. Zhang. Bayesian posterior contraction rates for linear severely ill-posed inverse problems. *Arxiv preprint: 1210.1563 (to appear, J. Inverse Ill-posed Probl.).*

[5] H. Attouch, G. Buttazzo, and G. Michaille. *Variational Analysis in Sobolev and BV Spaces: Applications to PDEs and Optimization.* SIAM, 2006.

[6] H. T. Banks and K. Kunisch. *Estimation techniques for distributed parameter systems.* Birkhauser Boston, 1989.

[7] J. M. Bardsley. MCMC-based image reconstruction with uncertainty quantification. *SIAM J. Sci. Comput.*, 34(3):A1316–A1332, 2012.

[8] J. M. Bardsley. Gaussian Markov Random Field priors for inverse problems. *Inverse Problems and Imaging*, 7(2):397–416, 2013.

[9] A. Barron, M. J. Schervish, and L. Wasserman. The consistency of posterior distributions in nonparametric problems. *The Annals of Statistics*, 27(2):536–561, 1999.

[10] A. Beskos, F. J. Pinski, J. M. Sanz-Serna, and A. M. Stuart. Hybrid Monte Carlo on Hilbert spaces. *Stochastic Processes and their Applications*, 121(10):2201–2230, 2011.

[11] A. Beskos, G. Roberts, and A. Stuart. Optimal scalings for local Metropolis-Hastings chains on nonproduct targets in high dimensions. *Ann. Appl. Probab.*, 19(3):863–898, 2009.

[12] N. Bochkina. Consistency of the posterior distribution in generalized linear inverse problems. *Inverse Problems*, 29(9), 2013.

[13] V. I. Bogachev. *Gaussian measures*, volume 62. Amer Mathematical Society, 1998.

[14] L Cavalier. Nonparametric statistical inverse problems. *Inverse Problems*, 24(3):034004, 2008.

[15] S. Cotter, G. Roberts, A. Stuart, and D. White. Mcmc methods for functions: modifying old algorithms to make them faster. *arXiv preprint arXiv:1202.0709*, 2012.

[16] N. Cressie, A. S. Davis, J. L. Folks, and G. E. Policello II. The moment-generating function and negative integer moments. *The American Statistician*, 35(3):148–150, 1981.

[17] G. Da Prato. *An introduction to infinite-dimensional analysis*. Springer, 2005.

[18] G. Da Prato and J. Zabczyk. Stochastic equations in infinite dimensions, volume 44 of encyclopedia of mathematics and its applications, 1992.

[19] M. Dashti, S. Harris, and A. Stuart. Besov priors for Bayesian inverse problems. *Inverse Probl. Imaging*, 6(2):183–200, 2012.

[20] P. Diaconis and D. Freedman. On the consistency of Bayes estimates. *The Annals of Statistics*, pages 1–26, 1986.

[21] J. L. Doob. Application of the theory of martingales. *Le calcul des probabilites et ses applications*, pages 23–27, 1949.

[22] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of inverse problems*, volume 375. Springer, 1996.

[23] S. N. Ethier and T. G. Kurtz. Markov processes: Characterization and convergence. *New York: Wiley*, 1986.

[24] J. P. Florens and A. Simoni. Regularizing priors for linear inverse problems. *TSE Working Paper*, 10, 2010.

[25] J. N. Franklin. Well-posed stochastic extensions of ill-posed linear problems. *Journal of Mathematical Analysis and Applications*, 31(3):682–716, 1970.

[26] D. A. Freedman. On the asymptotic behavior of Bayes' estimates in the discrete case. *The Annals of Mathematical Statistics*, 34(4):1386–1403, 1963.

[27] S. Ghosal, J. K. Ghosh, and R. V. Ramamoorthi. Posterior consistency of Dirichlet mixtures in density estimation. *Annals of Statistics*, 27(1):143–158, 1999.

[28] S. Ghosal, J. K. Ghosh, and A. W. van der Vaart. Convergence rates of posterior distributions. *Annals of Statistics*, 28(2):500–531, 2000.

[29] E. Giné and R. Nickl. Rates of contraction for posterior distributions in L^r-metrics, $1 \leq r \leq \infty$. *Ann. Statist.*, 39(6):2883–2911, 2011.

[30] M. Hairer, J. C. C Mattingly, and M. Scheutzow. Asymptotic coupling and a general form of Harris' theorem with applications to stochastic delay equations. *Probability Theory and Related Fields*, 149(1-2):223–259, 2011.

