

ZURICH LECTURES IN ADVANCED MATHEMATICS



Richard Nickl

# Bayesian Non-linear Statistical Inverse Problems



EM  
S  
PRESS



**EM**  
**S** ■  
**PRESS**

## Zurich Lectures in Advanced Mathematics

Edited by Habib Ammari, Alexander Gorodnik (Managing Editor), Urs Lang (Managing Editor), and Michael Struwe

Mathematics in Zurich has a long and distinguished tradition, in which the writing of lecture notes volumes and research monographs plays a prominent part. The *Zurich Lectures in Advanced Mathematics* series aims to make some of these publications better known to a wider audience. The series has three main constituents: lecture notes on advanced topics given by internationally renowned experts, in particular lecture notes of "Nachdiplomvorlesungen", organized jointly by the Department of Mathematics and the Institute for Research in Mathematics (FIM) at ETH, graduate text books designed for the joint graduate program in Mathematics of the ETH and the University of Zürich, as well as contributions from researchers in residence. Moderately priced, concise and lively in style, the volumes of this series will appeal to researchers and students alike, who seek an informed introduction to important areas of current research.

Previously published in this series:

- Y. B. Pesin, *Lectures on partial hyperbolicity and stable ergodicity*  
S.-Y. A. Chang, *Non-linear Elliptic Equations in Conformal Geometry*  
S. B. Kuksin, *Randomly forced nonlinear PDEs and statistical hydrodynamics in 2 space dimensions*  
P. Etingof, *Calogero–Moser systems and representation theory*  
G. Balkema, P. Embrechts, *High Risk Scenarios and Extremes – A geometric approach*  
D. Christodoulou, *Mathematical Problems of General Relativity I*  
C. De Lellis, *Rectifiable Sets, Densities and Tangent Measures*  
P. Seidel, *Fukaya Categories and Picard–Lefschetz Theory*  
A. H. W. Schmitt, *Geometric Invariant Theory and Decorated Principal Bundles*  
M. Farber, *Invitation to Topological Robotics*  
A. Barvinok, *Integer Points in Polyhedra*  
C. Lubich, *From Quantum to Classical Molecular Dynamics: Reduced Models and Numerical Analysis*  
S. Onn, *Nonlinear Discrete Optimization – An Algorithmic Theory*  
K. Nakanishi, W. Schlag, *Invariant Manifolds and Dispersive Hamiltonian Evolution Equations*  
E. Faou, *Geometric Numerical Integration and Schrödinger Equations*  
A.-S. Sznitman, *Topics in Occupation Times and Gaussian Free Fields*  
F. Labourie, *Lectures on Representations of Surface Groups*  
I. Gallagher, L. Saint-Raymond, B. Texier, *From Newton to Boltzmann: Hard Spheres and Short-range Potentials*  
R. J. Marsh, *Lecture Notes on Cluster Algebras*  
E. Hebey, *Compactness and Stability for Nonlinear Elliptic Equations*  
S. Serfaty, *Coulomb Gases and Ginzburg–Landau Vortices*  
A. Figalli, *The Monge–Ampère Equation and Its Applications*  
W. Schachermayer, *Asymptotic Theory of Transaction Costs*  
A. Thomas, *Geometric and Topological Aspects of Coxeter Groups and Buildings*  
T. Fisher, B. Hasseblatt, *Hyperbolic Flows*  
S. Baader, *Geometry and Topology of Surfaces*  
N. Anantharaman, *Quantum Ergodicity and Delocalization of Schrödinger Eigenfunctions*  
X. Fernández-Real, X. Ros-Oton, *Regularity Theory for Elliptic PDE*  
A. Polishchuk,  *$A_\infty$ -Structures and Moduli Spaces*

Published with the support of the Huber-Kudlich-Stiftung, Zürich



Richard Nickl

# Bayesian Non-linear Statistical Inverse Problems



**EM**  
**S** ■  
PRESS

## Author

Richard Nickl  
Department of Pure Mathematics and Mathematical Statistics  
University of Cambridge  
Cambridge CB3 0WB, UK  
Email: [nickl@maths.cam.ac.uk](mailto:nickl@maths.cam.ac.uk)

**2020 Mathematics Subject Classification:** Primary 62G20; Secondary 35R30, 65C05

**Keywords:** Statistical inverse problem, Darcy's problem, Langevin algorithms, Bayesian non-parametric statistics

ISBN 978-3-98547-053-2, eISBN 978-3-98547-553-7, DOI 10.4171/ZLAM/30

**Bibliographic information published by the Deutsche Nationalbibliothek**  
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie;  
detailed bibliographic data are available on the Internet at <http://dnb.dnb.de>.

Published by EMS Press, an imprint of the  
European Mathematical Society – EMS – Publishing House GmbH  
Institut für Mathematik  
Technische Universität Berlin  
Straße des 17. Juni 136  
10623 Berlin, Germany  
<https://ems.press>

© 2023 EMS Press

Typesetting using the author's LaTeX sources: Nikos Tymis, Thessaloniki, Greece  
Printed in Germany  
© Printed on acid free paper

Für Julian



# Preface

Die Schriftsteller  
auch wenn sie Wissenschaftler sind  
sind Übertreibungsspezialisten

---

T. Bernhard

These notes are based on graduate (‘Nachdiplom’) lectures given in spring term 2022 at the Department of Mathematics, ETH Zürich, Switzerland. I am very grateful to the Forschungsinstitut für Mathematik (FIM) at ETH Zürich for hosting me, as well as to Afonso Bandeira and Sara van der Geer for their hospitality during my visit.

Chapters 1 and 2 develop a framework of Bayesian Gaussian process methods in non-linear random design regression models and give sufficient conditions to obtain global convergence guarantees for posterior measures in PDE-type inverse problems. Chapters 3 to 5 develop the local theory about fluctuations and shape of posterior measures in high dimensions, and touch on the related issue of convergence properties of gradient-based MCMC algorithms. The reader will require background in real analysis, measure-theoretic probability and stochastic convergence theory – one may consult [46] or similar texts. We will frequently use mathematical techniques from high-dimensional statistics and probability as developed in [60]. Relevant material from elliptic partial differential equations, functional analysis and stochastic calculus will be reviewed in Appendices A and B.

I would like to thank the many colleagues whose feedback and input over the past years have helped shape my understanding of this material. Let me mention Francis Bach, Afonso Bandeira, Ismaël Castillo, Victor Chernozhukov, Marc Hoffmann, Ankur Moitra, Francois Monard, Kolyan Ray, Sebastian Reich, Markus Reiß, Philippe Rigollet, Judith Rousseau, Christoph Schwab, Vladimir Spokoiny, Botond Szabo, Andrew Stuart, Edriss Titi, Sasha Tsybakov, Sara van de Geer, Aad van der Vaart, Harry van Zanten and Martin Wainwright; as well as current, former and visiting members of my research group, among them Kweku Abraham, Randolf Altmeyer, Jan Bohr, Matteo Giordano, Hanne Kekkonen, Aksel K. Rasmussen, and Sven Wang. I would particularly like to thank Gabriel P. Paternain for the many inspiring discussions about non-linear inverse problems over the past years – these notes owe much to them.

Cambridge, UK, March 2023

Richard Nickl



# Contents

<b>Preface</b> .....	vii
<b>1 Non-linear statistical inverse problems</b> .....	1
1.1 Model examples .....	2
1.1.1 Boundary measurements in tomography .....	2
1.1.2 Parameter identification for elliptic PDEs .....	3
1.1.3 Data assimilation .....	6
1.2 Bayesian regression .....	7
1.2.1 The forward map $\mathcal{G}$ .....	7
1.2.2 A random design regression model with normal errors .....	7
1.2.3 The Bayesian prior and posterior distribution .....	9
1.2.4 Posterior computation by MCMC .....	10
1.3 The frequentist perspective .....	14
1.3.1 Information distances for random design regression .....	14
1.3.2 A first posterior contraction theorem .....	17
1.4 Notes .....	20
1.4.1 Exercises .....	20
1.4.2 Remarks and comments .....	20
<b>2 Global stability and posterior consistency</b> .....	23
2.1 Analytical hypotheses and PDEs .....	24
2.1.1 Forward regularity conditions for $\mathcal{G}$ .....	24
2.1.2 Injectivity and stability estimates .....	27
2.2 Regularisation with Gaussian process priors .....	31
2.3 Convergence of posterior measure and mean .....	35
2.4 Notes .....	38
2.4.1 Exercises .....	38
2.4.2 Remarks and comments .....	39
<b>3 Information operators and curvature</b> .....	41
3.1 Information geometry .....	41
3.1.1 The LAN expansion .....	42
3.1.2 Cramer–Rao bounds and inverse information .....	45
3.2 Gradient stability and concentration .....	49
3.2.1 Convexity of $-\ell_N$ and the gradient of $\mathcal{G}$ .....	49
3.2.2 A concentration result for the empirical Hessian .....	51

3.3	Information operators for elliptic PDEs . . . . .	58
3.3.1	Schrödinger equation . . . . .	58
3.3.2	Diffusion equation . . . . .	59
3.3.3	Injectivity and local identifiability . . . . .	61
3.4	Notes . . . . .	63
3.4.1	Exercises . . . . .	63
3.4.2	Remarks and comments . . . . .	64
<b>4</b>	<b>Bernstein–von Mises theorems . . . . .</b>	<b>67</b>
4.1	Gaussian asymptotics for cylindrical laws . . . . .	68
4.1.1	Asymptotic normality of linear functionals of the posterior . . . . .	68
4.1.2	Asymptotic distribution of the posterior mean . . . . .	77
4.1.3	Applications to uncertainty quantification . . . . .	80
4.2	Solving information equations in PDE models . . . . .	81
4.2.1	A Bernstein–von Mises theorem for the Schrödinger equation . . . . .	82
4.2.2	Impossibility of the BvM-phenomenon for Darcy’s problem . . . . .	84
4.3	Notes . . . . .	88
4.3.1	Exercises . . . . .	88
4.3.2	Remarks and comments . . . . .	89
<b>5</b>	<b>Posteriors are probably log-concave . . . . .</b>	<b>91</b>
5.1	Wasserstein approximation of the posterior . . . . .	92
5.1.1	Construction of a log-concave surrogate posterior . . . . .	94
5.1.2	The log-concave approximation theorem . . . . .	95
5.2	Computational complexity of MCMC in high dimensions . . . . .	101
5.2.1	Gradient methods for approximately log-concave posteriors . . . . .	101
5.2.2	On failure of ‘cold start’ MCMC in high dimensions . . . . .	106
5.3	Application to PDE models . . . . .	113
5.3.1	Darcy’s problem on the eigen-spaces of the Laplacian . . . . .	114
5.3.2	Polynomial time computation of the posterior mean . . . . .	119
5.4	Notes . . . . .	121
5.4.1	Exercises . . . . .	121
5.4.2	Remarks and comments . . . . .	121
<b>A</b>	<b>Analytical background . . . . .</b>	<b>125</b>
A.1	Sobolev and related spaces . . . . .	125
A.2	Elliptic second-order differential operators . . . . .	126
A.3	Orthonormal discretisation of $L^2$ and metric entropy . . . . .	128
A.4	Feynman–Kac formulæ . . . . .	132
A.5	Elliptic regularity estimates . . . . .	133

<b>B Further auxiliary results</b> .....	139
B.1 Results from Gaussian process theory .....	139
B.2 A concentration inequality for empirical processes .....	143
B.3 Mixing time bounds for Langevin diffusions .....	145
B.4 A characterisation of vanishing efficient information .....	147
<b>References</b> .....	151
<b>Index</b> .....	159



## Chapter 1

# Non-linear statistical inverse problems

The study of inverse problems forms an active field at the interface of applied and pure mathematics as well as the statistical and physical sciences. Prototypical examples include parameter identification in partial differential equations (PDEs) as well as various tomography and data assimilation tasks. The Bayesian approach to such problems has seen substantial activity in the last decade after seminal work by Andrew Stuart [118] – the recent references [37], [80, 108], [40], [25], [5], among many others, survey the relevance of such inference techniques in various areas of applied mathematics. Some of these ideas can be traced back as far as Henri Poincaré – see the discussion in Persi Diaconis’ early contribution [42], which is itself a relevant predecessor to the use of Bayesian thinking in this field.

The recent surge of interest stems from the inherent ability of the Bayesian approach to automatically and simultaneously address a variety of challenges of contemporary data science. These include algorithmic feasibility even for non-convex problems in high- or infinite-dimensional parameter spaces (via Markov chain Monte Carlo methods) as well as the provision of uncertainty quantification methods and ‘error bars’ for algorithmic outputs. A scientist using such methodology needs to specify a numerically tractable likelihood-function modelling the statistical measurement process and a reasonable prior distribution (such as a Gaussian random process) on the parameter, and in return obtains a posterior distribution that in principle can be used to solve all inferential tasks.

As much as the Bayesian ‘package’ is attractive in applications, a framework providing rigorous statistical and algorithmic *guarantees* for such methods – such as a convergence analysis in the large sample size scenario – has been developed only recently. For linear inverse problems the Bayesian approach is fairly well understood (e.g., [76, 107] and also [61, 64, 89] and references therein) and can be related to the well established ‘regularisation’ literature [51], [73], [17], where a convex penalised least squares fit functional is minimised to reconstruct the parameter of interest. But for *non-linear* problems Bayesian methods are genuinely distinct from optimisation based approaches and hence the theory requires different ideas. The purpose of these notes is to lay out some of the main mathematical mechanisms underpinning a body of recent theoretical work [93], [90, 91], [101], [1, 23, 62, 75, 100] on this subject. It builds on previous work on the understanding of Bayesian procedures in general high- and infinite-dimensional models by Aad van der Vaart and co-authors (e.g., [128] and the monograph [58]).

A common setting involves a non-linear ‘forward map’  $\mathcal{G}$  defined on some infinite-dimensional ‘parameter’ space  $\Theta$  and taking values in some Hilbert space of functions

or operators. We can think of  $\mathcal{G}(\theta)$  describing the solution of a PDE in dependence of an unknown function (or ‘coefficient’)  $f_\theta$  modelled by some parameter  $\theta$ . The inverse problem is to recover  $\theta$  from ‘data’  $\mathcal{G}(\theta)$ . Even if  $\theta \mapsto \mathcal{G}(\theta)$  is proved to be injective, so that an inverse  $\mathcal{G}^{-1}$  exists in principle, solution maps  $\mathcal{G}$  of PDEs are typically ‘smoothing’ (compactifying), which in infinite-dimensions means that  $\mathcal{G}^{-1}$  may not be continuous in relevant topologies. Also, since  $\mathcal{G}$  is non-linear, we typically do not have an explicit representation of the inverse map and analytical formulæ reconstructing  $\theta$  from  $\mathcal{G}$  (such as, e.g., for Fourier, Laplace and Radon transforms available in linear settings) are out of reach. And the problem becomes even less tractable if we explicitly acknowledge discretisation and statistical measurement error. Nevertheless, we will show that a principled Bayesian ‘likelihood’ approach may overcome such obstacles and, under certain conditions, provide provably valid algorithmic solutions for non-linear inverse problems.

## 1.1 Model examples

Non-linear forward maps  $\mathcal{G}$  arise in a variety of settings with spectral inverse problems, X-ray transforms, boundary rigidity problems, data assimilation and more generally in parameter identification problems with PDEs – see [74], [118], [73], [80, 108], [68] and also the recent monograph [104], by Gabriel Paternain, Mikko Salo and Gunther Uhlmann. We first describe some classical examples.

### 1.1.1 Boundary measurements in tomography

**1.1.1.1 The Calderón problem.** Let  $\mathcal{X} \subset \mathbb{R}^d$ ,  $d \geq 2$ , be a bounded domain with smooth boundary  $\partial\mathcal{X}$  and  $\theta: \mathcal{X} \rightarrow [\theta_{\min}, \infty)$  a differentiable map. Consider unique solutions  $u = u_{\theta,h}$  to the PDE

$$\begin{aligned} \nabla \cdot (\theta \nabla u) &= 0 && \text{in } \mathcal{X}, \\ u &= h && \text{on } \partial\mathcal{X}, \end{aligned}$$

where  $h: \partial\mathcal{X} \rightarrow \mathbb{C}$  prescribes some boundary values. (Existence of such solutions is discussed in more detail after (1.2) below.) If we cannot measure  $u$  inside of  $\mathcal{X}$ , we may record boundary information of  $u_{\theta,h}$  for sufficiently many distinct  $h$ . Specifically, one can measure the *Neumann (boundary) data* as

$$\Lambda_\theta(h) := \left. \frac{\partial u_{\theta,h}}{\partial \nu} \right|_{\partial\mathcal{X}},$$

where  $\partial/\partial\nu$  denotes the outward normal derivative on  $\partial\mathcal{X}$  (to be understood in a trace sense). The inverse problem here is to recover  $\theta$  from knowledge of the boundary operator  $\mathcal{G}(\theta) = \Lambda_\theta$  only – this is known as the Calderón problem, which has

fundamental applications in *electric impedance tomography*, where  $\theta$  models a conductivity inside of  $\mathcal{X}$  to be recovered from non-invasive measurements. See [123] for an overview of the theory and applications of this problem (first studied by Alberto Calderón [26]), and [92, 119] for landmark injectivity theorems for  $\mathcal{G}$ .

**1.1.1.2 Non-Abelian X-ray transforms.** Let  $M \subset \mathbb{R}^2$  be the closed unit disk with boundary  $\partial M$ . Consider lines in the plane (i.e., geodesics) parameterised by  $\gamma(t) = x + tv$ , where  $x \in \mathbb{R}^2$  and  $v \in S^1$ . We further introduce the influx boundary as

$$\partial_+ SM = \{(x, v) \in \partial M \times S^1 : x \cdot v \leq 0\},$$

where ‘ $\cdot$ ’ is the standard dot product in the plane. If we take  $(x, v) \in \partial_+ SM$ , then the line  $\gamma(t) = x + tv$  will exit the disk in finite time  $\tau_\gamma = \tau(x, v) := -2x \cdot v$ .

Let  $\theta: M \rightarrow \mathbb{C}^{n \times n}$  be a continuous matrix field. Given a geodesic  $\gamma: [0, \tau_\gamma] \rightarrow M$  with endpoints  $\gamma(0), \gamma(\tau_\gamma) \in \partial M$ , we consider the matrix ODE

$$\dot{U} + \theta(\gamma(t))U = 0, \quad U(\tau_\gamma) = \text{Id}.$$

We define the boundary scattering data of  $\theta$  on  $\gamma$  to be  $C_\theta(\gamma) := U(0)$ . This problem, backward in time for convention here, is well-posed and leads to a unique definition of  $U(0)$ , containing information about  $\theta$  along the geodesic  $\gamma$ . Note that when  $\theta$  is scalar ( $n = 1$ ), we obtain

$$\log U(0) = \int_0^{\tau_\gamma} \theta(\gamma(t)) dt,$$

which is the classical X-ray/Radon transform [106] of  $\theta$  along the ray  $\gamma$ . Considering the collection of all such data makes up the *non-Abelian X-ray transform* of  $\theta$ , viewed here as the non-linear map

$$\mathcal{G}(\theta) \equiv C_\theta: \partial_+ SM \rightarrow \mathbb{C}^{n \times n}, \quad (1.1)$$

and the goal is to recover  $\theta$  from  $C_\theta$ . We refer to [102, 103] and references therein for injectivity theorems for  $\mathcal{G}$ . A concrete physical application arises when  $\theta$  takes values in the Lie algebra  $\mathfrak{so}(n)$  of skew-symmetric matrices associated to the special orthogonal group  $\text{SO}(n)$ . In this case the scattering data  $C_\theta$  maps into  $\text{SO}(n)$ , and for  $n = 3$ , this is relevant in *neutron spin tomography*, see [67, 114] and also Figure 1.1 below.

## 1.1.2 Parameter identification for elliptic PDEs

The two preceding examples are within the scope of the general theory to be developed here – see Remark 2.3.4 and the notes to subsequent chapters for more details. But in these lecture notes we will focus exclusively on two non-linear model examples

with *elliptic* PDEs that arise from steady state measurements of heat-type equations. Their relative analytical simplicity allows us to give a largely self-contained treatment of the material, yet at the same time these examples provide a ‘pedagogical’ template for more complicated settings.

To this end, let us introduce our elliptic base examples. Let  $\mathcal{X}$  be a bounded domain in  $\mathbb{R}^d$  with smooth boundary  $\partial\mathcal{X}$ . For a linear elliptic differential operator  $\mathcal{L}_f$  indexed by some function  $f: \mathcal{X} \rightarrow \mathbb{R}$ , we consider solutions  $u = u_f$ ,

$$\begin{aligned} \mathcal{L}_f u &= g \quad \text{on } \mathcal{X}, \\ u &= h \quad \text{on } \partial\mathcal{X}, \end{aligned} \tag{1.2}$$

where  $g$  is a given smooth source function and  $h$  prescribes smooth boundary values. For the examples of  $\mathcal{L}_f$  studied in these notes, unique solutions  $u_f$  to the boundary value problem (1.2) exist by standard results for linear elliptic PDEs (see after (A.10) in Appendix A). We will later parameterise  $f = f_\theta$  by a function  $\theta$  varying in a linear space but suppress this in the notation for now.

If  $\nabla \cdot$ ,  $\nabla$  and  $\Delta$  denote the divergence, gradient and Laplace operations respectively, our first example will be the divergence form operator  $\mathcal{L}_f$  described by the action on sufficiently regular functions  $u: \mathcal{X} \rightarrow \mathbb{R}$  via

$$\mathcal{L}_f u = \nabla \cdot (f \nabla u) = \nabla f \cdot \nabla u + f \Delta u \quad \text{for } f \geq f_{\min} > 0. \tag{1.3}$$

Here  $f$  models a real-valued *diffusivity* (or conductivity) of the medium  $\mathcal{X}$  and as explained in more detail in the next subsection, (1.2) describes ‘steady state’ solutions of a *diffusion equation*. Determining  $f$  from solutions  $u_f$  to (1.2) has various applications in physics and engineering, for instance, in the context of ‘groundwater flow’ it is sometimes called *Darcy’s problem* (see [118, Section 3.7]). Unlike in the Calderón problem from earlier which deals with the same PDE, we will assume that we can measure  $u_f$  *throughout*  $\mathcal{X}$  and not just at  $\partial\mathcal{X}$  – see Remark 2.3.4 below for a discussion of the resulting differences.

The second example we shall examine closely is a *Schrödinger operator*  $\mathcal{L}_f$  given by the action

$$\mathcal{L}_f u = \frac{\Delta u}{2} - f u \quad \text{for } f \geq 0, \tag{1.4}$$

where the absorption potential  $f$  models an attenuation in the diffusivity (e.g., a cooling effect in the heat flow) – in this case we speak of (1.2) as the (steady state) *Schrödinger equation* – even though we do not consider the more complicated *complex* valued version of (1.2) relevant in quantum physics. Determining  $f$  from solutions  $u_f$  to (1.2) for such  $\mathcal{L}_f$  is a problem appearing, for instance, in photo-acoustic tomography [8,9]. In (1.4),  $\Delta$  could be replaced by any given (known) elliptic second-order differential operator but we stick to the standard Laplacian for (mostly notational) simplicity. The Schrödinger model is in a certain sense atypical for most

non-linear inverse problems as it gives a direct inversion formula for  $f = \Delta u/2u$  (provided  $u$  does not vanish on  $\mathcal{X}$ ). But it will be very useful to succinctly explain several mathematical ideas in these notes related to the ‘local curvature’ of typical non-linear forward maps  $f \mapsto u_f$  encountered in the theory. It also provides a relevant template for general operators of the form  $\mathcal{L}_f = \mathcal{D} + f$ , with  $f$  an order-zero perturbation of a given (not necessarily elliptic) differential operator  $\mathcal{D}$ . One example would be the one underlying (1.1), where  $\mathcal{D}$  arises from the geodesic vector field. See the notes to this section for more discussion.

We note that while the operator  $\mathcal{L}_f$  in the preceding two examples is linear, the forward map  $\mathcal{G}$  arising from  $f \mapsto u_f$  is not, see Exercise 1.4.1. It will however be injective (under conditions to be detailed below), and the problem of inferring  $f$  from  $u_f$  is then a well-posed non-linear inverse problem.

**1.1.2.1 From heat to steady state equation.** As we will later consider noisy ‘real world’ measurements of ‘steady state’ solutions  $u_f$  of (1.2), we digress briefly and explain a way to acquire approximate measurements of  $u_f$  from observations of the process of *diffusion* over a time interval  $(0, T]$ . Consider solutions  $\{u = u_f(x, t) : x \in \mathcal{X}, t \in (0, T]\}$  of the time evolution equation

$$\begin{aligned} \frac{\partial u}{\partial t} &= \mathcal{L}_f u \quad \text{on } \mathcal{X} \times (0, T], \\ u(\cdot, t) &= 0 \quad \text{on } \partial\mathcal{X}, \text{ for all } t \in (0, T], \\ u(\cdot, 0) &= g \quad \text{on } \mathcal{X}, \end{aligned} \tag{1.5}$$

where  $g$  is some initial condition. It is well known (see Appendix A.2 for details) that the elliptic operators  $\mathcal{L}_f$  from (1.3), (1.4) equipped with Dirichlet boundary conditions (i.e., on the space  $H_0^1$  from Appendix A.1) have a spectral representation in terms of eigen-pairs

$$(\lambda_{j,f}, e_{j,f}) \in (0, \infty) \times H_0^1(\mathcal{X}), \quad j \in \mathbb{N},$$

of the form

$$\mathcal{L}_f = - \sum_{j \geq 1} \lambda_{j,f} e_{j,f} \langle e_{j,f}, \cdot \rangle_{L^2} \tag{1.6}$$

where  $\langle \cdot, \cdot \rangle_{L^2}$  is the standard inner product of the Hilbert space  $L^2(\mathcal{X}) = L^2(\mathcal{X}, dx)$ , with  $dx$  being the Lebesgue measure. The unique solutions to the heat equation (1.5) with initial condition  $g$  then have the well known representation

$$u_f(\cdot, t)(g) = \sum_j e^{-t\lambda_{j,f}} e_{j,f} \langle e_{j,f}, g \rangle_{L^2}, \quad t \geq 0.$$

Formally, integrating term by term,

$$\int_a^b e^{-\lambda_{j,f}s} ds = -\frac{1}{\lambda_{j,f}} [e^{-\lambda_{j,f}b} - e^{-\lambda_{j,f}a}] \quad \text{for all } a < b, j \in \mathbb{N},$$

and taking limits  $a \rightarrow 0$ ,  $b \rightarrow \infty$ , we see that time averages of  $u_f(\cdot, s)(g)$ ,

$$-\int_0^\infty u_f(\cdot, s)(g) ds = -\sum_{j \geq 1} \lambda_{j,f}^{-1} e_{j,f} \langle e_{j,f}, g \rangle_{L^2} = \mathcal{L}_f^{-1}(g),$$

provide the solution to (1.2) with  $h = 0$  (Dirichlet boundary conditions). So a Riemann sum approximation

$$-\sum_l u_f(\cdot, s_l)(g)[s_l - s_{l-1}], \quad (1.7)$$

where  $u_f(\cdot, s_l)(g)$  are measurements of the solution of (1.5) at discrete equally spaced times  $s_l \in (0, \infty)$ , will give an approximate measurement of the ‘steady state’ solution  $u_f$  of (1.2). (Boundary values different from  $h = 0$  can be considered likewise by modifying the base heat equation (1.5) accordingly.)

### 1.1.3 Data assimilation

While steady state approximations to time evolution equations such as those just described are commonly used in applications, sometimes modelling time horizon dynamics is the explicit task of statistical inference – for instance, in ‘data assimilation’ problems [36, 80, 108]. A prototypical example arising in atmospheric sciences and fluid dynamics is tracking a velocity field describing the solutions to the non-linear (incompressible) *Navier–Stokes equations* with unknown initial condition. If  $\nu > 0$  denotes a scalar viscosity parameter, we can consider two-dimensional vector fields  $u = u_\theta = (u_\theta(t, x) : t \in (0, T], x \in \Omega)$ ,  $\Omega \subset \mathbb{R}^2$ , solving the system of non-linear partial differential equations given by

$$\begin{aligned} \frac{\partial}{\partial t} u - \nu \Delta u + u \cdot \nabla u + \nabla p &= 0 & \text{on } \Omega \times (0, T], \\ \nabla \cdot u &= 0 & \text{on } \Omega \times (0, T], \\ u &= 0 & \text{on } \partial\Omega, \\ u(0, \cdot) &= u_0 & \text{on } \Omega, \end{aligned}$$

where  $u(0)$  is a divergence free initial condition and  $\nabla p$  a pressure term, see [35] for details. One regards the initial condition  $u(0) \equiv \theta$  as the unknown parameter and the solutions  $\mathcal{G}(\theta) = u_\theta$  measured at discrete points  $(t_i, X_i)$  in time and space as the ‘forward’ data. The aim is to infer  $\theta$  and then also all subsequent states  $(u_\theta(t, x) : t > 0, x \in \Omega)$  of the system.

In the above example the non-linearity of  $\mathcal{G}$  arises from the underlying PDE – had we used the standard heat equation (1.5) instead, the resulting solution map  $\theta \mapsto u_\theta$  would have been linear. With heat-equation type data arising from the generator  $\mathcal{L}_f$ , non-linearities still arise for the coefficient to solution maps  $f \mapsto u_f(t, \cdot)$ .

For some observational settings, steady state approximations can be infeasible, specifically when measurements are available only at low ‘observation frequency’ in time  $t$ . In this case, knowledge of the Markovian time evolution structure should be exploited to solve the inverse problem. We will not prove rigorous theorems about such data assimilation problems here, but the key ideas and techniques apply to these settings as well – we refer to [63, 94, 96, 99] for some first results in this direction.

## 1.2 Bayesian inference in non-linear regression models

### 1.2.1 The forward map $\mathcal{G}$

We now introduce a notational framework designed to accommodate a large variety of examples of non-linear inverse problems. Let  $(\mathcal{X}, \mathcal{A})$  and  $(\mathcal{Z}, \mathcal{B})$  be measurable spaces equipped with measures  $\lambda$  and  $\zeta$ , respectively. We assume that  $\lambda$  is a probability measure and that  $\zeta$  is a finite measure. Let further  $V, W$  be vector spaces of fixed finite dimensions  $p_V, p_W \in \mathbb{N}$ , with inner products  $\langle \cdot, \cdot \rangle_W, \langle \cdot, \cdot \rangle_V$  and norms  $|\cdot|_W, |\cdot|_V$ , respectively. Then

$$L^\infty(\mathcal{X}), \quad L^2(\mathcal{X}) = L^2_\lambda(\mathcal{X}, V) \quad \text{and} \quad L^\infty(\mathcal{Z}), \quad L^2(\mathcal{Z}) = L^2_\zeta(\mathcal{Z}, W)$$

denote the bounded measurable, and  $\lambda$ - or  $\zeta$ - square integrable,  $V$ - or  $W$ -valued vector fields defined on  $\mathcal{X}, \mathcal{Z}$ , respectively. Denote by

$$\|\cdot\|_{L^2_\zeta(\mathcal{Z})}, \quad \|\cdot\|_{L^2_\lambda(\mathcal{X})}, \quad \langle \cdot, \cdot \rangle_{L^2(\mathcal{Z})}, \quad \langle \cdot, \cdot \rangle_{L^2(\mathcal{X})}$$

the usual  $L^2$ -norms and Hilbert space inner products on these spaces. We write generically  $\|\cdot\|_\infty$  for the supremum norm of a function over its domain.

We shall consider parameter spaces  $\Theta$  that are (Borel-measurable) subsets of  $L^2(\mathcal{Z}, W)$  – later these will often have to be *linear* spaces so that Gaussian measures can be defined on them, but for now this is not necessary. On  $\Theta$ , measurable ‘forward maps’

$$\theta \mapsto \mathcal{G}(\theta), \quad \mathcal{G} : \Theta \rightarrow L^2_\lambda(\mathcal{X}, V), \quad (1.8)$$

are defined. For the PDEs introduced in Section 1.1.2 we will regard the coefficient  $f = f_\theta$  as being parameterised by  $\theta \in \Theta$ , and  $\mathcal{G}$  will be the solution map  $\theta \mapsto u_{f_\theta}$  of (1.2). In this case we have  $V = W = \mathbb{R}$  and  $\mathcal{X} = \mathcal{Z}$ , but in many other inverse problems such as those from Sections 1.1.1, 1.1.3 (or see [74, 104]), a more flexible choice of  $V, W, \mathcal{X}, \mathcal{Z}$  is of interest, which is why we present the general theory in this way.

### 1.2.2 A random design regression model with normal errors

Real-world measurements in inverse problems arising from forward data  $\mathcal{G}(\theta)$  in (1.8) are discrete and subject to observational noise. For instance, with solutions  $u_f$  of (1.2)

one typically discretises  $\mathcal{X}$  into a finite set  $X_1, \dots, X_N$  for which measurements  $u_f(X_i)$  are taken. We regard the  $X_i$ 's as chosen *at random* from  $\mathcal{X}$  according to the probability distribution  $\lambda$ . This viewpoint of ‘probabilistic numerics’ (see the notes to this section for discussion) combines naturally with our second source of randomness – the additive measurement error  $\varepsilon_i$  occurring when performing a measurement of  $\mathcal{G}(\theta)(X_i)$ . We will follow Gauss [56] and assume that  $\varepsilon_i$  follows a normal (Gaussian) distribution, which has a solid probabilistic foundation – for instance, if our observations arise from time averages (1.7) of solutions  $u_f$  of the underlying heat equation, then each  $\varepsilon_i$  itself is already a cumulative sum of independent measurement errors, which by the central limit theorem of probability will be approximately normally distributed. We note, however, that the assumption that the  $\varepsilon_i$  be Gaussian is *not* necessary, see the notes to this chapter.

To fix ideas suppose that observations arise as the jointly independent and identically distributed (i.i.d.) random vectors  $(Y_i, X_i)_{i=1}^N$  of the form

$$Y_i = \mathcal{G}(\theta)(X_i) + \varepsilon_i, \quad \varepsilon_i \stackrel{\text{i.i.d.}}{\sim} N(0, I_V), \quad i = 1, \dots, N, \quad (1.9)$$

where  $\mathcal{G}$  is as in (1.8), and the  $X_i$ 's are random i.i.d. covariates drawn from law  $\lambda$  on  $\mathcal{X}$ , independent of the noise  $\varepsilon_i$ . We assume that the inner product of  $V$  is chosen such that the covariance matrix  $I_V$  of each noise vector  $\varepsilon_i \in V$  is diagonal.

The joint law of the random variables  $(Y_i, X_i)_{i=1}^N$  in (1.9) defines a product probability measure on  $(V \times \mathcal{X})^N$ , and it will be denoted by  $P_\theta^N = \bigotimes_{i=1}^N P_\theta^i$ , where we note  $P_\theta^i = P_\theta^1$  for all  $i$ . We write  $P_\theta = P_\theta^i$  for the law of a ‘generic copy’  $(Y, X)$ , which has probability density (Radon–Nikodym derivative)

$$\begin{aligned} \frac{dP_\theta}{d\mu}(y, x) &\equiv p_\theta(y, x) \\ &= \frac{1}{(2\pi)^{p_V/2}} \exp\left\{-\frac{1}{2}|y - \mathcal{G}(\theta)(x)|_V^2\right\}, \quad y \in V, x \in \mathcal{X}, \end{aligned} \quad (1.10)$$

for dominating measure  $d\mu = dy \times d\lambda$ , with  $dy$  Lebesgue measure on  $V$ .

The infinite product probability measure  $\bigotimes_{i=1}^\infty P_\theta$  describing the law of all possible infinite sequences of observations (in  $(V \times \mathcal{X})^\mathbb{N}$ ) will be denoted by  $P_\theta^\mathbb{N}$ . The usual stochastic  $O_P, o_P$  notation will be used throughout for  $P = P_\theta^N$  or equivalently  $P_\theta^\mathbb{N}$  – see [126] for the standard theory of stochastic convergence. The expectation operator of  $P_\theta^\mathbb{N}$  is denoted by  $E_\theta$  and we sometimes write  $E_\lambda$  for the expectation under the  $X_i$ 's only, while  $E_\varepsilon$  denotes expectation under the noise variables  $\varepsilon_i$  only. We also write shorthand

$$D_N := \{Y_1, \dots, Y_N, X_1, \dots, X_N\}, \quad N \in \mathbb{N}, \quad (1.11)$$

for the full data vector.

### 1.2.3 The Bayesian prior and posterior distribution

We now let  $\Theta$  be a measurable subset of  $L^2_\xi(\mathcal{Z}, W)$  with trace Borel- $\sigma$ -algebra  $\mathcal{F}$  – in fact all that follows in this subsection works for parameter spaces  $\Theta$  which are Polish spaces equipped with their Borel  $\sigma$ -algebra  $\mathcal{F}$ . Let  $\Pi$  be a probability measure on  $(\Theta, \mathcal{F})$  that we will often just call the ‘prior’. On the product space  $\Theta \times (V \times \mathcal{X})$  with tensor Borel  $\sigma$ -algebra  $\mathcal{F} \otimes (\mathcal{B}_V \otimes \mathcal{A})$  we can consider a probability measure  $Q$  given by the density

$$dQ(\theta, (y, x)) = p_\theta(y, x) d\mu(y, x) d\Pi(\theta), \quad y \in V, x \in \mathcal{X}, \theta \in \Theta,$$

where  $p_\theta$  is as in (1.10), and where we assume that  $(\theta, x) \mapsto \mathcal{G}(\theta)(x)$  is jointly  $\mathcal{F} \otimes \mathcal{A} - \mathcal{B}_\mathbb{R}$  measurable. This space formally supports the Bayesian model: the random variables  $(Y, X)|\theta$  have conditional  $\mu$ -densities

$$\frac{p_\theta(y, x) d\Pi(\theta)}{\int_{V \times \mathcal{X}} p_\theta(y, x) d\mu(y, x) d\Pi(\theta)} = p_\theta(y, x),$$

as in (1.10), whereas  $\theta|(Y, X)$  has ‘posterior’ density

$$\frac{p_\theta(Y, X) d\Pi(\theta)}{\int_\Theta p_\theta(Y, X) d\Pi(\theta)}.$$

Then for an i.i.d. sample from this model we have likewise

$$(Y_i, X_i)_{i=1}^N | \theta \stackrel{\text{i.i.d.}}{\sim} P_\theta^N$$

with *posterior distribution*

$$\Pi(\cdot | D_N) = \Pi(\cdot | (Y_i, X_i)_{i=1}^N)$$

of  $\theta|(Y_i, X_i)_{i=1}^N$  on  $\Theta$  now given by the ratio

$$d\Pi(\theta | D_N) = \frac{\prod_{i=1}^N p_\theta(Y_i, X_i) d\Pi(\theta)}{\int_\Theta \prod_{i=1}^N p_\theta(Y_i, X_i) d\Pi(\theta)} = \frac{e^{\ell_N(\theta)} d\Pi(\theta)}{\int_\Theta e^{\ell_N(\theta)} d\Pi(\theta)}, \quad \theta \in \Theta, \quad (1.12)$$

where we introduce the notation

$$\ell_N(\theta) = \sum_{i \leq N} \ell_i(\theta) \quad \text{for } \ell_i(\theta) = -\frac{1}{2} |Y_i - \mathcal{G}(\theta)(X_i)|_V^2, \quad (1.13)$$

which is, up to additive constants, the log-likelihood function  $\log dP_\theta^N(D_N)$  of the data  $D_N$ . At this point the preceding ratios are well defined at least  $Q$ -almost surely because they are conditional distributions/densities on the product space  $\Theta \times (V \times \mathcal{X})$ . A second ‘frequentist’ way to make sense of these ratios is discussed before Lemma 1.3.3 below. See also [58, Chapter 1] for more discussion of the latter aspect.

### 1.2.4 Posterior computation by MCMC

A Bayesian statistician will base its inferences about  $\theta$  on the posterior distribution  $\Pi(\cdot|D_N)$  from (1.12). In our context where  $\mathcal{G}$  is *non-linear*, the expression in (1.12) remains abstract and in absence of analytical formulae it is unclear how we can extract practical statistical information from it, specifically if we wish to avoid computation of the normalising factor which would involve the evaluation of a possibly intractable integral over  $\Theta$ . It is here where Markov chain Monte Carlo (MCMC) methods enter the stage. The main idea behind MCMC methods is to generate a Markov chain that has the posterior distribution as invariant measure. This can often be done in a way that requires at each step a single evaluation of the likelihood function  $\ell_N(\theta)$  from (1.13), which in turn involves an evaluation of  $\mathcal{G}$  rather than of its typically much more complicated inverse  $\mathcal{G}^{-1}$ , and also bypasses calculation of the normalising factor in (1.12). If such a Markov chain  $(\vartheta_k)$  is ergodic and ‘mixes well’ we can collect averages  $(1/J) \sum_{k=1}^J \vartheta_k$  and use them to numerically approximate the posterior mean or quantiles – see Figure 1.1 below for an illustration.

We will turn to the study of performance guarantees for MCMC methods later, but for now let us give a concept proof for how to setup Markov chains that have prescribed posterior distributions  $\Pi(\cdot|D_N)$  as invariant measures. Recall that the distribution of a Markov chain is characterised by its initial condition  $\vartheta_0$  and by its transition probabilities  $\Pr(\vartheta_m \in B | \vartheta_{m-1} = t)$ ,  $t \in \Theta$ ,  $m \in \mathbb{N}$ , where  $B$  is any (measurable) subset of  $\Theta$ . By the Markov property, the preceding transition probabilities are the same for all  $m$ , with Markov kernel  $K(t, \cdot)$ ,

$$\Pr(\vartheta_1 \in B | \vartheta_0 = t) = K(t, B) \quad \text{for } t \in \Theta, B \subset \Theta \text{ (measurable)}.$$

A probability measure  $\mu$  on  $\Theta$  is called *invariant* (or *stationary*) for the Markov kernel  $K$  if

$$\int_{\Theta} K(t, B) d\mu(t) = \mu(B) \quad \text{for all } B \subset \Theta \text{ (measurable)}.$$

In numerical practice we will work with a high-dimensional Euclidean ‘approximation’ space  $\mathbb{R}^D$  of  $\Theta$ , and wish to generate samples  $\vartheta_1, \vartheta_2, \dots, \vartheta_M, \dots$  from a Markov chain  $\{\vartheta_m : m \in \mathbb{N}\}$  taking values in  $\Theta = \mathbb{R}^D$  that has as invariant measure  $\mu = \Pi(\cdot|D_N)$  the posterior distribution.

**1.2.4.1 The pCN algorithm as a Metropolis–Hastings algorithm.** A general class of so-called Metropolis–Hastings algorithms arises in settings where we can evaluate ratios  $\mu(s)/q(s|t)$  for some auxiliary conditional probability density function  $q(\cdot|t)$  on  $\Theta$ . In this case we can generate a Markov chain  $\{\vartheta_m : m \in \mathbb{N}\}$  as follows:

- (1) For  $m \in \mathbb{N}$  and given  $\vartheta_m$ , generate a new draw (the ‘proposal’)

$$s_m \sim q(\cdot|\vartheta_m).$$

(2) Define

$$\vartheta_{m+1} = \begin{cases} s_m, & \text{with probability } \rho(\vartheta_m, s_m), \\ \vartheta_m, & \text{with probability } 1 - \rho(\vartheta_m, s_m), \end{cases}$$

with ‘acceptance’ probability

$$\rho(t, s) = \min\left\{\frac{\mu(s) q(t|s)}{\mu(t) q(s|t)}, 1\right\}.$$

The following proposition is standard and left as Exercise 1.4.2.

**Proposition 1.2.1.** *Let the Markov chain  $\{\vartheta_m : m \in \mathbb{N}\}$  be generated as above and suppose  $\mu, q(\cdot|t), t \in \Theta$ , are strictly positive throughout  $\Theta$ . Then  $\mu$  is an invariant measure for the Markov chain.*

When the prior  $\Pi$  giving rise to (1.12) is a Gaussian  $N(0, \Sigma_D)$ -distribution on  $\Theta = \mathbb{R}^D$  with non-singular covariance matrix  $\Sigma_D$ , we can use it to construct the ‘proposal distribution’  $q$ . This leads to the following ‘pCN’ Markov chain:

(1) For  $m \in \mathbb{N}$  and given  $\vartheta_m$ , generate a random vector  $\xi \sim N(0, \Sigma_D)$  and set

$$s_m = \sqrt{1 - 2\delta} \vartheta_m + \sqrt{2\delta} \xi$$

for some fixed ‘step size’  $\delta > 0$ .

(2) Define

$$\vartheta_{m+1} = \begin{cases} s_m, & \text{with probability } \rho(\vartheta_m, s_m), \\ \vartheta_m, & \text{with probability } 1 - \rho(\vartheta_m, s_m), \end{cases}$$

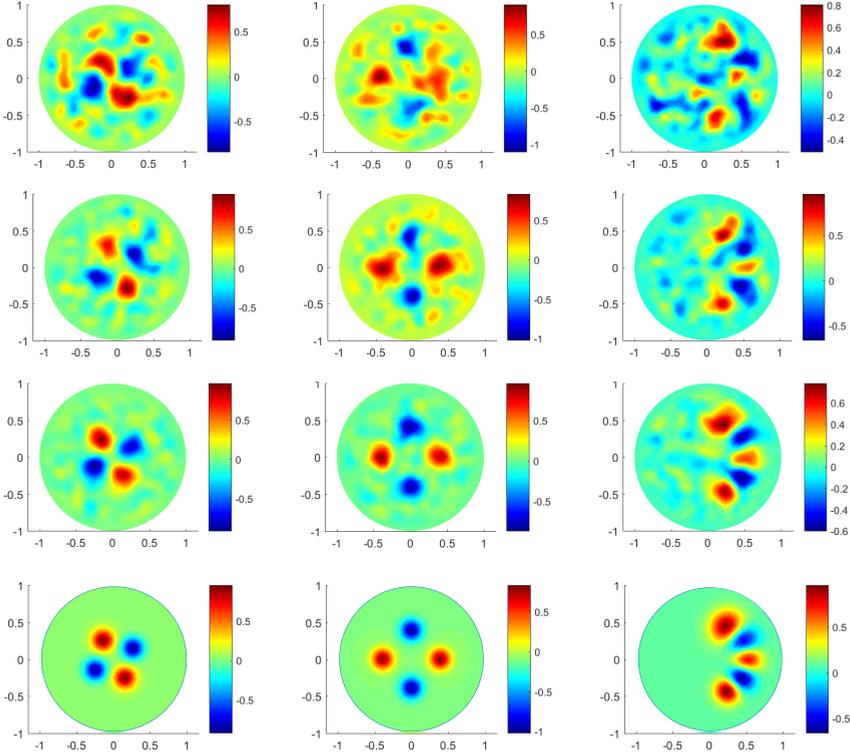
with ‘acceptance probabilities’

$$\rho(\vartheta_m, s_m) = \min\{e^{\ell_N(s_m) - \ell_N(\vartheta_m)}, 1\}.$$

An implementation of pCN for the non-linear inverse problem from (1.1) can be found in Figure 1.1 below – this is taken from [90], where further details about the precise statistical measurement model, step size and choice of prior parameters can be found.

The statistical intuition behind this algorithm is simple: We compute the new (conditionally Gaussian) proposal  $s_m$  and perform a likelihood ratio test against the previous position  $\vartheta_m$  of the Markov chain. If the likelihood is increased, we always accept the proposal. If the likelihood is decreased, we still accept with a certain probability to ensure the Markov chain does *not* stick to a local optimum. The data  $D_N$  only enters the algorithm when computing the likelihood ratios. This method targets the correct posterior distribution:

**Proposition 1.2.2.** *The preceding pCN Markov chain  $(\vartheta_m)$  has the posterior distribution (1.12) as invariant measure whenever the prior  $\Pi$  follows a  $N(0, \Sigma_D)$  distribution on  $\mathbb{R}^D = \Theta$ .*



**Figure 1.1.** Top 3 rows: The 3-dimensional posterior mean field  $(1/J) \sum_{k=1}^J \vartheta_k$  for sample sizes  $N = 200, 400, 800$ , arising from a (so(3)-valued) Matern Gaussian process prior, applied to a non-Abelian X-ray transform (1.1). The number of MCMC iterations via pCN is  $J = 100000$ . Bottom: the true matrix field  $\theta_0$ .

*Proof.* Note that for our choice of prior the target distribution (1.12) is of the form

$$\mu(\theta) \propto \exp\left\{\ell_N(\theta) - \frac{\theta^\top \Sigma_D^{-1} \theta}{2}\right\}, \quad \theta \in \mathbb{R}^D,$$

where  $\ell_N$  is as in (1.13). We will show that the pCN chain is a special case of a Metropolis–Hastings algorithm with proposal density

$$q(\cdot|t) \sim N(\sqrt{1-2\delta}t, 2\delta\Sigma_D) \quad \text{on } \Theta = \mathbb{R}^D, \quad t \in \Theta.$$

To simplify the algebra we set  $\Sigma_D = I_D$ , so that the proposal densities are

$$q(s|t) \propto \exp\left\{-\frac{1}{4\delta}\|s - \sqrt{1-2\delta}t\|^2\right\}, \quad s, t \in \mathbb{R}^D,$$

with  $\|\cdot\|$  the Euclidean norm on  $\mathbb{R}^D$  – but the proof works for any non-singular  $\Sigma_D$  just as well. Conditional sampling from  $q(\cdot|t)$  amounts simply to drawing a normal

random vector. Basic calculations give

$$\begin{aligned} \frac{\mu(s) q(t|s)}{\mu(t) q(s|t)} &= \exp\left\{\ell_N(s) - \frac{\|s\|^2}{2} - \ell_N(t) + \frac{\|t\|^2}{2}\right\} \\ &\quad \times \exp\left\{-\frac{1}{4\delta}\|t - \sqrt{1-2\delta}s\|^2 + \frac{1}{4\delta}\|s - \sqrt{1-2\delta}t\|^2\right\} \\ &= e^{\ell_N(s) - \ell_N(t)}, \end{aligned}$$

showing that the pCN scheme indeed has the appropriate ‘Metropolis–Hastings’ correction, and Proposition 1.2.1 applies. ■

We conclude that for Gaussian priors, approximate posterior computation by MCMC is in principle possible so long as we can evaluate  $\ell_N(\theta)$  at each iteration.

**1.2.4.2 Langevin algorithm.** Another approach to sample from a prescribed target measure  $\mu$  is based on ideas from stochastic differential equations. Here  $\mu$  does not need to arise from a Gaussian prior as in Proposition 1.2.2 but rather  $\mu = \mu_U$  is a general Borel probability measure on  $\Theta = \mathbb{R}^D$  which has a Lebesgue density proportional to  $e^{-U}$  for some potential  $U: \mathbb{R}^D \rightarrow \mathbb{R}$ , specifically,

$$\mu_U(B) = \frac{\int_B e^{-U(\theta)} d\theta}{\int_{\mathbb{R}^D} e^{-U(\theta)} d\theta}, \quad B \subseteq \mathbb{R}^D \text{ measurable.} \quad (1.14)$$

If the gradient  $\nabla U$  of  $U$  is Lipschitz continuous, we define the continuous time *Langevin diffusion process* as the unique strong solution  $(L_t : t \geq 0)$  of the stochastic differential equation (SDE)

$$dL_t = -\nabla U(L_t) dt + \sqrt{2} dW_t, \quad t \geq 0, L_t \in \mathbb{R}^D, \quad (1.15)$$

where  $(W_t : t \geq 0)$  is a  $D$ -dimensional standard Brownian motion.

**Proposition 1.2.3.** *The above SDE has a path-wise solution  $(L_t : t \geq 0)$  which is a continuous time Markov process with invariant measure  $\mu_U$  from (1.14).*

*Proof.* This is a standard result in stochastic calculus. The generator of this Markov process equals the second-order elliptic operator  $\mathcal{L}_U = \Delta + \nabla U \cdot \nabla$ . Then  $\mu_U$  solves the PDE  $\mathcal{L}_B^* \mu_U = 0$ , where  $\mathcal{L}_B^*$  is the adjoint of  $\mathcal{L}_B$  for  $dx$ , and hence is an invariant measure. See [7, p. 46f] or also [13] for details. ■

The Euler–Maruyama discretisation of the dynamics (1.15) gives rise to the discrete-time Markov chain  $(\vartheta_k : k \geq 0)$ ,

$$\vartheta_{k+1} = \vartheta_k - \gamma \nabla U(\vartheta_k) + \sqrt{2\gamma} \xi_{k+1} \quad \text{for all } k \geq 0, \quad (1.16)$$

where  $(\xi_k : k \geq 1)$  form an i.i.d. sequence of  $D$ -dimensional standard Gaussian

$N(0, I_{D \times D})$  vectors,  $\gamma > 0$  is some fixed *step size*, and  $\vartheta_0$  an initial value. By choosing for  $U = \ell_N + \log d\Pi$  the appropriate potential from (1.12), we can thus target the posterior measure by the Markov chain  $(\vartheta_k : k \in \mathbb{N})$ .

Just as in the pCN scheme, running this Markov chain requires drawing a Gaussian variable at each step as well as an evaluation of the gradient of  $\ell_N$  (and hence of  $\mathcal{G}, \nabla\mathcal{G}$ ), but neither the computation of  $\mathcal{G}^{-1}$  nor of the normalising factor in (1.12). Note that in contrast to pCN, the above algorithm does not include an ‘accept/reject’ Metropolis–Hastings correction step – the invariant measure of the discrete scheme is thus slightly different from  $\mu_U$  even though the misspecification error decreases with the discretisation error/step size  $\gamma$ , so that the method is still practical in numerical computation. We will refer to  $(\vartheta_k)$  as the *unadjusted Langevin algorithm* (ULA) in what follows, and its performance will be studied in detail in Chapter 5.

### 1.3 The frequentist perspective on posterior inference

While we rely on the Bayesian formalism in order to construct the posterior distribution  $\Pi(\cdot|D_N)$  from which we generate statistical inference algorithms, to obtain performance guarantees, we will instead take an ‘objective’ point of view and study the behaviour of  $\Pi(\cdot|D_N)$  under the ‘frequentist hypothesis’ that the  $X_1, X_2, \dots$ , are drawn from the infinite product measure  $P_{\theta_0}^{\mathbb{N}}$ , where  $\theta_0 \in \Theta$  is the ‘ground truth parameter’. This view-point dates back to Laplace [78] and is well-studied in mathematical statistics [58, 60, 82, 126] – see also the notes to this section for more on the history of this problem. It also appears most reasonable in the scientific applications of PDEs and inverse problems we have in mind, where the prior should be regarded as a regularisation tool rather than a subjective belief. In this section we give a first fundamental result about the frequentist behaviour of posterior measures in the setting of the regression model (1.9). In the proofs we will require some information theoretic inequalities which we derive first.

#### 1.3.1 Information distances for random design regression

Various features of statistical models can be encoded in so-called *information distances* on the laws  $\{P_\theta : \theta \in \Theta\}$  indexing the statistical ‘experiment’. In our model (1.9) these laws are equivalently represented by the probability densities  $p_\theta$  from (1.10). The first and most classical information distance is the ‘*Kullback–Leibler*’ (KL) *divergence*, or *entropy*, given by

$$\text{KL}(P_\theta, P_\vartheta) := E_\theta \left[ \log \frac{p_\theta}{p_\vartheta}(Y, X) \right],$$

a quantity that is always non-negative by an application of Jensen’s inequality. Sometimes we need to control higher (say, second) moments of the log-ratio in the last

expectation, so we introduce

$$V(P_\theta, P_\vartheta) := E_\theta \left[ \log \frac{p_\theta}{p_\vartheta}(Y, X) \right]^2.$$

The Hellinger distance  $h$  is given by

$$h^2(p_\theta, p_\vartheta) := \int_{V \times \mathcal{X}} [\sqrt{p_\theta} - \sqrt{p_\vartheta}]^2 d\mu$$

The following proposition relates these distances to more analytical aspects of the ‘forward’ maps  $\mathfrak{G}(\theta)$  appearing in our regression model (1.9), assuming the latter are uniformly bounded by a fixed constant  $U$ .

**Proposition 1.3.1.** *Suppose that for a subset  $\Theta \subset L^2(\mathcal{Z}, W)$  and some finite constant  $U = U_{\mathfrak{G}, \Theta, V} > 0$  we have*

$$\sup_{\theta \in \Theta} \|\mathfrak{G}(\theta)(\cdot)\|_V \leq U. \quad (1.17)$$

For the model densities from (1.10) arising from (1.9), we have for every  $\theta, \vartheta \in \Theta$ ,

$$\text{KL}(P_\theta, P_\vartheta) = \frac{1}{2} \|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathbf{x}, V)}^2, \quad (1.18)$$

$$V(P_\theta, P_\vartheta) \leq 2(U^2 + 1) \|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathbf{x}, V)}^2, \quad (1.19)$$

as well as

$$C_U \|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathbf{x}, V)}^2 \leq h^2(p_\theta, p_\vartheta) \leq \frac{1}{4} \|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathbf{x}, V)}^2, \quad (1.20)$$

where

$$C_U = \frac{1 - e^{-U^2/2}}{2U^2}.$$

*Proof.* For  $(Y, X) \sim P_\theta$ , from (1.9) with noise variable  $\varepsilon \sim N(0, I_V)$  and using (1.10), we see that

$$\log \frac{p_\theta}{p_\vartheta}(Y, X) = \frac{1}{2} |\mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X)|_V^2 + \langle \varepsilon, \mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X) \rangle_V, \quad (1.21)$$

and the  $E_\theta = E_\lambda E_\varepsilon$  expectation of the right-hand side is precisely (1.18) as  $E_\varepsilon \varepsilon = 0$ .

Next, squaring the last displayed identity and using  $(a + b)^2 \leq 2(a^2 + b^2)$ ,

$$\begin{aligned} E_\theta \left[ \log \frac{p_\theta}{p_\vartheta}(Y, X) \right]^2 &\leq \frac{1}{2} E_\lambda [|\mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X)|_V^2]^2 \\ &\quad + 2E_\lambda E_\varepsilon \langle \varepsilon, \mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X) \rangle_V^2 \\ &\leq 2(U^2 + 1) \|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathbf{x}, V)}^2, \end{aligned}$$

where we have used (1.17) and the fact that conditionally on  $X$ , the random variable  $(\varepsilon, \mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X))_V$  is  $N(0, |\mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X)|_V^2)$  distributed. This proves inequality (1.19).

It remains to prove (1.20). We first bound the so-called *Hellinger affinity* defined as

$$\rho(p_\theta, p_\vartheta) := \int_{V \times \mathcal{X}} \sqrt{p_\theta p_\vartheta} d\mu = 1 - \frac{1}{2} h^2(p_\theta, p_\vartheta), \quad (1.22)$$

where the second identity is easily verified. In view of (1.10), by writing  $\phi_V = (2\pi)^{-p_V/2} e^{-1 \cdot | \cdot |_V^2 / 2}$  for the (Lebesgue-) probability density of a standard multivariate normal variable  $Z$  on the finite-dimensional vector space  $V$ , and using the standard identity  $E e^{(u, Z)_V} = e^{|u|_V^2 / 2}$ ,  $u \in V$ , for Laplace transforms of such variables,

$$\begin{aligned} & \rho(p_\theta, p_\vartheta) \\ &= \frac{1}{(2\pi)^{p_V/2}} \int_{V \times \mathcal{X}} \exp\left\{-\frac{1}{4}|y - \mathfrak{G}(\theta)(x)|_V^2 - \frac{1}{4}|y - \mathfrak{G}(\vartheta)(x)|_V^2\right\} d\mu(y, x) \\ &= \int_{\mathcal{X}} \exp\left\{-\frac{1}{4}(|\mathfrak{G}(\theta)(x)|_V^2 + |\mathfrak{G}(\vartheta)(x)|_V^2)\right\} \int_V e^{\frac{1}{2}(y, \mathfrak{G}(\theta)(x) + \mathfrak{G}(\vartheta)(x))_V} \phi_V(y) dy d\lambda(x) \\ &= \int_{\mathcal{X}} \exp\left\{-\frac{2}{8}(|\mathfrak{G}(\theta)(x)|_V^2 + |\mathfrak{G}(\vartheta)(x)|_V^2) + \frac{1}{8}|\mathfrak{G}(\theta)(x) + \mathfrak{G}(\vartheta)(x)|_V^2\right\} d\lambda(x) \\ &= E_\lambda \exp\left\{-\frac{1}{8}|\mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X)|_V^2\right\}. \end{aligned}$$

By Jensen's inequality the right-hand side is lower bounded by

$$\exp\left\{-\frac{1}{8}\|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathcal{X}, V)}^2\right\}.$$

Combined with (1.22) and the standard bound  $1 - e^{-z} \leq z$  for  $z \geq 0$ , this gives

$$\frac{h^2(p_\theta, p_\vartheta)}{2} \leq 1 - \exp\left\{-\frac{1}{8}\|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathcal{X}, V)}^2\right\} \leq \frac{1}{8}\|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathcal{X}, V)}^2,$$

so that the second inequality in (1.20) follows. For the first inequality we use the basic consequence of convexity of  $e^{-z}$ ,

$$e^{-z_1} \leq \frac{e^{-z_2} - 1}{z_2} z_1 + 1 \quad \text{for } 0 \leq z_1 < z_2,$$

with choices

$$z_1 = \frac{1}{8}|\mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X)|_V^2, \quad z_2 = \frac{U^2}{2},$$

permitted, in view of (1.17), to the effect

$$E_\lambda \exp\left\{-\frac{1}{8}|\mathfrak{G}(\theta)(X) - \mathfrak{G}(\vartheta)(X)|_V^2\right\} \leq \frac{e^{-U^2/2} - 1}{4U^2} \|\mathfrak{G}(\theta) - \mathfrak{G}(\vartheta)\|_{L_\lambda^2(\mathcal{X}, V)}^2 + 1.$$

Combined with the displayed identity for  $\rho(p_\theta, p_\vartheta)$  and (1.22), this implies

$$1 - \frac{1}{2} h^2(p_\theta, p_\vartheta) \leq \frac{e^{-U^2/2} - 1}{4U^2} \|\mathcal{G}(\theta) - \mathcal{G}(\vartheta)\|_{L_\lambda^2(\mathcal{X}, V)}^2 + 1$$

from which the left-hand side of (1.20) follows immediately.  $\blacksquare$

We see from this proposition that relevant information distances are closely related to the standard  $L_\lambda^2(\mathcal{X}, V)$ -norms, as long as the  $\mathcal{G}(\theta)$  are uniformly bounded (without the latter assumption the story can be quite different, see [20]). In the PDE models relevant here, the uniform boundedness hypothesis will be satisfied.

### 1.3.2 A first posterior contraction theorem

Let us introduce the (semi-) metric

$$d_{\mathcal{G}}(\theta, \theta') := \|\mathcal{G}(\theta) - \mathcal{G}(\theta')\|_{L_\lambda^2(\mathcal{X}, V)}$$

on  $\Theta$ . For any subset  $\Theta' \subset \Theta$ , we denote by  $N(\Theta', d_{\mathcal{G}}, \eta)$ ,  $\eta > 0$ , the minimal number of balls of  $d_{\mathcal{G}}$ -radius  $\eta$  needed to cover  $\Theta'$ .

The proof of the next theorem is based on a general principle that ensures that posterior measures in i.i.d. models contract about Hellinger neighbourhoods of the ground truth  $\theta_0 \in \Theta$  that generated the data, as long as the prior is not too rough (i.e., concentrates on sets of not too large metric capacity) and charges a ‘neighbourhood’ of  $\theta_0$  for the KL- and  $V$ -distance of the model with sufficiently high probability. By Proposition 1.3.1 this translates into a comparable result in the  $d_{\mathcal{G}}$ -distance when data arises from (1.9) with uniformly bounded regression functions  $\mathcal{G}(\theta)$ .

**Theorem 1.3.2.** *Let  $\Pi_N$  be a sequence of prior Borel probability measures on some Borel subset  $\Theta \subset L_\xi^2(\mathcal{Z}, W)$ , and let  $\Pi_N(\cdot | (Y_i, X_i)_{i=1}^N) = \Pi_N(\cdot | D_N)$  be the resulting posterior distribution (1.12) arising from observations in model (1.9) with forward map  $\mathcal{G}: \Theta \rightarrow L_\lambda^2(\mathcal{X}, V)$ . Assume that for some fixed  $\theta_0 \in \Theta$ , envelope constants  $U = U_N \geq 1$ ,  $\|\mathcal{G}(\theta_0)\|_\infty \leq U$ , and a sequence  $\delta_N \rightarrow 0$  such that  $N\delta_N^2 \geq 1$ ,  $\sqrt{N}\delta_N/U_N \rightarrow \infty$  as  $N \rightarrow \infty$ , the sets*

$$\mathcal{B}_N := \{\theta \in \Theta : d_{\mathcal{G}}(\theta, \theta_0) \leq \delta_N, \|\mathcal{G}(\theta)\|_\infty \leq U\} \quad (1.23)$$

satisfy, for all  $N$  large enough,

$$\Pi_N(\mathcal{B}_N) \geq e^{-AN\delta_N^2} \quad \text{for some } A > 0. \quad (1.24)$$

Further, assume that there exists a sequence of Borel sets  $\Theta_N \subset \Theta$  for which

$$\Pi_N(\Theta_N^c) \leq e^{-BN\delta_N^2} \quad \text{for some } B > A + 2, \quad (1.25)$$

and such that, for all  $\bar{m} > 0$  large enough,

$$\log N(\Theta_N, d_{\mathcal{G}}, \bar{m}\delta_N) \leq N\delta_N^2. \quad (1.26)$$

Then, for all  $0 < b < B - A - 2$ , we can choose  $L = L(B, \bar{m}, b)$  large enough such that

$$P_{\theta_0}^N(\Pi_N(\theta \in \Theta_N, d_{\mathcal{G}}(\theta, \theta_0) \leq L\delta_N C_V^{-1} | D_N) \leq 1 - e^{-bN\delta_N^2}) \rightarrow 0 \quad (1.27)$$

as  $N \rightarrow \infty$ , where the constant  $C_V$  is as after (1.20), with

$$V = V_N = \max(U, \sup_{\theta \in \Theta_N} \|\mathcal{G}(\theta)\|_{\infty}).$$

*Proof.* We first need a lemma that we shall also use in different contexts later, and that also clarifies that the denominator in the formula for the posterior measure (1.12) is bounded away from zero on events of high  $P_{\theta_0}^N$ -probability, so that the posterior ratio is well-defined not just  $\mathcal{Q}$ -almost surely but also in the frequentist sense.

**Lemma 1.3.3.** *Let  $\mathcal{G}$  be as in the theorem and let  $\nu$  be a probability measure on some (measurable) subset  $B_N \subseteq \mathcal{B}_N$ . Then, for  $\ell_N$  from (1.13) and all  $K \geq 2$ , we have*

$$P_{\theta_0}^N\left(\int_{B_N} e^{\ell_N(\theta) - \ell_N(\theta_0)} d\nu(\theta) \leq e^{-KN\delta_N^2}\right) \leq \frac{8(U^2 + 1)}{K^2 N \delta_N^2}.$$

*Proof.* Notice first that, by (1.19),  $\log(p_{\theta}/p_{\theta_0})(Y, X)$  is square-integrable for the product measure  $P_{\theta_0} \otimes \nu$ . Then from Jensen's inequality (applied to  $\log$  and  $\int(\cdot) d\nu$ ) and recalling (1.10), the probability in question is bounded by

$$P_{\theta_0}^N\left(\int \int_{B_N} \log \frac{p_{\theta}}{p_{\theta_0}} d\nu(\theta) d(P_N - P_{\theta_0}) \leq -K\delta_N^2 - \int \int_{B_N} \log \frac{p_{\theta}}{p_{\theta_0}} d\nu(\theta) dP_{\theta_0}\right),$$

where  $P_N = (1/N) \sum_{i=1}^N \delta_{(Y_i, X_i)}$  is the empirical measure associated to the sample from  $P_{\theta_0}$ . Now as in (1.21) and by the definition of  $\mathcal{B}_N$ , for all  $\theta \in B_N$ ,

$$-\int \log \frac{p_{\theta}}{p_{\theta_0}} dP_{\theta_0} = \frac{1}{2} \|\mathcal{G}(\theta) - \mathcal{G}(\theta_0)\|_{L_{\lambda}^2(\mathbf{x}, \nu)}^2 \leq \delta_N^2,$$

so that for  $K \geq 2$  and using also Fubini's theorem, the last probability can be bounded by

$$P_{\theta_0}^N\left(\sqrt{N} \int \int_{B_N} \log \frac{p_{\theta_0}}{p_{\theta}} d\nu(\theta) d(P_N - P_{\theta_0}) \geq \frac{K\sqrt{N}\delta_N^2}{2}\right).$$

Now using the inequalities of Chebyshev and Jensen, Fubini's theorem, (1.19) from Proposition 1.3.1 and again the definition of  $\mathcal{B}_N \supset B_N$ , the last probability is bounded as

$$\frac{4E_{\theta_0}\left[\int_{B_N} \log(p_{\theta_0}/p_{\theta})(Y, X) d\nu(\theta)\right]^2}{K^2 N \delta_N^4} \leq \frac{4 \int_{B_N} V(P_{\theta_0}, P_{\theta}) d\nu(\theta)}{K^2 N \delta_N^4} \leq \frac{8(U^2 + 1)}{K^2 N \delta_N^2},$$

completing the proof of the lemma.  $\blacksquare$

We now prove the theorem. In view of the last lemma with  $\nu = \Pi(\cdot)/\Pi(B_N)$ ,  $B_N = \mathcal{B}_N$ ,  $K = A + 2$ , and (1.24), we can restrict to events

$$A_N = \left\{ \int_{\Theta} e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta) \geq e^{-(A+2)N\delta_N^2} \right\} \quad (1.28)$$

of probability  $P_{\theta_0}^N(A_N) \rightarrow 1$  as  $N \rightarrow \infty$ . Moreover, using (1.26) and the right-hand side in (1.20) we can verify the metric entropy condition in [60, Theorem 7.1.4] with choices  $\varepsilon_0 = m'\delta_N$  and  $\log N(\varepsilon) = cN\delta_N^2$ ,  $c > 0$ , constant in  $\varepsilon > \varepsilon_0$ , to deduce that for every  $k > 1$ , we can choose  $m' < m$  large enough and find ‘tests’  $\Psi_N = \Psi: (V \times \mathcal{X})^N \rightarrow \{0, 1\}$  for which

$$P_{\theta_0}^N(\Psi_N = 1) \xrightarrow{N \rightarrow \infty} 0 \quad \text{and} \quad \sup_{\substack{\theta \in \Theta_N \\ h(p_\theta, p_{\theta_0}) > m\delta_N}} E_\theta^N(1 - \Psi_N) \leq e^{-kN\delta_N^2}. \quad (1.29)$$

Now let us write

$$\bar{\Theta}_N = \Theta_N \cap \{h(p_\theta, p_{\theta_0}) \leq m\delta_N\}$$

and  $\bar{\Theta}_N^c$  for its complement in  $\Theta$ . By (1.12), as  $N \rightarrow \infty$ ,

$$\begin{aligned} & P_{\theta_0}^N(\Pi(\bar{\Theta}_N^c | D_N) \geq e^{-bN\delta_N^2}) \\ &= P_{\theta_0}^N \left( \frac{\int_{\bar{\Theta}_N^c} e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta)}{\int_{\Theta} e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta)} \geq e^{-bN\delta_N^2}, \Psi_N = 0, A_N \right) + o(1) \\ &\leq P_{\theta_0}^N \left( \int_{\bar{\Theta}_N^c} e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta) (1 - \Psi_N) \geq e^{-(b+A+2)N\delta_N^2} \right) + o(1). \end{aligned}$$

By Markov’s inequality, decomposing

$$\bar{\Theta}_N^c = \Theta_N^c \cup \{h(p_\theta, p_{\theta_0}) > m\delta_N\},$$

and using Fubini’s theorem as well as

$$E_{\theta_0}^N \prod_{i=1}^N \frac{p_\theta}{p_{\theta_0}} (1 - \Psi_N) = E_\theta^N (1 - \Psi_N) \leq 1, \quad (1.30)$$

we further bound the last probability as

$$\begin{aligned} & e^{(b+A+2)N\delta_N^2} \int_{\bar{\Theta}_N^c} E_\theta^N (1 - \Psi_N) d\Pi(\theta) \\ &\leq e^{(b+A+2)N\delta_N^2} \left( 2\Pi(\Theta_N^c) + \int_{\theta \in \Theta_N: h(p_\theta, p_{\theta_0}) > m\delta_N} E_\theta^N (1 - \Psi_N) d\Pi(\theta) \right) \\ &\lesssim e^{(b+A+2-B)N\delta_N^2} + e^{(b+A+2-k)N\delta_N^2} \xrightarrow{N \rightarrow \infty} 0, \end{aligned}$$

where we have used (1.25) and (1.29) with  $k$  and then  $m$  large enough. The proof of the theorem is completed after noting that the Hellinger-distance contraction result implies the one in  $d_{\mathcal{G}}$ -distance, by virtue of the first inequality in (1.20) with appropriate envelope constant. ■

## 1.4 Notes

### 1.4.1 Exercises

**Exercise 1.4.1.** Show that the solution maps  $f \mapsto u_f$  for (1.2) with  $\mathcal{L}_f$  given by (1.3), (1.4), respectively, are *not linear*. [Hint: Initially consider the case  $d = 1$  and then for  $d > 1$  take  $f(x_1, \dots, x_d) = f(x_1)$ . For the Schrödinger equation this can also be inferred from the Feynman–Kac formula (A.23).]