[31] M. Hairer, A. M. Stuart, and S. J. Vollmer. Spectral gaps for metropolis-hastings algorithms in infinite dimensions. Submitted.

[32] M. Hairer, A. M. Stuart, and J. Voss. Analysis of SPDEs arising in path sampling. II. The nonlinear case. *Ann. Appl. Probab.*, 17(5-6):1657–1706, 2007.

[33] G. Helmberg. *Introduction to spectral theory in Hilbert space*. North-Holland Amsterdam, 1969.

[34] A. Hofinger and H. K. Pikkarainen. Convergence rate for the Bayesian approach to linear inverse problems. *Inverse Problems*, 23(6):2469, 2007.

[35] A. Hofinger and H. K. Pikkarainen. Convergence rates for linear inverse problems in the presence of an additive normal noise. *Stochastic Analysis and Applications*, 27(2):240–257, 2009.

[36] T. Hohage. Regularization of exponentially ill-posed problems. *Numerical functional analysis and optimization*, 21(3-4):439–464, 2000.

[37] T. M. Huang. Convergence rates for posterior distributions and adaptive estimation. *The Annals of Statistics*, 32(4):1556–1593, 2004.

[38] N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous univariate distributions. Vol. 1.* Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1994. A Wiley-Interscience Publication.

[39] N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous univariate distributions. Vol. 2.* Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1995. A Wiley-Interscience Publication.

[40] J. P. Kaipio and E. Somersalo. *Statistical and computational inverse problems*, volume 160. Springer, 2005.

[41] O. Kallenberg. *Foundations of modern probability.* Springer, 2002.

[42] A. Kirsch. *An introduction to the mathematical theory of inverse problems*, volume 120. Springer, 2011.

[43] B. T. Knapik, B. T. Szabó, A. W. van der Vaart, and J. H. van Zanten. Bayes procedures for adaptive inference in nonparametric inverse problems. *arXiv:1209.3628*, 2012.

[44] B. T. Knapik, A. W. van Der Vaart, and J. H. van Zanten. Bayesian inverse problems with Gaussian priors. *The Annals of Statistics*, 39(5):2626–2657, 2011.

[45] B. T. Knapik, A. W. van der Vaart, and J. H. van Zanten. Bayesian recovery of the initial condition for the heat equation. *Communications in Statistics-Theory and Methods*, 42(7):1294–1313, 2013.

[46] S. Lasanen. Measurements and infinite-dimensional statistical inverse theory. *PAMM*, 7(1):1080101–1080102, 2007.

[47] S. Lasanen. Non-Gaussian statistical inverse problems. Part I: Posterior distributions. *Inverse Problems and Imaging*, 6(2):215–266, 2012.

[48] S. Lasanen. Non-Gaussian statistical inverse problems. Part II: Posterior convergence for approximated unknowns. *Inverse Problems and Imaging*, 6(2):267–287, 2012.

[49] M. Lassas, E. Saksman, and S. Siltanen. Discretization-invariant Bayesian inversion and Besov space priors. *Inverse Probl. Imaging*, 3(1):87–122, 2009.

[50] P. D. Lax. *Functional analysis*. Pure and Applied Mathematics (New York). Wiley-Interscience [John Wiley & Sons], New York, 2002.

[51] P. D. Lax. *Linear algebra and its applications*. Pure and Applied Mathematics (Hoboken). Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, second edition, 2007.

[52] M. S. Lehtinen, L. Paivarinta, and E. Somersalo. Linear inverse problems for generalised random variables. *Inverse Problems*, 5(4):599, 1989.

[53] A. Lunardi. *Interpolation theory*. Edizioni Della Normale, 2009.

[54] A. Mandelbaum. Linear estimators and measurable linear transformations on a Hilbert space. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 65(3):385–397, 1984.

[55] J. C. Mattingly, N. S. Pillai, and A. M. Stuart. Diffusion limits of the random walk Metropolis algorithm in high dimensions. *The Annals of Applied Probability*, 22(3):881–930, 2012.

[56] F. van der Meulen, M. Schauer, and J. H. van Zanten. Reversible jump MCMC for nonparametric drift estimation for diffusion processes. *Computational Statistics & Data Analysis*, 2013.

[57] S. S. P Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Cambridge University Press, 2009.

[58] A. Neubauer and H. K. Pikkarainen. Convergence results for the Bayesian inversion theory. *Journal of Inverse and Ill-posed Problems*, 16(6):601–613, 2008.

[59] O. Papaspiliopoulos, Y. Pokern, G. O. Roberts, and A. M. Stuart. Nonparametric estimation of diffusions: a differential equations approach. *Biometrika*, 99(3):511–531, 2012.

[60] O. Papaspiliopoulos, G. O. Roberts, and M. Sköld. A general framework for the parametrization of hierarchical models. *Statistical Science*, pages 59–73, 2007.

[61] D. L. Phillips. A technique for the numerical solution of certain integral equations of the first kind. *Journal of the ACM (JACM)*, 9(1):84–97, 1962.

[62] N. S. Pillai, A. M. Stuart, and A. H. Thiery. On the random walk Metropolis algorithm for Gaussian random field priors and the gradient flow. *Arxiv preprint*, 2011.

[63] N. S. Pillai, A. M. Stuart, and A. H. Thiéry. Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions. *The Annals of Applied Probability*, 22(6):2320–2356, 2012.

[64] Y. Pokern, A. M. Stuart, and J. H. van Zanten. Posterior consistency via precision operators for Bayesian nonparametric drift estimation in SDEs. *Stochastic Process. Appl.*, 123(2):603–628, 2013.

[65] K. Ray. Bayesian inverse problems with non-conjugate priors. *arXiv preprint arXiv:1209.6156*, 2012.

[66] M. Reed and B. Simon. *Functional analysis*, volume 1. Elsevier, 1981.

[67] G. O. Roberts, A. Gelman, and W. R. Gilks. Weak convergence and optimal scaling of random walk Metropolis algorithms. *The annals of applied probability*, 7(1):110–120, 1997.

[68] G. O. Roberts and J. S. Rosenthal. Geometric ergodicity and hybrid markov chains. *Electron. Comm. Probab*, 2(2):13–25, 1997.

[69] G. O. Roberts and J. S. Rosenthal. Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(1):255–268, 1998.

[70] G. O. Roberts and J. S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.*, 16(4):351–367, 2001.

[71] G. O Roberts and J. S Rosenthal. General state space Markov chains and MCMC algorithms. *Probability Surveys*, 1:20–71, 2004.

[72] G. O. Roberts and O. Stramer. On inference for partially observed nonlinear diffusion models using the Metropolis–Hastings algorithm. *Biometrika*, 88(3):603–621, 2001.

[73] J. C. Robinson. *Infinite-dimensional dynamical systems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, Cambridge, 2001. An introduction to dissipative parabolic PDEs and the theory of global attractors.

[74] J. Rousseau. Rates of convergence for the posterior distributions of mixtures of betas and adaptive nonparametric estimation of the density. *The Annals of Statistics*, 38(1):146–180, 2010.

[75] F. Schwabl. *Quantum mechanics*. Springer, Berlin, fourth edition, 2007.

[76] L. Schwartz. On Bayes procedures. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 4(1):10–26, 1965.

[77] X. Shen and L. Wasserman. Rates of convergence of posterior distributions. *Ann. Statist.*, 29(3):687–714, 2001.

[78] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numer.*, 19:451–559, 2010.

[79] B. T. Szabó, A. W. van der Vaart, and J. H. van Zanten. Empirical Bayes scaling of Gaussian priors in the white noise model. *Electronic Journal of Statistics*, 7:991–1018, 2013.

[80] A. Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005.

[81] A. N. Tikhonov. Solution of incorrectly formulated problems and the regularization method. In *Soviet Math. Dokl.*, volume 4, pages 1035–1038, 1963.

[82] A. W. van der Vaart and J. H. van Zanten. Bayesian inference with rescaled Gaussian process priors. *Electronic Journal of Statistics*, 1:433–448, 2007.

[83] A. W. van der Vaart and J. H. van Zanten. Rates of contraction of posterior distributions based on Gaussian process priors. *The Annals of Statistics*, pages 1435–1463, 2008.

[84] A. W. van der Vaart and J. H. van Zanten. Adaptive Bayesian estimation using a Gaussian random field with inverse Gamma bandwidth. *The Annals of Statistics*, 37(5B):2655–2675, 2009.

[85] C. R. Vogel. *Computational methods for inverse problems*, volume 10. Siam, 2002.

[86] S. J. Vollmer. Posterior consistency for Bayesian inverse problems through stability and regression results. *arXiv preprint arXiv:1302.4101*, 2013.

[87] K. Yosida. *Functional analysis*. Classics in Mathematics. Springer-Verlag, Berlin, 1995. Reprint of the sixth (1980) edition.

[88] Y. X. Zhang, C. L. Fu, and Z. L. Deng. An a posteriori truncation method for some Cauchy problems associated with Helmholtz-type equations. *Inverse Problems in Science and Engineering*, (ahead-of-print):1–18, 2012.