**Exercise 1.4.2.** Prove Proposition 1.2.1. (See, e.g., [112].)

### 1.4.2 Remarks and comments

Inverse problems and parameter identification problems for PDEs have been studied for a long time – see [51, 68, 73, 74, 104, 118] and references therein. The PDE (1.2) with diffusion operator  $\mathcal{L}_f$  from (1.3) serves as ‘fruitfly’ example throughout large parts of the literature (see [118, Section 3.7] and [62] for many more references). The example arising with Schrödinger operators (1.4) has its own physical motivation (see [8, 9, 75]) but equally importantly provides a template for other ‘perturbed’ differential operators of the form  $\mathcal{D} + f$ , where  $\mathcal{D}$  is not necessarily elliptic or of second order. In fact, the theory laid out in these notes has been demonstrated to be compatible with the analytically more challenging example of *non-Abelian X-ray transforms* from (1.1), see [90], [91], [23] for references. Here one can regard  $\mathcal{D}$  to be the geodesic vector field on a simple manifold. Many more such ‘geometric’ non-linear inverse problems are discussed in detail in [74, 104] and serve as further examples. But the focus of these notes is primarily on *statistical* aspects of the problem and the analytically ‘softer’ elliptic model examples already exhibit many of the main statistical features of the theory.

Random design regression models (1.9) are convenient in our setting because of the i.i.d. structure of the samples  $(Y_i, X_i)$ , which will later on permit the use of techniques from empirical process theory and concentration of measure [47, 60, 120, 129, 131]. They also provide a natural link between standard information distances and the  $L^2$ -structure on the ‘forward data’  $\mathcal{G}(\theta)$  – Proposition 1.3.1 is essentially due to [20]. Regarding the  $X_i$  as random is a popular approach studied independently in ‘probabilistic numerics’, see [42] and more recently [25]. Random design regression models can be shown to be ‘asymptotically equivalent’ to most other commonly used

non-parametric regression models, see [109], and our findings thus also inform such measurement settings. We further note that the theory in this manuscript works as well if the noise vector  $\varepsilon_i$  is not Gaussian but has mean zero  $E\varepsilon_i = 0$  and sub-Gaussian moments. As long as we still use the (now mis-specified) ‘quasi-likelihood’ model arising from Gaussian densities (1.10), one can prove a version of Theorem 1.3.2 following ideas in [58, Section 8.5.2]. A similar remark about the noise variables applies to developments in later chapters and will not be repeated.

Key references for the Bayesian approach to inverse problems were discussed at the beginning of this chapter. MCMC methods have been highly influential in statistical science in the last few decades, see [112] for an introductory text and [37], [65], [18] in the context of the infinite-dimensional models relevant here. See specifically [65] for a more detailed theoretical treatment of the pCN algorithm in this context. Langevin MCMC methods also have a long history [113] and recent interest has been triggered in high-dimensional settings by the articles [39, 48] which will be relevant in Chapter 5.

Theorem 1.3.2 is a variant of a by now classical contraction rate result for posterior distributions in infinite-dimensional statistical parameter spaces first given in the landmark paper [57]. An in depth account of this theory can be found in [58] and what is relevant for the proofs here can also be found in [60, Chapter 7.3]. The proofs naturally exploit properties of the Hellinger distance which plays a central role throughout mathematical statistics, see [19, 58, 60, 82, 124]. In traditional ‘direct’ models, such posterior ‘consistency’ theorems have a long history with important contributions by Doob, Le Cam, Schwartz, Freedman, Diaconis, Barron, Shervish, Wasserman, Ghosal, van der Vaart among others, see [45], [81], [54], [116], [43], [12], [57] and also [58, Section 6.9] for a precise historical discussion.



## Chapter 2

### Global stability and posterior consistency

In this section we investigate when the Bayesian posterior distribution solves a non-linear problem successfully in the sense that it will converge in the large sample limit ( $N \rightarrow \infty$ ) to a Dirac measure at the ‘ground truth’ parameter  $\theta_0$  generating the data – a notion called *posterior consistency*. We will quantify that rate of convergence in norms on the parameter space  $\Theta$  of interest. Just as in Theorem 1.3.2, this will involve probabilistic statements under the law  $P_{\theta_0}^{\mathbb{N}}$ .

The main challenge compared to classical ‘direct’ statistical problems here is the presence of the *non-linear* map  $\mathcal{G}$ : Firstly, Bayesian inference can be expected to recover  $\mathcal{G}(\theta)$  if an appropriate prior for the regression function  $\mathcal{G}(\theta)$  is used (verifying (1.24)). In our setting one will initially employ a – say, Gaussian – prior for the parameter  $\theta$  of interest. It is not clear whether the implied prior for the regression functions  $\mathcal{G}(\theta)$  is still adequate, for instance, the prior for  $\mathcal{G}(\theta)$  is not Gaussian any longer when  $\mathcal{G}$  is non-linear. This issue will be addressed by a ‘forward regularity’ condition which basically requires Lipschitz continuity of  $\mathcal{G}$  for correct norms on a sufficiently large portion  $\Theta'$  of the parameter space  $\Theta$ , and which can be verified for the examples considered here by regularity estimates for inhomogeneous elliptic PDEs, as we will show.

The second challenge goes more directly to the heart of the subject: If the ‘forward map’  $\mathcal{G}$  satisfies what is commonly called a *stability estimate*

$$\|\theta_1 - \theta_2\| \leq c(\Theta') \|\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)\|^\eta \quad \text{for any } \theta_1, \theta_2 \in \Theta', \quad (2.1)$$

for suitable norms, sufficiently large subsets  $\Theta' \subset \Theta$  and ‘stability’ Hölder exponent  $\eta > 0$ , then there may be hope that a Bayes method that successfully infers the regression maps  $\mathcal{G}(\theta)$  also recovers the parameter  $\theta$  of interest in the inverse problem. Such stability estimates can be regarded as ‘quantitative statements’ about the injectivity of the map  $\theta \mapsto \mathcal{G}(\theta)$  and hence require a deeper understanding of the analytical properties of the inverse problem at hand. They imply in particular statistical *identifiability* of the model  $\{P_\theta : \theta \in \Theta\}$ , but for non-linear  $\mathcal{G}$  the implied constant  $c(\Theta')$  typically grows with the diameter of  $\Theta'$  and the stability estimate does not hold on the entire parameter space  $\Theta$ . This causes difficulties in the application of techniques from Bayesian non-parametric statistics with the unbounded parameter spaces that naturally support Gaussian priors.

We will show how the preceding programme can be made to work nevertheless for a class of Gaussian process priors (infinite-dimensional normal distributions). We start with the analytical properties that the forward map  $\mathcal{G}$  will be required to verify and check them for our main PDE examples. We then draw from the theory

of Gaussian measures in infinite-dimensions to apply Theorem 1.3.2 with appropriate ‘regularisation sets’  $\Theta_N$ , and ultimately deduce a variety of posterior consistency theorems.

## 2.1 Analytical hypotheses on the forward map and PDE examples

Recall the forward map  $\mathcal{G}: \Theta \rightarrow L_\lambda^2(\mathcal{X}, V)$  from (1.8), where the domain  $\Theta$  is assumed to be a (measurable) subset of  $L_\zeta^2(\mathcal{Z}, W)$ . We introduce a norm  $\|\cdot\|_{\mathcal{R}}$  of some linear *regularisation subspace*  $\mathcal{R}$  of  $L_\zeta^2(\mathcal{Z}, W)$  whose choice will later be related to the prior  $\Pi$  used. One could think of  $\mathcal{R}$  as a Sobolev or Hölder space of functions over a domain  $\mathcal{Z}$ . We write

$$B_{\mathcal{R}}(M) = \{\theta \in \mathcal{R} : \|\theta\|_{\mathcal{R}} \leq M\} \quad (2.2)$$

for the ball of radius  $M$  in the space  $\mathcal{R}$ .

### 2.1.1 Forward regularity conditions for $\mathcal{G}$

For  $\mathcal{Z}$  a bounded smooth domain in  $\mathbb{R}^d$ , the standard Sobolev spaces  $H^\kappa(\mathcal{Z})$  of  $W$ -valued vector fields are naturally defined, and so are their topological dual spaces  $(H^\kappa(\mathcal{Z}))^*$ , see Appendix A.1. When  $\kappa = 0$ , we have  $H^0(\mathcal{Z}) = L_\zeta^2(\mathcal{Z}, W)$  and note that then  $\|\cdot\|_{(H^0(\mathcal{Z}))^*} = \|\cdot\|_{L_\zeta^2}$  in (2.4) by self-duality of Hilbert spaces.

**Condition 2.1.1.** Consider a parameter space  $\Theta \subseteq L_\zeta^2(\mathcal{Z}, W)$  and measurable map  $\mathcal{G}: \Theta \rightarrow L_\lambda^2(\mathcal{X}, V)$ . Let  $\kappa \geq 0$ . Suppose for all  $M > 0$  and some normed linear subspace  $(\mathcal{R}, \|\cdot\|_{\mathcal{R}})$  of  $L_\zeta^2(\mathcal{Z}, W)$ , there exist finite constants  $U \geq 1$  and  $L > 0$  (that may depend on  $M$ ) such that

$$\sup_{\theta \in \Theta \cap B_{\mathcal{R}}(M)} \sup_{x \in \mathcal{X}} |\mathcal{G}(\theta)(x)|_V \leq U \quad (2.3)$$

and

$$\|\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)\|_{L_\lambda^2(\mathcal{X}, V)} \leq L \|\theta_1 - \theta_2\|_{(H^\kappa(\mathcal{Z}))^*} \quad \text{for all } \theta_1, \theta_2 \in \Theta \cap B_{\mathcal{R}}(M). \quad (2.4)$$

The condition requires that on bounded subsets of  $\mathcal{R}$ , the forward map  $\mathcal{G}$  is uniformly bounded and Lipschitz from  $(H^\kappa)^*$  to  $L_\lambda^2$ . As we are taking a dual space, for  $\kappa > 0$  the dual norm is weaker than the standard  $L_\zeta^2(\mathcal{Z}, W)$ -norm on  $\Theta$  (when  $\zeta$  is Lebesgue measure on  $\mathcal{Z}$ ). One can regard  $\kappa$  as measuring the ‘forward smoothing’ nature of  $\mathcal{G}$ . Determining the correct value of  $\kappa$  can be challenging in non-linear inverse problems but the results that follow can be used for any choice of  $\kappa \geq 0$ . (When  $\kappa = 0$  we can slightly weaken the requirement (2.4), see Exercise 2.4.4.) The theory also easily generalises to  $\mathcal{G}$  that is only Hölder instead of Lipschitz continuous, but this will not be relevant in these notes.

**2.1.1.1 Condition 2.1.1 and PDE examples.** To use Gaussian process priors we will require  $\Theta$  in Condition 2.1.1 to be a *linear* space. For the model PDE from (1.2), the parameter  $f$  indexing the differential operators  $\mathcal{L}_f$  in (1.3) and (1.4) does not lie in a linear space (as  $f \geq f_{\min}$  or  $f \geq 0$  are required), but we can map a linear space of  $\theta$ 's into positive  $f$ 's by use of a ‘link’ function (such as  $f = f_\theta = e^\theta > 0$ ). The forward map  $\mathcal{G}$  is then the solution map of these PDEs composed with the link function.

As for these PDEs both parameters  $\theta$ ,  $f_\theta$  and solutions  $\mathcal{G}(\theta)$  are all real-valued functions defined on the same bounded smooth domain  $\mathcal{X} \subset \mathbb{R}^d$ , we will take  $\mathcal{X} = \mathcal{Z}$ ,  $V = W = \mathbb{R}$ , and  $\lambda = \zeta$  equal to Lebesgue measure normalised to  $\lambda(\mathcal{X}) = 1$  when verifying Condition 2.1.1 in what follows. We will use standard techniques from the theory of elliptic PDEs reviewed in Appendix A.

**2.1.1.2 Schrödinger equation.** Let us first show how this works in the elliptic PDE (1.2), where

$$\mathcal{L}_f = \frac{\Delta}{2} - f$$

is a Schrödinger operator (1.4). We choose a parameter space  $\Theta \subset H^\xi(\mathcal{X}) \subset C(\mathcal{X})$  for some  $\xi > d/2$ , and enforce positivity of the potential  $f \in H^\xi$  by re-parameterising  $f_\theta = e^\theta > 0$ . The forward map is given by the unique solution

$$\mathcal{G}(\theta) \equiv u_\theta, \quad \theta \in \Theta, \quad (2.5)$$

of (1.2) with  $\mathcal{L}_f = \mathcal{L}_{f_\theta}$ , smooth boundary ‘temperatures’  $h \geq h_{\min} > 0$  and with vanishing source  $g = 0$ . The Feynman–Kac formula (A.23) (for standard Brownian motion  $\gamma = 1/2$  and  $V = f_\theta \geq 0$ ) implies (cf. (A.26))

$$\sup_{\theta \in \Theta} \|\mathcal{G}(\theta)\|_\infty = \sup_{\theta \in \Theta} \|u_\theta\|_\infty \leq \|h\|_\infty \equiv U < \infty, \quad (2.6)$$

which verifies (2.3) without even specifying a regularisation space  $\mathcal{R}$ . To check the Lipschitz property (2.4), note that for any  $\theta_i \in C(\mathcal{X})$  we necessarily have  $u_{\theta_1} - u_{\theta_2} = h - h = 0$  on  $\partial\mathcal{X}$  as well as

$$\mathcal{L}_{f_{\theta_1}}[u_{\theta_1} - u_{\theta_2}] = (\mathcal{L}_{f_{\theta_2}} - \mathcal{L}_{f_{\theta_1}})u_{\theta_2} = (f_{\theta_1} - f_{\theta_2})u_{\theta_2} \quad \text{on } \mathcal{X},$$

in other words,  $w = \mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)$  solves the inhomogeneous Schrödinger equation  $\mathcal{L}_{f_{\theta_1}} w = (f_{\theta_1} - f_{\theta_2})u_{\theta_2}$  on  $\mathcal{X}$  subject to Dirichlet boundary conditions  $w = 0$  on  $\partial\mathcal{X}$ . The inverse Schrödinger operator  $\mathcal{L}_{f_{\theta_1}}^{-1}$  for these boundary conditions is a Lipschitz operator on  $L_\lambda^2(\mathcal{X})$  with uniform in  $f \geq 0$  Lipschitz-constant  $c > 0$ , see (A.12) with  $\gamma = 1/2$ ,  $V = f$ . Hence, using also (2.6),

$$\begin{aligned} \|\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)\|_{L_\lambda^2} &= \|\mathcal{L}_{f_{\theta_1}}^{-1}[(f_{\theta_1} - f_{\theta_2})u_{\theta_2}]\|_{L_\lambda^2} \\ &\leq c \|u_{\theta_2}\|_\infty \|f_{\theta_1} - f_{\theta_2}\|_{L_\lambda^2} \\ &\leq C(c, M) \|h\|_\infty \|\theta_1 - \theta_2\|_{L_\lambda^2}, \end{aligned}$$

where  $M \geq \max_i \|\theta_i\|_\infty$  is required only in the last step to control the Lipschitz constant of  $e^\theta$ . (Choosing a positive ‘link’ function  $\phi: \mathbb{R} \rightarrow [0, \infty)$  with  $\|\phi'\|_\infty < \infty$  [instead of the exponential map], one can obtain  $\mathcal{G}$  that is in fact globally Lipschitz on  $\Theta = H^\xi$ .) Let us summarise this in the following result.

**Proposition 2.1.2.** *Let  $\Theta \subset H^\xi(\mathcal{X})$  for some  $\xi > d/2$ . The forward map  $\mathcal{G}$  from (2.5) satisfies Condition 2.1.1 for  $\kappa = 0$  and for any regularisation norm  $\|\cdot\|_{\mathcal{R}}$  that dominates the  $\|\cdot\|_\infty$ -norm.*

What precedes is already useful and general and will be used in later chapters, even though in terms of forward Lipschitz estimates it is not fully sharp: The forward map  $\mathcal{G}$  is actually smoothing of order two and Condition 2.1.1 can be checked for  $\kappa = 2$  instead of just  $\kappa = 0$  (unless the boundary temperatures  $h$  are ‘irregular’). In this case, however, the non-linearity of the map becomes more apparent and stronger regularisation norms  $\|\cdot\|_{\mathcal{R}}$  are required to control the Lipschitz constants in the preceding estimates. We examine how this works for the diffusion equation and leave the details for the Schrödinger case to the reader, see Exercise 2.4.1 of this section.

**2.1.1.3 Diffusion equation.** We now study solutions to the equation (1.2) with divergence form operator  $\mathcal{L}_f = \nabla \cdot (f \nabla)$  from (1.3). As parameter space we consider a subset of the Sobolev space  $\Theta \subset H^\beta(\mathcal{X})$ ,  $\beta > 1 + d/2$ . While weaker hypotheses on  $\beta$  are conceivable, we assume sufficient regularity to streamline the exposition. For a fixed positive scalar  $f_{\min} > 0$ , we then parameterise positive conductivities as

$$f_\theta := f_{\min} + e^\theta \quad \text{for } \theta \in H^\beta(\mathcal{X}), \quad (2.7)$$

ensuring strict ellipticity of the operator  $\mathcal{L}_{f_\theta}$ . The forward map is then given by the solution

$$\mathcal{G}(\theta) \equiv u_{f_\theta}, \quad \theta \in \Theta, \quad (2.8)$$

of (1.2) with  $\mathcal{L}_f = \mathcal{L}_{f_\theta}$  and smooth  $g, h$ .

To check Condition 2.1.1, we can infer the uniform boundedness (2.3) from the Feynman–Kac representation (A.23) with  $\gamma = f_\theta$ ,  $V = 0$  (see also (A.26)),

$$\|u_{f_\theta}\|_\infty \leq c(\mathcal{X}, d, f_{\min}) \|g\|_\infty + \|h\|_\infty \equiv U < \infty. \quad (2.9)$$

For (2.4) notice that for any distinct  $\theta_i \in \Theta$  and on  $\mathcal{X}$ ,

$$\mathcal{L}_{f_{\theta_1}} [\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)] = g - g + (\mathcal{L}_{f_{\theta_2}} - \mathcal{L}_{f_{\theta_1}})[u_{f_{\theta_2}}] = \nabla \cdot ((f_{\theta_2} - f_{\theta_1}) \nabla u_{f_{\theta_2}}),$$

and also that  $\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2) = h - h = 0$  on  $\partial\mathcal{X}$ . Thus  $\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)$  itself solves an equation of the type (1.2) with operator  $\mathcal{L}_f = \mathcal{L}_{f_{\theta_1}}$ , boundary values  $h = 0$  and source  $g$  given in the right-hand side in the last display and can hence be represented via the inverse  $\mathcal{L}_{f_{\theta_1}}^{-1}$  of  $\mathcal{L}_{f_{\theta_1}}$  for Dirichlet boundary conditions. Using the elliptic

regularity estimate from Proposition A.5.1 (with  $\gamma = f_{\theta_1}$ ,  $V = 0$ ), we have

$$\begin{aligned} \|\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)\|_{L_\lambda^2} &= \|\mathcal{L}_{f_{\theta_1}}^{-1}[\nabla \cdot ((f_{\theta_2} - f_{\theta_1})\nabla u_{f_{\theta_2}})]\|_{L_\lambda^2} \\ &\leq C(M)\|\nabla \cdot ((f_{\theta_2} - f_{\theta_1})\nabla u_{f_{\theta_2}})\|_{(H_0^2)^*}, \end{aligned} \quad (2.10)$$

so long as the regularisation norm  $\|\cdot\|_{\mathcal{R}}$  bounds the  $C^1$ -norm of  $\theta$  to the effect that (using also (2.7))  $\sup_{\theta \in B_{\mathcal{R}}(M)} \|f_\theta\|_{C^1} \leq C(M) < \infty$ .

Note that Proposition A.5.2 with  $\gamma = f_\theta \in H^\beta$ , for  $\beta > 1 + d/2$  and  $V = 0$ , combined with the Sobolev imbedding  $H^{\beta+1} \subset C^2$  imply that  $\sup_{\theta \in B_{\mathcal{R}}(M)} \|u_{f_\theta}\|_{C^2} < \infty$  whenever the  $\|\cdot\|_{\mathcal{R}}$ -norm dominates that  $H^\beta$ -norm. Then, applying the divergence theorem (A.7) to the vector field  $(f_{\theta_2} - f_{\theta_1})\nabla u_{f_{\theta_2}}$ , we have

$$\begin{aligned} \|\nabla \cdot ((f_{\theta_2} - f_{\theta_1})\nabla u_{f_{\theta_2}})\|_{(H_0^2)^*} &= \sup_{\varphi \in H_0^2, \|\varphi\|_{H^2} \leq 1} \left| \int_{\mathcal{X}} \varphi \nabla \cdot ((f_{\theta_2} - f_{\theta_1})\nabla u_{f_{\theta_2}}) \right| \\ &= \sup_{\varphi \in H_0^2, \|\varphi\|_{H^2} \leq 1} \left| \int_{\mathcal{X}} (f_{\theta_2} - f_{\theta_1})\nabla \varphi \cdot \nabla u_{f_{\theta_2}} \right| \\ &\leq \|f_{\theta_2} - f_{\theta_1}\|_{(H^1)^*} \sup_{\|\varphi\|_{H^2} \leq 1} \|\nabla \varphi \cdot \nabla u_{f_{\theta_2}}\|_{H^1} \\ &\lesssim \|u_{f_{\theta_2}}\|_{C^2} \|f_{\theta_2} - f_{\theta_1}\|_{(H^1)^*}, \end{aligned}$$

using also (A.3). Finally, for  $\max_i \|\theta_i\|_{C^1} \leq M$ , the series  $\sum_{k=1}^{\infty} (\theta_1 - \theta_2)^{k-1}/k!$  converges (absolutely) to a function that is uniformly bounded in  $C^1$ , and by (A.3) we see

$$\|f_{\theta_2} - f_{\theta_1}\|_{(H^1)^*} = \sup_{\|\varphi\|_{H^1} \leq 1} \left| \int_{\mathcal{X}} (\theta_2 - \theta_1)\varphi e^{\theta_1} \sum_{k=1}^{\infty} \frac{(\theta_2 - \theta_1)^{k-1}}{k!} \right| \lesssim \|\theta_1 - \theta_2\|_{(H^1)^*}.$$

In summary this gives

$$\|\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)\|_{L_\lambda^2} \lesssim c(M)\|\theta_1 - \theta_2\|_{(H^1)^*},$$

and since the  $H^\beta$ -norm ( $\beta > 1 + d/2$ ) dominates the  $C^1$ -norm by (A.4) we have:

**Proposition 2.1.3.** *Let  $\Theta \subset H^\beta(\mathcal{X})$  for some  $\beta > 1 + d/2$ . The forward map  $\mathcal{G}$  from (2.8) satisfies Condition 2.1.1 for  $\kappa = 1$  and for any regularisation norm  $\|\cdot\|_{\mathcal{R}}$  that dominates the  $H^\beta(\mathcal{X})$ -norm.*

## 2.1.2 Injectivity and stability estimates

A statistical model of probability distributions  $\{P_\theta : \theta \in \Theta\}$  is called *identifiable* if  $P_\theta = P_{\theta'}$  can only occur whenever  $\theta = \theta'$ . In the case of observations arising

from (1.9), with uniformly bounded  $\mathcal{G}(\theta)$  as in (2.3), this is equivalent to injectivity of the map  $\theta \mapsto \mathcal{G}(\theta)$  on the parameter space  $\Theta$  (as follows, e.g., from (1.20)). To obtain statistical guarantees for posterior distributions, mere injectivity is not enough and we need to quantify in a certain sense ‘how injective’ the map  $\mathcal{G}$  is.

Consider first the forward map (2.5) arising from solutions  $u_f$  to the Schrödinger equation. As a consequence (A.25) of the Feynman–Kac formula, there exists  $c = c(\mathcal{X}, d)$  such that  $u_f \geq h_{\min} e^{-c\|f\|_\infty} > 0$ , and so we can divide by  $u_f$  and represent

$$f = \frac{\Delta u_f}{2u_f} \quad \text{on } \mathcal{X}, \quad (2.11)$$

in particular, if  $u_f = u_g$ , then  $f = g$  and the map  $f \mapsto u_f$  is injective (from  $H^\xi \rightarrow H^{\xi+2}$ , say, cf. (A.11)). If  $\theta_0 = \log f_0$ ,  $f_0 > 0$ , is the true parameter, the last identity suggests a stability estimate of the form

$$\|\theta - \theta_0\|_{L^2} \lesssim \|f_\theta - f_0\|_{L^2} \lesssim \|u_{f_\theta} - u_{f_{\theta_0}}\|_{H^2} \quad \text{for any } \theta \in \Theta = H^\xi, \xi > d/2, \quad (2.12)$$

but the constants in these inequalities are not uniform in  $\theta \in H^\xi$ , among other reasons because the lower bound for  $u_{f_\theta}$  deteriorates as  $e^{e\|\theta\|_\infty}$  for large  $\|\theta\|_\infty$ . This illustrates why regularisation is important: if we know  $\theta$  is bounded in a suitable  $\mathcal{R}$ -norm, we can control the constants in these inequalities and then interpolate the  $H^2$ -norms (using (A.5)) to bound them by a constant factor times the  $L^2$ -norm raised to some Hölder exponent  $\eta < 1$ . This motivates the following hypothesis on the ‘inverse modulus of continuity’ of  $\mathcal{G}$  as in (2.1).

**Condition 2.1.4** (Stability estimate). Let  $B_{\mathcal{R}}(M)$  be as in (2.2) for the regularity space  $\mathcal{R}$ . For some  $\eta > 0$  and all  $M > 0$ , suppose there exists a constant  $L' = L'_{\mathcal{G}}$  such that for all  $\delta > 0$  small enough,

$$\sup\{\|\theta - \theta_0\|_{L^2_\xi} : \theta \in \Theta \cap B_{\mathcal{R}}(M), \|\mathcal{G}(\theta) - \mathcal{G}(\theta_0)\|_{L^2_\lambda} \leq \delta\} \leq L'\delta^\eta. \quad (2.13)$$

For the Schrödinger equation this condition can be checked (see Exercise 2.4.1) using the preceding ideas. Condition 2.1.4 is relevant in non-linear inverse problems more generally as only a ‘stability inequality’ (2.13) is required rather than an inversion formula for  $\mathcal{G}^{-1}$  such as (2.11).

**2.1.2.1 Stability estimate for Darcy’s problem.** We will prove a stability estimate for the forward map (2.8) arising with the diffusion operator from (1.3). The inverse problem can be cast as determining a solution  $f = f_\theta$  in the first-order PDE

$$\nabla f \cdot v + af = g \quad \text{on } \mathcal{X}, \quad (2.14)$$

given  $v = \nabla u$ ,  $a = \Delta u$ ,  $g$  and boundary values  $u = 0$  (say). When  $\nabla u$  does not vanish (e.g., as in Exercise 2.4.2), we could use the method of characteristics to solve this

equation in  $f$  along the flow lines (integral curves) of the vector field  $\nabla u$  (where it becomes an ODE). The hypothesis that  $\nabla u$  does not vanish is however rather atypical given that  $u$  arises as a solution to (1.2), and when  $\nabla u$  has zeros this approach is not appropriate. While this issue will resurface in later chapters, for now we note that when  $g > 0$ , then for any  $x \in \mathcal{X}$  where  $\nabla u(x) = 0$ , we must have  $\Delta u(x) > 0$ , and this can be used to show injectivity of  $\mathcal{G}$  by an argument using ‘integrating factors’. To accommodate such scenarios in a unified way we will require the hypothesis (2.15), discussed in some more detail after the proof.

**Proposition 2.1.5.** *Let  $\theta_1, \theta_2 \in H^\beta(\mathcal{X})$ ,  $\beta > 1 + d/2$ , be such that  $\|\theta_i\|_{C^1} \leq B$ ,  $\theta_1 = \theta_2$  on  $\partial\mathcal{X}$ , and denote by  $u_{f_\theta}$  the corresponding solutions to (1.2) with  $\mathcal{L}_{f_\theta}$  from (1.3),  $f_\theta$  as in (2.7), and for smooth  $g, h$ . Assume, moreover,*

$$\inf_{x \in \mathcal{X}} \left[ \frac{1}{2} \Delta u_{f_\theta}(x) + \mu |\nabla u_{f_\theta}(x)|_{\mathbb{R}^d}^2 \right] \geq c_0 > 0 \quad (2.15)$$

holds for  $\theta = \theta_1$  and some  $\mu, c_0 > 0$ . Then for some finite constant

$$C = C(B, \mu, c_0, \mathcal{X}, g, h, f_{\min}) > 0,$$

we have

$$\|\theta_1 - \theta_2\|_{L^2} \leq C \|u_{f_{\theta_1}} - u_{f_{\theta_2}}\|_{H^2}.$$

*Proof.* We will use throughout that  $u_{f_\theta} \in H^{\beta+1}(\mathcal{X})$  defines an element of  $C^2(\mathcal{X})$  in view of (A.4) and (A.11) with  $V = 0$ . Define the operator

$$h \mapsto T_\theta(h) := \nabla \cdot (h \nabla u_{f_\theta}), \quad T_\theta : H_c^1(\mathcal{X}) \rightarrow L^2(\mathcal{X}).$$

**Lemma 2.1.6.** *Under the hypotheses of Proposition 2.1.5 with  $\theta = \theta_1$ , we have*

$$\|T_\theta(h)\|_{L^2} = \|\nabla \cdot (h \nabla u_{f_\theta})\|_{L^2} \geq c \|h\|_{L^2}$$

for all  $h \in H_c^1(\mathcal{X})$  and some constant  $c = c(\mu, U, c_0) > 0$ , where  $U$  is from (2.9).

*Proof.* Let us write  $u_\theta = u_{f_\theta}$  in the proof. Applying the divergence theorem (A.7) to any  $v \in C^1(\mathcal{X})$  vanishing at  $\partial\mathcal{X}$  gives

$$\langle \Delta u_\theta, v^2 \rangle_{L^2} + \frac{1}{2} \langle \nabla u_\theta, \nabla(v^2) \rangle_{L^2} = \frac{1}{2} \langle \Delta u_\theta, v^2 \rangle_{L^2}.$$

Consider first  $h \in C^\infty(\mathcal{X})$  of compact support in  $\mathcal{X}$  (and  $h \neq 0$  without loss of generality). Scaling  $h$  by integrating factors, we set  $v = e^{-\mu u_\theta} h$  with  $\mu > 0$  to be chosen and obtain

$$\frac{1}{2} \int_{\mathcal{X}} \nabla(v^2) \cdot \nabla u_\theta = - \int_{\mathcal{X}} \mu |\nabla u_\theta|_{\mathbb{R}^d}^2 v^2 + \int_{\mathcal{X}} v e^{-\mu u_\theta} \nabla h \cdot \nabla u_\theta,$$

so that, by the Cauchy–Schwarz inequality,

$$\begin{aligned} \left| \int_{\mathcal{X}} \left( \frac{1}{2} \Delta u_{\theta} + \mu |\nabla u_{\theta}|_{\mathbb{R}^d}^2 \right) v^2 \right| &= \left| \langle (\Delta u_{\theta} + \mu |\nabla u_{\theta}|_{\mathbb{R}^d}^2), v^2 \rangle_{L^2} + \frac{1}{2} \langle \nabla u_{\theta}, \nabla(v^2) \rangle_{L^2} \right| \\ &= \left| \langle h \Delta u_{\theta} + \nabla h \cdot \nabla u_{\theta}, h e^{-2\mu u_{\theta}} \rangle_{L^2} \right| \\ &\leq \bar{\mu} \|\nabla \cdot (h \nabla u_{\theta})\|_{L^2} \|h\|_{L^2} \end{aligned} \quad (2.16)$$

for  $\bar{\mu} = \exp(2\mu \|u_{\theta}\|_{\infty})$ . We next lower bound the multipliers of  $v^2$  in the left-hand side of (2.16). By (2.15),

$$\left| \int_{\mathcal{X}} \left( \frac{1}{2} \Delta u_{\theta} + \mu |\nabla u_{\theta}|_{\mathbb{R}^d}^2 \right) v^2 \right| \geq c_0 \int_{\mathcal{X}} v^2,$$

and combining this with (2.16) we deduce, for appropriate  $c' > 0$ ,

$$\|\nabla \cdot (h \nabla u_{\theta})\|_{L^2} \|h\|_{L^2} \geq c' \|v\|_{L^2}^2 \gtrsim \|h\|_{L^2}^2,$$

which proves the lemma for all  $h \in C^{\infty}(\mathcal{X})$  of compact support. If a sequence  $h_n$  of such functions converges in  $H^1(\mathcal{X})$ -norm to some  $h$  as  $n \rightarrow \infty$ , then  $h_n$  and using (A.3) also  $T_{\theta}(h_n)$  converge in  $L^2$  to  $h$  and  $T_{\theta}(h)$ , respectively, and hence the inequality of the lemma is preserved after taking limits, so that the result extends to the completion  $H_c^1(\mathcal{X})$  of  $C_c^{\infty}(\mathcal{X})$  for the  $\|\cdot\|_{H^1}$ -norm.  $\blacksquare$

To deduce Proposition 2.1.5, we first notice that by the mean value theorem, for  $\tilde{\theta}(x) \in [\theta_1(x), \theta_2(x)]$ ,

$$f_{\theta_1} - f_{\theta_2} = e^{\theta_1} - e^{\theta_2} = e^{\tilde{\theta}}(\theta_1 - \theta_2),$$

and since the  $\theta_i$  are bounded in  $C^1$ , they and then also  $\tilde{\theta}$  are uniformly bounded by a constant depending only on  $B$ , whence

$$\|\theta_1 - \theta_2\|_{L^2} \leq c(B) \|f_{\theta_1} - f_{\theta_2}\|_{L^2}.$$

Then let us write  $h = f_{\theta_1} - f_{\theta_2}$  which defines an element of  $H_0^1(\mathcal{X}) = H_c^1(\mathcal{X})$  under the hypotheses maintained. By (1.2) we have

$$\nabla \cdot (h \nabla u_{\theta_1}) = \nabla \cdot (f_{\theta_2} \nabla(u_{\theta_2} - u_{\theta_1})) = \mathcal{L}_{f_{\theta_2}}(u_{\theta_2} - u_{\theta_1}),$$

and hence for the constants on the right-hand side depending only on an upper bound for  $\|f_{\theta_2}\|_{C^1} \lesssim c(B) < \infty$ , we obtain

$$\|\nabla \cdot (h \nabla u_{\theta_1})\|_{L^2} \lesssim \|u_{\theta_2} - u_{\theta_1}\|_{H^2}.$$

By Lemma 2.1.6, the left-hand side is lower bounded by a constant multiple of  $\|h\|_{L^2} = \|f_{\theta_1} - f_{\theta_2}\|_{L^2}$ , so that the result follows.  $\blacksquare$

Condition (2.15) can be verified for a large class of constellations of  $h, g$  in (1.2), with  $\mathcal{L}_f$  from (1.3). It clearly holds whenever the gradient of  $u_\theta$  does not vanish (for examples, see Exercise 2.4.2). More generically if  $h = 0$  and the ‘initial temperatures’ in (1.5) are strictly positive,  $g \geq g_{\min} > 0$ , then

$$0 < g_{\min} \leq g = f_\theta \Delta u_{f_\theta} + \nabla f_\theta \cdot \nabla u_{f_\theta} \quad \text{on } \mathcal{X}, \quad (2.17)$$

so that either  $\Delta u_{f_\theta} \geq g_{\min}/(2\|f_\theta\|_\infty)$  or  $|\nabla u_{f_\theta}(x)|_{\mathbb{R}^d} \geq g_{\min}/(2\|f_\theta\|_{C^1})$  has to hold on  $\mathcal{X}$ , which allows one to check Condition (2.15).

To conclude, the verification of Condition 2.1.4 for  $\mathcal{R} = H^\beta(\mathcal{X})$ ,  $\beta > 1 + d/2$ , now follows from Proposition 2.1.5, an elliptic regularity estimate for solutions  $u_{f_\theta}$  of (1.2), and interpolation. This is summarised in the final result of this subsection.

**Proposition 2.1.7.** *Let  $\theta, \theta_0 \in H^\beta(\mathcal{X})$  for some  $\beta > 1 + d/2$  satisfy  $\theta = \theta_0 = 0$  on  $\partial\mathcal{X}$  and let  $B \geq \max(\|\theta\|_{H^\beta}, \|\theta_0\|_{H^\beta})$ . Let  $\mathfrak{G}(\theta) = u_\theta$  be the forward map from (2.8), and suppose further (2.15) holds for  $\theta = \theta_0$  and some  $\mu, c_0$ . Then there exists a constant  $c = c(B, \mu, c_0, \mathcal{X}, g, h, f_{\min}, \beta)$  such that*

$$\|\theta - \theta_0\|_{L^2} \leq c \|\mathfrak{G}(\theta) - \mathfrak{G}(\theta_0)\|_{L^2}^\eta \quad \text{for } \eta = \frac{\beta - 1}{\beta + 1},$$

in particular, Condition 2.1.4 holds with this choice of  $\eta$  and  $L' = L'(B)$  if the regularisation norm  $\|\cdot\|_{\mathcal{R}}$  dominates the  $H^\beta$ -norm.

*Proof.* Since the  $H^\beta$  norms bound the  $C^1$ -norms by (A.4), we deduce from Proposition 2.1.5 and (A.5) that

$$\|\theta - \theta_0\|_{L^2} \lesssim \|u_\theta - u_{\theta_0}\|_{H^2} \lesssim \|u_\theta - u_{\theta_0}\|_{L^2}^{(\beta-1)/(\beta+1)} \|u_\theta - u_{\theta_0}\|_{H^{\beta+1}}^{2/(\beta+1)},$$

with constants only depending on permitted quantities. The last factor is uniformly bounded in view of the elliptic regularity estimate in Proposition A.5.2 with  $\gamma = f_\theta, f_{\theta_0}, V = 0$ , so that the result follows.  $\blacksquare$

## 2.2 Regularisation with Gaussian process priors

In this section  $\mathcal{Z}$  will be a bounded domain in  $\mathbb{R}^d$  with smooth boundary. (For  $\mathcal{Z}$  equal to a  $d$ -dimensional Torus, the proofs that follow work as well if one considers periodic Sobolev spaces.) We will apply Theorem 1.3.2 to prior measures  $\Pi$  that arise from Gaussian processes  $(X(z) : z \in \mathcal{Z})$  indexed by the domain  $\mathcal{Z}$ . That is,  $\Pi$  will be an infinite-dimensional Gaussian distribution that is the law of a random field taking values in a closed linear subspace  $\Theta$  of the Hilbert space  $L^2_\zeta(\mathcal{Z})$ , see Appendix B.1 for a review of these concepts.

Normed linear subspaces  $\mathcal{R} \subseteq \Theta$  where the measure  $\Pi$  is supported, i.e.,

$$\Pi(\mathcal{R}) = \Pi(\theta \in \Theta : \|\theta\|_{\mathcal{R}} < \infty) = 1,$$

will be shown to serve as choices for the ‘regularisation spaces’ in Conditions 2.1.1 and 2.1.4 – they often encode ‘path regularity’ of the Gaussian process on its index set  $\mathcal{Z}$ . Further features of these prior measures are related to so-called *reproducing kernel Hilbert spaces* (RKHS)  $\mathcal{H}$  of the prior. These are also subspaces of  $\Theta$  but do not (in infinite dimensions) support  $\Pi$ . Rather they describe geometric ‘covariance-related’ properties of  $\Pi$ . In the theory that follows we can think of  $\mathcal{R} = H^\beta(\mathcal{Z})$  and  $\mathcal{H} = H^\alpha(\mathcal{Z})$  with  $\alpha > \beta + d/2$ , perhaps endowed with further boundary conditions (via  $H_c^\beta(\mathcal{Z})$  or  $H_0^\beta(\mathcal{Z})$  spaces, see Appendix A.1).

**Condition 2.2.1.** Let  $\Pi'$  be a centred Gaussian Borel probability measure on the linear space  $\Theta \subseteq L^2_\zeta(\mathcal{Z}, W)$  with RKHS  $\mathcal{H}$ . Suppose further that  $\Pi'(\mathcal{R}) = 1$  for some separable normed linear subspace  $(\mathcal{R}, \|\cdot\|_{\mathcal{R}})$  of  $\Theta$ .

In Appendix B.1.3 we review some standard examples of Gaussian processes satisfying Condition 2.2.1 (see also Exercise 2.4.3). The prior  $\Pi = \Pi_N$  arises from the base prior  $\Pi'$  as the law  $\mathcal{L}(\theta)$  after ‘rescaling’

$$\theta = N^{-d/(4\alpha+4\kappa+2d)}\theta', \quad \theta' \sim \Pi', \quad (2.18)$$

for ‘regularity’ parameter  $\alpha > 0$  to be chosen and where  $\kappa$  is the ‘forward smoothing degree’ of  $\mathfrak{G}$  from (2.4). In other words, we shrink the random Gaussian series  $\theta'$  towards zero – quantitatively this entails (in our context, *essential*) extra regularisation of the posterior measure since one can heuristically (say if  $\mathcal{H}$  is finite-dimensional) think of the ‘reweighting’ factor in (1.12) as being of Gaussian density form

$$d\Pi(\cdot) \propto \exp\left\{-\frac{N^{d/(2\alpha+2\kappa+d)}\|\cdot\|_{\mathcal{H}}^2}{2}\right\},$$

thus penalising large values of  $\|\theta\|_{\mathcal{H}}$  more than one would do without re-scaling. While the catonic (subjective) Bayesian may wish to avoid such  $N$ -dependent priors, penalties of this type are widely used in the analysis of regularised least squares procedures (e.g., [124, p. 174]) and also in the Bayesian context [127].

**Theorem 2.2.2.** *Suppose Condition 2.1.1 holds for the forward map  $\mathfrak{G}$ , some  $\kappa \geq 0$ ,  $U < \infty$ , regularisation space  $\mathcal{R}$  and bounded domain  $\mathcal{Z} \subset \mathbb{R}^d$  with smooth boundary. Let the prior  $\Pi'$  satisfy Condition 2.2.1 for this choice of  $\mathcal{R}$  and suppose its RKHS  $\mathcal{H}$  satisfies the continuous imbedding*

$$\mathcal{H} \subset H_c^\alpha(\mathcal{Z}) \quad \text{if } \kappa \geq 1/2 \quad \text{or} \quad \mathcal{H} \subset H^\alpha(\mathcal{Z}) \quad \text{if } \kappa < 1/2,$$

and denote by  $\Pi_N$  the law of the rescaled prior from (2.18). If  $\theta_0 \in \mathcal{H} \cap \mathcal{R}$ , then the hypotheses (1.24), (1.25) and (1.26) in Theorem 1.3.2 hold for some  $A > 0$ , for

sequence

$$\delta_N = N^{-(\alpha+\kappa)/(2\alpha+2\kappa+d)}, \quad (2.19)$$

and any  $B$  if the ‘regularisation sets’ are chosen as

$$\Theta_N = \{\theta \in \mathcal{R} : \theta = \theta_1 + \theta_2, \|\theta_1\|_{(H^\kappa)^*} \leq M\delta_N, \|\theta_2\|_{\mathcal{H}} \leq M, \|\theta\|_{\mathcal{R}} \leq M\}, \quad (2.20)$$

for  $M = M(B)$  sufficiently large. In particular, (1.27) holds with  $V = U < \infty$ .

*Proof.* The proof is organised in four steps. The small ball computation in the second step determines the sequence  $\delta_N$ , and the first and third steps combined verify the deviation bound (1.25), while the fourth step checks the complexity condition (1.26). We note that the proof in fact works for any subset  $\bar{\Theta}_N$  of  $\Theta_N$  as long as  $\Pi(\bar{\Theta}_N^c) \leq e^{-BN\delta_N^2}$  holds for some  $B > A + 2$ . We use throughout that  $N^{d/(4\alpha+4\kappa+2d)} = \sqrt{N}\delta_N$  for our choice of  $\delta_N$ .

*Step 1.* As  $\mathcal{R}$  is separable the Hahn–Banach theorem implies that its norm can be represented as

$$\|\theta'\|_{\mathcal{R}} = \sup_{T \in \mathcal{T}} |T(\theta')|, \quad \theta' \in \mathcal{R},$$

where  $\mathcal{T}$  is a countable family of continuous linear forms on  $(\mathcal{R}, \|\cdot\|_{\mathcal{R}})$ . In particular, for  $\theta' \sim \Pi'$ , the collection of random variables  $\{T(\theta') : T \in \mathcal{T}\}$  defines a centred Gaussian process with countable index set and  $\Pr(\|\theta'\|_{\mathcal{R}} = \sup_{T \in \mathcal{T}} |T(\theta')| < \infty) = 1$  by hypothesis. Thus Fernique’s theorem [60, Theorem 2.1.20] implies initially that  $E\|\theta'\|_{\mathcal{R}} \leq D$  for some constant  $D$  depending only on the base prior  $\Pi'$ , and then gives the bound

$$\begin{aligned} \Pi(\|\theta\|_{\mathcal{R}} > M) &= \Pi'(\|\theta'\|_{\mathcal{R}} > M\sqrt{N}\delta_N) \\ &\leq \Pi'\left(\|\theta'\|_{\mathcal{R}} - E\|\theta'\|_{\mathcal{R}} > \frac{M\sqrt{N}\delta_N}{2}\right) \\ &\leq e^{-cM^2N\delta_N^2} \\ &\leq \frac{1}{2} e^{-BN\delta_N^2} \end{aligned} \quad (2.21)$$

for all  $M$  large enough (and since  $\sqrt{N}\delta_N \geq 1$ ).

*Step 2.* We next address the small ball computation for (1.24) with sets  $\mathcal{B}_N$  from (1.23). Since  $\theta_0$  is assumed to belong to  $\mathcal{R}$ , we have for all  $M$  and  $\bar{M} = \bar{M}(M)$  large enough that

$$\|\theta - \theta_0\|_{\mathcal{R}} \leq M \Rightarrow \|\theta\|_{\mathcal{R}} \leq M + \|\theta_0\|_{\mathcal{R}} \equiv \bar{M},$$

and by (2.3) in Condition 2.1.1 then also  $\|\mathcal{G}(\theta)\|_{\infty} \leq U$  for some constant  $U = U_{\mathcal{G}}(\bar{M})$ . Next, using also (2.4) in Condition 2.1.1, [60, Corollary 2.6.18] and the

Gaussian correlation inequality, Theorem B.1.2, we have

$$\begin{aligned}
& \Pi(\theta : \|\mathcal{G}(\theta) - \mathcal{G}(\theta_0)\|_{L^2} \leq \delta_N, \|\mathcal{G}(\theta)\|_\infty \leq U) \\
& \geq \Pi(\theta : \|\mathcal{G}(\theta) - \mathcal{G}(\theta_0)\|_{L^2} \leq \delta_N, \|\theta - \theta_0\|_{\mathcal{R}} \leq M) \\
& \geq \Pi\left(\theta : \|\theta - \theta_0\|_{(H^\kappa)^*} \leq \frac{\delta_N}{L_{\mathcal{G}}(\bar{M})}, \|\theta - \theta_0\|_{\mathcal{R}} \leq M\right) \\
& \geq e^{-\frac{1}{2}N\delta_N^2 \|\theta_0\|_{\mathcal{H}}^2} \Pi\left(\theta : \|\theta\|_{(H^\kappa)^*} \leq \frac{\delta_N}{L_{\mathcal{G}}(\bar{M})}\right) \cdot \Pi(\|\theta\|_{\mathcal{R}} \leq M). \quad (2.22)
\end{aligned}$$

The last probability equals

$$\Pi'(\|\theta'\|_{\mathcal{R}} \leq M\sqrt{N}\delta_N) = 1 - \Pi'(\|\theta'\|_{\mathcal{R}} > M\sqrt{N}\delta_N) \geq 1 - e^{-cMN\delta_N^2} \geq \frac{1}{2},$$

by virtue of (2.21) in Step 1 (and for  $M$  large enough). So it remains to lower bound the middle factor in (2.22), which equals

$$\Pi'\left(\theta' : \|\theta'\|_{(H^\kappa)^*} \leq \frac{\sqrt{N}\delta_N^2}{L_{\mathcal{G}}(\bar{M})}\right).$$

We use (A.21) to the effect that a ball  $h_c^\alpha(r)$  of radius  $r$  in  $\mathcal{H} \subset H_c^\alpha$  has log- $\epsilon$ -covering numbers bounded as

$$\log N(h_c^\alpha(r), \|\cdot\|_{(H^\kappa)^*}, \epsilon) \lesssim \epsilon^{-d/(\alpha+\kappa)} \quad \text{for any } 0 < \epsilon < r, \quad (2.23)$$

and this remains valid for  $\kappa < 1/2$  also with  $H_c^\alpha$  replaced by  $H^\alpha$  as discussed after (A.21). Hence an application of Theorem B.1.1 in the ambient Banach space  $B = (H^\kappa(\mathcal{Z}))^* \supset L^2(\mathcal{Z})$  supporting  $\Pi'$  gives, after some basic calculations,

$$\Pi'\left(\theta' : \|\theta'\|_{(H^\kappa)^*} \leq \frac{\sqrt{N}\delta_N^2}{L_{\mathcal{G}}(\bar{M})}\right) \geq e^{-\bar{a}N\delta_N^2} \quad (2.24)$$

for some  $\bar{a} > 0$ . Overall the right-hand side in (2.22) is lower bounded by

$$\frac{1}{2} \exp\left\{-\left(\bar{a} + \frac{\|\theta_0\|_{\mathcal{H}}^2}{2}\right)N\delta_N^2\right\} \geq e^{-AN\delta_N^2}$$

for  $A = \bar{a} + \|\theta_0\|_{\mathcal{H}}^2/2 > 0$ .

*Step 3.* We next show (1.25). By Step 1, it suffices to prove

$$\Pi(\theta : \theta = \theta_1 + \theta_2 : \|\theta_1\|_{(H^\kappa)^*} \leq M\delta_N, \|\theta_2\|_{\mathcal{H}} \leq M) \geq 1 - \frac{1}{2} \exp\{-BN\delta_N^2\}$$

for  $M$  large enough. We can ignore the  $1/2$  factor by increasing the constant  $B$ . For the rescaled prior measures this requires a lower bound for

$$\Pi'(\theta' : \theta' = \theta'_1 + \theta'_2, \|\theta'_1\|_{(H^\kappa)^*} \leq M\sqrt{N}\delta_N^2, \|\theta'_2\|_{\mathcal{H}} \leq M\sqrt{N}\delta_N).$$

Using Theorem B.1.1 as before (2.24), we deduce that for some  $c > 0$ ,

$$-\log \Pi'(\theta' : \|\theta'\|_{(H^\kappa)^*} \leq M\sqrt{N}\delta_N^2) \leq c(M\sqrt{N}\delta_N^2)^{-2d/(2\alpha+2\kappa-d)},$$

so that taking any  $M > (B/c)^{-(2\alpha+2\kappa-d)/(2d)}$  implies

$$-\log \Pi'(\theta' : \|\theta'\|_{(H^\kappa)^*} \leq M\sqrt{N}\delta_N^2) \leq B(\sqrt{N}\delta_N^2)^{-2d/(2\alpha+2\kappa-d)} = BN\delta_N^2. \quad (2.25)$$

Next, denote

$$B_N = -2\Phi^{-1}(e^{-BN\delta_N^2}),$$

where  $\Phi$  is the standard normal cumulative distribution function. Then, by standard inequalities for  $\Phi^{-1}$  (e.g., [58, Lemma K.6]), we have  $B_N \simeq \sqrt{BN}\delta_N$  as  $N \rightarrow \infty$ , so that for  $M \geq \sqrt{B}$ ,

$$\begin{aligned} \Pi'(\theta' : \theta' = \theta'_1 + \theta'_2, \|\theta'_1\|_{(H^\kappa)^*} \leq M\sqrt{N}\delta_N^2, \|\theta'_2\|_{\mathcal{H}} \leq M\sqrt{N}\delta_N) \\ \geq \Pi'(\theta' : \theta' = \theta'_1 + \theta'_2, \|\theta'_1\|_{(H^\kappa)^*} \leq M\sqrt{N}\delta_N^2, \|\theta'_2\|_{\mathcal{H}} \leq B_N). \end{aligned}$$

By (2.25) and the isoperimetric inequality for Gaussian measures, see [60, Theorem 2.6.12], the last probability is then lower bounded by

$$\begin{aligned} \Phi(\Phi^{-1}[\Pi'_N(\|\theta'\|_{(H^\kappa)^*} \leq M\sqrt{N}\delta_N^2)] + B_N) &\geq \Phi(\Phi^{-1}[e^{-BN\delta_N^2}] + B_N) \\ &= 1 - e^{-BN\delta_N^2}, \end{aligned}$$

concluding the proof.

*Step 4.* It remains to prove (1.26). By definition of the set  $\Theta_N$  and (2.4) from Condition 2.1.1, it suffices to construct a  $\bar{m}\delta_N/4$ -covering in  $(H^\kappa)^*$ -distance of a ball in  $\mathcal{H}$  intersected with a ball in  $\mathcal{R}$ , for all  $\bar{m} = \bar{m}(M, L)$  large enough. The bound required in (1.26) then follows directly from (2.23) with  $\epsilon = (\bar{m}/4)\delta_N$  and  $\bar{m}$  large enough. (The case  $\kappa < 1/2$  again follows similarly, cf. after (A.21)). ■

The last theorem applies readily to choices for  $\Pi'$  discussed in Appendix B.1.3. For high-dimensional ‘sieved’ priors expressed in appropriate basis functions of the Dirichlet Laplacian, the proofs require minor adjustments and they also permit some refinements regarding the ‘regularisation sets’  $\Theta_N$ , see Exercise 2.4.3. We also note that when  $\kappa = 0$ , it suffices to require Condition 2.1.1 with the stronger  $\|\cdot\|_\infty$ -norm replacing the  $L^2 = H^0$  norm on the right-hand side in (2.4), see Exercise 2.4.4.

## 2.3 Convergence of posterior measure and mean

From the preceding developments we can now state the following posterior contraction rate result in the inverse problem at hand.

**Theorem 2.3.1.** *Suppose the Gaussian process prior  $\Pi_N$  and forward map  $\mathcal{G}$  satisfy the conditions of Theorem 2.2.2 with  $\Theta_N$  as in (2.20), some regularisation norm  $\mathcal{R}$ ,  $\delta_N = N^{-(\alpha+\kappa)/(2\alpha+2\kappa+d)}$  and some  $\kappa \geq 0$ . If  $\Pi(\cdot | (Y_i, X_i)_{i=1}^N) = \Pi(\cdot | D_N)$  is the posterior distribution from (1.12) arising from observations in the model (1.9), then for all  $b > 0$ , we can choose  $m$  large enough such that as  $N \rightarrow \infty$ ,*

$$P_{\theta_0}^N(\Pi_N(\theta \in \Theta_N, \|\mathcal{G}(\theta) - \mathcal{G}(\theta_0)\|_{L_\lambda^2(\mathcal{X}, \mathcal{V})} \leq m\delta_N | D_N) \leq 1 - e^{-bN\delta_N^2}) \rightarrow 0.$$

Moreover, assuming also Condition 2.1.4 for the present choice of  $\mathcal{R}$  and some  $\eta, L' > 0$ , we deduce that

$$P_{\theta_0}^N(\Pi_N(\theta \in \Theta_N, \|\theta - \theta_0\|_{L_\xi^2} \leq L'(m\delta_N)^\eta | D_N) \leq 1 - e^{-bN\delta_N^2}) = o(1). \quad (2.26)$$

*Proof.* From Theorems 1.3.2 and 2.2.2 we deduce directly the contraction rate (2.27) on the ‘forward level’. Next Condition 2.1.4 implies the set inclusion

$$\{\theta \in \Theta_N, \|\mathcal{G}(\theta) - \mathcal{G}(\theta_0)\|_{L_\lambda^2(\mathcal{X}, \mathcal{V})} \leq m\delta_N\} \subset \{\|\theta - \theta_0\|_{L_\xi^2} \leq L'(m\delta_N)^\eta\},$$

so that (2.26) follows directly from (2.27). ■

To construct a concrete statistical estimator (or ‘algorithm’), we extract a specific feature of the posterior distribution such as the mean  $E^\Pi[\theta | D_N]$  or posterior mode (maximiser of the posterior surface). The former is well defined as a Bochner integral and a natural Bayesian estimator, not the least since we can form ergodic averages of the MCMC outputs to numerically approximate it. The following theorem shows how the contraction rate from (2.3.1) translates into a convergence rate for  $E^\Pi[\theta | D_N]$ .

**Theorem 2.3.2.** *Under the conditions of (2.26) in Theorem 2.3.1, we have that*

$$\|E^\Pi[\theta | D_N] - \theta_0\|_{L_\xi^2} = O_{P_{\theta_0}^N}(\delta_N^\eta).$$

*Proof.* Let us write  $L^2 = L_\xi^2$  in this proof and set  $\eta_N = L'(m\delta_N)^\eta$ , with constant  $m$  to be chosen large enough. Then, by the inequalities of Jensen and Cauchy–Schwarz,

$$\begin{aligned} & \|E^\Pi[\theta | D_N] - \theta_0\|_{L^2} \\ & \leq E^\Pi[\|\theta - \theta_0\|_{L^2} | D_N] \\ & \leq \eta_N + E^\Pi[\|\theta - \theta_0\|_{L^2} 1\{\|\theta - \theta_0\|_{L^2} > \eta_N\} | D_N] \\ & \leq \eta_N + E^\Pi[\|\theta - \theta_0\|_{L^2}^2 | D_N]^{1/2} \Pi(\|\theta - \theta_0\|_{L^2} > \eta_N | D_N)^{1/2}, \end{aligned}$$

and we now show that the last term is  $O_{P_{\theta_0}^N}(\eta_N)$  to prove the theorem. We recall the sets  $A_N$  from (1.28) which satisfy  $P_{\theta_0}^N(A_N) \rightarrow 1$  as  $N \rightarrow \infty$ . Then using (2.26), Markov’s inequality, Fubini’s theorem, (1.30) and that the Gaussian measure  $\Pi'$  is

supported in  $L^2$  and hence integrates  $\|\cdot\|_{L^2}^2$  to a finite constant,

$$\begin{aligned}
 & P_{\theta_0}^N(E^\Pi[\|\theta - \theta_0\|_{L^2}^2 | D_N] \times \Pi(\|\theta - \theta_0\|_{L^2} > \eta_N | D_N) > \eta_N^2) \\
 & \leq P_{\theta_0}^N(E^\Pi[\|\theta - \theta_0\|_{L^2}^2 | D_N] e^{-bN\delta_N^2} > \eta_N^2) + o(1) \\
 & \leq P_{\theta_0}^N\left(e^{-bN\delta_N^2} \frac{\int \|\theta - \theta_0\|_{L^2}^2 e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta)}{\int e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta)} > \eta_N^2, A_N\right) + o(1) \\
 & \leq e^{(A+2-b)N\delta_N^2} \eta_N^{-2} \int \|\theta - \theta_0\|_{L^2}^2 E_{\theta_0}^N[e^{\ell_N(\theta) - \ell_N(\theta_0)}] d\Pi(\theta) + o(1) \\
 & \leq e^{(A+2-b)N\delta_N^2} \eta_N^{-2} \int \|\theta - \theta_0\|_{L^2}^2 d\Pi(\theta) + o(1) \xrightarrow{N \rightarrow \infty} 0,
 \end{aligned}$$

for  $m$  and then  $b$  large enough.  $\blacksquare$

Putting all what precedes together we can obtain posterior contraction theorems for the non-linear inverse problem arising from the model PDE (1.2) with  $\mathcal{L}_f$  from (1.3) or (1.4). We formulate a detailed result for Darcy's problem and leave the Schrödinger case to Exercise 2.4.1.

Recall from (2.17) that simple sufficient conditions for the hypothesis (2.15) in the following theorem can be given in terms of the source function  $g$  in (1.2), and that natural priors satisfying the required hypotheses exist (see Appendix B.1.3).

**Theorem 2.3.3** (Bayes solution of Darcy's problem). *Consider a Gaussian process prior  $\Pi'$  verifying Condition 2.2.1 with  $\Theta = \mathcal{R} = H_c^\beta(\mathcal{X})$  for some  $\beta > 1 + d/2$  and RKHS  $\mathcal{H} \subset H_c^\alpha(\mathcal{X})$  for some  $\alpha > \beta + d/2$ . Let the rescaled prior  $\Pi_N$  arise from (2.18) for some  $0 \leq \kappa \leq 1$ . Consider the forward map  $\mathcal{G}: \Theta \rightarrow L_\lambda^2(\mathcal{X})$  from (2.8) and assume that (2.15) is satisfied for  $\theta = \theta_0 \in \mathcal{H}$ . Let  $\Pi(\cdot | (Y_i, X_i)_{i=1}^N)$  be the posterior distribution from (1.12) arising from observations in the model (1.9). Let  $\delta_N = N^{-(\alpha+\kappa)/(2\alpha+2\kappa+d)}$ . Then for all  $b > 0$  we can choose  $m$  large enough such that as  $N \rightarrow \infty$ ,*

$$\begin{aligned}
 & P_{\theta_0}^N(\Pi_N(\theta, \|\mathcal{G}(\theta) - \mathcal{G}(\theta_0)\|_{L_\lambda^2} \leq m\delta_N, \|\theta\|_{H^\beta} \leq m | D_N) \leq 1 - e^{-bN\delta_N^2}) \\
 & = o(1).
 \end{aligned} \tag{2.27}$$

Moreover, for  $\eta = (\beta - 1)/(\beta + 1)$  and constant  $M > 0$  large enough, we also have

$$P_{\theta_0}^N(\Pi_N(\theta : \|\theta - \theta_0\|_{L_\lambda^2} > M\delta_N^\eta | D_N) \geq e^{-bN\delta_N^2}) = o(1)$$

as  $N \rightarrow \infty$ , and likewise the posterior mean satisfies

$$\|E^\Pi[\theta | D_N] - \theta_0\|_{L_\lambda^2} = O_{P_{\theta_0}^N}(\delta_N^\eta). \tag{2.28}$$

*Proof.* Proposition 2.1.3 verifies Conditions 2.1.1 for  $\kappa = 1$  and then for any  $0 \leq \kappa \leq 1$ . Moreover, for all  $\theta \in \Theta = H_c^\beta \subset H_0^\beta$  (see Appendix A.1) the boundary values  $(f_\theta)|_{\partial\mathcal{X}} = 1 + f_{\min}$  co-incide at  $\partial\mathcal{X}$ , so that Proposition 2.1.7 verifies Condition 2.1.4. The result then follows from Theorems 2.3.1 and 2.3.2.  $\blacksquare$

One shows easily that the same convergence rate toward  $f_{\theta_0}$  is inherited by the induced estimates  $f_{\theta} = f_{\min} + e^{\theta}$ ,  $\theta \sim \Pi(\cdot | (Y_i, X_i)_{i=1}^N)$  of the conductivity, since  $e^{\theta}$  is Lipschitz on bounded sets of  $H^{\beta}$ . Inspection of the proof further shows that when  $0 \leq \kappa < 1/2$ , one can replace  $H_c^{\beta}$ ,  $H_c^{\alpha}$  by  $H_0^{\beta}$ ,  $H_0^{\alpha}$ , respectively, in the hypotheses of the preceding theorem. The preceding theorem shows that the posterior contracts towards  $\theta_0$  at a rate as close to  $1/\sqrt{N}$  as desired in the ‘smooth case’ where  $\alpha, \beta \rightarrow \infty$ , which is near-optimal in this situation. A discussion of more general optimality considerations can be found in the notes below.

**Remark 2.3.4** (Relationship to other inverse problems). A classical non-linear inverse problem (the Calderón problem) arising with the diffusion operator  $\mathcal{L}_f$  from (1.3) was discussed already briefly in Section 1.1.1, where only *boundary measurements* of  $u_f$  are available. We refer to [1] for an account of the precise forms of statistical measurement models one can consider. As mentioned earlier the resulting forward map  $\mathcal{G}$  is injective [92, 119] even in absence of interior measurements, but it is also ‘severely ill-posed’ in the sense that the recovery guarantees are at best of logarithmic order  $(\log N)^{-a}$ ,  $a > 0$ , in natural statistical measurement settings and with parameter spaces for  $f$  paralleling those in Theorem 2.3.3. A proof of this fact as well as many more relevant references can be found in [1]. We can thus conclude from the previous theorem that the availability of ‘interior’ measurements of the solutions  $u_f$  of (1.2) in Darcy’s problem substantially accelerates these convergence rates to algebraic ones (i.e.,  $N^{-a}$  for some  $a > 0$ ). Note however that in other settings where only boundary measurements (‘scattering data’) are available, algebraic rates can be achieved – see [90] for the non-linear inverse problem with non-Abelian X-ray transforms from (1.1), with proofs that follow the same lines as those laid out in this chapter. The information-theoretic complexity of statistical recovery rates thus depends on possibly subtle interactions between the types of measurements one takes and the underlying partial differential equation. For data assimilation problems discussed in Section 1.1.3, see [94, 99].

## 2.4 Notes

### 2.4.1 Exercises

**Exercise 2.4.1.** Consider the Schrödinger forward map  $\mathcal{G}$  from (2.5) and let  $\theta_0 \in H^{\xi}$  for some  $\xi > 2 + d/2$ . Show that  $\mathcal{G}$  satisfies the forward regularity Condition 2.1.1 with  $\kappa = 2$  and  $\mathcal{R} = H^{\xi}$ . [Hint: adapt the arguments from (2.10) and below it, or see the proof of [100, Theorem 12].] Show further that  $\mathcal{G}$  satisfies the stability Condition 2.1.4 with this choice of  $\mathcal{R}$  and  $\eta = \xi/(2 + \xi)$ . [Hint: Track the constants in (2.12) and then use interpolation and elliptic regularity as in the proof of Proposition 2.1.7; alternatively see [100, Lemma 28].] Then prove an analogue of Theorem 2.3.3 with contraction rates  $\delta_N^{\eta}, \delta_N = N^{-(\alpha+2)/(2\alpha+4+d)}$ .

**Exercise 2.4.2.** For  $\mathcal{L}_f$  from (1.3) with  $f = 1$  (the standard Laplacian), consider solutions  $u_f$  to (1.2) with  $g = 2$  identically and  $h = (|\cdot|_{\mathbb{R}^d}^2 - 1)/d$ . Show that if  $\mathcal{X}$  is a domain separated away from the origin, then  $|\nabla u_1| \geq c > 0$  on  $\mathcal{X}$ . Show by perturbation that this lower bound extends to  $\nabla u_f$  whenever  $\|f - 1\|_{H^\beta} < \epsilon$ ,  $\epsilon$  small enough, and  $\beta > 1 + d$ . Then verify Condition (2.15) for such  $f = f_\theta$ .

**Exercise 2.4.3.** For eigen-pairs  $(\lambda_j, e_j) \in (0, \infty) \times H_0^1$  of the Dirichlet Laplacian  $-L_{1,0}$  from (A.9) and  $\delta_N$  as in (2.19) with  $\kappa = 0$ , consider a prior  $\Pi_N = \text{Law}(\theta)$  as

$$\theta = \frac{1}{\sqrt{N\delta_N}} \sum_{j=1}^D g_j \lambda_j^{-\alpha/2} e_j \quad \text{for } D \lesssim N\delta_N^2 \text{ and } \alpha > 1 + \frac{d}{2}.$$

For  $\theta \in \Theta$  denote its projection onto the linear span  $E_D$  of  $\{e_j : j \leq D\}$  by  $\theta_D$ . Suppose for  $\mathcal{G}$  satisfying Condition 2.1.1 and ground truth  $\theta_0 \in H_0^\alpha$ , we have

$$\|\mathcal{G}(\theta_0) - \mathcal{G}(\theta_{0,D})\|_{L^2} \leq \frac{\delta_N}{2}. \quad (2.29)$$

Show that the conclusions of Theorem 2.2.2 hold true with regularisation sets

$$\Theta_N \subset \{\theta \in H_0^\alpha \cap E_D : \|\theta\|_{H^\alpha} \leq M\}$$

and that hence Theorem 2.3.3 holds with  $\beta$  replaced by  $\alpha$  for this prior. [Hint: If  $\Pi'$  is the law of the infinite series from (B.1) then its RKHS  $\mathcal{H}$  equals  $\tilde{H}^\alpha \subset H_0^\alpha$ . We also have  $\Pi'(\mathcal{R}) = 1$  for  $\mathcal{R} = H_0^\beta$  just as in the last step of the proof of Theorem B.1.3. We can then follow the steps of the proof of Theorem 2.2.2 for this choice of  $\mathcal{H}$ ,  $\mathcal{R}, \kappa = 0$ , and with  $\Pi_N$  the law of  $\theta = \theta'_D / (\sqrt{N}\delta_N)$ ,  $\theta' \sim \Pi'$ . One uses that the projection onto  $E_D$  does not increase  $\mathcal{H}$ ,  $\mathcal{R}$ -norms, and in the small ball calculation (2.22) we use (2.29). On the sets  $\Theta_N$  from (2.20), we can then use the definition of the  $\tilde{H}^\alpha$ -norm in (A.15), Weyl's law (A.14) and Parseval's identity, to estimate

$$\|\theta_1\|_{H^\alpha} \lesssim \|\theta_1\|_{\tilde{H}^\alpha} \lesssim D^{\alpha/d} \|\theta_1\|_{L^2} \leq (N\delta_N^2)^{\alpha/d} \delta_N \leq M.]$$

**Exercise 2.4.4.** Prove Theorem 2.2.2 for the case  $\kappa = 0$  and under the weaker hypothesis where  $\|\theta_1 - \theta_2\|_\infty$  replaces  $\|\theta_1 - \theta_2\|_{(H^0)^*}$  on the right-hand side of (2.4) in Condition 2.1.1. [Hint: The proof is the same, but instead of (2.23), one uses the entropy bound (4.184) in [60] and applies Theorem B.1.1 in the Banach space  $L^\infty(\mathcal{Z})$ .]

## 2.4.2 Remarks and comments

Conditions such as 2.1.1 and 2.1.4 have been used implicitly or explicitly in the statistical analysis of non-linear inverse problems for a while, we mention [93, 97, 100, 133] among others. Their verification for the elliptic PDE models here is fairly routine

but does require some attention to the exact dependence of the constants on the parameter space employed: Proposition 2.1.5 is from [100] – see also [69, 95] – but results of this flavour go back as far as [110], with a recent reference on the topic being [24]. The main ideas combine with more complex non-linear inverse problems where such stability estimates are often also available (but more difficult to obtain), see, e.g., [90, 103], [21, 23] and also more generally [104].

The main mathematical mechanisms behind Theorem 2.2.2 have been introduced in the important contribution [128] for ‘direct’ statistical inference problems, see also [58]. The present form of this theorem relevant for non-linear inverse problems via ‘rescaled’ Gaussian priors that ensure sufficient ‘posterior regularity’ to permit the application of stability estimates is due to [90]. That reference [90] also contains the first proofs of results such as Theorems 2.3.1 and 2.3.2 in the non-linear setting. In the context of Darcy’s problem, the consistency Theorem 2.3.3 was first obtained in [62] by adapting the ideas from [90]. For specific (‘sieved’) Gaussian process priors one can give improved results for the posterior regularity sets  $\Theta_N$  from (2.20), see Exercise 2.4.3 and also [23, 101].

Regarding optimality of the results, one can prove that the rate  $\delta_N$  obtained in (2.27) for the ‘PDE-constrained regression problem’ is optimal in a minimax sense (see [100, Theorem 10]), which implies in particular that Proposition 2.1.3 cannot hold for  $\kappa > 1$ . In contrast, the rate  $\delta_N^\eta$  from (2.28) is unlikely to be sharp, since we have not attempted to optimise the exponent  $\eta$  in the stability estimate in Proposition 2.1.7. For ‘smooth’ conductivities where  $\alpha, \beta \rightarrow \infty$ , the rate in (2.28) *does* approach the optimal convergence rate  $1/\sqrt{N}$  of finite-dimensional models. But the optimal reconstruction rate for Darcy’s problem with general Sobolev regularity remains open at present. For the Schrödinger equation, the minimax optimal rate for inferring  $f$  in  $L^2$ -distance is  $N^{-\alpha/(2\alpha+4+d)}$ , see [93] and Exercises 2.4.1, 2.4.3. Our focus here is not on ‘optimal’ rates but to provide sufficient conditions for posterior consistency with ‘algebraic’ rates  $N^{-a}$ ,  $a > 0$ , for a flexible class of commonly used Gaussian process priors. Such contraction rate results constitute the key ‘localisation’ step in the analysis of posterior measures in subsequent chapters.

## Chapter 3

# Information operators and curvature

The contraction rate theorems from Section 2.3 provide ‘global’ conditions on forward maps  $\mathcal{G}$  that ensure that posterior distributions  $\Pi(\cdot|D_N)$  (and their means) arising from Gaussian process priors concentrate on  $L^2_\lambda$ -balls centred at the ground truth  $\theta_0$ . The radius of these balls was shown to shrink at rate  $N^{-a}$  for some  $a > 0$ , with  $N$  being the sample size, and  $a$  depending on the Sobolev regularity of the prior  $\Pi$ , on analytical properties of the map  $\mathcal{G}$ , and on  $\theta_0$ .

Now suppose one ‘zooms in’ to such a local neighbourhood of  $\theta_0$  of the posterior measure, and that  $\mathcal{G}$  admits a good (say quadratic) approximation by its linearisation operator (derivative)  $D\mathcal{G}_{\theta_0}$  at  $\theta_0$ . In this case the posterior may start to display some remarkable universality features which are the subject of the remainder of these lecture notes. We will see that these features depend on the inverse problem only via  $D\mathcal{G}_{\theta_0}$  and on an appropriate adjoint operator  $D\mathcal{G}_{\theta_0}^*$ , both to be defined carefully below. In analogy to the notion of the ‘information matrix’ used in classical (semi-) parametric statistics [126], one may call  $D\mathcal{G}_{\theta_0}^* D\mathcal{G}_{\theta_0}$  the *information operator*. Both the information operator and  $D\mathcal{G}_{\theta_0}^*$  encode a degree of ‘local identifiability’ of the statistical model near  $\theta_0$ , and also quantify the ‘local curvature’ at  $\theta_0$  of the ‘asymptotic least squares fit functional’. In this chapter we will review and lay out some of these ideas from semi-parametric statistics and then connect them to the high-dimensional inverse problems/PDE settings relevant here.

### 3.1 Information geometry for Gaussian non-linear regression models

The idea to describe ‘information theoretic’ features of statistical models  $\{P_\theta : \theta \in \Theta\}$  by the derivative of the log-likelihood  $(d/d\theta) \log dP_\theta$  at the ‘true value’  $\theta_0$  is at least hundred years old and was advocated perhaps most prominently by R. A. Fisher. See [126] for an account of the classical theory, whose main ideas and terminology we will follow here. In our non-linear regression model (1.10) these derivatives depend on the inverse problem via the linearisation of the forward map  $\mathcal{G}: \Theta \rightarrow L^2_\lambda(\mathcal{X}, V)$  from (1.8), where we recall that the domain  $\Theta$  is assumed to be a (measurable) subset of  $L^2_\xi(\mathcal{Z}, W)$ .

The following condition stipulates the existence of a continuous derivative  $D\mathcal{G}_{\theta_0}$  of  $\mathcal{G}$  at  $\theta_0$  in given directions  $h \in H$ , where  $H$  is a linear *tangent space* of admissible directions. When the parameter space  $\Theta$  is a linear space (as is relevant when using Gaussian process priors), it is natural to choose  $\Theta = H$ , but see after Definition 3.1.2 for more discussion.

**Condition 3.1.1.** Let  $\theta_0 \in \Theta$  and let  $H \subset L^\infty(\mathcal{Z}, W)$  be a linear space such that the paths  $\{\theta_0 + h, \|h\|_\infty < \epsilon, h \in H\}$ , lie in  $\Theta$  for some  $\epsilon > 0$ . Suppose the map  $\mathcal{G}$  from (1.8) satisfies as  $\|h\|_\infty \rightarrow 0$ ,

$$\|\mathcal{G}(\theta_0 + h) - \mathcal{G}(\theta_0) - D\mathcal{G}_{\theta_0}[h]\|_{L_\lambda^2(\mathcal{X}, V)} \equiv \rho_{\theta_0}[h] = o(\|h\|_\infty) \quad (3.1)$$

for some continuous linear operator

$$\mathbb{I}_{\theta_0} \equiv D\mathcal{G}_{\theta_0} : (H, \langle \cdot, \cdot \rangle_{L_\zeta^2(\mathcal{Z}, W)}) \rightarrow L_\lambda^2(\mathcal{X}, V). \quad (3.2)$$

Variations of the previous condition can be considered too but the current form will be convenient for the PDE settings we have in mind.

By the Riesz representation theorem, the operator  $\mathbb{I}_{\theta_0}$  from (3.2) has a Hilbert space adjoint operator

$$\mathbb{I}_{\theta_0}^* : L_\lambda^2(\mathcal{X}, V) \rightarrow \overline{(H, \langle \cdot, \cdot \rangle_{L_\zeta^2(\mathcal{Z}, W)})} \equiv \bar{H} \quad (3.3)$$

(with  $\bar{H}$  denoting Hilbert-space completion), for which

$$\langle \mathbb{I}_{\theta_0} h, g \rangle_{L_\lambda^2} = \langle h, \mathbb{I}_{\theta_0}^* g \rangle_{L_\zeta^2}, \quad h \in H, g \in L_\lambda^2.$$

This then leads to the following definition:

**Definition 3.1.2.** For a forward map  $\mathcal{G}$  satisfying Condition 3.1.1 for tangent space  $H$  and operator  $\mathbb{I}_\theta$  from (3.2), the information operator is defined as

$$\mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0} : (H, \langle \cdot, \cdot \rangle_{L_\zeta^2(\mathcal{Z}, W)}) \rightarrow \overline{(H, \langle \cdot, \cdot \rangle_{L_\zeta^2(\mathcal{Z}, W)})}.$$

Of course we can extend  $\mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0}$  to all of  $\bar{H}$  by continuity. We note that at this stage, the choice of  $\Theta$  is not just one of a natural domain for the map  $\mathcal{G}$ , but also represents a commitment on the statistical model as it determines the (maximal) tangent space  $H$ . For instance, if  $\Theta = H$ , the ‘complexity’ of the set  $H$  of paths near  $\theta_0$  is determined by the complexity of the parameter space  $\Theta$ . In this case, restricting a priori to a  $L_\zeta^2$ -closed subspace  $\Theta_0 \subset \Theta$  changes the adjoint operator  $\mathbb{I}_{\theta_0}^*$  and then also the information operator. But in our infinite-dimensional setting,  $L_\zeta^2$ -dense subspaces give rise to the same definition – for example, choosing  $\Theta = H^\beta(\mathcal{Z})$  or  $\Theta = L^\infty(\mathcal{Z})$  makes no difference as their completions for the  $L_\zeta^2$ -norms are both  $L_\zeta^2$  (when  $\mathcal{Z}$  is a bounded smooth domain in  $\mathbb{R}^d$  and  $\zeta$  is the Lebesgue measure, say).

### 3.1.1 The LAN expansion

A large class of ‘regular’ statistical models  $\{P_\theta : \theta \in \Theta\}$  admit what is called a ‘locally asymptotically normal’ (LAN) expansion of the logarithms of the likelihood ratios

$$\log \frac{dP_{\theta_0+h/\sqrt{N}}}{dP_{\theta_0}}(D_N) \approx N \left( -\frac{1}{2} \|h\|_{\text{LAN}}^2, \|h\|_{\text{LAN}}^2 \right), \quad \theta_0 \in \Theta, \quad (3.4)$$

along local paths  $\theta_0 + h/\sqrt{N} \in \Theta$ ,  $h \in H$ , assuming data  $D_N \sim P_{\theta_0}^N$ , and where  $\langle \cdot, \cdot \rangle_{\text{LAN}}$  defines an inner product on the tangent space  $H$  with ‘LAN-norm’  $\|\cdot\|_{\text{LAN}}$ . The distribution on the right-hand side is not just normal but perhaps more importantly also coincides with the distribution of the log-likelihood ratio (Radon–Nikodym density)

$$\log \frac{dP_h}{dP_0}(Z) \quad \text{for } Z \sim dP_0, \quad (3.5)$$

where  $dP_h$  is the distribution of a Gaussian variable  $Z = h + \mathbb{W}$  with unknown mean ‘shift’ parameter  $h \in H$  and  $\mathbb{W}$  a centred Gaussian white noise process on  $(H, \langle \cdot, \cdot \rangle_{\text{LAN}})$ . (We refer to [60] for the definition of such processes as well as to the Cameron–Martin theorem required to show that the distribution of (3.5) indeed equals the right-hand side on (3.4).) The gist of Le Cam theory [82, 126, 129] is that in LAN settings the asymptotic inference problem in the model  $\{P_\theta : \theta \in \Theta\}$  can locally near  $\theta_0$  be regarded as the much simpler one of inferring the shift  $h$  from a Gaussian observation  $Z \sim dP_h$ .

For the Gaussian regression model (1.9), such a LAN expansion holds as soon as the forward map  $\mathcal{G}$  satisfies some natural conditions.

**Theorem 3.1.3.** *Suppose that Condition 3.1.1 holds for some  $\theta_0 \in \Theta$ , and fix any  $h \in H \subset L^\infty(\mathcal{Z})$ . Assume further that for some constant  $C = C_{\mathcal{G}, h, \epsilon}$ ,*

$$\|\mathcal{G}(\theta_0 + sh) - \mathcal{G}(\theta_0)\|_\infty \leq C|s| \quad \text{for all } |s| \text{ small enough.} \quad (3.6)$$

*Then the log-likelihood ratio process in the model (1.9) with marginal densities given in (1.10) satisfies, as  $N \rightarrow \infty$ , the asymptotic expansion*

$$\log \frac{dP_{\theta_0+h/\sqrt{N}}^N}{dP_{\theta_0}^N}(D_N) = W_N(h) - \frac{1}{2} \|\mathbb{I}_{\theta_0}[h]\|_{L_\lambda^2(\mathbf{x}, \mathcal{V})}^2 + o_{P_{\theta_0}^N}(1) \quad (3.7)$$

*for random variables*

$$W_N \equiv -\frac{1}{\sqrt{N}} \sum_{i=1}^N \langle \mathbb{I}_{\theta_0}[h](X_i), \varepsilon_i \rangle_{\mathcal{V}} \xrightarrow{d} N(0, \|\mathbb{I}_{\theta_0}[h]\|_{L_\lambda^2(\mathbf{x}, \mathcal{V})}^2), \quad (3.8)$$

*where  $\xrightarrow{d}$  denotes convergence in distribution (under  $P_{\theta_0}^N$ ) of real random variables. In particular, the model satisfies the LAN-approximation (3.4):*

$$\log \frac{dP_{\theta_0+h/\sqrt{N}}^N}{dP_{\theta_0}^N}(D_N) \xrightarrow{d} N\left(-\frac{1}{2} \|\mathbb{I}_{\theta_0}[h]\|_{L_\lambda^2(\mathbf{x}, \mathcal{V})}^2, \|\mathbb{I}_{\theta_0}[h]\|_{L_\lambda^2(\mathbf{x}, \mathcal{V})}^2\right) \quad (3.9)$$

*as  $N \rightarrow \infty$ .*

*Proof.* We recall the notation (1.13) so that  $\ell_N(\theta)$  equals  $\log dP_\theta^N$  up to universal additive constants. Then under  $P_{\theta_0}^N$  we have

$$\begin{aligned}
& \ell_N(\theta_0) - \ell_N\left(\theta_0 + \frac{h}{\sqrt{N}}\right) \\
&= -\frac{1}{2} \sum_{i=1}^N \left( |Y_i - \mathfrak{G}(\theta_0)(X_i)|_V^2 - \left| Y_i - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) \right|_V^2 \right) \\
&= -\frac{1}{2} \sum_{i=1}^N \left( |\varepsilon_i|_V^2 - \left| \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) + \varepsilon_i \right|_V^2 \right) \\
&= \sum_{i=1}^N \left\langle \varepsilon_i, \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) \right\rangle_V \\
&\quad + \frac{1}{2} \sum_{i=1}^N \left| \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) \right|_V^2 \equiv \text{I} + \text{II}.
\end{aligned}$$

We can write

$$\begin{aligned}
\text{I} &= \frac{1}{\sqrt{N}} \sum_{i=1}^N \langle \varepsilon_i, D\mathfrak{G}_{\theta_0}[h](X_i) \rangle_V \\
&\quad + \sum_{i=1}^N \left\langle \varepsilon_i, \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) - D\mathfrak{G}_{\theta_0}\left[\frac{h}{\sqrt{N}}\right](X_i) \right\rangle_V.
\end{aligned}$$

The first summand amounts to the term  $W_N$  in (3.7), and (3.8) then follows from  $\mathbb{I}_{\theta_0}[h] \in L_\lambda^2(\mathcal{X}, V)$  and the central limit theorem. Now, by independence and Condition 3.1.1, the last summand has variance bounded by the second moment

$$NE_\lambda \left| \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) - D\mathfrak{G}_{\theta_0}\left[\frac{h}{\sqrt{N}}\right](X_i) \right|_V^2 = N\rho_{\theta_0}^2 \left[ \frac{h}{\sqrt{N}} \right] = o(1),$$

hence, by Markov's inequality, it is  $o_{P_{\theta_0}^N}(1)$  and negligible in the asymptotic distribution. For term II, we first centre at its  $E_\lambda$  expectation

$$\begin{aligned}
& \frac{1}{2} \sum_{i=1}^N \left( \left| \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) \right|_V^2 - E_\lambda \left| \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) \right|_V^2 \right) \\
& \quad + \frac{N}{2} \left\| \mathfrak{G}(\theta_0) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right) \right\|_{L^2(\mathcal{X})}^2. \tag{3.10}
\end{aligned}$$

For the first term we can bound the variance of the summands

$$Z_i = \left| \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right)(X_i) \right|_V^2$$

via the local Lipschitz property (3.6) as

$$EZ_i^2 \leq \left\| \mathfrak{G}(\theta_0) - \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right) \right\|_{\infty}^4 \lesssim \frac{1}{N^2},$$

so that by independence and Chebyshev's inequality  $\text{Var}(\sum_{i=1}^N Z_i/2) \lesssim NN^{-2} \rightarrow 0$  for fixed  $h$  and hence  $\sum_i Z_i/2 = o_{P_{\theta_0}^N}(1)$  is asymptotically negligible. We finally note that

$$\begin{aligned} & \left\| \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right) - \mathfrak{G}(\theta_0) - D\mathfrak{G}_{\theta_0}\left[\frac{h}{\sqrt{N}}\right] + D\mathfrak{G}_{\theta_0}\left[\frac{h}{\sqrt{N}}\right] \right\|_{L^2(\mathfrak{X})}^2 \\ &= \left\| D\mathfrak{G}_{\theta_0}\left[\frac{h}{\sqrt{N}}\right] \right\|_{L^2(\mathfrak{X})}^2 \\ & \quad + \frac{2}{\sqrt{N}} \left\langle D\mathfrak{G}_{\theta_0}[h], \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right) - \mathfrak{G}(\theta_0) - D\mathfrak{G}_{\theta_0}\left[\frac{h}{\sqrt{N}}\right] \right\rangle_{L^2(\mathfrak{X})} \\ & \quad + \left\| \mathfrak{G}\left(\theta_0 + \frac{h}{\sqrt{N}}\right) - \mathfrak{G}(\theta_0) - D\mathfrak{G}_{\theta_0}\left[\frac{h}{\sqrt{N}}\right] \right\|_{L^2(\mathfrak{X})}^2 \\ &= \frac{1}{N} \|D\mathfrak{G}_{\theta_0}[h]\|_{L^2(\mathfrak{X})}^2 + o\left(\frac{1}{N}\right), \end{aligned}$$

where we have used Condition 3.1.1. Therefore the second term in (3.10) features the second term required in the right-hand side of (3.7) plus asymptotically negligible terms, completing the proof of (3.9). ■

### 3.1.2 Cramer–Rao bounds and inverse information

We can analyse finer properties of inference problems with infinite-dimensional parameter spaces  $\Theta$  by studying linear functionals  $\langle \theta, \psi \rangle_{L^2_{\xi}}$ ,  $\theta \in \Theta$ , where  $\psi$  runs through a collection of ‘test functions’. Following standard ideas from functional analysis, we can ‘learn’ various properties of an infinite-dimensional problem by understanding sufficiently large collections of such one-dimensional subproblems. In this section we use ideas from semi-parametric statistics to derive what ‘optimal’ procedures can attain in the LAN model from Theorem 3.1.3 in dependence of analytical properties of the test function  $\psi$  and the operator  $\mathbb{I}_{\theta_0}$ . These ideas set the stage for studying universal Gaussian approximations to Bayesian posterior measures in Chapter 4.

**3.1.2.1 The local asymptotic minimax bound for linear functionals.** For  $H$  the tangent space from Condition 3.1.1, suppose for now that a linear functional  $\Psi : \Theta \rightarrow \mathbb{R}$  can be represented as

$$\Psi(h) = \langle \tilde{\psi}_{\theta_0}, \mathbb{I}_{\theta_0} h \rangle_{L^2_{\lambda}} \quad \text{for all } h \in H, \text{ and for some } \tilde{\psi}_{\theta_0} \in L^2_{\lambda}(\mathfrak{X}, V). \quad (3.11)$$

Notice that by the Cauchy–Schwarz inequality such  $\Psi$  is continuous on  $H$  for the LAN-norm  $\|\mathbb{I}_{\theta_0}(\cdot)\|_{L^2_{\lambda}}$ . Whenever such  $\tilde{\psi}_{\theta_0}$  exists, we can take it to belong to the

closure  $\bar{\mathcal{J}}_H$  of the linear space

$$\mathcal{J}_H = \{w = \mathbb{I}_{\theta_0}(h) : h \in H\}$$

in  $L_\lambda^2$ , by orthogonal projection onto  $\bar{\mathcal{J}}_H$  if necessary. Assuming (3.11) and the LAN expansion from Theorem 3.1.3, the following is a classical result from mathematical statistics [126, 129]. We assume that  $\mathbb{I}_{\theta_0}$  is injective but this is not necessary (see Exercise 3.4.1).

**Theorem 3.1.4.** *Consider data  $(Y_i, X_i)_{i=1}^N$  in the model (1.9), where  $\mathcal{G}$  satisfies the hypotheses of Theorem 3.1.3 with tangent space  $H$ . Suppose that  $\mathbb{I}_{\theta_0}: H \rightarrow L_\lambda^2$  is injective and that the functional  $\Psi: \Theta \rightarrow \mathbb{R}$  satisfies (3.11). Then the local asymptotic minimax risk for estimating  $\Psi$  at  $\theta_0$  is lower bounded as*

$$\liminf_{N \rightarrow \infty} \inf_{\tilde{\psi}_N: (V \times \mathcal{X})^N \rightarrow \mathbb{R}} \sup_{\substack{h \in H \\ \|h\|_{L_\lambda^2} \leq 1/\sqrt{N}}} NE_{\theta_0+h}^N (\tilde{\psi}_N - \Psi(\theta_0 + h))^2 \geq \|\tilde{\psi}_{\theta_0}\|_{L_\lambda^2}^2. \quad (3.12)$$

*Proof.* The result follows from [129, Theorem 3.11.5] applied to the LAN model from Theorem 3.1.3. In the former theorem we choose quadratic loss  $\ell = (\cdot)^2$ , Banach space  $\mathbf{B} = \mathbb{R}$  (so existence of a tight Gaussian limit is implied by the finiteness of  $\|\tilde{\psi}_{\theta_0}\|_{L_\lambda^2}$ ) and with Hilbert space norm  $\|\cdot\| = \|\mathbb{I}_{\theta_0}(\cdot)\|_{L_\lambda^2}$  on the tangent space  $H$ , noting also that the map  $\Psi$  satisfies, by (3.11) and the Cauchy–Schwarz inequality,

$$|\Psi(h)| \leq \|\tilde{\psi}_{\theta_0}\|_{L_\lambda^2} \|h\|,$$

and it is hence continuous from  $H \rightarrow \mathbf{B}$ . ■

In particular, for functionals  $\Psi$  as in (3.11), there is hope for the existence of efficient  $\sqrt{N}$ -consistent estimators that attain this lower bound. Let us shed some more light on the condition (3.11). If  $\Psi$  is of the more natural form

$$\Psi(\theta) = \langle \theta, \psi \rangle_{L_\xi^2} \quad \text{for some } \psi \in L_\xi^2, \quad (3.13)$$

then (3.11) can be rewritten as requiring that

$$\psi = \mathbb{I}_{\theta_0}^* \tilde{\psi}_{\theta_0} \quad \text{for some } \tilde{\psi}_{\theta_0} \in L_\lambda^2, \quad (3.14)$$

in other words, that  $\psi$  lies in the range  $R(\mathbb{I}_{\theta_0}^*)$  of the adjoint operator  $\mathbb{I}_{\theta_0}^*$ . This is related to the local identifiability of our model via the standard Hilbert space identity ( $\perp$  denoting orthogonal complement)

$$\ker(\mathbb{I}_{\theta_0}) = R(\mathbb{I}_{\theta_0}^*)^\perp, \quad (3.15)$$

so it requires, in particular, that  $\psi$  does not lie in the kernel of  $\mathbb{I}_{\theta_0}$  (acting on  $L_\xi^2$ ). But note that unless the range of  $\mathbb{I}_{\theta_0}^*$  is closed, even if  $\psi \notin \ker(\mathbb{I}_{\theta_0})$ , so that  $\psi \in \overline{R(\mathbb{I}_{\theta_0}^*)}$ , it does not need to belong to  $R(\mathbb{I}_{\theta_0}^*)$ .

A sufficient (but not necessary) condition for (3.14) is that  $\psi$  lies in the range of the information operator  $\mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0}$ . Indeed if  $\tilde{\psi} \in H$  solves the *information equation*

$$\mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0} \tilde{\psi} = \psi \quad (3.16)$$

for the given  $\psi$ , then we can use  $\tilde{\psi}_{\theta_0} = \mathbb{I}_{\theta_0} \tilde{\psi} \in L_{\lambda}^2$  in (3.11). We will turn to this issue again in Section 4.1.

**3.1.2.2 The information number and its inverse.** When (3.11) holds, the lower bound in Theorem 3.1.4 has variational characterisation

$$\|\tilde{\psi}_{\theta_0}\|_{L_{\lambda}^2}^2 = \sup_{0 \neq w = \mathbb{I}_{\theta_0}(h), h \in H} \frac{\langle \tilde{\psi}_{\theta_0}, w \rangle_{L_{\lambda}^2}^2}{\langle w, w \rangle_{L_{\lambda}^2}}. \quad (3.17)$$

Indeed, the Cauchy–Schwarz inequality gives one side of the inequality and the converse follows from taking a sequence  $w \in \mathcal{J}_H$  approaching  $\tilde{\psi}_{\theta_0} \in \bar{\mathcal{J}}_H$ .

If the target of inference is a functional of the form  $\Psi(\theta) = \langle \theta, \psi \rangle_{L^2}$  from (3.13), then the characterisation (3.17) can be related to the *efficient information number*  $i_{\theta_0, H, \psi}$  for (unbiased) estimation of  $\Psi$  over the sub-models  $\{\theta_0 + sh : |s| < \epsilon, h \in H\}$  appearing in the classical Cramér–Rao theorem: We first set

$$i_{\theta_0, h, \psi} := \frac{\|\mathbb{I}_{\theta_0} h\|_{L_{\lambda}^2}^2}{\langle \psi, h \rangle_{L_{\xi}^2}^2}, \quad h \in H,$$

and define

$$i_{\theta_0, H, \psi} \equiv \inf_{h \in H, \langle \psi, h \rangle_{L_{\xi}^2} \neq 0} \frac{\|\mathbb{I}_{\theta_0} h\|_{L_{\lambda}^2}^2}{\langle \psi, h \rangle_{L_{\xi}^2}^2}. \quad (3.18)$$

Next, if the range condition (3.14) holds, then we can rewrite the denominator as  $\langle \mathbb{I}_{\theta_0}^* \tilde{\psi}_{\theta_0}, h \rangle_{L_{\xi}^2}^2 = \langle \tilde{\psi}_{\theta_0}, \mathbb{I}_{\theta_0} h \rangle_{L_{\lambda}^2}^2$ . Also (3.15) implies that  $\psi$  is then necessarily orthogonal on the kernel of  $\mathbb{I}_{\theta_0}$  and we deduce

$$\|\tilde{\psi}_{\theta_0}\|_{L_{\lambda}^2}^2 = \sup_{h \in H, \mathbb{I}_{\theta_0} h \neq 0} \frac{\langle \tilde{\psi}_{\theta_0}, \mathbb{I}_{\theta_0} h \rangle_{L_{\lambda}^2}^2}{\langle \mathbb{I}_{\theta_0} h, \mathbb{I}_{\theta_0} h \rangle_{L_{\lambda}^2}} = i_{\theta_0, H, \psi}^{-1}. \quad (3.19)$$

**3.1.2.3 Non-existence of the inverse information.** When the information number  $i_{\theta_0, H, \psi}$  is non-zero so that the quantity (3.19) is finite, the lower bound from Theorem 3.1.4 can often be ‘matched’ by non-parametric Bayesian methods by (i) solving the information equation (3.16) and by (ii) making the techniques underlying the proof of Theorem 3.1.3 quantitative and ‘uniform’ in relevant classes of alternatives  $h$  – this will be the content of Section 4.1.

In contrast, if the lower bound in Theorem 3.1.4 is *infinite*, this introduces fundamental limitations to the proof methods arising from LAN expansions. Specifically, when considering inference on functionals  $\Psi$  from (3.13), the following theorem shows that the adjoint range condition (3.14) is necessary for the existence of a non-zero information number (3.18), and for the possibility of locally uniformly  $\sqrt{N}$ -consistent estimation of the scalar parameter  $\Psi(\theta)$ . This will become relevant in Chapter 4, too.

**Theorem 3.1.5.** *Consider data  $(Y_i, X_i)_{i=1}^N$  in the model (1.9), where  $\mathcal{G}$  satisfies the hypotheses of Theorem 3.1.3 with tangent space  $H$  and adjoint score operator  $\mathbb{I}_{\theta_0}^*$  from (3.3). Consider a functional  $\Psi(\theta) = \langle \theta, \psi \rangle_{L_\xi^2}$  for some  $\psi \in L_\xi^2$  for which*

$$\psi \notin R(\mathbb{I}_{\theta_0}^*) := \{h = \mathbb{I}_{\theta_0}^* g \text{ for some } g \in L_\lambda^2(\mathcal{X}, V)\}. \quad (3.20)$$

Then the efficient information number from (3.18) equals zero,  $i_{\theta_0, H, \psi} = 0$ , and

$$\liminf_{N \rightarrow \infty} \inf_{\bar{\psi}_N : (V \times \mathcal{X})^N \rightarrow \mathbb{R}} \sup_{\substack{h \in H \\ \|h\|_{L_\xi^2} \leq 1/\sqrt{N}}} NE_{\theta_0+h}^N(\bar{\psi}_N - \Psi(\theta_0 + h))^2 = \infty. \quad (3.21)$$

*Proof.* That  $i_{\theta_0, H, \psi} = 0$  follows from a functional analytic argument showing that the ratio in (3.18) necessarily vanishes if (3.20) holds – see Theorem B.4.1 in the appendix for the proof.

Then assuming  $i_{\theta_0, H, \psi} = 0$ , we augment the observation space to include measurements  $(Z_i, Y_i, X_i)_{i=1}^N \sim \bar{P}_{\theta_0}^N$ , where the

$$Z_i \stackrel{\text{i.i.d.}}{\sim} N(\langle \theta_0, \psi \rangle_{L_\xi^2}, \sigma^2)$$

are independent of the  $(Y_i, X_i)$ 's, and where  $\sigma^2 > 0$  is known but arbitrary. The new statistical model  $\{\bar{P}_\theta^N : \theta \in \Theta\}$  again verifies the LAN approximation (3.4) with ‘augmented’ LAN norm

$$\|h\|_{\text{LAN}}^2 = \|\bar{\mathbb{I}}_{\theta_0} h\|_{L_\lambda^2}^2 = \|\mathbb{I}_{\theta_0} h\|_{L_\lambda^2}^2 + \sigma^{-2} \langle \psi, h \rangle_{L_\xi^2}^2, \quad h \in H, \quad (3.22)$$

as follows from Theorem 3.1.3, a standard tensorisation argument for the LAN property in independent sample spaces, and the fact that a  $\mathcal{N}(\langle \theta_0, \psi \rangle_{L_\xi^2}, \sigma^2)$  model has LAN ‘norm’  $\sigma^{-2} \langle \psi, h \rangle_{L_\xi^2}^2$ , by a simple direct calculation with Gaussian densities. The new efficient information from (3.18) for estimating  $\Psi(\theta)$  from the augmented data (and with identical tangent space) is now of the form

$$\bar{i}_{\theta_0, H, \psi} = \inf_h \frac{\|\bar{\mathbb{I}}_{\theta_0} h\|_{L_\lambda^2}^2 + \sigma^{-2} \langle \psi, h \rangle_{L_\xi^2}^2}{\langle \psi, h \rangle_{L_\xi^2}^2} = i_{\theta_0, H, \psi} + \sigma^{-2} = \sigma^{-2} > 0,$$

where we have used the fact that  $i_{\theta_0, H, \psi} = 0$ .

To proceed we note that the linear functional  $\Psi(\cdot)$  is continuous on  $H$  for the augmented LAN-norm (3.22) so that we can invoke the Riesz representation theorem on  $\tilde{H}$  (or, if  $\bar{\mathbb{I}}_{\theta_0}$  is not injective on  $\tilde{H}$ , on the quotient space of  $\tilde{H}$  for the kernel of  $\bar{\mathbb{I}}_{\theta_0}$ ) to the effect that

$$\Psi(h) = \langle \bar{\mathbb{I}}_{\theta_0} \tilde{h}, \bar{\mathbb{I}}_{\theta_0} h \rangle_{L^2_\lambda}, \quad h \in H, \text{ for some } \tilde{\psi}_{\theta_0} \equiv \bar{\mathbb{I}}_{\theta_0} \tilde{h} \in \overline{(\tilde{J}_H)}.$$

In particular,  $\Psi$  now verifies (3.11) in the augmented setting and one then shows (see Exercise 3.4.2) that mutatis mutandis, (3.12) and (3.19) hold in the augmented model as well, with linearisation operator  $\bar{\mathbb{I}}_{\theta_0}$  and tangent space  $H$ . In particular,

$$\liminf_{N \rightarrow \infty} \inf_{\tilde{\psi}_N: (\mathbb{R} \times V \times \mathcal{X})^N \rightarrow \mathbb{R}} \sup_{\substack{h \in H \\ \|h\|_{L^2_\xi} \leq 1/\sqrt{N}}} NE_{\theta_0+h}^N(\tilde{\psi}_N - \Psi(\theta_0 + h))^2 \geq \bar{\tau}_{\theta_0, H, \psi}^{-1} = \sigma^2$$

for estimators  $\tilde{\psi}$  based on the more informative data. The asymptotic local minimax risk in (3.21) exceeds the quantity in the last display, and letting  $\sigma^2 \rightarrow \infty$  implies the result.  $\blacksquare$

## 3.2 Gradient stability, local convexity, and concentration

In this section we look at another aspect that is determined by the linearisation operator  $D\mathcal{G}_{\theta_0} = \mathbb{I}_{\theta_0}$  of  $\mathcal{G}$  at the ground truth  $\theta_0$  – namely the *local curvature* of the log-likelihood function  $\ell_N(\theta)$  from (1.13) featuring in the expression of the posterior measure (1.12). We will show that under certain conditions on  $D\mathcal{G}_{\theta_0}$ , the ‘least squares fit’  $-\ell_N(\theta)$  is locally convex near the ground truth  $\theta_0$  with high  $P_{\theta_0}^N$ -probability. This implies that the posterior measure  $\Pi(\cdot|D_N)$  arising from Gaussian priors will be locally *log-concave* in that region, which will be exploited further in Chapter 5.

### 3.2.1 Convexity of $-\ell_N$ and the gradient of $\mathcal{G}$

We introduce *finite-dimensional* spaces

$$\mathbb{R}^D \simeq \text{span}\{e_j : 1 \leq j \leq D\} \equiv E_D \subset \Theta, \quad (3.23)$$

arising from a set of orthonormal basis functions  $e_j \in L^2_\xi(\mathcal{Z}, W)$  used to ‘discretise’ the ambient linear parameter space  $\Theta \subset L^2_\xi$ . We denote the Euclidean norm on  $E_D \simeq \mathbb{R}^D$  by  $\|\cdot\|_{E_D}$ . A prototypical example pursued below is to take the eigenfunctions of the Dirichlet Laplacian  $\Delta$  on a bounded smooth domain  $\mathcal{Z}$  as in Exercise 2.4.3 from earlier. While  $D$  is finite, we think of it as ‘high-dimensional’,

since letting  $D \rightarrow \infty$  will reduce the approximation error. The gradient vectors  $\nabla \mathcal{G}(\theta)$  of  $\mathcal{G}(\cdot)(x): E_D \rightarrow \mathbb{R}$ ,  $x \in \mathcal{X}$ , are now defined via partial derivatives

$$\frac{\partial}{\partial \theta_j} \mathcal{G}(\theta) = \frac{\partial}{\partial t_j} \mathcal{G}\left(\sum_{j=1}^D t_j e_j\right), \quad \text{with } t_j = \langle \theta, e_j \rangle_{L_\xi^2}, \quad (3.24)$$

where  $\theta = \sum_{j=1}^D t_j e_j \in E_D$ .

In view of the non-linearity of  $\mathcal{G}$ , the ‘least squares fit’  $-\ell_N(\theta) = -\sum_{i=1}^N \ell_i(\theta)$  from (1.13) is neither globally nor locally convex for the given observation vector. Indeed, assuming  $\mathcal{G}$  is sufficiently smooth, the empirical Hessian of  $-\ell_N(\theta)$  equals the sum of all

$$-\nabla^2 \ell_i(\theta) = [\nabla \mathcal{G}(\theta)(X_i)][\nabla \mathcal{G}(\theta)(X_i)]^\top + [\mathcal{G}(\theta)(X_i) - Y_i][\nabla^2 \mathcal{G}(\theta)(X_i)]. \quad (3.25)$$

Quadratic forms of the last expression may have an arbitrary sign, different for distinct  $\theta$ , and so the Hessian matrix is not necessarily positive semi-definite. But if we assume for the moment that  $\theta_0 \in E_D$  (later we will consider its projection  $\theta_{0,D}$  onto  $E_D$ ) and take  $E_{\theta_0}$  expectations under the model equation (1.9) in the last identity, we see that the second term in (3.25) vanishes at  $\theta = \theta_0$ , and that hence for any  $v \in E_D \simeq \mathbb{R}^D$ ,

$$v^\top [-E_{\theta_0} \nabla^2 \ell_i(\theta_0)] v = \|\nabla \mathcal{G}(\theta_0)^\top v\|_{L_\lambda^2}^2.$$

We deduce that if the gradient  $\nabla \mathcal{G}(\theta)$  satisfies the ‘gradient stability’ estimate

$$\|\nabla \mathcal{G}(\theta_0)^\top v\|_{L_\lambda^2}^2 \gtrsim \|v\|_{E_D}^2, \quad (3.26)$$

then *on average* the negative log-likelihood function has a positive definite Hessian at  $\theta_0$ , and we expect this bound to extend (by continuity) to  $-E_{\theta_0} \nabla^2 \ell(\theta)$  at least for  $\theta$  in a neighbourhood of  $\theta_0$ . We can then use tools from high-dimensional probability to deduce that the discretely sampled functions  $\ell_N$  concentrate around their statistical means sufficiently well to inherit the local curvature.

The hypothesis (3.26) is a quantitative injectivity condition on the linearisation operator  $D\mathcal{G}_{\theta_0}$  from Condition 3.1.1 expressed in terms of its action on the discretisation space  $E_D$ . Under Condition 3.1.1 with  $H \supset E_D$ , the requirement (3.26) becomes

$$\langle \nabla \mathcal{G}(\theta_0)^\top v, \nabla \mathcal{G}(\theta_0)^\top v \rangle_{L_\lambda^2} = \langle D\mathcal{G}_{\theta_0}[v], D\mathcal{G}_{\theta_0}[v] \rangle_{L_\lambda^2} = \langle v, D\mathcal{G}_{\theta_0}^* D\mathcal{G}_{\theta_0}[v] \rangle_{L_\xi^2},$$

and therefore amounts to a lower bound on the minimal eigenvalue of the matrix  $\nabla \mathcal{G}(\theta_0) \nabla \mathcal{G}(\theta_0)^\top$  arising from the  $L_\xi^2$ -action of the information operator  $\mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0}$  from (3.1.2) on the finite-dimensional subspace  $E_D$ . Such a result is implied in particular by an infinite-dimensional stability estimate

$$\|D\mathcal{G}_{\theta_0}[h]\|_{L_\lambda^2} = \|\mathbb{I}_{\theta_0}[h]\|_{L_\lambda^2} \gtrsim \|h\|_+ \quad \text{for } h \in H, \quad (3.27)$$

for an appropriate  $\|\cdot\|_+$  norm that is ‘compatible’ with the  $E_D$  scale. When  $E_D$  is the eigen-basis of the Dirichlet Laplacian we will verify this condition for our model PDEs – see Theorem 5.3.2 for Darcy’s problem and (3.48) as well as Exercise 5.4.5 for the Schrödinger equation.

### 3.2.2 A concentration result for the empirical Hessian

We will require ‘local’ conditions near a fixed  $\theta_0 \in \Theta \subset L^2_\xi(\mathcal{Z}, W)$ , in fact near its  $L^2_\xi$ -orthonormal projection  $\theta_{0,D}$  onto  $E_D$ . For radius  $\mathbf{r} > 0$  to be chosen, set

$$\mathcal{B} := \{\theta \in E_D : \|\theta - \theta_{0,D}\|_{E_D} < \mathbf{r}\}, \quad (3.28)$$

where we write  $\|\cdot\|_{E_D} \simeq \|\cdot\|_{\mathbb{R}^D}$  for the  $L^2_\xi$ -norm on  $E_D$  in this subsection.

The following local regularity condition on the map  $\mathcal{G}$  is satisfied, for instance, as soon as  $\mathcal{G}$  is  $C^3$  on  $\mathcal{B}$  with local  $C^3$ -norm constants growing at most *polynomially* in dimension  $D$ . To formulate it let us define the following local  $C^{2,1}$ -norm for maps  $F: \mathcal{B} \rightarrow V$ ,  $F(\theta) = (F_1(\theta), \dots, F_{d_V}(\theta))$ ,

$$\begin{aligned} \|F\|_{C^{2,1}(\mathcal{B}, V)} := & \max_{1 \leq k \leq d_V} \left\{ \|F_k\|_\infty + \|\nabla F_k\|_{L^\infty(\mathcal{B}, \mathbb{R}^D)} + \|\nabla^2 F_k\|_{L^\infty(\mathcal{B}, \mathbb{R}^{D \times D})} \right. \\ & \left. + \sup_{\theta, \theta' \in \mathcal{B}, \theta \neq \theta'} \frac{|\nabla^2 F_k(\theta) - \nabla^2 F_k(\theta')|_{\text{op}}}{\|\theta - \theta'\|_{E_D}} \right\}, \end{aligned}$$

where  $\|\cdot\|_{L^\infty(\mathcal{B}, U)}$  are the supremum norms of maps defined on  $\mathcal{B}$  taking values in a normed vector space  $U$ , the space  $\mathbb{R}^{D \times D}$  is equipped with the usual operator norm  $|\cdot|_{\text{op}}$ , and  $\nabla = \nabla_\theta$  and  $\nabla^2$  denote the gradient and ‘Hessian’ operator, respectively, arising via (3.24).

**Condition 3.2.1** (Local regularity). Let  $\mathcal{B}$  be given in (3.28) for some  $\mathbf{r} > 0$ , and suppose that for all  $x \in \mathcal{X}$ , the map  $\theta \mapsto \mathcal{G}(\theta)(x)$  from (1.8) is in  $C^{2,1}(\mathcal{B}, V)$  and satisfies  $\sup_{x \in \mathcal{X}} \|\mathcal{G}(\cdot)(x)\|_{C^{2,1}(\mathcal{B}, V)} \leq c_2 D^{\kappa_2}$  for some  $c_2 \geq 1$  and  $\kappa_2 \geq 0$ .

For the next condition let us define

$$\ell(\theta, (Y, X)) := -\frac{1}{2} |Y - \mathcal{G}(\theta)(X)|_V^2, \quad \theta \in \Theta, \quad (3.29)$$

which equals  $\ell_N$  from (1.13) for a single ‘generic’ observation  $(Y, X)$  from model (1.9). Regarding  $\ell(\cdot, (Y, X))$  as a real-valued map defined on  $E_D$ , Condition 3.2.1 and the chain rule imply that the gradient vector  $\nabla \ell(\theta, (Y, X))$  and Hessian matrix  $\nabla^2 \ell(\theta, (Y, X))$  with respect to  $\theta \in E_D$  exist for every  $(Y, X)$ . In particular, the  $D \times D$  matrix  $E_{\theta_0}[-\nabla^2 \ell(\theta, (Y, X))]$  is symmetric and  $\lambda_{\min}(A)$  will denote the smallest eigenvalue of a symmetric matrix  $A$ .

**Condition 3.2.2** (Local average curvature). Let  $\mathcal{B}$  be given in (3.28) for some  $\mathbf{r} > 0$ . Assume that for some  $c_0 > 0, c_1 \geq 1, \kappa_0, \kappa_1 \geq 0$  and all  $D \in \mathbb{N}$ ,

$$\inf_{\theta \in \mathcal{B}} \lambda_{\min}(E_{\theta_0}[-\nabla^2 \ell(\theta, (Y, X))]) \geq c_0 D^{-\kappa_0} \quad (3.30)$$

and

$$\begin{aligned} \sup_{\theta \in \mathcal{B}} [ |E_{\theta_0} \ell(\theta, (Y, X))| + \|E_{\theta_0}[\nabla \ell(\theta, (Y, X))]\|_{\mathbb{R}^D} \\ + \|E_{\theta_0}[\nabla^2 \ell(\theta, (Y, X))]\|_{\text{op}} ] \leq c_1 D^{\kappa_1}. \end{aligned} \quad (3.31)$$

As discussed in the preceding subsection, the verification of (3.30) for  $\mathcal{B}$  a neighbourhood of  $\theta_{0,D}$  can be reduced to the condition (3.26) combined with Lipschitz continuity of the Hessian with respect to the  $\theta$ -variable. Condition (3.31) is implied by Condition 3.2.1 for  $\kappa_1 = \kappa_2$  but we still record it separately to permit a potentially better constant  $\kappa_1 < \kappa_2$ .

The following theorem is the main result of this subsection. It is based on tools from high-dimensional probability. The constraint  $\mathbf{r} \leq 1$  is not necessary but convenient to permit a simple expression in (3.32) for the concentration exponent  $\mathcal{R}_N$ , and also natural for the applications we have in mind. A more detailed version of this result can be found in [101].

**Theorem 3.2.3.** *Let  $\ell_N: E_D \rightarrow \mathbb{R}$  be given by (1.13). Suppose Conditions 3.2.1 and 3.2.2 are satisfied for  $\mathcal{B}$  from (3.28) and some  $\mathbf{r} \leq 1$ . There exists a constant  $C = C(c_0, c_1, c_2) > 0$  such that if*

$$\mathcal{R}_N := CND^{-2\kappa_0-4\kappa_2}, \quad (3.32)$$

then for any  $D, N \geq 1$  satisfying  $D \leq \mathcal{R}_N$ , we have, for constants  $c, c'$  depending only on  $d_V$ ,

$$P_{\theta_0}^N \left( \inf_{\theta \in \mathcal{B}} \lambda_{\min}[-\nabla^2 \ell_N(\theta)] < \frac{c_0}{2} ND^{-\kappa_0} \right) \leq ce^{-\mathcal{R}_N}, \quad (3.33)$$

as well as

$$\begin{aligned} P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} [ |\ell_N(\theta)| + \|\nabla \ell_N(\theta)\|_{\mathbb{R}^D} + \|\nabla^2 \ell_N(\theta)\|_{\text{op}} ] > N(5c_1 D^{\kappa_1} + 1) \right) \\ \leq c'(e^{-\mathcal{R}_N} + e^{-N/8}). \end{aligned} \quad (3.34)$$

*Proof.* It suffices to prove the assertion for  $\mathcal{R}_N \geq 1$ . We also restrict to the case  $V = W = \mathbb{R}$  relevant in these notes – considering general vector-spaces  $V, W$  of fixed finite dimension is only notationally different (e.g., [23, Lemma 5.6]).

For the proof we require some more notation: For  $x \in \mathcal{X}$ , we write shorthand  $\mathfrak{G}^x(\theta) := \mathfrak{G}(\theta)(x)$  and  $Z := (Y, X) \sim P_{\theta_0}$ . Throughout,  $P_N := N^{-1} \sum_{i=1}^N \delta_{Z_i}$  denotes

the empirical measure induced by  $Z_i \equiv (Y_i, X_i)_{i=1}^N$ , which acts linearly on measurable functions  $h: \mathbb{R} \times \mathcal{X} \rightarrow \mathbb{R}$  via

$$P_N(h) = \int_{\mathbb{R} \times \mathcal{X}} h dP_N = \frac{1}{N} \sum_{i=1}^N h(Z_i).$$

The following standard identities follow from the chain rule and will be used repeatedly in the proofs below – recall that  $\nabla$  and  $\nabla^2$  act on the  $\theta$ -variable in  $E_D$ :

$$\begin{aligned} -\ell(\theta, Z) &= \frac{1}{2}[Y - \mathcal{G}^X(\theta)]^2 = \frac{1}{2}[\mathcal{G}^X(\theta_0) + \varepsilon - \mathcal{G}^X(\theta)]^2, \\ -\nabla\ell(\theta, Z) &= [\mathcal{G}^X(\theta) - Y]\nabla\mathcal{G}^X(\theta), \\ -\nabla^2\ell(\theta, Z) &= \nabla\mathcal{G}^X(\theta)\nabla\mathcal{G}^X(\theta)^\top + [\mathcal{G}^X(\theta) - Y]\nabla^2\mathcal{G}^X(\theta), \\ -E_{\theta_0}[\ell(\theta, Z)] &= \frac{1}{2} + \frac{1}{2}E^X[\mathcal{G}^X(\theta_0) - \mathcal{G}^X(\theta)]^2. \end{aligned} \tag{3.35}$$

When no confusion can arise, we will suppress the second argument  $Z$  and write  $\ell(\theta)$  for  $\ell(\theta, Z)$ . We also set  $\bar{\ell}_N := \ell_N/N$ .

By Condition 3.2.1, the matrix  $-\nabla^2\bar{\ell}_N(\theta)$  is symmetric and by a standard inequality for eigenvalues due to Weyl and Condition 3.2.2, we have for any  $\theta \in \mathcal{B}$  that

$$\begin{aligned} \lambda_{\min}[-\nabla^2\bar{\ell}_N(\theta)] &\geq \lambda_{\min}(E_{\theta_0}[-\nabla^2\ell(\theta)]) - \|\nabla^2\bar{\ell}_N(\theta) - E_{\theta_0}[\nabla^2\ell(\theta)]\|_{\text{op}} \\ &\geq c_0D^{-\kappa_0} - \|\nabla^2\bar{\ell}_N(\theta) - E_{\theta_0}[\nabla^2\ell(\theta)]\|_{\text{op}}. \end{aligned}$$

Hence we deduce

$$\begin{aligned} P_{\theta_0}^N \left( \inf_{\theta \in \mathcal{B}} \lambda_{\min}[-\nabla^2\ell_N(\theta, Z)] < \frac{Nc_0D^{-\kappa_0}}{2} \right) \\ &\leq P_{\theta_0}^N \left( \|\nabla^2\bar{\ell}_N(\theta) - E_{\theta_0}[\nabla^2\ell(\theta)]\|_{\text{op}} \geq \frac{c_0D^{-\kappa_0}}{2} \text{ for some } \theta \in \mathcal{B} \right) \\ &\leq P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} \sup_{v: \|v\|_{\mathbb{R}^D} \leq 1} |v^\top(\nabla^2\bar{\ell}_N(\theta) - E_{\theta_0}[\nabla^2\ell(\theta)])v| \geq \frac{c_0D^{-\kappa_0}}{2} \right). \end{aligned} \tag{3.36}$$

The next step is to reduce the supremum over  $\{v : \|v\|_{\mathbb{R}^D} \leq 1\}$  to a suitable finite maximum over grid points  $v_i$  by a contraction argument (commonly used in high-dimensional probability, e.g., [131]). For  $\rho > 0$ , let  $N(\rho)$  denote the minimal number of balls of  $\|\cdot\|_{\mathbb{R}^D}$ -radius  $\rho$  required to cover  $\{v : \|v\|_{\mathbb{R}^D} \leq 1\}$ , and let  $v_i, \|v_i\|_{\mathbb{R}^D} \leq 1$ , be the centre points of a minimal covering. Thus for any  $v \in \mathbb{R}^D$ , there exists an index  $i$  such that  $\|v - v_i\|_{\mathbb{R}^D} \leq \rho$ . Hence, writing shorthand

$$M_\theta = \nabla^2\bar{\ell}_N(\theta) - E_{\theta_0}[\nabla^2\ell(\theta)], \quad \theta \in \mathcal{B},$$

we have by the Cauchy–Schwarz inequality and the symmetry of the matrix  $M_\theta$ ,

$$\begin{aligned} v^\top M_\theta v &= v_i^\top M_\theta v_i + (v - v_i)^\top M_\theta v + v_i^\top M_\theta (v - v_i) \\ &\leq v_i^\top M_\theta v_i + \|v - v_i\|_{\mathbb{R}^D} \|M_\theta v\|_{\mathbb{R}^D} + \|v - v_i\|_{\mathbb{R}^D} \|M_\theta v_i\|_{\mathbb{R}^D} \\ &\leq v_i^\top M_\theta v_i + 2\rho \sup_{v: \|v\|_{\mathbb{R}^D} \leq 1} v^\top M_\theta v. \end{aligned}$$

Choosing  $\rho = 1/4$  and taking suprema it follows that for any  $\theta \in \mathcal{B}$ ,

$$\sup_{v: \|v\|_{\mathbb{R}^D} \leq 1} v^\top M_\theta v \leq 2 \max_{i=1, \dots, N(1/4)} v_i^\top M_\theta v_i.$$

Since the covering  $(v_i)$  is independent of  $\theta$ , we can further estimate the right-hand side of (3.36) by a union bound to the effect that

$$\begin{aligned} &P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} \sup_{v: \|v\|_{\mathbb{R}^D} \leq 1} |v^\top M_\theta v| \geq \frac{c_0 D^{-\kappa_0}}{2} \right) \\ &\leq N \left( \frac{1}{4} \right) \cdot \sup_{v: \|v\|_{\mathbb{R}^D} \leq 1} P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} |v^\top M_\theta v| \geq \frac{c_0 D^{-\kappa_0}}{4} \right) \\ &\leq N \left( \frac{1}{4} \right) \cdot \sup_{v: \|v\|_{\mathbb{R}^D} \leq 1} \left[ P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} |P_N(g_{v, \theta} - g_{v, \theta_{0,D}})| \geq \frac{c_0 D^{-\kappa_0}}{8} \right) \right. \\ &\quad \left. + P_{\theta_0}^N \left( |P_N(g_{v, \theta_{0,D}})| \geq \frac{c_0 D^{-\kappa_0}}{8} \right) \right], \end{aligned} \quad (3.37)$$

where we have defined

$$g_{v, \theta}(\cdot) := v^\top (\nabla^2 \ell(\theta, \cdot) - E_{\theta_0}[\nabla^2 \ell(\theta)])v, \quad v \in \mathbb{R}^D.$$

and where we recall that  $\theta_{0,D}$  is the centre-point of the set  $\mathcal{B}$  from (3.28). For the rest of the proof, we fix any  $v \in \mathbb{R}^D$  with  $\|v\|_{\mathbb{R}^D} \leq 1$ . Next, we use (3.35) to decompose the ‘uncentred’ part of  $g_{v, \theta}$  as

$$\begin{aligned} &-v^\top \nabla^2 \ell(\theta, Z)v \\ &= v^\top [\nabla \mathcal{G}^X(\theta) \nabla \mathcal{G}^X(\theta)^\top + [\mathcal{G}^X(\theta) - \mathcal{G}^X(\theta_0)] \nabla^2 \mathcal{G}^X(\theta)]v - \varepsilon v^\top \nabla^2 \mathcal{G}^X(\theta)v \\ &=: \tilde{g}_{v, \theta}^I(X) + \varepsilon g_{v, \theta}^{\text{II}}(X), \end{aligned}$$

such that

$$g_{v, \theta}(z) = \tilde{g}_{v, \theta}^I(x) + \varepsilon g_{v, \theta}^{\text{II}}(x),$$

where we have defined the centred version of  $\tilde{g}_{v, \theta}^I$  as

$$g_{v, \theta}^I(x) = \tilde{g}_{v, \theta}^I(x) - E_{\theta_0}[\tilde{g}_{v, \theta}^I(X)], \quad x \in \mathcal{X}.$$

We can therefore bound the right-hand side of (3.37) by

$$\begin{aligned}
 & N\left(\frac{1}{4}\right) \cdot \sup_{v: \|v\|_{\mathbb{R}^D} \leq 1} \left[ P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} \left| \frac{1}{N} \sum_{i=1}^N (g_{v,\theta}^I - g_{v,\theta_0,D}^I)(X_i) \right| \geq \frac{c_0 D^{-\kappa_0}}{16} \right) \right. \\
 & + P_{\theta_0}^N \left( \left| \frac{1}{N} \sum_{i=1}^N g_{v,\theta_0,D}^I(X_i) \right| \geq \frac{c_0 D^{-\kappa_0}}{16} \right) \\
 & + P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} \left| \frac{1}{N} \sum_{i=1}^N \varepsilon_i (g_{v,\theta}^{\text{II}} - g_{v,\theta_0,D}^{\text{II}})(X_i) \right| \geq \frac{c_0 D^{-\kappa_0}}{16} \right) \\
 & \left. + P_{\theta_0}^N \left( \left| \frac{1}{N} \sum_{i=1}^N \varepsilon_i g_{v,\theta_0,D}^{\text{II}}(X_i) \right| \geq \frac{c_0 D^{-\kappa_0}}{16} \right) \right] =: N\left(\frac{1}{4}\right) \cdot [\text{I} + \text{II} + \text{III} + \text{IV}].
 \end{aligned}$$

We now use empirical process techniques (Lemma B.2.1 and also Hoeffding's inequality) to bound the preceding probabilities.

**Terms I and II.** In order to apply Lemma B.2.1 to term I, we require some preparations. By the definition of  $\tilde{g}_{v,\theta}^I$  and of the operator norm  $\|\cdot\|_{\text{op}}$ , using the elementary identity  $v^\top(aa^\top - bb^\top)v = v^\top(a+b)(a-b)^\top v$  for any  $v, a, b \in \mathbb{R}^D$ , we have that for any  $\theta, \bar{\theta} \in \mathcal{B}$ ,

$$\begin{aligned}
 \|\tilde{g}_{v,\theta}^I - \tilde{g}_{v,\bar{\theta}}^I\|_\infty & \leq \left\| [\nabla \mathcal{G}(\theta) \nabla \mathcal{G}(\theta)^\top + [\mathcal{G}(\theta) - \mathcal{G}(\theta_0)] \nabla^2 \mathcal{G}(\theta)] \right. \\
 & \quad \left. - [\nabla \mathcal{G}(\bar{\theta}) \nabla \mathcal{G}(\bar{\theta})^\top + [\mathcal{G}(\bar{\theta}) - \mathcal{G}(\theta_0)] \nabla^2 \mathcal{G}(\bar{\theta})] \right\|_{L^\infty(\mathcal{X}, \mathbb{R}^{D \times D})} \\
 & \leq \left\| [\nabla \mathcal{G}(\theta) - \nabla \mathcal{G}(\bar{\theta})] [\nabla \mathcal{G}(\theta) + \nabla \mathcal{G}(\bar{\theta})]^\top \right\|_{L^\infty(\mathcal{X}, \mathbb{R}^{D \times D})} \\
 & \quad + \left\| [\mathcal{G}(\theta) - \mathcal{G}(\bar{\theta})] \nabla^2 \mathcal{G}(\theta) \right\|_{L^\infty(\mathcal{X}, \mathbb{R}^{D \times D})} \\
 & \quad + \left\| [\mathcal{G}(\bar{\theta}) - \mathcal{G}(\theta_0)] [\nabla^2 \mathcal{G}(\theta) - \nabla^2 \mathcal{G}(\bar{\theta})] \right\|_{L^\infty(\mathcal{X}, \mathbb{R}^{D \times D})} \\
 & \leq 2C_{\mathcal{G},D} |\theta - \bar{\theta}|_{E_D}, \tag{3.38}
 \end{aligned}$$

where Condition 3.2.1 furnishes the Lipschitz constant

$$C_{\mathcal{G},D} := c_2^2 D^{2\kappa_2}. \tag{3.39}$$

The hypothesis  $\mathbf{r} \leq 1$  also implies that the  $\tilde{g}_{v,\theta}^I - \tilde{g}_{v,\bar{\theta}}^I$  are uniformly bounded by  $4C_{\mathcal{G},D}$ . We then introduce the rescaled function class

$$h_\theta^I := \frac{g_{v,\theta}^I - g_{v,\theta_0,D}^I}{16C_{\mathcal{G},D}}, \quad \mathcal{H}^I = \{h_\theta^I : \theta \in \mathcal{B}\},$$

which has envelope and variance proxy bounded as

$$\sup_{\theta \in \mathcal{B}} \|h_\theta^I\|_\infty \leq \frac{1}{4} \equiv U, \quad \sup_{\theta \in \mathcal{B}} (E_{\theta_0}[h_\theta^I(X)^2])^{1/2} \leq \frac{1}{4} \equiv \sigma.$$

Next, if

$$d_2^2(\theta, \bar{\theta}) = E_{\theta_0}[(h_\theta^1(X) - h_{\bar{\theta}}^1(X))^2], \quad d_\infty(\theta, \bar{\theta}) = \|h_\theta^1 - h_{\bar{\theta}}^1\|_\infty \quad \text{for } \theta, \bar{\theta} \in \mathcal{B},$$

then using (3.38) we have that

$$d_2(\theta, \bar{\theta}) \leq d_\infty(\theta, \bar{\theta}) \leq |\theta - \bar{\theta}|_{E_D} \quad \text{for all } \theta, \bar{\theta} \in \mathcal{B}.$$

Thus for any  $\rho \in (0, 1)$ , using [60, Proposition 4.3.34], we obtain that

$$N(\mathcal{H}^1, d_2, \rho) \leq N(\mathcal{H}^1, d_\infty, \rho) \leq N(\mathcal{B}, |\cdot|_{E_D}, \rho) \leq \left(\frac{3}{\rho}\right)^D.$$

For any  $A \geq 2$  we have

$$\int_0^1 \log\left(\frac{A}{x}\right) dx = \log(A) + 1, \quad \int_0^1 \sqrt{\log\left(\frac{A}{x}\right)} dx \leq \frac{2 \log A}{2 \log A - 1} \sqrt{\log(A)}$$

(see [60, p. 190] for the latter standard inequality), and hence, using  $A = 3$ , we can respectively bound the  $L^\infty$  and  $L^2$  metric entropy integrals of  $\mathcal{H}^1$  by

$$\begin{aligned} \mathcal{J}_\infty(\mathcal{H}^1) &= \int_0^{4U} \log N(\mathcal{H}^1, d_\infty, \rho) d\rho \lesssim D, \\ \mathcal{J}_2(\mathcal{H}^1) &\leq \int_0^{4\sigma} \sqrt{\log N(\mathcal{H}^1, d_2, \rho)} d\rho \lesssim \sqrt{D}. \end{aligned}$$

Now, an application of Lemma B.2.1 implies that for any  $x \geq 1$  and some universal constant  $L' > 0$ , we have that

$$P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} \frac{1}{\sqrt{N}} \left| \sum_{i=1}^N h_\theta^1(X_i) \right| \geq L' \left[ \sqrt{D} + \sqrt{x} + \frac{D+x}{\sqrt{N}} \right] \right) \leq 2e^{-x}. \quad (3.40)$$

We also have by the definition of  $g_{v, \theta_0, D}^1$  that

$$\|g_{v, \theta_0, D}^1\|_\infty \leq 2 \|\tilde{g}_{v, \theta_0, D}^1\|_\infty \leq 2C_{\mathfrak{G}, D},$$

with  $C_{\mathfrak{G}, D}$  from (3.39), and hence, by Hoeffding's inequality [60, Theorem 3.1.2],

$$\Pi \leq 2 \exp\left(-\frac{N(c_0 D^{-\kappa_0})^2}{512C_{\mathfrak{G}, D}^2}\right). \quad (3.41)$$

Thus, from (3.39),  $D \leq \mathcal{R}_N$  and choosing  $x = 4\mathcal{R}_N$ , we have

$$L' \left[ \sqrt{D} + \sqrt{x} + \frac{D+x}{\sqrt{N}} \right] \leq \frac{c_0 D^{-\kappa_0} \sqrt{N}}{256C_{\mathfrak{G}, D}}, \quad 4\mathcal{R}_N \leq \frac{N(c_0 D^{-\kappa_0})^2}{512C_{\mathfrak{G}, D}^2},$$

whenever  $C > 0$  in (3.32) is small enough. Therefore, combining (3.40) and (3.41),

and using the definitions of the term I and of  $h_\theta^I$ , we obtain

$$\text{II} + \text{I} \leq 2e^{-4\mathcal{R}_N} + P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} \frac{1}{\sqrt{N}} \left| \sum_{i=1}^N h_\theta^I(X_i) \right| \geq \frac{c_0 D^{-\kappa_0} \sqrt{N}}{256 C_{\mathcal{G}, D}} \right) \leq 4e^{-4\mathcal{R}_N}. \quad (3.42)$$

**Terms III and IV.** Let us now treat the empirical process indexed by the functions  $\{g_{v,\theta}^{\text{II}} : \theta \in \mathcal{B}\}$ . Since  $\|v\|_{\mathbb{R}^D} \leq 1$ , we have for  $c_{\mathcal{G}, D} = \sqrt{C_{\mathcal{G}, D}} = c_2 D^{\kappa_2}$  and any  $\theta, \bar{\theta} \in \mathcal{B}$ ,

$$\|g_{v,\theta}^{\text{II}} - g_{v,\bar{\theta}}^{\text{II}}\|_\infty \leq \|\nabla^2 \mathcal{G}(\theta) - \nabla^2 \mathcal{G}(\bar{\theta})\|_{L^\infty(\mathcal{X}, \mathbb{R}^{D \times D})} \leq c_{\mathcal{G}, D} |\theta - \bar{\theta}|_{E_D} \leq c_{\mathcal{G}, D},$$

using also  $\mathbf{r} \leq 1$  to obtain the envelope bound in the last inequality. Now the rescaled function class

$$h_\theta^{\text{II}} := \frac{g_{v,\theta}^{\text{II}} - g_{v,\theta_0,D}^{\text{II}}}{4c_{\mathcal{G}, D}}, \quad \mathcal{H}^{\text{II}} = \{h_\theta^{\text{II}} : \theta \in \mathcal{B}\},$$

admits envelopes

$$\sup_{\theta \in \mathcal{B}} \|h_{v,\theta}^{\text{II}}\|_\infty \leq \frac{1}{4} \equiv U, \quad \sup_{\theta \in \mathcal{B}} (E_{\theta_0} [h_{v,\theta}^{\text{II}}(X)^2])^{1/2} \leq \frac{1}{4} \equiv \sigma,$$

and defining

$$d_2^2(\theta, \bar{\theta}) := E_{\theta_0} [(h_{v,\theta}^{\text{II}}(X) - h_{v,\bar{\theta}}^{\text{II}}(X))^2], \quad d_\infty(\theta, \bar{\theta}) = \|h_{v,\theta}^{\text{II}} - h_{v,\bar{\theta}}^{\text{II}}\|_\infty, \quad \theta, \bar{\theta} \in \mathcal{B},$$

we have

$$d_2(\theta, \bar{\theta}) \leq d_\infty(\theta, \bar{\theta}) \leq |\theta - \bar{\theta}|_{E_D} \quad \text{for all } \theta, \bar{\theta} \in \mathcal{B}.$$

Therefore, just as with the entropy integral bounds obtained for term I above, we have  $N(\mathcal{H}^{\text{II}}, d_2, \rho) \leq (3/\rho)^D$  and thus, by Lemma B.2.1,

$$P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} \frac{1}{\sqrt{N}} \left| \sum_{i=1}^N \varepsilon_i h_\theta^{\text{II}}(X_i) \right| \geq L' \left[ \sqrt{D} + \sqrt{x} + \frac{D+x}{\sqrt{N}} \right] \right) \leq 2e^{-x} \quad \text{for } x \geq 1.$$

Moreover, by the hypotheses,  $\|g_{v,\theta}^{\text{II}}\|_\infty \leq c_{\mathcal{G}, D}$ , and hence, invoking the Bernstein inequality (B.4) with  $U = \sigma \equiv c_{\mathcal{G}, D}$ , we obtain that

$$P_{\theta_0}^N \left( \left| \frac{1}{\sqrt{N}} \sum_{i=1}^N \varepsilon_i g_{v,\theta}^{\text{II}}(X_i) \right| \geq c_{\mathcal{G}, D} \left( \sqrt{2x} + \frac{x}{3\sqrt{N}} \right) \right) \leq 2e^{-x} \quad \text{for } x > 0.$$

Choosing  $x = 4\mathcal{R}_N$  in the preceding displays, we obtain that for  $C > 0$  small enough in (3.32), any  $D \leq \mathcal{R}_N$ , and by similar calculations as before (3.42),

$$\begin{aligned} \text{III} + \text{IV} &\leq P_{\theta_0}^N \left( \sup_{\theta \in \mathcal{B}} \frac{1}{\sqrt{N}} \left| \sum_{i=1}^N \varepsilon_i h_\theta^{\text{II}}(X_i) \right| \geq \frac{c_0 D^{-\kappa_0} \sqrt{N}}{96m_2} \right) \\ &\quad + P_{\theta_0}^N \left( \left| \frac{1}{\sqrt{N}} \sum_{i=1}^N \varepsilon_i g_{v,\theta}^{\text{II}}(X_i) \right| \geq \frac{c_0 D^{-\kappa_0} \sqrt{N}}{16} \right) \leq 4e^{-4\mathcal{R}_N}. \end{aligned} \quad (3.43)$$

**Combining the terms.** By combining the bounds (3.36), (3.37), (3.42), (3.43) and using that  $N(1/4) \leq 9^D \leq e^{3D}$  (cf. [60, Proposition 4.3.34]), we obtain for  $D \leq \mathcal{R}_N$  and our choice of  $\mathcal{R}_N$  with  $C$  small enough, the final bound

$$\begin{aligned} P_{\theta_0}^N \left( \inf_{\theta \in \mathcal{B}} \lambda_{\min}(-\nabla^2 \ell_N(\theta, Z)) < \frac{Nc_0 D^{-\kappa_0}}{2} \right) &\leq N \left( \frac{1}{4} \right) \cdot (\text{I} + \text{II} + \text{III} + \text{IV}) \\ &\leq 8e^{-\mathcal{R}_N}, \end{aligned}$$

completing the proof of (3.33). The bound (3.34) is proved in a similar (in fact simpler) way and left to Exercise 3.4.3.  $\blacksquare$

### 3.3 Information operators for elliptic PDEs

We now return to the model examples with forward maps (2.5), (2.8) arising from the PDE (1.2) with elliptic operator  $\mathcal{L}_{f_\theta}$  from (1.4), (1.3), respectively. For both PDEs we will derive the linearisation  $\mathbb{I}_\theta$  from Condition 3.1.1, its adjoint for natural infinite-dimensional tangent spaces  $H$ , and the resulting information operator from Definition 3.1.2. Recall that for these examples we take  $V = W = \mathbb{R}$  and  $\zeta = \lambda$  equal to Lebesgue measure on the bounded smooth domain  $\mathcal{X} = \mathcal{Z}$  in the general setting of the preceding subsection.

The main idea behind checking Condition 3.1.1 is to write

$$\mathfrak{G}(\theta_0 + h) - \mathfrak{G}(\theta_0) - D\mathfrak{G}_{\theta_0}[h], \quad h \in H,$$

as the solution of a particular *inhomogeneous* elliptic PDE for appropriate choice of  $D\mathfrak{G}_{\theta_0}$ , and to use regularity estimates for the solution operators  $\mathcal{L}_{f_\theta}^{-1}$  to verify (3.1).

#### 3.3.1 Schrödinger equation

For the forward map  $\mathfrak{G}(\theta) = u_\theta$  from (2.5) arising with the Schrödinger operator  $\mathcal{L}_{f_\theta}$  from (1.4), let us verify Condition 3.1.1 with  $\Theta = H = H^\xi(\mathcal{X})$ ,  $\xi > d/2$ , and

$$\mathbb{I}_\theta[h] \equiv D\mathfrak{G}_\theta[h] = \mathcal{L}_{f_\theta}^{-1}[u_\theta e^\theta h], \quad h \in H, \quad (3.44)$$

where  $\mathcal{L}_{f_\theta}^{-1}$  is the inverse Schrödinger operator for Dirichlet boundary conditions (given in (A.10) for  $\gamma = 1/2$ ,  $V = f_\theta$ ). We see from (A.12), (A.26), (A.4) that  $h \mapsto \mathbb{I}_\theta(h)$  is a bounded and continuous operator on  $L_\lambda^2$  for every fixed  $\theta$ . To check that  $\mathbb{I}_\theta$  is the linearisation  $\mathfrak{G}$  at a point  $\theta \in \Theta$ , consider a perturbation  $\theta + h$  in any direction  $h$  in the tangent space  $H$ . Then on  $\mathcal{X}$ ,

$$\begin{aligned} \mathcal{L}_{f_{\theta+h}}[u_{\theta+h} - u_\theta - \mathbb{I}_\theta[h]] &= 0 + (\mathcal{L}_{f_\theta} - \mathcal{L}_{f_{\theta+h}})[u_\theta + \mathbb{I}_\theta[h]] - \mathcal{L}_{f_\theta}[\mathbb{I}_\theta[h]] \\ &= (f_{\theta+h} - f_\theta)u_\theta + (f_{\theta+h} - f_\theta)\mathbb{I}_\theta[h] - u_\theta e^\theta h \\ &= (e^{\theta+h} - e^\theta - e^\theta h)u_\theta - (e^{\theta+h} - e^\theta)\mathbb{I}_\theta[h] \equiv R(h), \end{aligned}$$

while at the boundary  $\partial\mathcal{X}$  we necessarily have  $u_{\theta+h} - u_\theta - \mathbb{I}_\theta[h] = 0$ . Hence for every  $h \in H$ ,

$$\mathcal{G}(\theta + h) - \mathcal{G}(\theta) - D\mathcal{G}_\theta[h] = u_{\theta+h} - u_\theta - \mathbb{I}_\theta[h] = \mathcal{L}_{f_{\theta+h}}^{-1} [R(h)].$$

Now using the PDE estimates (A.12), (A.26) both with constant  $c$  depending only on  $\mathcal{X}$ , and standard properties of the exponential map for bounded  $\theta$ ,  $\|\theta\|_\infty \lesssim \|\theta\|_{H^\xi}$  (in view of (A.4)), we see

$$\begin{aligned} \|\mathcal{G}(\theta + h) - \mathcal{G}(\theta) - D\mathcal{G}_\theta[h]\|_{L^2} &= \|\mathcal{L}_{f_{\theta+h}}^{-1} [R(h)]\|_{L^2} \lesssim \|R(h)\|_{L^2} \\ &\lesssim \|(e^{\theta+h} - e^\theta - e^\theta h)u_\theta\|_{L^2} + \|(e^{\theta+h} - e^\theta)\mathbb{I}_\theta[h]\|_{L^2} \\ &\lesssim \|h\|_\infty^2 + \|h\|_\infty \|\mathbb{I}_\theta[h]\|_{L^2} \lesssim \|h\|_\infty^2. \end{aligned}$$

Summarising what precedes we have proved the following:

**Theorem 3.3.1.** *Let  $\Theta = H^\xi(\mathcal{X})$  for some  $\xi > d/2$ . The forward map  $\mathcal{G}$  from (2.5) satisfies Condition 3.1.1 for any  $\theta = \theta_0 \in \Theta$ , tangent space  $H = H^\xi(\mathcal{X})$ , operator  $\mathbb{I}_\theta$  from (3.44) and  $\rho_\theta(h) \leq C \|h\|_\infty^2$  with some constant  $C$  that depends only on  $\theta$  and  $\mathcal{X}$ .*

In fact, for the forward map  $\mathcal{G}$  from (2.5), it can be shown that it verifies (3.1) with weaker norms than  $\|\cdot\|_\infty$  measuring the size of the perturbation  $h$  (see, e.g., [93]), but for the present purposes the above proposition is sufficient.

If we choose as tangent space  $H = H^\xi(\mathcal{X})$  or just  $H = H_c^\xi(\mathcal{X})$ , then its completion in  $L^2(\mathcal{X})$  equals  $\bar{H} = L^2(\mathcal{X})$  itself, and the adjoint operator of  $\mathbb{I}_\theta$  then takes a particularly simple form

$$\mathbb{I}_\theta^* g = u_\theta e^\theta \mathcal{L}_{f_\theta}^{-1} [g]. \quad (3.45)$$

Indeed, using that  $\mathcal{L}_{f_\theta}^{-1}$  is self-adjoint for the  $L_\lambda^2(\mathcal{X})$ -inner product (cf. (A.10)),

$$\langle \mathbb{I}_\theta[h], g \rangle_{L^2} = \langle \mathcal{L}_{f_\theta}^{-1} [u_\theta e^\theta h], g \rangle_{L^2} = \langle h, u_\theta e^\theta \mathcal{L}_{f_\theta}^{-1} [g] \rangle_{L^2} = \langle h, \mathbb{I}_\theta^* g \rangle_{L^2}.$$

One shows further that the information operator  $\mathbb{I}_\theta^* \mathbb{I}_\theta$  arising from (2.5) as in Definition 3.1.2 is a self-adjoint and compact operator on  $L^2$  (see Exercise 3.4.4).

### 3.3.2 Diffusion equation

In the following theorem  $\mathcal{L}_{f_\theta}^{-1}$  denotes the inverse of the divergence form operator  $\mathcal{L}_{f_\theta}$  from (1.3) for Dirichlet boundary conditions – see (A.10) (with  $\gamma = f_\theta$  and  $V = 0$ ).

**Theorem 3.3.2.** *Let  $\Theta = H^\beta(\mathcal{X})$ ,  $\beta > 1 + d/2$ . The forward map  $\mathcal{G}(\theta) = u_\theta = u_{f_\theta}$  from (2.8) satisfies Condition 3.1.1 for any  $\theta = \theta_0 \in \Theta$ , tangent space  $H = H^\beta(\mathcal{X})$ , operator*

$$\mathbb{I}_\theta(h) \equiv D\mathcal{G}_\theta[h] = -\mathcal{L}_{f_\theta}^{-1} [\nabla \cdot (e^\theta h \nabla u_{f_\theta})], \quad h \in H, \quad (3.46)$$

and  $\rho_\theta(h) \leq C \|h\|_\infty^2$  for some constant  $C = C(\mathcal{X}, f_{\min}, \theta)$ .

*Proof.* Fix  $\theta \in \Theta$ , let  $h \in H$  be arbitrary, and write  $u_\theta$  for  $u_{f_\theta}$  in this proof. We first observe that

$$\begin{aligned}
& \mathcal{L}_{f_{\theta+h}}(u_{\theta+h} - u_\theta - \mathbb{I}_\theta(h)) \\
&= g - (\mathcal{L}_{f_{\theta+h}} - \mathcal{L}_{f_\theta})(u_\theta + \mathbb{I}_\theta(h)) - \mathcal{L}_{f_\theta}(u_\theta + \mathbb{I}_\theta(h)) \\
&= -(\nabla \cdot (e^{\theta+h} - e^\theta) \nabla(u_\theta + \mathbb{I}_\theta(h))) + \nabla \cdot (e^\theta h \nabla u_\theta) \\
&= -\nabla \cdot ((e^{\theta+h} - e^\theta - e^\theta h) \nabla u_\theta) - \nabla \cdot ((e^{\theta+h} - e^\theta) \nabla \mathbb{I}_\theta(h)) \\
&\equiv R_1(h) + R_2(h),
\end{aligned}$$

as well as  $u_{\theta+h} - u_\theta - \mathbb{I}_\theta(h) = 0$  on  $\partial\mathcal{X}$ . Therefore, we can write

$$\begin{aligned}
\|\mathcal{G}(\theta + h) - \mathcal{G}(\theta) - D\mathcal{G}_\theta[h]\|_{L^2} &= \|u_{\theta+h} - u_\theta - \mathbb{I}_\theta[h]\|_{L^2} \\
&= \|\mathcal{L}_{f_{\theta+h}}^{-1}[R_1(h) + R_2(h)]\|_{L^2} \\
&\leq \|\mathcal{L}_{f_{\theta+h}}^{-1}[R_1(h)]\|_{L^2} + \|\mathcal{L}_{f_{\theta+h}}^{-1}[R_2(h)]\|_{L^2}.
\end{aligned}$$

We now show that the right-hand side is  $O(\|h\|_\infty^2)$ . For the first term we have from the regularity estimate (A.32), with  $\gamma = f_\theta$ ,  $V = 0$  and Lipschitz constant  $\bar{c} = \bar{c}(\mathcal{X}, f_{\min})$ , as well as the divergence theorem (A.7),

$$\begin{aligned}
\|\mathcal{L}_{f_{\theta+h}}^{-1}[R_1(h)]\|_{L^2} &\leq \bar{c} \|\nabla \cdot ((e^{\theta+h} - e^\theta - e^\theta h) \nabla u_\theta)\|_{(H_0^1)^*} \\
&\lesssim \|(e^{\theta+h} - e^\theta - e^\theta h) \nabla u_\theta\|_{L^2} \lesssim \|h\|_\infty^2.
\end{aligned}$$

using also that  $u_\theta \in H^1$  under the maintained hypotheses (e.g., Proposition A.5.2). For the second term we have similarly, using (A.32) and (A.8) this time,

$$\begin{aligned}
\|\mathcal{L}_{f_{\theta+h}}^{-1}[R_2(h)]\|_{L^2} &\lesssim \|\nabla \cdot ((e^{\theta+h} - e^\theta) \nabla \mathbb{I}_\theta(h))\|_{(H_0^1)^*} \\
&\lesssim \|h\|_\infty \|\nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta h \nabla u_\theta)]\|_{L^2} \\
&\lesssim \|h\|_\infty \|\nabla \cdot (e^\theta h \nabla u_\theta)\|_{(H_0^1)^*} \\
&\lesssim \|h\|_\infty^2 \|u_\theta\|_{H^1} \lesssim \|h\|_\infty^2.
\end{aligned}$$

The last two inequalities also imply the continuity of  $\mathbb{I}_\theta$  as an operator on  $L^2$ , completing the proof.  $\blacksquare$

Just as with the Schrödinger equation, if we take the tangent space  $H = H^\beta$ , or just  $H_c^\beta$ , then its closure  $\bar{H}$  for the  $L^2$ -norm equals all of  $L^2$  and the adjoint in (3.3) is computed for the standard  $L^2$ -inner product.

**Proposition 3.3.3.** *In the setting of Theorem 3.3.2, the adjoint  $\mathbb{I}_\theta^*: L_\lambda^2(\mathcal{X}) \rightarrow L_\lambda^2(\mathcal{X})$  of  $\mathbb{I}_\theta: L_\lambda^2(\mathcal{X}) \rightarrow L_\lambda^2(\mathcal{X})$  is given by*

$$\mathbb{I}_\theta^*[g] = e^\theta \nabla u_\theta \cdot \nabla \mathcal{L}_{f_\theta}^{-1}[g] \quad \text{for } g \in L_\lambda^2(\mathcal{X}).$$

*Proof.* Since  $\mathbb{I}_\theta$  from (3.46) defines a bounded linear operator on the Hilbert space  $L_\lambda^2(\mathcal{X})$ , a unique adjoint operator  $I_\theta^*$  exists by the Riesz representation theorem. Let us first show that

$$\langle h, (I_\theta^* - \mathbb{I}_\theta^*)g \rangle_{L^2} = 0 \quad \text{for all } h, g \in C_0^\infty(\mathcal{X}). \quad (3.47)$$

Indeed, since  $\mathcal{L}_{f_\theta}^{-1}$  from (A.10) is self-adjoint on  $L_\lambda^2$  and maps into  $H_0^1(\mathcal{X})$ , we can apply Theorem 3.3.2 and the divergence theorem (A.7) with vector field  $e^\theta h \nabla u_\theta$  to deduce

$$\begin{aligned} \langle h, I_\theta^* g \rangle_{L^2(\mathcal{X})} &= \langle \mathbb{I}_\theta h, g \rangle_{L^2(\mathcal{X})} = -\langle \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta h \nabla u_\theta)], g \rangle_{L^2(\mathcal{X})} \\ &= -\int_{\mathcal{X}} [\nabla \cdot (e^\theta h \nabla u_\theta)] \mathcal{L}_{f_\theta}^{-1}[g] d\lambda \\ &= \int_{\mathcal{X}} h e^\theta \nabla u_\theta \cdot \nabla \mathcal{L}_{f_\theta}^{-1}[g] d\lambda = \langle h, \mathbb{I}_\theta^* g \rangle_{L^2(\mathcal{X})}, \end{aligned}$$

so that (3.47) follows. Since  $C_0^\infty(\mathcal{X})$  is dense in  $L_\lambda^2(\mathcal{X})$  and since  $I_\theta^*, \mathbb{I}_\theta^*$  are continuous on  $L_\lambda^2(\mathcal{X})$  (by construction in the former case and by (A.13) and Proposition A.5.2 in the latter case), the identity (3.47) extends to all  $g \in L_\lambda^2(\mathcal{X})$  and hence  $I_\theta^* = \mathbb{I}_\theta^*$ , as desired.  $\blacksquare$

From the previous two results one can define the information operator  $\mathbb{I}_\theta^* \mathbb{I}_\theta$  acting continuously (in fact compactly, see Exercise 3.4.4) on  $L^2$ .

### 3.3.3 Injectivity and local identifiability

As mentioned at the outset of this chapter, the operators  $\mathbb{I}_{\theta_0}, \mathbb{I}_{\theta_0}^*, \mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0}$  encode local identifiability properties of a statistical model  $\{dP_\theta : \theta \in \Theta\}$  at  $\theta_0 \in \Theta$ . In particular, Theorem 3.1.5 implies that  $\sqrt{N}$ -consistent inference on functionals  $\langle \theta, \psi \rangle_{L^2}$  at  $\theta$  is a fortiori only possible if  $\mathbb{I}_\theta$  is injective along tangent space directions  $\psi$ : If  $\psi$  is to lie in the range  $R(\mathbb{I}_\theta^*)$  of the adjoint of  $\mathbb{I}_\theta^*$ , then it cannot lie in the kernel of  $\mathbb{I}_\theta$  in view of (3.15). So it is natural to investigate the injectivity properties of  $\mathbb{I}_\theta$  first, and in doing that we can also shed some light on the basic requirement (3.27) necessary to verify the local curvature hypothesis underpinning Theorem 3.2.3.

**3.3.3.1 Schrödinger equation.** Let us start again with the example of the Schrödinger equation. We show that the operator  $\mathbb{I}_\theta$  given in Theorem 3.3.1 is injective on all of  $L^2(\mathcal{X})$ : We have from (3.44) and Proposition A.5.1 (with  $\gamma = 1/2$ ,  $V = f_\theta$ ) that

$$\|\mathbb{I}_\theta h\|_{L^2} = \|\mathcal{L}_{f_\theta}^{-1}[u_\theta e^\theta h]\|_{L^2} \gtrsim \|u_\theta e^\theta h\|_{(H_0^2)^*} \quad (3.48)$$

for any  $h \in L_\lambda^2(\mathcal{X})$ . Under the hypotheses of Theorem 3.3.1, the function  $u_\theta e^\theta$  is bounded away from zero throughout  $\mathcal{X}$  and contained in  $H^2(\mathcal{X})$  (cf. (A.25) and

Proposition A.5.2). Thus for any  $\bar{\varphi} \in C_c^\infty(\mathcal{X})$  the function  $\varphi = \bar{\varphi}/(e^\theta u_\theta)$  belongs to  $H_0^2(\mathcal{X})$  and hence (3.48) implies that  $\int \bar{\varphi} h = 0$  for any such  $\bar{\varphi}$  whenever  $\mathbb{I}_\theta h = 0$ . But this implies  $h = 0$  almost everywhere on  $\mathcal{X}$  and hence that  $\mathbb{I}_\theta$  is injective on  $L^2(\mathcal{X})$ .

What precedes implies as well that the information operator arising from  $\mathbb{I}_\theta$  is injective on  $L^2(\mathcal{X})$ . Indeed, if  $\mathbb{I}_\theta^* \mathbb{I}_\theta h = 0$ , then

$$0 = \langle h, \mathbb{I}_\theta^* \mathbb{I}_\theta h \rangle_{L^2} = \|\mathbb{I}_\theta h\|_{L^2}^2, \quad (3.49)$$

and thus  $\mathbb{I}_\theta h = 0$ . We further notice that (3.48) verifies (3.27) for  $\|\cdot\|_+ = \|\cdot\|_{(H_0^2)^*}$ , which readily gives a recipe to check the local curvature condition (3.26) if we can adapt the basis  $E_D$  to the dual norm  $(H_0^2)^*$ . This will be pursued further in Chapter 5.

**3.3.3.2 Diffusion equation.** For the divergence form equation, the injectivity of  $\mathbb{I}_\theta$  on  $L^2(\mathcal{X})$  is a more subtle question. For the operator  $\mathbb{I}_\theta$  from Theorem 3.3.2 we can use Proposition A.5.1 (with  $\gamma = f_\theta$ ,  $V = 0$ ) to obtain

$$\|\mathbb{I}_\theta h\|_{L^2} \gtrsim \|\nabla \cdot (e^\theta h \nabla u_{f_\theta})\|_{(H_0^2)^*}, \quad h \in L^2.$$

Hence if  $\mathbb{I}_\theta h = 0$ , then by testing against all smooth  $\varphi \in C_c^\infty \subset H_0^2$ , we can deduce that also  $\nabla \cdot (e^\theta h \nabla u_{f_\theta}) = 0$ . Whether this implies that  $h$  vanishes depends on the precise form of the PDE (1.2).

If we impose the natural hypothesis (2.15) from earlier ensuring injectivity of the non-linear map  $\mathcal{G}$ , then Lemma 2.1.6 does imply that  $\mathbb{I}_\theta$  (and as in (3.49) also  $\mathbb{I}_\theta^* \mathbb{I}_\theta$ ) are injective on the space  $H_0^1(\mathcal{X})$ . A slight modification of this argument combined with Lemma 2.1.6 will allow us to verify the curvature Condition (3.26) for appropriate choice of  $E_D \subset H_0^1$  as we will see in Theorem 5.3.2 below.

But on the other hand, Exercise 3.4.5 gives an example where  $\mathbb{I}_\theta$  does satisfy (2.15) but is *not* injective on *all* of  $L^2(\mathcal{X})$ . This leads to information-theoretic obstructions in view of Theorem 3.1.5 and since the kernel in (3.15) is calculated on  $L^2(\mathcal{X})$  and not on  $H_0^1(\mathcal{X})$  – see Section 4.2 for more on this.

In some situations, however,  $\mathbb{I}_\theta$  can be injective on all of  $L^2(\mathcal{X})$ , as in the following model example for the standard Laplacian  $\theta = 0$  (and  $f_{\min} = 0$  in slight abuse of notation), which will also be of interest again in Section 4.2.

**Proposition 3.3.4.** *Let  $\mathcal{X}$  equal to the unit disk in  $\mathbb{R}^2$  centred at  $(0, 0)$  and let  $\mathbb{I}_\theta$  be as in Theorem 3.3.2, where  $\mathcal{G}$  from (2.8) arises with source  $g = 2$  and boundary values  $h = (|\cdot|_{\mathbb{R}^2}^2 - 1)/2$  ( $= 0$  on  $\partial\mathcal{X}$ ). Then for  $f_{\min} = 0$ ,  $\theta = 0$ , the map  $\mathbb{I}_0: L^2(\mathcal{X}) \rightarrow L^2(\mathcal{X})$  is injective.*

*Proof.* Suppose  $\mathbb{I}_0(f) = 0$  for  $f \in L^2(\mathcal{X})$ . Then for any  $g \in C^\infty(\mathcal{X})$ , we have, by Proposition 3.3.3,

$$0 = \langle \mathbb{I}_0 f, g \rangle_{L^2(\mathcal{X})} = \langle f, \mathbb{I}_0^* g \rangle_{L^2(\mathcal{X})} = \langle f, X \mathcal{L}_1^{-1}[g] \rangle_{L^2(\mathcal{X})},$$

with vector field

$$X = \nabla u_1 \cdot \nabla(\cdot) = x_1 \partial x_1 + x_2 \partial x_2, \quad (x_1, x_2) \in \mathcal{X}.$$

Choosing  $g = \Delta \bar{g}$  for any smooth  $\bar{g}$  of compact support, we deduce that

$$\int_{\mathcal{X}} X(\bar{g}) f d\lambda = 0 \quad \text{for all } \bar{g} \in C_c^\infty(\mathcal{X}), \quad (3.50)$$

and we now show that this implies  $f = 0$ . A somewhat informal dynamical argument would say that (3.50) asserts that  $f d\lambda$  is an invariant density under the flow of  $X$ . Since the flow of  $X$  in backward time has a sink at the origin, the density can only be supported at  $(x_1, x_2) = 0$  and thus  $f = 0$ .

One can give a distributional argument as follows. Suppose we consider polar coordinates  $(r, \vartheta) \in (0, 1) \times S^1$  and functions  $\bar{g}$  of the form  $\phi(r)\psi(\vartheta)$ , where  $\phi \in C_0^\infty(0, 1)$  and  $\psi \in C^\infty(S^1)$ . In polar coordinates,  $X = r \partial_r$  and hence we may write (3.50) as

$$\int_0^1 \left( r^2 \left( \int_0^{2\pi} f(r, \vartheta) \psi(\vartheta) d\vartheta \right) \partial_r \phi \right) dr = 0. \quad (3.51)$$

By Fubini's theorem, for each  $\psi$  we have an integrable function

$$F_\psi(r) := \int_0^{2\pi} f(r, \vartheta) \psi(\vartheta) d\vartheta,$$

and thus  $r^2 F_\psi$  defines an integrable function on  $(0, 1)$  whose distributional derivative satisfies  $\partial_r(r^2 F_\psi) = 0$  by virtue of (3.51). Thus  $r^2 F_\psi = c_\psi$  (using that a distribution on  $(0, 1)$  with zero derivative must be a constant). Now consider  $\psi \in C^\infty(S^1)$  also as a function in  $L^2(\mathcal{X})$  and compute the pairing

$$(f, \psi)_{L^2(\mathcal{X})} = \int_0^1 r F_\psi(r) dr = c_\psi \int_0^1 r^{-1} dr = \pm\infty,$$

unless  $c_\psi = 0$ . Thus  $f = 0$ . ■

By perturbation and the Morse lemma, one can show further that the preceding injectivity result extends at least to  $\mathbb{I}_\theta$  for an appropriate neighbourhood of  $\theta$ 's near the standard Laplacian  $\theta = 0$ , see Exercise 3.4.6.

## 3.4 Notes

### 3.4.1 Exercises

**Exercise 3.4.1.** Show that Theorem 3.1.4 remains valid if  $\mathbb{I}_\theta$  is not injective by repeating its proof on the quotient space of  $\tilde{H}$  for the kernel of  $\mathbb{I}_\theta$ .

**Exercise 3.4.2.** Show that Theorem 3.1.4 and (3.19) remain valid in the augmented model from the proof of Theorem 3.1.5. [Hint: Proceed as in the proof of Theorem 3.1.4 but apply [129, Theorem 3.11.5] with the LAN expansion of the augmented model.]

**Exercise 3.4.3.** Provide a proof of (3.34). [Follow the proof of (3.33) in Lemma 3.2.3, or see [101, Lemma 3.4].]

**Exercise 3.4.4.** Show that the information operators  $\mathbb{I}_\theta^* \mathbb{I}_\theta$  from Definition 3.1.2 arising from  $\mathbb{I}_\theta$  given in Theorems 3.3.1 and 3.3.2, respectively, are *compact* operators on  $L_\lambda^2(\mathcal{X})$ .

**Exercise 3.4.5.** Consider  $\mathcal{G}, h, g$  as in Proposition 3.3.4 but now with  $\mathcal{X}$  a smooth bounded domain in  $\mathbb{R}^2$  that is *separated away from the origin*. Show that  $\nabla u_1$  does not vanish on  $\mathcal{X}$  and hence satisfies (2.15), but that  $\mathbb{I}_0$  is *not* injective on  $L_\lambda^2(\mathcal{X})$ . [Hint: See [95, Section 3.4.1].]

**Exercise 3.4.6.** Show that Proposition 3.3.4 extends to all  $\theta$  such that  $\|\theta\|_{H^\beta} < \eta$  for  $\beta > 1 + d/2$  and  $\eta > 0$  small enough. [Hint: Show first by perturbation that  $\|u_{f_\theta} - u_1\|_{C^1} \rightarrow 0$  as  $\eta \rightarrow 0$  and that hence  $\nabla u_\theta$  can only vanish at a single point for  $\eta$  sufficiently small. Then repeat the proof of Proposition 3.3.4.]

### 3.4.2 Remarks and comments

The results from Section 3.1 constitute by now classical material from semi-parametric statistics, we refer to [125, 126] where these ideas are laid out in general likelihood models of which (1.9) is just a special case. The preceding references develop a formal ‘score operator calculus’ in general statistical models that are ‘differentiable in quadratic mean (DQM)’, which for our (1.9) leads initially to slightly different expressions for  $\mathbb{I}_\theta$  involving a further ‘projection’ step onto the regression residuals  $Y - \mathcal{G}(\theta)(X) = \varepsilon$ . This is easily reconciled with the form of  $\mathbb{I}_\theta$  given here as discussed in [95]. Since the LAN property in Theorem 3.1.3 with LAN norm  $\|\mathbb{I}_{\theta_0}(\cdot)\|_{L^2}$  can be proved directly in our Gaussian regression model, we can avoid introducing the DQM property altogether. This expedites the proofs of Theorems 3.1.4 and 3.1.5, but we again refer to [95] for a discussion of how the model (1.9) does fit into the abstract framework from [125, 126].

The connection between local average curvature of the log-likelihood function and the information matrix is well known in ‘parametric’ finite-dimensional statistical models – the idea to exploit stability estimates (3.27) in conjunction with high-dimensional concentration of measure arguments in Theorem 3.2.3 has been developed in this form in [101], see also [23].

Obtaining the linearisation of the non-linear forward maps (2.5) and (2.8) via arguments such as those given in Section 3.3 is a basic application of perturbative

methods for inhomogeneous linear PDEs. The current proofs are taken from [93], [91], [95] where these problems were apparently first investigated in the setting of LAN expansions and information operators, with no claim of priority. Proposition 3.3.4 is taken from [95]. The techniques for the Schrödinger equation again extend to more complex inverse problems where a base differential operator  $\mathcal{D}$  is perturbed by some potential, see, e.g., [91], [23] for the case where  $\mathcal{D}$  is the geodesic vector field on the unit disk related to (1.1).



## Chapter 4

### Bernstein–von Mises theorems

A classical result of mathematical statistics due to Laplace, Bernstein, von Mises, Le Cam and van der Vaart gives general conditions under which posterior distributions  $\Pi(\cdot|D_N)$  on parameter spaces  $\Theta$  of fixed finite dimension are asymptotically approximated by a normal distribution centred at an efficient estimator (say the maximum likelihood estimator, or the posterior mean) and with ‘inverse information covariance matrix’, see [82, 126]. This phenomenon occurs for general priors and likelihood models and can hence be understood as an expression of a ‘universality’ principle that posterior measures necessarily resemble the shape of Gaussians at least for large sample sizes. While this can be observed empirically in real data, the mathematical results substantiating it are called Bernstein–von Mises (BvM) theorems (to be distinguished from ‘Laplace approximations’ discussed in the next chapter). A main appeal of rigorous such theorems is that they provide objective statistical guarantees for posterior credible sets and hence for the Bayesian approach to *uncertainty quantification* (UQ) and the construction of ‘error bars’ for algorithmic outputs – see Section 4.1.3 for more on this.

While the BvM theorem is well understood in ‘parametric’ finite-dimensional models, obtaining high- or infinite-dimensional versions of it poses fundamental challenges. Already in basic conjugate Gaussian sequence space models examples for the invalidity of a Bernstein–von Mises approximation can be given, see [38, 53, 83] and also [71] for high-dimensional regression. However, following ideas from Le Cam theory and semi-parametric statistics, [28, 29, 31] developed techniques that allow to obtain BvM type results also in infinite-dimensional models, at least in ‘weak enough topologies’. In this chapter we lay out some main ideas of this approach in the Gaussian regression model (1.9) and in the context of PDE inverse problems. Roughly speaking the goal is to describe conditions such that the ‘cylindrical’ laws of posterior measures characterised by all one-dimensional statistics  $\langle \theta, \psi \rangle_{L^2}$  for a convergence-determining class of smooth test functions  $\psi$  are approximated by an infinite-dimensional Gaussian measure with a covariance structure that is information-theoretically optimal (in the sense of Theorem 3.1.4). (This convergence can be upgraded to be ‘uniform’ in  $\psi$  at the expense of technicalities that we wish to avoid here, but see after the proof of Theorem 4.2.1 for discussion.) The new analytical issue that arises is the inversion of the information operator appearing in the information equation (3.16). We will show that a semi-parametric BvM theorem can be proved when sufficiently regular solutions to the information equation exist. This in turn leads to new PDE questions of its own and we spell out what can happen for our two model examples with the Schrödinger and divergence form equations (2.5) and (2.8).

## 4.1 Gaussian asymptotics for cylindrical laws

In this section we will show that if we can find sufficiently regular solutions to the ‘information equation’ from (3.16) for any given test function  $\psi \in \Theta$ , then as  $N \rightarrow \infty$ , a Gaussian approximation for the one-dimensional posterior statistics  $\langle \theta, \psi \rangle_{L^2}$ ,  $\theta \sim \Pi(\cdot | (Y_i, X_i)_{i=1}^N)$ , holds true under some natural additional quantitative assumptions on the forward map  $\mathcal{G}$  and the Gaussian process prior  $\Pi$ .

Colloquially, one says that ‘random via the data’ probability measures  $\mu_N \equiv \mu_N | (Y_i, X_i)_{i=1}^N$  converge weakly in probability to a limiting normal distribution  $\mu$ . To make mathematical sense of this, one takes a metric  $d_{\text{weak}}$  for weak convergence of probability measures on the underlying space (see [46]) and shows the convergence to zero in  $P_{\theta_0}^N$ -probability of the real random variable  $d_{\text{weak}}(\mu_N, \mu)$ . For random variables  $Z_N \sim \mu_N$ ,  $Z \sim \mu$ , we equivalently say that ‘ $Z_N \xrightarrow{d} Z$  converges in distribution in  $P_{\theta_0}^N$ -probability’ if

$$d_{\text{weak}}(\mu_N, \mu) \xrightarrow{P_{\theta_0}^N} 0 \quad \text{as } N \rightarrow \infty, \quad (4.1)$$

in the remainder of this section.

### 4.1.1 Asymptotic normality of linear functionals of the posterior

We consider again the general Gaussian regression model (1.9) with linear parameter space  $\Theta \subset L^2_{\xi}(\mathcal{Z}, W)$ , where  $\mathcal{Z}$  is a bounded domain in  $\mathbb{R}^d$  with smooth boundary. We enforce on the forward map  $\mathcal{G}: \Theta \rightarrow L^2_{\lambda}(\mathcal{X}, V)$  the Conditions 2.1.1 and 2.1.4 from earlier. We also suppose the linearisation Condition 3.1.1 holds for a linear operator  $\mathbb{I}_{\theta_0}$  and with tangent space  $H = \Theta$ .

To simplify some proofs we slightly strengthen the forward regularity hypotheses to a  $L^{\infty}$ -Lipschitz property.

**Condition 4.1.1.** For all  $M$  and  $\theta_1, \theta_2 \in \Theta$  such that  $\|\theta_i\|_{\mathcal{R}} \leq M$ , there exists a constant  $L$  such that

$$\|\mathcal{G}(\theta_1) - \mathcal{G}(\theta_2)\|_{\infty} \leq L \|\theta_1 - \theta_2\|_{\infty}.$$

Moreover, the linear operator  $\mathbb{I}_{\theta_0}$  is continuous from  $(\Theta, \|\cdot\|_{\infty})$  to  $L^{\infty}(\mathcal{X}, V)$ .

This condition can be verified along similar lines as in Propositions 2.1.2 and 2.1.3, involving a few extra technicalities. For  $\mathcal{G}$  arising with the Schrödinger equation, see Exercise 4.3.1.

When considering inference on linear functionals  $\langle \psi, \theta \rangle_{L^2(\mathcal{Z})}$  of  $\theta$ , we require the existence of sufficiently regular solutions of the information equation (3.16) arising from the ‘information’ operator  $\mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0}$  from Definition 3.1.2 with  $\Theta = H$ . Formally we assume the following:

**Condition 4.1.2.** Given  $\psi \in \Theta$ , suppose that there exists  $\bar{\psi} = \bar{\psi}_{\theta_0} \in \Theta$  such that  $\mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0} \bar{\psi} = \psi$ , that is,  $\langle \mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0} \bar{\psi} - \psi, h \rangle_{L_\xi^2(\mathcal{Z}, \mathcal{W})} = 0$  for all  $h \in \Theta$ .

As prior  $\Pi = \Pi_N$  we take a (rescaled) Gaussian probability measure on  $\Theta$  from Theorem 2.2.2. Next for the ‘regularisation sets’  $\Theta_N$  from that theorem, the proofs that follow require an initial posterior contraction rate  $\bar{\delta}_N$  in  $\|\cdot\|_\infty$ -norm rather than just the  $L_\xi^2$ -norm. If the regularisation space  $\mathcal{R}$  satisfies  $\mathcal{R} \subseteq H^\beta$  for some  $\beta > d/2$ , then such a (not necessarily sharp) contraction rate can be obtained by interpolation and the Sobolev imbedding.

**Proposition 4.1.3.** *Suppose  $\mathcal{G}, \Pi', \Pi_N, \mathcal{R}, \theta_0, \delta_N$  are as in Theorem 2.2.2 with regularisation space  $\mathcal{R}$  continuously embedded into  $H^\beta(\mathcal{Z})$  for some  $\beta > d/2$ . Suppose further that Condition 2.1.4 holds for such  $\mathcal{R}$  and some  $\eta > 0$ . Define*

$$\Theta_{N,M,\infty} = \{\theta \in \Theta : \|\theta\|_{\mathcal{R}} \leq M, \|\theta - \theta_0\|_\infty \leq M\bar{\delta}_N\}, \quad (4.2)$$

where  $M$  is a fixed constant and

$$\bar{\delta}_N = \delta_N^{\eta(\beta-\beta')/\beta}, \quad \text{any } \beta' \in (d/2, \beta). \quad (4.3)$$

Then for all  $b > 0$ , we can choose  $M$  large enough such that as  $N \rightarrow \infty$ ,

$$P_{\theta_0}^N(\Pi_N(\theta \in \Theta_{N,M,\infty} | D_N) \leq 1 - e^{-bN\delta_N^2}) \rightarrow 0.$$

*Proof.* By the Sobolev imbedding (A.4) and interpolation (A.5) we have

$$\|\theta - \theta_0\|_\infty \lesssim \|\theta - \theta_0\|_{H^\beta} \lesssim \|\theta - \theta_0\|_{L^2}^\vartheta \|\theta - \theta_0\|_{H^\beta}^{1-\vartheta},$$

where  $\vartheta = (\beta - \beta')/\beta$ . By hypothesis on  $\mathcal{R}$ , the norms  $\|\theta - \theta_0\|_{H^\beta}$  are bounded by  $M + \|\theta_0\|_{\mathcal{R}} < \infty$ , so the limit (2.26) from Theorem 2.3.1 implies the desired result.  $\blacksquare$

The regularisation set  $\Theta_{N,M,\infty}$  will play a quantitative role in the proofs to follow. It permits control of the non-linearity of the log-likelihood function (1.13), the discretisation errors arising from statistical sampling, and – as will be relevant – the sensitivity of  $\Pi_N$  with respect to small perturbations in  $\bar{\psi}$ -directions. To this end, let  $\mathcal{J}_N$  be an upper bound for the following (‘Dudley’-type) integral of the Kolmogorov metric entropy of  $\Theta_{N,M,\infty}$ :

$$\mathcal{J}_N(s, t) \geq \int_0^s \sqrt{\log 2N(\Theta_{N,M,\infty}, \|\cdot\|_\infty, t\epsilon)} d\epsilon, \quad s, t > 0, \quad (4.4)$$

where  $N(\Theta_{N,M,\infty}, \|\cdot\|_\infty, \epsilon)$  denotes the minimal number of  $\|\cdot\|_\infty$ -norm balls of radius  $\epsilon$  required to cover the set  $\Theta_{N,M,\infty}$ .

**Condition 4.1.4.** Assume  $\sqrt{N}\delta_N^2 \rightarrow 0$  and that for  $\bar{\delta}_N$  as in (4.3),

$$\sqrt{N}\bar{\delta}_N^2 \mathcal{J}_N(1, \bar{\delta}_N^2) \rightarrow 0 \quad \text{as } N \rightarrow \infty. \quad (4.5)$$

Further suppose that  $\bar{\psi}_{\theta_0}$  from Condition 4.1.2 belongs to  $\mathcal{H} \cap \mathcal{R}$  and let  $\rho_{\theta_0}$  be the error term from Condition 3.1.1. For  $\sigma_N$  a sequence such that for all  $N$  large enough and all  $t \in \mathbb{R}$  fixed,

$$\sigma_N \geq \sup_{\theta \in \Theta_{N,M,\infty}} \rho_{\theta_0} \left[ \theta - \theta_0 - \frac{t}{\sqrt{N}} \bar{\psi}_{\theta_0} \right],$$

assume that as  $N \rightarrow \infty$ ,

$$\max \left( N(\sigma_N^2 + \sigma_N \bar{\delta}_N), \sqrt{N} \mathcal{J}_N(\sigma_N, 1), \frac{\bar{\delta}_N \sqrt{\log N} \mathcal{J}_N^2(\sigma_N, 1)}{\sigma_N^2} \right) \rightarrow 0. \quad (4.6)$$

When the approximation in Condition 3.1.1 is quadratic ( $\rho_{\theta_0}(h) = O(\|h\|_\infty^2)$ ), then (4.5) and (4.6) reduce to the much simpler conditions  $N \bar{\delta}_N^3 \rightarrow 0$ ,  $\beta > 2d$ , in light of  $\mathcal{R} \subset H^\beta(\mathcal{Z})$ , see the proof of Theorem 4.2.1 below.

Our first theorem concerns an asymptotic approximation of the induced laws of the posterior distribution of  $\langle \psi, \theta \rangle_{L^2}$  by a Gaussian distribution centred at

$$\hat{\Psi}_N = \langle \psi, \theta_0 \rangle_{L^2(\mathcal{Z})} + \frac{1}{N} \sum_{i=1}^N \langle \mathbb{I}_{\theta_0} \bar{\psi}_{\theta_0}(X_i), \varepsilon_i \rangle_V. \quad (4.7)$$

We recall that the next limit is to be understood in the sense of (4.1).

**Theorem 4.1.5.** *Suppose  $\Pi_N, \theta_0$  are as in Proposition 4.1.3 and that the forward map  $\mathcal{G}$  satisfies Conditions 2.1.1, 2.1.4, 3.1.1 with  $H = \Theta$  and 4.1.1. Assume further that  $\psi \in \Theta$  is such that Condition 4.1.2 holds for some  $\bar{\psi}_{\theta_0}$  for which Condition 4.1.4 can be verified. Let  $\theta \sim \Pi(\cdot | (Y_i, X_i)_{i=1}^N)$  be a posterior draw. Then we have, as  $N \rightarrow \infty$  and in  $P_{\theta_0}^N$ -probability,*

$$\sqrt{N} (\langle \theta, \psi \rangle_{L^2(\mathcal{Z})} - \hat{\Psi}_N) | (Y_i, X_i)_{i=1}^N \xrightarrow{d} N(0, \|\mathbb{I}_{\theta_0} \bar{\psi}_{\theta_0}\|_{L^2_\lambda(\mathcal{X}, V)}^2).$$

*Proof.* The plan is to prove convergence of the moment generating functions (Laplace transforms) of  $\sqrt{N} (\langle \theta, \psi \rangle_{L^2(\mathcal{Z})} - \hat{\Psi}_N) | (Y_i, X_i)_{i=1}^N$  which implies weak convergence by standard arguments. We recall the notation  $D_N := (Y_i, X_i)_{i=1}^N$  for the data vector. We will also write shorthand  $\bar{\psi} = \bar{\psi}_{\theta_0}$  in this proof. The proof is split into four separate steps.

**4.1.1.1 Localisation of the posterior measure.** By Condition 4.1.4 the function  $\bar{\psi}$  from Condition 4.1.2 defines an element of the RKHS  $\mathcal{H}_N = \sqrt{N} \delta_N \mathcal{H}$  of  $\Pi_N$  with RKHS norm

$$\|\cdot\|_{\mathcal{H}_N} = \sqrt{N} \delta_N \|\cdot\|_{\mathcal{H}}.$$

If  $\theta \sim \Pi_N$ , then by the definition of the RKHS, the random variable  $\langle \theta, \bar{\psi} \rangle_{\mathcal{H}_N}$  has a  $N(0, \|\bar{\psi}\|_{\mathcal{H}_N}^2)$  distribution. Hence if we define

$$T_N = \left\{ \theta \in \Theta : \frac{|\langle \theta, \bar{\psi} \rangle_{\mathcal{H}_N}|}{\|\bar{\psi}\|_{\mathcal{H}_N}} > \sqrt{B} \sqrt{N} \delta_N \right\} \quad \text{for } B > 0,$$

then the tail inequality for standard normal random variables implies that  $\Pi(T_N) \leq e^{-BN\delta_N^2}$ . So, by Theorem 2.2.2 (and the remark at the beginning of its proof), and for  $\Theta_{N,M,\infty}$  from Proposition 4.1.3, if we set

$$\bar{\Theta}_N := \Theta_{N,M,\infty} \cap T_N^c, \quad \text{then} \quad \Pi(\bar{\Theta}_N^c | D_N) = O_{P_{\theta_0}^N}(e^{-bN\delta_N^2}) = o_{P_{\theta_0}^N}(1) \quad (4.8)$$

for any  $b > 0$  and as  $N \rightarrow \infty$ , as long as we choose  $M, B$  large enough. In the proofs that follow we consider  $\theta \sim \Pi^{\bar{\Theta}_N}(\cdot | D_N)$ , where the posterior (1.12) is taken to arise from prior probability measure

$$\Pi^{\bar{\Theta}_N} \equiv \frac{\Pi(\cdot \cap \bar{\Theta}_N)}{\Pi(\bar{\Theta}_N)}$$

equal to  $\Pi$  restricted to  $\bar{\Theta}_N$  from (4.2) and renormalised. Standard arguments (Exercise 4.3.2) then imply, for  $\|\cdot\|_{TV}$  the total variation distance on probability measures on  $\Theta$ , that as  $N \rightarrow \infty$ ,

$$\|\Pi(\cdot | D_N) - \Pi^{\bar{\Theta}_N}(\cdot | D_N)\|_{TV} \leq 2\Pi(\bar{\Theta}_N^c | D_N) \xrightarrow{P_{\theta_0}^N} 0, \quad (4.9)$$

and as a consequence for any metric  $d_{\text{weak}}$  for weak convergence also

$$d_{\text{weak}}(\Pi(\cdot | D_N), \Pi^{\bar{\Theta}_N}(\cdot | D_N)) \xrightarrow{P_{\theta_0}^2} 0.$$

It hence suffices to prove Theorem 4.1.5 for  $\Pi^{\bar{\Theta}_N}(\cdot | D_N)$  instead of  $\Pi(\cdot | D_N)$ .

#### 4.1.1.2 Perturbation of the posterior Laplace transform.

**Proposition 4.1.6.** *For  $\theta, \psi \in \Theta$  and  $\bar{\psi} = \bar{\psi}_{\theta_0}$  from Condition 4.1.2, define*

$$\theta_{(t)} = \theta - t \frac{\bar{\psi}_{\theta_0}}{\sqrt{N}} \quad \text{for } t \in \mathbb{R}.$$

Let  $\hat{\Psi}_N$  be as in (4.7),  $\bar{\Theta}_N$  as in (4.8), and  $\ell_N$  as in (1.13). Then we have for every fixed  $t \in \mathbb{R}$  and a sequence  $R_N = o_{P_{\theta_0}^N}(1)$  as  $N \rightarrow \infty$ ,

$$\begin{aligned} & E^{\Pi^{\bar{\Theta}_N}} \left[ \exp\{t\sqrt{N}(\langle \theta, \psi \rangle_{L^2(\mathbb{Z})} - \hat{\Psi}_N)\} | D_N \right] \\ &= e^{\frac{t^2}{2} \|\mathbb{I}_{\theta_0} \bar{\psi}\|_{L^2(\mathbb{X})}^2} \times \frac{\int_{\bar{\Theta}_N} e^{\ell_N(\theta_{(t)})} d\Pi(\theta)}{\int_{\bar{\Theta}_N} e^{\ell_N(\theta)} d\Pi(\theta)} \times e^{R_N}. \end{aligned}$$

*Proof.* For  $W_N$  as in (3.8) with  $h = \bar{\psi}$ , the posterior Laplace transform equals

$$\begin{aligned} & E^{\Pi^{\bar{\Theta}_N}} \left[ e^{t\sqrt{N}(\langle \theta, \psi \rangle_{L^2(\mathbb{Z})} - \hat{\Psi}_N)} | D_N \right] \\ &= \frac{\int_{\bar{\Theta}_N} e^{t\sqrt{N}(\langle \theta - \theta_0, \psi \rangle_{L^2(\mathbb{Z})} - tW_N + \ell_N(\theta) - \ell_N(\theta_{(t)}) + \ell_N(\theta_{(t)})} d\Pi(\theta)}{\int_{\bar{\Theta}_N} e^{\ell_N(\theta)} d\Pi(\theta)} \end{aligned}$$

The main step in the proof is a uniform in  $\theta \in \bar{\Theta}_N$  perturbation expansion of the log-likelihood ratios under  $P_{\theta_0}^N$ . We can write

$$\begin{aligned}
& \ell_N(\theta) - \ell_N(\theta_{(t)}) \\
&= -\frac{1}{2} \sum_{i=1}^N (|Y_i - \mathfrak{G}(\theta)(X_i)|_V^2 - |Y_i - \mathfrak{G}(\theta_{(t)})(X_i)|_V^2) \\
&= -\frac{1}{2} \sum_{i=1}^N (|\mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}(\theta)(X_i) + \varepsilon_i|_V^2 - |\mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}(\theta_{(t)})(X_i) + \varepsilon_i|_V^2) \\
&= -\sum_{i=1}^N (\langle \varepsilon_i, \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}(\theta)(X_i) \rangle_V - \langle \varepsilon_i, \mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}(\theta_{(t)})(X_i) \rangle_V) \\
&\quad - \frac{1}{2} \sum_{i=1}^N (|\mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}(\theta)(X_i)|_V^2 - |\mathfrak{G}(\theta_0)(X_i) - \mathfrak{G}(\theta_{(t)})(X_i)|_V^2) \equiv \text{I} + \text{II}.
\end{aligned}$$

About term I, we ‘linearise’ the map  $\mathfrak{G}$  at  $\theta_0$  in each inner product to obtain

$$\begin{aligned}
\text{I} &= \sum_{i=1}^N \langle \varepsilon_i, D\mathfrak{G}_{\theta_0}(X_i)[\theta - \theta_{(t)}] \rangle_V \\
&\quad + \sum_{i=1}^N \langle \varepsilon_i, \mathfrak{G}(\theta)(X_i) - \mathfrak{G}(\theta_0)(X_i) - D\mathfrak{G}_{\theta_0}(X_i)[\theta - \theta_0] \rangle_V \\
&\quad - \sum_{i=1}^N \langle \varepsilon_i, \mathfrak{G}(\theta_{(t)})(X_i) - \mathfrak{G}(\theta_0)(X_i) - D\mathfrak{G}_{\theta_0}(X_i)[\theta_{(t)} - \theta_0] \rangle_V,
\end{aligned}$$

which we write as

$$\frac{t}{\sqrt{N}} \sum_{i=1}^N \langle \varepsilon_i, D\mathfrak{G}_{\theta_0}(X_i)[\bar{\psi}] \rangle_V + R_{(0)}(\theta) - R_{(t)}(\theta) = tW_N + R_{(0)}(\theta) - R_{(t)}(\theta),$$

noting that  $\theta_{(0)} = \theta$  and where the ‘remainder empirical processes’ are given by

$$R_{(t)} \equiv \sum_{i=1}^N \langle \varepsilon_i, \mathfrak{G}(\theta_{(t)})(X_i) - \mathfrak{G}(\theta_0)(X_i) - D\mathfrak{G}_{\theta_0}(X_i)[\theta_{(t)} - \theta_0] \rangle_V.$$

We show in Lemma 4.1.7 below that for all  $t \in \mathbb{R}$  fixed,

$$\sup_{\theta \in \Theta_{N,M,\infty}} |R_{(t)}(\theta)| = o_{P_{\theta_0}^N}(1), \tag{4.10}$$

so that these terms form a part of the sequence  $R_N$ .

For term II, with  $E_\lambda$  denoting expectation under the  $X_i$ 's only, we have

$$\begin{aligned} & -\frac{1}{2} \sum_{i=1}^N (|\mathcal{G}(\theta_0)(X_i) - \mathcal{G}(\theta)(X_i)|_V^2 - E_\lambda |\mathcal{G}(\theta_0)(X_i) - \mathcal{G}(\theta)(X_i)|_V^2 \\ & - |\mathcal{G}(\theta_0)(X_i) - \mathcal{G}(\theta_t)(X_i)|_V^2 + E_\lambda |\mathcal{G}(\theta_0)(X_i) - \mathcal{G}(\theta_t)(X_i)|_V^2) \\ & - \frac{N}{2} \|\mathcal{G}(\theta_0) - \mathcal{G}(\theta)\|_{L^2(\mathcal{X})}^2 + \frac{N}{2} \|\mathcal{G}(\theta_0) - \mathcal{G}(\theta_t)\|_{L^2(\mathcal{X})}^2. \end{aligned}$$

The sums in the first two lines are empirical processes and are shown in Lemma 4.1.8 below to be  $o_{P_{\theta_0}^N}(1)$  uniformly in  $\theta \in \Theta_{N,M,\infty}$  for every fixed  $t$ , and can thus also be absorbed into  $R_N$ .

For the terms in the last line of the last display, we can further decompose

$$\begin{aligned} & \|\mathcal{G}(\theta_0) - \mathcal{G}(\theta_t)\|_{L^2(\mathcal{X})}^2 \\ & = \|\mathcal{G}(\theta_t) - \mathcal{G}(\theta_0) - D\mathcal{G}_{\theta_0}[\theta_t - \theta_0] + D\mathcal{G}_{\theta_0}[\theta_t - \theta_0]\|_{L^2(\mathcal{X})}^2 \\ & = \|D\mathcal{G}_{\theta_0}[\theta_t - \theta_0]\|_{L^2(\mathcal{X})}^2 \\ & \quad + 2\langle D\mathcal{G}_{\theta_0}[\theta_t - \theta_0], \mathcal{G}(\theta_t) - \mathcal{G}(\theta_0) - D\mathcal{G}_{\theta_0}[\theta_t - \theta_0] \rangle_{L^2(\mathcal{X})} \\ & \quad + \|\mathcal{G}(\theta_t) - \mathcal{G}(\theta_0) - D\mathcal{G}_{\theta_0}[\theta_t - \theta_0]\|_{L^2(\mathcal{X})}^2, \end{aligned}$$

including also the case  $\theta = \theta_0$  by convention for  $t = 0$ . Now using Conditions 3.1.1 and 4.1.4 and the Cauchy–Schwarz inequality, the last two remainder terms are bounded by a constant multiple of

$$\sup_{\theta \in \Theta_{N,M,\infty}} [\rho_{\theta_0}^2(\theta_t - \theta_0) + \|\theta_t - \theta_0\|_{L^2} \rho_{\theta_0}(\theta_t - \theta_0)] \lesssim \sigma_N^2 + \sigma_N \bar{\delta}_N = o\left(\frac{1}{N}\right).$$

The remaining terms in the expansion are

$$\begin{aligned} & \frac{N}{2} \left( \left\| D\mathcal{G}_{\theta_0} \left[ \theta - \theta_0 - \frac{t}{\sqrt{N}} \bar{\psi} \right] \right\|_{L^2(\mathcal{X},V)}^2 - \|D\mathcal{G}_{\theta_0}[\theta - \theta_0]\|_{L^2(\mathcal{X},V)}^2 \right) \\ & = -t\sqrt{N} \langle D\mathcal{G}_{\theta_0}[\theta - \theta_0], D\mathcal{G}_{\theta_0}[\bar{\psi}] \rangle_{L^2(\mathcal{X},V)} + \frac{t^2}{2} \|D\mathcal{G}_{\theta_0}[\bar{\psi}]\|_{L^2(\mathcal{X},V)}^2 \\ & = -t\sqrt{N} \langle \theta - \theta_0, \mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0} \bar{\psi} \rangle_{L^2(\mathcal{Z},W)} + \frac{t^2}{2} \|\mathbb{I}_{\theta_0}[\bar{\psi}]\|_{L^2(\mathcal{X},V)}^2, \end{aligned}$$

which, combined with Condition 4.1.2, the bounds from term I and the identity in the first display in this proof, implies the result.  $\blacksquare$

**4.1.1.3 Stochastic bounds on remainder terms and discretisation error.** The following two lemmas use tools from infinite-dimensional probability to bound the collections of empirical processes appearing as remainder terms in the proof of Proposition 4.1.6. While that proposition considers localisation to the sets  $\bar{\Theta}_N$ , the following bounds actually hold uniformly in the larger classes  $\Theta_{N,M,\infty}$  from (4.2).

**Lemma 4.1.7.** *We have (4.10).*

*Proof.* For  $t$  fixed, define new functions  $g_\theta: \mathcal{X} \rightarrow V$  as

$$g_\theta = \mathcal{G}(\theta_{(t)})(\cdot) - \mathcal{G}(\theta_0)(\cdot) - D\mathcal{G}_{\theta_0}(\cdot)[\theta_{(t)} - \theta_0].$$

Then the remainder term from (4.10), viewed as a stochastic process indexed by  $\theta \in \Theta_{N,M,\infty}$ , equals a centred (since  $E\varepsilon_i = 0$ ) empirical process for the jointly i.i.d. variables  $(X_i, \varepsilon_i)$  of the form

$$|R_{(t)}(\theta)| \equiv \left| \sum_{i=1}^N \langle \varepsilon_i, g_\theta(X_i) \rangle_V \right| \equiv \left| \sum_{j=1}^{p_V} \sum_{i=1}^N \varepsilon_{i,j} g_{\theta,j}(X_i) \right| \leq \sum_{j=1}^{p_V} \left| \sum_{i=1}^N \varepsilon_{i,j} g_{\theta,j}(X_i) \right|.$$

Here  $g_{\theta,j}$  are the entries of the vector field  $g_\theta \in V$ , and the  $\varepsilon_{i,j}$  are all i.i.d.  $N(0, 1)$  variables. We will now bound the supremum over  $\Theta_{N,M,\infty}$  of the each of the last  $p_V$  summands by using a moment inequality for the empirical process  $\{ \sum_{i=1}^N f_\theta(Z_i) : f \in \mathcal{F} \}$ , where for every  $1 \leq j \leq p$  fixed (and with  $e$  denoting a real variable in this proof in slight abuse of notation),

$$f_\theta \in \mathcal{F} \equiv \mathcal{F}_j = \{ f_\theta(z) = e g_{\theta,j}(x) : \theta \in \Theta_{N,M,\infty} \}, \quad z = (e, x) \in \mathbb{R} \times \mathcal{X},$$

and  $Z_1, \dots, Z_N$  are i.i.d. copies of the variables  $Z = (\varepsilon, X) \sim N(0, 1) \times \lambda = P$ .

We will apply [60, Theorem 3.5.4] but to do so we need to calculate some preliminary bounds. First, by independence of  $X, \varepsilon$ , the ‘weak’ variances of  $\mathcal{F}$  are of order

$$\sup_{\theta \in \Theta_{N,M,\infty}} E f_\theta^2(Z) = \sup_{\theta \in \Theta_{N,M,\infty}} E g_{\theta,j}^2(X) \leq \sup_{\theta \in \Theta_{N,M,\infty}} \rho_{\theta_0}^2(\theta_{(t)} - \theta_0) \leq \sigma_N^2,$$

by Conditions 3.1.1 and 4.1.4. Next, by Condition 4.1.1 and the definition of  $\Theta_{N,M,\infty}$ , we have

$$\sup_{\theta \in \Theta_{N,M,\infty}} \|g_{\theta,j}\|_\infty \lesssim \|\theta_{(t)} - \theta_0\|_\infty \lesssim \bar{\delta}_N(1 + \|\bar{\psi}\|_\infty) \lesssim \bar{\delta}_N.$$

As a consequence, the preceding empirical process has point-wise envelopes

$$\sup_{\theta \in \Theta_{N,M,\infty}} |f_\theta(e, x)| \lesssim |e| \bar{\delta}_N \equiv F_N(e, x) \quad \text{for all } (e, x) \in \mathbb{R} \times \mathcal{X},$$

in particular,  $F_N > 0$   $P$ -a.s. and

$$\begin{aligned} \|F_N\|_{L^2(P)}^2 &:= \int_{\mathbb{R} \times \mathcal{X}} F_N^2(z) dP(z) \lesssim \bar{\delta}_N^2, \\ \|F_N\|_{L^2(Q)}^2 &:= \int_{\mathbb{R} \times \mathcal{X}} F_N^2(z) dQ(z) \simeq \bar{\delta}_N^2 s_Q^2, \end{aligned}$$

where, for any (discrete, finitely supported) probability measure  $Q$  on  $\mathbb{R} \times \mathcal{X}$ , we

have set  $s_Q^2 := \int_{\mathbb{R} \times \mathcal{X}} e^2 dQ(e, x)$ . Finally, we have again from Condition 4.1.1, for any  $\theta, \theta' \in \Theta$  and some fixed constant  $c_0$ , that

$$\begin{aligned} \|f_\theta - f_{\theta'}\|_{L^2(Q)} &:= \sqrt{\int_{\mathbb{R}} \int_{\mathcal{X}} e^2 (g_{\theta,j}(x) - g_{\theta',j}(x))^2 dQ(e, x)} \\ &\leq s_Q \|g_{\theta,j} - g_{\theta',j}\|_\infty \\ &\leq s_Q (\|\mathcal{G}(\theta_{(t)}) - \mathcal{G}(\theta'_{(t)})\|_\infty + \|\mathbb{I}_{\theta_0}[\theta_{(t)} - \theta'_{(t)}]\|_\infty) \\ &\leq \frac{c_0}{\bar{\delta}_N} \|F_N\|_{L^2(Q)} \|\theta - \theta'\|_\infty. \end{aligned}$$

We conclude that any  $\bar{\delta}_N \epsilon / c_0$ -covering of  $\Theta_{N,M,\infty}$  for the norm  $\|\cdot\|_\infty$  induces a  $\|F_N\|_{L^2(Q)} \epsilon$ -covering of  $\mathcal{F}$  for the  $L^2(Q)$  norm, and so  $J(\mathcal{F}, F, s)$  in [60, equation (3.169)] is bounded by a constant multiple of our  $\mathcal{J}_N(s, \bar{\delta}_N)$  (using also [60, Lemma 3.5.3 (a)]). With these preparations, we can now apply [60, Theorem 3.5.4] where for our choice of envelope  $F_N$  we can take  $\|U\|_{L^2(P)}$  in that theorem bounded by a constant multiple of  $\sqrt{\log N} \bar{\delta}_N$  (using independence of  $X, \varepsilon$  and also [60, Lemma 2.3.3]). The upper bound (3.171) in [60] then implies that

$$\begin{aligned} E \sup_{\theta \in \Theta_{N,M,\infty}} \left| \sum_{i=1}^N f_\theta(Z_i) \right| \\ \lesssim \sqrt{N} \max \left[ \bar{\delta}_N \mathcal{J}_N \left( \frac{\sigma_N}{\bar{\delta}_N}, \bar{\delta}_N \right), \frac{\sqrt{\log N} \bar{\delta}_N^3 \mathcal{J}_N^2(\sigma_N / \bar{\delta}_N, \bar{\delta}_N)}{\sqrt{N} \sigma_N^2} \right], \end{aligned}$$

which in turn, using the substitution  $\bar{\delta}_N \epsilon = \rho$  in (4.4), is bounded by a constant multiple of the maximum of the second and third terms appearing in (4.6). Hence the remainder terms from (4.10) converge to zero in expectation, and then also in probability (by Markov's inequality). (Let us finally note that, strictly speaking, the application of [60, Theorem 3.5.4] requires  $0 \in \mathcal{F}$  and  $\mathcal{F}$  countable: If  $\|\theta_0\|_{\mathcal{R}} < M$ , then  $g_\theta = 0$  for  $\theta = \theta_0 - (t/\sqrt{N})\bar{\psi} \in \Theta_{N,M,\infty}$  and  $N$  large enough, so  $0 \in \mathcal{F}$ . Otherwise, we can recentre  $f_\theta$  at  $f_{\theta_*}$  for some arbitrary  $\theta_*$  and use a standard (one-dimensional) moment bound for  $E |\sum_{i=1}^N f_{\theta_*}(Z_i)| \leq \sqrt{N} \sigma_N \rightarrow 0$ . One then applies the previous argument to the class  $\mathcal{F} - f_{\theta_*}$ , so that the same overall bound holds true also in this case. Finally, by continuity of  $\theta \mapsto g_{\theta,j}$  on the totally bounded set  $\Theta_{N,M,\infty}$ , the supremum of the empirical process can be realised over a countable dense subset of  $\Theta_{N,M,\infty}$ , so the assumption that  $\mathcal{F}$  be countable can be met, too.) ■

**Lemma 4.1.8.** *We have, for any  $t \in \mathbb{R}$ , that*

$$\begin{aligned} \sup_{\theta \in \Theta_{N,M,\infty}} \left| \sum_{i=1}^N (|\mathcal{G}(\theta_0)(X_i) - \mathcal{G}(\theta_{(t)})(X_i)|_V^2 \right. \\ \left. - E_\lambda |\mathcal{G}(\theta_0)(X_i) - \mathcal{G}(\theta)(X_i)|_V^2) \right| = o_{P_{\theta_0}^N}(1). \end{aligned}$$

*Proof.* For the supremum of the empirical process  $\{\sum_{i=1}^N (f(X_i) - Ef(X_i)) : f \in \mathcal{F}\}$ , we will obtain a bound, this time with indexing class

$$\mathcal{F} = \{f_\theta = |\mathcal{G}(\theta_0)(\cdot) - \mathcal{G}(\theta_t)(\cdot)|_V^2 : \theta \in \Theta_{N,M,\infty}\}.$$

Using Condition 4.1.1, the envelopes of  $\mathcal{F}$  can be taken to be

$$\sup_{\theta \in \Theta_{N,M,\infty}} \|\mathcal{G}(\theta_0) - \mathcal{G}(\theta_t)\|_\infty^2 \lesssim \sup_{\theta \in \Theta_{N,M,\infty}} \|\theta_0 - \theta_t\|_\infty^2 \lesssim \bar{\delta}_N^2 \equiv F,$$

and we also have, since  $\|\mathcal{G}(\theta)\|_\infty \leq U$  for  $\theta \in \Theta_{N,M,\infty}$  by Condition 2.1.1, that

$$\|f_\theta - f_{\theta'}\|_\infty \lesssim \|\theta - \theta'\|_\infty \quad \text{for all } \theta, \theta' \in \Theta_{N,M,\infty}.$$

This implies, similar to the proof in the previous lemma, that a  $c_0 \bar{\delta}_N^2 \epsilon$ -covering of  $\Theta_{N,M,\infty}$  for the  $\|\cdot\|_\infty$ -norm (and  $c_0$  a small but fixed constant) induces a  $\|F\|_{L^2(Q)} \epsilon$ -covering of  $\mathcal{F}$  for the  $L^2(Q)$ -norm ( $Q$  any probability measure), and that the functional  $J(\mathcal{F}, F, s)$  in [60, equation (3.169)] is bounded by a constant multiple of our  $J(s, \bar{\delta}_N^2)$ . The convergence to zero required in the lemma now follows from [60, Theorem 3.5.4], in fact Remark 3.5.5 after it, the requirement (4.5) from Condition 4.1.4, and Markov's inequality.  $\blacksquare$

**4.1.1.4 Gaussian change of variables.** We now control the ratio of Gaussian integrals appearing in Proposition 4.1.6.

**Proposition 4.1.9.** *As  $N \rightarrow \infty$ , we have for any fixed  $t \in \mathbb{R}$  that*

$$\frac{\int_{\bar{\Theta}_N} e^{\ell_N(\theta_{(t)})} d\Pi(\theta)}{\int_{\bar{\Theta}_N} e^{\ell_N(\theta)} d\Pi(\theta)} \xrightarrow{P_{\theta_0}^N} 1.$$

*Proof.* If we denote by  $\Pi_t$  the Gaussian law of  $\theta_{(t)} = \theta - (t/\sqrt{N})\bar{\psi}$ , then the Cameron–Martin theorem (e.g., [60, Theorem 2.6.13]) provides the formula for the Radon–Nikodym density of

$$\frac{d\Pi_t}{d\Pi}(\theta) = \exp\left\{\frac{t}{\sqrt{N}} \langle \theta, \bar{\psi} \rangle_{\mathcal{H}_N} - \frac{t^2}{2N} \|\bar{\psi}\|_{\mathcal{H}_N}^2\right\}, \quad \theta \sim \Pi, \quad \bar{\psi} \in \mathcal{H}_N.$$

The ratio in the proposition thus equals

$$\frac{e^{-\frac{t^2}{2N} \|\bar{\psi}\|_{\mathcal{H}_N}^2} \int_{\bar{\Theta}_{N,t}} e^{\ell_N(\vartheta)} e^{\frac{t}{\sqrt{N}} \langle \vartheta, \bar{\psi} \rangle_{\mathcal{H}_N}} d\Pi(\vartheta)}{\int_{\bar{\Theta}_N} e^{\ell_N(\theta)} d\Pi(\theta)}, \quad \text{with } \bar{\Theta}_{N,t} = \{\vartheta = \theta_{(t)} : \theta \in \bar{\Theta}_N\}.$$

Uniformly in  $\theta \in T_N^c \subset \bar{\Theta}_N$ , from (4.8) we have, as  $N \rightarrow \infty$ , that

$$\frac{|t|}{\sqrt{N}} |\langle \theta, \bar{\psi} \rangle_{\mathcal{H}_N}| \lesssim \delta_N \|\bar{\psi}\|_{\mathcal{H}_N} \lesssim \sqrt{N} \delta_N^2 \|\bar{\psi}\|_{\mathcal{H}} \rightarrow 0,$$

by Condition 4.1.4, which also implies that  $(t^2/N)\|\bar{\psi}\|_{\mathcal{H}_N}^2 = o(1)$  since  $1/\sqrt{N} = o(\delta_N)$ . Now since

$$\frac{|t|}{\sqrt{N}} \sup_{\vartheta \in \bar{\Theta}_{N,t}} |\langle \vartheta, \bar{\psi} \rangle_{\mathcal{H}_N}| \leq \frac{|t|}{\sqrt{N}} \sup_{\theta \in T_N^c} |\langle \theta, \bar{\psi} \rangle_{\mathcal{H}_N}| + \frac{t^2}{N} \|\bar{\psi}\|_{\mathcal{H}_N}^2,$$

we deduce from what precedes that the last ratio of integrals equals

$$e^{o(1)} \times \frac{\int_{\bar{\Theta}_{N,t}} e^{\ell_N(\vartheta)} d\Pi(\vartheta)}{\int_{\bar{\Theta}_N} e^{\ell_N(\theta)} d\Pi(\theta)} = e^{o(1)} \times \frac{\Pi(\bar{\Theta}_{N,t}|D_N)}{\Pi(\bar{\Theta}_N|D_N)}.$$

The denominator converges to 1 in  $P_{\theta_0}^N$ -probability by (4.8), and so does then the numerator, using again (4.8) and that  $t\|\bar{\psi}\|_{\infty}/\sqrt{N} = o(\delta_N)$  and  $t\|\bar{\psi}\|_{\mathcal{R}}/\sqrt{N} = o(1)$  under the maintained assumptions. ■

Combining Propositions 4.1.6 and 4.1.9, we have for all  $t \in \mathbb{R}$ , as  $N \rightarrow \infty$ ,

$$E^{\Pi^{\bar{\Theta}_N}} \left[ \exp\{t\sqrt{N}(\langle \theta, \psi \rangle_{L_{\xi}^2(\mathcal{Z})} - \hat{\Psi}_N)\} | D_N \right] \rightarrow \exp\left\{ \frac{t^2}{2} \|\mathbb{I}_{\theta_0} \bar{\psi}\|_{L_{\lambda}^2(\mathcal{X})}^2 \right\} \quad (4.11)$$

in  $P_{\theta_0}^N$ -probability, and therefore, using also (4.9), for  $\theta \sim \Pi(\cdot|D_N)$ ,

$$\sqrt{N}(\langle \theta, \psi \rangle_{L_{\xi}^2(\mathcal{Z})} - \hat{\Psi}_N) | D_N \xrightarrow{d} N(0, \|\mathbb{I}_{\theta_0} \bar{\psi}\|_{L_{\lambda}^2(\mathcal{X})}^2), \quad (4.12)$$

by the in  $P_{\theta_0}^N$ -probability version of the usual implication that convergence of Laplace transforms implies convergence in distribution (see Exercise 4.3.3). This completes the proof of Theorem 4.1.5. ■

## 4.1.2 Asymptotic distribution of the posterior mean

To use an approximation as the one from Theorem 4.1.5 for applications to uncertainty quantification (see Section 4.1.3), we need to choose a *feasibly computable centring statistic* instead of the (infeasible)  $\hat{\Psi}_N$ . A desirable choice, both for inference and computation via MCMC, is

$$\langle \bar{\theta}_N, \psi \rangle_{L_{\xi}^2(\mathcal{Z}, \mathcal{W})}, \quad \text{where } \bar{\theta}_N = E^{\Pi}[\theta | (Y_i, X_i)_{i=1}^N]$$

is the mean of the posterior distribution. Under the maintained hypotheses, the Bochner integral  $E^{\Pi}[\theta | (Y_i, X_i)_{i=1}^N]$  can be shown to exist for any given data vector  $(Y_i, X_i)_{i=1}^N$ .

**Theorem 4.1.10.** *In the setting of Theorem 4.1.5, if  $\bar{\theta}_N = E^{\Pi}[\theta | (Y_i, X_i)_{i=1}^N]$  denotes the posterior mean, then we have, as  $N \rightarrow \infty$ ,*

$$\sqrt{N} \langle \theta - \bar{\theta}_N, \psi \rangle_{L_{\xi}^2(\mathcal{Z})} | (Y_i, X_i)_{i=1}^N \xrightarrow{d} N(0, \|\mathbb{I}_{\theta_0} \bar{\psi}_{\theta_0}\|_{L_{\lambda}^2(\mathcal{X}, \mathcal{V})}^2) \quad \text{in } P_{\theta_0}^N\text{-probability.}$$

Moreover, as  $N \rightarrow \infty$ , we also have

$$\sqrt{N} \langle \bar{\theta}_N - \theta_0, \psi \rangle_{L^2_{\xi}(\mathcal{Z}, \mathcal{W})} \xrightarrow{d} N(0, \|\mathbb{I}_{\theta_0} \bar{\psi}_{\theta_0}\|_{L^2_{\lambda}(\mathcal{X}, \mathcal{V})}^2).$$

The proof consists of a ‘quantitative’ uniform integrability argument employing the following lemma which provides a stochastic bound on the posterior second moments.

**Lemma 4.1.11.** *Under the hypotheses of Theorem 4.1.10, we have*

$$NE^{\Pi}[\langle (\theta, \psi)_{L^2_{\xi}(\mathcal{Z})} - \hat{\Psi}_N \rangle^2 | D_N] = O_{P_{\theta_0}^N}(1).$$

*Proof.* The left-hand side in the last display is bounded by

$$2NE^{\Pi}[\langle \theta - \theta_0, \psi \rangle_{L^2_{\xi}(\mathcal{Z})}^2 | D_N] + 2N(\hat{\Psi}_N - \langle \theta_0, \psi \rangle_{L^2_{\xi}(\mathcal{Z})})^2,$$

and in view of (4.7), the second term in the last decomposition is bounded in  $P_{\theta_0}^N$ -probability by the central limit theorem applied to  $W_N$  from (3.8) with  $h = \bar{\psi}_{\theta_0}$  (one also applies the continuous mapping theorem for  $x \mapsto x^2$  and Prohorov’s theorem [126] to deduce from convergence in distribution of  $NW_N^2$  that it is uniformly tight.)

It hence remains to bound the first term in the last decomposition. Define  $C_N = \{\|\theta - \theta_0\|_{\infty} \leq M\bar{\delta}_N\} \subset \Theta$  for  $M$  to be chosen, and write the first quantity in the last display as (two times)

$$NE^{\Pi}[\langle \theta - \theta_0, \psi \rangle_{L^2_{\xi}(\mathcal{Z})}^2 1_{C_N} | D_N] + NE^{\Pi}[\langle \theta - \theta_0, \psi \rangle_{L^2_{\xi}(\mathcal{Z})}^2 1_{C_N^c} | D_N] = \text{I} + \text{II}. \quad (4.13)$$

To deal with term II, we apply the Cauchy–Schwarz inequality to obtain the bound

$$N \sqrt{E^{\Pi}[\langle \theta - \theta_0, \psi \rangle_{L^2(\mathcal{Z})}^4 | D_N]} \sqrt{\Pi(\|\theta - \theta_0\|_{\infty} > M\bar{\delta}_N | D_N)},$$

and we now show that this term is bounded in  $P_{\theta_0}^N$ -probability: Using Lemma 1.3.3 and the arguments from Step 2 in the proof of Theorem 2.2.2 to bound the probability of the sets  $A_N$  from (1.28), Markov’s inequality and  $E_{\theta_0}^N e^{\ell_N(\theta) - \ell_N(\theta_0)} = 1$ , as well as Proposition 4.1.3 with  $M$  large enough so that  $b > A + 3$ , we obtain

$$\begin{aligned} & P_{\theta_0}^N(E^{\Pi}[\langle \theta - \theta_0, \psi \rangle^4 | D_N] \Pi(\|\theta - \theta_0\|_{\infty} > M\bar{\delta}_N | D_N)) > N^{-2}) \\ & \leq P_{\theta_0}^N(E^{\Pi}[\langle \theta - \theta_0, \psi \rangle^4 | D_N] e^{-bN\bar{\delta}_N^2} > N^{-2}) + o(1) \\ & \leq P_{\theta_0}^N\left(\frac{\int_{\Theta} \langle \theta - \theta_0, \psi \rangle^4 e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta)}{\int_{\Theta} e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta)} > e^{bN\bar{\delta}_N^2} N^{-2}, A_N\right) + o(1) \\ & \leq \|\psi\|_{L^2(\mathcal{Z})}^4 e^{(A+2-b)N\bar{\delta}_N^2} N^2 \int_{\Theta} \|\theta - \theta_0\|_{L^2(\mathcal{Z})}^4 E_{\theta_0}^N e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta) + o(1) \\ & \lesssim N^2 e^{-N\bar{\delta}_N^2} + o(1) \rightarrow 0 \end{aligned}$$

as  $N \rightarrow \infty$ , by hypothesis on  $\delta_N, \Pi_N$ . Collecting what precedes implies that the term II in (4.13) is indeed  $O_{P_{\theta_0}^N}^N(1)$ .

The next step is to bound the term I in (4.13). Recalling that  $\Pi^{\bar{\Theta}^N}[\cdot|D_N]$  denotes the posterior distribution arising from prior restricted and renormalised to  $\bar{\Theta}_N$ , we decompose

$$\begin{aligned} & NE^\Pi[\langle \theta - \theta_0, \psi \rangle_{L^2(\mathcal{Z})}^2 1_{C_N} | D_N] \\ &= NE^{\Pi^{\bar{\Theta}^N}}[\langle \theta - \theta_0, \psi \rangle_{L^2(\mathcal{Z})}^2 1_{C_N} | D_N] \\ &\quad + NE^\Pi[\langle \theta - \theta_0, \psi \rangle_{L^2(\mathcal{Z})}^2 1_{C_N} | D_N] - NE^{\Pi^{\bar{\Theta}^N}}[\langle \theta - \theta_0, \psi \rangle_{L^2(\mathcal{Z})}^2 1_{C_N} | D_N] \\ &= A + B. \end{aligned}$$

For term  $A$ , using  $x^2 \leq 2e^x$ ,  $x \geq 0$ , the definition of  $\hat{\Psi}_N$  from (4.7) and  $W_N = O_{P_{\theta_0}^N}(1)$  with  $h = \bar{\psi}_{\theta_0}$  from (3.8), the limit (4.11) at  $t = 1$  implies that for all  $N$  large enough and some  $r_N = o_{P_{\theta_0}^N}(1)$ ,

$$A \leq 2e^{W_N + r_N} e^{\frac{1}{2} \|\mathbb{I}_{\theta_0} \bar{\psi}\|_{L^2(\mathcal{X}, \mathcal{V})}^2},$$

and hence this term is stochastically bounded.

Finally, by definition of the events  $C_N$ , the term  $|B|$  can be written as

$$\begin{aligned} & N \left| \int_{C_N} \langle \theta - \theta_0, \psi \rangle_{L^2(\mathcal{Z})}^2 [d\Pi(\theta|D_N) - d\Pi^{\bar{\Theta}^N}(\theta|D_N)] \right| \\ & \leq MN \bar{\delta}_N^2 \|\psi\|_{L^1(\mathcal{Z})}^2 \|\Pi(\cdot|D_N) - \Pi^{\bar{\Theta}^N}(\cdot|D_N)\|_{TV} \\ & \lesssim N \bar{\delta}_N^2 \Pi(\bar{\Theta}_N^c | D_N) \lesssim N \bar{\delta}_N^2 O_{P_{\theta_0}^N}(e^{-bN\bar{\delta}_N^2}) = o_{P_{\theta_0}^N}(1), \end{aligned}$$

where we have used (4.9) and (4.8), completing the proof of the lemma.  $\blacksquare$

Now to prove the theorem, note that by (4.12) and (4.1), for

$$Z_n | D_N \equiv \sqrt{N}(\langle \theta, \psi \rangle_{L^2(\mathcal{Z})} - \hat{\Psi}_N) | D_N, \quad Z \sim N(0, \|\mathbb{I}_{\theta_0} \bar{\psi}\|_{L^2_\lambda(\mathcal{X})}^2),$$

and  $d_{\text{weak}}$  any metric for weak convergence of laws  $\mathcal{L}(\cdot)$  on  $\mathbb{R}$ , we have

$$d_{\text{weak}}(\mathcal{L}(Z_N | D_N), \mathcal{L}(Z)) \xrightarrow[N \rightarrow \infty]{P_{\theta_0}^N} 0. \quad (4.14)$$

The idea of the proof to follow is that the previous lemma implies (by uniform integrability) convergence of moments in the last limit (4.14), and thus that, since  $EZ = 0$ , the posterior mean equals  $\hat{\Psi}_N$  up to a stochastic term of order  $o(1/\sqrt{N})$ . However, as the probability measures  $\mathcal{L}(Z_N | D_N)$  to which this argument is applied are *random* via the data  $D_N$ , the proof requires some care. We will employ a contradiction

argument: To prove Theorem 4.1.10, it suffices by Theorem 4.1.5, Slutsky’s lemma and (3.8) with  $h = \bar{\psi}_{\theta_0}$  to prove that as  $N \rightarrow \infty$ ,

$$\sqrt{N}(\langle E^\Pi[\theta|D_N], \psi \rangle_{L^2_\xi(\mathcal{Z})} - \hat{\Psi}_N) \xrightarrow{\text{Pr}} 0, \quad (4.15)$$

where we write  $\text{Pr}$  for the probability measure  $P_{\theta_0}^\mathbb{N}$  on the underlying measurable space  $(\Omega, \mathcal{S}) := ((V \times \mathcal{X})^\mathbb{N}, \mathcal{S})$  supporting all data variables  $(D_N, N \in \mathbb{N})$ . Suppose the last limit does not hold true. Then there exists  $\Omega' \in \mathcal{S}$  of positive probability  $\text{Pr}(\Omega') > \tau$  and  $\zeta' > 0$  such that along a subsequence of  $N$  (still denoted by  $N$ ), we have

$$|\sqrt{N}(\langle E^\Pi[\theta|D_N(\omega)], \psi \rangle_{L^2_\xi(\mathcal{Z})} - \hat{\Psi}_N(\omega))| \geq \zeta' > 0 \quad \text{for } \omega \in \Omega'. \quad (4.16)$$

Now since convergence in  $\text{Pr}$ -probability implies  $\text{Pr}$ -almost sure convergence along a subsequence, we can extract a further subsequence of  $N$  such that (4.14) holds almost surely, that is, on an event  $\Omega_0 \subset \Omega$  such that  $\text{Pr}(\Omega_0) = 1$ . For each fixed  $\omega \in \Omega_0$ , we can use the Skorohod imbedding [46, Theorem 11.7.2] to construct (if necessary on a new probability space) new real random variables  $\tilde{Z}_N, \tilde{Z}$  such that their laws satisfy

$$\mathcal{L}(\tilde{Z}_N) = \mathcal{L}(Z_N|D_N(\omega)), \quad \mathcal{L}(\tilde{Z}) = \mathcal{L}(Z), \quad \tilde{Z}_N \xrightarrow[N \rightarrow \infty]{\text{a.s.}} \tilde{Z},$$

and we also know by Lemma 4.1.11 that  $E\tilde{Z}_N^2 = E[Z_N^2|D_N(\omega)] = O(1)$  for all  $\omega \in \Omega'_0 \subset \Omega_0$  of probability  $\text{Pr}(\Omega'_0) > 1 - \tau$  as close to one as desired. But this implies that the  $(\tilde{Z}_N : N \in \mathbb{N})$  are uniformly integrable real random variables so that almost sure convergence implies convergence of first moments [46, Theorem 10.3.6], that is,

$$E|Z_n|D_N(\omega) - Z| = E|\tilde{Z}_N - \tilde{Z}| \xrightarrow[N \rightarrow \infty]{} 0$$

for all  $\omega \in \Omega'_0$ . In particular, using also Fubini’s theorem,

$$\begin{aligned} & \sqrt{N}(\langle E^\Pi[\theta|D_N(\omega)], \psi \rangle_{L^2_\xi(\mathcal{Z})} - \hat{\Psi}_N(\omega)) \\ &= E^\Pi[\sqrt{N}(\langle \theta, \psi \rangle - \hat{\Psi}_N)|D_N(\omega)] \rightarrow EZ = 0 \end{aligned}$$

for  $\omega \in \Omega'_0$ . But if the last limit holds for all  $\omega \in \Omega'_0$  with probability  $\text{Pr}(\Omega'_0) > 1 - \tau$  we have a contradiction to (4.16) (as then  $\text{Pr}(\Omega) \geq \text{Pr}(\Omega') + \text{Pr}(\Omega'_0) > 1 - \tau + \tau = 1$ ), completing the proof of (4.15) and thus of the theorem.

### 4.1.3 Applications to uncertainty quantification

Let us illustrate a statistical application of the preceding theorems to Bayesian uncertainty quantification (UQ) for functionals  $\langle \theta, \psi \rangle_{L^2_\xi(\mathcal{Z})}$ . Consider a level  $1 - \xi$  Bayesian credible interval

$$C_N = \left\{ v \in \mathbb{R} : |v - \langle \bar{\theta}, \psi \rangle_{L^2_\xi(\mathcal{Z})}| \leq \frac{R_N}{\sqrt{N}} \right\}, \quad \Pi(C_N|D_N) = 1 - \xi,$$

where  $\bar{\theta} = E^\Pi[\theta|D_N]$  is the posterior mean given the data  $D_N = (Y_i, X_i)_{i=1}^N$  and  $0 < \xi < 1$  a given coverage level. Construction of the interval  $C_N$  requires MCMC computation of that mean and of the quantiles  $R_N$  of the posterior distribution, but not of the asymptotic variance appearing in Theorem 4.1.10.

Now let  $Q(t) = \Pr(|Z| \leq t)$ ,  $t \in \mathbb{R}$ , where  $Z$  is the limiting normal distribution occurring in Theorem 4.1.10. Assume that  $Z$  is not degenerate (i.e., of non-zero variance) so that  $Q$  is continuous and strictly increasing. Then by Theorem 4.1.10 and the last part of Exercise 4.3.3 we see that

$$Q(R_N) = Q(R_N) - \Pi(C_N|D_N) + (1 - \xi) \xrightarrow{P_{\theta_0}^N} 1 - \xi$$

as  $N \rightarrow \infty$ , and since  $Q$  is invertible with continuous inverse, the continuous mapping theorem implies

$$R_N \xrightarrow{P_{\theta_0}^N} Q^{-1}(1 - \xi).$$

By Slutsky's lemma and again Theorem 4.1.10, we then have as  $N \rightarrow \infty$ ,

$$\frac{Q^{-1}(1 - \xi)}{R_N} \sqrt{N} \langle \bar{\theta}_N - \theta_0, \psi \rangle_{L_\xi^2(\mathcal{Z}, \mathcal{W})} \xrightarrow{d} Z.$$

Next, let us compute the frequentist coverage probability of the credible set  $C_N$  for the true parameter  $\theta_0$ . We have from the last limit and the continuous mapping theorem for  $|\cdot|$  that

$$\begin{aligned} P_{\theta_0}^N(\langle \theta_0, \psi \rangle_{L_\xi^2} \in C_N) &= P_{\theta_0}^N(\sqrt{N} |\langle \theta_0 - \bar{\theta}_N, \psi \rangle_{L_\xi^2}| \leq R_N) \\ &= P_{\theta_0}^N\left(\frac{Q^{-1}(1 - \xi)}{R_N} \sqrt{N} |\langle \bar{\theta}_N - \theta_0, \psi \rangle_{L_\xi^2}| \leq Q^{-1}(1 - \xi)\right) \\ &\rightarrow Q(Q^{-1}(1 - \xi)) = 1 - \xi, \end{aligned}$$

showing that  $C_N$  is indeed a precise asymptotic level  $1 - \xi$  confidence set of stochastic diameter  $O(1/\sqrt{N})$ . One can obtain ‘simultaneous’ versions of these results for appropriate collections of  $\psi$ 's too if an appropriate ‘uniform’ BvM theorem is at hand, see [60, Theorem 7.3.23].

## 4.2 Solving information equations in PDE models

The new condition encountered in this chapter concerns the solvability of the information equation (3.16). The corresponding Condition 4.1.2 does not only stipulate that a solution  $\bar{\psi}$  exists, but also that it has sufficient regularity to belong to the space  $\Theta$ . Whether this is possible or not is a priori not clear and depends on the particular inverse problem via the information operator (cf. Definition 3.1.2) induced by  $\mathcal{G}$ .

For the PDE examples studied in these notes, our path reaches a perhaps surprising fork at this point: We will show in this section that for the Schrödinger model (2.5), the information operator is itself a composition of elliptic operators and can be inverted on the space  $C_c^\infty(\mathcal{X})$  of smooth functions of compact support in  $\mathcal{X}$  (or in fact on appropriate Sobolev spaces). As a consequence, a Bernstein–von Mises theorem can be proved for  $\langle \theta, \psi \rangle_{L^2}$  with  $\theta \sim \Pi(\cdot | D_N)$  a posterior draw and for all  $\psi \in C_c^\infty(\mathcal{X})$ . These ideas extend to ‘perturbed’ differential operators  $\mathcal{D} - f$  of the form mentioned after (1.4) – see the notes to this chapter for more discussion. In contrast, for Darcy’s problem (2.8), we will show that the information equation can *not* be solved for typical elements of  $C_c^\infty(\mathcal{X})$  (such as all *positive* smooth  $\psi$ ). In particular, in view of Theorem 3.1.5, this implies that the Gaussian BvM-type approximations predicted by Le Cam theory for posterior functionals  $\langle \theta, \psi \rangle$  do in fact *not* hold true for this PDE.

#### 4.2.1 A Bernstein–von Mises theorem for the Schrödinger equation

We consider the Schrödinger forward map from (2.5) with potential  $f = f_\theta = e^\theta$ ,  $\theta \in \Theta$ , and parameter space  $C_c^\infty \subset \Theta \subset H^\xi(\mathcal{X})$ ,  $\xi > d/2$ . The information operator  $\mathbb{I}_\theta^* \mathbb{I}_\theta$  is given by Theorem 3.3.1 and (3.44), (3.45). Now for  $\psi \in C_c^\infty(\mathcal{X})$  and recalling the Schrödinger operator  $\mathcal{L}_{f_\theta} = \frac{1}{2} \Delta - f_\theta$ , we can define

$$\bar{\psi}_{\theta_0} = \frac{\mathcal{L}_{f_{\theta_0}} \mathcal{L}_{f_{\theta_0}} \left[ \frac{\psi}{u_{f_{\theta_0}} f_{\theta_0}} \right]}{u_{f_{\theta_0}} f_{\theta_0}}, \quad (4.17)$$

where we recall that  $\mathfrak{G}(\theta_0) = u_{f_{\theta_0}}$  is bounded away from zero on  $\mathcal{X}$  (since  $h \geq h_{\min} > 0$  on  $\mathcal{X}$  and using (A.25)). If we assume further that  $\theta_0 \in C^\infty(\mathcal{X})$ , then Proposition A.5.2 and (4.17) imply that  $\bar{\psi}_{\theta_0} \in C_c^\infty(\mathcal{X}) \in \Theta$ .

Since  $\mathcal{L}_{f_{\theta_0}}^{-1}$  is the inverse of  $\mathcal{L}_{f_{\theta_0}}$  for Dirichlet boundary conditions (cf. (1.6) and (A.10)), we obtain from (3.44), (3.45) that

$$\begin{aligned} \mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0} \bar{\psi}_{\theta_0} &= \mathbb{I}_{\theta_0}^* \left[ \mathcal{L}_{f_{\theta_0}}^{-1} \mathcal{L}_{f_{\theta_0}} \mathcal{L}_{f_{\theta_0}} \left[ \frac{\psi}{u_{f_{\theta_0}} f_{\theta_0}} \right] \right] \\ &= u_{f_{\theta_0}} f_{\theta_0} \mathcal{L}_{f_{\theta_0}}^{-1} \mathcal{L}_{f_{\theta_0}} \left[ \frac{\psi}{u_{f_{\theta_0}} f_{\theta_0}} \right] = \psi \end{aligned} \quad (4.18)$$

for any  $\psi \in C_c^\infty(\mathcal{X})$ . That is,  $\bar{\psi} = \bar{\psi}_{\theta_0}$  solves the information equation (3.16) for such  $\psi$ , and Condition 4.1.2 is verified. In particular,  $(\mathbb{I}_{\theta_0}^* \mathbb{I}_{\theta_0})^{-1}$  is then a proper inverse on  $C_c^\infty(\mathcal{X})$ . Collecting these findings we have effectively proved the following Bernstein–von Mises theorem.

**Theorem 4.2.1.** Consider a Gaussian prior  $\Pi_N$  from (2.18) with  $\kappa = 0$  and base prior  $\Pi'$  as in Condition 2.2.1, where

$$C_c^\infty \subset \Theta, \quad \mathcal{H} \subset H^\alpha(\mathcal{X}), \quad \mathcal{R} = H^\xi(\mathcal{X}), \quad \xi > \alpha + \frac{d}{2},$$

and with  $\alpha$  satisfying

$$\frac{\alpha}{2\alpha + d} \frac{\alpha - d}{\alpha + 2 - d/2} > \frac{1}{3}. \quad (4.19)$$

Let  $\theta \sim \Pi(\cdot | D_N)$ , where  $\Pi(\cdot | D_N)$  is the posterior measure (1.12) on  $\Theta$  arising from observations  $D_N = (Y_i, X_i)_{i=1}^N$  in model (1.9) with  $\mathcal{G}(\theta)$  from (2.5). Denote the posterior mean by  $\bar{\theta}_N = E^\Pi[\theta | D_N]$ , and let  $\psi \in C_c^\infty(\mathcal{X})$ . Assume  $\theta_0 \in C^\infty(\mathcal{X}) \cap \mathcal{H}$  is such that  $\inf_{x \in \mathcal{X}} f_{\theta_0}(x) > f_{\min}$ . Then we have

$$\sqrt{N} \langle \theta - \bar{\theta}_N, \psi \rangle_{L_\lambda^2(\mathcal{X})} | D_N \xrightarrow{d} N(0, \sigma^2(f_{\theta_0}, \psi)) \quad \text{in } P_{\theta_0}^N\text{-probability,}$$

as  $N \rightarrow \infty$ , and moreover that

$$\sqrt{N} \langle \bar{\theta}_N - \theta_0, \psi \rangle_{L_\lambda^2(\mathcal{X})} \xrightarrow{d} N(0, \sigma^2(f_{\theta_0}, \psi)),$$

where the asymptotic variance is given by

$$\sigma^2(f_{\theta_0}, \psi) = \left\| \mathcal{L}_{f_{\theta_0}} \left[ \frac{\psi}{u_{f_{\theta_0}} f_{\theta_0}} \right] \right\|_{L_\lambda^2(\mathcal{X})}^2. \quad (4.20)$$

*Proof.* We need to check the conditions from Theorem 4.1.5. Condition 2.1.1 was checked for  $\kappa = 0$  in Proposition 2.1.2, while Condition 2.1.4 was checked with  $\eta = \xi/(2 + \xi)$  in Exercise 2.4.1. Condition 3.1.1 was verified in Theorem 3.3.1 while Condition 4.1.1 follows from Exercise 4.3.1. Condition 4.1.2 was checked just before the statement of the theorem. We can invoke Proposition 4.1.3 with  $\beta = \xi$ ,  $\kappa = 0$ , and verify Condition 4.1.4 as follows: The  $\|\cdot\|_\infty$ -covering numbers of a  $\beta$ -Sobolev ball in dimension  $d$  are of the order

$$\log N(\Theta_{N,M,\infty}, \|\cdot\|_\infty, \epsilon) \lesssim \left(\frac{1}{\epsilon}\right)^{d/\beta}, \quad \beta > 0,$$

see [60, equation (4.184)] for the case when the functions are defined over  $[0, 1]^d$ , and this bound applies to our setting by a standard extension argument (and regarding  $\mathcal{X}$  as a subset of  $[0, 1]^d$  without loss of generality). Also for  $\theta \in \Theta_{N,M,\infty}$ , we can take  $\rho_{\theta_0}(\theta - \theta_0 + (t/\sqrt{N})\bar{\psi}_{\theta_0}) \lesssim \bar{\delta}_N^2 \equiv \sigma_N$ . We first note that the quantity in (4.5) is bounded by

$$\sqrt{N} \bar{\delta}_N^2 \int_0^1 (\bar{\delta}_N^2 \epsilon)^{-d/(2\beta)} d\epsilon \lesssim \sqrt{N} \bar{\delta}_N^{2-d/\beta}, \quad (4.21)$$

since  $\beta > d/2$ . We will eventually show that the last bound converges to zero as  $N \rightarrow \infty$ , which also implies  $N\sigma_N^2 \lesssim N\bar{\delta}_N^4 \rightarrow 0$ . The middle term in the maximum

in (4.6) can similarly be bounded by

$$\sqrt{N} \mathcal{J}_N(\sigma_N, 1) = \sqrt{N} \int_0^{\sigma_N} \epsilon^{-d/(2\beta)} d\epsilon \lesssim \sqrt{N} \bar{\delta}_N^{2-d/\beta},$$

and hence is of the same order as the one in (4.21). For the third member in the maximum (4.6), by a similar calculation, we have

$$\frac{\bar{\delta}_N \sqrt{\log N}}{\sigma_N^2} \mathcal{J}_N^2(\sigma_N, 1) \lesssim \sqrt{\log N} \bar{\delta}_N^{1-2d/\beta}.$$

We can conclude from what precedes that it suffices to show that

$$\max(\sqrt{N} \bar{\delta}_N^{2-d/\beta}, N \bar{\delta}_N^3, \sqrt{\log N} \bar{\delta}_N^{1-2d/\beta}) \rightarrow 0 \quad \text{as } N \rightarrow \infty.$$

This requires  $\beta > 2d$  and then simplifies to the basic requirement  $N \bar{\delta}_N^3 \rightarrow 0$ . Using the convergence rate  $\bar{\delta}_N$  provided by Proposition 4.1.3 and (4.19), Condition 4.1.4 is verified. The final result now follows from Theorems 4.1.5 and 4.1.10 and since, arguing as in the first step in (4.18),

$$\tilde{\psi}_{\theta_0} = \mathbb{I}_{\theta_0} \bar{\psi}_{\theta_0} = \mathcal{L}_{f_{\theta_0}} \left[ \frac{\psi}{u_{f_{\theta_0}} e^{\theta_0}} \right]. \quad (4.22)$$

The proof is complete. ■

We note that the asymptotic variance in (4.20) is just the squared  $L_\lambda^2$ -norm of  $\tilde{\psi}_{\theta_0}$  from (4.22), and hence precisely equals the lower bound from Theorem 3.1.4.

The previous result can in principle be made *uniform* in collections  $\mathcal{W}$  of  $\psi$ 's that are bounded in suitable Sobolev norms. If we wish to obtain such uniformity to hold for weak convergence (in  $P_{\theta_0}^N$ -probability) of the posterior measure towards the limiting centred Gaussian process  $(\mathbb{X}(\psi) : \psi \in \mathcal{W})$  with covariance

$$E\mathbb{X}(\psi)\mathbb{X}(\psi') = \left\langle \mathcal{L}_{f_{\theta_0}} \left[ \frac{\psi}{u_{f_{\theta_0}} f_{\theta_0}} \right], \mathcal{L}_{f_{\theta_0}} \left[ \frac{\psi'}{u_{f_{\theta_0}} f_{\theta_0}} \right] \right\rangle_{L_\lambda^2}, \quad \psi, \psi' \in \mathcal{W}, \quad (4.23)$$

from (4.20), then this forces  $\|\psi\|_{H^\beta} \leq C$  for some  $C > 0$  and  $\beta > 2 + d/2$ . The reason is that otherwise the limiting process would not be a tight random variable, so that by Prohorov's theorem, laws cannot converge weakly towards it – see [93, Proposition 6] for a proof of this fact, and the notes to this section for more on the subject of ‘uniform in  $\psi$ ’ Bernstein–von Mises theorems.

## 4.2.2 Impossibility of the BvM-phenomenon for Darcy's problem

We now show that for the forward map arising from the divergence form equation (2.8), even when considering the most regular parameter  $\theta \equiv \text{const}$  corresponding

to the standard Laplacian on a  $2d$ -disk, the information equation (3.16) can in fact *not be solved* for generic classes of smooth  $\psi$ 's. The problem was already foreshadowed in the discussion before Proposition 2.1.5, only that now the Laplacian term that helped crucially in (2.14) has disappeared – the vector field  $\nabla u_\theta$  then entirely determines the solvability of the relevant information equation, and the counter-examples that follow exhibit that obstructions can arise from this in natural settings.

We will prove more than non-solvability of the information equation – we show that even the range condition  $\psi \in R(\mathbb{I}_\theta^*)$  is typically violated and that hence, in view of Theorem 3.1.5,  $\sqrt{N}$ -consistent estimation of parameters  $\langle \theta, \psi \rangle$  is not possible for a large class of smooth  $\psi$ 's. In particular, *a fortiori*, no semi-parametric Bernstein–von Mises theorem can hold true, as this would entail the existence of an efficient estimator (whose asymptotic distribution coincides with the centring  $\hat{\Psi}_N$  from (4.7)).

**4.2.2.1 The range of  $\mathbb{I}_\theta^*$  and transport PDEs.** We consider the forward map  $\mathcal{G}$  from (2.8) arising from solutions  $\mathcal{G}(\theta) = u_\theta$  of the PDE (1.2) with divergence form operator (1.3) for  $f = f_\theta = f_{\min} + e^\theta$ ,  $\theta \in \Theta = H_c^\beta(\mathcal{X})$ . From Proposition 3.3.3 the range of the adjoint operator is given by

$$R(\mathbb{I}_\theta^*) = \{\psi = e^\theta \nabla u_\theta \cdot \nabla \mathcal{L}_{f_\theta}^{-1}[g] \text{ for some } g \in L_\lambda^2(\mathcal{X})\}.$$

The operator  $\mathcal{L}_{f_\theta}^{-1}$  maps  $L_\lambda^2(\mathcal{X})$  into  $H_0^2(\mathcal{X}) = \{y \in H^2 : y|_{\partial\mathcal{X}} = 0\}$  and hence if  $\psi$  lies in the range of  $\mathbb{I}_\theta^*$ , then the equation

$$\begin{aligned} \nabla u_\theta \cdot \nabla y &= e^{-\theta} \psi && \text{on } \mathcal{X}, \\ y &= 0 && \text{on } \partial\mathcal{X}, \end{aligned} \tag{4.24}$$

necessarily has a solution  $y = y_\psi \in H_0^2$ . The existence of solutions to the transport PDE (4.24) depends in a possibly intricate way on the compatibility of  $\psi$  with geometric properties of the vector field  $\nabla u_\theta$ , which in turn is determined by the geometry of the forward map  $\mathcal{G}$  (via  $f, g, \theta$ ) in the base PDE (1.2).

We now illustrate the obstructions that can arise from the above transport problem in the representative model case where  $\mathcal{L}_{f_\theta}$  is the standard Laplacian  $\Delta$ , i.e.,  $\theta = 0$ ,  $f_{\min} = 0$ , and with source  $g = 2$ , boundary values  $h = (|\cdot|_{\mathbb{R}^2}^2 - 1)/2$ . As underlying domain, we choose either

- (i)  $\mathcal{X}$  equal to a unit disk in  $\mathbb{R}^2$  *separated away* from the origin  $(0, 0)$ , or
- (ii)  $\mathcal{X}$  equal to the unit disk in  $\mathbb{R}^2$  centred at  $(0, 0)$ .

We will refer to them as Example (i) and (ii), respectively, and note that the injectivity of  $\mathbb{I}_\theta$  has already been studied in Exercise 3.4.5 and Proposition 3.3.4, respectively. By a simple perturbation argument, the findings below extend to a neighbourhood of  $\theta = 0$ , see Exercise 4.3.4.

**4.2.2.2 Example (i).** In (i), the solution to (1.2) with  $f_\theta = 1$  equals  $u_\theta = h$  and hence  $\nabla u_\theta = x$  does not vanish on  $\mathcal{X}$  for  $\theta = 0$ . The last property extends to any  $\theta$  in a  $H^\beta$ -neighbourhood of the standard Laplacian  $\theta = 0$  by Exercise 4.3.4.

The integral curves  $\gamma(t)$  in  $\mathcal{X}$  associated to the smooth vector field  $\nabla u_\theta \neq 0$  are given near  $x \in \mathcal{X}$  as the unique solutions (e.g., [121, p. 9]) of the vector ODE

$$\frac{d\gamma}{dt} = \nabla u_\theta(\gamma), \quad \gamma(0) = x.$$

Since  $\nabla u_\theta$  does not vanish we obtain through each  $x \in \mathcal{X}$  a unique curve  $(\gamma(t) : 0 \leq t \leq T_\gamma)$  originating and terminating at the boundary  $\partial\mathcal{X}$ , with finite ‘travel time’  $T_\gamma \leq T(\mathcal{X}, c_\nabla) < \infty$ . Along this curve the chain rule implies

$$\frac{d}{dt} y(\gamma(t)) = \frac{d\gamma(t)}{dt} \cdot \nabla y(\gamma(t)) = (\nabla u_\theta \cdot \nabla y)(\gamma(t)) \quad \text{for } 0 < t < T_\gamma,$$

and the PDE (4.24) reduces to the ODE

$$\frac{dy}{dt} = \psi. \tag{4.25}$$

Now suppose  $\psi \in R(\mathbb{I}_\theta^*)$ . Then a solution  $y \in H_0^2$  to (4.24) satisfying  $y|_{\partial\mathcal{X}} = 0$  must exist. Such  $y$  then also solves the ODE (4.25) along each curve  $\gamma$ , with initial and terminal values  $y(0) = y(T_\gamma) = 0$ . By the fundamental theorem of calculus (and uniqueness of solutions) this forces

$$\int_0^{T_\gamma} \psi(\gamma(t)) dt = 0 \tag{4.26}$$

to vanish. In other words,  $\psi$  permits a solution  $y$  to (4.24) only if  $\psi$  integrates to zero along each integral curve (orbit) induced by the vector field  $\nabla u_\theta$ . Now consider any smooth (non-zero) *nonnegative*  $\psi \in C_c^\infty(\mathcal{X}) \subset \Theta$ , and take  $x \in \mathcal{X}$  such that  $\psi \geq c > 0$  near  $x$ . For  $\gamma$  the integral curve passing through  $x$ , we then cannot have (4.26) as the integrand never takes negative values while it is positive and continuous near  $x$ . We conclude by way of contradiction that  $\psi \notin R(\mathbb{I}_\theta^*)$ . Applying Theorem 3.1.5 with tangent space  $H = C_c^\infty$  (and with its hypotheses verified in Theorem 3.3.2 and Exercise 4.3.1), we have proved the following theorem.

**Theorem 4.2.2.** *Consider estimation of the functional  $\Psi(\theta) = \langle \theta, \psi \rangle_{L^2(\mathcal{X})}$  from data  $(Y_i, X_i)_{i=1}^N$  drawn i.i.d. from  $P_\theta^N$  in the model (1.9) with forward map  $\mathcal{G}$  from (2.8) with  $f_{\min} = 1$ , where  $g, h$  in (1.2) are chosen as in Proposition 3.3.4, with the domain  $\mathcal{X}$  separated away from the origin, and with tangent space  $H = H_c^\beta(\mathcal{X})$ . Suppose  $0 \neq \psi \in C_0^\infty(\mathcal{X})$  satisfies  $\psi \geq 0$  on  $\mathcal{X}$ . Then for  $\theta = 0$  the efficient information for estimating  $\Psi(\theta)$  satisfies*

$$\inf_{h \in H, \langle h, \psi \rangle_{L^2} \neq 0} \frac{\|\mathbb{I}_0 h\|_{L_\lambda^2}^2}{\langle \psi, h \rangle_{L_\lambda^2}^2} = 0. \tag{4.27}$$

In particular,

$$\liminf_{N \rightarrow \infty} \inf_{\bar{\psi}_N: (\mathbb{R} \times \mathcal{X})^N \rightarrow \mathbb{R}} \sup_{\substack{\theta' \in \Theta \\ \|\theta'\|_{L^2_{\xi}} \leq 1/\sqrt{N}}} NE_{\theta'}^N(\bar{\psi}_N - \Psi(\theta'))^2 = \infty. \quad (4.28)$$

The proof extends straightforwardly to a (fixed) neighbourhood of the Laplacian  $\theta = 0$  as the key property used in the previous proof (that  $\nabla u_{\theta}$  is a non-vanishing vector field) extends by perturbation, see Exercise 4.3.4.

Let us notice that one can further show that (4.26) is also a *sufficient* condition for  $\psi$  to lie in the range of  $\mathbb{I}_{\theta}^*$  (provided  $\psi$  is smooth and with compact support in  $\mathcal{X}$ ). As this condition strongly depends on  $\theta$  via the vector field  $\nabla u_{\theta}$ , it seems difficult to describe any choices of  $\psi$  that lie in  $\bigcap_{\theta \in \Theta} R(\mathbb{I}_{\theta}^*)$ .

The problem in this example is that while  $\mathbb{I}_{\theta}$  is injective on  $H_0^1 \supset \Theta$ , it is not injective any longer on the  $L^2$ -closure of  $\Theta$ , which is the relevant issue for efficient estimation of functionals  $\langle \theta, \psi \rangle$ . The next counter-example is more subtle in this regard, as  $\mathbb{I}_{\theta}$  is injective on all of  $L^2$  in this case (see Proposition 3.3.4).

**4.2.2.3 Example (ii).** We now consider the same setting as in Theorem 4.2.2 but with  $\mathcal{X}$  the unit disk centred at the origin, so that the vector field  $\nabla u_{\theta} = \nabla u_0 = x$  has a critical point at the origin when  $\theta = 0$  (again this property of an isolated zero extends by perturbation to a neighbourhood of  $\theta$ , see Exercise 4.3.4).

We showed in Proposition 3.3.4 that for this example, the operator  $\mathbb{I}_{\theta}$  is injective on all of  $L^2(\mathcal{X})$ , and hence any  $\psi \in L^2(\mathcal{X})$  lies in the *closure* of the range of  $\mathbb{I}_{\theta}^*$  (recalling the Hilbert space identity  $\overline{R(\mathbb{I}_{\theta}^*)} = \ker(\mathbb{I}_{\theta})^{\perp}$ ). Nevertheless, there are many relevant  $\psi$ 's that are not contained in  $R(\mathbb{I}_{\theta}^*)$  (before taking the closure). In the present example, the gradient of  $u_{\theta}$  vanishes and the integral curves  $\gamma$  associated to  $\nabla u_{\theta} = (x_1, x_2)$  emanate along straight lines from  $(0, 0)$  towards boundary points  $(z_1, z_2) \in \partial\mathcal{X}$ , where necessarily  $y((z_1, z_2)) = 0$  if  $y$  is to solve (4.24). If we parameterise these lines as

$$\{(z_1 e^t, z_2 e^t) : -\infty < t \leq 0\},$$

then, as after (4.25), we see that if a solution  $y \in H_0^2$  to (4.24) exists, then  $\psi$  must necessarily satisfy

$$\int_{-\infty}^0 \psi(z_1 e^t, z_2 e^t) dt = 0 - y(0) = \text{const} \quad \text{for all } (z_1, z_2) \in \partial\mathcal{X}.$$

This again cannot happen, for example, for any non-negative non-zero  $\psi \in H$  that vanishes along a given curve  $\gamma$  (for instance, if it is zero in any given quadrant of  $\mathcal{X}$ ), as this forces  $\text{const} = 0$ . Theorem 3.1.5 again as in the preceding theorem yields the following:

**Theorem 4.2.3.** *Consider the setting of Theorem 4.2.2 but where now  $\mathcal{X}$  is the unit disk centred at  $(0, 0)$ , and where  $0 \leq \psi \in C_0^\infty(\mathcal{X})$ ,  $\psi \neq 0$ , vanishes along some straight ray from  $(0, 0)$  to the boundary  $\partial\mathcal{X}$ . Then (4.27) and (4.28) hold at  $\theta = 0$ .*

Arguing as after Proposition 3.3.4, the result can be extended to any  $\theta$  near the standard Laplacian by an application of the Morse lemma, see Exercise 4.3.4.

The last example reveals that it is the lack of a closed range of  $\mathbb{I}_{\theta_0}$  that creates the obstruction for efficient inference here – in particular, the non-existence of  $\sqrt{N}$ -consistent estimators for  $\langle \theta, \psi \rangle$  is *not* a consequence of the lack of injectivity of  $\mathbb{I}_\theta$ . In the asymptotic setting of Le Cam theory we thus encounter fundamental limitations unless the tangent space  $H$  is strongly constrained (to a closed proper subspace of  $L_\lambda^2$  such as  $E_D$  from (5.3) for  $D$  fixed, cf. also the proof of Theorem 5.3.2 below). In the next chapter we will provide some partial ‘non-asymptotic’ remedies for this issue.

## 4.3 Notes

### 4.3.1 Exercises

**Exercise 4.3.1.** Show that Condition 4.1.1 holds for the forward map (2.5) arising with the Schrödinger equation. [Hint: Argue as in Proposition 2.1.2 and Theorem 3.3.1 but replace the  $L^2$ -continuity estimates for  $\mathcal{L}_f^{-1}$  by  $L^\infty$ -Lipschitz estimates from (A.27).] Show further that the inequality (3.6) is satisfied for the forward map (2.8) with  $d \leq 3$  and any  $h \in C^1$ . [Hint: Adapt the estimate (2.10) via the Sobolev imbedding  $H^2 \subset L^\infty$  and the elliptic regularity estimate (A.31).]

**Exercise 4.3.2.** Prove (4.9). [See [126, p. 142].]

**Exercise 4.3.3.** For  $\mu_N, \mu$  random probability measures on  $\mathbb{R}$  defined on some probability space, suppose  $\int_{\mathbb{R}} e^{tx} d\mu_N(x) \rightarrow \int_{\mathbb{R}} e^{tx} d\mu(x)$  in probability for all  $t \in \mathbb{R}$ . Show that  $d_{\text{weak}}(\mu_N, \mu) \rightarrow 0$  in probability, and that the last limit also implies convergence to zero of the Kolmogorov distance  $\sup_{x \in \mathbb{R}} |\mu_N((-\infty, x]) - \mu((-\infty, x])| \rightarrow 0$ . [Hint: adapt the standard proof for non-random measures along almost surely convergent subsequences; also see the appendix of [31].]

**Exercise 4.3.4.** Suppose  $\|\theta - 1\|_{H^\beta} < \eta$  for some  $\beta > 1 + d/2$ . Show that  $\|u_\theta - u_1\|_\infty \rightarrow 0$  as  $\eta \rightarrow 0$ . [Hint: By the usual perturbation argument and the Sobolev imbedding,

$$\begin{aligned} \|u_\theta - u_1\|_{C^1} &\lesssim \|V_1[\nabla \cdot ((\theta - 1)\nabla u_\theta)]\|_{H^b} \lesssim \|(\theta - 1)\nabla u_\theta\|_{H^{b-1}} \\ &\lesssim \|\theta - 1\|_{H^{b-1}} \|u_\theta\|_{C^b} \leq \|\theta - 1\|_{H^\beta} \|u_\theta\|_{H^\beta} < C\eta. \end{aligned}$$

Show further that if  $\nabla u_1$  has exactly one isolated critical point in  $\mathcal{X}$ , then so has  $u_\theta$  for such  $\theta$  and  $\eta$  small enough.

### 4.3.2 Remarks and comments

The results from this chapter on Bernstein–von Mises theorems are taken from [91] which build on earlier work in [93] as well as [89] and more generally with ideas for ‘direct’ models going back to [28, 29, 31] – see also [96, 98, 111]. The essential observation in [28, 29] to avoid the negative results due to [53, 71] was that in infinite-dimensions, Bernstein–von Mises theorems only hold in a weak enough topology, specifically weak enough such that the limiting Gaussian process is well defined, cf. the discussion around (4.23). For inverse problems such ‘functional’ (uniform in  $\psi$ ) results in negative order Sobolev-type spaces are obtained in [93] and [96, 98] based on earlier work in direct models [29] (see also [27, 30, 32] for work in settings not related to PDEs). The ‘smooth test function’ approach here with  $\psi \in C_c^\infty$  allows for reasonably short proofs if smooth solutions to the information equation (3.16) exist (for smooth  $\psi$ ), as is the case in Theorem 4.2.1.

The counter-examples for solvability of the information equation in Darcy’s problem are due to [95]. They show that whether or not BvM theorems hold true depends in a delicate way on the PDE driving the inverse problem even in settings where ground truth and test function  $\theta_0, \psi$  are both smooth ( $C^\infty$ ).

The invertibility of the information operator for the Schrödinger equation was shown in [93] and this extends to other examples where a base differential operator  $\mathcal{D}$  is perturbed by some zero-order term (‘potential’). A more involved example for which the theory of this chapter was shown to work in [91] are again the non-Abelian X-ray transforms from (1.1): For these  $\mathbb{I}_\theta^*$  also turns out to be elliptic [105] at least in the interior of  $\mathcal{X}$ . Dealing with boundary issues at  $\partial\mathcal{X}$  is a non-trivial task but the invertibility of the information operator (with appropriate adjoint) was proved in [88, 89, 91].

Generally, when the information operator is elliptic then  $\mathbb{I}_\theta^*$  can be expected to have closed range which allows the Le Cam machinery underlying the proof of Theorem 4.1.5 to proceed, whereas in other settings (such as Darcy’s problem), the lack of closedness of the information operator appears to be at the heart of the issues pointed at here. These problems do not surface in the classical theory with fixed *finite-dimensional* tangent spaces  $H$ , as the range of a linear operator is then always closed and it suffices to prove injectivity of  $\mathbb{I}_\theta$ . In PDE inverse problems such a finite-dimensional setting can mask many subtleties encountered in high- and infinite-dimensions, see, e.g., [22] where it is shown that in the context of the Calderón problem from Section 1.1.1, when restricting to a fixed finite-dimensional space of piece-wise constant conductivities, a Bernstein–von Mises approximation does hold true at rate  $1/\sqrt{N}$  (even though the non-parametric convergence rates for Hölder smooth conductivities are at best logarithmic [1]).



## Chapter 5

# Posteriors are probably log-concave

We now describe a distinct non-asymptotic perspective on the phenomenon that for large sample sizes, posterior measures arising from Gaussian process priors may resemble ‘bell-shaped’ or ‘normal-type’ distributions even in settings where  $\mathcal{G}$  is non-linear and  $\Theta$  is high-dimensional. We also investigate consequences of rigorous such results on the possibility of efficient computation of posterior distributions by MCMC methods.

We wish to remove some of the limitations of the Bernstein–von Mises type results from the previous chapter. Theorems 4.2.2 and 4.2.3 show that in natural infinite-dimensional models  $\Theta$  and representative PDE settings such as with Darcy’s problem, smooth linear statistics of posterior measures may well *not* have an asymptotically Gaussian distribution. And even in settings where the posterior has an asymptotically normal distribution such as in the Schrödinger model considered in Theorem 4.2.1, the limiting Gaussian process obtained along the ‘convergence determining class’ of  $\psi \in C_c^\infty(\mathcal{X})$  has a covariance structure that prevents the posterior measure to converge to its Gaussian limit except for in norms that are much weaker than  $L^2$ -norms – see the paragraph containing (4.23) for more discussion. So while such approximations are very exact and precise results, they hold only in possibly too restrictive scenarios.

It pays off though to take a second look at the issue: In the ‘locally asymptotically normal (LAN)’ world of Le Cam theory and semi-parametric statistics, one requires injectivity of  $\mathbb{I}_\theta$  on all of  $L_\xi^2$  (see (3.15)) and needs exact solutions to the related information equation (3.16). On the other hand, from a non-asymptotic point of view, as discussed in Section 3.2, injectivity of  $\mathbb{I}_\theta$  on an appropriate high-dimensional approximation space  $E_D \subset \Theta$  is sufficient to obtain (average) local curvature of the log-likelihood function  $\ell_N$  and then also of the posterior distribution. In particular, an inequality such as (3.27) is sufficient and no exact solution of the information equation is required. This can make an essential difference as we illustrate in this section with Darcy’s problem (where this issue was already partly foreshadowed in Exercise 3.4.5 and the discussion preceding Proposition 3.3.4).

A classical way to take advantage of such local curvature of  $\ell_N$  would be to Taylor expand  $\ell_N$  around an appropriate centring and to use the resulting quadratic approximation to construct a Gaussian proxy measure for  $\Pi(\cdot|D_N)$  – these are called ‘Laplace approximations’ and will be discussed in more detail in Remark 5.1.4. We will take a different approach to this problem however and approximate the posterior density from the larger (itself infinite-dimensional) class of ‘unimodal’ *log-concave* distributions. These also explain the predicted ‘bell-shaped’ nature of pos-

teriors but allow for faster approximation rates (measured in the relevant Wasserstein distance on probability measures) than Laplace approximations. They are also naturally compatible with gradient-based MCMC methods, as we will see. Indeed, if a sufficiently good estimate of the centring is available, one can devise a Langevin MCMC algorithm that samples from the log-concave proxy and the Wasserstein approximation Theorem 5.1.3 implies that such methods perform well also as samplers for the actual posterior measure. In particular, computation of the posterior mean will be shown to be possible in ‘polynomial run-time’ in relevant parameters. This holds for ‘warm start’ MCMC methods – the question of initialisation is discussed in the notes to this section. We also provide a result that illustrates why the problem of computing target measures by ‘local’ MCMC methods that are initialised outside a region of log-concavity is generally ‘exponentially hard’ in high-dimensional settings without additional structural assumptions. The results in this chapter show how one can leverage PDE structure and local log-concavity to break such computational hardness barriers.

## 5.1 Wasserstein approximation of the posterior

The idea in this section is to exploit the ‘local curvature’ Theorem 3.2.3 and contraction rate theorems such as Theorem 2.3.1 to show that the posterior measure from (1.12) is very well approximated by a log-concave (unimodal) distribution on high-dimensional approximation spaces  $E_D \simeq \mathbb{R}^D$  of  $\Theta$  that is concentrated near  $\theta_0$  (or rather its projection  $\theta_{0,D}$  onto  $E_D$ ). This will be true on the ‘frequentist’ event (subset of  $(V \times \mathcal{X})^N$ )

$$\begin{aligned} \mathcal{E}_{\text{conv}} := & \left\{ \inf_{\theta \in \mathcal{B}} \lambda_{\min}[-\nabla^2 \ell_N(\theta)] \geq \frac{c_0 N D^{-\kappa_0}}{2} \right\} \\ & \cap \left\{ \sup_{\theta \in \mathcal{B}} [|\ell_N(\theta)| + \|\nabla \ell_N(\theta)\|_{\mathbb{R}^D} + \|\nabla^2 \ell_N(\theta)\|_{\text{op}}] \leq N(5c_1 D^{\kappa_1} + 1) \right\}, \end{aligned} \quad (5.1)$$

whose probability was bounded in Theorem 3.2.3, and where we recall the set

$$\mathcal{B} = \mathcal{B}_{\mathbf{r}} = \left\{ \theta \in E_D : \|\theta - \theta_{0,D}\|_{E_D} < \mathbf{r} \right\} \quad (5.2)$$

from (3.28) for some radius  $\mathbf{r}$  to be specified as well as the finite-dimensional approximation space  $E_D \subset \Theta \subset L^2(\mathcal{Z})$  from (3.23). In fact, for conciseness in the proofs that follow we will choose the eigenbasis of the Dirichlet Laplacian  $\Delta$  on the bounded smooth domain  $\mathcal{Z} \subset \mathbb{R}^d$ ; for eigen-pairs  $(\lambda_j, e_j) \in (0, \infty) \times H_0^1(\mathcal{Z})$  of the operator  $-L_{1,0}$  from (A.9), we take

$$E_D = \left\{ \vartheta = \sum_{j=1}^D t_j e_j, t_j = \langle \vartheta, e_j \rangle_{L^2(\mathcal{Z})} \right\} \subset \Theta := \tilde{H}^\alpha(\mathcal{Z}) \quad \text{for } D \in \mathbb{N}, \quad (5.3)$$

with the generalised Sobolev scale  $\tilde{H}^\alpha(\mathcal{Z}) \subset H_0^\alpha(\mathcal{Z})$  defined in (A.15) below. While the results in this section *do* depend quantitatively on the action of the information operator from Definition 3.1.2 on the discretisation space  $E_D$ , generalisations beyond the choice of the eigenfunctions of the Dirichlet Laplacian (e.g., to wavelet bases on domains, or singular value decompositions of other operators with non-standard boundary conditions) are within reach as long as condition (3.30) can be checked. See the notes to this section for more discussion.

The finite-dimensional approximation of  $\Theta$  can be modelled explicitly in a Bayesian way by considering a ‘sieved’ prior truncated at some large enough  $D \in \mathbb{N}$ . For  $\delta_N = N^{-\alpha/(2\alpha+d)}$ , consider a Gaussian probability measure  $\Pi_N = \text{Law}(\theta)$  on  $E_D$  arising as

$$\theta = \frac{1}{\sqrt{N\delta_N}} \sum_{j=1}^D g_j \lambda_j^{-\alpha/2} e_j \quad \text{for } D \leq N\delta_N^2 \text{ and } \alpha > 1 + d/2, \quad (5.4)$$

where the  $g_j$ ’s are i.i.d.  $N(0, 1)$ . The posterior distribution from data in (1.9) then arises as in (1.12) and the contraction rate of  $\Pi_N(\cdot | D_N)$  for general forward maps was already studied in Exercise 2.4.3.

We will require the ‘average’ curvature and regularity of the log-likelihood function that underpins Theorem 3.2.3 from earlier to hold on a sufficiently large neighbourhood of the projection  $\theta_{0,D}$  of  $\theta_0$  onto  $E_D$

**Condition 5.1.1.** For  $\mathcal{R} = \Theta = \tilde{H}^\alpha(\mathcal{Z}) \subset H_0^\alpha(\mathcal{Z})$  from (A.15) with  $\alpha > 1 + d/2$ , let  $\mathcal{G}$  satisfy Condition 2.1.1 with  $\kappa = 0$  and Condition 2.1.4 for some  $0 < \eta \leq 1$ . Suppose further that Conditions 3.2.1 and 3.2.2 hold for some  $\kappa_0, \kappa_1, \kappa_2 > 0$ ,  $c_0, c_1, c_2 > 0$ , and  $\mathcal{B}, E_D \subset \Theta$  as in (5.2), (5.3), respectively, and with a choice of radius  $\mathbf{r}$  satisfying, for  $\delta_N = N^{-\alpha/(2\alpha+d)}$ , that

$$\mathbf{r} \geq (\log N) \tilde{\delta}_N, \quad \text{where } \tilde{\delta}_N \equiv (\log N) \max(\delta_N^\eta, D^{\kappa_0+\kappa_2} \delta_N). \quad (5.5)$$

The radius  $\mathbf{r}$  on which ‘average’ curvature of  $\ell_N$  is required for the main Theorem 5.1.3 below is thus determined by the ‘local’ and ‘global’ ill-posedness of the map  $\mathcal{G}$ , measured by the parameters  $\kappa_0$  and  $\eta$ , respectively. In fact, under somewhat stronger hypotheses (and by more involved proofs than those that follow), the term  $D^{\kappa_0+\kappa_2}$  can be replaced by the smaller term  $D^{\kappa_0/2}$  in (5.5), see the notes for discussion.

Under the above conditions it was shown in Theorem 3.2.3 that the event  $\mathcal{E}_{\text{conv}}$  from (5.1) has  $P_{\theta_0}^N$ -probability of order  $1 - O(\exp\{-CND^{-2\kappa_0-4\kappa_2}\})$ . In other words, the negative log-likelihood  $-\ell_N(\theta)$  is strongly convex on  $\mathcal{B}$  with high probability as long as the model dimension is not too large (so that  $ND^{-2\kappa_0-4\kappa_2} \rightarrow \infty$ ). Note that while one can choose  $D$  to diverge with  $N$ , the results in this section are ‘non-asymptotic’ and hold for all  $D$  bounded by constants to be specified. In particular, the notation  $D_N$  remains reserved for the data vector (1.11).

### 5.1.1 Construction of a log-concave surrogate posterior

Given a choice of  $\mathbf{r}$  from Condition 5.1.1 and the projection  $\theta_{0,D}$  of  $\theta_0$  onto  $E_D$ , let  $\theta_{\text{init}}$  be any fixed vector in  $E_D$  such that

$$\|\theta_{\text{init}} - \theta_{0,D}\|_{E_D} \leq \frac{\mathbf{r}}{8}. \quad (5.6)$$

For the purposes of log-concave approximation theorems one may take  $\theta_{\text{init}} = \theta_{0,D}$  itself but since  $\theta_{0,D}$  is unknown, for applications to bounds on mixing times of ‘feasible’ MCMC type algorithms, a separate ‘data-driven’ initialiser  $\theta_{\text{init}}$  is typically required. See the notes to this section for more discussion.

We require two auxiliary functions,  $g_{\mathbf{r}}$  (with strongly convex tails) and  $\alpha_{\mathbf{r}}$  (cut-off function): For some smooth and symmetric (about 0)  $\varphi: \mathbb{R} \rightarrow [0, \infty)$  supported in  $[-1, 1]$  and integrating to  $\int_{\mathbb{R}} \varphi(x) dx = 1$ , let us define mollifiers  $\varphi_h(x) := h^{-1}\varphi(x/h)$ ,  $h > 0$ . Then we define  $\tilde{\gamma}_{\mathbf{r}}, \gamma_{\mathbf{r}}: \mathbb{R} \rightarrow \mathbb{R}$  by

$$\tilde{\gamma}_{\mathbf{r}}(t) := \begin{cases} 0 & \text{if } t < 5\mathbf{r}/8, \\ (t - 5\mathbf{r}/8)^2 & \text{if } t \geq 5\mathbf{r}/8, \end{cases} \quad \gamma_{\mathbf{r}}(t) := [\varphi_{\mathbf{r}/8} * \tilde{\gamma}_{\mathbf{r}}](t),$$

where  $*$  denotes convolution. Further, let  $\alpha: [0, \infty) \rightarrow [0, 1]$  be smooth and satisfy  $\alpha(t) = 1$  for  $t \in [0, 3/4]$  and  $\alpha(t) = 0$  for  $t \in [7/8, \infty)$ . We then define  $g_{\mathbf{r}}: E_D \rightarrow [0, \infty)$  and  $\alpha_{\mathbf{r}}: E_D \rightarrow [0, 1]$  as

$$g_{\mathbf{r}}(\theta) := \gamma_{\mathbf{r}}(\|\theta - \theta_{\text{init}}\|_{E_D}), \quad \alpha_{\mathbf{r}}(\theta) = \alpha\left(\frac{\|\theta - \theta_{\text{init}}\|_{E_D}}{\mathbf{r}}\right). \quad (5.7)$$

Now for  $\ell_N$  as in (1.13) and  $K$  to be specified, define the ‘surrogate’ likelihood function

$$\tilde{\ell}_N(\theta) = \alpha_{\mathbf{r}}(\theta)\ell_N(\theta) - Kg_{\mathbf{r}}(\theta), \quad \theta \in E_D, \quad (5.8)$$

which ‘convexifies’  $-\ell_N$  in the ‘tails’ as Proposition 5.1.2 shows. The proof is not difficult but somewhat technical and relegated to Exercise 5.4.1.

**Proposition 5.1.2.** *Let  $\tilde{\ell}_N$  from (5.8) be defined for some constant  $K$  satisfying*

$$K \geq CN(c_1 D^{\kappa_1} + 1) \cdot (1 + \mathbf{r}^{-2}), \quad (5.9)$$

with  $c_1, \kappa_1$  from Condition 5.1.1 and  $C > 1$  depending only on the choice of the function  $\alpha$  in (5.7). We then have

$$\ell_N(\theta) = \tilde{\ell}_N(\theta) \quad \text{for all } \theta \in E_D \text{ such that } \|\theta - \theta_{0,D}\|_{E_D} \leq \frac{3\mathbf{r}}{8}. \quad (5.10)$$

Moreover, on the event  $\mathcal{E}_{\text{conv}}$  from (5.1) we have that  $\tilde{\ell}_N \in C^2(\mathbb{R}^D)$  (where  $\mathbb{R}^D \simeq E_D$ ) and

$$\inf_{\theta \in E_D} \lambda_{\min}(-\nabla^2 \tilde{\ell}_N(\theta)) \geq \frac{c_0 ND^{-\kappa_0}}{2},$$

as well as

$$\|\nabla \tilde{\ell}_N(\theta) - \nabla \tilde{\ell}_N(\bar{\theta})\|_{E_D} \leq 7K \|\theta - \bar{\theta}\|_{E_D} \quad \text{for all } \theta, \bar{\theta} \in E_D.$$

The resulting renormalised probability measure  $\tilde{\Pi}(\cdot | (Y_i, X_i)_{i=1}^N) = \tilde{\Pi}(\cdot | D_N)$  will be our ‘surrogate’ posterior distribution. In light of the preceding proposition, on the high probability event  $\mathcal{E}_{\text{conv}}$  this measure has a log-concave probability density on  $E_D$  given by

$$\tilde{\pi}(\theta | (Y_i, X_i)_{i=1}^N) = \frac{e^{\tilde{\ell}_N(\theta)} \pi(\theta)}{\int_{E_D} e^{\tilde{\ell}_N(\theta)} \pi(\theta)}, \quad \theta \in E_D, \quad (5.11)$$

where  $\pi$  is the probability density of the prior from (5.4).

### 5.1.2 The log-concave approximation theorem

The Wasserstein distance  $W_2$  between probability measures  $\mu$  and  $\nu$  on  $E_D$  is

$$W_2^2(\mu, \nu) = \inf_{\tau} \int_{E_D \times E_D} |\theta - \theta'|_{E_D}^2 d\tau(\theta, \theta'), \quad (5.12)$$

where the infimum is taken over all ‘couplings’  $\tau$  of  $\mu$  and  $\nu$ , that is, probability measures  $\tau$  on  $E_D \times E_D$  whose marginals are  $\mu$  and  $\nu$ , respectively.

The following theorem establishes a quantitative bound on the log-concavity of  $\tilde{\Pi}(\cdot | D_N)$ , that the gradient of its log-density coincides with the gradient log-posterior near  $\theta_{0,D}$  (and hence induces the same Langevin dynamics there, as is relevant in Section 5.2.1 to follow), and that, in particular, the posterior measure  $\Pi(\cdot | D_N)$  is well approximated in Wasserstein distance by  $\tilde{\Pi}(\cdot | D_N)$ , with high probability under the law  $P_{\theta_0}^N$  of the data  $D_N = (Y_i, X_i)_{i=1}^N$ . The ground truth  $\theta_0 \in \Theta$  is assumed to belong to the Sobolev scale  $\tilde{H}^\alpha(\mathcal{Z}) \supset E_D$  from (A.15), and it is further assumed to be well approximated by its  $L^2$ -projection  $\theta_{0,D}$  onto  $E_D$ :

$$\|\mathcal{G}(\theta_0) - \mathcal{G}(\theta_{0,D})\|_{L_\lambda^2} \leq \frac{\delta_N}{2}. \quad (5.13)$$

The latter condition can be checked easily using the  $L^2$ -Lipschitz property of  $\mathcal{G}$  from Condition 2.1.1 combined with  $\theta_0 \in \tilde{H}^\alpha(\mathcal{Z})$ , if either  $\theta_0$  is smooth and/or  $D$  is large enough.

**Theorem 5.1.3** (Log-concave Wasserstein approximation). *Let the spaces  $E_D$  and prior  $\Pi = \Pi_N$ ,  $\alpha$  be as in (5.3) and (5.4), respectively. Let the posterior distribution  $\Pi(\cdot | D^N)$  from (1.12) arise from data  $D_N = (Y_i, X_i)_{i=1}^N \sim P_{\theta_0}^N$  from (1.9) with forward map  $\mathcal{G}$  from (1.8) satisfying Condition 5.1.1 for some  $\mathbf{r} \leq 1$ . Assume further that  $D \leq A_1 N^{d/(2\alpha+d)}$  for some  $A_1 < \infty$  and that  $\theta_0 \in \tilde{H}^\alpha(\mathcal{Z})$  satisfies (5.13). Let the surrogate posterior measure  $\tilde{\Pi}(\cdot | D_N)$  be as in (5.11) for the given value of  $\mathbf{r}$  and any  $K \geq N(\log N)D^{\kappa_1}/\mathbf{r}^2$ .*

We then have the following:

(A) For all  $\theta \in \{\theta \in E_D : \|\theta - \theta_{0,D}\|_{E_D} \leq 3\mathbf{r}/8\}$  we have  $\tilde{\ell}_N(\theta) = \ell_N(\theta)$  and

$$\nabla \log \tilde{\pi}(\cdot | D_N)(\theta) = \nabla \log \pi(\cdot | D_N)(\theta).$$

(B) On an event  $\mathcal{E}_N \subset (V \times \mathcal{X})^N$  of  $P_{\theta_0}^N$ -probability at least  $1 - O(e^{-CND^{-2\kappa_0 - 4\kappa_2}})$  for some  $C > 0$ , the probability measure  $\tilde{\Pi}(\cdot | D_N)$  is globally log-concave on  $E_D$ , specifically (with  $c_0$  from Condition 3.2.2)

$$\inf_{\theta \in E_D} \lambda_{\min}(-\nabla^2 \log \tilde{\pi}(\theta | D_N)) \geq \frac{c_0}{2} ND^{-\kappa_0},$$

and  $\nabla \tilde{\ell}_N$  is globally Lipschitz on  $E_D$  with Lipschitz constant  $7K$ .

(C) On an event  $\mathcal{E}_N \subset (V \times \mathcal{X})^N$  of  $P_{\theta_0}^N$ -probability at least  $1 - O(e^{-bN\delta_N^2 - e^{-CND^{-2\kappa_0 - 4\kappa_2}}})$  for some  $b > 0, C > 0$ , we have the Wasserstein approximation

$$W_2^2(\Pi(\cdot | D_N), \tilde{\Pi}(\cdot | D_N)) \leq e^{-N^{d/(2\alpha+d)}}.$$

The theorem and the constants implicit in the  $O$ -notation are ‘non-asymptotic’ in the sense that whenever the hypotheses hold for pairs  $(D, N) \in \mathbb{N}^2$ , then so do the conclusions, (which are informative only when  $N \rightarrow \infty$ ).

**Remark 5.1.4** (Laplace approximations). Let us remark that the log-concave approximations provided by the previous theorem are qualitatively different from commonly used ‘Laplace approximations’ for numerical integration of Gibbs measures (1.12). In such approximations the posterior is replaced by the Gaussian measure centred at a maximiser  $\bar{\theta}$  of  $\ell_N$  with inverse covariance equal to the Hessian of  $\ell_N$  at  $\bar{\theta}$  – see [134] and also [14, 66, 87, 115, 117] for recent references. In other words one replaces  $\Pi(\cdot | D_N)$  by a measure with probability density

$$n(\theta | D_N) \propto \exp\left\{\frac{1}{2}(\theta - \bar{\theta}_N)^\top \nabla^2 \ell_N(\bar{\theta})(\theta - \bar{\theta}_N)\right\} \pi(\theta), \quad \theta \in E_D,$$

which on the high probability event  $\mathcal{E}_{\text{conv}}$  from (5.1) and for  $\bar{\theta} \in \mathcal{B}$  gives a well-defined Gaussian approximation. Unlike in our log-concave approximation, the Hessian of  $\log n(\cdot | D_N)$  is constant in  $\theta$ . The numerical validity of such Laplace approximations is related to the stochastic order of the higher order terms in the Taylor series of  $\ell_N$  about  $\bar{\theta}$  and cannot in general be expected to result in convergence rates that are faster than algebraic in  $1/N$ . The added flexibility of being able to approximate from the (infinite-dimensional) class of *log-concave measures* instead of just from normal distributions hence can lead to substantially faster approximation rates such as those in Theorem 5.1.3 (C). Moreover, by virtue of Part (A), the gradient of the log posterior remains the same in the region  $\mathcal{B}$  where the bulk of the mass of  $\Pi(\cdot | D_N)$  concentrates – the key statistical features of the posterior measure are hence not affected by

the approximation. In contrast Laplace approximations do modify the posterior also in the core region  $\mathcal{B}$ .

**5.1.2.1 Proof of Theorem 5.1.3.** As mentioned at the outset of this section, there are two main ideas underpinning the proof of this theorem: One is Theorem 3.2.3 which implies that  $\ell_N$  and then also the posterior distribution are *locally* log-concave near  $\theta_{0,D}$ , with high probability. The second ingredient is that the posterior measure is statistically consistent for  $\theta_0$  under  $P_{\theta_0}^N$  (see Section 2.3, and Exercise 2.4.3 for the particular prior employed) and hence, by (5.5), charges most of its mass precisely to the neighbourhood of  $\theta_{0,D}$  where  $\Pi(\cdot | (Y_i, X_i)_{i=1}^N)$  is log-concave. In particular, the surrogate posterior differs from the real posterior only in an area of insignificant posterior mass, hence providing an accurate approximation. Making the last statement quantitative in high dimensions is the main challenge in the proof.

Part (A) follows directly from (5.10) and the definition of  $\tilde{\pi}(\cdot | (Y_i, X_i)_{i=1}^N)$  in (5.11). We next note that on the event  $\mathcal{E}_{\text{conv}}$  from (5.1) and for our choice of  $K$  we can invoke Proposition 5.1.2 to verify the required gradient Lipschitz property in Part (B) as well as

$$\sup_{\theta, \vartheta \in \mathbb{R}^D, \|\vartheta\|_{\mathbb{R}^D} = 1} \vartheta^\top [\nabla^2 \tilde{\ell}_N(\theta)] \vartheta \leq -\frac{c_0}{2} ND^{-\kappa_0}. \quad (5.14)$$

Part (B) of Theorem 5.1.3 now follows from Theorem 3.2.3 bounding the probability of the event  $\mathcal{E}_{\text{conv}}$  (cf. also after Condition 5.1.1). The rest of the proof is concerned with part (C), using Part (A) and (B) and without loss of generality for all  $N$  large enough.

**5.1.2.2 Decomposition of the Wasserstein distance.** For this proof let us parametrise the Euclidean balls

$$\mathcal{B}(r) = \{\theta \in E_D : \|\theta - \theta_{0,D}\|_{E_D} \leq r\}$$

by their radius  $r > 0$ . By [132, Theorem 6.15] with  $x_0 = \theta_{0,D}$ ,

$$W_2^2(\tilde{\Pi}(\cdot | D_N), \Pi(\cdot | D_N)) \leq 2 \int_{E_D} \|\theta - \theta_{0,D}\|_{E_D}^2 d|\tilde{\Pi}(\cdot | D_N) - \Pi(\cdot | D_N)|(\theta).$$

This can be further bounded, for  $m > 0$  to be chosen and  $\tilde{\delta}_N$  from Condition 5.1.1, by

$$\begin{aligned} &\leq 2m^2 \tilde{\delta}_N^2 \int_{\mathcal{B}(m\tilde{\delta}_N)} d|\Pi(\cdot | D_N) - \tilde{\Pi}(\cdot | D_N)|(\theta) \\ &\quad + 2 \int_{\|\theta - \theta_{0,D}\|_{E_D} > m\tilde{\delta}_N} \|\theta - \theta_{0,D}\|_{E_D}^2 d\tilde{\Pi}(\theta | D_N) \\ &\quad + 2 \int_{\|\theta - \theta_{0,D}\|_{E_D} > m\tilde{\delta}_N} \|\theta - \theta_{0,D}\|_{E_D}^2 d\Pi(\theta | D_N) \equiv \text{I} + \text{II} + \text{III}, \end{aligned}$$

and we now bound I, II, III in separate steps.

**Term II.** We write the surrogate posterior density (5.11) as

$$\tilde{\pi}(\theta|D_N) = \frac{e^{\tilde{\ell}_N(\theta) - \tilde{\ell}_N(\theta_{0,D})} \pi(\theta)}{\int_{E_D} e^{\tilde{\ell}_N(\theta) - \tilde{\ell}_N(\theta_{0,D})} \pi(\theta) d\theta}, \quad \theta \in E_D.$$

Since  $\delta_N = o(\tilde{\delta}_N)$ , we have for  $\mathbf{r}$  from Condition 5.1.1 and some  $c = c(L) > 0$  small enough (with  $L = L_{\mathcal{G}}(1)$  the local Lipschitz constant from (2.4)), the set inclusion

$$\begin{aligned} B_N &:= \{\|\theta - \theta_{0,D}\|_{E_D} \leq c\delta_N\} \\ &\subset \left\{ \|\theta - \theta_{0,D}\|_{E_D} \leq \frac{3\mathbf{r}}{8} \right\} \cap \left\{ \|\mathcal{G}(\theta) - \mathcal{G}(\theta_0)\|_{L_\lambda^2} \leq \frac{\delta_N}{2} \right\}, \end{aligned}$$

using also (5.13). In particular,  $\ell_N(\theta) = \tilde{\ell}_N(\theta)$  on  $B_N$  (cf. (5.10)). Next note that a version of Lemma 1.3.3 with  $\ell_N(\theta_0)$  replaced by  $\ell_N(\theta_{0,D})$  is proved analogously, using (5.13) to control projection terms. Applying the so modified lemma with  $\nu = \Pi(\cdot)/\Pi(B_N)$  as well as the small ball estimate given in Step 2 in the proof of Theorem 2.2.2 (with  $\kappa = 0$  and adapted to the sieved case as in Exercise 2.4.3), we deduce that for  $\bar{c} = 2 + A$ ,

$$\begin{aligned} \int_{E_D} e^{\tilde{\ell}_N(\theta) - \tilde{\ell}_N(\theta_{0,D})} d\Pi(\theta) &\geq \int_{B_N} e^{\tilde{\ell}_N(\theta) - \tilde{\ell}_N(\theta_{0,D})} d\Pi(\theta) \\ &= \int_{B_N} e^{\ell_N(\theta) - \ell_N(\theta_{0,D})} d\nu(\theta) \times \Pi(B_N) \\ &\geq e^{-\bar{c}N\delta_N^2} \end{aligned} \tag{5.15}$$

on events  $\mathcal{A}_N$  of  $P_{\theta_0}^N$ -probability approaching one (in fact, of the required order  $1 - O(e^{-bN\delta_N^2})$ , if we increase  $\bar{c}$  appropriately, see Exercise 5.4.2).

Then on the preceding events the term II can be bounded, using a second-order Taylor expansion of  $\tilde{\ell}_N(\theta)$  around  $\theta_{0,D}$  with  $\bar{\theta}$  lying on the line segment connecting  $\theta$  and  $\theta_{0,D}$ , as

$$\begin{aligned} &\int_{\|\theta - \theta_{0,D}\|_{E_D} > m\tilde{\delta}_N} \|\theta - \theta_{0,D}\|_{E_D}^2 \tilde{\pi}(\theta|D_N) d\theta \\ &\leq e^{\bar{c}N\delta_N^2} \int_{\|\theta - \theta_{0,D}\|_{E_D} > m\tilde{\delta}_N} \|\theta - \theta_{0,D}\|_{E_D}^2 e^{\tilde{\ell}_N(\theta) - \tilde{\ell}_N(\theta_{0,D})} \pi(\theta) d\theta \\ &= e^{\bar{c}N\delta_N^2} \int_{\|\theta - \theta_{0,D}\|_{E_D} > m\tilde{\delta}_N} \|\theta - \theta_{0,D}\|_{E_D}^2 \\ &\quad \times e^{\nabla \tilde{\ell}_N(\theta_{0,D})^\top (\theta - \theta_{0,D}) + (\theta - \theta_{0,D})^\top \nabla^2 \tilde{\ell}_N(\bar{\theta}) (\theta - \theta_{0,D}) / 2} \pi(\theta) d\theta. \end{aligned}$$

For the first term in the exponent in the last integral we have (since  $\tilde{\ell}_N = \ell_N$  at  $\theta_{0,D}$ ), using the Cauchy–Schwarz inequality,

$$\begin{aligned} |\nabla \tilde{\ell}_N(\theta_{0,D})^\top (\theta - \theta_{0,D})| &\leq (\|\nabla \ell_N(\theta_{0,D}) - E_{\theta_0} \nabla \ell_N(\theta_{0,D})\|_{\mathbb{R}^D} \\ &\quad + \|E_{\theta_0} \nabla \ell_N(\theta_{0,D})\|_{\mathbb{R}^D}) \|\theta - \theta_{0,D}\|_{E_D}. \end{aligned} \tag{5.16}$$

In view of (3.35) and Condition 3.2.1 we can apply Bernstein's inequality (B.4) to the first summand to deduce

$$\|\nabla \ell_N(\theta_{0,D}) - E_{\theta_0} \nabla \ell_N(\theta_{0,D})\|_{\mathbb{R}^D}^2 \lesssim DN, \quad (5.17)$$

with sufficiently high probability (see Exercise 5.4.2 for details). Likewise,  $E_{\theta_0}$ -integrating the identity (3.35) and using Conditions 3.2.1, (5.13), the second summand is of order

$$\|E_{\theta_0} \nabla \ell_N(\theta_{0,D})\|_{\mathbb{R}^D} \lesssim ND^{\kappa_2} \|\mathcal{G}(\theta_{0,D}) - \mathcal{G}(\theta_0)\|_{L_\lambda^2} \lesssim ND^{\kappa_2} \delta_N.$$

Hence, using the hypothesis on  $D$ , the right-hand side in (5.16) is further upper bounded by

$$(\sqrt{DN} + ND^{\kappa_2} \delta_N) \|\theta - \theta_{0,D}\|_{E_D} \lesssim ND^{\kappa_2} \delta_N \|\theta - \theta_{0,D}\|_{E_D}.$$

Now by virtue of (5.14) the exponent in the integrand above (5.16) is bounded by

$$\left( \bar{L} \frac{ND^{\kappa_2} \delta_N}{\|\theta - \theta_{0,D}\|_{\mathbb{R}^D}} - \frac{c_0}{2} ND^{-\kappa_0} \right) \|\theta - \theta_{0,D}\|_{E_D}^2$$

for some  $\bar{L} > 0$ . Then for  $m$  large enough and by definition of  $\tilde{\delta}_N$  from Condition 5.1.1, we can bound this further from above by

$$-\frac{c_0}{4} ND^{-\kappa_0} \|\theta - \theta_{0,D}\|_{E_D}^2$$

for  $\theta$  in the relevant domain of integration, so that the last displayed term above (5.16) is bounded by

$$e^{\tilde{c}N\delta_N^2} e^{-\frac{c_0 m^2}{4} ND^{-\kappa_0} \tilde{\delta}_N^2} E^\Pi[\|\theta - \theta_{0,D}\|_{E_D}^2].$$

The second moments of the Gaussian measure  $\Pi_N$  are uniformly bounded, and  $\|\theta_{0,D}\|_{E_D}$  is uniformly bounded by a constant depending only on  $\|\theta_0\|_{L^2}$ . From Condition 5.1.1 we conclude that the term II does not exceed  $e^{-N\delta_N^2}/4$  with sufficiently high probability.

**Term III.** We first note that from Exercise 2.4.3 and (5.13) we can prove (2.27) and then (2.26) with  $\theta_{0,D}$  replacing  $\theta_0$  there, that is, the posterior contracts about  $\theta_{0,D}$  just as about  $\theta_0$ . This observation and (5.5) then imply that for every  $a > 0$ , we can find  $m$  large enough such that

$$\Pi(\|\theta - \theta_{0,D}\|_{E_D} > m\tilde{\delta}_N | D_N) \leq e^{-aN\delta_N^2}$$

on events  $\mathcal{S}'_N$  of sufficiently high  $P_{\theta_0}^N$ -probability. We can further restrict to  $\mathcal{A}_N$  on which the last inequality in (5.15) holds (see also Exercise 5.4.2). Then using the Cauchy–Schwarz inequality, the identity  $E_{\theta_0}^N e^{\ell_N(\theta) - \ell_N(\theta_0)} = 1$  as in (1.30) and

Markov's inequality, we have

$$\begin{aligned}
 & P_{\theta_0}^N \left( \mathcal{A}_N \cap \mathcal{S}'_N, \int_{\|\theta - \theta_{0,D}\|_{E_D} > m\tilde{\delta}_N} \|\theta - \theta_{0,D}\|_{E_D}^2 d\Pi(\cdot|D_N) > \frac{e^{-N\delta_N^2}}{8} \right) \\
 & \leq P_{\theta_0}^N \left( \mathcal{A}_N \cap \mathcal{S}'_N, \right. \\
 & \quad \left. \Pi(\|\theta - \theta_{0,D}\|_{E_D} > m\tilde{\delta}_N|D_N) E^\Pi[\|\theta - \theta_{0,D}\|_{E_D}^4|D_N] > \frac{e^{-2N\delta_N^2}}{64} \right) \\
 & \leq P_{\theta_0}^N \left( \mathcal{S}'_N, e^{(2+\bar{c}-a)N\delta_N^2} \int_{E_D} \|\theta - \theta_{0,D}\|_{E_D}^4 e^{\ell_N(\theta) - \ell_N(\theta_0)} d\Pi(\theta) > \frac{1}{64} \right) \\
 & \lesssim e^{(2+\bar{c}-a)N\delta_N^2} \int_{E_D} (1 + \|\theta\|_{E_D}^4) d\Pi(\theta) = O(e^{-bN\delta_N^2})
 \end{aligned}$$

whenever  $m$  and then  $a$  are large enough, since  $\Pi$  has uniformly bounded fourth moments and since  $\|\theta_{0,D}\|_{E_D}$  is uniformly bounded by a constant depending only on  $\|\theta_0\|_{L^2}$ . This shows that term III is of the required order with sufficiently high probability.

**Term I.** We have from Condition 5.1.1 that for fixed  $m > 0$  and all  $N$  large enough,

$$\mathcal{B}(m\tilde{\delta}_N) \subseteq \left\{ \theta : \|\theta - \theta_{0,D}\|_{E_D} \leq \frac{3\mathbf{r}}{8} \right\}.$$

On the latter set, by (5.10), the probability measures  $\tilde{\Pi}(\cdot|D_N)$  and  $\Pi(\cdot|D_N)$  coincide up to a normalising factor, and thus we can represent their Lebesgue densities as

$$\tilde{\pi}(\theta|D_N) = p_N \pi(\theta|D_N) \quad \text{for } \theta \in \mathcal{B}(m\tilde{\delta}_N),$$

for some  $0 < p_N < \infty$ . Moreover, by the preceding estimates for terms II and III (which hold just as well without the integrating factors  $n(\theta) := \|\theta - \theta_{0,D}\|_{\mathbb{R}^D}^2$ , as only  $E^\Pi n(\theta) \leq \text{const}$  was used), we have

$$\begin{aligned}
 p_N \Pi(\mathcal{B}(m\tilde{\delta}_N)|D_N) &= \tilde{\Pi}(\mathcal{B}(m\tilde{\delta}_N)|D_N) \geq 1 - \frac{e^{-N\delta_N^2}}{8} \Rightarrow 1 - \frac{e^{-N\delta_N^2}}{8} \leq p_N, \\
 p_N^{-1} \tilde{\Pi}(\mathcal{B}(m\tilde{\delta}_N)|D_N) &= \Pi(\mathcal{B}(m\tilde{\delta}_N)|D_N) \geq 1 - \frac{e^{-N\delta_N^2}}{8} \Rightarrow 1 - \frac{e^{-N\delta_N^2}}{8} \leq \frac{1}{p_N}
 \end{aligned}$$

on events of sufficiently high  $P_{\theta_0}^N$ -probability. On these events necessarily

$$p_N \in \left[ 1 - \frac{e^{-N\delta_N^2}}{8}, \left[ 1 - \frac{e^{-N\delta_N^2}}{8} \right]^{-1} \right],$$

and so for  $N$  large enough,

$$\begin{aligned}
 \int_{\mathcal{B}(m\tilde{\delta}_N)} d|\Pi(\cdot|D_N) - \tilde{\Pi}(\cdot|D_N)|(\theta) &= |1 - p_N| \int_{\mathcal{B}(m\tilde{\delta}_N)} \pi(\theta|D_N) d\theta \\
 &\leq |1 - p_N| \leq \frac{e^{-N\delta_N^2}}{4},
 \end{aligned}$$

which is obvious for  $p_N \leq 1$  as then  $|1 - p_n| = 1 - p_n$  and follows from the mean value theorem applied to  $f(x) = (1 - x)^{-1}$  near  $x = 0$  also for  $p_N > 1$ . Collecting all the bounds we have shown that  $\text{I} + \text{II} + \text{III} \leq e^{-N\delta_N^2}$  with the desired  $P_{\theta_0}^N$ -probability, completing the proof of part (C) of Theorem 5.1.3.

## 5.2 Computational complexity of MCMC in high dimensions

### 5.2.1 Gradient methods for approximately log-concave posteriors

We now consider the problem of generating random samples from a high-dimensional posterior measure

$$\Pi(B|D_N) = \frac{\int_B e^{\ell_N(\theta)} d\Pi(\theta)}{\int_{E_D} e^{\ell_N(\theta)} d\Pi(\theta)}, \quad B \subseteq E_D \text{ measurable}, \quad (5.18)$$

arising from data (1.9) with log-likelihood function (1.13) and a general Gaussian  $N(0, \Sigma)$  prior  $\Pi$  of density  $\pi$  on  $\mathbb{R}^D \simeq E_D \subseteq \Theta$ . We specifically have the prior (5.4) in mind where  $\Sigma = \text{diag}(\lambda_j^\alpha : j \leq D)$ , but the results in this subsection hold for general Gaussian priors with positive definite covariance matrix  $\Sigma \in \mathbb{R}^{D \times D}$  and do not depend on the particular choice of  $E_D$  from (5.3) either.

We use the stochastic gradient method (1.16) obtained from discretisation of the  $D$ -dimensional Langevin diffusion (1.15) with drift vector field  $\nabla(\tilde{\ell}_N + \log \pi)$  based on the *surrogate* likelihood function  $\tilde{\ell}_N$  from (5.8), centred at  $\theta_{\text{init}}$  from (5.6). More precisely, for *stepsize*  $\gamma > 0$  and auxiliary variables  $\xi_k \stackrel{\text{i.i.d.}}{\sim} N(0, I_{D \times D})$ , define a Markov chain

$$\begin{aligned} \vartheta_0 &= \theta_{\text{init}}, \\ \vartheta_{k+1} &= \vartheta_k + \gamma[\nabla\tilde{\ell}_N(\vartheta_k) - \Sigma^{-1}\vartheta_k] + \sqrt{2\gamma}\xi_{k+1} \quad \text{for } k = 0, 1, \dots \end{aligned} \quad (5.19)$$

Probabilities and expectations with respect to the law of this Markov chain (random only through the  $\xi_k$ , conditional on the data  $D_N$ ) will be denoted by  $\mathbf{P}_{\theta_{\text{init}}}$ ,  $\mathbf{E}_{\theta_{\text{init}}}$  respectively. The invariant measure of the underlying continuous time Langevin diffusion equals the *surrogate posterior distribution*  $\tilde{\Pi}(\cdot|D_N)$  with density  $\tilde{\pi}(\cdot|D_N)$  given in (5.11).

The idea of the above algorithm is that if we can initialise into the region  $\mathcal{B}$  of local average log-concavity from (5.2), then we can force the gradient method back to  $\mathcal{B}$  whenever it exists that region, by virtue of the dynamics of the ‘surrogate’ vector field  $\nabla\tilde{\ell}_N$  (cf. Proposition 5.1.2). By construction, the dynamics of the above MCMC scheme coincides exactly with the standard Langevin method targeting  $\Pi(\cdot|D_N)$  in the region  $\mathcal{B}$  where the bulk of the posterior mass lies.

The mixing time of this Markov chain towards its (with high probability, log-concave) surrogate target measure measured in Wasserstein distance can be controlled

by high-dimensional Markov chain theory (see Appendix B.3), and Theorem 5.1.3 controls the approximation error induced by the surrogate approximation. Formally, the results that follow will apply the results from Appendix B.3 on an intersection

$$\mathcal{E} := \mathcal{E}_{\text{conv}} \cap \mathcal{E}_{\text{wass}}(\rho), \quad (5.20)$$

of frequentist ‘high probability’ events (i.e., measurable subsets of  $(V \times \mathcal{X})^N$ ): First we restrict to the ‘curvature’ event  $\mathcal{E}_{\text{conv}}$  from (5.1) controlled in Theorem 3.2.3 and we further assume that the Wasserstein distance  $W_2$  between  $\tilde{\Pi}(\cdot|D_N)$  and  $\Pi(\cdot|D_N)$  can be bounded, specifically, for any  $\rho > 0$ , let us define the event

$$\mathcal{E}_{\text{wass}}(\rho) := \left\{ W_2^2(\Pi(\cdot|D_N), \tilde{\Pi}(\cdot|D_N)) \leq \frac{\rho}{2} \right\}.$$

See Theorem 5.1.3 (C) for sufficient conditions for this to be the case with ‘small’  $\rho$  decaying exponentially in  $N$ .

Our first result consists of a global Wasserstein approximation of  $\Pi(\cdot|D_N)$  by the law  $\mathcal{L}(\vartheta_k)$  on  $\mathbb{R}^D$  of the  $k$ -th iterate  $\vartheta_k$  arising from (5.19). Note that the event  $\mathcal{E}$  from Theorem 3.2.3 on which it holds implicitly stipulates choices of the constants  $\kappa_i$ ,  $i = 0, 1, 2$ , describing analytical properties of the average log-likelihood function.

**Theorem 5.2.1** (Wasserstein mixing time). *Suppose that  $D, N \in \mathbb{N}$  are such that  $D \leq \mathcal{R}_N$  with  $\mathcal{R}_N = CND^{-2\kappa_0-4\kappa_2}$  from (3.32) and let  $K$  be as in (5.9). Further, if  $\lambda_{\min}(\cdot)$  and  $\lambda_{\max}(\cdot)$  define the smallest and largest eigenvalues of a symmetric matrix, respectively, define the constants*

$$m := \frac{Nc_0D^{-\kappa_0}}{2} + \lambda_{\min}(\Sigma^{-1}), \quad \Lambda := 7K + \lambda_{\max}(\Sigma^{-1}).$$

*Then for any  $0 < \gamma \leq 1/\Lambda$  and any  $\rho > 0$ , the algorithm  $(\vartheta_k : k \geq 0)$  from (5.19) initialised at  $\theta_{\text{init}}$  and satisfying (5.6) satisfies, on the event  $\mathcal{E}$  from (5.20) and for all  $k \geq 0$ ,*

$$W_2^2(\mathcal{L}(\vartheta_k), \Pi(\cdot|D_N)) \leq \rho + b(\gamma) + 4\left(\tau(\Sigma, R) + \frac{D}{m}\right)\left(1 - \frac{\gamma m}{2}\right)^k, \quad (5.21)$$

*where, for some universal constants  $c_1, c_2 > 0$ , any  $R \geq \|\theta_{0,D}\|_{E_D}$  and  $\kappa(\Sigma) = \lambda_{\max}(\Sigma)/\lambda_{\min}(\Sigma)$ ,*

$$b(\gamma) = c_1 \left[ \frac{\gamma D \Lambda^2}{m^2} + \frac{\gamma^2 D \Lambda^4}{m^3} \right], \quad \tau(\Sigma, R) = c_2 \kappa(\Sigma) [1 + \mathbf{r}^2 + R^2].$$

*Proof.* For any  $\theta, \bar{\theta} \in E_D$ , we have for the log-prior density that

$$\begin{aligned} \|\nabla \log \pi(\theta) - \nabla \log \pi(\bar{\theta})\|_{E_D} &= \|\Sigma^{-1}(\theta - \bar{\theta})\|_{E_D} \leq \lambda_{\max}(\Sigma^{-1})\|\theta - \bar{\theta}\|_{E_D}, \\ \lambda_{\min}(-\nabla^2 \log \pi(\theta)) &\geq \lambda_{\min}(\Sigma^{-1}), \end{aligned}$$

and for the likelihood surrogate  $\tilde{\ell}_N$ , by Proposition 5.1.2, and on the event  $\mathcal{E}$  from (5.20), that

$$\begin{aligned} \|\nabla \tilde{\ell}_N(\theta) - \nabla \tilde{\ell}_N(\bar{\theta})\|_{E_D} &\leq 7K\|\theta - \bar{\theta}\|_{E_D}, \\ \lambda_{\min}(-\nabla^2 \tilde{\ell}_N(\theta)) &\geq \frac{Nc_0 D^{-\kappa_0}}{2}. \end{aligned}$$

Combining the last two displays, and on the event  $\mathcal{E}$ , we can verify Assumption B.3.1 for  $-\log d\tilde{\Pi}(\cdot|D_N)$  with constants

$$m = \frac{Nc_0 D^{-\kappa_0}}{2} + \lambda_{\min}(\Sigma^{-1}), \quad \Lambda = 7K + \lambda_{\max}(\Sigma^{-1}).$$

We may thus apply Proposition B.3.2 to obtain

$$\begin{aligned} W_2^2(\mathcal{L}(\vartheta_k), \Pi(\cdot|D_N)) &\leq 2W_2^2(\Pi(\cdot|D_N), \tilde{\Pi}(\cdot|D_N)) + 2W_2^2(\mathcal{L}(\vartheta_k), \tilde{\Pi}(\cdot|D_N)) \\ &\leq \rho + b(\gamma) + 4\left(1 - \frac{m\gamma}{2}\right)^k \left[\|\theta_{\text{init}} - \theta_{\text{max}}\|_{E_D}^2 + \frac{D}{m}\right], \end{aligned}$$

where  $\theta_{\text{max}}$  denotes the unique maximiser of  $\log d\tilde{\Pi}(\cdot|D_N)$  over  $\mathbb{R}^D$  (which exists on the event  $\mathcal{E}_{\text{conv}}$ , by virtue of strong concavity).

We conclude by an estimate for  $\|\theta_{\text{init}} - \theta_{\text{max}}\|_{E_D}$ . One may expect this norm to be small compared to  $\mathbf{r}$ , but at this point a crude estimate will suffice. From (5.7) we have for all  $\|\theta - \theta_{\text{init}}\|_{E_D} \geq 2\mathbf{r}$  that  $g_{\mathbf{r}}(\theta) \geq (\|\theta - \theta_{\text{init}}\|_{E_D} - \mathbf{r})^2 \geq \frac{1}{4}\|\theta - \theta_{\text{init}}\|_{E_D}^2$ . Thus, for  $C$  from (5.9) and any  $\theta \in E_D$  satisfying

$$\|\theta - \theta_{\text{init}}\|_{E_D}^2 \geq \frac{20}{C} + 4\mathbf{r}^2,$$

and using also the upper bound for  $|\ell_N(\theta)|$  in the definition of  $\mathcal{E}_{\text{conv}}$ , we obtain

$$\begin{aligned} -\tilde{\ell}_N(\theta) = Kg_{\mathbf{r}}(\theta) &\geq CN(c_1 D^{\kappa_1} + 1)\left(1 + \frac{1}{\mathbf{r}^2}\right) \cdot \frac{\|\theta - \theta_{\text{init}}\|_{E_D}^2}{4} \\ &\geq \frac{C}{4}N(c_1 D^{\kappa_1} + 1)\|\theta - \theta_{\text{init}}\|_{E_D}^2 \\ &\geq 5N(c_1 D^{\kappa_1} + 1) \geq -\tilde{\ell}_N(\theta_{\text{init}}). \end{aligned}$$

This implies that necessarily the unique maximiser  $\theta_{\tilde{\ell}}$  of the (on  $\mathcal{E}_{\text{conv}}$ ) strongly concave map  $\tilde{\ell}_N$  over  $\mathbb{R}^D$  satisfies  $\|\theta_{\tilde{\ell}} - \theta_{\text{init}}\|_{\mathbb{R}^D}^2 \leq 20/C + 4\mathbf{r}^2$ . Moreover, in view of the definition of  $\mathcal{B}$ , we have that

$$\|\theta_{\text{init}}\|_{E_D} \leq \|\theta_{\text{init}} - \theta_{0,D}\|_{E_D} + \|\theta_{0,D}\|_{E_D} \leq \|\theta_{\text{init}} - \theta_{0,D}\|_{E_D} + R \leq \mathbf{r} + R,$$

which also allows us to deduce

$$\|\theta_{\tilde{\ell}}\|_{E_D} \leq \|\theta_{\tilde{\ell}} - \theta_{\text{init}}\|_{E_D} + \|\theta_{\text{init}}\|_{E_D} \leq \sqrt{\frac{20}{C}} + 3\mathbf{r} + R.$$

We further have that  $\theta_{\max}^\top \Sigma^{-1} \theta_{\max} \leq \theta_{\tilde{\ell}}^\top \Sigma^{-1} \theta_{\tilde{\ell}}$  (otherwise  $\theta_{\max}$  would not maximise  $\log d\tilde{\Pi}(\cdot|D_N)$ ) and thus, for  $\kappa(\Sigma)$  the condition number of  $\Sigma$ ,

$$\|\theta_{\max}\|_{E_D}^2 \leq \frac{1}{\lambda_{\min}(\Sigma^{-1})} \theta_{\max}^\top \Sigma^{-1} \theta_{\max} \leq \frac{1}{\lambda_{\min}(\Sigma^{-1})} \theta_{\tilde{\ell}}^\top \Sigma^{-1} \theta_{\tilde{\ell}} \leq \kappa(\Sigma) \|\theta_{\tilde{\ell}}\|_{E_D}^2.$$

Combining the preceding displays, the proof is now completed as follows:

$$\begin{aligned} \|\theta_{\max} - \theta_{\text{init}}\|_{E_D}^2 &\lesssim \|\theta_{\max}\|_{E_D}^2 + \|\theta_{\text{init}}\|_{E_D}^2 \\ &\lesssim \kappa(\Sigma) \|\theta_{\tilde{\ell}}\|_{E_D}^2 + \mathbf{r}^2 + R^2 \\ &\lesssim \kappa(\Sigma)[1 + \mathbf{r}^2 + R^2]. \end{aligned} \quad \blacksquare$$

From the previous theorem and concentration inequalities for Markov chains one can obtain the following bound on the computation of posterior statistics by ergodic averages of  $\vartheta_k$  collected after some burn-in time  $J_{\text{in}} \in \mathbb{N}$ . Specifically, if we define, for any  $H: E_D \rightarrow \mathbb{R}$  integrable with respect to  $\Pi(\cdot|D_N)$ , the random variable

$$\hat{\pi}_{J_{\text{in}}}^J(H) = \frac{1}{J} \sum_{k=J_{\text{in}}+1}^{J_{\text{in}}+J} H(\vartheta_k), \quad (5.22)$$

we obtain the following non-asymptotic concentration bound, where  $\|H\|_{\text{Lip}}$  denotes the Lipschitz constant  $\sup_{\theta \neq \theta' \in E_D} [|H(\theta) - H(\theta')| / \|\theta - \theta'\|_{E_D}]$  of  $H$ .

**Theorem 5.2.2** (Lipschitz functionals). *In the setting of the previous theorem, there exist further constants  $c_3, c_4 > 0$  such that for any  $\rho > 0$ , any burn-in period*

$$J_{\text{in}} \geq \frac{c_3}{m\gamma} \times \log\left(1 + \frac{1}{\rho + b(\gamma)} + \tau(\Sigma, R) + \frac{D}{m}\right), \quad (5.23)$$

any  $J \in \mathbb{N}$ , any Lipschitz function  $H: \mathbb{R}^D \rightarrow \mathbb{R}$ , any  $t \geq \sqrt{8} \|H\|_{\text{Lip}} \sqrt{\rho + b(\gamma)}$  and on the event  $\mathcal{E}$  from (5.20), we have

$$\mathbf{P}_{\theta_{\text{init}}}(|\hat{\pi}_{J_{\text{in}}}^J(H) - E^\Pi[H|D_N]| \geq t) \leq 2 \exp\left(-c_4 \frac{t^2 m^2 J \gamma}{\|H\|_{\text{Lip}}^2 (1 + 1/(mJ\gamma))}\right).$$

*Proof.* For any  $t \geq 0$  and any Lipschitz function  $H: E_D \rightarrow \mathbb{R}$ , we have

$$\begin{aligned} &\mathbf{P}_{\theta_{\text{init}}}(|\hat{\pi}_{J_{\text{in}}}^J(H) - E^\Pi[H|D_N]| \geq t) \\ &\leq \mathbf{P}_{\theta_{\text{init}}}(|\hat{\pi}_{J_{\text{in}}}^J(H) - \mathbf{E}_{\theta_{\text{init}}}[\hat{\pi}_{J_{\text{in}}}^J(H)]| \geq t - |\mathbf{E}_{\theta_{\text{init}}}[\hat{\pi}_{J_{\text{in}}}^J(H)] - E^\Pi[H|D_N]|). \end{aligned} \quad (5.24)$$

To further estimate the right-hand side, note that for  $c_3$  large enough and any  $k \geq J_{\text{in}}$ , by (5.23) and Theorem 5.2.1, we have

$$W_2^2(\mathcal{L}(\vartheta_k), \Pi(\cdot|D_N)) \leq 2(\rho + b(\gamma)).$$

Noting that (B.9) in fact holds for any probability measure  $\mu$  and thus in particular for  $\mu = \Pi(\cdot|D_N)$ , it follows that for any Lipschitz function  $H: E_D \rightarrow \mathbb{R}$ ,

$$(\mathbf{E}_{\theta_{\text{init}}}[\hat{\pi}_{J_{\text{in}}}^J(H)] - E^\Pi[H|D_N])^2 \leq 2\|H\|_{\text{Lip}}^2(\rho + b(\gamma)).$$

Thus if  $t \geq 0$  satisfies our hypothesis, then applying Proposition B.3.3 to both  $H$  and  $-H$  yields that the right-hand side in (5.24) is further bounded by

$$\mathbf{P}_{\theta_{\text{init}}}\left(|\hat{\pi}_{J_{\text{in}}}^J(H) - \mathbf{E}_{\theta_{\text{init}}}[\hat{\pi}_{J_{\text{in}}}^J(H)]| \geq \frac{t}{2}\right) \leq 2 \exp\left(-c \frac{t^2 m^2 J \gamma}{\|H\|_{\text{Lip}}^2(1 + 1/(mJ\gamma))}\right). \blacksquare$$

From the last theorem one can obtain as a direct consequence the following guarantee for computation of the posterior mean  $E^\Pi[\theta|D_N]$  by the ergodic average accrued along the Markov chain.

**Corollary 5.2.3.** *In the setting of Theorem 5.2.2, if we define the ergodic average*

$$\bar{\theta}_{J_{\text{in}}}^J = \frac{1}{J} \sum_{k=J_{\text{in}}+1}^{J_{\text{in}}+J} \vartheta_k,$$

then on the event  $\mathcal{E}$  from (5.20) and for  $t \geq \sqrt{8}\sqrt{\rho + b(\gamma)}$ , we have for some constant  $c_5 > 0$  that

$$\mathbf{P}_{\theta_{\text{init}}}\left(\|\bar{\theta}_{J_{\text{in}}}^J - E^\Pi[\theta|D_N]\|_{E_D} \geq t\right) \leq 2D \exp\left(-c_5 \frac{t^2 m^2 J \gamma}{D(1 + 1/(mJ\gamma))}\right).$$

*Proof.* We first estimate the probability to be bounded by

$$\mathbf{P}_{\theta_{\text{init}}}\left(\|\bar{\theta}_{J_{\text{in}}}^J - \mathbf{E}_{\theta_{\text{init}}}[\bar{\theta}_{J_{\text{in}}}^J]\|_{E_D} \geq t - \|\mathbf{E}_{\theta_{\text{init}}}[\bar{\theta}_{J_{\text{in}}}^J] - E^\Pi[\theta|D_N]\|_{E_D}\right).$$

Next, for any  $k \geq 1$ , let  $\nu_k$  denote an optimal coupling between  $\mathcal{L}(\vartheta_k)$  and  $\Pi(\cdot|D_N)$  (cf. [132, Theorem 4.1]). Then by Jensen's inequality and the definition of  $W_2$  from (5.12),

$$\begin{aligned} \|\mathbf{E}_{\theta_{\text{init}}}[\bar{\theta}_{J_{\text{in}}}^J] - E^\Pi[\theta|D_N]\|_{E_D}^2 &= \left\| \frac{1}{J} \sum_{k=J_{\text{in}}+1}^{J_{\text{in}}+J} \int_{E_D \times E_D} (\theta - \theta') d\nu_k(\theta, \theta') \right\|_{E_D}^2 \\ &= \sum_{j=1}^D \left( \frac{1}{J} \sum_{k=J_{\text{in}}+1}^{J_{\text{in}}+J} \int_{E_D \times E_D} (\theta_j - \theta'_j) d\nu_k(\theta, \theta') \right)^2 \\ &\leq \frac{1}{J} \sum_{k=J_{\text{in}}+1}^{J_{\text{in}}+J} \int_{E_D \times E_D} \sum_{j=1}^D (\theta_j - \theta'_j)^2 d\nu_k(\theta, \theta') \\ &= \frac{1}{J} \sum_{k=J_{\text{in}}+1}^{J_{\text{in}}+J} W_2^2(\mathcal{L}(\vartheta_k), \Pi[\cdot|D_N]). \end{aligned}$$

Thus we obtain from (5.21), (5.23) (as after (5.24)) that

$$\|E_{\theta_{\text{init}}}[\bar{\theta}_{J_{\text{in}}}^J] - E^\Pi[\theta|D_N]\|_{E_D} \leq \sqrt{2}\sqrt{\rho + b(\gamma)}.$$

Now for any  $j = 1, \dots, d$ , let us write  $H_j: E_D \rightarrow \mathbb{R}$ ,  $\theta \mapsto \theta_j$ , for the  $j$ -the coordinate projection map, of Lipschitz constant 1. Then in the notation (5.22) we can write

$$[\bar{\theta}_{J_{\text{in}}}^J]_j = \hat{\pi}_{J_{\text{in}}}^J(H_j) \quad \text{for } j = 1, \dots, D.$$

For  $t \geq \sqrt{8(\rho + b(\gamma))}$ , by applying Proposition B.3.3 as in the proof of Theorem 5.2.2 as well as a union bound gives

$$\begin{aligned} & \mathbf{P}_{\theta_{\text{init}}}(\|\bar{\theta}_{J_{\text{in}}}^J - E^\Pi[\theta|D_N]\|_{E_D} \geq t) \\ & \leq \mathbf{P}_{\theta_{\text{init}}}\left(\|\bar{\theta}_{J_{\text{in}}}^J - \mathbf{E}_{\theta_{\text{init}}}[\bar{\theta}_{J_{\text{in}}}^J]\|_{E_D} \geq \frac{t}{2}\right) \\ & = \mathbf{P}_{\theta_{\text{init}}}\left(\sum_{j=1}^D [\hat{\pi}_{J_{\text{in}}}^J(H_j) - \mathbf{E}_{\theta_{\text{init}}}[\hat{\pi}_{J_{\text{in}}}^J(H_j)]]^2 \geq \frac{t^2}{4}\right) \\ & \leq \sum_{j=1}^D \mathbf{P}_{\theta_{\text{init}}}\left([\hat{\pi}_{J_{\text{in}}}^J(H_j) - \mathbf{E}_{\theta_{\text{init}}}[\hat{\pi}_{J_{\text{in}}}^J(H_j)]]^2 \geq \frac{t^2}{4D}\right) \\ & \leq 2D \exp\left(-c \frac{t^2 m^2 J \gamma}{D[1 + 1/(mJ\gamma)]}\right). \quad \blacksquare \end{aligned}$$

The two previous results imply that one can compute the posterior mean (or posterior functionals  $E^\Pi[H|D_N]$  with  $\|H\|_{\text{Lip}} \leq 1$ ) within precision  $\varepsilon > 0$  as long as  $\varepsilon \gtrsim \sqrt{\rho}$ . For instance, if  $\gamma$  is chosen as

$$\gamma \simeq \min\left\{\frac{\varepsilon^2 m^2}{D\Lambda^2}, \frac{\varepsilon m^{3/2}}{D^{1/2}\Lambda^2}\right\}, \quad (5.25)$$

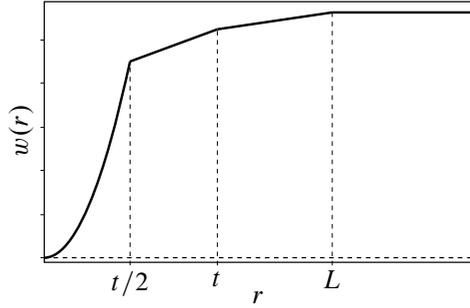
then the overall number of required MCMC iterations  $J_{\text{in}} + J$  depends polynomially on the quantities  $N, D, m^{-1}, \Lambda, \varepsilon^{-1}$ . When the latter three constants also exhibit at most polynomial growth in  $N, D$ , we can deduce that polynomial-time computation of such posterior characteristics is feasible, on the event  $\mathcal{E}$  from (5.20) at computational cost

$$J_{\text{in}} + J = O(N^{b_1} D^{b_2} \varepsilon^{-b_3}) \quad \text{for some } b_1, b_2, b_3 > 0,$$

with  $\mathbf{P}_{\theta_{\text{init}}}$ -probability as close to one as desired.

## 5.2.2 On failure of ‘cold start’ MCMC in high dimensions

In Section 5.3.2 to follow we will use the preceding theorems to show that properly initialised Langevin-type MCMC methods can compute posterior characteristics in



**Figure 5.1.** The graph of  $w(r)$ .

Darcy’s inverse problem in a run-time that scales at most *polynomially* in the key parameters  $D$ ,  $N$  and the desired precision level. This is a non-trivial result in view of the widely accepted ‘folklore’ that sampling from high-dimensional target measures is ‘exponentially hard’ in absence of strong structural assumptions. For instance, a Markov chain may require a long time to ‘escape’ from a local optimum of the posterior density, and if such ‘energy wells’ have depth  $N$ , the algorithmic run-time towards the invariant measure scales exponentially in  $N$ . Also, the mixing time of Markov chains for convergence to equilibrium may scale exponentially in dimension  $D$  without further special assumptions.

In this subsection we give a concrete construction of such a ‘computational hardness phenomenon’ in a context with Gaussian process priors and average ‘population’ log-likelihoods that are not globally log-concave. Specifically, we consider sampling from a probability measure

$$d\bar{\Pi}(\theta) \propto e^{-Nw(\|\theta\|_{\mathbb{R}^D})} d\Pi(\theta), \quad \theta \in \mathbb{R}^D, \quad (5.26)$$

arising from a ‘radial’, expected log-likelihood  $\ell(\theta) = NE_{\theta_0} \ell_N(\theta) = -Nw(\|\theta\|_{\mathbb{R}^D})$  with scalar function  $w$  to be specified below, and a centred Gaussian prior  $\Pi = \text{Law}(\theta)$  on  $\mathbb{R}^D$  given by

$$\theta = (\theta_j)_{j=1}^D, \quad \theta_j = j^{-\alpha/d} g_j \quad \text{for } g_j \stackrel{\text{i.i.d.}}{\sim} N(0, 1) \text{ and } \alpha > d/2. \quad (5.27)$$

In view of (A.14) this is akin to the prior used in (5.4) but without the  $N$ -dependent ‘rescaling’. The latter could be considered as well in the results that follow (by just increasing the constant  $b$  to be introduced by  $(d/2)/(2\alpha + d)$  in the proofs). The results presented further extend without difficulty (see [11]) to the ‘empirical’ versions  $\ell_N$  replacing  $N\ell(\theta)$  but to expedite proofs we consider the ‘average’ case only, which is sufficient to illustrate the main ideas. We will also assume  $\theta_0 = 0$  so that the prior is already centred at the true parameter.

To construct  $w$ , let us set  $b = \alpha/d - 1/2 > 0$  for  $\alpha$  the regularity level from (5.27). Then define (see Figure 5.1), for any  $r \geq 0$ ,

$$\begin{aligned} w(r) &= 4(Tr)^2 1_{[0,t/2)}(r) + \left[ (Tt)^2 + T\left(r - \frac{t}{2}\right) \right] 1_{[t/2,t)}(r) \\ &\quad + \left[ (Tt)^2 + \left(\frac{Tt}{2}\right) + \rho(r-t) \right] 1_{[t,L)}(r) \\ &\quad + \left[ (Tt)^2 + \left(\frac{Tt}{2}\right) + \rho(L-t) \right] 1_{[L,\infty)}(r), \end{aligned}$$

where, for any choice of  $\rho \in (0, 1]$  and  $0 < t_b < 1/2$ , we set

$$t = t_b N^{-b}, \quad T = T_b N^b, \quad (5.28)$$

with  $T_b > 0$  a fixed constant to be chosen, and arbitrary fixed  $L > t$ . Note that such  $w$  is ‘unimodal’ around 0 and the resulting Gibbs measure  $\bar{\Pi}$  is locally (but not globally) log-concave near its unique maximum at  $\theta = 0$ ; the density of  $\bar{\Pi}$  is strictly decreasing along each half-line departing from the origin.

The theorem that follows will imply that in high-dimensional models with  $D \rightarrow \infty$ , a badly initialised ‘local’ MCMC algorithm may never (in ‘finite time’) visit the region where the target measure  $\bar{\Pi}$  places the bulk of its mass, i.e., where it is statistically informative in the sense that the ‘log-likelihood’ term  $\ell(\theta)$  dominates the prior. As a consequence the inference obtained is no better than a random number generator that ignores the observed data. The hardness of sampling from  $\bar{\Pi}$  arises here not from a ‘well’ (local optimum) but from a combination of high dimensionality and lack of ‘signal in the tails’. We refer to [11] for a more detailed evaluation of these negative results. They showcase the need to exploit specific structure of non-linear models when studying algorithms in high-dimensions (as we will do in Section 5.3).

To formulate the main result of this subsection, let us define balls

$$B_r = \{\theta \in \mathbb{R}^D : \|\theta\|_{\mathbb{R}^D} \leq r\}, \quad r > 0,$$

centred at  $\theta_0 = 0$ , as well as the annuli

$$\Theta_{r,\varepsilon} = \{\theta \in \mathbb{R}^D : \|\theta\|_{\mathbb{R}^D} \in (r, r + \varepsilon)\}, \quad \varepsilon > 0. \quad (5.29)$$

We consider MCMC methods given by Markov chains  $(\vartheta_k : k \in \mathbb{N})$  that target the probability measure  $\bar{\Pi}$  from (5.26), more precisely, that have  $\bar{\Pi}$  as their invariant measure. We also assume that their dynamics are ‘local’ in the sense that the step-sizes from  $\vartheta_k$  to  $\vartheta_{k+1}$  at each iterate  $k$  cannot exceed a Euclidean distance  $\eta$ .

**Theorem 5.2.4.** *Consider the probability measure  $\bar{\Pi}$  from (5.26) with prior  $\Pi$  on  $\mathbb{R}^D$  from (5.27) for some  $\alpha > d/2$  and  $w$  as in and before (5.28). Assume  $D/N \simeq \kappa$  as  $N \rightarrow \infty$  for some  $\kappa > 0$ , and set  $b = \alpha/d - 1/2 > 0$ . Then there exist fixed constants  $s_b \in (t_b, 1/2)$  and  $T_b > 0$  for which the following statements hold true:*

- (i) We have  $\bar{\Pi}(B_{s_b N^{-b}}) = 1 - O(e^{-N})$  as  $N \rightarrow \infty$ .
- (ii) There exists  $\varepsilon, \nu > 0$  depending on  $\kappa, \alpha, d$  such that for any Markov chain  $(\vartheta_k : k \in \mathbb{N})$  of step-size  $\eta \in (0, s_b N^{-b})$  and with invariant measure  $\bar{\Pi}$ , we can find an initialisation point  $\vartheta_0 \in \Theta_{N^{-b}, \varepsilon N^{-b}}$  for which the hitting time  $\tau_{B_s}$  for  $\vartheta_k$  to reach  $B_{s_b N^{-b}}$  is lower bounded as

$$\tau_{B_s} = \inf\{k : \vartheta_k \in B_{s_b N^{-b}}\} \geq \exp\{\nu D\},$$

with probability approaching one under the law of the Markov chain.

The proof will show that the result holds for  $\vartheta_0$  drawn at random from an absolutely continuous distribution on  $\Theta_{N^{-b}, \varepsilon N^{-b}}$ , and then employs the probabilistic method to deduce existence of a fixed ‘worst case’ initialiser.

The theorem is stated for Markov chains with ‘deterministic’ step-sizes, whereas many concrete MCMC methods such as ULA have *random* step-sizes, that however do not exceed  $\eta$  with high probability under the randomisation scheme of the algorithm. The previous theorem *does* extend to such settings, and in particular to methods such as pCN, MALA – as is shown in [11] – but we will only prove the above version here for simplicity. Note that while the step sizes have to be sufficiently small, they are of polynomial order in  $1/N$ , whereas the hitting time lower bound scales exponentially in  $N \simeq D$ .

Theorem 5.2.4 does require non-degenerate asymptotics  $D/N \simeq \kappa > 0$ , to permit ‘volumetric’ arguments and ‘unimodal’ target measures  $\bar{\Pi}$ . If in contrast  $\bar{\Pi}$  has a ‘well’ (i.e.,  $w$  has a local optimum away from 0), then the ideas underlying the proof that follows can be adapted to give an exponential lower bound  $\tau_{B_s} \geq e^{\nu N}$  with high probability also when  $D/N \rightarrow 0$ .

**5.2.2.1 Proof of Theorem 5.2.4.** We first show Part (ii), which relies on the following ‘bottleneck’ argument for Markov chains, essentially due to [70]. The result is general but we present it in our setting for  $(\vartheta_k)$  moving between annuli from (5.29) and targeting  $\bar{\Pi}$  from (5.26). Notice that we always have  $\bar{\Pi}(\Theta_{\sigma, \varepsilon}) > 0$  in this case.

**Proposition 5.2.5.** *Consider a Markov chain  $(\vartheta_k : k \in \mathbb{N})$  on  $\mathbb{R}^D$  with invariant measure  $\bar{\Pi}$  and such that*

$$\frac{\bar{\Pi}(\Theta_{s, \eta})}{\bar{\Pi}(\Theta_{\sigma, \varepsilon})} \leq e^{-\nu N} \tag{5.30}$$

for some  $\nu, \varepsilon > 0$  and  $0 < s < \sigma, \eta < \sigma - s$ . Suppose  $\vartheta_0$  is started in  $\Theta_{\sigma, \varepsilon}$ , drawn from the conditional distribution  $\bar{\Pi}(\cdot | \Theta_{\sigma, \varepsilon})$ , and denote by  $\tau_s$  the hitting time of the Markov chain onto  $\Theta_{s, \eta}$ , that is, the number  $\tau_s$  of iterates required until  $\vartheta_k$  visits the set  $\Theta_{s, \eta}$ . Then

$$\Pr(\tau_s \leq K) \leq K e^{-\nu N} \quad \text{for all } K > 0.$$

*Proof.* The probability we wish to bound equals

$$\begin{aligned} & \Pr(\vartheta_k \in \Theta_{s,\eta} \text{ for some } 1 \leq k \leq K | \vartheta_0 \in \Theta_{\sigma,\epsilon}) \\ &= \frac{\Pr(\vartheta_0 \in \Theta_{\sigma,\epsilon}, \vartheta_k \in \Theta_{s,\eta} \text{ for some } 1 \leq k \leq K)}{\bar{\Pi}(\Theta_{\sigma,\epsilon})} \\ &\leq \frac{\sum_{k \leq K} \Pr(\vartheta_k \in \Theta_{s,\eta})}{\bar{\Pi}(\Theta_{\sigma,\epsilon})} \leq K \frac{\bar{\Pi}(\Theta_{s,\eta})}{\bar{\Pi}(\Theta_{\sigma,\epsilon})} \leq K e^{-\nu N}, \end{aligned}$$

since the marginal distributions of  $\vartheta_k$  are all  $\bar{\Pi}$  when  $\vartheta_0 \sim \bar{\Pi}$ .  $\blacksquare$

The last proposition holds ‘on average’ for initialisers  $\vartheta_0 \sim \bar{\Pi}(\cdot | \Theta_{\sigma,\epsilon})$ . But since  $\Pr = E_{\bar{\Pi}(\cdot | \Theta_{\sigma,\epsilon})} \Pr_{\vartheta_0}$ , where  $\Pr_{\vartheta_0}$  is the law of the Markov chain started at  $\vartheta_0$ , the hitting time inequality holds at least for one point in  $\Theta_{\sigma,\epsilon}$ , since we have  $\inf_{\vartheta_0} \Pr_{\vartheta_0} \leq E_{\bar{\Pi}(\cdot | \Theta_{\sigma,\epsilon})} \Pr_{\vartheta_0}$  (‘the probabilistic method’).

In our setting,  $\eta < s_b N^{-b} = s$  with  $t_b < s_b < 1/2$ , and we will apply Proposition 5.2.5 with  $\sigma = N^{-b}$ ,  $\epsilon = \varepsilon N^{-b}$ . If we initialise at  $\vartheta_0 \in \Theta_{N^{-b}, \varepsilon N^{-b}}$ , then under our step-size condition the Markov chain needs to visit  $\Theta_{s_b N^{-b}, \eta}$  in order to reach  $B_{s_b N^{-b}}$ . Therefore the hitting time for the ball  $B_{s_b N^{-b}}$  is lower bounded by the hitting time for the ‘intermediate’ set  $\Theta_{s_b N^{-b}, \eta}$ , and Theorem 5.2.4 (ii) with  $v = \nu/2$  will follow if we can verify (5.30) for these choices of  $s, \eta, \sigma$  and some  $\varepsilon, \nu$ . Taking logarithms, the required inequality (5.30) becomes

$$\frac{1}{N} \log \bar{\Pi}(\Theta_{s_b N^{-b}, \eta}) - \frac{1}{N} \log \bar{\Pi}(\Theta_{N^{-b}, \varepsilon N^{-b}}) \leq -\nu,$$

which in view of (5.26) and the choice of  $w$  from before (5.28) is implied by

$$\frac{1}{N} \log \frac{\Pi(\Theta_{s_b N^{-b}, \eta})}{\Pi(\Theta_{N^{-b}, \varepsilon N^{-b}})} \leq -\nu - \rho N^{-b} (1 + \varepsilon - s_b). \quad (5.31)$$

Annuli farther away from the origin will have larger prior volume and in high dimensions  $D/N \rightarrow \kappa$ . This induces a ‘free energy barrier’, permitting one to verify (5.31), as we now show.

**Lemma 5.2.6.** *For every  $s_b > t_b$  we can choose  $\nu, \varepsilon > 0$ , depending only on  $s_b, \kappa, \alpha, d$ , such that (5.31) holds.*

*Proof.* Lemma 5.2.7 and the hypotheses on  $\eta$  imply

$$\begin{aligned} \Pi(\Theta_{s_b N^{-b}, \eta}) &= \Pi(\|\theta\|_{\mathbb{R}^D} \in (s_b N^{-b}, s_b N^{-b} + \eta)) \\ &\leq \Pi(\|\theta\|_{\mathbb{R}^D} \leq 2s_b N^{-b}) \leq e^{-c_0 N (2s_b + \kappa^{-\alpha/d} (2s_b)^{-\tau/2})^{-\tau}}. \end{aligned}$$

To lower bound  $\Pi(\Theta_{N^{-b}, (1+\varepsilon)N^{-b}})$ , we choose  $\varepsilon$  large enough such that

$$\bar{c}_0 (1 + \varepsilon)^{-\tau} < c_0 (1 + \kappa^{-\alpha/d})^{-\tau},$$

so that for all  $N$  large enough,

$$\begin{aligned} & \Pi(\|\theta\|_{\mathbb{R}^D} \in (N^{-b}, (1+\varepsilon)N^{-b})) \\ &= \Pi(\|\theta\|_{\mathbb{R}^D} \leq (1+\varepsilon)N^{-b}) - \Pi(\|\theta\|_{\mathbb{R}^D} \leq N^{-b}) \\ &\geq e^{-\bar{c}_0(1+\varepsilon)^{-\tau}N} - e^{-c_0(1+\kappa^{-\alpha/d})^{-\tau}N} \geq e^{-2\bar{c}_0(1+\varepsilon)^{-\tau}N}. \end{aligned}$$

The required bound (5.31) will thus follow if

$$c_0(2s_b + \kappa^{-\alpha/d}(2s_b)^{-\tau/2})^{-\tau} \geq 2\bar{c}_0(1+\varepsilon)^{-\tau} + \nu + \rho N^{-b}(1+\varepsilon - s_b). \quad (5.32)$$

If we define  $\nu$  to equal to 1/2 of the left-hand side of the last display, then (5.32) will follow for any  $s_b \leq 1/2$ ,  $\rho \leq 1$  and the given  $\kappa, \alpha, d$ , by choosing  $\varepsilon$  large enough and whenever  $N$  is large enough. ■

To complete the proof of the theorem we now show Part (i) for some fixed choice of  $t_b < s_b < 1/2$  and  $T_b$ , let us write

$$\mathbb{G}(A) := \int_A e^{-Nw(\|\theta\|_{\mathbb{R}^D})} d\Pi(\theta)$$

for any measurable subset  $A$  of  $\mathbb{R}^D$ . Using Lemma 5.2.7, we have

$$\frac{1}{N} \log \mathbb{G}(B_{t_b N^{-b}/2}) \geq -w\left(\frac{t_b N^{-b}}{2}\right) - \bar{c}_0\left(\frac{t_b}{2}\right)^{-\tau}$$

as well as

$$\frac{1}{N} \log \mathbb{G}(B_{s_b N^{-b}}^c) \leq -w(s_b N^{-b}) + \frac{1}{N} \log \Pi(B_{s_b N^{-b}}^c).$$

Now  $\Pi(B_{s_b N^{-b}}^c) \rightarrow 1$  by Lemma 5.2.7, and hence by the choice of  $w$  from before (5.28),

$$\begin{aligned} \frac{1}{N} \log \frac{\mathbb{G}(B_{t_b N^{-b}/2})}{\mathbb{G}(B_{s_b N^{-b}}^c)} &\geq -(T_b t_b)^2 - c_0\left(\frac{t_b}{2}\right)^{-\tau} \\ &\quad + \left[ (T_b t_b)^2 + \left(\frac{T_b t_b}{2}\right) + \rho N^{-b}(s_b - t_b) \right] + \frac{1}{N} \log \Pi(B_s^c) \\ &\geq \frac{T_b t_b}{2} - c_0\left(\frac{t_b}{2}\right)^{-\tau} + o(1), \end{aligned}$$

as  $N \rightarrow \infty$ . Now, for  $t_b < s_b$  fixed, we can choose  $T_b$  large enough such that the last quantity exceeds 1. Therefore

$$\frac{\mathbb{G}(B_{t_b N^{-b}/2})}{\mathbb{G}(B_{s_b N^{-b}}^c)} \geq e^N \times (1 + o(1)).$$

For  $M_N = \{\theta : t_b N^{-b}/2 < \|\theta\|_{\mathbb{R}^D} \leq s_b N^{-b}\}$  this further implies

$$Z \equiv \frac{\mathbb{G}(B_{t_b N^{-b}/2}) + \mathbb{G}(M_N)}{\mathbb{G}(B_{s_b N^{-b}}^c)} \geq e^N \times (1 + o(1)),$$

and Part (i) then follows from this bound and

$$\begin{aligned}\bar{\Pi}(B_{s_b N^{-b}}) &= \frac{\mathbb{G}(B_{t_b N^{-b}/2}) + \mathbb{G}(M_N)}{\mathbb{G}(B_{t_b N^{-b}/2}) + \mathbb{G}(M_N) + \mathbb{G}(B_{s_b N^{-b}}^c)} \\ &= \frac{\mathbb{G}(B_{t_b N^{-b}/2}) + \mathbb{G}(M_N)}{(\mathbb{G}(B_{t_b N^{-b}/2}) + \mathbb{G}(M_N))(1 + Z^{-1})} = \frac{1}{1 + Z^{-1}}.\end{aligned}$$

We conclude this section with the proof of the following key volumetric lemma on the small-deviation landscape of ellipsoidally supported high-dimensional Gaussian measures.

**Lemma 5.2.7.** *Let  $\Pi$  be the Gaussian prior on  $\mathbb{R}^D$  arising from (5.27). Fix  $z > 0$  and  $\kappa > 0$ , and set*

$$b = \frac{\alpha}{d} - \frac{1}{2}, \quad \tau = \frac{1}{b} = \frac{2d}{2\alpha - d}.$$

*Then if  $D/N \simeq \kappa > 0$ , there exist constants  $\bar{c}_0 > c_0$  (depending on  $b, \kappa$ ) such that for all  $N (\geq N_0(z, b))$  large enough,*

$$c_0(z + \kappa^{-\alpha/d} z^{-\tau/2})^{-\tau} \leq -\frac{1}{N} \log \Pi(\|\theta\|_{\mathbb{R}^D} \leq zN^{-b}) \leq \bar{c}_0 z^{-\tau}. \quad (5.33)$$

*Proof.* Consider first the ‘untruncated prior’ (5.27) with series extending to infinity, convergent almost surely in the space  $\ell_2$  with norm  $\|\theta\|_{\ell_2}^2 = \sum_{j=1}^{\infty} \theta_j^2$ , since  $\alpha > d/2$ . Using Proposition A.3.1, Remark A.3.2 and Theorem B.1.1 with  $a = 2d/(2\alpha - d)$ , Banach space  $B = \ell_2$  and RKHS

$$\mathcal{H} = h^\alpha := \left\{ \theta : \|\theta\|_{h^\alpha}^2 = \sum_j j^{2\alpha/d} \theta_j^2 < \infty \right\},$$

we deduce the two-sided estimate

$$-\log \Pi(\|\theta\|_{\ell_2} \leq \gamma) \simeq \gamma^{-2d/(2\alpha-d)} = \gamma^{-\tau} \quad \text{as } \gamma \rightarrow 0.$$

Here, restricting to  $\gamma \in (0, 1)$ , the two-sided equivalence constants depend only on  $\alpha, d$ . Setting

$$\gamma = zN^{-b}, \quad z > 0,$$

and noting that  $b\tau = 1$ , we hence obtain that for some constants  $c_l, c_u > 0$ ,

$$e^{-c_l z^{-\tau} N} \leq \Pi(\|\theta\|_{\ell_2} \leq zN^{-b}) \leq e^{-c_u z^{-\tau} N} \quad \text{for any } z > 0. \quad (5.34)$$

For the projected prior on  $\mathbb{R}^D$ , we have, by set inclusion and projection,

$$\Pi(\|\theta\|_{\mathbb{R}^D} \leq zN^{-b}) \geq \Pi(\|\theta\|_{L^2} \leq zN^{-b}),$$

and hence it only remains to show the first inequality in equation (5.33). The Gaussian

isoperimetric theorem [60, Theorem 2.6.12] and (5.34) imply that for  $m \geq 4\sqrt{c_l}$  and some  $c > 0$ , we have (with  $\Phi$  denoting the cumulative distribution function for  $N(0, 1)$ )

$$\begin{aligned} \Pi(\theta = \theta_1 + \theta_2, \|\theta_1\|_{\ell_2} \leq zN^{-b}, \|\theta_2\|_{h^\alpha} \leq mz^{-\tau/2}\sqrt{N}) \\ \geq \Phi(\Phi^{-1}(\Pi(\{\theta : \|\theta\|_{\ell_2} \leq zN^{-b}\})) + mz^{-\tau/2}\sqrt{N}) \\ \geq \Phi(-\sqrt{2c_l}z^{-\tau/2}\sqrt{N} + mz^{-\tau/2}\sqrt{N}) \geq 1 - e^{-cz^{-\tau}N}. \end{aligned}$$

Then if the event in the last probability is denoted by  $I$ , we have

$$\Pi(\|\theta_D\|_{\mathbb{R}^D} \leq zN^{-b}) \leq \Pi(\|\theta_D\|_{\mathbb{R}^D} \leq zN^{-b}, I) + e^{-cz^{-\tau}N}.$$

On  $I$ , if  $D/N \rightarrow \kappa > 0$  and by the usual tail estimate for vectors  $\theta_2 \in h^\alpha$ , we have for some  $c' > 0$  the bound

$$\|\theta - \theta_D\|_{\ell_2} \leq \|\theta_1\|_{\ell_2} + c'D^{-\alpha/d}z^{-\tau/2}\sqrt{N} \leq zN^{-b} + c'\kappa^{-\alpha/d}z^{-\tau/2}N^{-b},$$

so that for any  $z > 0$ ,

$$\begin{aligned} \Pi(\|\theta_D\|_{\mathbb{R}^D} \leq zN^{-b}) &\leq \Pi(\|\theta\|_{\ell_2} \leq zN^{-b} + \|\theta - \theta_D\|_{\ell_2}, I) + e^{-cz^{-\tau}N} \\ &\leq \Pi(\|\theta\|_{\ell_2} \leq (2z + c'\kappa^{-\alpha/d}z^{-\tau/2})N^{-b}) + e^{-cz^{-\tau}N} \\ &\leq e^{-c_u(2z + c'\kappa^{-\alpha/d}z^{-\tau/2})^{-\tau}N} + e^{-cz^{-\tau}N}, \end{aligned}$$

and hence the lemma follows by appropriate choice of  $c_0 > 0$ . ■

### 5.3 Application to PDE models

We now explain how the ideas from this chapter apply to Darcy's problem, that is, the PDE (1.2) with diffusion operator  $\mathcal{L}_{f_\theta}$  from (1.3) resulting in the forward map  $\mathcal{G}$  from (2.8). The theory for the Schrödinger equation is similar and discussed in the notes (see also Exercise 5.4.5) along with further relevant examples.

We first need to derive some analytical properties of the forward map  $\mathcal{G}: \Theta \rightarrow L^2_\lambda(\mathcal{X})$  from (2.8) restricted to the finite-dimensional approximation space  $E_D$  from (5.3) arising as the eigen-spaces  $E_D$  of the Dirichlet Laplacian on  $\mathcal{X}$ . As ambient parameter space  $\Theta$  we take  $\tilde{H}^\alpha(\mathcal{X}) \subset H_0^\alpha(\mathcal{X})$  from (A.15) for some  $\alpha > 0$  to be specified – these spaces contain  $E_D$  for all  $D$ .

The key result will be Theorem 5.3.2 which verifies the average local curvature condition (3.30) for this forward map on  $E_D$ . It is based on the ideas from Section 3.2 and the observation that the stability condition (3.27) for  $D\mathcal{G}_\theta$  indeed holds for  $\theta \in \Theta$  by virtue of (a) the isomorphism properties of general elliptic operators between Sobolev spaces and their characterisation in the basis spanning  $E_D$  and

(b) the stability Lemma 2.1.6. Note that all  $\theta \in E_D \subset \tilde{H}^\alpha(\mathcal{X})$  vanish at the boundary  $\partial\mathcal{X}$ , as is required to apply Lemma 2.1.6. In the present context this is essential as without this restriction  $D\mathcal{G}_\theta$  may fail to be injective (recall Exercise 3.4.5). Our construction avoids the counter-examples from Section 4.2.2 – for the ‘non-asymptotic’ theorems here we never need to consider the  $L^2$ -completion of our tangent (or parameter) space  $\Theta$ , nor does the range of the adjoint of  $D\mathcal{G}_\theta$  play a role, as was the case in the ‘asymptotic’ Theorems 3.1.5, 4.2.2 and 4.2.3.

### 5.3.1 Darcy’s problem on the eigen-spaces of the Laplacian

To check Conditions 3.2.1 and 3.2.2 we require bounds on the gradient and Hessian of the real-valued map  $\mathcal{G}^x(\theta) \equiv \mathcal{G}(\theta)(x)$  with respect to the variable  $\theta \in E_D \simeq \mathbb{R}^D$  (for fixed  $x \in \mathcal{X}$ ).

We have already computed the  $L^2_\lambda$ -linearisation  $D\mathcal{G}_{\theta_0}$  of  $\mathcal{G}$  at  $\theta_0$  in Theorem 3.3.2, and the proof there can be upgraded to  $\|\cdot\|_\infty$ -differentiability if the perturbations  $h$  are bounded in sufficiently strong Sobolev norms. We restrict to  $d \leq 3$  for simplicity – this permits a straightforward application of the Sobolev imbedding theorem  $H^2 \subset L^\infty$  from (A.4), removing some technicalities in the proof of the general case.

**Proposition 5.3.1.** *Let  $d \leq 3$ . For  $\mathcal{G}$  from (2.8) and  $D\mathcal{G}_\theta$  from (3.46), let  $\theta, h$  belong to  $\Theta = \tilde{H}^\alpha$ ,  $\alpha > 1 + d/2$ . Then we have, as  $\|h\|_{H^\alpha} \rightarrow 0$ , that*

$$\|\mathcal{G}(\theta + h) - \mathcal{G}(\theta) - D\mathcal{G}_\theta[h]\|_\infty = O(\|h\|_{H^\alpha}^2).$$

In particular, the map  $\mathcal{G}^x = \mathcal{G}(\cdot)(x): E_D \rightarrow \mathbb{R}$  has gradient

$$v^\top \nabla \mathcal{G}^x(\theta) = -\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})](x) \quad \text{for } x \in \mathcal{X} \text{ and } v \in E_D,$$

where  $\mathcal{L}_{f_\theta}^{-1}$  is the inverse of the operator  $\mathcal{L}_{f_\theta}$  from (1.3) for Dirichlet boundary conditions (i.e., equal to  $L_{\gamma, V}^{-1}$  given in (A.10) with  $\gamma = f_\theta$ ,  $V = 0$ ).

*Proof.* Arguing just as in the proof of Theorem 3.3.2 and using the Sobolev imbedding  $H^2 \subset L^\infty$  with  $d \leq 3$  (see (A.4)), we obtain

$$\begin{aligned} \|\mathcal{G}(\theta + h) - \mathcal{G}(\theta) - D\mathcal{G}_\theta[h]\|_\infty &\lesssim \|\mathcal{L}_{f_{\theta+h}}^{-1}[R_1(h)]\|_{H^2} + \|\mathcal{L}_{f_{\theta+h}}^{-1}[R_2(h)]\|_{H^2} \\ &\lesssim \|R_1(h)\|_{L^2} + \|R_2(h)\|_{L^2}, \end{aligned}$$

where we have used (A.31) and that  $f_\theta, f_{\theta+h}$  are bounded in  $C^1$  again by the Sobolev imbedding. The right-hand side is shown to be  $O(\|h\|_{C^1}^2) = O(\|h\|_{H^\alpha}^2)$  by appropriate modifications of the arguments from the proof of Theorem 3.3.2, see Exercise 5.4.3 for details. The formula for the gradient then follows directly since  $v \in E_D \subset H^\alpha$ . ■

The previous proof can be iterated (using also the chain and product rule) to show that  $\mathcal{G}^x$  is in fact smooth on  $E_D$ . One obtains the formula for ‘second derivatives’  $D^2\mathcal{G}_\theta$  acting as bilinear forms

$$\begin{aligned} D^2\mathcal{G}_\theta[v_1, v_2] &= -\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v_1 v_2 \nabla u_{f_\theta})] \\ &\quad + \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v_1 \nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v_2 \nabla u_{f_\theta})])] \\ &\quad + \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v_2 \nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v_1 \nabla u_{f_\theta})])] \end{aligned}$$

for  $v_1, v_2 \in E_D$ . In particular, for  $v \in E_D$  and  $x \in \mathcal{X}$ , we have

$$\begin{aligned} v^\top \nabla^2 \mathcal{G}^x(\theta) v &= -\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v^2 \nabla u_{f_\theta})](x) \\ &\quad + 2\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})])](x). \end{aligned} \quad (5.35)$$

The details are only of a technical nature and left to Exercise 5.4.3.

**5.3.1.1 Average curvature for Darcy’s problem on  $E_D$ .** The following lemma checks the local curvature hypothesis (3.30) for the forward map from (2.8) if the natural ‘identifiability’ condition (2.15) is satisfied – cf. also (2.17) for sufficient conditions on  $g, h$  ensuring (2.15). A result comparable to the next theorem holds for the Schrödinger model (2.5) as well – see Exercise 5.4.5 – and also for other PDE models discussed in the notes to this section.

To concentrate on the main ideas we assume that  $\theta_0$  is sufficiently regular in a Sobolev sense and do not attempt to optimise the powers of  $D$  in the choice of radius  $\mathbf{r}$  of  $\mathcal{B}_r$  or in (5.36). Neither do we make any claims about the optimality of the hypothesis  $\alpha \geq 5$ .

**Theorem 5.3.2.** *Let  $\ell(\theta) = \ell(\theta, (Y, X))$  be as in (3.29) with  $\mathcal{G}: E_D \rightarrow \mathbb{R}$  from (2.8),  $d \leq 3$ , with  $E_D \subset \Theta$  as in (5.3) for  $\alpha \geq 5$ , and let  $\mathcal{B}_r$  be as in (5.2) with  $\mathbf{r} = rD^{-w}$ , where  $w = 8/d$  and  $r > 0$ . Suppose (2.15) holds for all  $\theta \in \mathcal{B}_r$ . Let  $\theta_0 \in \tilde{H}^5(\mathcal{X})$  satisfy  $\|\theta_0\|_{\tilde{H}^5} \leq S$  for some  $S > 0$ . Then there exist constants  $0 < r_S \leq 1$ ,  $c_1, c_2 > 0$  such that if also  $\|\mathcal{G}(\theta_0) - \mathcal{G}(\theta_{0,D})\|_{L_\lambda^2} \leq c_1 D^{-w}$ , then for all  $D \in \mathbb{N}$  and all  $r \leq r_S$ ,*

$$\inf_{\theta \in \mathcal{B}_r} \lambda_{\min}(E_{\theta_0}[-\nabla^2 \ell(\theta)]) \geq c_2 D^{-6/d}. \quad (5.36)$$

*Proof.* As a preliminary remark we record here that for  $\theta \in \mathcal{B}_r$ , we have

$$\|\theta\|_{H^5} \lesssim \|\theta\|_{\tilde{H}^5} \leq \|\theta_{0,D}\|_{\tilde{H}^5} + \|\theta - \theta_{0,D}\|_{\tilde{H}^5} \leq S + D^{5/d} r D^{-w} \lesssim 1, \quad (5.37)$$

where we have used the estimate  $\|v\|_{\tilde{H}^5} \lesssim D^{5/d} \|v\|_{E_D}$  for  $v \in E_D$ , valid by (A.14), (A.15) below. By (A.4) this implies  $\|\theta\|_\infty$  is uniformly bounded. Hence, by the chain rule, we can further restrict to  $f_\theta = e^\theta$  bounded in  $H^5$  norm and by Proposition A.5.2 then also to  $u_{f_\theta}$  bounded in  $H^6$ -norm. By the Sobolev imbedding (A.4) (with  $d \leq 3$ ),

this implies

$$\|u_{f_\theta}\|_{C^4} + \|f_\theta\|_{C^3} \leq \|u_{f_\theta}\|_{H^6} + \|f_\theta\|_{H^5} \leq B \quad (5.38)$$

for a constant  $B$  that is uniform in  $\mathcal{B}_r$ . The preceding estimates will be used tacitly in the proof.

We have from (3.35)

$$-\nabla^2 \ell(\theta) = \nabla \mathcal{G}^X(\theta) \nabla \mathcal{G}^X(\theta)^\top - (Y - \mathcal{G}^X(\theta)) \nabla^2 \mathcal{G}^X(\theta).$$

Using this, Proposition 5.3.1 and (5.35), we obtain that for any  $v \in E_D$ ,

$$\begin{aligned} v^\top E_{\theta_0}[-\nabla^2 \ell(\theta)]v &= \|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})]\|_{L_\lambda^2}^2 \\ &\quad - \langle \mathcal{G}(\theta_0) - \mathcal{G}(\theta), 2\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})])] \rangle_{L_\lambda^2} \\ &\quad + \langle \mathcal{G}(\theta_0) - \mathcal{G}(\theta), \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v^2 \nabla u_{f_\theta})] \rangle_{L_\lambda^2} \\ &=: \text{I} - \text{II} + \text{III}. \end{aligned} \quad (5.39)$$

We lower bound  $\lambda_{\min}(E_{\theta_0}[-\nabla^2 \ell(\theta)])$  by lower bounding the last quantity by a constant multiple times  $\|v\|_{E_D}^2$  for all  $v \in E_D$  such that  $\|v\|_{E_D} = 1$ .

(a) *Lower bound for I.* This is the key step. By the interpolation inequality (A.5) for Sobolev norms, we have, for any  $0 \neq w \in H^4(\mathcal{X})$ , that

$$\|w\|_{H^2} \lesssim \|w\|_{L^2}^{1/2} \|w\|_{H^4}^{1/2} \implies \|w\|_{L^2} \gtrsim \frac{\|w\|_{H^2}^2}{\|w\|_{H^4}}.$$

We also note that, by (A.15), (1.6), (A.18) and the remarks after (A.16), we have

$$\|w\|_{H^2} \simeq \|w\|_{\tilde{H}^2} \simeq \|\mathcal{L}_{f_\theta} w\|_{L^2}, \quad \text{any } w \in H_0^2(\mathcal{X}).$$

Applying what precedes to  $w = \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})] \in H_0^2(\mathcal{X})$  with  $v \neq 0$  gives, using also Lemma 2.1.6 and Proposition A.5.3 with  $\beta = 3$ ,

$$\begin{aligned} \|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})]\|_{L^2} &\gtrsim \frac{\|\nabla \cdot (e^\theta v \nabla u_{f_\theta})\|_{L^2}^2}{\|\nabla \cdot (e^\theta v \nabla u_{f_\theta})\|_{H^2}} \\ &\gtrsim \frac{e^{-2\|\theta\|_\infty} \|v\|_{L^2}^2}{\|e^\theta\|_{C^3} \|u_{f_\theta}\|_{C^4} \|v\|_{H^3}} \gtrsim D^{-3/d} \|v\|_{L^2}, \end{aligned}$$

in view of the estimate  $\|v\|_{H^3} \lesssim \|v\|_{\tilde{H}^3} \lesssim D^{3/d} \|v\|_{L^2}$  for  $v \in E_D \subset \tilde{H}^3 \subset H^3$ , valid by (A.15) and (A.14). In conclusion, we have proved that for all  $0 \neq v \in E_D$ ,

$$\|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})]\|_{L_\lambda^2}^2 \gtrsim D^{-6/d} \|v\|_{E_D}^2.$$

(b) *Upper bounds for II, III.* For  $\theta \in \mathcal{B}_r$ , the Cauchy–Schwarz inequality and the local  $E_D \rightarrow L_\lambda^2$ -Lipschitz property of  $\mathcal{G}$  (Proposition 2.1.3 with  $\kappa = 0$  and (5.37))

imply that the terms II, III are both bounded by  $(Lr + c_1)D^{-w}$  (i + ii), and where i is bounded by

$$\begin{aligned} \|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})])]\|_{L_\lambda^2} &\lesssim \|e^\theta v \nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})]\|_{L_\lambda^2} \\ &\lesssim \|v\|_\infty \|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})]\|_{H^1} \\ &\lesssim \|v\|_\infty \|e^\theta v \nabla u_{f_\theta}\|_{L^2} \\ &\leq \|v\|_\infty \|v\|_{L^2} \|f_\theta\|_\infty \|u_{f_\theta}\|_{C^1} \\ &\lesssim D^{(1/2)+\epsilon} \|v\|_{L^2}^2 \end{aligned}$$

using (A.32), (A.8), the divergence theorem (A.7) as well as

$$\|v\|_\infty \lesssim \|v\|_{H^{d(1/2+\epsilon)}} \lesssim D^{(1/2)+\epsilon} \|v\|_{L^2}, \quad v \in E_D,$$

again by (A.15) and (A.14); while ii is bounded by

$$\begin{aligned} \|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v^2 \nabla u_{f_\theta})]\|_{L_\lambda^2} &\lesssim \|e^\theta v^2 \nabla u_{f_\theta}\|_{L^2} \\ &\leq \|f_\theta\|_\infty \|u_{f_\theta}\|_{C^1} \|v\|_\infty \|v\|_{L_\lambda^2} \\ &\lesssim D^{(1/2)+\epsilon} \|v\|_{L^2}^2, \end{aligned}$$

using similar arguments.

To conclude, since  $w = 8/d$  and  $d < 4$ , we obtain the overall bound

$$\begin{aligned} \text{I} - \text{II} + \text{III} &\geq A_1 D^{-6/d} \|v\|_{E_D}^2 - A_2 (r + c_1) D^{-w} D^{(1/2)+\epsilon} \|v\|_{E_D}^2 \\ &\geq \frac{A_1}{2} D^{-6/d} \|v\|_{E_D}^2 \end{aligned}$$

for some constants  $A_1, A_2 > 0$ , all  $r, c_1$  small enough, and all  $0 \neq v \in E_D$ , proving the theorem.  $\blacksquare$

Using minor variations of the arguments from the preceding proof, one can verify the second part of Condition 3.2.2.

**Proposition 5.3.3.** *Under the hypotheses of Theorem 5.3.2, the bound (3.31) in Condition 3.2.2 holds for  $\kappa_1 = 0$ .*

*Proof.* We only consider the operator norm of the expected Hessian, the other terms are of smaller order and bounded in a similar way. We can again use the decomposition (5.39) and the previous proof already showed that  $\text{II} + \text{III} \lesssim D^{-w} D^{1/2+\epsilon} \lesssim 1$ . To upper bound I we use (A.32) and (5.38) and the divergence theorem to obtain

$$\sqrt{\text{I}} \leq \bar{c} \|\nabla \cdot (e^\theta v \nabla u_{f_\theta})\|_{(H_0^1)^*} \leq \|f_\theta\|_\infty \|u_{f_\theta}\|_{C^1} \|v\|_{L^2} \leq c \|v\|_{E_D},$$

which after collecting terms implies the result.  $\blacksquare$

**5.3.1.2 Verifying Condition 3.2.1.** There are several terms to check when verifying Condition 3.2.1, all following the same standard proof template. We encounter similar terms as in the proof of Theorem 5.3.2, but now need to bound supremum norms instead of  $L^2_\lambda$ -norms, or, by the Sobolev imbedding  $H^2 \subset L^\infty$  ( $d \leq 3$ ), corresponding  $H^2$ -norms.

For instance, for the gradient  $\nabla \mathcal{G}^x(\theta)$  from Proposition 5.3.1, we can use (A.31), (5.37), (5.38) for  $\theta \in \mathcal{B}_r$  as in Theorem 5.3.2, and (A.15), (A.14) to obtain

$$\begin{aligned} \sup_{x \in \mathcal{X}} |v^\top \nabla \mathcal{G}^x(\theta)| &\lesssim \|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})]\|_{H^2} \lesssim \|e^\theta v \nabla u_{f_\theta}\|_{H^1} \\ &\lesssim \|f_\theta\|_{C^1} \|u_{f_\theta}\|_{C^2} \|v\|_{H^1} \lesssim D^{1/d} \|v\|_{E_D}, \end{aligned}$$

while the second term featuring in the Hessian from (5.35) is of the order

$$\begin{aligned} \|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})])]\|_{H^2} &\lesssim \|e^\theta v \nabla \mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})]\|_{H^1} \\ &\lesssim \|\theta\|_\infty \|v\|_{C^1} \|\mathcal{L}_{f_\theta}^{-1}[\nabla \cdot (e^\theta v \nabla u_{f_\theta})]\|_{H^2} \\ &\lesssim \|v\|_{C^1} \|e^\theta v \nabla u_{f_\theta}\|_{H^1} \\ &\lesssim \|v\|_{C^1} \|v\|_{H^1} \|f_\theta\|_{C^1} \|u_{f_\theta}\|_{C^2} \\ &\lesssim \|v\|_{H^3} \|v\|_{H^1} \lesssim D^{4/d} \|v\|_{E_D}^2, \end{aligned}$$

where we have used the Sobolev imbedding  $H^3 \subset C^1$  for  $d \leq 3$  and again (5.37), (5.38), (A.15), (A.14).

Repeating arguments of this kind one can verify Condition 3.2.1 for Darcy's problem – the ‘critical’ term arises from the Lipschitz constant of the operator norms of  $\nabla^2 \mathcal{G}^x$ , and the details of the proof of the following proposition are left to Exercise 5.4.4.

**Proposition 5.3.4.** *Under the hypotheses of Theorem 5.3.2, Condition 3.2.1 is satisfied for  $\kappa_2 = 7/d$ .*

No claim is made about the optimality of our bound for  $\kappa_2$  – we only wish to exhibit its growth of a polynomial in  $D$  order to apply Theorem 5.1.3.

We have now everything at hand to apply Theorem 5.1.3 to Darcy's problem. The smoothness  $\alpha > 21$  of the model required in the following theorems is not optimal and arises as such because we have repeatedly opted for simpler proofs in the preceding results. Specifically we have not appealed to ‘Schauder-theory’ for  $C^\alpha$ -regularity of solutions to elliptic PDEs (e.g., [59, Section 6]) but only to Sobolev regularity bounds via  $L^2$ -energy estimates. When proofs require  $C^\alpha$ -bounds, the use of the Sobolev imbedding then increases relevant constants. We did this to focus on the main ideas rather than on technicalities. The resulting assumption on  $\alpha$  translates into a growth hypothesis on the dimension  $D$ , and the theorems that follow should hence be construed as holding in ‘moderately high-dimensional’ models – the precise value

of  $\alpha$  is not of central importance. Sharper versions are possible (and discussed in the notes) but do not impact the main insight drawn from the following theorem, that for sufficiently regular models, a log-concave approximation theorem is true for Darcy's problem while the Bernstein–von Mises theorem fails for it.

**Theorem 5.3.5.** *Suppose that  $\theta_0 \in \tilde{H}^\alpha(\mathcal{X})$  for  $\alpha > 21$ ,  $d \leq 3$ , that (5.13) holds with  $\delta_N = N^{-\alpha/(2\alpha+d)}$ , and that (2.15) is satisfied for all  $\theta \in \Theta \equiv \tilde{H}^\alpha$ . Then the forward map  $\mathcal{G}$  from (2.8) satisfies Condition 5.1.1 for  $\mathbf{r} = rD^{-8/d}$  and all small enough  $r > 0$ , spaces  $E_D$  from (5.3),*

$$\eta = \frac{\alpha - 1}{\alpha + 1}, \quad \kappa_0 = \frac{6}{d}, \quad \kappa_1 = 0, \quad \kappa_2 = \frac{7}{d}$$

and some  $c_i > 0$ ,  $i = 0, 1, 2$ , as long as  $D \leq A_1 N^{d/(2\alpha+d)}$  for some  $A_1 < \infty$ . In particular, the log-concave approximation from Theorem 5.1.3 holds true for the chosen constants and the corresponding surrogate posterior from (5.11) and with  $e^{-CND^{-2\kappa_0-4\kappa_2}} = o(1)$ .

*Proof.* Note that Propositions 2.1.3 and 2.1.7 with  $\beta = \alpha$  apply to verify Condition 2.1.1 for  $\kappa = 1$  (and then also  $\kappa = 0$ ) and Condition 2.1.4 with  $\eta = (\alpha - 1)/(\alpha + 1)$ . We take  $\beta = \alpha$  here, permitted in view of Exercise 2.4.3 for this particular prior (5.4). Conditions 3.2.1 and 3.2.2 were checked in the previous subsection, using also that (5.13) verifies the bound on  $\|\mathcal{G}(\theta_0) - \mathcal{G}(\theta_{0,D})\|_{L^2}$  for some  $c_1$  and the choice of  $D$  to follow. Collecting the required bounds, we need  $D^{-w} \gtrsim D^{(6+7)/d} \delta_N$  to verify (5.5), which is satisfied for our choice of  $D, \alpha$ . ■

### 5.3.2 Polynomial time computation of the posterior mean

We can now combine the previous theorem with the computation time bounds from Section 5.2.1 for gradient-based Langevin Markov chains. We only give a result for the polynomial time computability of the posterior mean vector – versions of Theorems 5.2.1 and 5.2.2 can be obtained as well. We consider a ‘warm start’ of the algorithm, assuming an appropriate initialiser  $\theta_{\text{init}}$  into the region  $\mathcal{B}_r$  exists. When  $D \approx N^a$  so that  $\mathbf{r} \simeq N^{-\bar{a}}$  for some  $a, \bar{a} > 0$ , it is itself a non-trivial task to find a polynomial time initialiser – see the notes for discussion.

We recall that ‘polynomial run-time’ in the following theorem refers to the dimension  $D$  of the state space one is using, the ‘informativeness’  $N$  of the posterior surface (5.18), as well as the desired precision level  $\varepsilon$ . Computational complexity is measured in terms of the number of required iterates of the MCMC scheme (5.19), including the burn-in time. Any MCMC step requires the evaluation of the gradient  $\nabla \ell_N$  and hence of  $\mathcal{G}, \nabla \mathcal{G}$ , each amounting to the solution of an elliptic boundary value problem in  $\mathcal{X} \subset \mathbb{R}^d$  (in the case of Darcy's problem). Such an operation is feasible (in polynomial time) by standard numerical PDE solvers, but we do not concern ourselves here

with the computational cost of this operation in itself. The following theorem shows that computational barriers in  $D, N$  such as those described in Section 5.2.2 can be overcome for Darcy's problem by 'warm start' gradient type MCMC methods.

**Theorem 5.3.6.** *Suppose that data  $D_N = (Y_i, X_i)_{i=1}^N$  arise from (1.9) in Darcy's problem with forward map  $\mathcal{G}$  from (2.8),  $d \leq 3$ , and such that (2.15) is satisfied for all  $\theta \in \Theta \equiv \tilde{H}^\alpha$ ,  $\alpha > 21$ . Consider the prior  $\Pi$  on  $E_D \subset \Theta$  given by (5.4) with dimension  $D \leq A_1 N^{d/(2\alpha+d)}$  for some  $A_1 > 0$ . Suppose that (5.13) holds for  $\theta_0 \in \Theta$ . Consider the computation of the mean vector*

$$E^\Pi[\theta|D_N] = \int_{E_D} \theta d\Pi(\theta|D_N) d\theta$$

of the posterior distribution  $\Pi(\cdot|D_N)$  from (5.18). For some arbitrary fixed  $P > 0$ , let the precision level  $\varepsilon$  satisfy  $\varepsilon \geq N^{-P}$ . For an initialiser  $\theta_{\text{init}}$  satisfying  $\|\theta_{\text{init}} - \theta_{0,D}\|_{E_D} \leq D^{-8/d}/(8 \log N)$ , let the MCMC average  $\bar{\theta}_{J_{\text{in}}}^J$  be as in Corollary 5.2.3 with  $\gamma$  as in (5.25) and corresponding choices of  $m, \Lambda$  in Theorem 5.2.1 provided by Theorem 5.3.5. Then there exist positive constants  $a, b_1, b_2, b_3$  such that with  $P_{\theta_0}^N \times \mathbf{P}_{\theta_{\text{init}}}$ -probability at least  $1 - e^{-N^a}$ ,

$$\|\bar{\theta}_{J_{\text{in}}}^J - E^\Pi[\theta|D_N]\|_{E_D} \leq \varepsilon,$$

and with burn-in and run-times  $J_{\text{in}}, J$  of order at most polynomial order,

$$J_{\text{in}} + J = O(N^{b_1} D^{b_2} \varepsilon^{-b_3}). \quad (5.40)$$

*Proof.* The result follows from Theorem 5.3.5 and Corollary 5.2.3 after bounding all relevant constants as discussed after (5.25). Note that the constants involving eigenvalues of the prior inverse covariance  $\Sigma^{-1}$  for our choice  $\Sigma = \text{diag}(\lambda_j^\alpha : j \leq D)$  grow at most polynomially in  $D$  in view of (A.14). ■

Note that the stepsizes  $\gamma$  are of the form  $O(D^{-b_4} N^{-b_5} \varepsilon^{b_6})$  for some  $b_i > 0$  in the previous algorithm.

We can combine the previous computational result with the consistency Theorem 2.3.3 for the posterior mean (or in fact Exercise 2.4.3) to further show that the MCMC averages also recover the ground truth  $\theta_0$  generating the data, that is, with  $P_{\theta_0}^N \times \mathbf{P}_{\theta_{\text{init}}}$ -probability at least  $1 - e^{-N^a}$ , we have

$$\|\bar{\theta}_{J_{\text{in}}}^J - \theta_0\|_{L^2} \leq \varepsilon \quad \text{whenever } \varepsilon \gtrsim \max(N^{-P}, \delta_N^\eta),$$

after a number  $J + J_{\text{in}}$  of iterates of the MCMC chain that grows at most polynomially in relevant quantities (cf. (5.40)). This shows that the ground truth diffusivity  $\theta_0$  in Darcy's problem can be recovered by a constructive polynomial time MCMC algorithm, assuming warm start initialisation as in (5.6) is possible.

## 5.4 Notes

### 5.4.1 Exercises

**Exercise 5.4.1.** Prove Proposition 5.1.2. [Hint: A proof of a more general result can be found in [101, Proposition 3.6].]

**Exercise 5.4.2.** Show that the probability bound in Lemma 1.3.3 can be taken of order  $1 - O(e^{-bN\delta_N^2})$  for some  $b > 0$  if  $K$  is chosen large enough. [Hint: Use (B.4) or see [101, Lemma 4.15].] Further show that (5.17) holds.

**Exercise 5.4.3.** Complete the proof of Proposition 5.3.1 and show further that the map  $\theta \mapsto \mathcal{G}^x(\theta)$  is  $C^3$  on the set  $\mathcal{B} = \mathcal{B}_r$  from Theorem 5.3.2. [See also [4].]

**Exercise 5.4.4.** Prove Proposition 5.3.4. [See also [4].]

**Exercise 5.4.5.** Prove an analogue of Theorem 5.3.2 for the forward map  $\mathcal{G}$  from the Schrödinger model (2.5), with  $w = 4/d$  and lower bound in (5.36) of order  $c_2 D^{-4/d}$ . [Hint: Start from (3.48), or see [101, Lemma 4.7].]

### 5.4.2 Remarks and comments

The main ideas of this section were developed in the recent contribution [101] and in the follow up paper [23], even though the particular proofs given here follow a slightly different route inspired by the very recent article [4] which avoids the analysis of ‘MAP’ estimators from [23, 101] at the expense of slightly stronger conditions on the model parameters in Condition 5.1.1. The mixing time results from Section 5.2.1 follow ideas from [39, 48, 49] for strongly log-concave targets (reviewed in the appendix).

The results in [101] and [23] consider different PDEs than Darcy’s problem, specifically the Schrödinger equation from (2.5) and non-Abelian X-ray transforms from (1.1), where results similar to Theorems 5.3.5 and 5.3.6 are obtained. In particular, these references establish versions of the local average curvature Theorem 5.3.2 for these PDEs (with appropriate choices of  $\kappa_\theta$ ,  $\mathbf{r}$ ,  $E_D$ ). The underlying gradient stability requirement seems more feasible to check than solving information equations (3.16), as highlighted by Darcy’s problem. We also note that the choice of the discretisation spaces  $E_D$  needs to be adapted to the mapping properties of  $D\mathcal{G}_\theta$  and, for instance, for X-ray transforms, an eigen-basis different from the standard Laplacian has to be considered to adjust for boundary behaviour, see [23], building on ideas from [88, 91].

As mentioned above, the proofs of the analogues of Theorem 5.1.3 in [23, 101] are slightly different in that they require a separate convergence analysis of the MAP estimator  $\hat{\theta}_{\text{MAP}}$  (i.e., the maximiser of the posterior density of (5.18) over  $E_D$ ). The benefit of this is that when dealing with term II in the decomposition of the

Wasserstein distance in the proof of Theorem 5.1.3, one can expand the log-posterior around the global maximum  $\hat{\theta}_{\text{MAP}}$  (instead of  $\theta_{0,D}$ ), avoiding the gradient term in (5.16). This allows to weaken the quantitative requirements in Condition 5.1.1 – specifically the factor  $D^{\kappa_0+\kappa_2}$  in the definition of  $\tilde{\delta}_N$  can be replaced by just  $D^{\kappa_0/2}$ , see [23, Condition 3.5] – but at the cost of introducing slightly stronger analytical conditions on  $\mathcal{G}$  required to obtain convergence rates for  $\hat{\theta}_{\text{MAP}}$ . In this context [101] also provide computational guarantees for optimisation based estimators and further clarify that the unimodal ‘surrogate posterior’ actually has  $\hat{\theta}_{\text{MAP}}$  as its mode, as one may expect. While these stronger sets of conditions on  $\mathcal{G}$  can often be checked for concrete PDEs, the proofs presented in the notes here are ‘entirely Bayesian’ and do not require to set up the machinery from  $M$ -estimation [124], [100] to study optimisation based estimators.

Log-concave approximation theorems such as Theorem 5.1.3 provide a conceptual alternative to Bernstein–von Mises approximations studied in Section 4.1 and also to Laplace approximations discussed in Remark 5.1.4. They are useful to establish computational guarantees for gradient-based MCMC methods as in Theorem 5.3.6 – a main challenge here is to show that the MCMC method requires at most a *polynomial* run-time both in  $D$  and  $N$  simultaneously. Results that scale uniformly in  $D$  (but exponentially in  $N$ ) can be obtained by infinite-dimensional Harris-type theorems (see [65]) but these do not give polynomial time guarantees for ‘informative’ posteriors when  $N \rightarrow \infty$ . See [101] for more discussion. We note also that log-concave approximation theorems *appear to fall short* (at least in our current understanding) of providing *statistical* guarantees for posterior-based uncertainty quantification such as those given in Section 4.1.3. The reason for this is (a) that Theorem 5.1.3 is not compatible with the weak convergence arguments relevant in Section 4.1.3 (as no  $N$ -independent limiting measure is provided), and (b) that as such the Wasserstein distance does not dominate the Kolmogorov ‘quantile’ distance for general pairs of  $N$ -dependent probability measures. Thus for purposes of uncertainty quantification, the BvM theorems from Section 4.1 are currently still the only available results, despite their limitations.

The results in Section 5.2.2 about hardness of local MCMC methods are partly inspired by ideas from [15, 16] (see also [10]) laid out in very different high-dimensional statistical physics models. Theorem 5.2.4 presented here is a simplified ‘average log-likelihood’ versions of one of the main results from [11], where also other priors and random step size MCMC methods such as MALA and pCN are explicitly considered.

We finally touch on the issue of initialisation in (5.6), specifically with an eye on data-driven choices for  $\theta_{\text{init}}$  as is relevant in Theorem 5.3.6 on the performance of gradient MCMC. We are not aware of a general strategy for initialisation other than perhaps in settings where the curvature hypothesis in Condition 3.2.2 holds on a fixed ball  $\mathcal{B}$  (of constant in  $D, N$  radius) when a standard grid-search can be used. For the

Schrödinger equation, an ad hoc strategy for initialisation can be provided [101], but for other PDEs this poses novel challenges that are not sufficiently well understood at the time of this writing.



## Appendix A

### Analytical background

In this section, unless mentioned otherwise,  $\mathcal{X}$  is always a bounded domain in  $\mathbb{R}^d$  with smooth boundary and the space  $L^2 = L^2(\mathcal{X}) = L^2_\lambda(\mathcal{X})$  denotes the standard Hilbert space of square integrable functions on  $\mathcal{X}$  with respect to Lebesgue measure  $\lambda$ , with inner product  $\langle \cdot, \cdot \rangle_{L^2}$  and norm  $\| \cdot \|_{L^2} = \sqrt{\langle \cdot, \cdot \rangle_{L^2}}$ . The relation  $a \lesssim b$  denotes an inequality  $a \leq Cb$  that holds up to fixed constants  $C > 0$ , and the corresponding convention is used for  $\gtrsim$ .

#### A.1 Sobolev and related spaces

We collect here without proofs some standard facts about Sobolev spaces that can be found in any of the standard references on the subject, for instance, [86, Chapter 1], [121, Chapter 4], [59, Chapter 7] or [2].

For a multi-index  $i = (i_1, \dots, i_d)$ , let  $D^i$  denote the  $i$ -th (weak) partial derivative operator of order  $|i|$ . Then for an integer  $\alpha \geq 0$ , the Sobolev spaces are

$$H^\alpha(\mathcal{X}) := \{f \in L^2(\mathcal{X}) \mid \text{for all } |i| \leq \alpha, D^i f \text{ exists and } D^i f \in L^2(\mathcal{X})\}, \quad (\text{A.1})$$

normed by

$$\|f\|_{H^\alpha(\mathcal{X})}^2 = \sum_{|i| \leq \alpha} \|D^i f\|_{L^2(\mathcal{X})}^2.$$

For non-integer real values  $\alpha \geq 0$ , one can define  $H^\alpha(\mathcal{X})$  by interpolation. For  $\alpha = 0$  the convention  $H^0(\mathcal{X}) = L^2(\mathcal{X})$  will be used.

The space of bounded and continuous functions on  $\mathcal{X}$  is denoted by  $C(\mathcal{X})$ , equipped with the supremum norm  $\| \cdot \|_\infty$ . For  $\eta \in \mathbb{N}$ , the space of  $\eta$ -times differentiable functions on  $\mathcal{X}$  with (bounded) uniformly continuous derivatives is denoted by  $C^\eta(\mathcal{X})$ , with norm  $\|f\|_{C^\eta(\mathcal{X})} = \sum_{|i| \leq \eta} \|D^i f\|_\infty$ . For  $\eta > 0$ ,  $\eta \notin \mathbb{N}$ , we say  $f \in C^\eta(\mathcal{X})$  if for all multi-indices  $\beta$  with  $|\beta| \leq \lfloor \eta \rfloor$  (the integer part of  $\eta$ ),  $D^\beta f$  exists and is  $(\eta - \lfloor \eta \rfloor)$ -Hölder continuous. The norm on  $C^\eta(\mathcal{X})$  for such  $\eta$  is

$$\|f\|_{C^\eta(\mathcal{X})} = \sum_{\beta: |\beta| \leq \lfloor \eta \rfloor} \|D^\beta f\|_\infty + \sum_{\beta: |\beta| = \lfloor \eta \rfloor} \sup_{x, y \in \mathcal{X}, x \neq y} \frac{|D^\beta f(x) - D^\beta f(y)|}{|x - y|^{\eta - \lfloor \eta \rfloor}}.$$

We also define the set of smooth functions as  $C^\infty(\mathcal{X}) = \bigcap_{\eta > 0} C^\eta(\mathcal{X})$  and its subspace  $C_c^\infty(\mathcal{X})$  of functions compactly supported in  $\mathcal{X}$ . All definitions so far make sense for any open subset  $\mathcal{X} \subseteq \mathbb{R}^d$ .

Now for  $\mathcal{X} \subset \mathbb{R}^d$  a bounded domain with smooth boundary  $\partial\mathcal{X}$ , the completion of  $C_c^\infty(\mathcal{X})$  for the  $H^\alpha$ -norm is denoted by  $H_c^\alpha(\mathcal{X})$ , and this is a closed subspace of  $H^\alpha(\mathcal{X})$  (equal to  $H^\alpha(\mathcal{X})$  only if  $\alpha < 1/2$ ). We also define the spaces  $H_0^\alpha(\mathcal{X})$ ,  $C_0^\infty(\mathcal{X})$  of elements of  $H^\alpha(\mathcal{X})$ ,  $C^\infty(\mathcal{X})$  that vanish at  $\partial\mathcal{X}$  (in a trace sense if necessary). We notice that  $H_0^1(\mathcal{X}) = H_c^1(\mathcal{X})$  but otherwise ( $\alpha > 1$ ) the inclusion  $H_c^\alpha(\mathcal{X}) \subset H_0^\alpha(\mathcal{X})$  is strict as the former spaces require normal derivatives up to order  $\alpha - 1/2$  to vanish at  $\partial\mathcal{X}$ . One also defines topological dual spaces  $H^{-\alpha}(\mathcal{X}) = (H_c^\alpha(\mathcal{X}))^*$  for  $\alpha \geq 0$ . For some results we also need Sobolev spaces  $H^\alpha(\partial\mathcal{X})$  defined similarly on the closed ‘boundary-less’ manifold  $\partial\mathcal{X}$ , as well as  $C^\infty(\partial\mathcal{X}) = \bigcap_{\alpha \geq 0} H^\alpha(\partial\mathcal{X})$ . See, e.g., [121, Chapter 4.3] for details.

All preceding spaces and norms can be defined for vector fields  $f: \mathcal{X} \rightarrow \mathbb{R}^d$  with standard modification of the norms, by requiring each of the coordinate functions  $f_i(\cdot)$ ,  $i = 1, \dots, d$ , to belong to the corresponding space of real-valued maps.

We will repeatedly use the standard multiplicative inequalities

$$\|fg\|_{H^\alpha} \lesssim \|f\|_{H^\alpha} \|g\|_{H^\alpha}, \quad \alpha > d/2, \quad (\text{A.2})$$

$$\|fg\|_{H^\alpha} \lesssim \|f\|_{C^\alpha} \|g\|_{H^\alpha}, \quad \alpha \geq 0, \quad (\text{A.3})$$

for all  $f, g$  in the appropriate function spaces. Also, for any  $\alpha > d/2$  and  $0 < \eta < \alpha - d/2$ , the Sobolev embedding  $H^\alpha \subseteq C^\eta$  holds, with corresponding norm estimates

$$\|f\|_\infty \lesssim \|f\|_{C^\eta} \lesssim \|f\|_{H^\alpha} \quad \text{for all } f \in H^\alpha. \quad (\text{A.4})$$

We also need the following interpolation inequality for Sobolev norms. For all  $\beta_1, \beta_2 \geq 0$  and  $\theta \in [0, 1]$ ,

$$\|u\|_{H^{\theta\beta_1 + (1-\theta)\beta_2}} \lesssim \|u\|_{H^{\beta_1}}^\theta \|u\|_{H^{\beta_2}}^{1-\theta} \quad \text{for all } u \in H^{\beta_1} \cap H^{\beta_2}. \quad (\text{A.5})$$

## A.2 Elliptic second-order differential operators

We now consider a second-order elliptic operator

$$L_{\gamma, V}(u) = \nabla \cdot (\gamma \nabla u) - Vu, \quad u \in C^\infty(\mathcal{X}),$$

where  $\gamma: \mathcal{X} \rightarrow [\gamma_{\min}, \infty)$  is a diffusion coefficient bounded from below by a scalar  $\gamma_{\min} > 0$ ,  $V: \mathcal{X} \rightarrow [0, \infty)$  is a non-negative potential, and  $\nabla \cdot v = \sum_{j=1}^d \partial v / \partial v_j$  denotes the divergence of a vector field  $v \in H^1(\mathcal{X})$ . We will assume throughout this section that  $\gamma \in H^\beta(\mathcal{X})$  for some  $\beta > 1 + d/2$  so that in particular  $\gamma \in C^\eta(\mathcal{X})$  for  $\eta > 1$  by (A.4). We also assume that  $V \in H^{\beta-1}(\mathcal{X}) \subset C^{\eta-1}(\mathcal{X})$ . Weaker regularity hypotheses on  $\gamma, V$  are possible but we do not pursue those to facilitate the exposition.

With this setup we can represent the divergence form operator  $\mathcal{L}_f$  from (1.3) as  $L_{f,0}$  and the Schrödinger operator  $\mathcal{L}_f$  from (1.4) as  $L_{1/2,f}$ . The PDE (1.2) is thus a special case of the more general PDE

$$\begin{aligned} L_{\gamma,V} v &= g \quad \text{on } \mathcal{X}, \\ v &= h \quad \text{on } \partial\mathcal{X}, \end{aligned} \tag{A.6}$$

whose solutions  $v$  we now study. We assume again that  $g, h$  are given smooth ( $C^\infty$ -) functions on  $\mathcal{X}, \partial\mathcal{X}$ , respectively.

For a smooth vector field  $v: \mathcal{X} \rightarrow \mathbb{R}^d$  the divergence theorem (see [121, p. 143 Proposition 2.3] or [52, p. 714f]) implies that for any  $u \in C_0^\infty(\mathcal{X})$ ,

$$\int_{\mathcal{X}} u \nabla \cdot v \, d\lambda = - \int_{\mathcal{X}} \nabla u \cdot v \, d\lambda, \tag{A.7}$$

and this extends to  $v \in H^1(\mathcal{X}), u \in H_0^1(\mathcal{X})$ , by approximation.

Now using (A.7) and the Poincaré inequality (e.g., [121, p. 340]), for any  $u \in H_0^1(\mathcal{X})$ ,

$$-\langle L_{\gamma,V} u, u \rangle_{L^2(\mathcal{X})} = \langle \gamma \nabla u, \nabla u \rangle_{L^2} + \langle Vu, u \rangle_{L^2} \geq \gamma_{\min} \|\nabla u\|_{L^2}^2 \geq C(\gamma_{\min}, \mathcal{X}) \|u\|_{H^1}^2.$$

Using the dual characterisation of the norm of  $H^{-1}(\mathcal{X}) = (H_0^1(\mathcal{X}))^*$  and testing with  $\phi = -u$ , we deduce

$$\|L_{\gamma,V} u\|_{H^{-1}} = \sup_{\|\phi\|_{H_0^1} \leq 1} \langle L_{\gamma,V} u, \phi \rangle_{L^2} \geq C(\gamma_{\min}, \mathcal{X}) \|u\|_{H^1}. \tag{A.8}$$

Similar arguments imply a converse inequality of (A.8) and hence that  $L_{\gamma,V}$  continuously maps  $H_0^1(\mathcal{X})$  into  $H^{-1}(\mathcal{X})$  with operator norm bounded by some  $D = D(\|\gamma\|_\infty, \|V\|_\infty)$ . From this, one then deduces the following facts, see e.g., [121, p. 355f]: the operator  $L_{\gamma,V}$  realises an isomorphism of  $H_0^1(\mathcal{X})$  onto  $H^{-1}(\mathcal{X})$  with a well-defined inverse  $L_{\gamma,V}^{-1}$  of  $L_{\gamma,V}$  that defines a compact negative self-adjoint operator on the Hilbert space  $L_\lambda^2(\mathcal{X})$ , mapping into  $H_0^1(\mathcal{X})$ . Compactness and the spectral theorem furnish us with eigen-pairs

$$(\lambda_j, e_j) = (\lambda_{\gamma,V,j}, e_{j,\gamma,V}) \in (0, \infty) \times H_0^1(\mathcal{X}), \quad j \in \mathbb{N}, \tag{A.9}$$

of  $-L_{\gamma,V}$  that depend on  $\gamma, V$  (suppressed in the notation unless necessary). In particular, the action of  $L_{\gamma,V}$  on sufficiently regular elements of  $L^2(\mathcal{X})$  can be represented as in (1.6), and for  $h = 0$  and  $g \in L^2(\mathcal{X})$ , we can represent the unique solutions  $v \in H_0^1(\mathcal{X})$  of (A.6) as

$$v = L_{\gamma,V}^{-1}[g] = - \sum_j \lambda_j^{-1} e_j \langle e_j, g \rangle_{L^2}. \tag{A.10}$$

Unique solutions  $v$  in  $H^1(\mathcal{X})$  for (A.6) with  $h \neq 0$  can be constructed likewise (see [59, Chapter 8] or arguing as in [121, p. 359f]), and for smooth  $g, h$  such  $v$  will belong to a Sobolev space if  $\gamma, V$  do: Theorems 8.12 and 8.13 in [59] imply that

$$v \in H^{k+1}(\mathcal{X}) \quad \text{if } \gamma \in H^k(\mathcal{X}), V \in H^{k-1}(\mathcal{X}), k \in \mathbb{N}, \quad (\text{A.11})$$

and that this remains true as long as the source  $g$  belongs to  $H^{k-1}$  (rather than being smooth).

We will show below (Proposition A.5.3) that  $L_{\gamma, V}^{-1}$  is in fact a Lipschitz operator between such Sobolev spaces but need to be careful to track the constants in dependence on  $V, \gamma, k$ . Some first Lipschitz estimates that are uniform in  $V, \gamma$  already follow from (A.8) with  $u = L_{\gamma, V}^{-1}[g]$ , specifically since

$$\|u\|_{L^2} \leq \|u\|_{H^1} \leq c \|L_{\gamma, V} u\|_{H^{-1}} \leq c \|L_{\gamma, V} u\|_{L^2} \quad \text{for } u \in H_0^1,$$

for some  $c = c(\gamma_{\min}, \mathcal{X}) < \infty$ , we see that for any  $g \in L^2$ ,

$$\|L_{\gamma, V}^{-1}[g]\|_{L^2} \leq c \|g\|_{L^2}, \quad (\text{A.12})$$

and we also have the stronger estimate

$$\|L_{\gamma, V}^{-1}[g]\|_{H^1} \leq c \|g\|_{L^2}. \quad (\text{A.13})$$

### A.3 Orthonormal discretisation of $L^2$ and metric entropy

#### A.3.1 The spectrally defined Sobolev-type spaces $\tilde{H}^s$

The spectral theorem furnishing the representation (A.10) in fact implies that  $\{e_j : j \in \mathbb{N}\}$  is a complete  $\langle \cdot, \cdot \rangle_{L^2}$ -orthonormal basis of the Hilbert space  $L^2 = L^2_\lambda(\mathcal{X})$ . Setting  $\gamma = 1, V = 0$ , the well-known Weyl-asymptotics for the eigenfunctions of the standard Laplacian  $-\Delta = -L_{1,0}$  are

$$\lambda_j \simeq j^{2/d} \quad \text{as } j \rightarrow \infty, \quad (\text{A.14})$$

see [122], Corollary 8.3.5. For these eigen-pairs of  $\Delta$ , new Sobolev-type spaces are then defined as

$$\tilde{H}^s(\mathcal{X}) = \left\{ f : \|f\|_{\tilde{H}^s(\mathcal{X})}^2 := \sum_{j \in \mathbb{N}} \lambda_j^s \langle f, e_j \rangle_{L^2}^2 < \infty \right\}, \quad s \in \mathbb{R}, \quad (\text{A.15})$$

which carry a natural (Hilbert-space) inner product

$$\langle f, g \rangle_{\tilde{H}^s} = \sum_j \lambda_j^s \langle f, e_j \rangle_{L^2} \langle g, e_j \rangle_{L^2}.$$

For  $s \geq 0$ , the above spaces consist of elements of  $L^2(\mathcal{X})$  but for  $s < 0$ , these are generalised functions (in the Schwartz sense) and a basic sequence space duality argument shows that

$$\tilde{H}^s(\mathcal{X}) = (\tilde{H}^{-s}(\mathcal{X}))^* \quad \text{for } s < 0, \tag{A.16}$$

we refer to [121, p. 473] for details.

One can define the  $\tilde{H}^s(\gamma)$  spaces also for the spectrum of the  $L_{\gamma,0}$ -operators instead of the standard Laplacian  $\gamma = 1$  – the sequence norms of these spaces  $\tilde{H}^s(\gamma)$ ,  $s = 0, 1, 2$ , can all be shown to be equivalent with equivalence constants depending only on  $U \geq \|\gamma\|_{C^1}$  and  $\gamma_{\min} \leq \gamma \leq \gamma_{\max}$ . Further properties of these spaces are discussed after the next subsection.

### A.3.2 The metric entropy bound in sequence space

A lemma used repeatedly in these notes is the following combinatorial result for the metric complexity of balls in  $\tilde{H}^\alpha$  spaces.

For a totally bounded subset  $A$  of a metric space  $(X, d_X)$ , the number  $N(A, d_X, \epsilon)$  denotes the minimal number of  $d_X$ -balls of radius  $\epsilon > 0$  required to cover  $A$ , and  $\log N(A, d_X, \epsilon)$  is called the *metric entropy* of  $A$ . When  $d_X$  arises from a norm  $\|\cdot\|$ , we just write  $N(A, \|\cdot\|, \epsilon)$  instead.

**Proposition A.3.1.** *Let  $\alpha \geq 0$  and  $-\infty < k < \alpha$ . Then the  $\epsilon$ -log-covering numbers of the ball  $\tilde{h}^\alpha(B)$  of  $\tilde{H}^\alpha(\mathcal{X})$ -norm radius  $B$  for the  $\tilde{H}^k(\mathcal{X})$ -distance satisfy*

$$\log N(\tilde{h}^\alpha(B), \|\cdot\|_{\tilde{H}^k}, \epsilon) \leq \left(\frac{AB}{\epsilon}\right)^{d/(\alpha-k)} \quad \text{for } 0 < \epsilon < AB, \tag{A.17}$$

where  $A = A(d, \alpha, k) < \infty$  is a fixed constant.

*Proof.* We first prove the case  $k = 0$ . A standard scaling argument for norms allows to restrict to  $B = 1$ . Let us write  $f_n = \langle f, e_n \rangle_{L^2}$  for the ‘Fourier’ coefficients of  $f \in L^2(\mathcal{X})$  in the basis  $e_n$  throughout the proof. We also note that (A.14) provides a constant  $b_1 > 0$  such that  $b_1 j^{2/d} \leq \lambda_j$  for all  $j \in \mathbb{N}$ .

For any  $f$  contained in

$$\tilde{h}^\alpha(1) = \left\{ f : \sum_n \lambda_n^\alpha f_n^2 \leq 1 \right\},$$

using (A.14) and Parseval’s identity, we can estimate the error of  $L_\xi^2$ -approximation  $P_{2^{\ell_0}} f$  of  $f$  from all frequencies up to the largest integer  $n < 2^{\ell_0}$ , as

$$\|f - P_{2^{\ell_0}} f\|_{L_\xi^2}^2 = \sum_{n \geq 2^{\ell_0}} f_n^2 \lambda_n^\alpha \lambda_n^{-\alpha} \leq b_1^{-\alpha} 2^{-2\alpha\ell_0/d} \leq \left(\frac{\epsilon}{4}\right)^2,$$

where we have chosen  $\ell_0$  as

$$\ell_0 = \frac{d}{\alpha} \log_2 \left( \frac{4 \cdot 2^{\alpha/d+1} b_1^{-\alpha/2}}{\epsilon} \right)$$

The remaining indices  $n$ ,  $0 \leq n < 2^{\ell_0}$ , are now decomposed into dyadic brackets

$$N_l = \{n : 2^{l-1} \leq n \leq 2^l - 1\}, \quad l = 1, \dots, |N_l| = 2^{l-1},$$

where by convention we set  $N_0 = \{0\}$ . Thus, using again Parseval's identity and the preceding estimate, we can bound, for any  $f, g \in \tilde{h}^\alpha(1)$ ,

$$\begin{aligned} \|f - g\|_{L_\xi^2} &\leq \sqrt{\sum_{0 \leq l \leq \ell_0+1} \sum_{n \in N_l} \lambda_n^{-\alpha} \lambda_n^\alpha (f_n - g_n)^2} + \frac{2\epsilon}{4} \\ &\leq 2^{\alpha/d} b_1^{-\alpha/2} \sum_{0 \leq l \leq \ell_0+1} 2^{-l\alpha/d} \sqrt{\sum_{n \in N_l} (\lambda_n^{\alpha/2} f_n - \lambda_n^{\alpha/2} g_n)^2} + \frac{\epsilon}{2}. \end{aligned}$$

Now by the definition of  $\tilde{h}^\alpha(1)$ , for every  $l \in \mathbb{N} \cup \{0\}$  the vectors  $\{\lambda_n^{\alpha/2} f_n : n \in N_l\}$ ,  $\{\lambda_n^{\alpha/2} g_n : n \in N_l\}$  lie in the unit ball of a Euclidean space of dimension  $|N_l| \leq 2^l$ . The Euclidean  $\epsilon'$ -covering numbers of such a unit ball  $B_l$  for the standard Euclidean norm are

$$N(l) \equiv N(B_l, \|\cdot\|_{\mathbb{R}^{|N_l|}}, \epsilon') \leq \left(\frac{3}{\epsilon'}\right)^{2^l} \quad \text{for } 0 < \epsilon' < 1$$

(see, e.g., [60, Proposition 4.3.34]). By choosing radius  $\epsilon'_l = 2^{l(\alpha/d+1)} 2^{-\ell_0(\alpha/d+1)}$  coverings of  $B_l$  for each  $l$  centred at points  $(f_{n,l,i} : n \in N_l, l \leq \ell_0 + 1)$ , and setting  $f_{n,l,i} = 0$  for  $n \geq 2^{\ell_0+1}$ , we obtain a covering

$$\left(\bar{f}^{(i)} = \lambda_n^{-\alpha/2} f_{n,l,i} : i = 1, \dots, \prod_{0 \leq l \leq \ell_0+1} N(l)\right)$$

of  $\tilde{h}^\alpha(1)$  of radius bounded by

$$\begin{aligned} \|f - \bar{f}^{(i)}\|_{L_\xi^2} &\leq \frac{\epsilon}{2} + 2^{\alpha/d} b_1^{-\alpha/2} \sum_{0 \leq l \leq \ell_0+1} 2^{-l\alpha/d} 2^{l(\alpha/d+1)} 2^{-\ell_0(\alpha/d+1)} \\ &\leq \frac{\epsilon}{2} + 2^{\alpha/d+2} b_1^{-\alpha/2} 2^{-\alpha\ell_0/d} \leq \epsilon, \end{aligned}$$

by choice of  $\ell_0$ . We thus obtain, for some  $c' = c'(\alpha, d) > 0$ ,

$$\begin{aligned} \log_2 N(\tilde{h}^\alpha(1), \|\cdot\|_{L_\xi^2}, \epsilon) &\leq \sum_{0 \leq l \leq \ell_0+1} \log_2 N(l) \\ &\leq \sum_{0 \leq l \leq \ell_0+1} 2^l \left[ \log_2 3 + (\ell_0 - l) \left( \frac{\alpha}{d} + 1 \right) \right] \leq c' 2^{\ell_0} \leq \left(\frac{A}{\epsilon}\right)^{d/\alpha}, \end{aligned}$$

so that (A.17) follows for  $k = 0$ .

For  $-\infty < k < \alpha$ , if  $f, g \in \tilde{h}^\alpha(B)$ , we can write

$$\|f - g\|_{\tilde{H}^k}^2 = \sum_n (f_n \lambda_n^{k/2} - g_n \lambda_n^{k/2})^2 = \|\tilde{f} - \tilde{g}\|_{L^2}^2,$$

where  $\tilde{f}, \tilde{g}$  are defined via ‘Fourier coefficients’  $\lambda_n^{k/2} f_n, \lambda_n^{k/2} g_n, n \in \mathbb{N}$ , which belong to a bounded subset of  $\tilde{H}^{\alpha-k}$ . Hence a  $\|\cdot\|_{L^2}$ -covering of  $\{\tilde{f} : f \in \tilde{h}^\alpha(B)\} \subset \tilde{h}^{\alpha-k}(B)$  induces a  $\tilde{H}^k$ -covering of  $h^\alpha(B)$ , and the bound for  $k = 0$  just established applies with  $\alpha - k$  in place of  $\alpha$ , completing the proof. ■

**Remark A.3.2.** One can show that the inequality (A.17) is optimal in the sense that a corresponding lower bound holds with  $A$  replaced by some  $A' > 0$  depending only on  $d, \alpha, k$ , see [50, Chapter 3] for a proof.

### A.3.3 Metric entropy estimates for standard Sobolev spaces

The spaces  $\tilde{H}^\alpha(\mathcal{X})$  have to be understood relative to the ‘eigen-pairs’  $(e_n, \lambda_n)$  of the Dirichlet Laplacian  $-\Delta = L_{1,0}$  on  $\mathcal{X}$  and (unlike on boundary-less manifolds) are *not* equal to the standard spaces  $H^\alpha(\mathcal{X})$ , due to the presence of a boundary. However, arguing as in [121, Section 5.A], one shows the following facts: We have

$$\tilde{H}^\alpha(\mathcal{X}) = H_0^\alpha(\mathcal{X}), \quad \alpha = 1, 2; \quad \tilde{H}^\alpha(\mathcal{X}) \subset H_0^\alpha(\mathcal{X}), \quad \alpha \in \mathbb{N}, \quad (\text{A.18})$$

and the  $\|\cdot\|_{H^\alpha}$  and  $\|\cdot\|_{\tilde{H}^\alpha}$ -norms are Lipschitz equivalent on  $\tilde{H}^\alpha$ . Moreover, any  $f \in H^\alpha(\mathcal{X})$  that is supported in some compact subset  $K \subset \mathcal{X}$  belongs to  $\tilde{H}^\alpha(\mathcal{X})$  with norm estimate  $\|f\|_{\tilde{H}^\alpha} \leq c_K \|f\|_{H^\alpha}$ . This is clear for  $\alpha = 1, 2$  and extends to  $\alpha \in \mathbb{N}$  by induction as  $\tilde{H}^\alpha(\mathcal{X})$  is the image of  $H^{\alpha-2}(\mathcal{X})$  under an application of  $L_{1,0}^{-1}$ , cf. (A.28). It extends in fact also to  $\alpha < 0$  in that any element of  $H^\alpha(\mathcal{X}) := (H_c^{-\alpha}(\mathcal{X}))^*$  that is compactly supported in  $\mathcal{X}$  in the sense of Schwartz distributions belongs to  $\tilde{H}^\alpha(\mathcal{X})$ , see [121, p. 474].

From what precedes and Proposition A.3.1, one can obtain the classical  $L^2$ -metric entropy inequality for balls  $h^\alpha(B)$  of radius  $B$  in  $H^\alpha(\mathcal{X})$ ,

$$\log N(h^\alpha(B), \|\cdot\|_{L^2(\mathcal{X})}, \epsilon) \lesssim \left(\frac{B}{\epsilon}\right)^{d/\alpha}, \quad 0 < \epsilon < B, \quad (\text{A.19})$$

by an extension argument: Let  $\mathcal{Y} \supset \mathcal{X}$  be a bounded smooth domain in  $\mathbb{R}^d$  such that the closure of  $\mathcal{X}$  is contained in the interior of  $\mathcal{Y}$ . Elements  $f$  of  $H^\alpha(\mathcal{X})$  can be regarded as restrictions of functions  $\tilde{f} : \mathcal{Y} \rightarrow \mathbb{R}$  in  $H^\alpha(\mathcal{Y})$  that are compactly supported in  $\mathcal{Y}$  and for which

$$\|\tilde{f}\|_{\tilde{H}^\alpha(\mathcal{Y})} \lesssim \|\tilde{f}\|_{H^\alpha(\mathcal{Y})} \lesssim \|f\|_{H^\alpha(\mathcal{X})};$$

the second inequality follows from the usual extension theorem for Sobolev functions (see [86, Section 1.8]) and by employing a ‘cut-off’ function in  $C_c^\infty(\mathcal{Y})$  that equals

one identically on  $\mathcal{X}$ . Thus an  $\epsilon$ - $L^2(\mathcal{Y})$ -covering of a sufficiently large ball in  $\tilde{H}^\alpha(\mathcal{Y})$  induces a covering of a ball in  $H^\alpha(\mathcal{X})$  of  $L^2(\mathcal{X})$  radius

$$\|f - g\|_{L^2(\mathcal{X})} \leq \|\bar{f} - \bar{g}\|_{L^2(\mathcal{Y})} \leq \epsilon,$$

and (A.19) then follows from Proposition A.3.1.

When covering a ball in  $H_c^\alpha(\mathcal{X})$ , the previous extension argument can be refined to cover a weaker class of norms (relevant in Condition 2.1.1). Elements  $f \in H_c^\alpha(\mathcal{X})$  vanish at  $\partial\mathcal{X}$  in a sufficiently regular way that their extension  $\tilde{f}$  by zero outside of  $\mathcal{X}$  belongs to  $H^\alpha(\mathcal{Y})$  (see [86, Theorem 1.11.4]). If we consider the topological dual space  $(\tilde{H}^\kappa(\mathcal{X}))^*$  for  $\kappa \geq 0$ , then

$$\|f\|_{(H^\kappa(\mathcal{X}))^*} = \sup_{\|\phi\|_{H^\kappa(\mathcal{X})} \leq 1} \left| \int_{\mathcal{X}} \phi f \right| \leq \sup_{\|\bar{\phi}\|_{H^\kappa(\mathcal{Y})} \leq C} \left| \int_{\mathcal{Y}} \bar{\phi} \tilde{f} \right| \lesssim \|\tilde{f}\|_{\tilde{H}^{-\kappa}(\mathcal{Y})}, \quad (\text{A.20})$$

where the last supremum ranges over those  $\bar{\phi}$  which satisfy  $\bar{\phi} = \phi$  on  $\mathcal{X}$  and are compactly supported in  $\mathcal{Y}$ , whence  $\|\bar{\phi}\|_{\tilde{H}^\kappa(\mathcal{Y})} \leq C'$  for some  $C' > 0$ , and where the last inequality follows from (A.16). We therefore deduce from Proposition A.3.1 that a ball  $h_c^\alpha(B)$  of radius  $B$  in  $H_c^\alpha(\mathcal{X})$  satisfies the metric entropy inequality

$$\log N(h_c^\alpha(B), \|\cdot\|_{(H^\kappa(\mathcal{X}))^*}, \epsilon) \lesssim \left(\frac{B}{\epsilon}\right)^{d/(\alpha+\kappa)}, \quad 0 < \epsilon < B. \quad (\text{A.21})$$

For  $0 \leq \kappa < 1/2$ , the last inequality holds for  $h^\alpha(B) \supset h_c^\alpha(B)$  – the boundary constraint can be relaxed as we can then use  $\bar{f}$  instead of  $\tilde{f}$  in (A.20) together with the fact that zero extensions  $\tilde{\phi}$  of  $\phi \in H^\kappa(\mathcal{X})$  belong to  $H^\kappa(\mathcal{Y})$  and then also  $\tilde{H}^\kappa(\mathcal{Y})$  for such  $\kappa$  (see [86, Theorem 1.11.4]).

Results of the type given in this subsection for general function spaces can be found in [50], but we prefer to give self-contained proofs of the precise bounds required in our setting.

## A.4 Feynman–Kac formulæ

Solutions  $v$  of elliptic PDEs (A.6) have probabilistic representations by virtue of integrals against sample paths of diffusion processes, e.g., [13, 55]. For the applications of diffusion process techniques, it is convenient to ensure that  $C^2$ -solutions  $v$  of (A.6) exist. In our setting where  $(\gamma, V) \in H^\beta \times H^{\beta-1}$  for some  $\beta > 1 + d/2$ , the regularity result (A.11) with  $k = \beta$  and the Sobolev imbedding (A.4) are sufficient to ensure  $v \in H^{\beta+1} \subset C^2$ .

From the theory of stochastic differential equations (e.g., [13, p. 10], and noting  $\gamma \in H^\beta \subset C^1$  so that also  $\sqrt{\gamma} \in C^1$  as  $\gamma \geq \gamma_{\min} > 0$ ), there exists a unique Markov diffusion process  $(X_t : t \geq 0)$  on  $\mathbb{R}^d$  which has  $\mathcal{L}_{\gamma,0} = \nabla \cdot (\gamma \nabla)$  as a generator,

corresponding to path-wise solutions to the stochastic differential equation

$$dX_t = \nabla\gamma(X_t) dt + \sqrt{2\gamma(X_t)} dW_t, \quad t \geq 0, \quad (\text{A.22})$$

where  $(W_t : t \geq 0)$  is a  $d$ -dimensional Brownian motion. Then, by [55, p. 127 Theorem 2.1], the  $C^2$ -solution  $v$  to the elliptic PDE (A.6) has representation

$$v(x) = -E_\gamma^x \left[ \int_0^{\tau_{\mathcal{X}}} g(X_t) e^{-\int_0^t V(X_s) ds} dt \right] + E_\gamma^x [h(X_{\tau_{\mathcal{X}}}) e^{-\int_0^{\tau_{\mathcal{X}}} V(X_s) ds}], \quad (\text{A.23})$$

for every  $x \in \mathcal{X}$ , where  $E_\gamma^x$  is the expectation under the law of the process  $(X_t : t \geq 0)$  solving (A.22) started at  $x$  with exit time  $\tau_{\mathcal{X}}$  from the domain  $\mathcal{X}$ .

**Lemma A.4.1.** *We have  $\sup_{x \in \mathcal{X}} E_\gamma^x \tau_{\mathcal{X}} \leq c$  for some  $c = c(\mathcal{X}, \gamma_{\min})$ .*

*Proof.* Using [13, Section VII, Theorem 4.3] and [41, Corollary 3.2.8] for the precise form of the constants, there exist non-negative transition densities  $p_\gamma$  of  $(X_s : s \geq 0)$  which satisfy the estimate

$$p_\gamma(t, x, y) \leq c_1 t^{-d/2}, \quad t > 0, x, y \in \mathbb{R}^d, \quad (\text{A.24})$$

with  $c_1$  depending only on the lower bound  $\gamma_{\min} \leq \gamma$ . Then, arguing as in the proof of [34, Theorem 1.17], with (A.24) replacing the standard heat kernel estimate for the Brownian motion ( $\gamma = 1/2$ ) used in [34], we obtain the desired inequality. ■

From what precedes and Jensen's inequality we deduce that if  $g = 0$ ,  $h \geq h_{\min} > 0$ , then, under the maintained assumptions, solutions  $v$  to (A.6) are bounded away from zero,

$$\inf_{x \in \mathcal{X}} v(x) \geq h_{\min} e^{-\|V\|_\infty \sup_{x \in \mathcal{X}} E_\gamma^x \tau_{\mathcal{X}}} \geq h_{\min} e^{-c\|V\|_\infty} > 0. \quad (\text{A.25})$$

In a similar vein, for non-negative potentials  $V \geq 0$ , the solutions  $v$  to (A.6) satisfy

$$\|v\|_\infty \leq \|g\|_\infty \sup_{x \in \mathcal{X}} E_\gamma^x \tau_{\mathcal{X}} + \|h\|_\infty \leq c\|g\|_\infty + \|h\|_\infty. \quad (\text{A.26})$$

In particular, the solution operator  $L_{\gamma, V}^{-1}$  of (A.6) with Dirichlet boundary conditions  $h = 0$  is linear and Lipschitz for the  $\|\cdot\|_\infty$ -norm,

$$\|L_{\gamma, V}^{-1}[g]\|_\infty \leq c\|g\|_\infty \quad \text{for all } g \in C^\infty(\mathcal{X}), \quad (\text{A.27})$$

with Lipschitz constant  $c$  from Lemma A.4.1, upgrading the  $L^2$ -estimate (A.12) to one for the uniform norm.

## A.5 Elliptic regularity estimates

The Lipschitz estimates (A.12), (A.13), (A.27) for the elliptic solution operator  $L_{\gamma, V}^{-1}$  are remarkably 'universal' in the conductivity  $\gamma$  and potential  $V$  in that they only

depend on  $\gamma, V$  via the lower bounds on these. However, from elliptic PDE theory we expect the solution operator to be 2-smoothing in a natural Sobolev scale. The constants in such stronger estimates depend on higher regularity of the coefficients  $\gamma, V$  and require different proofs.

A starting point is the corresponding result for the standard Laplacian  $\Delta = L_{1,0}$  which establishes a linear topological isomorphism

$$u \mapsto (\Delta u, \text{tr}[u]) : H^{k+1}(\mathcal{X}) \rightarrow H^{k-1}(\mathcal{X}) \times H^{k+1/2}(\partial\mathcal{X}), \quad (\text{A.28})$$

where  $\mathcal{X}$  is a smooth bounded domain in  $\mathbb{R}^d$ ,  $\text{tr} = \text{tr}_{\partial\mathcal{X}}$  the boundary trace operator, and integer  $k \geq 1$ . See [86, Theorem II.5.4] for a proof. The isomorphism furnishes us with constants  $k_1, k_2$  that depend only on  $k, \mathcal{X}, d$  for which

$$k_1 \|\Delta u\|_{H^{k-1}} \leq \|u\|_{H^{k+1}} \leq k_2 \|\Delta u\|_{H^{k-1}} \quad \text{if } u \in H_0^{k+1}(\mathcal{X}), \quad k \in \mathbb{N}. \quad (\text{A.29})$$

Now let us assume  $(\gamma, V) \in H^\beta \times H^{\beta-1}$  for  $\beta > 1 + d/2$  and write

$$\Delta u = \gamma^{-1}(L_{\gamma,V}u - \nabla\gamma \cdot \nabla u + Vu). \quad (\text{A.30})$$

Then for  $g \in L^2(\mathcal{X})$ , we have  $u = L_{\gamma,V}^{-1}[g] \in H_0^2(\mathcal{X})$  in view of (A.11) with  $k = 1$ , so the inequality (A.29) gives

$$\begin{aligned} \|L_{\gamma,V}^{-1}[g]\|_{H^2} &\leq k_2 \|\Delta L_{\gamma,V}^{-1}[g]\|_{L^2} = k_2 \|\gamma^{-1}(g - \nabla\gamma \cdot \nabla L_{\gamma,V}^{-1}[g] + VL_{\gamma,V}^{-1}[g])\|_{L^2} \\ &\leq k_2 \gamma_{\min}^{-1} (\|g\|_{L^2} + (\|\gamma\|_{C^1} + \|V\|_\infty) \|L_{\gamma,V}^{-1}[g]\|_{H^1}) \\ &\leq C(\gamma_{\min}, B) \|g\|_{L^2}, \end{aligned} \quad (\text{A.31})$$

where we have used (A.13) in the last line, and where  $B$  is an upper bound for  $\max(\|\gamma\|_{C^1}, \|V\|_\infty)$ .

The ‘dual’ version of the estimate (A.31) and its converse are as follows.

**Proposition A.5.1.** *Suppose  $(\gamma, V) \in H^\beta \times H^{\beta-1}$  for some  $\beta > 1 + d/2$  and that  $\gamma \geq \gamma_{\min} > 0$ . Then we have for all  $g \in L^2(\mathcal{X})$  that*

$$c(B) \|g\|_{(H_0^2)^*} \leq \|L_{\gamma,V}^{-1}[g]\|_{L^2} \leq C(B, \gamma_{\min}) \|g\|_{(H_0^2)^*},$$

where  $B$  is an upper bound for  $\max(\|\gamma\|_{C^1}, \|V\|_\infty)$ .

*Proof.* Representing the  $L^2$ -norm by duality and using the self-adjointness of  $L_{\gamma,V}^{-1}$  as well as (A.31), we obtain

$$\begin{aligned} \|L_{\gamma,V}^{-1}[g]\|_{L^2} &= \sup_{\|\varphi\|_{L^2} \leq 1} \left| \int_{\mathcal{X}} L_{\gamma,V}^{-1}[g]\varphi \right| = \sup_{\|\varphi\|_{L^2} \leq 1} \left| \int_{\mathcal{X}} g L_{\gamma,V}^{-1}[\varphi] \right| \\ &\leq C(\gamma_{\min}, B) \sup_{\varphi \in H_0^2, \|\varphi\|_{H^2} \leq 1} \left| \int_{\mathcal{X}} g\varphi \right| \\ &= C(\gamma_{\min}, B) \|g\|_{(H_0^2)^*}, \end{aligned}$$

which completes the proof of the upper bound. For the lower bound, notice first that by the divergence theorem (A.7) and the Cauchy–Schwarz inequality, for any  $w \in H_0^2$ ,

$$\begin{aligned} \|L_{\gamma,V}w\|_{(H_0^2)^*} &= \sup_{\phi \in H_0^2: \|\phi\|_{H^2} \leq 1} \left| \int w L_{\gamma,V} \phi \right| \\ &\leq \bar{c}(B) \|w\|_{L^2} \sup_{\phi \in H_0^2: \|\phi\|_{H^2} \leq 1} [\|\Delta \phi\|_{L^2} + \|\nabla \phi\|_{L^2} + \|\phi\|_{L^2}] \\ &\lesssim \bar{c}(B) \|w\|_{L^2}. \end{aligned}$$

Now in light of (A.31) we can insert  $w = L_{\gamma,V}^{-1}[g] \in H_0^2$  in the preceding inequality which yields the desired lower bound.  $\blacksquare$

It is sometimes convenient to remove the dependence of the constant on  $B$  in the upper bound in the previous result, which is possible by considering slightly weaker dual norms. Indeed, replacing (A.31) by (A.13) in the previous proof immediately gives, for  $\bar{c} = \bar{c}(\mathcal{X}, \gamma_{\min})$ ,

$$\|L_{\gamma,V}^{-1}[g]\|_{L^2} \leq \bar{c} \|g\|_{(H_0^1)^*}. \quad (\text{A.32})$$

A general bound on the Sobolev regularity of solutions of the PDE (A.6) – quantifying (A.11) – is now given in the following proposition. We remark that  $g \in C^\infty$  can be relaxed to  $g \in H^{\beta-1}$  by simple adjustments to the proof.

**Proposition A.5.2.** *Let  $\beta > 1 + d/2$ . Suppose  $v = v_{\gamma,V}$  is the unique solution of (A.6) for  $\gamma \in H^\beta(\mathcal{X})$  such that  $\gamma \geq \gamma_{\min} > 0$ ,  $V \in H^{\beta-1}(\mathcal{X})$  such that  $V \geq 0$ , and smooth  $g, h$ . Then  $v \in H^{\beta+1}(\mathcal{X})$  and for every  $B > 0$ , there exists a constant  $C = C(\beta, d, \mathcal{X}, \gamma_{\min}, g, h, B) > 0$  such that*

$$\sup_{\|\gamma\|_{H^\beta} + \|V\|_{H^{\beta-1}} \leq B} \|v_{\gamma,V}\|_{H^{\beta+1}} \leq C < \infty. \quad (\text{A.33})$$

*Proof.* We know that  $v \in H^{\beta+1}(\mathcal{X})$  by (A.11) with  $k = \beta$  and that

$$c_k(h) := \|h\|_{H^{k+1/2}(\partial\mathcal{X})} < \infty$$

for all  $k$  since  $h$  is smooth. Then (A.28), (A.30) and the multiplicative inequality (A.2) give

$$\begin{aligned} \|v\|_{H^{\beta+1}} &\lesssim c_\beta(h) + \|\gamma^{-1}(L_{\gamma,V}v - \nabla\gamma \cdot \nabla v + Vv)\|_{H^{\beta-1}} \\ &\lesssim c_\beta(h) + \|\gamma^{-1}\|_{H^{\beta-1}} (\|g\|_{H^{\beta-1}} + (\|\gamma\|_{H^\beta} + \|V\|_{H^{\beta-1}}) \|v\|_{H^\beta}). \end{aligned}$$

By the chain rule and (A.2), we can bound

$$\sup_{\|\gamma\|_{H^\beta} \leq B, \|V\|_{H^{\beta-1}} \leq B} \max(\|\gamma^{-1}\|_{H^{\beta-1}}, \|\gamma\|_{H^\beta}, \|V\|_{H^{\beta-1}}) \leq C(B) < \infty,$$

and then using also the interpolation inequality (A.5), we obtain

$$\begin{aligned} \|v\|_{H^{\beta+1}} &\lesssim c_\beta(h) + \|g\|_{H^{\beta-1}} + \|v\|_{H^\beta} \\ &\lesssim c_\beta(h) + \|g\|_{H^{\beta-1}} + \|v\|_{H^{\beta+1}}^{\beta/(\beta+1)} \|v\|_{L^2}^{1/(\beta+1)}. \end{aligned}$$

When  $\|v\|_{H^{\beta+1}} \leq 1$ , we use (A.26) to deduce

$$\begin{aligned} \|v\|_{H^{\beta+1}} &\lesssim c_\beta(h) + \|g\|_{H^{\beta-1}} + \|v\|_\infty^{1/(\beta+1)} \\ &\lesssim c_\beta(h) + \|g\|_{H^{\beta-1}} + (\|g\|_\infty + \|h\|_\infty)^{1/(\beta+1)}, \end{aligned}$$

and when  $\|v\|_{H^{\beta+1}} \geq 1$ , then dividing both sides by  $\|v\|_{H^{\beta+1}}^{\beta/(\beta+1)}$  and using again (A.26) yields

$$\begin{aligned} \|v\|_{H^{\beta+1}}^{1/(\beta+1)} &\lesssim c_\beta(h) + \|g\|_{H^{\beta-1}} + \|v\|_\infty^{1/(\beta+1)} \\ &\leq c_\beta(h) + \|g\|_{H^{\beta-1}} + (\|g\|_\infty + \|h\|_\infty)^{1/(\beta+1)}. \end{aligned}$$

Raising the last inequality to the power  $\beta + 1$  and taking maxima from the last two inequalities implies (A.33).  $\blacksquare$

Finally, we give an explicit Lipschitz estimate for the solution operator  $L_{\gamma,V}^{-1}$  between general Sobolev spaces.

**Proposition A.5.3.** *Suppose  $0 < \gamma_{\min} \leq \gamma$  and that  $\|\gamma\|_{C^\beta} + \|V\|_{C^{\beta-1}} \leq B$  for some  $\beta \geq 1$ . Then for some  $C = C(B, \beta, d, \mathcal{X}, \gamma_{\min})$ , we have, for all  $g \in H^{\beta-1}$ , that*

$$\|L_{\gamma,V}^{-1}[g]\|_{H^{\beta+1}} \leq C \|g\|_{H^{\beta-1}}.$$

*Proof.* By (A.11) we know that  $L_{\gamma,V}^{-1}[g] \in H_0^{\beta+1}$ . From (A.29) and (A.30), we see, for any  $u \in H_0^{\beta+1}$ , that

$$\begin{aligned} \|u\|_{H^{\beta+1}} &\leq k_2 \|\Delta u\|_{H^{\beta-1}} \\ &\leq \|\gamma_{\min}^{-1}\|_{C^{\beta-1}} \|L_{\gamma,V} u\|_{H^{\beta-1}} + \|\gamma\|_{C^\beta} \|u\|_{H^\beta} + \|V\|_{C^{\beta-1}} \|u\|_{H^{\beta-1}}. \end{aligned}$$

By the hypotheses,  $\|\gamma_{\min}^{-1}\|_{C^{\beta-1}} + \|\gamma\|_{C^\beta} + \|V\|_{C^{\beta-1}}$  is bounded by a fixed constant,  $B'$  say. We will show that the previous inequality implies that

$$\|u\|_{H^{\beta+1}} \leq C \|L_{\gamma,V} u\|_{H^{\beta-1}} \quad \text{for all } u \in H_0^{\beta+1}, \quad (\text{A.34})$$

which then also implies the proposition by taking  $u = L_{\gamma,V}^{-1}[g]$ . To prove (A.34) suppose by way of contradiction that the inequality is not true. Then there exists a sequence  $u_m \in H_0^{\beta+1}$  such that  $\|u_m\|_{H^{\beta+1}} = 1$  but  $\|L_{\gamma,V} u_m\|_{H^{\beta-1}} \rightarrow 0$  as  $m \rightarrow \infty$ . By compactness of the imbedding  $H_0^{\beta+1} \subset H_0^\beta \subset H_0^1 = H_c^1$  (the latter being a Banach space), this sequence converges (if necessary by passing to a subsequence)

to some  $u \in H_0^1$  satisfying  $L_{\gamma, \nu} u = 0$ . The last but one displayed inequality implies that the sequence  $u_m$  is Cauchy in  $H^{\beta+1}$ , and hence its limit  $u_0$  must also satisfy  $\|u_0\|_{H^{\beta+1}} = 1$ , in particular,  $u_0 \neq 0$ . But from the results from Appendix A.2, we know that the unique solution in  $H_0^1$  to  $L_{\gamma, \nu} u = 0$  equals  $u_0 = 0$ , a contradiction. ■



## Appendix B

### Further auxiliary results

#### B.1 Results from Gaussian process theory

A Borel probability measure  $\mu$  on a separable Banach space  $B$  is called *centred Gaussian* if the image measures  $\mu \circ L^{-1}$  are mean zero normal distributions on  $\mathbb{R}$  for all  $L$  in the topological dual space  $B^*$  of  $B$ . A (Borel-) random variable  $X$  defined on some probability space  $(\Omega, \mathcal{A}, \Pr)$  and taking values in a separable Banach space  $(B, \|\cdot\|_B)$  is called centred Gaussian if its law  $\mu_X = \Pr \circ X^{-1}$  defines a centred Gaussian probability measure on  $B$ .

A *centred Gaussian process* with index set  $T$  is a collection  $\{X(t) : t \in T\}$  of random variables such that any finite selection  $\{X(t_1), \dots, X(t_d)\}$ ,  $t_i \in T$ ,  $d \in \mathbb{N}$ , is multivariate centred normal in  $\mathbb{R}^d$ .

Any Gaussian random variable in a Banach space can be viewed as a Gaussian process indexed by  $B^*$ , and conversely often Gaussian processes can be realised as  $B$ -valued random variables for appropriate  $B$ , e.g., the Brownian motion  $(X(t) : t \in [0, 1])$  is a Gaussian process which (by sample-path continuity) also defines a Gaussian random variable in the space of continuous functions  $C([0, 1])$ . So the theories of Gaussian measures and processes inform each other and tools from both fields are useful when deployed appropriately. See [60, Section 2] for an account of the basic theory and many references.

There are some classical results about properties of Gaussian measures and processes, such as *Fernique's theorem*, the *Karhunen–Loève expansion*, the *Sudakov–Cirelson* and *log-Sobolev inequalities*, *Borell's isoperimetric inequality* and the *Cameron–Martin theorem*. The last two theorems involve the notion of a *reproducing kernel Hilbert space* (RKHS)  $\mathcal{H}$  of a Gaussian measure  $\mu$  whose inner product  $\langle \cdot, \cdot \rangle_{\mathcal{H}}$  encodes the geometry of the covariance structure of Gaussian variables  $X \sim \mu$ . We recall that  $\mathcal{H}$  is a subspace of  $B$  that is compactly embedded in the sense that the identity map

$$\text{id} : (\mathcal{H}, \|\cdot\|_{\mathcal{H}}) \rightarrow (B, \|\cdot\|_B)$$

is compact (in particular continuous). One can also define a ‘process-RKHS’ for Gaussian processes (instead of measures), and when the law of the process defines a Gaussian probability measure on  $B$ , these RKHS will coincide under mild regularity conditions (e.g., [58, Lemma 11.14]). We will admit these facts and theorems without stating them here, even though we will give precise references to the literature when they are used in our proofs.

There are a couple of results on Gaussian measures that are not covered in [60, Chapter 2] that we review here now.

### B.1.1 Small deviation estimates for Gaussian measures

The following theorem characterises the ‘small ball’ (or ‘small deviation’) asymptotics of a Gaussian measure  $\mu$  on a separable Banach space  $B$  quantified by the function

$$\phi(\varepsilon) = -\log \mu(x \in B : \|x\|_B \leq \varepsilon) \quad \text{for } \varepsilon > 0,$$

in terms of the metric entropy function

$$\mathcal{H}(\varepsilon) = \log N(\{x \in B : \|x\|_{\mathcal{H}} \leq 1\}, \|\cdot\|_B, \varepsilon) \quad \text{for } \varepsilon > 0,$$

which equals the logarithm of the minimal number of  $B$ -balls of radius  $\varepsilon > 0$  required to cover the (compact) unit ball of  $\mathcal{H} \subset B$ .

**Theorem B.1.1** (Kuelbs–Li–Linde). *Let  $a > 0$ . Then as  $\varepsilon \rightarrow 0$ ,*

$$\phi(\varepsilon) \approx \varepsilon^{-a} \quad \text{if and only if} \quad \mathcal{H}(\varepsilon) \approx \varepsilon^{-2a/(2+a)},$$

where  $\approx$  denotes two-sided inequalities up to multiplicative constants (that may depend on  $B$ ).

For a proof of this theorem, see [84, Theorem 1.2] with  $\alpha = a$  and  $\beta = 0$ , with a key earlier reference being [77]. A refinement of this result at the ‘log-scale’ ( $\beta > 0$  in [84, Theorem 1.2]) is possible but not required in the present context. We also refer to [85] for more on the topic of small ball deviation estimates for Gaussian processes.

### B.1.2 The Gaussian correlation inequality

A result in the theory of Gaussian processes that has been proved only recently is the Gaussian correlation inequality (formerly conjecture!) – the original proof of Royen is spelt out in [79].

**Theorem B.1.2** (Gaussian correlation inequality). *Let  $\mu$  be a centred Gaussian measure on a separable Banach space  $B$  and let  $C_1, C_2$  be convex closed symmetric subsets of  $B$ . Then*

$$\mu(C_1 \cap C_2) \geq \mu(C_1)\mu(C_2).$$

A proof for finite-dimensional Gaussian distributions ( $\dim(B) < \infty$ ) is given in [79]. As the conclusion of the result does not depend on  $\dim(B)$ , the Karhunen–Loève theorem combined with a standard limiting argument extends the result to infinite-dimensions, see [63, Appendix A.4] for details.

### B.1.3 Whittle–Matérn-type Gaussian processes

We can construct a centred Gaussian process  $X = \{X(z), z \in \mathcal{Z}\}$  indexed by an arbitrary subset  $\mathcal{Z} \subset \mathbb{R}^d$  by prescribing a *positive definite function* (covariance kernel)

$K: \mathbb{R}^d \rightarrow \mathbb{R}$  such that  $EX(s)X(t) = K(s-t)$ ,  $s, t \in \mathcal{Z}$ , see [60, Proposition 2.1.10]. One such choice (with  $|\cdot| = \sqrt{\langle \cdot, \cdot \rangle_{\mathbb{R}^d}}$  the Euclidean norm) is

$$K_\alpha(x) = \int_{\mathbb{R}^d} e^{i\langle x, \xi \rangle_{\mathbb{R}^d}} d\bar{\mu}(\xi), \quad d\bar{\mu}(\xi) = (1 + |\xi|^2)^{-\alpha} d\xi, \quad x \in \mathcal{Z},$$

whenever  $\alpha > d/2$ . That  $K_\alpha$  is positive definite is the easy part of ‘Bochner’s theorem’: For any finite collection  $\beta_j \in \mathbb{C}$ ,  $s_j \in \mathcal{Z}$ , we have

$$\begin{aligned} \sum_{k,j} \beta_k \overline{\beta_j} K_\alpha(s_k - s_j) &= \int_{\mathbb{R}^d} \sum_k \beta_k e^{i\langle s_k, \xi \rangle_{\mathbb{R}^d}} \sum_j \overline{\beta_j e^{i\langle s_j, \xi \rangle_{\mathbb{R}^d}}} d\bar{\mu}(\xi) \\ &= \int_{\mathbb{R}^d} \left| \sum_j \beta_j e^{i\langle s_j, \xi \rangle_{\mathbb{R}^d}} \right|^2 d\bar{\mu}(\xi) \geq 0. \end{aligned}$$

Such (stationary) Gaussian process priors are popular in numerical analysis, spatial statistics and machine learning, and are sometimes called ‘Whittle–Matérn’ processes. The following theorem permits their use in the setting of these notes.

**Theorem B.1.3** (Existence of Whittle–Matérn prior measures). *Let  $\mathcal{Z}$  be a bounded domain in  $\mathbb{R}^d$  with smooth boundary. Then the law of the Gaussian process  $X = \{X(z), z \in \mathcal{Z}\}$  with covariance function  $K_\alpha$ ,  $\alpha > d/2$ , defines a centred Gaussian probability measure  $\mu_X$  on the Banach space  $C_u(\mathcal{Z})$  of bounded uniformly continuous function on  $\mathcal{Z}$  with reproducing kernel Hilbert space  $\mathcal{H} = H^\alpha(\mathcal{Z})$ . We further have  $\mu_X(H^\beta(\mathcal{Z})) = 1$  for every  $0 \leq \beta < \alpha - d/2$ .*

*Proof.* The process-RKHS  $\mathcal{H}_{\alpha, \mathcal{Z}}$  of  $\{X(z) : z \in \mathcal{Z}\}$  can be characterised as the restriction to  $\mathcal{Z}$  of functions in

$$\mathcal{H}_{\bar{\mu}} = \left\{ g : \mathbb{R}^d \rightarrow \mathbb{R} \text{ Borel measurable} : \|g\|_{\mathcal{H}_{\bar{\mu}}}^2 \equiv \int_{\mathbb{R}^d} |\hat{g}|^2 d\bar{\mu} < \infty \right\},$$

with  $\hat{g} = \int_{\mathbb{R}^d} e^{i\langle \cdot, x \rangle_{\mathbb{R}^d}} g(x) dx$  the Fourier transform of  $g$ , equipped with the quotient norm

$$\|h\|_{\mathcal{H}_{\alpha, \mathcal{Z}}} = \inf_{g \in \mathcal{H}_{\bar{\mu}}, g=h \text{ on } \mathcal{Z}} \|g\|_{\mathcal{H}_{\bar{\mu}}},$$

see [58, Lemma 11.35] for a proof of this fact. Then, by, e.g., [86, Theorem 1.9.2] and for  $\mathcal{Z}$  a bounded smooth domain in  $\mathbb{R}^d$ , the norm of  $\mathcal{H}_{\alpha, \mathcal{Z}}$  is Lipschitz equivalent to the classically defined Sobolev norm (cf. (A.1)) of  $H^\alpha(\mathcal{Z})$ , and hence the process RKHS of  $\{X(z), z \in \mathcal{Z}\}$  coincides with  $H^\alpha(\mathcal{Z})$  for such  $\mathcal{Z}$ .

The Gaussian process  $(X(z) : z \in \mathcal{Z})$  has intrinsic covariance metric

$$d_X(s, t) = \sqrt{E(X(s) - X(t))^2}, \quad s, t \in \mathcal{Z},$$

satisfying for some  $0 < \gamma < \alpha - d/2$  the estimate

$$d_X^2(s, t) = EX^2(s) + EX^2(t) - 2EX(s)X(t) = 2(K_\alpha(0) - K_\alpha(s-t)) \leq c|s-t|^\gamma$$

because  $e^{-i(\cdot, \xi)_{\mathbb{R}^d}}$  is  $\gamma$ -Hölder with Hölder constant  $|\xi|^\gamma$  and since  $(1 + |\xi|^2)^{-\alpha} |\xi|^\gamma$  is integrable on  $\mathbb{R}^d$ . This implies (using, e.g., [60, Proposition 4.3.34]) that the bounded subset  $\mathcal{Z}$  of  $\mathbb{R}^d$  can be covered by at most  $N(\epsilon) = (A/\epsilon)^{2d/\gamma}$  balls of  $d_X$ -radius  $\epsilon > 0$ . Since  $\sqrt{\log N(\epsilon)}$  is  $d\epsilon$ -integrable at zero, [60, Theorem 2.3.7] implies that there exists a version of  $X$  that is uniformly continuous on  $\mathcal{Z}$  for the  $d_X$  and then also the  $|\cdot|$ -metric. By [60, Proposition 2.1.5], the law of  $X$  then defines a Gaussian Borel probability measure  $\mu_X$  on the separable Banach space  $C_u(\mathcal{Z})$  of bounded uniformly continuous functions on  $\mathcal{Z}$ , and the process RKHS  $\mathcal{H}_{\alpha, \mathcal{Z}}$  of  $(X(z) : z \in \mathcal{Z})$  co-incides with the RKHS  $\mathcal{H}$  of  $\mu_X$  by virtue of [58, Lemma 11.14].

To see that  $\mu_X(\|X\|_{H^\beta} < \infty) = 1$  whenever  $0 < \beta < \alpha - d/2$ , use the Karhunen–Loève theorem [60, Theorem 2.6.10] to represent  $X = \sum_i g_i h_i$  almost surely in  $C(\mathcal{Z})$ , where the  $(h_i)$  form an orthonormal basis of the RKHS  $\mathcal{H}^\alpha(\mathcal{Z})$ , and the  $g_i$  are iid standard normal  $N(0, 1)$  variables. The identity imbedding map  $H^\alpha(\mathcal{Z}) \rightarrow H^\beta(\mathcal{Z})$  is Hilbert-Schmidt for any  $\beta < \alpha - d/2$ , in particular,

$$\sum_{i=1}^{\infty} \|h_i\|_{H^\beta}^2 < \infty,$$

see [2, Theorem 6.61]. By the theory of radonifying maps this implies  $\Pr(X \in H^\beta) = 1$ , see [130, Theorem 3.20 and Proposition 13.5]. ■

We note that the previous proof shows that the theorem remains valid for arbitrary open bounded subsets  $\mathcal{Z}$  (possibly with irregular boundary) of  $\mathbb{R}^d$  as long as  $\mathcal{H}$  is taken to equal the space  $\mathcal{H}_{\alpha, \mathcal{Z}}$  introduced in the proof.

One can obtain stronger (than  $H^\beta$ ) almost sure Hölder regularity properties of the random function  $z \mapsto X(z)$  by more refined arguments (e.g., [58, Proposition I.4] or [3, p. 22f]). In these notes we confine ourselves to work with Sobolev regularity – this combines nicely with regularity estimates for elliptic PDEs used elsewhere and also permits a ‘soft’ proof of path regularity using only the Karhunen–Loève theorem and Hilbert–Schmidt embeddings.

The preceding arguments show that for any  $\alpha > d/2$ , we can find a Gaussian prior  $\Pi = \mu_X$  that has  $H^\alpha(\mathcal{Z})$  as RKHS. In some of the results in these notes we wish to enforce a certain boundary behaviour of the prior  $\Pi$ . We can achieve this by taking a cutoff-function  $\xi \in C^\infty(\mathcal{Z})$  of compact support in  $\mathcal{Z}$  and consider a new Gaussian process  $(X_\xi(z) = X(z)\xi(z) : z \in \mathcal{Z})$ . The RKHS of this new process then equals

$$\mathcal{H}_\xi = \{g = \xi h : h \in H^\alpha(\mathcal{Z})\} \subset H_c^\alpha(\mathcal{Z})$$

with norm  $\|\xi^{-1} \cdot\|_{H^\alpha}$ . We can also take a different approach by starting with a Gaussian series expansion

$$X(z) = \sum_{j=1}^{\infty} g_j \lambda_j^{-\alpha/2} e_j(z), \quad z \in \mathcal{Z}, \tag{B.1}$$

of the eigenfunctions of the Dirichlet Laplacian  $L_{1,0}$  from (A.9). Their laws define Gaussian probability measures on  $C(\mathcal{Z})$  and as in [60, Example 2.6.15] one shows that their RKHS are the closed subspaces  $\tilde{H}^\alpha(\mathcal{Z}) \subset H_0^\alpha(\mathcal{Z})$  of  $H^\alpha(\mathcal{Z})$  from (A.15).

## B.2 A concentration inequality for empirical processes

Inequalities for suprema of empirical process were used at several places in these notes. These results have a long history (see [60, Chapter 3]) and are particularly well developed for uniformly bounded variables. In the random design setting here, the processes are often not uniformly bounded due to the presence of the Gaussian noise variables  $\varepsilon_i$ . This can be dealt with as in the proof of Lemma 4.1.7 under integrability conditions only on the ‘square-root’ metric entropy in (4.4). However, if one is willing to work with stronger hypothesis of integrability of the metric entropy (without the square-root), a more basic concentration argument is sometimes equally useful. The following lemma is based on a chaining argument for stochastic processes with a mixed tail (cf. [120, Theorem 2.2.28] and also [44, Theorem 3.5]). For us it will be sufficient to control the ‘generic chaining’ functionals employed in these references by suitable metric entropy integrals. For any (semi-)metric  $d$  on a metric space  $T$ , we denote by  $N = N(T, d, \rho)$  the minimal cardinality of a covering of  $T$  by balls with centres  $(t_i : i = 1, \dots, N) \subset T$  such that for all  $t \in T$ , there exists  $i$  such that  $d(t, t_i) < \rho$ . Below we require the index set  $\Theta$  to be countable (to avoid measurability issues). Whenever we apply Lemma B.2.1 in these notes with an uncountable set  $\Theta$ , one can show that the supremum can be realised as one over a countable subset of it.

**Lemma B.2.1.** *Let  $\Theta$  be a countable set. Suppose a class of real-valued measurable functions*

$$\mathcal{H} = \{h_\theta : \mathcal{X} \rightarrow \mathbb{R}, \theta \in \Theta\}$$

*defined on a probability space  $(\mathcal{X}, \mathcal{A}, P^X)$  is uniformly bounded by  $U \geq \sup_\theta \|h_\theta\|_\infty$  and has variance envelope  $\sigma^2 \geq \sup_\theta E^X h_\theta^2(X)$ , where  $X \sim P^X$ . Define metric entropy integrals*

$$J_2(\mathcal{H}) = \int_0^{4\sigma} \sqrt{\log N(\mathcal{H}, d_2, \rho)} d\rho \quad \text{for } d_2(\theta, \theta') := \sqrt{E^X [h_\theta(X) - h_{\theta'}(X)]^2},$$

$$J_\infty(\mathcal{H}) = \int_0^{4U} \log N(\mathcal{H}, d_\infty, \rho) d\rho \quad \text{for } d_\infty(\theta, \theta') := \|h_\theta - h_{\theta'}\|_\infty.$$

*For  $X_1, \dots, X_N$  drawn i.i.d. from  $P^X$  and  $\varepsilon_i \stackrel{\text{i.i.d.}}{\sim} N(0, 1)$  independent of all the  $X_i$ ’s, consider empirical processes arising either as*

$$Z_N(\theta) = \frac{1}{\sqrt{N}} \sum_{i=1}^N h_\theta(X_i) \varepsilon_i, \quad \theta \in \Theta,$$

or as

$$Z_N(\theta) = \frac{1}{\sqrt{N}} \sum_{i=1}^N (h_\theta(X_i) - E h_\theta(X)), \quad \theta \in \Theta.$$

We then have for some universal constant  $L > 0$  and all  $x \geq 1$ ,

$$\Pr \left( \sup_{\theta \in \Theta} |Z_N(\theta)| \geq L \left[ J_2(\mathcal{H}) + \sigma \sqrt{x} + \frac{J_\infty(\mathcal{H}) + Ux}{\sqrt{N}} \right] \right) \leq 2e^{-x}.$$

*Proof.* We only prove the case where  $Z_N(\theta) = \sum_i h_\theta(X_i) \varepsilon_i / \sqrt{N}$ ; the simpler case without Gaussian multipliers is proved in the same way. We will apply [44, Theorem 3.5], whose condition (3.8) we need to verify. First, we note that for  $|\lambda| < 1/\|h_\theta - h_{\theta'}\|_\infty$ , and  $E^\varepsilon$  denoting the expectation with respect to  $\varepsilon$ ,

$$\begin{aligned} & E \exp\{\lambda \varepsilon (h_\theta - h_{\theta'})(X)\} \\ & \leq 1 + \sum_{k=2}^{\infty} \frac{|\lambda|^k E^\varepsilon |\varepsilon|^k E^X |h_\theta - h_{\theta'}|^k(X)}{k!} \\ & \leq 1 + \lambda^2 E^X [h_\theta(X) - h_{\theta'}(X)]^2 \sum_{k=2}^{\infty} \frac{E^\varepsilon |\varepsilon|^k}{k!} (|\lambda| \|h_\theta - h_{\theta'}\|_\infty)^{k-2} \\ & \leq \exp\left\{ \frac{\lambda^2 d_2^2(\theta, \theta')}{1 - |\lambda| d_\infty(\theta, \theta')} \right\}, \end{aligned} \tag{B.2}$$

where we have used the basic fact  $E^\varepsilon |\varepsilon|^k / k! \leq 1$ . By the i.i.d. hypothesis we then also have

$$E \exp\{\lambda (Z_N(\theta) - Z_N(\theta'))\} \leq \exp\left\{ \frac{\lambda^2 d_2^2(\theta, \theta')}{1 - |\lambda| d_\infty(\theta, \theta') / \sqrt{N}} \right\}.$$

An application of the exponential Chebyshev inequality (and optimisation in  $\lambda$ , as in the proof of [60, Proposition 3.1.8]) then implies that condition (3.8) in [44] holds for the stochastic process  $Z_N(\theta)$  with metrics  $\bar{d}_2 = 2d_2$  and  $\bar{d}_1 = d_\infty / \sqrt{N}$ . In particular, the  $\bar{d}_2$ -diameter  $\Delta_2(\mathcal{H})$  of  $\mathcal{H}$  is at most  $4\sigma$  and the  $\bar{d}_1$ -diameter  $\Delta_1(\mathcal{H})$  of  $\mathcal{H}$  is bounded by  $4U/\sqrt{N}$ . (These bounds are chosen so that they remain valid for the process without Gaussian multipliers as well.) Theorem 3.5 in [44] now gives, for some universal constant  $M$ , and any  $\theta_{\dagger} \in \Theta$  that

$$\Pr \left( \sup_{\theta \in \Theta} |Z_N(\theta) - Z_N(\theta_{\dagger})| \geq M \left( \gamma_2(\mathcal{H}) + \gamma_1(\mathcal{H}) + \sigma \sqrt{x} + \frac{Ux}{\sqrt{N}} \right) \right) \leq e^{-x},$$

where the ‘generic chaining’ functionals  $\gamma_1, \gamma_2$  are upper bounded by the respective metric entropy integrals of the metric spaces  $(\mathcal{H}, \bar{d}_i)$ ,  $i = 1, 2$ , up to universal constants (see [44, estimate (2.3)]). For  $\gamma_1$  also notice that a simple substitution

$\rho' = \rho\sqrt{N}$  implies that

$$\int_0^{4U/\sqrt{N}} \log N(\mathcal{H}, \bar{d}_1, \rho) d\rho = \frac{1}{\sqrt{N}} \int_0^{4U} \log N(\mathcal{H}, d_\infty, \rho') d\rho',$$

and we hence deduce that

$$\Pr\left(\sup_{\theta \in \Theta} |Z_N(\theta) - Z_N(\theta_\dagger)| \geq \bar{L} \left[ J_2(\mathcal{H}) + \sigma\sqrt{x} + \frac{J_\infty(\mathcal{H}) + Ux}{\sqrt{N}} \right] \right) \leq e^{-x} \quad (\text{B.3})$$

for some universal constant  $\bar{L}$ .

Now what precedes also implies the classical Bernstein inequality

$$\Pr\left(|Z_N(\theta)| \geq \sigma\sqrt{2x} + \frac{Ux}{3\sqrt{N}}\right) \leq 2e^{-x} \quad \text{for } x > 0, \quad (\text{B.4})$$

for any fixed  $\theta \in \Theta$ ,  $U \geq \|h_\theta\|_\infty$  and  $\sigma^2 \geq E^X h_\theta^2(X)$ , proved as (3.24) in [60], using (B.2). Applying this with  $\theta_\dagger$  and using (B.3), the final result follows now from

$$\begin{aligned} & \Pr\left(\sup_{\theta \in \Theta} |Z_N(\theta)| > 2\tau(x)\right) \\ & \leq \Pr\left(\sup_{\theta \in \Theta} |Z_N(\theta) - Z_N(\theta_\dagger)| > \tau(x)\right) + \Pr(|Z_N(\theta_\dagger)| > \tau(x)) \leq 2e^{-x}, \end{aligned}$$

for any  $x \geq 1$ , where  $\tau(x) = \bar{L}[J_2(\mathcal{H}) + \sigma\sqrt{x} + (J_\infty(\mathcal{H}) + Ux)/\sqrt{N}]$  and  $L \geq 2\bar{L} > 0$  is large enough.  $\blacksquare$

### B.3 Mixing time bounds for Langevin diffusions

In this section we collect some results about convergence guarantees for an Unadjusted Langevin Algorithm (ULA) from (1.16) for sampling from *strongly log-concave target probability measures*. Early results of this kind are due to [113] but the recent focus has been on non-asymptotic results in high-dimensional settings, see [39, 48, 49]. Some of the ideas underpinning the fast mixing of ‘hypercontractive’ Langevin diffusions (1.15) in high-dimensions go back a long way further, we mention [6], [72], the recent monograph [7], and further references below.

Suppose that  $\mu$  is a Borel probability measure on  $\mathbb{R}^D$  which has a Lebesgue density proportional to  $e^{-U}$  for some potential  $U: \mathbb{R}^D \rightarrow \mathbb{R}$ , specifically

$$\mu(B) = \frac{\int_B e^{-U(\theta)} d\theta}{\int_{\mathbb{R}^D} e^{-U(\theta)} d\theta}, \quad B \subseteq \mathbb{R}^D \text{ measurable.} \quad (\text{B.5})$$

Following [49] (cf. H1 and H2 there) we will assume that the potential  $U$  has a  $\Lambda$ -Lipschitz gradient and is  $m$ -strongly convex, in particular,  $\mu$  is then log-concave.

**Assumption B.3.1.** (1) The function  $U: \mathbb{R}^D \rightarrow \mathbb{R}$  is continuously differentiable and there exists a constant  $\Lambda \geq 0$  such that for all  $\theta, \bar{\theta} \in \mathbb{R}^D$ ,

$$\|\nabla U(\theta) - \nabla U(\bar{\theta})\|_{\mathbb{R}^D} \leq \Lambda \|\theta - \bar{\theta}\|_{\mathbb{R}^D}.$$

(2) There exists a constant  $0 < m \leq \Lambda$  such that for all  $\theta, \bar{\theta} \in \mathbb{R}^D$ , we have

$$U(\bar{\theta}) \geq U(\theta) + \langle \nabla U(\theta), \bar{\theta} - \theta \rangle_{\mathbb{R}^D} + \frac{m}{2} \|\theta - \bar{\theta}\|_{\mathbb{R}^D}^2.$$

Under Assumption B.3.1, the potential  $U$  has a unique minimiser over  $\mathbb{R}^D$ , which we shall denote by  $\theta_U$ .

Now if  $(L_t : t \geq 0)$  is the  $D$ -dimensional Langevin Markov diffusion process from (1.15) with potential  $U$  satisfying Assumption B.3.1, and if  $P_t^{\theta_{\text{init}}}$  is its distribution at time  $t$  when started at  $\theta_{\text{init}} \in \mathbb{R}^D$ , then one can show a mixing time inequality, with  $W_2$  from (5.12) and  $E_D = \mathbb{R}^D$ ,

$$W_2(P_t^{\theta_{\text{init}}}, \mu) \leq e^{-mt} \left[ \|\theta_{\text{init}} - \theta_U\|_{\mathbb{R}^D} + \sqrt{\frac{D}{m}} \right], \tag{B.6}$$

which exhibits only a weak dependence on dimension  $D$ . A direct proof of (B.6) based on coupling ideas [33] can be found in [49, Proposition 1]. A functional analytic interpretation is that log-concave measures  $\mu$  satisfy a  $(1/m)$ -log-Sobolev inequality [7, Proposition 5.7.1] and the resulting dimension-free mixing bound for the entropic distance [7, p. 244] carries over to the  $W_2$ -distance by the quadratic transportation cost inequality of Otto–Villani [7, Theorem 9.6.1].

Now recall the Euler–Maruyama discretisation (1.16) of the Markov diffusion process (1.15) giving rise to the discrete-time Markov chain  $(\vartheta_k : k \geq 0)$ . We denote by  $\mathbf{P}_{\theta_{\text{init}}}, \mathbf{E}_{\theta_{\text{init}}}$  the law and expectation operator, respectively, of the Markov chain  $(\vartheta_k : k \geq 1)$  when started at a deterministic point  $\vartheta_0 = \theta_{\text{init}}$ . We also write  $\mathcal{L}(\vartheta_k)$  for the (marginal) distribution of the  $k$ -th iterate  $\vartheta_k$ . The following result now gives a ‘discrete’ analogue of (B.6).

**Proposition B.3.2.** *Suppose that  $U$  satisfies Assumption B.3.1. Then we have for  $\mu$  as in (B.5) that*

$$W_2^2(\mathcal{L}(\vartheta_k), \mu) \leq 2 \left(1 - \frac{m\gamma}{2}\right)^k \left[ \|\theta_{\text{init}} - \theta_U\|_{\mathbb{R}^D}^2 + \frac{D}{m} \right] + \frac{b(\gamma)}{2} \quad \text{for } k \geq 0, \tag{B.7}$$

where

$$b(\gamma) = 36 \frac{\gamma D \Lambda^2}{m^2} + 12 \frac{\gamma^2 D \Lambda^4}{m^3}, \tag{B.8}$$

*Proof.* The bound (B.7) follows from an application of [49, Theorem 5] with fixed step size  $\gamma > 0$ , where in our case, noting again that  $\kappa \in [m, 2m]$ , the expression  $u_n^{(1)}(\gamma)$  is upper bounded by  $2(1 - m\gamma/2)^k$  and the expression  $u_n^{(2)}(\gamma)$  is upper

bounded by (using that  $\gamma \leq \min\{2/\Lambda, 1/m\} \leq \min\{2/\Lambda, 2/\kappa\}$ )

$$\begin{aligned} & \Lambda^2 D\gamma^2 (\kappa^{-1} + \gamma) \left(2 + \frac{\Lambda^2 \gamma}{m} + \frac{\Lambda^2 \gamma^2}{6}\right) \sum_{i=1}^k \left(1 - \frac{\kappa \gamma}{2}\right)^{k-i} \\ & \leq \Lambda^2 D\gamma^2 (\kappa^{-1} + \gamma) \left(2 + \frac{\Lambda^2 \gamma}{m} + \frac{\Lambda^2 \gamma^2}{6}\right) \frac{2}{\kappa \gamma} \\ & \leq \Lambda^2 D\gamma \left(\kappa^{-2} + \frac{\gamma}{\kappa}\right) \left(6 + \frac{2\Lambda^2 \gamma}{m}\right) \\ & \leq \Lambda^2 D\gamma m^{-2} \left(18 + \frac{6\Lambda^2 \gamma}{m}\right), \end{aligned}$$

which equals (B.8). ■

For any measurable function  $H: \mathbb{R}^D \rightarrow \mathbb{R}$  and any  $J_{\text{in}}, J \geq 0$ , let us define the average of  $H$  along an ULA trajectory after ‘burn-in’ period  $J_{\text{in}}$  by

$$\hat{\mu}_{J_{\text{in}}}^J(H) = \frac{1}{J} \sum_{k=J_{\text{in}}+1}^{J_{\text{in}}+J} H(\vartheta_k).$$

**Proposition B.3.3.** *Suppose that  $U$  satisfies Assumption B.3.1 and suppose  $\gamma \leq 2/(m + \Lambda)$ . Then for all  $J, J_{\text{in}} \geq 1, x > 0$  and any Lipschitz function  $H: \mathbb{R}^D \rightarrow \mathbb{R}$ , we have the concentration inequality*

$$\mathbf{P}_{\theta_{\text{mit}}}(\hat{\mu}_{J_{\text{in}}}^J(H) - \mathbf{E}_{\theta_{\text{mit}}}[\hat{\mu}_{J_{\text{in}}}^J(H)] \geq x) \leq \exp\left(-\frac{J\gamma x^2 m^2}{16\|H\|_{\text{Lip}}^2(1 + 2/(mJ\gamma))}\right).$$

Moreover,

$$(\mathbf{E}_{\theta_{\text{mit}}}[\hat{\mu}_{J_{\text{in}}}^J(H)] - E_{\mu}H)^2 \leq \|H\|_{\text{Lip}}^2 \frac{1}{J} \sum_{k=J_{\text{in}}+1}^{J_{\text{in}}+J} W_2^2(\mathcal{L}(\vartheta_k), \mu). \quad (\text{B.9})$$

*Proof.* The first statement follows from [49, Theorem 17], by  $\kappa = 2m\Lambda/(m + \Lambda) \in [m, 2m]$  and the fact that the constant  $v_{N,n}(\gamma)$  from (28) of [49] can be upper bounded by

$$1 + \frac{m^{-1} + 2/(m + \Lambda)}{\gamma J} \leq 1 + \frac{2}{m\gamma J}.$$

The display (B.9) is derived in (27) of [49]. ■

## B.4 A characterisation of vanishing efficient information

The following theorem is due to [125].

**Theorem B.4.1.** *In the setting of Theorem 3.1.5, assume that  $\psi$  satisfies (3.20). Then  $i_{\theta, H, \psi} = 0$ .*

*Proof.* We prove the result for the case when the information operator  $\mathbb{I}_\theta^* \mathbb{I}_\theta$  is compact on the Hilbert space  $\bar{H} = \overline{(H, \langle \cdot, \cdot \rangle_{L_\lambda^2})}$ , relevant in the setting of inverse problems in these notes. (The general result is proved in [125, Theorem 4.1].)

Let us write  $I \equiv \mathbb{I}_\theta$ ,  $L^2 = L_\lambda^2(\mathcal{X})$  in this proof, and let  $\ker(I^*I) = \{h \in \bar{H} : I^*Ih = 0\}$ . If  $I^*I$  is a compact operator on  $\bar{H}$ , then by the spectral theorem for self-adjoint operators, there exists an orthonormal system of  $\bar{H}$  of eigenvectors  $\{e_k : k \in \mathbb{N}\}$  spanning  $\bar{H} \ominus \ker(I^*I)$  corresponding to eigenvalues  $\lambda_k > 0$  so that

$$I^*Ie_k = \lambda_k e_k \quad \text{and} \quad I^*Ih = \sum_k \lambda_k \langle h, e_k \rangle_{L_\xi^2} e_k \quad \text{for } h \in \bar{H}.$$

We can then define the usual square-root operator  $(I^*I)^{1/2}$  by

$$(I^*I)^{1/2}h = \sum_k \lambda_k^{1/2} \langle h, e_k \rangle_{L_\xi^2} e_k, \quad h \in \bar{H}.$$

If we denote by  $P_0$  the  $L_\xi^2$ -projection onto  $\ker(I^*I)$ , then the range of  $(I^*I)^{1/2}$  equals

$$R((I^*I)^{1/2}) = \left\{ g \in \bar{H} : P_0(g) = 0, \sum_k \lambda_k^{-1} \langle e_k, g \rangle_{L_\xi^2}^2 < \infty \right\}. \quad (\text{B.10})$$

Indeed, using standard Hilbert space arguments, (a) since  $P_0(e_k) = 0$  for all  $k$ , for any  $h \in \bar{H}$  the element  $g = (I^*I)^{1/2}h$  belongs to the right-hand side in the last display, and conversely (b) if  $g$  satisfies  $P_0(g) = 0$  and  $\sum_k \lambda_k^{-1} \langle e_k, g \rangle_{L_\xi^2}^2 < \infty$ , then  $h = \sum_k \lambda_k^{-1/2} \langle e_k, g \rangle_{L_\xi^2} e_k$  belongs to  $\bar{H}$  and  $(I^*I)^{1/2}h = g$ .

Next, [125, Lemma A.3] implies that  $R(I^*) = R((I^*I)^{1/2})$ . Now suppose  $\psi \in \bar{H}$  is such that  $\psi \notin R(I^*)$  and hence  $\psi \notin R((I^*I)^{1/2})$ . Then from (B.10), either  $P_0(\psi) \neq 0$  or  $\sum_k \lambda_k^{-1} \langle e_k, \psi \rangle_{L_\xi^2}^2 = \infty$  (or both). In the first case, let  $\bar{h} = P_0(\psi)$  so

$$\|I\bar{h}\|_{L^2} = \|I(P_0(\psi))\|_{L^2} = \langle I^*I(P_0(\psi)), P_0(\psi) \rangle_{L_\xi^2} = 0$$

but  $\langle \psi, \bar{h} \rangle_{L_\xi^2} = \|P_0\psi\|_{L_\xi^2}^2 = \delta$  for some  $\delta > 0$ . Since  $H$  is dense in  $\bar{H}$ , for any  $\epsilon$ ,  $0 < \epsilon < \min(\delta/(2\|\psi\|_{L_\xi^2}), \delta^2/4)$ , we can find  $h \in H$  such that  $\|h - \bar{h}\|_{L_\xi^2} < \epsilon$  and, by continuity, also  $\|I(h - \bar{h})\|_{L^2} < \epsilon$ . Then

$$\sqrt{i_{\theta, h, \psi}} = \frac{\|Ih\|_{L^2}}{|\langle \psi, h \rangle_{L_\xi^2}|} \leq 2 \frac{\epsilon}{\delta} \leq \sqrt{\epsilon}.$$

We conclude that  $i_{\theta, H, \psi} < \epsilon$  in (3.18), so that the result follows since  $\epsilon$  was arbitrary. In the second case we have  $\sum_k \lambda_k^{-1} \langle e_k, \psi \rangle_{L_\xi^2}^2 = \infty$  and define

$$\psi_N = \sum_{k \leq N} \lambda_k^{-1} e_k \langle e_k, \psi \rangle_{L_\xi^2}, \quad N \in \mathbb{N},$$

which defines an element of  $\bar{H}$ . By density we can choose  $h_N \in H$  such that  $\|h_N - \psi_N\|_{L_\xi^2} < 1/\|\psi\|_{L_\xi^2}$  as well as  $\|I(h_N - \psi_N)\|_{L^2} < 1$ , for every  $N$  fixed. Next observe that

$$\begin{aligned} \langle \psi, \psi_N \rangle_{L_\xi^2} &= \sum_{k \leq N} \lambda_k^{-1} \langle e_k, \psi \rangle_{L_\xi^2}^2 \equiv M_N, \\ \|I(\psi_N)\|_{L^2}^2 &= \langle I^* I(\psi_N), \psi_N \rangle_{L_\xi^2} = \sum_{k \leq N} \lambda_k^{-1} \langle e_k, \psi \rangle_{L_\xi^2}^2 = M_N, \end{aligned}$$

and that  $M_N \rightarrow \infty$  as  $N \rightarrow \infty$ . Then by our choice of  $h_N \in H$  and if  $M_N \geq 2$ , we have, by the triangle inequality,

$$\begin{aligned} |\langle \psi, h_N \rangle_{L_\xi^2}| &\geq |\langle \psi, \psi_N \rangle_{L_\xi^2}| - |\langle \psi, \psi_N - h_N \rangle_{L_\xi^2}| \geq M_N - 1 \geq \frac{M_N}{2}, \\ \|I(h_N)\|_{L^2} &\leq \|I(\psi_N)\|_{L^2} + \|I(h_N - \psi_N)\|_{L^2} \leq \sqrt{M_N} + 1 \leq 2\sqrt{M_N}. \end{aligned}$$

From this we conclude that the inverse of (3.18) satisfies

$$i_{\theta, H, \psi}^{-1} \geq \frac{\langle \psi, h_N \rangle_{L_\xi^2}^2}{\|I h_N\|_{L^2}^2} \geq \frac{1}{16} \frac{M_N^2}{M_N} \geq \frac{M_N}{16}.$$

As  $N$  was arbitrary and  $M_N \xrightarrow{N \rightarrow \infty} \infty$  we must have  $i_{\theta, H, \psi} = 0$ , as desired. ■



## References

- [1] K. Abraham and R. Nickl, On statistical Calderón problems. *Math. Stat. Learn.* **2** (2019), no. 2, 165–216
- [2] R. A. Adams and J. J. F. Fournier, *Sobolev spaces*. Second edn., Pure and Applied Mathematics (Amsterdam) 140, Elsevier/Academic Press, Amsterdam, 2003
- [3] R. J. Adler and J. E. Taylor, *Random fields and geometry*. Springer Monographs in Mathematics, Springer, New York, 2007
- [4] R. Altmeyer, Polynomial time guarantees for sampling based posterior inference in high-dimensional generalised linear models. Preprint 2022, [arXiv:2208.13296v1](https://arxiv.org/abs/2208.13296v1)
- [5] S. Arridge, P. Maass, O. Öktem, and C.-B. Schönlieb, Solving inverse problems using data-driven models. *Acta Numer.* **28** (2019), 1–174
- [6] D. Bakry and M. Émery, Diffusions hypercontractives. In *Séminaire de probabilités, XIX, 1983/84*, pp. 177–206, Lecture Notes in Math. 1123, Springer, Berlin, 1985
- [7] D. Bakry, I. Gentil, and M. Ledoux, *Analysis and geometry of Markov diffusion operators*. Grundlehren der mathematischen Wissenschaften 348, Springer, Cham, 2014
- [8] G. Bal and K. Ren, Multi-source quantitative photoacoustic tomography in a diffusive regime. *Inverse Problems* **27** (2011), article no. 7, 075003
- [9] G. Bal and G. Uhlmann, Inverse diffusion theory of photoacoustics. *Inverse Problems* **26** (2010), article no. 8, 085010
- [10] A. S. Bandeira, A. El Alaoui, S. Hopkins, T. Schramm, A. S. Wein, and I. Zadik, *The Franz–Parisi criterion and computational trade-offs in high dimensional statistics*. Advances in Neural Information Processing Systems 35, pp. 33831–33844, Curran Associates, Inc., 2022
- [11] A. S. Bandeira, A. Maillard, R. Nickl, and S. Wang, On free energy barriers in Gaussian priors and failure of cold start MCMC for high-dimensional unimodal distributions. *Phil. Trans. Royal Society* **381** (2023), article no. 20220150
- [12] A. Barron, M. J. Schervish, and L. Wasserman, The consistency of posterior distributions in nonparametric problems. *Ann. Statist.* **27** (1999), no. 2, 536–561
- [13] R. F. Bass, *Diffusions and elliptic operators*. Probability and its Applications (New York), Springer-Verlag, New York, 1998
- [14] A. Belloni and V. Chernozhukov, On the computational complexity of MCMC-based estimators in large samples. *Ann. Statist.* **37** (2009), no. 4, 2011–2055
- [15] G. Ben Arous, R. Gheissari, and A. Jagannath, Algorithmic thresholds for tensor PCA. *Ann. Probab.* **48** (2020), no. 4, 2052–2087
- [16] G. Ben Arous, A. S. Wein, and I. Zadik, Free energy wells and overlap gap property in sparse PCA. *Comm. Pure Appl. Math.* online first (2022), DOI [10.1002/cpa.22083](https://doi.org/10.1002/cpa.22083)
- [17] M. Benning and M. Burger, Modern regularization methods for inverse problems. *Acta Numer.* **27** (2018), 1–111

- [18] A. Beskos, M. Girolami, S. Lan, P. E. Farrell, and A. M. Stuart, Geometric MCMC for infinite-dimensional inverse problems. *J. Comput. Phys.* **335** (2017), 327–351
- [19] L. Birgé, Approximation dans les espaces métriques et théorie de l’estimation. *Z. Wahrsch. Verw. Gebiete* **65** (1983), no. 2, 181–237
- [20] L. Birgé, Model selection for Gaussian regression with random design. *Bernoulli* **10** (2004), no. 6, 1039–1051
- [21] J. Bohr, Stability of the non-abelian  $X$ -ray transform in dimension  $\geq 3$ . *J. Geom. Anal.* **31** (2021), no. 11, 11226–11269
- [22] J. Bohr, A Bernstein–von-Mises theorem for the Calderón problem with piecewise constant conductivities. *Inverse Problems* **39** (2023), no. 1, article no. 015002
- [23] J. Bohr and R. Nickl, On log-concave approximations of high-dimensional posterior measures and stability properties in non-linear inverse problems. To appear in *Ann. Inst. Henri Poincaré Probab. Stat.* 2023; Preprint 2021, [arXiv:2105.07835](https://arxiv.org/abs/2105.07835)
- [24] A. Bonito, A. Cohen, R. DeVore, G. Petrova, and G. Welper, Diffusion coefficients estimation for elliptic partial differential equations. *SIAM J. Math. Anal.* **49** (2017), no. 2, 1570–1592
- [25] F.-X. Briol, C. J. Oates, M. Girolami, M. A. Osborne, and D. Sejdinovic, Probabilistic integration: a role in statistical computation? *Statist. Sci.* **34** (2019), no. 1, 1–22
- [26] A.-P. Calderón, On an inverse boundary value problem. In *Seminar on Numerical Analysis and its Applications to Continuum Physics (Rio de Janeiro, 1980)*, pp. 65–73, Soc. Brasil. Mat., Rio de Janeiro, 1980
- [27] I. Castillo, Pólya tree posterior distributions on densities. *Ann. Inst. Henri Poincaré Probab. Stat.* **53** (2017), no. 4, 2074–2102
- [28] I. Castillo and R. Nickl, Nonparametric Bernstein–von Mises theorems in Gaussian white noise. *Ann. Statist.* **41** (2013), no. 4, 1999–2028
- [29] I. Castillo and R. Nickl, On the Bernstein–von Mises phenomenon for nonparametric Bayes procedures. *Ann. Statist.* **42** (2014), no. 5, 1941–1969
- [30] I. Castillo and V. Ročková, Uncertainty quantification for Bayesian CART. *Ann. Statist.* **49** (2021), no. 6, 3482–3509
- [31] I. Castillo and J. Rousseau, A Bernstein–von Mises theorem for smooth functionals in semiparametric models. *Ann. Statist.* **43** (2015), no. 6, 2353–2383
- [32] I. Castillo and S. van der Pas, Multiscale Bayesian survival analysis. *Ann. Statist.* **49** (2021), no. 6, 3559–3582
- [33] M. F. Chen and S. F. Li, Coupling methods for multidimensional diffusion processes. *Ann. Probab.* **17** (1989), no. 1, 151–177
- [34] K. L. Chung and Z. X. Zhao, *From Brownian motion to Schrödinger’s equation*. Grundlehren der mathematischen Wissenschaften 312, Springer-Verlag, Berlin, 1995
- [35] P. Constantin and C. Foias, *Navier–Stokes equations*. Chicago Lectures in Mathematics, University of Chicago Press, Chicago, IL, 1988

- [36] S. L. Cotter, M. Dashti, J. C. Robinson, and A. M. Stuart, Bayesian inverse problems for functions and applications to fluid mechanics. *Inverse Problems* **25** (2009), no. 11, article no. 115008
- [37] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White, MCMC methods for functions: modifying old algorithms to make them faster. *Statist. Sci.* **28** (2013), no. 3, 424–446
- [38] D. D. Cox, An analysis of Bayesian inference for nonparametric regression. *Ann. Statist.* **21** (1993), no. 2, 903–923
- [39] A. S. Dalalyan, Theoretical guarantees for approximate sampling from smooth and log-concave densities. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** (2017), no. 3, 651–676
- [40] M. Dashti and A. M. Stuart, The Bayesian approach to inverse problems. In *Handbook of uncertainty quantification*, pp. 311–428, Springer, Cham, 2017
- [41] E. B. Davies, *Heat kernels and spectral theory*. Cambridge Tracts in Mathematics 92, Cambridge University Press, Cambridge, 1990
- [42] P. Diaconis, Bayesian numerical analysis. In *Statistical decision theory and related topics, IV, Vol. 1 (West Lafayette, Ind., 1986)*, pp. 163–175, Springer, New York, 1988
- [43] P. Diaconis and D. Freedman, On the consistency of Bayes estimates. *Ann. Statist.* **14** (1986), no. 1, 1–67
- [44] S. Dirksen, Tail bounds via generic chaining. *Electron. J. Probab.* **20** (2015), no. 53
- [45] J. L. Doob, Application of the theory of martingales. In *Le Calcul des Probabilités et ses Applications*, pp. 23–27, Colloques Internationaux du Centre National de la Recherche Scientifique, no. 13, Centre National de la Recherche Scientifique, Paris, 1949
- [46] R. M. Dudley, *Real analysis and probability*. Cambridge Studies in Advanced Mathematics 74, Cambridge University Press, Cambridge, 2002
- [47] R. M. Dudley, *Uniform central limit theorems*. Second edn., Cambridge Studies in Advanced Mathematics 142, Cambridge University Press, New York, 2014
- [48] A. Durmus and É. Moulines, Nonasymptotic convergence analysis for the unadjusted Langevin algorithm. *Ann. Appl. Probab.* **27** (2017), no. 3, 1551–1587
- [49] A. Durmus and É. Moulines, High-dimensional Bayesian inference via the unadjusted Langevin algorithm. *Bernoulli* **25** (2019), no. 4A, 2854–2882
- [50] D. E. Edmunds and H. Triebel, *Function spaces, entropy numbers, differential operators*. Cambridge Tracts in Mathematics 120, Cambridge University Press, Cambridge, 1996
- [51] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of inverse problems*. Mathematics and its Applications 375, Kluwer Academic Publishers Group, Dordrecht, 1996
- [52] L. C. Evans, *Partial differential equations*. Second edn., Graduate Studies in Mathematics 19, American Mathematical Society, Providence, RI, 2010
- [53] D. Freedman, On the Bernstein–von Mises theorem with infinite-dimensional parameters. *Ann. Statist.* **27** (1999), no. 4, 1119–1140
- [54] D. A. Freedman, On the asymptotic behavior of Bayes’ estimates in the discrete case. *Ann. Math. Statist.* **34** (1963), 1386–1403

- [55] M. Freidlin, *Functional integration and partial differential equations*. Annals of Mathematics Studies 109, Princeton University Press, Princeton, NJ, 1985
- [56] C. F. Gauß, *Theoria motus corporum coelestium in sectionibus conicis solem ambientium*. Frid. Perthes and I. H. Besser, Hamburg, 1809
- [57] S. Ghosal, J. K. Ghosh, and A. W. van der Vaart, Convergence rates of posterior distributions. *Ann. Statist.* **28** (2000), no. 2, 500–531
- [58] S. Ghosal and A. van der Vaart, *Fundamentals of nonparametric Bayesian inference*. Cambridge Series in Statistical and Probabilistic Mathematics 44, Cambridge University Press, Cambridge, 2017
- [59] D. Gilbarg and N. S. Trudinger, *Elliptic partial differential equations of second order*. Grundlehren der Mathematischen Wissenschaften 224, Springer-Verlag, Berlin-New York, 1977
- [60] E. Giné and R. Nickl, *Mathematical foundations of infinite-dimensional statistical models*. Cambridge Series in Statistical and Probabilistic Mathematics 40, Cambridge University Press, New York, 2016
- [61] M. Giordano and H. Kekkonen, Bernstein–von Mises theorems and uncertainty quantification for linear inverse problems. *SIAM/ASA J. Uncertain. Quantif.* **8** (2020), no. 1, 342–373
- [62] M. Giordano and R. Nickl, Consistency of Bayesian inference with Gaussian process priors in an elliptic inverse problem. *Inverse Problems* **36** (2020), no. 8, article no. 085001
- [63] M. Giordano and K. Ray, Nonparametric Bayesian inference for reversible multidimensional diffusions. *Ann. Statist.* **50** (2022), no. 5, 2872–2898
- [64] S. Gugushvili, A. van der Vaart, and D. Yan, Bayesian linear inverse problems in regularity scales. *Ann. Inst. Henri Poincaré Probab. Stat.* **56** (2020), no. 3, 2081–2107
- [65] M. Hairer, A. M. Stuart, and S. J. Vollmer, Spectral gaps for a Metropolis–Hastings algorithm in infinite dimensions. *Ann. Appl. Probab.* **24** (2014), no. 6, 2455–2490
- [66] T. Helin and R. Kretschmann, Non-asymptotic error estimates for the Laplace approximation in Bayesian inverse problems. *Numer. Math.* **150** (2022), no. 2, 521–549
- [67] A. Hilger, I. Manke, N. Kardjilov, M. Osenberg, H. Markötter, and J. Banhart, Tensorial neutron tomography of three-dimensional magnetic vector fields in bulk materials. *Nat. Commun.* **9** (2018), article no. 4023
- [68] J. Ilmavirta and F. Monard, Integral geometry on manifolds with boundary and applications. In *The Radon transform—the first 100 years and beyond*, pp. 43–113, Radon Ser. Comput. Appl. Math. 22, Walter de Gruyter, Berlin, 2019
- [69] K. Ito and K. Kunisch, On the injectivity and linearization of the coefficient-to-solution mapping for elliptic boundary value problems. *J. Math. Anal. Appl.* **188** (1994), no. 3, 1040–1066
- [70] M. Jerrum, *Counting, sampling and integrating: algorithms and complexity*. Lectures in Mathematics, ETH Zürich, Birkhäuser Verlag, Basel, 2003
- [71] I. M. Johnstone, High dimensional Bernstein–von Mises: simple examples. In *Borrowing strength: theory powering applications—a Festschrift for Lawrence D. Brown*, pp. 87–98, Inst. Math. Stat. Collect. 6, Inst. Math. Statist., Beachwood, OH, 2010

- [72] R. Jordan, D. Kinderlehrer, and F. Otto, The variational formulation of the Fokker-Planck equation. *SIAM J. Math. Anal.* **29** (1998), no. 1, 1–17
- [73] B. Kaltenbacher, A. Neubauer, and O. Scherzer, *Iterative regularization methods for nonlinear ill-posed problems*. Radon Series on Computational and Applied Mathematics 6, De Gruyter, Berlin, 2008
- [74] A. Katchalov, Y. Kurylev, and M. Lassas, *Inverse boundary spectral problems*. Chapman & Hall/CRC Monographs and Surveys in Pure and Applied Mathematics 123, Chapman & Hall/CRC, Boca Raton, FL, 2001
- [75] H. Kekkonen, Consistency of Bayesian inference with Gaussian process priors for a parabolic inverse problem. *Inverse Problems* **38** (2022), no. 3, article no. 035002
- [76] B. T. Knapik, A. W. van der Vaart, and J. H. van Zanten, Bayesian inverse problems with Gaussian priors. *Ann. Statist.* **39** (2011), no. 5, 2626–2657
- [77] J. Kuelbs and W. V. Li, Metric entropy and the small ball problem for Gaussian measures. *J. Funct. Anal.* **116** (1993), no. 1, 133–157
- [78] P. S. m. de Laplace, *Théorie analytiques des probabilités*. Courcier, Paris, 1812
- [79] R. Latała and D. Matlak, Royen’s proof of the Gaussian correlation inequality. In *Geometric aspects of functional analysis*, pp. 265–275, Lecture Notes in Math. 2169, Springer, Cham, 2017
- [80] K. Law, A. Stuart, and K. Zygalakis, *Data assimilation*. Texts in Applied Mathematics 62, Springer, Cham, 2015
- [81] L. Le Cam, On some asymptotic properties of maximum likelihood estimates and related Bayes’ estimates. *Univ. California Publ. Statist.* **1** (1953), 277–329
- [82] L. Le Cam, *Asymptotic methods in statistical decision theory*. Springer Series in Statistics, Springer-Verlag, New York, 1986
- [83] H. Leahu, On the Bernstein–von Mises phenomenon in the Gaussian white noise model. *Electron. J. Stat.* **5** (2011), 373–404
- [84] W. V. Li and W. Linde, Approximation, metric entropy and small ball estimates for Gaussian measures. *Ann. Probab.* **27** (1999), no. 3, 1556–1578
- [85] W. V. Li and Q.-M. Shao, Gaussian processes: inequalities, small ball probabilities and applications. In *Stochastic processes: theory and methods*, pp. 533–597, Handbook of Statist. 19, North-Holland, Amsterdam, 2001
- [86] J.-L. Lions and E. Magenes, *Non-homogeneous boundary value problems and applications. Vol. I*. Grundlehren der mathematischen Wissenschaften 181, Springer-Verlag, New York-Heidelberg, 1972
- [87] Y. Lu, A. Stuart, and H. Weber, Gaussian approximations for probability measures on  $\mathbb{R}^d$ . *SIAM/ASA J. Uncertain. Quantif.* **5** (2017), no. 1, 1136–1165
- [88] F. Monard, Functional relations, sharp mapping properties, and regularization of the X-ray transform on disks of constant curvature. *SIAM J. Math. Anal.* **52** (2020), no. 6, 5675–5702
- [89] F. Monard, R. Nickl, and G. P. Paternain, Efficient nonparametric Bayesian inference for X-ray transforms. *Ann. Statist.* **47** (2019), no. 2, 1113–1147

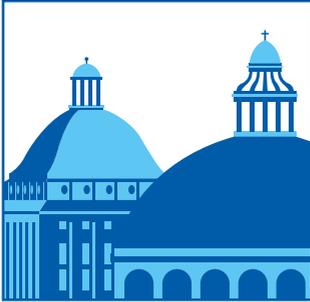
- [90] F. Monard, R. Nickl, and G. P. Paternain, Consistent inversion of noisy non-Abelian X-ray transforms. *Comm. Pure Appl. Math.* **74** (2021), no. 5, 1045–1099
- [91] F. Monard, R. Nickl, and G. P. Paternain, Statistical guarantees for Bayesian uncertainty quantification in nonlinear inverse problems with Gaussian process priors. *Ann. Statist.* **49** (2021), no. 6, 3255–3298
- [92] A. I. Nachman, Global uniqueness for a two-dimensional inverse boundary value problem. *Ann. of Math. (2)* **143** (1996), no. 1, 71–96
- [93] R. Nickl, Bernstein–von Mises theorems for statistical inverse problems I: Schrödinger equation. *J. Eur. Math. Soc. (JEMS)* **22** (2020), no. 8, 2697–2750
- [94] R. Nickl, Inference for diffusions from low frequency measurements. Preprint 2022, [arXiv:2210.13008](https://arxiv.org/abs/2210.13008)
- [95] R. Nickl and G. P. Paternain, On some information-theoretic aspects of non-linear statistical inverse problems. To appear in Proc. Int. Cong. Math. 2022; Preprint 2021, [arXiv:2107.09488](https://arxiv.org/abs/2107.09488)
- [96] R. Nickl and K. Ray, Nonparametric statistical inference for drift vector fields of multi-dimensional diffusions. *Ann. Statist.* **48** (2020), no. 3, 1383–1408
- [97] R. Nickl and J. Söhl, Nonparametric Bayesian posterior contraction rates for discretely observed scalar diffusions. *Ann. Statist.* **45** (2017), no. 4, 1664–1693
- [98] R. Nickl and J. Söhl, Bernstein–von Mises theorems for statistical inverse problems II: compound Poisson processes. *Electron. J. Stat.* **13** (2019), no. 2, 3513–3571
- [99] R. Nickl and E. Titi, *Posterior consistency of Bayesian data assimilation with Gaussian process priors for the 2D Navier–Stokes equations*. Unpublished manuscript (2023)
- [100] R. Nickl, S. van de Geer, and S. Wang, Convergence rates for penalized least squares estimators in PDE constrained regression problems. *SIAM/ASA J. Uncertain. Quantif.* **8** (2020), no. 1, 374–413
- [101] R. Nickl and S. Wang, On polynomial-time computation of high-dimensional posterior measures by Langevin-type algorithms. *J. Eur. Math. Soc.* Online First (2022), DOI [10.4171/jems/1304](https://doi.org/10.4171/jems/1304)
- [102] G. P. Paternain and M. Salo, The non-Abelian X-ray transform on surfaces. Preprint (2020) [arXiv:2006.02257](https://arxiv.org/abs/2006.02257); *J. Differential Geom.*, to appear
- [103] G. P. Paternain, M. Salo, and G. Uhlmann, The attenuated ray transform for connections and Higgs fields. *Geom. Funct. Anal.* **22** (2012), no. 5, 1460–1489
- [104] G. P. Paternain, M. Salo, and G. Uhlmann, *Geometric inverse problems. With emphasis on two dimensions*. Cambridge Studies in Advanced Mathematics 204, Cambridge University Press, Cambridge, 2023
- [105] L. Pestov and G. Uhlmann, Two dimensional compact simple Riemannian manifolds are boundary distance rigid. *Ann. of Math. (2)* **161** (2005), no. 2, 1093–1110
- [106] J. Radon, Über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten. *Ber. Verh. K. Sächs. Ges. Wiss., Leipzig, Math.-Phys. Kl.* **69** (1917), 262–277.

- [107] K. Ray, Bayesian inverse problems with non-conjugate priors. *Electron. J. Stat.* **7** (2013), 2516–2549
- [108] S. Reich and C. Cotter, *Probabilistic forecasting and Bayesian data assimilation*. Cambridge University Press, New York, 2015
- [109] M. Reiß, Asymptotic equivalence for nonparametric regression with multivariate and random design. *Ann. Statist.* **36** (2008), no. 4, 1957–1982
- [110] G. R. Richter, An inverse problem for the steady state diffusion equation. *SIAM J. Appl. Math.* **41** (1981), no. 2, 210–221
- [111] V. Rivoirard and J. Rousseau, Bernstein–von Mises theorem for linear functionals of the density. *Ann. Statist.* **40** (2012), no. 3, 1489–1523
- [112] C. P. Robert and G. Casella, *Monte Carlo statistical methods*. Second edn., Springer Texts in Statistics, Springer-Verlag, New York, 2004
- [113] G. O. Roberts and R. L. Tweedie, Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2** (1996), no. 4, 341–363
- [114] M. Sales, M. Strobl, T. Shinohara, A. Tremsin, L. T. Kuhn, W. R. B. Lionheart, N. M. Desai, A. B. Dahl, and S. Schmidt, Three dimensional polarimetric neutron tomography of magnetic fields. *Sci. Rep.* **8** (2018), article no. 2214
- [115] C. Schillings, B. Sprungk, and P. Wacker, On the convergence of the Laplace approximation and noise-level-robustness of Laplace-based Monte Carlo methods for Bayesian inverse problems. *Numer. Math.* **145** (2020), no. 4, 915–971
- [116] L. Schwartz, On Bayes procedures. *Z. Wahrscheinlichkeitstheorie und Verw. Gebiete* **4** (1965), 10–26
- [117] V. Spokoiny, Dimension free non-asymptotic bounds on the accuracy of high dimensional Laplace approximation. Preprint 2022, [arXiv:2204.11038](https://arxiv.org/abs/2204.11038)
- [118] A. M. Stuart, Inverse problems: a Bayesian perspective. *Acta Numer.* **19** (2010), 451–559
- [119] J. Sylvester and G. Uhlmann, A global uniqueness theorem for an inverse boundary value problem. *Ann. of Math. (2)* **125** (1987), no. 1, 153–169
- [120] M. Talagrand, *Upper and lower bounds for stochastic processes*. Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge 60. xv, 626 p. Springer Berlin Heidelberg, 2014
- [121] M. E. Taylor, *Partial differential equations I. Basic theory*. Second edn., Applied Mathematical Sciences 115, Springer, New York, 2011
- [122] M. E. Taylor, *Partial differential equations II. Qualitative studies of linear equations*. Second edn., Applied Mathematical Sciences 116, Springer, New York, 2011
- [123] G. Uhlmann, Electrical impedance tomography and Calderón’s problem. *Inverse Problems* **25** (2009), no. 12, article no. 123011
- [124] S. A. van de Geer, *Applications of empirical process theory*. Cambridge Series in Statistical and Probabilistic Mathematics 6, Cambridge University Press, Cambridge, 2000
- [125] A. van der Vaart, On differentiable functionals. *Ann. Statist.* **19** (1991), no. 1, 178–204

- [126] A. W. van der Vaart, *Asymptotic statistics*. Cambridge Series in Statistical and Probabilistic Mathematics 3, Cambridge University Press, Cambridge, 1998
- [127] A. van der Vaart and H. van Zanten, Bayesian inference with rescaled Gaussian process priors. *Electron. J. Stat.* **1** (2007), 433–448
- [128] A. W. van der Vaart and J. H. van Zanten, Rates of contraction of posterior distributions based on Gaussian process priors. *Ann. Statist.* **36** (2008), no. 3, 1435–1463
- [129] A. W. van der Vaart and J. A. Wellner, *Weak convergence and empirical processes*. Springer Series in Statistics, Springer-Verlag, New York, 1996
- [130] J. van Neerven,  $\gamma$ -radonifying operators—a survey. In *The AMSI-ANU Workshop on Spectral Theory and Harmonic Analysis*, pp. 1–61, Proc. Centre Math. Appl. Austral. Nat. Univ. 44, Austral. Nat. Univ., Canberra, 2010
- [131] R. Vershynin, *High-dimensional probability*. Cambridge Series in Statistical and Probabilistic Mathematics 47, Cambridge University Press, Cambridge, 2018
- [132] C. Villani, *Optimal transport. Old and new*. Grundlehren der Mathematischen Wissenschaften 338. Springer-Verlag, Berlin, 2009
- [133] S. J. Vollmer, Posterior consistency for Bayesian inverse problems through stability and regression results. *Inverse Problems* **29** (2013), article no. 12, 125011
- [134] R. Wong, *Asymptotic approximations of integrals*. Classics in Applied Mathematics 34, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001

# Index

- adjoint score condition, 46
- asymptotic minimax lower bound, 46
  
- Bernstein–von Mises theorem, 67, 82, 84
  
- Calderón problem, 2, 38
- curvature, 41, 49, 52, 92, 93, 115
  
- Darcy’s problem, 4, 28, 37, 84, 89, 113, 114, 120
- diffusion equation, 4, 26
- divergence form operator, 4
- divergence theorem, 127
  
- elliptic regularity estimate, 134
  
- Feynman–Kac formula, 132
- forward map, 1, 7
- forward regularity, 24
  
- Gaussian correlation inequality, 140
- Gaussian process, 31, 139
- gradient stability, 50
  
- Hellinger distance, 15
- hitting time, 109
  
- information distances, 14
- information equation, 47, 69, 82, 85
- information operator, 42, 50, 58
- invariant measure, 10, 101, 109
- inverse information, 47, 48
- isoperimetric inequality, 35, 113
  
- Kullback–Leibler divergence, 14
  
- LAN expansion, 42
- Langevin algorithm, 13, 101, 119, 145
- Laplace approximation, 96
- log-concave, 91, 96, 145
- log-concave approximation, 95
  
- MCMC, 10, 21, 101, 106, 119
- metric entropy, 129
- mixing time, 102, 146
  
- non-Abelian X-ray transform, 3
  
- pCN algorithm, 10
- posterior contraction, 17, 35, 37, 69
- posterior distribution, 9
- posterior mean, 36, 77, 81, 83, 105, 120
  
- random design regression model, 7
- reproducing kernel Hilbert space, 32, 139
  
- Schrödinger equation, 4, 25, 28, 58, 61, 82
- Schrödinger operator, 4
- sieve prior, 39
- small ball probability, 140
- Sobolev space, 125
- stability estimate, 23, 28
  
- ULA algorithm, 14
- uncertainty quantification, 80
  
- Wasserstein distance, 95
- Whittle–Matérn process, 141



Richard Nickl

## Bayesian Non-linear Statistical Inverse Problems

Bayesian methods based on Gaussian process priors are frequently used in statistical inverse problems arising with partial differential equations (PDEs). They can be implemented by Markov chain Monte Carlo (MCMC) algorithms. The underlying statistical models are naturally high- or infinite-dimensional and the present book presents a rigorous mathematical analysis of the statistical performance, and algorithmic complexity, of such methods in a natural setting of non-linear random design regression.

Due to the non-linearity present in many of these inverse problems, natural least squares functionals are non-convex and the Bayesian paradigm presents an attractive alternative to optimisation-based approaches. This book develops a general theory of Bayesian inference for non-linear forward maps and rigorously considers two PDE model examples arising with Darcy's problem and a Schrödinger equation. The focus is initially on statistical consistency of Gaussian process methods, and then moves on to study local fluctuations and approximations of posterior distributions by Gaussian or log-concave measures whose curvature is described by PDE mapping properties of underlying 'information operators'. Applications to the algorithmic runtime of gradient-based MCMC methods are discussed as well as computation time lower bounds for worst case performance of some algorithms.

<https://ems.press>

ISBN 978-3-98547-053-2

